Optical Tracking for Performance Testing of Microsurgical Instruments

Robert MacLachlan
Cameron Riviere

CMU-RI-TR-07-01

January 2007

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Table of Contents

1

Abstract

ASAP is a position-measurement instrument developed to measure human hand tremor. It makes 1000 five-degree-of-freedom position measurements per second with 3 micron resolution over a 3cm workspace. The measurement principle is optical triangulation of modulated lights on the probe using position-sensitive-diode (PSD) sensors.

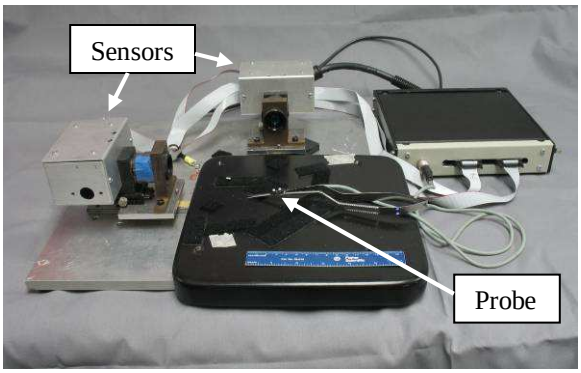

Figure 1: ASAP system

1 Introduction

ASAP (Apparatus for Sensing Accurate Positions) is an instrument for simultaneously measuring the 3D position of multiple modulated sources with micron resolution and a measurement rate of up to 1 kHz. This instrument was developed to support research on hand tremor during microsurgery[1] and the development of devices to compensate hand tremor**Error! Reference source not found.**[3]. Tremor does not contain significant energy above 20 Hz, so a sampling rate of 100 Hz was adequate for this work.

Recently, we have been using ASAP not only to provide experimental results for evaluation of canceling techniques, but also as part of the closed-loop tremor canceling control system itself. In order to exhibit significant cancellation, the loop bandwidth must substantially exceed the highest tremor frequencies, so a sampling rate of 1000 Hz has been used.

1.1 Mechanical configuration

The major functional components of the system are two 2D optical position sensors and a probe with two lights on it (Figure 1). The basic measurement principle is to triangulate the position of each light by using information from the two optical position sensors. With a two-

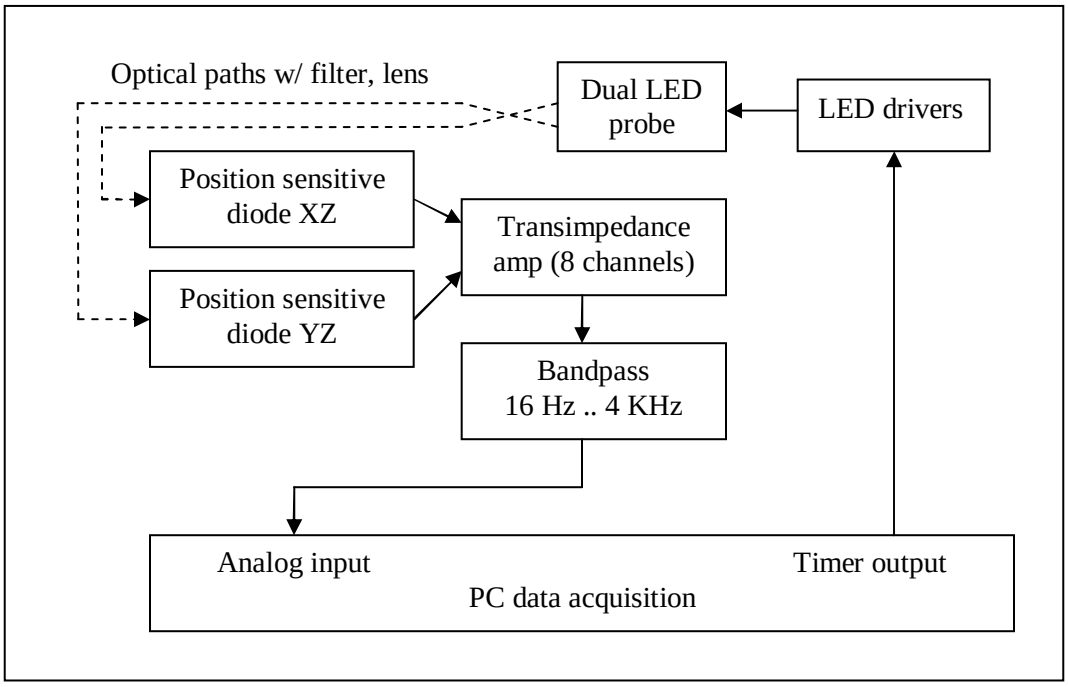

Figure 2: Hardware organization

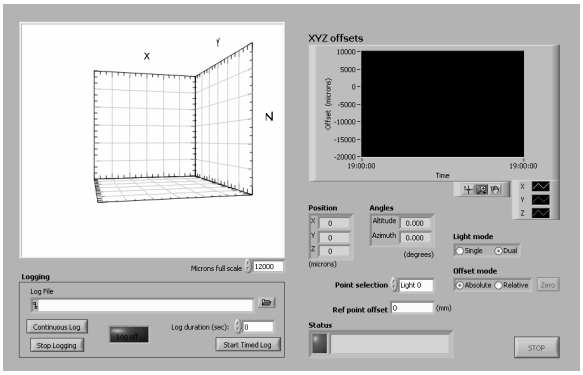

Figure 3: Labview user interface

light probe, the five-degree-of-freedom (5-DOF) position and orientation can be measured. The position sensors are arranged at right angles to each other. The measurement process does not depend on this special geometry, but this geometry is optimal in that it provides equal measurement resolution along all three axes.

1.2 System architecture

The instrument requires only a small amount of custom electronics (Figure 2). A commercial data acquisition board (National Instruments PCI-6289) digitizes the sensor signals and generates the excitation signal for the lights. The remainder of the signal processing and the instrument control interface are implemented in Labview software (Figure 3).

1.3 Position sensitive diodes

ASAP uses *position sensitive diodes* (PSDs) (Figure 4), also known as lateral-effect photodiodes, to measure the light positions. A PSD is a non-imaging analog position sensor. It is a large-area planar photodiode with high-resistivity contacts. The contact resistance forms a current divider that splits the photocurrent between the terminals, allowing the position on the active area to be determined.

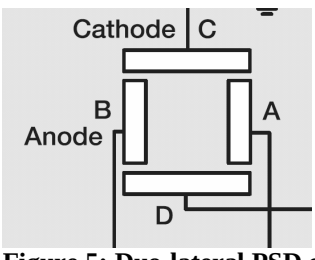

Figure 5: Duo-lateral PSD contact schematic

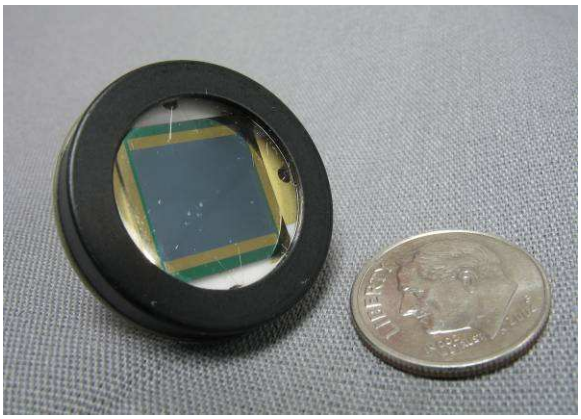

Figure 4: DL-10 position sensing diode

Because of the analog spatial averaging performed by the resistive contact, the PSD responds to the mean center of illumination (or centroid). The size and shape of the spot have little effect on the position measurement, so precise focusing is not required.

In the duo-lateral PSD, both the front and back (anode and cathode) contacts are high-resistivity sheets, with two contact busses at the opposite edges of each sheet. The front contacts are on the left and right edges of the sensor, and the rear contacts on the top and bottom edges. This allows the X and Y positions of the spot to be measured independently (Figure 5).

If the currents along an axis are X0 and X1, then the position on that axis is:

$$(X1 - X0)/(X1 + X0) \qquad (1)$$

This produces a normalized position over the range [-1, 1].

1.4 Lock-in amplification

ASAP uses lock-in amplification to eliminate various internal and external sources of noise and interference. Lock-in amplification is commonly used with high-performance PSD sensors, but ASAP is completely dependent on it for operation, since this is how the positions of two lights are independently measured by a single sensor.

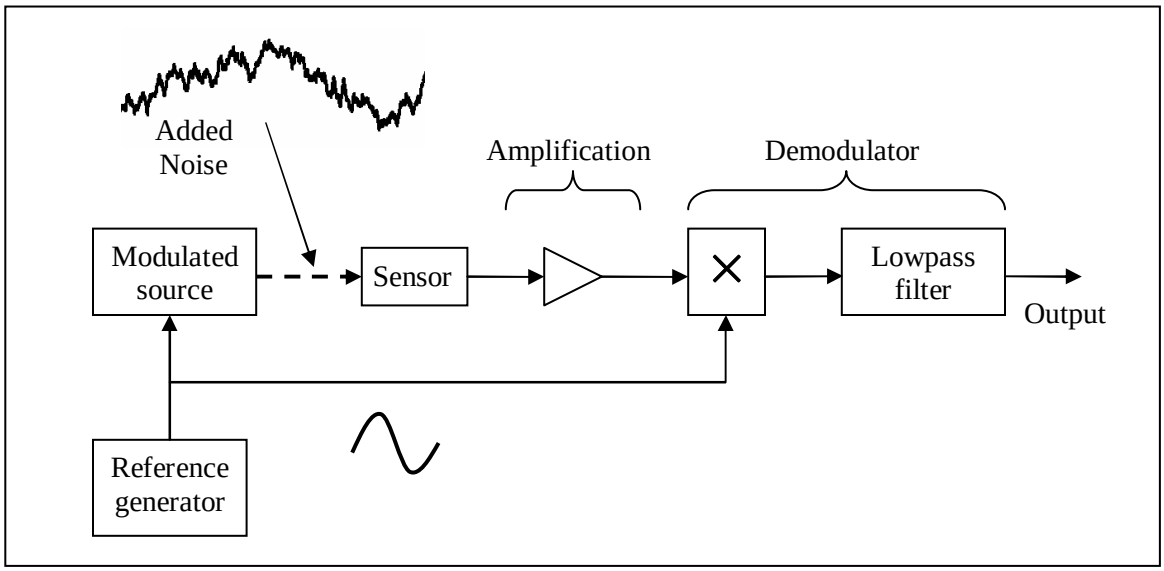

Figure 6: Lock-in amplification

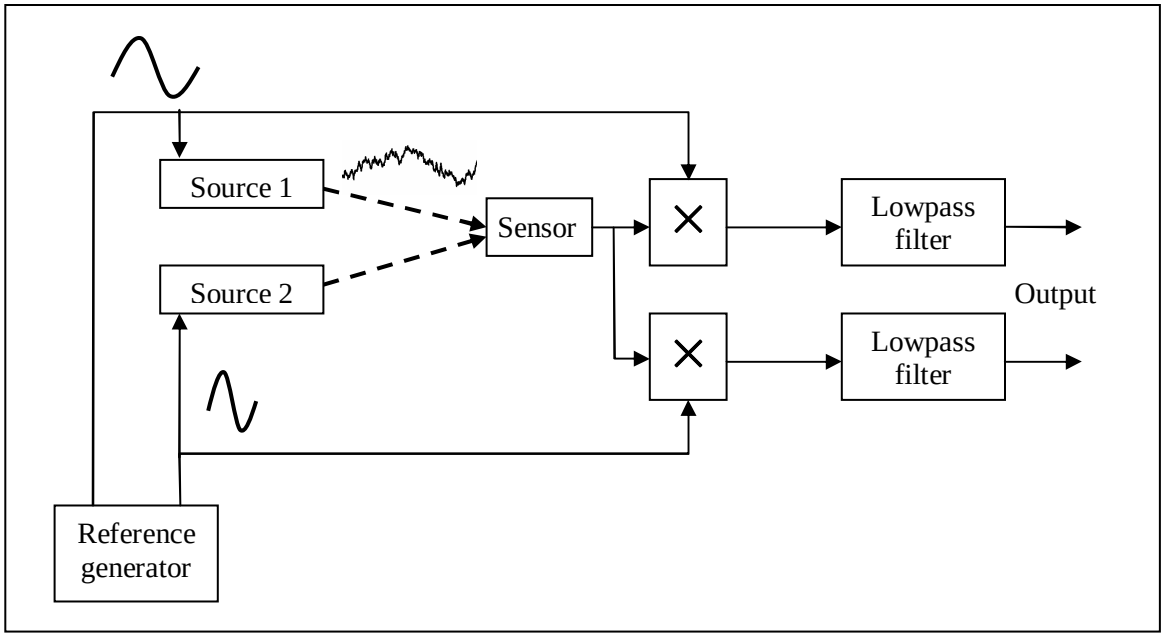

Figure 7: Lock in amplification with frequency domain multiplexing

The general scheme of lock-in amplification (Figure 6) involves periodic modulation of the source, combined with in-phase synchronous demodulation of the amplified sensor signal. The demodulator includes a lowpass filter that establishes the signal bandwidth.

The principle of operation of the demodulator is that multiplying by the reference signal creates frequency products proportional to the sum and difference between the components of the incoming signal and the reference. The sum products fall outside of the passband of the lowpass filter, leaving only the signal that falls within the filter bandwidth on either side of the reference frequency.

In addition to establishing the bandwidth, the filter also averages the $\sin^2$ harmonics of the modulated signal into the recovered signal amplitude. The lowpass filter must strongly suppress these harmonics to minimize ripple in the output signal.

The important properties of lock-in amplification are:

1. It easily creates an arbitrarily narrow receive bandwidth that tracks the frequency of the source modulation, removing most interfering signals.
2. The signal bandwidth no longer includes DC, removing offsets, drift and 1/f noise, regardless of whether they come from the source, signal propagation, sensor, or amplification.
3. The half of the noise which is out of phase with the reference is also rejected.

Lock-in amplification is very similar to the signal processing used in radio communication systems, leading to the observation that with minimal additional complexity we can use *frequency domain multiplexing* (Figure 7) to allow more than one signal to be transmitted simultaneously over the same medium. In other words, we deliberately introduce another potentially interfering signal which is sufficiently far separated in frequency that no significant interference takes place.

2 Related work

Many 3D measurement systems have been described in the literature, and there are also quite a few commercial products. These systems are based on a number of different measurement principles. Although several of the established measurement principles could be used to develop an instrument with the speed and resolution that we require, none of the existing systems that we considered did meet our requirements. Most available systems had a much larger workspace than we needed, and correspondingly coarse position resolution. Furthermore, most systems had too slow an update rate.

Given that we must develop our own hardware, an overriding consideration became to determine what measurement approach would result in the simplest hardware, and thus the least implementation effort.

2.1 Optical coordination sensor

The decision to use PSDs in the first version of ASAP [3] was strongly influenced by the successful use of PSDs in another sensor [4] at the Robotics Institute. This sensor had proven a design for the transimpedance amplifiers and the fairly complex analog demodulator; furthermore, spare circuit boards were available.

2.2 Original ASAP instrument

The first implementation of the ASAP instrument used the amplifier and analog demodulator cards described above. The optical and mechanical configuration of the current ASAP instrument was also developed at that time [5].

In addition to the modulation implementation, there were a number of other differences:

- Only a single 3D position was measured; no angular position information was available.
- A reflective probe with a modulated external source was used. This precluded frequency-domain multiplexing of the lights, degraded the signal-to-noise ratio due to lower brightness, and dramatically increased the undesirable influence of reflective objects other than the probe.
- The 3D localization algorithm neglected perspective effects, so was only accurate near the center of the workspace.

2.3 Magnetic systems

Several commercial products use a magnetic measurement principle (Polhemus *FasTrak*, Ascension *MicroBird*) where multiple modulated external magnetic fields are sensed by one or more pickup coils in a probe. The major advantage of this measurement principle is that no line of sight is required — the probe is easily sensed even inside the body. The major disadvantage is measurement errors due to conductive and (especially) ferro-magnetic objects in the vicinity.

Though a magnetic sensor optimized for small workspace and high update rate could likely achieve the required performance, the

commercial products provide insufficient speed and resolution.

2.4 Time domain multiplexed PSD systems

The alternative to frequency-domain multiplexing for multiple point measurement using PSDs is time-domain multiplexing. The lights are flashed in turn, and a single synchronized measurement of the PSD signals is made for each flash. The effect of drift, 1/f noise and background illumination can be reduced by taking a dark measurement shortly before or after the pulse and subtracting this from the lit measurement.

Reference [6] describes an interesting system that uses time-domain multiplexed PSDs. The "inside out" design (where the sensors are in the probe) makes the probe is too large for use in a microsurgery system. However, there is informative discussion of the measurement principle, and in particular, how to reduce measurement noise by using a Kalman filter.

In comparison to frequency-domain multiplexing, the time-domain approach has a fundamentally worse signal-to-noise ratio that degrades as the number of lights increases. This is because the bandwidth must be wide enough to pass the relatively short pulses. So the noise bandwidth must increase as the number of lights increases.

2.5 Camera systems

An obvious approach to 3D position tracking is to use two or more video cameras to triangulate the locations of reflective markers or active sources. Advantages are easy scaling to many points, suppression of secondary reflections via thresholding, and possible rejection of background illumination with background subtraction. Many commercial systems are available (such as the Northern Digital *Polaris*).

The main shortcomings for our purposes are:

- Commercial systems are optimized for capture of human gross motion, so have too low resolution.
- Update rate is limited to the camera speed, typically 60 Hz. Use of high-speed cameras could potentially increase this, but at the cost of a large data processing burden that would be difficult to sustain in real time.
- Resolution is limited by camera resolution, though interpolation allows sub-pixel resolution.

2.6 Linear CCD systems

A clever approach to 3D measurement of modulated sources uses 3 or more single-line CCD sensors with cylinder lenses. This greatly reduces the data-processing burden in comparison to the approach of using an imaging video camera, yet the same background subtraction and thresholding processing can be done.

Use of a linear CCD does require time-domain multiplexing of the lights, since only a single light can be resolved at one time. However, CCD dynamic range is limited mainly by quantum shot noise, and high-resolution linear CCDs are available. Even after allowing for degraded speed/resolution due to time-domain multiplexing, this measurement principle promises the highest dynamic range of those described here.

A commercial product that uses this principle is the Northern Digital *OptoTrak*. While the measurement speed is an adequate 1500 samples/sec, it has a 1 meter workspace, and a correspondingly large 100 micron resolution. The PhaseSpace *Impulse* also appears to use this principle, but has a marginal 480 samples/sec output rate, and is also designed for human motion capture over distances greater than 1 m.

Furthermore, turn-key commercial systems are frequently unsuitable for use inside a feedback loop due to significant processing delays in the proprietary software. For example, the PhaseSpace *Impulse* specifies < 10ms latency.

Since the commercial products are not suitable, the main question is whether this measurement principle would make sense in a custom system.

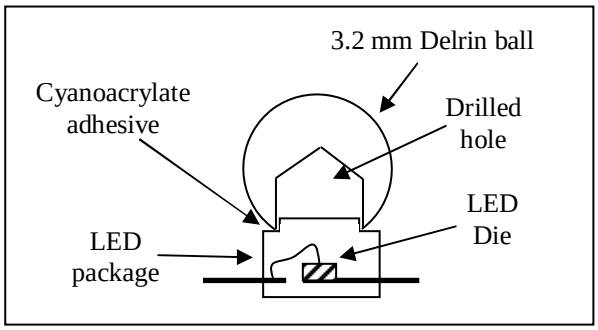

Figure 10: LED emitter construction

The main shortcoming is that the hardware complexity and signal processing demands seem considerably higher than in ASAP.

3 Construction details

This section details the construction of the major hardware components: the LED probe, the sensors, and the signal conditioning. The redesign incorporated several packaging improvements that improved performance, stability and reliability, such as moving the transimpedance amplifiers near the PSD sensor, using low noise power supplies and improving grounding.

3.1 LED probe

An ideal emitter would have rotational symmetry and uniform illumination such that the 2D centroid of the image is coincident with the 3D center of the emitter when viewed over a wide range of angles; any departures from this ideal causes a position error that varies with viewing angle.

Our emitter is an illuminated sphere created by attaching a Delrin ball to a packaged LED (Figure 10). The drilled hole serves two purposes: positive location of the ball on the LED and increased illumination uniformity. In order to maximize the symmetry of the emitter, light leakage from the clear body of the LED is masked by painting it silver, then black.

The LED is an HDSL-4400 (Avago Technologies). Its peak emission is at 875 nm.

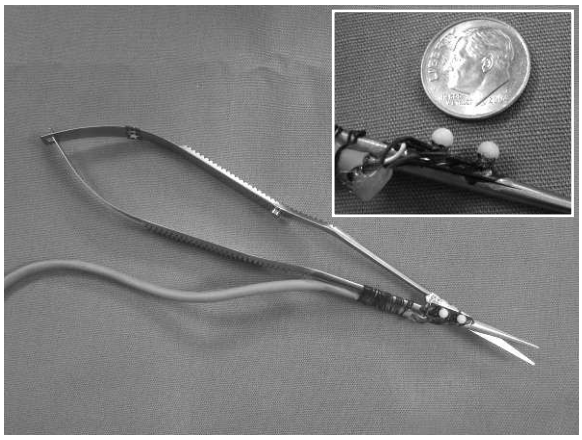

Figure 11: ASAP probe

The LED is pulsed using a 150 ma on/off square wave. The complete probe has two emitters along its axis, spaced 7 mm apart (Figure 11). Each LED is modulated with a distinct frequency (approximately 2 and 3 KHz).

3.2 Sensor design

Light is focused onto the PSD by a single-element double-convex lens (Figure 12). The lens is 12.5mm in diameter and has a 30mm focal length. It is constructed of fused quartz with an anti-reflective coating (Melles Griot). Before reaching the PSD, the light passes through a long-pass filter which minimizes interference by blocking visible light.

ASAP uses the OSI Optoelectronics DL-10 duo-lateral position sensing diode (Figure 4), which is a 1 cm square precision device with nonlinearity < 1% and drift of < 0.6 μ m/°C. The device is reverse-biased by 5 volts, primarily to increase the linearity under high illumination levels.

3.3 Signal conditioning

There are eight signal channels in all (four current terminals of each PSD, and two PSDs). Each signal channel has the same signal processing.

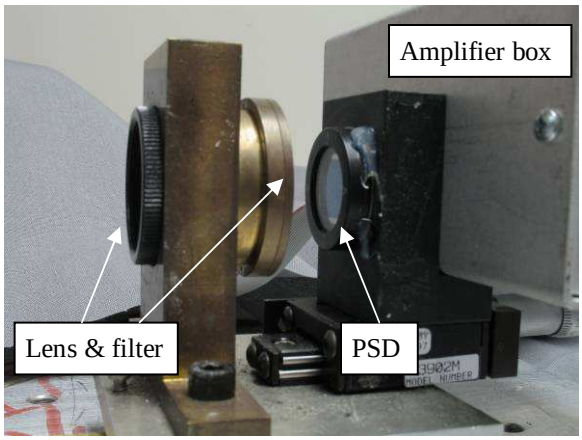

Figure 12: Sensor detail

Figure 13: LED drivers and antialias filters

First, there is a transimpedance amplifier (current to voltage converter) with a gain of 0.2 volts/μ A. Unusually for a photodiode, the PSD is a low impedance source due to the relatively low (10 KΩ) resistance of the sheet contacts. The resolution of the PSD is noise-limited, with the primary noise source being the Johnson noise of this source resistance[7]. To optimize the noise performance at this impedance, a low current noise bipolar input opamp (similar to the Linear Technologies LT1007) is used.

The transimpedance amplifiers are mounted in the sensor assemblies to minimize noise pickup and gain peaking due to wiring capacitance. The LED drivers and remaining signal conditioning are in a single box (Figure 13).

The voltage signal from the transimpedance amplifier is AC coupled through a single-pole 16 Hz highpass filter. This removes the DC bias, 1/f noise and any interference from gradually changing ambient lighting. The signal is then lowpass filtered though a 4 KHz two-pole Butterworth antialias filter. Finally, each signal is digitized with 18 bit resolution at approximately 30K samples/sec.

The bandwidth of the antialias filter must be wide enough to pass all of the modulation frequencies. Furthermore, in order to avoid drift and nonlinearity in the position measurement, it is desirable that the phase and gain of the filters be stable and well-matched across all of the channels. This encourages the use of a low-order, low-Q antialias filter with a corner frequency well above the highest modulation frequency. This is, of course, desirable from a hardware complexity perspective as well, but increases the need for oversampling in the data acquisition.

To maximize matching, the antialias filters were constructed with 1% capacitors and 0.1% resistors. The capacitors for the channel pairs for each axis were matched to 0.1%.

Mismatch of gains between the axes is not so problematic because it can be handled by the calibration procedure. If the gains of the two axis channels differ, then the gain differs from one side of the workspace to the other, which is not calibrated out.

There is some common-mode noise from the bias supply and the sensor itself. Maximizing matching in the channel pairs improves the rejection of this common-mode noise.

4 Signal processing

All of the software for ASAP is implemented in Labview, both for instrument control and signal processing. The major signal processing steps are: software implementation of lock-in amplification, sensor calibration and 3D localization.

4.1 Software implementation of lock-in amplification

Traditionally, demodulation for lock-in amplification has been implemented in

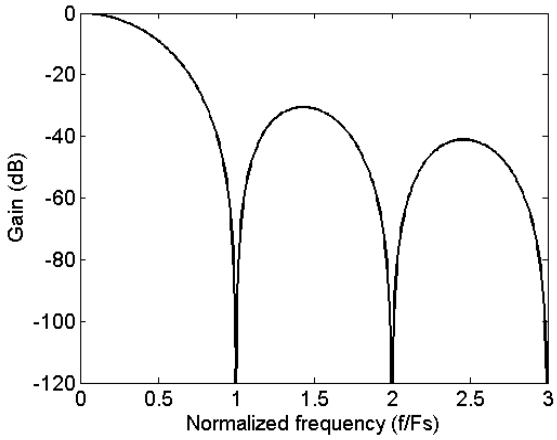

Figure 14: sin(x)/x response

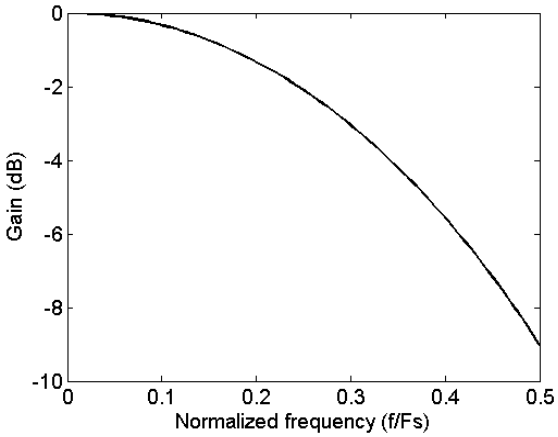

Figure 15: sin(x)/x response detail

hardware, usually with some sort of switching element in place of an analog multiplier. The first version ASAP (which had only a single light) used this approach. Due to the need to demodulate eight channels, this requires a significant amount of hardware, and the thought of further multiplying this hardware two or more times to support frequency domain multiplexing is daunting.

Instead, we opted to move as much of the signal processing into software as possible. Only current-to-voltage conversion and antialias filtering is done in hardware. All the demodulation is done in software. This does increase the demands on the analog-to-digital conversion in terms of conversion rate and dynamic range, and does increase the CPU demands, but given the low modulation frequency (a few kHz), this is well within the capability of standard desktop PCs and data acquisition hardware.

In order to support efficient digital processing, we devised an optimized modulation scheme. The important concept is that the modulation frequencies and the sampling frequency are constrained by integer ratios such that:

1. The input sampling rate is a multiple of the output rate (the *repeat interval*), and
2. The modulation precisely repeats at an interval equal to the output sample rate.

In other words, if the output sample rate is 1 kHz, then an integral number of cycles of all the modulation frequencies must fit in a 1ms interval.

In fact, the generation and digitization of the signals from a common timebase guarantees that the excitation will repeat at some interval, and normally decimation is by an integer ratio, but by forcing these ratios to coincide we allow for particularly efficient demodulation.

Data is processed in blocks that are a multiple of the repeat interval, and a lookup table is precomputed containing the sine reference signals over the repeat interval. Demodulation is then a single multiply/accumulate operation. The data block is multiplied by the reference block, and the result is summed and divided by the block size (computing the mean).

4.1.1 Decimation

When using digital demodulation, the lowpass filter becomes a *decimation* filter, that is, a lowpass filter where the output sample rate is lower than the input sample rate. We must do the analog/digital conversion at a high enough rate to capture all of the signal and noise that passes though the antialias filter with minimal aliasing of the noise, but after demodulation the sample rate only needs to be high enough to capture the measurement bandwidth on a single channel.

Taking the mean over the data block is the decimation filter, creating a sin(x)/x (or *sinc*) frequency response (Figure 14). In addition to the particularly efficient implementation, the sinc response has the very nice property of

perfectly suppressing all of the modulation frequencies and their harmonics. This can be easily seen in the time domain because we are summing over a whole number of cycles of all the modulation frequencies.

Unfortunately, other aspects of the sinc response are not so ideal: slow rolloff in the stopband, substantial droop in the passband (**Error! Reference source not found.**), and only 9 dB attenuation at the Nyquist frequency. The modest attenuation at and above the Nyquist frequency results in a noise bandwidth that is larger than the signal bandwidth, somewhat reducing the signal-to-noise ratio. When the passband droop is a problem, it can be minimized by choosing a higher output rate, then establishing the final noise bandwidth with a standard lowpass filter.

4.1.2 Frequency optimization

How do we choose the modulation frequencies, sampling rate and repeat interval?

- First, we must choose according to design constraints: the repeat interval and sampling rate must correspond to the desired output interval, the sampling rate cannot exceed the data acquisition card hardware limits, and the modulation frequencies must fall within the passband of the antialias filter.
- Less obviously, but quite helpful in reducing the search space, the modulation frequencies and the sample rate must all be generated from a single 10 MHz timebase by integer division.
- Lastly, the frequencies should be chosen to minimize the interference between the channels due to aliasing of the modulation frequencies and intermodulation products.

The task of minimizing interference is complicated by two other implementation decisions: the use of square wave drive to the LEDs and the choice of a low-order antialias filter. There is substantial harmonic content in the square waves, and the antialias filter provides only modest attenuation of the first few harmonics. This means that we must choose the frequencies so that the aliases of the first few harmonics do not fall within the passband of any other modulation frequency.

This is done using a program that generates all of the possible frequency combinations that satisfy the constraints, and then chooses the combination that minimizes the interference predicted by a model of the signal path. The model includes:

- Intermodulation between the channels (generating sum and difference frequencies) due to nonlinearity in the sensor or analog signal processing,
- The response of the antialias filter, and
- The sinc response of the decimation filter.

It turns out that the actual amount of intermodulation is quite low (< 0.1%), so aliasing of the harmonics is the dominant cause of interference between the channels.
With the current parameters for 1000 samples/sec output rate, the interference model predicts 0.3% interference between the channels. This is well less than optical nonidealities such as light from one source reflecting off of the diffuser for the other source.

The resulting frequencies are:

F1 = 2052.55 Hz;F2 = 3078.82 Hz;
Fs = 29761.90 samples/sec;
Repeat interval = 21 samples;
Output rate = 1026.2726 samples/sec.

4.1.3 Phase calibration

The sine lookup table for the demodulator is different for each channel (and frequency). This allows for the compensation of phase shifts from various sources:

- Non-simultaneous sampling of the channels due to multiplexed data acquisition,
- Phase shift in the LED driver and antialias filter, and
- Any channel mismatching.

In truth, due to the normalization that takes place in the PSD position computation (1), small phase errors don't have much effect on position error, especially if they are common to all channels. However, the complexity in accounting for these errors is not great, and the

computation cost is only incurred once when the lookup table is built.

Furthermore, we felt it was a good idea to write a special test application which measured the channel phase errors as well us other system health measures. This allows us to evaluate the correct functioning of the hardware and correct synchronization of the reference signal and data acquisition. Give the information on per-channel phase error it is a simple matter to incorporate this in the lookup table generation.

4.2 Sensor calibration

Once the four signals from each PSD are demodulated, they are converted into two axis positions using the idealized PSD position formula (1).

To maximize accuracy and to test performance we have developed a calibration procedure for ASAP. We use a precision linear stage to move the probe to known positions across the workspace, recording the axis positions at each point. We then use this data to compute two matrices for each sensor: **T** and **K**. In camera calibration terminology, **T** is the *external calibration* and **K** is the *internal calibration.*

T is a 4x4 matrix which represents the position and orientation of the sensor in linear homogenous coordinates. **T** transforms points or vectors in the sensor coordinates into the instrument coordinates.

K is a 4x4 matrix which represents the projection transform done by the lens/sensor combination, including focal length, aspect ratio and lens centering. **K** transforms a 2D sensor spot position [u, v] into the physical location of the point projection in sensor coordinates.

The main work of calibration is done using [8], which is a public domain implementation of the Tsai camera calibration algorithm[9]. We ignore the radial distortion, as this is small for ASAP.

The input needed by the next step (3D localization) is the ray from each sensor to the measured point. The ray is represented as a base point and a unit direction vector. The base point is fixed for each sensor, and is simply the sensor origin transformed into the instrument coordinates:

$$\mathbf{T}\,[0, 0, 0, 1]^T$$

The direction vector is obtained by using K to convert [u, v] to a physical vector, normalizing the result, and then using **T** to convert the vector into instrument coordinates:

$$\mathbf{T} \left| \mathbf{K}\,[u, v, 0, 1]^T \right|$$

4.3 3D localization

Given the rays from each sensor to each light, we find the 3D light position by intersecting the two rays. The problem is over-constrained since we have 4 measurements and 3 unknowns, so the intersection will be approximate. We use the midpoint between the two rays at their point of closest approach.

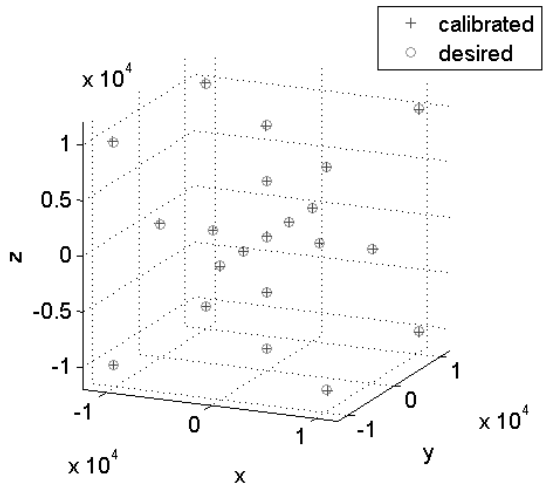

Figure 16: Translation error (microns)

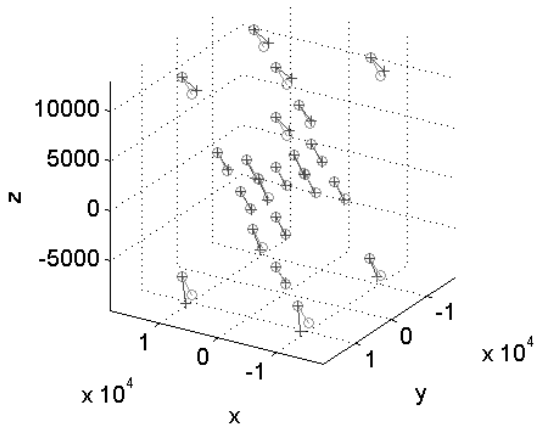

Figure 17: Angular error during translation (microns, 30x angular exaggeration)

The distance of closest approach (ideally 0) is an indicator of measurement quality. The measurement is considered invalid if the match error is greater than 500 microns, the signal amplitude is too low, or the light is too close to the edge of the sensor.

It is possible to compute the 3D position in real-time in Labview code using the geometric intersection formulas, but we were able to significantly reduce the runtime by precomputing a four-dimensional lookup table and computing the 3D position using multi-linear interpolation.

The two sensor Z axis measurements are almost identical, so we were able to significantly reduce the table size by making the last index be the difference between the two Z axis measurements. The lookup table is then 41 x 41 x 41 x 5, or 344605 entries.

5 Performance

Performance of ASAP has not been characterized completely, but sufficient testing has been done to verify that the performance meets the requirements.

5.1 Accuracy

The sensor and signal processing paths have good stability, but absolute accuracy and time stability have not been characterized as yet, as we are primarily concerned with relative motion at frequencies above 0.1 Hz. Relative motion accuracy requires only linearity and scale factor stability.

It is in any case difficult to determine whether micron-level drifts are due to the instrument or due to changes in the test fixture due to room temperature changes, etc. Consider that, when averaged over periods longer than a few milliseconds, the position resolution is finer than the wavelength of the IR light used.

For these reasons, we tested the relative accuracy by first applying a 6-DOF motion that minimizes the matching error.

Figure 16 shows the position error when the probe undergoes pure translation. These data were collected during a calibration procedure (as described above), but are from the other light that was *not* used to generate the calibration. Error is less than 2% of translation, and is hardly visible at this scale.

There is cross-coupling between rotation and translation due to sensor nonlinearity and the non-uniform illumination or partial occlusion of the LED diffusers.

Figure 17 shows the angular error as the probe undergoes pure translation. The peak error is less than 2°, and is exaggerated 30x so that it is

visible. Limited testing of the erroneous translation caused by pure rotation has shown less than 10 microns per degree of rotation.

5.2 Resolution and Speed

The output noise is white, so there is a noise-to-bandwidth tradeoff, with the position noise proportional to the square root of bandwidth. At 1000 Hz the noise is
0.5 microns RMS (3 microns peak-to-peak).

5.3 Workspace size

The workspace is approximately a 3 cm cube. Rotation is limited by the constraint that both LEDs must be visible in both sensors, with neither being obscured by the other or by the LED body and tool tip.

5.4 Interference susceptibility

The system works well under normal room lighting, but high levels of infrared interference such as from bright incandescent lights or IR emitters such as laptop IR ports can cause significant interference.

For best measurement stability, the sensors should view the LEDs against a non-reflective background, as secondary reflections off of clothing, etc., can shift the measurement by tens of microns.

6 Conclusions

We have demonstrated a measurement system for micron-scale 5 degree of freedom position measurement with a 1000 sample per second output rate. No other system we are aware of has this combination of specifications. Similar systems may prove useful for applications outside of the area of hand tremor measurement and cancellation.

We have found that frequency-domain multiplexing can be effectively used to simultaneously track multiple sources using a PSD without significantly compromising the signal-to-noise ratio.

Synchronous demodulation of PSD signals can easily be implemented in software. The advantage of the reduction in hardware complexity becomes especially compelling when more than one signal is being demodulated. The old single-channel demodulator required 960 cm^2 of board space, whereas the new signal conditioning needs only 250 cm^2, despite supporting an arbitrary number of channels.

More generally, our experience shows that PSDs are an attractive sensor for high-bandwidth position measurement because quite simple hardware can be used to implement a high-performance system.

References

[1] F. Peral-Gutierrez, A.L. Liao, and C. Riviere, "Static and dynamic accuracy of vitreoretinal surgeons," Proc. 26th Annu. Int. Conf. IEEE Eng. Med. Biol. Soc., vol. 4, 2004, pp. 2734-2737.

[2] C. N. Riviere, J. Gangloff, and M. de Mathelin, "Robotic compensation of biological motion to enhance surgical accuracy," *Proc. IEEE*, 94(9):1705-1716, 2006.

[3] C. N. Riviere, W. T. Ang, and P. K. Khosla, "Toward active tremor canceling in handheld microsurgical instruments," *IEEE Trans. Rob. Autom.*, 19(5):793-800, 2003.

[4] W. Ma, A. Rizzi, and R. Hollis, "Optical Coordination Sensor for Precision Cooperating Robots," Proc. IEEE Int. Conf. Rob. Autom., vol. 2, April 2000, pp. 1621-1626.

[5] L. Hotraphinyo, "Precision optical tracking of microsurgical instruments for performance evaluation," master's thesis, Dept. of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, Pa., December 2001.

[6] G. Welch, G. Bishop, L. Vicci, *et al*., "The HiBall tracker: high-performance wide-area tracking for virtual and augmented environments," Proc. ACM Symp. Virtual Real. Softw. Technol., London, December 1999, Addison-Wesley, pp. 1-11.

[7] A. Makynen, *Position-sensitive devices and sensor systems for optical tracking and displacement sensing applications*, doctoral

dissertation, University of Oulu, Finland, September 2000.
[8] Reg Willson, Tsai Camera Calibration C code, [http://www.cs.cmu.edu/~rgw/TsaiCode.html].
[9] R. Y. Tsai , "A versatile camera calibration technique for high-accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses", *IEEE J. Rob. Autom.*, vol. RA-3, no. 4, pp.323-344, August 1987.