RamanFlex
School of Engineering Science
Burnaby, BC
V5A 1S6

October 31, 2002

Dr. Andrew Rawicz
School of Engineering Science
Simon Fraser University
Burnaby, British Columbia
V5A 1S6

Dear Dr. Rawicz,

Enclosed are the design specifications for a Gas Analyzer based on measurements of the Raman effect. This document presents, in detail, the design for a device which fulfills the functional requirements projected for the prototype to be completed in December, 2002. Unlike the previous documents, which targeted commercial development as well as laboratory prototyping, we have focused exclusively on the completion of the short-term prototype.

In addition to the design itself, we have presented an analysis of the signal strengths as the Raman effect is created, measured, and quantified. This is by no means intended to be a complete analysis; we merely wished to obtain a figure which would help us evaluate the feasibility of continuing with this project. Our results, as you will find, are favourable.

If there are any questions, concerns, or comments, do not hestitate to contact any member of the design team. Our contact information has been included in the first page of our proposal.

Sincerely,

Graeme Smecher
Project Lead
RamanFlex

RamanFlex

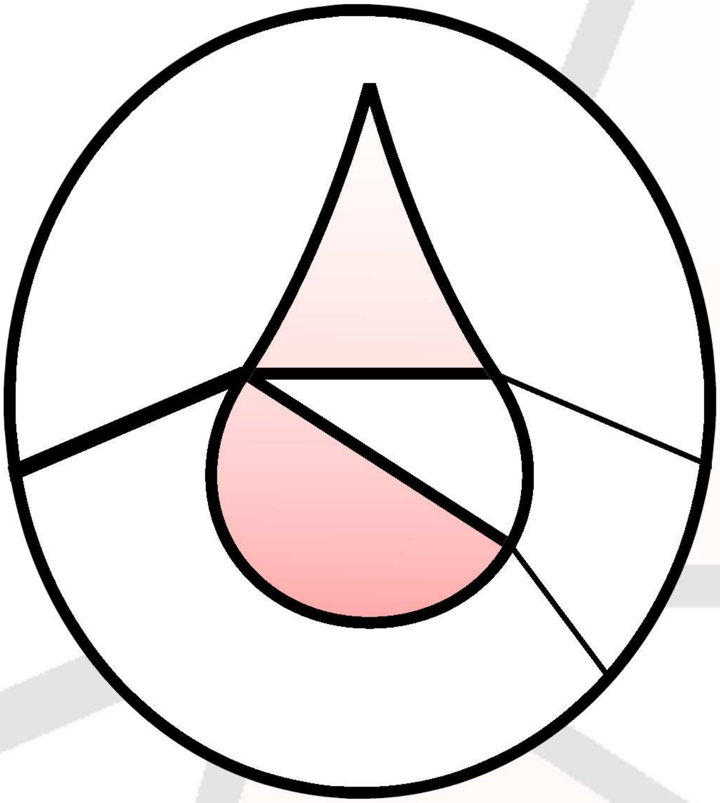

Design Specifications

Simon Laalo
Jon Jolivet
Graeme Smecher
Bernard Smit

The Fine Print

This document reflects the plan for a four-month project, and is accurate to the best of our abilities. The project is, however, a moving target. Therefore, this document is subject to revision.

The RamanFlex team may be contacted *en masse* at `ensc340-napiform@sfu.ca`. Alternately, our individual contact information is as follows:

name	**e-mail**	**home**	**cel / pager**
Simon Laalo	`sjlaalo@sfu.ca`	604-942-9935	
Jon Jolivet	`jjolivet@sfu.ca`	604-444-3431	604-473-1367
Graeme Smecher	`gsmecher@sfu.ca`	604-324-5055	
Bernard Smit	`bsmit@sfu.ca`	604-945-7155	604-837-9468

This is Version 1 of this document, prepared on October 31, 2002. Please check for newer versions with one of the above contacts.

Executive Summary

Since its introduction to the recreational market in the 1970s, SCUBA diving has enjoyed a steady increase in popularity. While in ordinary circumstances it is a safe sport, it can become extremely dangerous if the compressed air a diver uses is unsafe.

An increasingly large contingent of technically inclined divers have started mixing their own breathing gas. By modifying the balance of nitrogen or oxygen in their air, or by introducing foreign gases such as helium, divers can alter the way in which their bodies absorb the air they breathe. This practice, known as *mixed gas diving*, is extremely difficult to perform safely. If the balance of air is not strictly controlled, or if the concentrations are misjudged, divers become susceptible to a wide variety of fatal complications.

Our project consists of a gas analyzer which will initially measure the concentrations of Nitrogen and Oxygen in an air sample. For the purposes of ENSC 305/340, we intend to be able to measure gas concentrations to a precision of 10%. Eventually, in addition to measuring standard breathing gases, we hope to measure pollutants such as CO or hydrocarbon exhaust gases with much greater precision.

This design specification describes the technical details involved in the construction of a Raman spectrometer suitable for breathing gas analysis. The primary goal of the design detailed here is to meet the requirements addressed in the Functional Specifications document.

In the prior documents, we separated the project goals into a number of phases. The first of these phases was tailored for completion at the end of this December. For the design specifications, it no longer makes sense to perform a detailed analysis of subsequent phases, provided that we bear in mind the functional specifications for continued product development and do not preclude any of them. The purpose of this document is not to flush out a commercial prototype design; instead, we will focus almost exclusively on the scope of requirements for the first phase.

CONTENTS

List of Figures

List of Tables

1 INTRODUCTION

This design specification presents an implementation of a device which fulfills the functional contract laid out in the functional specification. Because it is technical in content, it is primarily designed for evaluation by an engineer or other technical reader. However, a metric of the success of this document is its comprehendibility by less technical readers; because it is to be used to demonstrate the worthiness of our project for further exploration, budget allocation, etc., it should remain broadly approachable.

In the functional specifications, some background knowledge was presented introducing the Raman effect. A more thorough description is presented here; however, for a basic description of the effect, readers are directed to the functional specifications. In addition to a more thorough description of the Raman effect, we present a mathematical model which gives a first-order approximation of the intensity of the effect we must measure.

2 DOCUMENT CONVENTIONS

There are a number of bookkeeping issues which must be addressed before the designs can be presented. The following subsections detail conventions and acronyms, and list references which were invaluable in the production of this document.

2.1 REQUIREMENT CROSS-REFERENCING

This document is intended to directly follow the functional specification produced earlier in our product design cycle[1]. The two documents are intended to be cross-referenced against one another to ensure that the requirements introduced in the functional specifications document are addressed here. However, a direct coorespondance between section headings would be clumsy at best; this document should naturally be organized around the subsystems of our product, while the prior document was organized around requirement categories such as environmental conditions, typical operating requirements, etc.

To maintain ease of reference without hampering flow, we will maintain the requirement reference numbers used in the previous document. Whenever a section concerns a particular requirement, it will be stated clearly. Therefore, this document should reference each of the applicable functional requirements at least once.

---

[1] The current release of our functional specifications document is entitled, “RamanFlex Functional Specifications”, revision 1.)

2.2 Acronyms

A large number of acronyms are referenced in this document. Because some of them may not be familiar to the technical reader, a brief definition of each within the context of this project is provided here.

AFE Analog Front-End. AFEs provide some of the complex analog-to-digital conversion functions required by CCDs.

CCD Charge Coupled Device. CCD sensors are digital image sensors, and are typically at the heart of digital cameras, scanners, and many other devices. They are incredibly sensitive. The concept of exposure time from regular photographic film carries over to CCD devices; to measure dim signals, a CCD can simply be exposed for a longer period of time.

CPLD Complex Programmable Logic Device. For our practical purposes, the description of FPGA is accurate for CPLDs.

DIP Dual In-Line Package. This is an aging standard package for integrated circuits.

FPGA Field-Programmable Gate Array. FPGAs are programmable logic devices; they can be reconfigured to perform virtually any logic function of their input and output pins. They present an extremely attractive platform for digital logic prototyping, because they are flexible and easy to use.

MOSFET Metal-Oxide Semiconductor Field-Effect Transistors. These transistors are at the heart of most digital logic designs, and form the basis for CCD image sensors.

RS-232 The standard PC serial interface. While this interface is being gradually retired in favor of USB or more modern standards, its ubiquity and simplicity make it appropriate for many projects.

SCBA / SCUBA Self Contained Breathing Apparatus / Self Contained Underwater Breathing Apparatus. This acronym refers to the system of regulators, hoses, and tanks that allow divers, firemen, and so forth to breathe compressed air.

2.3 References

Mcreery, Richard L. 2000. *Raman Spectroscopy for Chemical Analysis.* New York: Wiley-Interscience

Yamamoto, Hiroshi and Ozaki, Yukihiro. Highly Sensitive Gas Anslysis by Raman Spectroscopy and its Application. Review of Laser Engineering, 25:10, 1997.

2.4 DATASHEETS

Toshiba TCD1205D `http://www.toshiba.com/taec/components/Datasheet/TCD1205D.pdf`

National Semiconductor LM9823 http://www.national.com/ds/LM/LM9823.pdf

3 THE RAMAN EFFECT

The Raman effect is central to the operation of our device, and therefore we have included an overview of its foundations for the convenience of the reader.

The interactions between the molecules of a sample and photons from a light source form the basis for all analytical spectrometry techniques. This molecule-photon interaction manifests itself in a number of ways: Rayliegh scattering, Raman scattering, and absorption. Due to the diffuse nature of gasses, the majority of the photons merely pass through without any interaction with the molecules of the gas.

Rayleigh scattering is the result of an interaction of a molecule and a photon when the wavelength of the incident light is longer than the length of the molecule. Scattering may occur in any direction, but there is greater probability of scatterring occuring at 90^o to the angle of incidence. Very few photons interact with the sample molecules, so the intensity of the scattered Rayleigh beam is extremely small in comparison to the original incident light intensity.

During an interaction, some energy from the photon may be transferred to mechanical vibration of the molecule. Occasionally, the molecule will be in an excited vibrational state, and the interacting photon will actually gain energy from the molecule. Any loss or gain in the energy of a photon corresponds to a shift in its frequency. Positive shifts, which result from the photon absorbing energy from the molecule, are referred to as 'anti-Stokes' Raman scattering. Negative shifts, which occur when the photon loses energy to the molecule, are referred to as 'Stokes' Raman scattering.

Mechanical vibrations in a molecule effects the polarizability α. Ignoring quantum mechanical affect we assume the incident light can be expressed as: (McReery, 2000)

$$E = E_o cos(2\pi\nu_o t) \tag{1}$$

The polarisability is given by $P = \alpha E$. Usually, the molecular vibrations can be considered to be composed of normal modes:

$$Q_j = Q_j^o cos(2\pi\nu_j t) \tag{2}$$

The mechanical vibrations cause a modulation in the polarisability given by:

$$\alpha = \alpha_o + \left(\frac{\delta\alpha}{\delta Q_j} Q_j\right) + \cdots \tag{3}$$

Giving the following description for the polarisation in the molecule, and hence, a description of the Raman scattered light:

$$P = \alpha_o E_o cos(2\pi\nu_o t) + E_o Q_j^o \left(\frac{\delta\alpha}{\delta Q_j}\right) \frac{cos(2\pi(\nu_o + \nu_j)t) + cos(2\pi(\nu_o - \nu j)t)}{2} \tag{4}$$

Very sensitive light receptors are required to observe these frequency shifts. The intensity of the Raman scattering is very weak, with only 1 in 10^{10} photons exhibiting any form of Raman scattering (resulting in a 10^{10} signal attenuation), and even fewer photons exhibiting anti-Stokes Raman scattering.

The frequency shifts of Raman scattering depend directly on the material properties — particularly the molecular structure. Carbon-oxygen bonds produce a different Raman shift than carbon-carbon bonds, for example. Through tedious and exact scientific experimentation, the unique Raman spectra for various substances have been determined. This uniqueness of spectra allows a Raman spectrometer to identify an unknown sample purely from its Raman Spectrum. The Raman spectrum of air is presented in Figure 1.

The scattering efficiencies of oxygen and nitrogen are approximately the same. The obvious difference in intensities from each results from the differences in their concentrations. Therefore, by measuring the differences in the Raman intensities of these two gases, we may determine their relative concentrations.

4 Hardware Design

Our design consists of three systems: optics, electronics, and software. We require optics to generate the Raman effect, electronics to measure it, and software to analyze it.

The following sections present, in detail, the design for our Raman spectrometer. They describe both the design features required to address each subsystem's functional require-

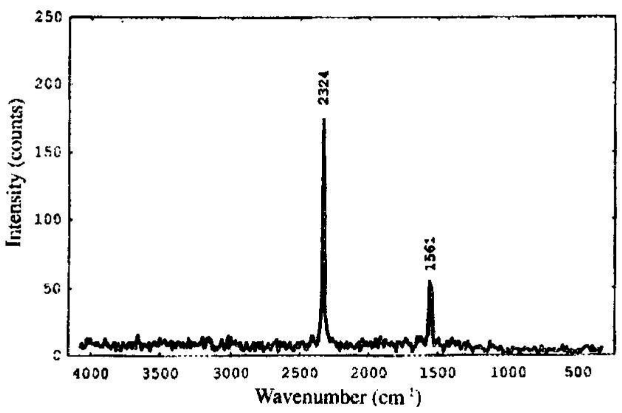

Figure 1: Measured Raman spectrum of air (Yamamoto and Ozaki, 1997)

ments, as well as any device parameters required to ensure interoperation between subsystems. The design is presented in the same order as logical signal flow; therefore, an estimate of signal intensity is developed as the systems are explained.

4.1 OPTICS

The optical components perform two functions. First, they must produce a measurable Raman effect; second, they must present this effect in a manner appropriate for electronic measurement.

There are five optical subsystems in total; the following sections explain their function and design in detail.

4.1.1 Laser

There are two excellent motivations for the use of a laser to stimulate Raman scattering. They are intensity and monochromaticity.

Firstly, the Raman effect is extremely faint — in general, the intensity of Raman scattered light is on the order of 10^{-10} times the intensity of the exciting light. Therefore, an extremely intense light source is vital to produce a measurable Raman effect.

Secondly, the Raman effect manifests itself in a shift in the colour of scattered light. Therefore, in order to measure a Raman shift for a particular gas, both the incident beam and the Raman light must be of known wavelengths. As soon as more than one incident beam is present, it becomes difficult to determine which Raman signal cooresponds to which incident beam.

We have selected a laser with an output power of approximately 10 mW at a wavelength of 650 nm. While we have a more powerful 300 mW laser at our disposal, we have opted for the less powerful laser. The weaker laser's operating characteristics more closely match those of a laser which could be used in a commercial product. Due to their low cost and ease of use, a semiconductor laser would be more appropriate. We have been unable to source a semiconductor laser for suitably low cost; because our 10 mW laser operates at the same frequency and provides a similar output power, it provides a suitable replacement for prototyping purposes.

The output from this section is, therefore, a 10 mW laser beam at 650 nm.

4.1.2 Laser Filtration

Lasers, unfortunately, are not perfect. The spectral purity of any laser is not ideal; the laser's output, ideally consisting of a single, stable wavelength, actually has a number of artifacts. The primary laser line does not occupy a single frequency but actually has a Gaussian-like profile. The center of this profile wanders up and down with variations in temperature.

Additionally, most lasers exhibit multimode or mode-hopping behavior, or both. For all practical purposes, lasers have more than one high-amplitude output frequency.

Hopefully, these effects are of sufficiently low order that they do not produce erroneous measurements. Should careful signal analysis not adequately resolve these problems, the laser might require prefiltering. We have a monochromator at our disposal; this device is extremely large, finicky and expensive, and we are reluctant to use it since its inclusion is not suitable for a commercial device.

For the purposes of ENSC 305/340, identification of N_2 and O_2 spectra is sufficient; unless laser noise overpowers the sensor or otherwise prevent the gathering of data from these two gases, the laser will not present difficulty for our exploratory prototype. Therefore, we will not include conditioning optics in our experimental setup unless absolutely neccessary.

4.1.3 Sample Chamber

Because the Raman signal is approximately 10^{-10} of the intensity of the incident beam, the sample chamber must be designed to capture as much of the Raman light as possible. Furthermore, the sample chamber must isolate the incident laser light from the Raman signal; otherwise, the laser light will saturate the sensor. Fortunately, for linearily polarized laser light, the bulk of Raman scattered light is at a 90 degree angle from the incident beam. The 10 mW laser is linearily polarized.

Because the interior of our sample cell will be traversed by beams of light, each of them will scatter Raman light in a cylindrical shell. Therefore, the entire volume of the sample chamber will essentially be irradiated with Raman light; our task is to collect as much of it as possible.

The sample chamber which we envision is depicted in Figure 2. Because it may take a number of reflections before Raman scattered light enters the analytical optics, high reflectivity is vital. Therefore, we plan to gold-plate the inside of our sample cavity; gold plating is not only easy to accomplish, but has an extremely high reflectivity in the low-visible and near-infrared regions.

The laser light is introduced near the top of the cylinder, slightly off of the radial plane of the cylinder. It is aimed slightly towards the other end of the cylinder, such that it will follow a helical path down the tube. This path should concentrate Raman scattered light along the longitudinal axis of the cylinder. At the opposite end of the cylinder from the laser (the bottom), a planar mirror reflects light back towards the top.

At the top of the cavity, a convex lens focuses the scattered light onto the analytical optics.

The intensity of the Raman light is typically about 10^{-10} times the intensity of the laser light; therefore, without optical losses, the exit energy will be about 1 pW. A pessimistic estimate for the number of photons actually collected is 1%; therefore, let us assume that the output power will be about 10 fW.

4.1.4 Analytical Optics

Once scattered light has been collected, the remaining task is to measure the intensities of different colours. The easiest way to accomplish this separation is to shine the light onto a diffraction grating. Then, the light which is refracted by the grating may be measured by the CCD sensor, which is described in the section on electronic modules.

At a laser wavelength of 650 nm, the Raman lines for oxygen and nitrogen appear at

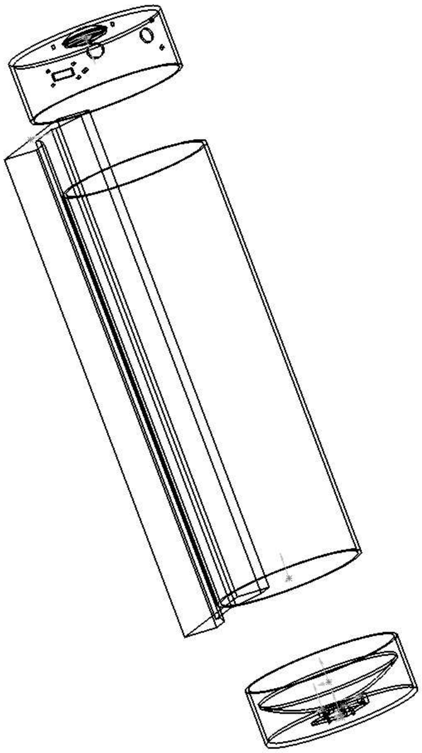

Figure 2: Sample chamber topology

723 nm and 766 nm respectively. We also wish to capture the laser on the CCD in order to be able to measure its wavelength spread. Using a fairly crude approximation, a diffraction grating with 600 lines per millimeter will capture a spread between 640 and 875 nm on the CCD, assuming the CCD is 10 cm away from the grating. This allows the laser, oxygen, and nitrogen lines to be captured; also, enough range is left to allow the measurement of water vapor and a number of pollutants.

Diffraction gratings typically transmit over 50% of the incident light; therefore, we will neglect optical losses here because such a pessimistic attenuation was assumed in the sample chamber.

4.1.5 Sensor

We plan to use a linear CCD sensor. The operation of this device is chiefly electronic; therefore, it is introduced in the following section.

4.2 ELECTRONICS

The electronics portion of the design of our Raman spectrometer bridges the gap between software and optics. It is not only responsible for providing a suitable platform for software development, but also for transferring the Raman signal into the digital domain in a format which is convenient for analysis.

There are five subsystems within the electronics module: the CCD, and its interface electronics; the host interface, and memory for storage; glue and miscellaneous components; and a power supply. Each subsystem is described in a subsection below.

4.2.1 CCD and Interface Electronics

The CCD device bridges the gulf between electronics and optics. While conceptually straightforward, these devices are difficult to interface to. Before considering the interface electronics, however, we present an introduction to CCD device operation.

4.2.1.1 CCD Device Basics

The active elements in a CCD sensor are similar to MOSFET transistors. For a digital camera, these elements are arranged in a grid (producing an area sensor); for our purposes, we require a linear CCD which only has a single stripe of active elements.

Collecting photons is a straightforward process. Under the correct biasing conditions, photons become trapped in the sensitive area; each CCD element behaves as a charge integrator which collects optically generated electrons. Neglecting noise, the largest portion of which results from thermally generated carriers, the sensor behaves just like a piece of film.

The readout process is not nearly so simple. After the sensor has been exposed to light for a period of time (known as the integration time) the sensor is ready to be read out. This typically happens one pixel at a time. Because the pixels are not read out simultaneously, the charges collected during the integration time must be removed from the sensitive area — otherwise, the pixels which are read out last will have had a longer integration time than those read out earlier.

Therefore, most linear CCD devices have an analog shift register arrangement. After the integration time has elapsed, all of the charges collected are emptied into a shift register. Each pixel is then read out sequentially.

4.2.1.2 CCD and CCD Interface

CCD devices are straightforward in operation, but are extremely tricky to interface with. They are a mix of analog and digital circuitry, and have extremely low noise tolerances. They also require extremely precise clocking.

Therefore, in addition to the CCD itself (which is typically available in a DIP package), two additional requirements are met by ICs. The first is an AFE (Analog Front-End) chip. An AFE takes care of digitizing the incoming signal, and performs some simple signal processing. It provides a series of 16-bit words, multiplexed onto an 8-bit bus.

The clocking requirements are met by a programmable logic device (either a FPGA or a CPLD.) Because we anticipate that successfully clocking the CCD will require several design iterations, we wish to avoid discrete clock circuitry at all costs. The relatively high speed of the FPGA/CPLD with respect to the readout rate of our CCD should ensure that we will not be taxing the FPGA/CPLD with inappropriately complex or stringent timing requirements.

The CCD devices which we are considering using (manufactured by Toshiba, e.g. the TCD1205D) operate at a typical readout rate of between 0.5 and 2.0 MHz. Because the timing signals must be carefully managed, an overall clock resolution of several times that would be required to ensure adequate timing resolution. Therefore, we intend to clock our FPGA/CPLD at a rate of above 10 MHz, allowing for a minimum phase resolution of greater than 5 FPGA/CPLD clocks per CCD clock. The greater the ratio of the clock rates, the more precisely we will be able to clock our CCD.

Detailed CCD state-transition data is included in the section on control, which may be found below. Detailed timing information is available in the timing section, also below.

4.2.1.3 CCD Sensitivity and Signal Strength

The TCD1205D has a sensitivity figure of 80 volts per lux second. One lux is approximately equivalent to 1.46 $\mathrm{mW/m^2}$; therefore, the CCD has a sensitivity of $5.5Vcm^2/J$. Assuming that the entire Raman energy falls onto 100 pixels, an integration time of 3 seconds will result in a Raman voltage signal of 6 mV on each active pixel. The dark voltage generated during that integration time is well within the saturation range of the CCD.

Dark voltage is noise which CCD devices generate when charge carriers are thermally generated within the active region of a pixel. For the TCD1205D, this voltage is given in terms of volts per integration second; however, dark voltage is not smooth. There is a reasonably large noise component which accompanies the steady bias increase.

The easiest way to reduce both dark voltage and the accompanying noise is to cool the CCD. Each 6^oC drop in temperature halves the dark voltage. By reducing the CCD to freezing, the dark voltage is reduced to 33mV, and the signal-to-noise ratio is increased accordingly.

Because we are using a 16 bit CCD controller, we can estimate the level of precision which we can expect our signals to have. The TCD1205D lists a saturation voltage of 800 mV; we may naïvely assume that 14 of the 16 bits are valid and that they measure the entire 800 mV range. These assumptions give a voltage resolution of around $50\mu V$, which means that our estimate on the Raman voltage will have about 8% accuracy. While this figure is quite close to the accuracy requirement in the functional specifications, the assumptions required to derive it were pessimistic enough that a reasonable margin for error exists. Additionally, the CCD controller which we anticipate using has a significant amount of circuitry dedicated to providing full accuracy in many conditions; assuming that the A/D converter measures the entire 800 mV span leaves an extremely wide safety margin.

4.2.2 Host Interface

We have discussed the circuit's interface to the optical portion of the project. On the other side of the electronic module is the host interface. According to our functional requirements, this device must have an RS-232 interface to a PC. By using the PC for most data analysis tasks, we maintain both the greatest flexibility and the largest ease of development.

We intend to use one of Microchip's PIC microcontrollers. The device which we will use will have a serial interface suitable for use with RS-232; therefore, in addition to the microcontroller and some discrete components, only a line driver (such as Dallas/Maxim's MAX232) required to translate $5V$ logic levels to RS-232 line levels. This line driver will also provide some protection; the RS-232 port is subject to ESD and other abuse from rough use.

4.2.3 Controller / State Machine Hardware

Both the host and CCD interfaces require complex control. While the host is ultimately responsible for controlling the measurement process, it cannot (and should not) meet the stringent timing requirements imposed by the CCD interface. The RS-232 interface between the spectrometer itself and a PC is simply too slow.

Therefore, the spectrometer must be able to function in a largely autonomous manner. While it will still accept configuration information from the PC, the PC itself should take on a supervisory role only. It will instruct the device when to make a measurement, and receive the data once measurements are complete. It will not be required to control or guide the measurement process in any way, apart from setting configuration registers (which specify, for example, the propogation delay.)

There are two classes of control requirements, both of which will need to be met by our project. The first is best addressed in hardware — clocking the CCD, for instance, would be extremely difficult to do in software. The second, of course, is software; it would, for example, be extremely cumbersome to implement the RS-232 interface with a host PC in logic.

In order to address both of these classes of hardware requirements, we propose to use both a microcontroller and a FPGA or CPLD as controller components. The state machines are detailed in the Control Architecture section below.

4.2.4 Memory

The rate at which data is read from the CCD is much greater than the rate at which the RS-232 interface to the host PC can operate. Therefore, it is impossible (and, regardless of feasibility, highly undesirable) to operate the CCD interface in lockstep with the RS-232 interface. A huge degree of flexibility is attained by storing the data from the CCD within the device until it is requested by the host PC.

It is obvious, therefore, that some memory is required. Because we anticipate a CCD with fewer than 4000 pixels, our memory requirements will only be 4000×16 bits per pixel, or about 8 kilobytes. Because RAM is so inexpensive, and because the additional flexibility could prove extremely useful, we have selected a 32 kilobyte RAM device. The additional data could be used, for example, for signal processing or averaging functions.

4.2.5 Power Supply

Because we are using both digital and analog signals, and because we are measuring very faint effects, a clean power supply is absolutely vital. Because we do not have the expertise or time required to develop a high-quality power supply, we will use a third-party supply.

For a commercial device, this power supply would need to power all of the components in our design; however, because our laser is powered off of a standard AC outlet, only the remaining devices use the power supply in the prototype device.

5 Software and Control Architecture

The behavior of the gas analyzer is defined by three parties, all of which are described above. The microcontroller and programmable logic device both directly control the operation of the device; their operation, in turn, is specified by the host PC via its RS-232 link.

Because all three of these subsystems are involved in the full control of the device, a full description of the states and dataflow is presented separately from the descriptions of each of its constituent subsystems. This section details these states and control requirements.

There are two primary state machines operating within the gas analyzer. The first is responsible for reading data from the CCD into memory. The second handles requests from the PC via the RS-232 port. While these two state machines are independent and perform different functions, they must operate in lockstep — when the CCD is being read out to memory, the PC should not attempt to access memory, since there is no guarantee that it would be able to read an entire dataset before the CCD state machine overwrites the memory.

5.1 Reading from the CCD

Figure 3 shows the state transitions involved in reading a line from the CCD and storing it in the device's memory. It is important to recognize that this state chart represents the states in the CCD readout mechanism only; indeed, while this state machine is in the idle state, the host PC may be reading sample data out from memory.

The following sections detail each state within this state machine.

5.1.1 Idle (State 1)

During this state, the readout portion of the design sits idle. The CCD state machine should be in this state when the host interface state machine is accessing the RAM; otherwise, concurrency issues could crop up. The entire situation is best avoided; therefore, the state

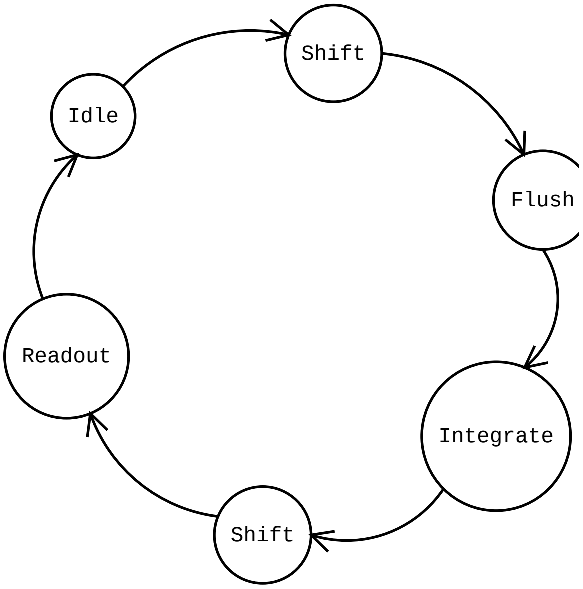

Figure 3: Readout state transitions

machines operate in lockstep. One may operate only if the other is in the idle state, unless their operation is noninterfering.

5.1.2 Shift (State 2)

Once the host computer has requested a that sample be taken, the readout state machine leaves the idle state. However, before the data in the CCD may be read out, any electrons currently in either the shift register or the sensitive area of the CCD must be flushed. This state lasts for only as long as the CCD requires to shift electrons out of the sensitive area — which is one cycle of the CCD clock.

5.1.3 Flush (State 3)

Once the active area has been flushed of electrons, the analog shift register used for read-out must be cleared as well. Because pixels are read on a one-by-one basis, this flushing process takes an amount of time proportional to the number of pixels.

5.1.4 Integrate (State 4)

Integration time cooresponds to the amount of time that the CCD has been collecting charges on its active surface. Since the active surface was last cleared via a shift during State 2, it has been collecting charges since then; this state, therefore, simply waits until the remainder of the desired integration time has elapsed.

5.1.5 Shift (State 5)

Once again, the data must be moved from the active element into the analog shift register. This is required both to stop more charges from being accumulated, and to allow the read-out electronics to output pixel data. As before, this shift state occupies one clock cycle of the CCD clock.

5.1.6 Read (State 6)

The final state is similar in operation to the Flush state. Once again, pixel data must be read out from the CCD one pixel at a time; however, this time, the data (which is digitized in the CCD controller) must be stored in RAM.

5.2 PC INTERFACE

The PC interface state machine is extremely simple; the host interrogates the gas analyzer via a command-response protocol. For the ENSC 340/305 prototype, no communications are initiated by the device itself. For later device models, there are situations in which the device may wish to notify the PC of some event; however, since modifying the behavior of this state machine is simply a matter of assembly code, assuming that the PC initiates all communications will greatly ease the development process without precluding project development beyond the end of this semester.

The protocol detailed here is extremely slow; we anticipate that it will be the largest contributor to the latency between when a sample is requested and when it is available on the PC. If response time requirements are difficult to fulfill, the serial interface detailed here will be the first candidate for reexamination.

The request/response protocol is simple. There are very few commands; each are initiated with a one-byte command code. The commands are as follows:

No-op No operation: ignore the command.

Read Byte Read a single byte from memory space.

Read Block Read a block of 256 bytes from memory space.

Write Byte Write a byte to memory space.

Configure CCD Reconfigure the CCD controller using configuration data stored in the on-board RAM.

Each of the three memory transactions begins with a one-byte command code, followed by two address bytes (high byte, then low byte.) Then, for write operations, the PC supplies a byte of data to write; for read operations, the device sends either one byte (Read Byte command) or a series of 256 bytes (Read Block command.)

The configuration transactions are single-byte commands; only the command byte itself is sent. Configuration data is written to memory using the existing memory access commands.

The remaining command is a no-op. This command can be sent repeatedly by the PC to ensure that the device is ready to accept a command, rather than being left in an interim state. Should communications between the PC and the device become desynchronized, repeated no-op commands can effectively bring the communicating parties back into sync.

Both sample and configuration memory is mapped into the address space specified by these commands. With a total of 16 address bits, a 64 kilobyte address space is available for PC inspection; therefore, even with the entire 32kb RAM mapped into memory, a large number of address locations are available for configuration or expansion use.

Assuming the specified minimum configuration rate of 9600 bits per second, we may make some estimates about the required time to transfer sample data to a PC from the device. If the CCD contains about 4,000 elements, then about 8 kilobytes of data must be transferred to the PC to obtain an entire sample. Assuming only block read commands are used, the overhead consists of three bytes per 256; because 8 kilobytes requires 32 blocks of 256 bytes each, a total of just under 10kb must be transferred. At 9600 baud, about 10 seconds will be used for transfer purposes. This protocol, therefore, consumes about 10 seconds; our maximum turnaround time for a measurement is 30 seconds. The remainder of the device must perform its functions in under 20 seconds to meet specifications.

It should be noted that these estimates are extremely conservative; the largest CCD which we are considering using has some 3,500 active elements. Rounding has also been extremely generous. The transfer speed which was used is the minimum that will meet our functional specifications. Minimum performance should therefore be better than this estimate.

5.3 Minor States

Apart from the two main operating states (for CCD read-out and RS-232 communications), there are a number of minor states. These states consist of configuration and set-up states. Because they are both extremely simple, and because they depend on the specific hardware used in the device, they are not detailed here. In terms of cooperation with other states, these states are required to supersede read-out or communications states.

5.4 Host Software

In addition to the microcontroller's software requirements, we must provide some means for data to be read off of the device onto a PC. This software will communicate with the gas analyzer using the request/response protocol detailed above. It will place sample data in a format convenient for analysis with third-party tools such as GNUplot. These tools will be used to produce a spectrograph.

6 Timing Requirements

Clocking the CCD and controller is a sufficiently complex topic that it requires special attention. This section details the timing requirements placed on the CCD readout circuitry.

The description of CCD operation provided in the Electronics section above neglects some of the clocking details. As described, there are two shift registers — one above the active area, and one below. When data is shifted out of the active area, even pixels are shifted into one shift register, while odd pixels are shifted into the other. Therefore, there are two separate shift controls: one for each shift register.

If these shift controls are clocked inversely with respect to each other, then data is shifted out on both high and low clock edges. Rising edges' signals come from one shift register, and falling edges' signals come from the other. The clocks for the shift registers are denoted ϕ_1 and ϕ_2.

In addition to these two clock signals, both the reset (RS) and boost gate (BT) signals must be clocked at the data output rate, which is twice the clock rate of the ϕ_1 and ϕ_2 signals.

The procedure which shifts data from the sensitive elements up to the analog shift register requires that the phase clock is temporarily stopped. During this pause in CCD clocking, raising the SH pin will shift all of the collected charges from the active sensor element to the analog shift register. After this shift, the readout process may begin.

During readout, clock signals must be kept at carefully matched phase shifts. The clocking diagram for reading out the CCD data is shown in Figure 4. (This figure is reproduced from the Toshiba TCD1205D datasheet without permission.)

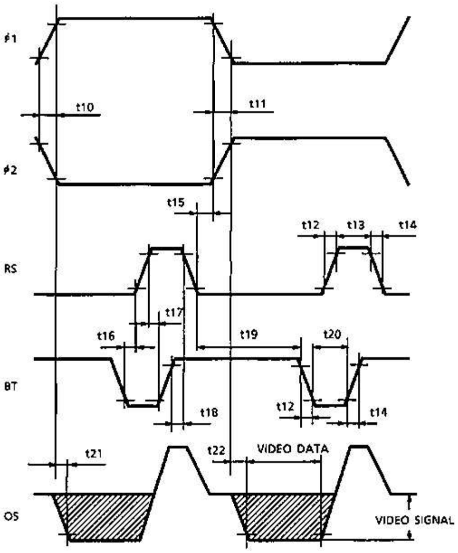

Figure 4: TCD1205D Readout Clocking

As mentioned above, the CCD readout speed for these CCDs is a maximum of 2.0 MHz. Therefore, clocking the FPGA/CPLD at 20 MHz or more will give sufficient precision to synthesize a suitable CCD clock.

The CCD controller IC, as well, requires careful clocking. Because data arrives on each rising edge of the CCD's RS or BT signals, the CCD must be clocked at a multiple of this rate. For a CCD front end such as National Semiconductor's LM9823, the controller is clocked at between twice and eight times the CCD data rate. For a 2 MHz readout rate, therefore, the CCD controller must be clocked at between 4 MHz and 16 MHz. Depending on our selection of timing generator frequency, our CCD controller clock rate may be constrained.

7 EVALUATION OF REQUIREMENTS

Throughout this document, the design process focused on the flow of data from one module to the next. By presenting the design in this format, a logical process of development should become obvious to the reader.

However, the eventual purpose of this document is to realize a design which fulfills all appropriate requirements from the functional specification document. Because the design has now been presented, an evaluation of how well it meets the functional requirements is now possible.

The following functional requirements apply to this phase of the project:

RN-13-123 The device shall have an IEEE 1284 (RS-232) serial connection via a 9-pin DB port (MALE).

RN-14-123 The RS-232 port shall provide sufficient driving voltage and current as per the IEEE 1284 standard.

RN-15-123 The RS 232 port shall be isolated and protected from the rest of the device circuitry.

We intend to use a MAX232 line driver which communicates directly with our microcontroller. This line driver connects directly to a DB-9 socket, contains isolators, and conforms to the IEEE 1284 standard.

RN-18-123 The device shall draw under 100W of power from the power outlet.

The CCD controller, CCD, microcontroller, and programmable logic devices each consume well under 1 watt each. The laser consumes under 10 watts. Should the CCD require cooling to lower its noise figures, a thermoelectric (Peltier) cooler might be required; these consume around 60W of power. Our total power consumption remains under 100 watts.

RN-19-1 The device shall provide percentage composition data to an accuracy of not less than 10%.

As calculated in the electronics section, a set of extremely pessimistic assumptions results in an accuracy of 8%.

RN-21-123 The sampling time required by the device shall be under 30 seconds.

The sum of integration and communications times is well under 30 seconds. Several states do not have calculated delays; they will be negligable in comparison to these two sources of delay.

RN-34-123 The device shall communicate with the PC at a data rate of not less than 9600 bps.

PIC microcontrollers communicate via RS-232 at data rates well in excess of 9600 bps. We will likely use 9600 bps unless this data rate proves unsuitable.

RN-42-1 The PC software shall display the data on screen as it is received from the device.

PC software will be developed which produces a graphical desplay of the spectrograph after it is received

8 CONCLUSION

This document demonstrates that the Raman effect is theoretically measurable with a commercially viable, inexpensive design. It also demonstrates that we have the ability to address this project's functional specifications in the same design. What remains to be seen is whether we can realize such a design in practice; not only does it involve technically complicated electronics, it also contains optical components. We have little experience with optics outside of a physics lab; therefore, it will be interesting how closely our eventual design matches our intended one.