Common Questions Answered

- Q: Does the D-1000 "work"
- A: YES

- Q: Does the D-1000 measure glucose
- A: YES

- Q: Does the D-1000 predict blood glucose
- A: YES (as an averaging device)

- Q: Do the averaged predictions correlate well with blood glucose
- A: YES

- Q: Does the D-1000 measure blood glucose
- A1: YES analytically
- A2: NO optically

- Q: Is it possible to measure blood glucose optically
- A: YES

Who Says It's Not Blood?

- High SNR singular values indicate ~ 200 μAU glucose signal
- Optical penetration depths are 1mm on average going as deep as 3mm
- Capillaries are ~ 15 μm in diameter and "web-like"
- This means that blood is about .5% of the light path or less
- Glucose signal contribution from blood then is .5% of 200 μAU or less (< 1μAU)
- D-1000 noise floor is ~ 5 μAU
- This means that we effectively optically measure skin glucose only

How then can we know if we're seeing blood glucose or not?

We know from Monte-Carlo simulations and from experimental determination that we have a mean penetration depth of about 1 mm into the skin with maximum depths as deep as 3 mm. We also know that capillaries are about 15 μm in diameter on average and have a web-like cross sectional layout as opposed to a blanket-like layout. This means that we can expect blood to be about 0.5% of our light path or less. If from our SNR rank determination we see that glucose has about 200 μAU of signal from our light path, then the contribution from blood glucose is about 0.5% of 200 μAU or 1 μAU. The D-1000 has a noise floor of about 5 μAU. Therefore we can safely conclude that the D-1000 does not "see" blood glucose and only "sees" skin glucose.

So … now we have to increase the D-1000 capabilities.

How do we do that?

First, let's make a brief note regarding blood glucose NIR spectroscopy. Experimentation performed by Ralf M. before he came here and other research groups since have interesting results. These research studies attempted to build calibrations using blood extracted from patients and measured in a spectroscopic device. Not only were they able to build successful calibrations on the blood samples but the calibrations were built on blood from multiple patients! This hints of the brass ring that we've titled the "universal calibration." The results of these studies seem to suggest that human blood has enough in common between people that the differences in individual's blood are not significant enough cause error in a glucose calibration.

The "Time-Delay" Problem

- If not blood, then what glucose are we measuring optically then?
- We measure interstitial or skin glucose
- Since capillaries, which contain blood, are the suppliers for the skin's glucose, skin and blood glucose are forced to be closely related (correlated)
- Skin and blood glucose concentrations have extremely similar time profiles
- The skin time profile tends to randomly lead or lag in time w.r.t. the blood
- What are the effects of this time delay?
  - The finger prick measurement is a "noisy" skin glucose reference (lower slope)
  - Skin glucose predictions are noisier
  - Instantaneous skin measurements may differ significantly from blood
  - The lead/lag randomness allows averaging to eliminate/reduce the prediction noise

Capabilities and limitations

The best starting point for this seminar is an overview of the problem commonly called the "time-delay" problem. As the answers to the common questions show, we do measure glucose, but we don't measure blood glucose. What glucose are we measuring then? We measure interstitial glucose, commonly referred to as skin glucose. The capillaries, which contain blood and carry it close to the skin surface, are the suppliers for the skin's glucose.

It this good or bad?
For us, it turns out to be a good thing, as long as we're aware of the problem. Here's how. We're measuring two signals; one is the glucose spectrum in the skin and the other is blood glucose reference from the finger prick monitor. Let's first assume that the blood and skin glucose levels are always the same or at least proportional (say a 2 to 1 ratio all the time). If this were the case, we'd do a least squares fit to the spectra and blood glucose reference values. When we'd measure a new skin spectrum, we'd apply the fit constants to the spectrum and come up with a glucose concentration that would be the blood glucose concentration. But in the real world situation, at the time of measurement, the skin and blood glucose levels are not always the same. Actually, the only time that they should be the same is after a long period of constant blood glucose concentration. The rest of the time, skin and blood glucose values can differ by as much as 90 mg/dL at any given instant!
Here's how this effects our least-squares approach. First, the blood glucose concentration is not the proper reference for each spectrum that we measure in calibration. This translates into having a noisy reference; reference meaning glucose reference. In terms of error grid analysis, this means a lower slope. For practical purposes, this means that we have more noisy SKIN glucose predictions. I say SKIN because that is the spectrum that we are measuring and not blood. The second effect is seen in prediction. Every time we measure a spectrum and calculate a glucose concentration, the glucose concentration we calculate is for SKIN. But we already know that skin and blood can have different concentrations at any given instant, so we see an instantaneous "bias" also.

Since the capillaries are the source for the skin's glucose, the skin and blood glucose concentrations have no choice but to be closely related. Actually, if you track both with time, the skin and blood glucose concentrations have the same profile, but one usually leads in front of or lags behind the other. This leading and lagging effect is not constant. Some times the blood will have higher concentrations and sometimes lower. This is what everyone is referring to as the time-delay problem.

[Hump Day and Cygnus plots go here]

To make use of our predictions, we average out the leading and lagging discrepancies. This is effectively done over a two-week time frame. So now, after averaging the skin glucose concentrations and blood glucose concentrations over a two-week period, the results are the same. This is how the D-1000 works.

Singular Value Decomposition (SVD)

- We measure 43 independent variables
- This allows for measuring 43 independent directions of variation or principal component (PC)
- Each PC has a 43 element spectrum associated with it
- These PC's can be ordered from the strongest to the weakest variation with a single value to represent the strength of variation

Okay ... so we say we see skin glucose, can we prove it? How do we know that we don't see blood glucose?

For this, we need to get a little technical. When we measure a spectrum, we measure 64 different wavelengths. Of these, we use 43. This means that we have 43 independent variables or a 43 dimensional space. What does that mean? Having 43 independent variables means that we can break up a set of spectra into 43 different independent variations. Each variation would have a 43 pixel-wavelength spectrum associated with it and all 43 variations are unrelated. These 43 variations can be ordered from the strongest variation to the weakest variation and be assigned a number representing the strength of the variation. This is a mathematical procedure called singular value decomposition. The "singular value" comes from the single number assigned to each variation and the "decomposition" comes from decomposing the data set into independent directions of variation.

SVD Proves We "See" Glucose

- Make predictions on each PC
- Determine how much of the predictions are related to glucose references (call this glucose signal)
- Determine how much of the predictions are not related to glucose references (call this glucose noise)
- Determine if any PC's have significantly higher SNR than others
- Determine if the high SNR PC's have realistic singular values
- Check to see if the spectral features of the high SNR PC's correspond to any glucose features

So how do we use SVD to prove that we see glucose? Without going into too much mathematical detail, here's how it's done:

Take each direction (called a principal component or PC or rank) and make predictions using coefficients from our least-squares fit (called the b-vector).

Determine how much of those predictions are related to our glucose reference values (this we'll call glucose signal)

Determine how much of those predictions are not related to out glucose reference values (this we'll call glucose noise)

Calculate the glucose signal to noise ratios.

Determine if any ranks have significantly higher SNR than others

Determine if the high SNR ranks have a realistic singular value for glucose

Determine if the spectral features of the PC correspond to any glucose features

Realistic Singular Value Criteria

- From theory and experimentation we know that
  - 100 mg/dL change = 160 μAU change @1600nm
- There can be multiple glucose peaks in a PC (at least 3 large peaks)
- Total glucose signal in a data set is the standard deviation of the Hemocues
- Realistic singular values then must be less than
  - 3 * STD(Hemocues) *160 μAU / 100 mg/dL

From theoretical and experimental calculations, we know that 100 mg/dL of glucose change corresponds to 160 μAU of absorbance change at 1600 nm. If the data that we're looking at is calibration data, then we can approximate what the singular values should be for a rank to actually be a glucose bearing rank. To do this, the standard deviation of the glucose reference values for the calibration time period is multiplied by 160 μAU. This is then multiplied by a factor of 2 or 3 because the value of 160 μAU is based only on the 1600 nm peak and there are other peaks in the spectra that we measure. This number is used to determine if high glucose SNR ranks are really glucose or if false correlations have been built into the calibration. After performing this technique on a number of patients, we see that when a good calibration is built, the high glucose SNR ranks fall within the singular value range that we expect. The spectra of these ranks typically reveal strong features centered on one or more glucose peak wavelengths. This acts as our proof of the D-1000's ability to "see" glucose.

[SNR and SVD plots go here]

Increasing Capabilities Beyond the D-1000

- Approach universal calibration by using time-resolved spectroscopy to measure blood-only spectra
  - Requires advancements in electronics, or
  - Advancements in optical signal filtering
  - Requires stronger signal
- Optically measure blood glucose by increasing SNR by minimum factor of 200
  - Improve detection system to reduce noise (SNR increase by factor 2 or less)
  - Increase total light signal
    - Currently @ .25 mW/nm using 30W bulb
    - Need 50 mW/nm
    - Use 6000 W bulb ?!?!
    - Delta-1 concept

So to increase the D-1000 capabilities, here's what we'd like to have. We'd first like to be able to increase the total light signal or blood path to the point that we would have at least 200 μAU of BLOOD glucose signal. And, we'd like to have a method for measuring only blood spectra.

The measurement of the blood spectra was described in the tech report. As it turns out, the devices currently do not exist for time resolved spectroscopy of the capillaries at this time, but there have been some suggestions since the release of the tech report that pose some possibilities of side-stepping the technological deficiencies. The key interest in being able to perform the time resolved spectroscopy is that a "universal" calibration may be possible.

By simply increasing the total signal by roughly a factor of say 200, we can build a calibration that predicts blood glucose directly even though each device may still have to be calibrated for individual patients. This factor of 200 roughly corresponds to a light source with a minimum output of 50 mW/nm. At first glance, one might want to replace the 30W bulb with a 6000W bulb. Besides requiring a small combustion engine to supply the power, this would end up frying any fleshy material that would get near the focal point of the illumination system. So how do we go from .25 mW/nm to 50 mW/nm without replacing the arm-tray with a bar-b-que grill and drip pan? We take the delta-1 design to the next conceptual level.

Delta-1 Concept

- Monochromatic or single wavelength system
- Two light sources for each wavelength
- One detector for light source reference
- One detector for skin signal
- Calculate absorbance using detector signals from each source to reference out source noise and light path
- Pros:
  - Extremely stable (RTB 5 μAU or less while being beaten by a hamme)
  - Rapid warm-up time
  - Low power consumption
- Cons:
  - Requires 17 different LED designs (presently not available)
  - Requires spectrograph
  - Insufficient light levels to optically measure blood glucose

The delta-1 is a monochromatic or single wavelength system. The idea is as follows:

Have two light sources for each wavelength
Have one detector for source reference

Have one detector for skin signal

Calculate absorbance using detector signals from each source to reference out source noise and the light path.

[Delta-1 design slide goes here]

Improving the Delta-1

- Find single material light source
  - Require light source to be narrow band (3 nm FWHM or less)
  - Require light source to be 50 mW/nm or more
  - Require light source to be minimal power consumer
  - Require light source to allow for rapid warm-up
- Insure that all wavelengths are narrow band to eliminate spectrograph

The delta one was designed with the use of LED's in mind. This forced the need for a spectrograph to make use of the broad wavelength output from each LED. The spectrograph would act as a wavelength filter on the LED to select the three wavelength bands from each LED (ie, each LED would be a source of three wavelengths of light). Then, there would be 17 different LED type, each one covering a different wavelength range for a total of 51 different wavelengths. This is a sound design for a spectrometer. The reference test bed (RTB), a mock up of a delta-1 system using silicon LED's and detectors instead of the InGaAs ones, was built and proven to operate at the 5μAU noise level or lower under even the severest condition of being beaten with a hammer while spectra were being recorded. Some of the drawbacks to the delta-1 design were the need for a spectrograph, the need for 17 different LED's, and the unavailability of InGaAs LED's. Besides these drawbacks, the delta-1 design as proposed is still at least and order of magnitude or more below the required power levels to be able to resolve blood glucose.

So how does the delta-1 help solve the blood glucose problem? The conceptual design is based on monochromatic light. It is also based on semi-conductor light sources (solid state). The solid state light source concept allows the device to be shut off and turned on with out degradation of the device performance since the source noise is referenced out. What is needed for the delta-1 design to be a success is a solid state monochromatic light source with output power 50 mW/nm. The device that meets these requirements is called the transversely pumped counter-propagating optical parametric oscillator (TCPOPO).

Optical Parametric Oscillation (OPO)

- Tunable laser light source
- Requires mirrored feedback cavity (low efficiency)
- Requires non-linear optical crystal
- Requires high-power mode-locked laser for pump
- Typically large laboratory instrument
- Output "tuned" by rotating crystal

Before TCPOPO is discussed, we first have to describe OPO. Optical parametric oscillation is achieved by making use of a property of some crystals called dielectric susceptibility. In some crystals, this property is such that a light of the proper wavelength and phase will experience amplitude amplification. By the same property, that light can be split into 2 different wavelengths. These crystals are called nonlinear optical crystals. Common examples are KTP and lithium niobate. With these crystals, setting up an OPO cavity that consists of 2 mirrors and the crystal can create a tunable laser light source. A laser pumps the cavity and the crystal. As the crystal is rotated, the direction of the dielectric susceptibility axis changes and the pump laser wavelength generates two other laser wavelengths based on the conservation of momentum. These devices are available as off the shelf items from laser manufacturers, but are costly and delicate. A typical OPO would occupy a 4' x 4' area. This clearly is not practical for use with the delta-1, but the tunable light source application is.

[OPO cavity design slide and OPO non-linear amplification slide go here]

Transversely Pumped Counter-propagating OPO (TPCOPO)

- Semiconductor version of an OPO
- Non-linear crystal and cavity replaced by alternating semiconductor layers and Bragg reflectors
- Can be pumped with as little as 100 μW
- Output tuned by input angle of pump
- Up to 25% efficient
- Tunable range depends on material and pump wavelength
- Can be designed for any range
- Developed by Dr. Yujie Ding

So, TCPOPO is the semiconductor version of an OPO. My thesis advisor, Dr. Yujie Ding, developed this technique, and was here last summer to give a presentation regarding it. Instead of using conventional non-linear crystals, semiconductor material is used. Instead of placing the crystal in a mirror cavity, it is grown between two Bragg reflectors that are matched for a designed pump wavelength. The key difference between TCPOPO and OPO besides being semiconductor based is that TCPOPO has output that comes out perpendicular to the pump light direction and is tuned depending on the angle of incidence of the pump.

[TCPOPO design goes here]

TPCOPO to Improve Delta-1

- TPCOPO can cover 1200 nm - 1850 nm with 260 individual wavelength bands
- Requires only one TPCOPO design (not 17) and one pump diode design
- 25% efficiency @ 2nm FWHM requires pump of 400 mW to achieve signals for resolving blood glucose
- Each element would have 50% duty cycle yielding a 200 mW light source
- Arrays could be etched using lithographic techniques
- Laser diodes could be mounted on each element at angles corresponding to required wavelengths
- One array would be about 5.2mm x 100μm x 100μm
- One 3 inch wafer could yield about 8770 arrays

How can this be used in the delta-1? First of all, one TCPOPO design can be created for tunability over our entire wavelength range. We can take the device and etch it to make an array. Then, laser diodes can be mounted on each element of the array at different angels corresponding to different wavelengths of output.

The design that we would use is approximately 25% efficient meaning that a 200 mW pump diode would give us the 50 mW of output power that we need. Since wavelength is determined by input angle, only one TCPOPO design needs to be constructed as opposed to the 17 different LED designs of the delta-1. Also, wavelength "line-width" for a TCPOPO is about 1 to 2 nm. This means that a spectrograph is not needed to filter wavelength.

[TCPOPO spec sheet and array design go here]

Possible Grant Funding for TPCOPO

- All TPCOPO designs mentioned here can be referenced in scientific journals
- TPCOPO can not be patented because of publications
- Dr. Ding's research is mainly DoD funded
- TPCOPO has many DoD applications especially in communications
- Characterization has to be completed before we can utilize the technology

  (about 2 - 3 years away)

The obstacle that currently stands in our way with using the TCPOPO is that they are still is the early development stages. This means that the theory is developed and some have been manufactured as R&D devices, but complete characterization has not been done yet. Dr. Ding has published everything that I described here about the TCPOPO except for its usage in the delta-1. This means that the technology can not be patented, but grants can be received to further and complete the work. Dr. Ding has mentioned that he'd be willing to act as a consultant for us regarding the TCPOPO development, but nothing has been formally done in this regard.

In Summary

- The D-1000 does optically measure glucose with a precision of about 35mg/dL RMS
- The D-1000 can only measure blood analytically at present, not optically
- Solutions do exist for optically measuring blood glucose
  - Incandescent light systems are not directly among the solutions
- TPCOPO device design would allow optical measurement of blood glucose at signal levels required to show performance equivalence to invasive devices
- TPCOPO device design could be pocket calculator size
- TPCOPO in conjunction with time-resolved techniques could allow for universal calibrations of glucose and other blood components

In summary:

The d-1000 does see glucose (roughly 35mg/dL RMS)

It does not see blood glucose

Seeing blood glucose is possible

It would be prohibitive to create an incandescent light based system to see blood glucose

TCPOPO device design would allow us to reasonably see blood glucose spectral signals at the level to required to be device equivalent to finger prick devices

TCPOPO device design could be made to pocket calculator size.