SO2 Camera Final Thesis Report

Jeff Brown & Mike Taras
April 22, 2009
Alex McGlashan

Table of Contents

Introduction / Overview:

The SO2 project expanded upon the principles of remote gas sensing, and a camera was developed capable of real-time detection and concentration of gases (primarily SO2 for the scope of this project). With such a device, SO2 can be detected and measured simply by aiming the camera system at a plume, such as one from an active volcano or industrial stack. The system uses readily available components and can be manufactured at a lower cost than similar (and very rare) systems currently available. As opposed to simply logging data for analysis at a later time, the device shows - in real-time - a false-colour image correlating to the concentration of the gas observed, overlaid onto a picture taken with 300 nm light of the object being studied. The result is then an overlay of the gas concentration onto the actual picture of the stack being analyzed. The device is small, portable, and utilizes a USB interface for convenient, quick setup and operation.

Expanding upon Dr. Morrow's (Resonance Ltd. President) work with remote gas sensing along a single path [1], we used a CCD camera to detect specific gases (SO2) in a 2D array. This would allow for spatial information to be gathered, and thus form a picture of the gas and its concentration. This therefore eliminates the need to scan a linear remote gas sensor across the object being studied to obtain a complete concentration profile of the plume.

The technique uses a cell containing a gas that is to be detected, and a reference cell. Looking at a plume or area of interest with a potential concentration of the gas in question, the two cells are switched in and out of the CCD detector, and an image is captured and saved for each. It should be noted here that the CCD is sensitive to only a small region of wavelengths that are known to be absorbed by the gas in question. This is accomplished by placing a band-pass filter in front of it. This filter will block out regions that do not have absorption peaks for the gas to be detected, and thus greatly reduce noise. Analyzing the difference in intensity for each pixel will correlate to the amount of light that was absorbed by the gas in the plume as compared to the gas in the cell. Therefore, the concentration of the gas in question is known for each pixel on the CCD array, giving a spatial 2D image showing gas concentration. The technology has been demonstrated using a single pixel as per W.H.Morrow research. [1] The SO2 camera uses the same technique but extends it to utilize a two-dimensional array of pixels.

The detection of SO2, the primary gas focused on for this project, required a CCD sensitive to light in the ultraviolet range (~300 nm). In order to obtain such a sensor, Lumogen, a fluorescent substance, was researched as a means to convert 300 nm light into visible light – as a coating on the CCD surface. The coating would fluoresce, emitting visible light detectable by the CCD when excited by the incident UV light [2], however problems applying the coating eventually rendered the Lumogen method impractical. Instead, the exposure time of the camera was extended to 250 ms to gather enough UV light to create a sharp image.

Purpose / Benefits / Practicality:

The following quote was taken directly from the Ontario ministry Of The Environment website:

> "Health effects caused by exposure to high levels of SO2 include breathing problems, respiratory illness, changes in the lung's defenses, and worsening respiratory and cardiovascular disease. People with asthma or chronic lung or heart disease are the most sensitive to SO2. It also damages trees and crops. SO2, along with nitrogen oxides, are the main precursors of acid rain. This contributes to the acidification of lakes and streams, accelerated corrosion of buildings and reduced visibility. SO2 also causes formation of microscopic acid aerosols, which have serious health implications as well as contributing to climate change." [3]

Ontario Sulphur Dioxide Emissions by Sector
(Emissions From Human Activity, 2000 Estimates)

Pie Chart

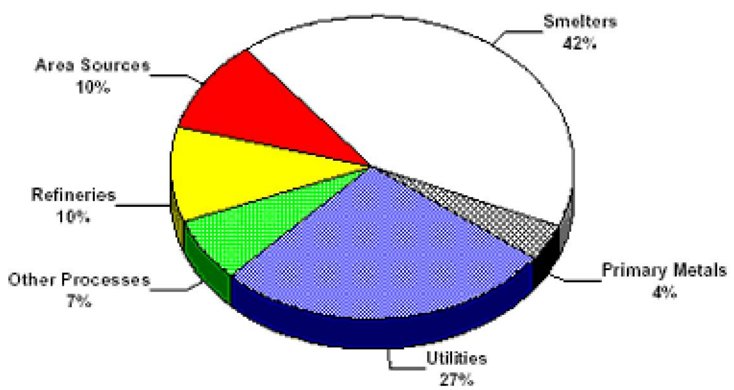

Category	Percent
Smelters	42%
Utilities	27%
Area Sources	10%
Refineries	10%
Other Processes	7%
Primary Metals	4%

The resulting SO_2 camera prototype will allow:

- Monitoring and regulate the amount of potentially harmful gases emitted from industrial plants (Nanticoke, Sudbury, various mines);
- Accumulation of data pertaining to various gas emissions across entire regions. If the flux of gases could be known entering various ecosystems, any effects in the future could be compared to historical gas levels, which have been recorded thanks to remote gas sensors such as the one proposed in this document;
- Monitoring and testing gas emissions from foreign vessels – such as those frequenting shipping lanes (St. Lawrence Seaway, shipping ports) and verifying they conform to Canadian environmental legislation;
- SO2 readings can help to identify acid rain patterns across Ontario.

SO2 Gas Camera Weekly Progress Reports

First Week Of January (Including End of December, 28th – 3rd)

In this week Jeff and I sat down to design the SO2 camera and discuss our options with the type of cameras to purchase and the predicted sensitivity of the device. We had a meeting with Bill on December 16th in Barrie, and researched the fluorescent properties of lumogen applied to glass slides and mixed with epoxy. It was found that the lumogen did not fluoresce after being dissolved in acetone and then let to dry. This is probably due to how the molecules are arranging themselves and settling after the acetone evaporates. This was an unfortunate discovery as dissolving the lumogen in acetone and then letting it evaporate produced very thin and uniform layers of lumogen. We discussed heavily the math behind the correlation technique, and another trip to Barrie to consult with Bill will be needed to fully understand the concepts.

For the camera prototype, two simple servos will be used to place the SO2 gas cells in and out of the optical path. To do this, and to synchronize it with the LabVIEW software interface (that still needed to be written at this point) I developed a servo controller and serial port interface on a PIC18f microcontroller. The code was written in assembly and takes commands from a serial port to move the servos, which place the gas cells in and out of the optical path. It then signals that this task has completed, where the LabVIEW control program would acquire an image from the camera, and this process would repeat. Current estimates place the framerate of this system at 0.5 – 1 frames per second.

Second Week Of January (4th – 10th)

With the lumogen research looking unlikely to be successful with spin coating, Jeff and I started to consider using thermal deposition to apply a layer of lumogen 20 – 30 nm thick, as outlined in technical papers detailing methods to apply lumogen layers. This method was previously explored, however the lack of a thermal evaporation system put this on hold. As a result, we started to consider purchasing a more expensive, but highly sensitive back-thinned CCD that would be 40% more efficient than our current camera in the 300nm range, effectively eliminating the need for the lumogen coating. The camera was thought to be roughly $4000 with all necessary electronics; however this quote turned out to be $2000 too low. Therefore, this system was unfortunately considered out of the realm of possibility within the constraints of this project.

A teleconference was held on Friday, January 9th with Bill from Resonance to discuss in further detail the expensive back-thinned option. Using a rectangular CCD, it is possible to have a system with no moving parts, having two small lenses side by side, each imaging on their respective half of the CCD, as in the diagram below.

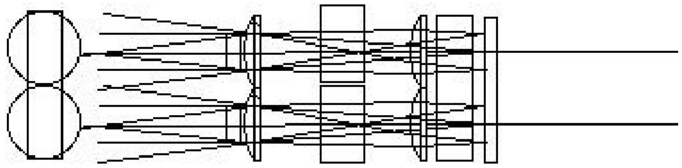

This method would have the advantage of being very small, less complex, and more robust. Unfortunately in the end the cost was too great and our research shifted back towards the monochrome cheap camera.

Third Week Of January (11th – 17th)

The Sony CCD monochrome "near-UV" camera was tested and compared to the Pulnix on January 14th. It was found to be equally as sensitive as the Pulnix, which indicates the CCD is likely no different from any other CCD, albeit having the window removed so UV isn't blocked. We also reconsidered thermal deposition of lumogen since Resonance currently has an unused small vacuum chamber with heating element that Jeff and I could use for deposition. Since we do not have a deposition rate monitor, we discussed a method to determine when the lumogen layer thickness has reached an optimal point. A photodiode, sensitive to the 633 nm (red) light given off from the lumogen will be placed beside the quartz slide to be coated within the thermal evaporator. This diode will not be sensitive to 300 nm UV light, which will be directed at it from a source supplied by resonance. The 300 nm source will also be outside of the deposition area so it does not get covered. Using a microcontroller to log the intensity output of the photodiode, the thermal deposition can begin, and as the diode is coated with lumogen the impinging 300nm radiation will start to fluoresce and red 630 nm light will be emitted, which the diode is sensitive to. Therefore, the point in time at which maximum signal is acquired from the photodiode dictates the optimal lumogen layer thickness.

Fourth Week Of January (18th – 24th)

Another teleconference with Bill was held on January 21st. In this meeting, we discussed the manufacture and shipment from Resonance to us two SO2 calibrated gas cells. The cells will be needed to test the sensitivity of the instrument. When gas cells are inserted into the field of view, a reduction in intensity should be seen. This intensity difference directly correlates to the resolution of the camera's SO2 readings. The greater the intensity difference, and the greater the bit-depth of the CCD, the more graduated steps in intensity will be apparent. Bill liked the idea of using his thermal evaporator and the diode sensor to determine optimal lumogen layer thickness.

Fifth Week Of January (25th – 31st)

In this week, focus was placed on the LabVIEW programming and acquisition of images for manipulation via USB framegrabber. Most of my time this week was spent on coding, and readying the software for addition of the correlation math needed to make the system work properly. As of now, image acquisition, subtraction, and intensity reading at every pixel is available. Math operations are easy to implement in LabVIEW, and serial port communication with the PIC controlling the servos has been completed. The false-color algorithm is also being developed with progress being made.

First Week Of February (1st – 7th)

During this week, a rough prototype of the system was planned and started under construction. Servos need to be ordered for the gas cell switcher as well as a voltage regulator and other miscellaneous electronic tidbits. Jeff is mounting a gas cell holder wheel to a stepper motor for trials, but a servo system will be used for now. The LabVIEW programming of the main software continues, and the system will be mounted on a portable block of wood so readings can be taken outside for development.

Second Week Of February (8th – 14th)

Jeff Brown started to gather parts and design a microcontroller based deposition rate monitor using a photodiode for the deposition of lumogen onto a substrate. The system, as explained in the section *Third Week Of January* will allow us to indirectly know the optimal layer thickness of lumogen to give maximum CCD sensitivity. I am continuing LabVIEW programming and also constructed a basic filter wheel for the SO2 camera using a servo and microcontroller. The LabVIEW interface gives serial commands to the filter wheel which switches gas cells in and out of the optical path.

Third Week Of February (15th – 21st)

The serial control interface for the gas cell servo switcher was completed. The setup is crude but perfectly acceptable for testing. A servo was made to have 270 degrees of rotation (done by supplying it with non-standard pulse-width modulated (PWM) signals from a microcontroller and therefore able to place three gas cell in and out of the camera's optical path (spaced 90 degrees apart). A PIC microcontroller was used for the serial interface and servo control, and accepts 1-byte commands from a computer running the LabVIEW software. The value 0 – 255 of the byte directly correlates to the servo's position, with 0 being full left rotation, and 255 full right. With this setup rudimentary testing of the SO2 camera may begin, with automated gas cell swapping. Gas cells are being shipped to us from Resonance, and we await their arrival before any experimentation can be performed.

Fourth Week Of February (22nd – 28th)

With no gas cells from Resonance yet, the LabVIEW code for the camera was refined further to make capturing images after the servo has moved into position an automated procedure. Careful timing was implemented such that the frame capture occurs just as the servo is done moving the cell into position. Immediately after a frame is successfully grabbed, the servo is commanded to switch the next cell in until all required images for the correlation are obtained, at which point the program loops. With the camera mounted on a wooden plank for testing, and a USB connection for the framegtrabber and serial cable for the PIC, the camera system is now mobile when used with a laptop and can be brought outside to obtain images of the sky. The gas cells from Resonance must arrive before any significant tests are performed. Talks with Resonance reveals that the gas cells are almost ready for shipment.

First Week Of March (1st – 7th)

A trip was planned to Resonance during this week to meet with Dr. Morrow and discuss the progress of the project and to work out the correlation algorithm to be used in the system. We also pressured Bill to have the gas cells ready and successfully obtained them. Much was discussed during the trip and some much needed experiments were performed. Jeff and Bill discussed using the vacuum system to thermally deposit lumogen on the Sony camera's CCD, however due to time constraints it was decided not to pursue this route until deemed absolutely necessary. Instead, the cameras were tested with a halogen light source to measure their sensitivity. At first glance, the sensitivity seemed acceptable for SO2 detection, however it was found that the 310 nm bandpass filter has an IR leak, and therefore the camera had a higher apparent sensitivity than first measured. Bill had an IR blocking filter that was used to ensure that only 310 nm UV light was making it to the CCD, and a noticeable drop in signal strength was observed. We noted that the maximum exposure time for the cameras (Sony and Pulnix) was only 1/60th of a second, and therefore decided to add multiple frames in LabVIEW to try and increase contrast and signal. This approach was successful; however the images remained noisy due to the readout noise incurred every time the CCD was read. We added 16 consecutive frames to obtain an image, and averaged the pixels. An acceptable result was obtained when looking at the halogen lamp (calibrated to be roughly the same output as the sun at 300 nm). We placed gas cells in the optical path and noted the drop in intensity to try and perform a basic correlation. It was found, unfortunately, that the correlation strength was only on the order of 2%, when at least 7 % was expected. This meant that the camera would be able to detect SO2, but wouldn't be sensitive enough to qualify as a sellable product. The low sensitivity has to do with how the CCD simply integrates all light that makes it in the camera, instead of looking at specific absorption wavelength peaks of the SO2 signature. Because approximately half of the light making it through to the CCD does not respond to SO2 being in the optical path, it simply adds noise. Increased sensitivity was obtained with a faster lens (F1 as opposed to F2.5) and talks of increasing the exposure time of the camera, or finding a different camera with a longer exposure time began.

Second Week Of March (8th – 14th)

Mike Taras worked on the LabVIEW interface for the camera and was able to successfully create false-colour images from grayscale image information. This meant that the false-colour representation originally intended for the SO2 camera would work in real-time. No processing overhead was observed for the image transformation.

Third Week Of March (15th – 21st)

Jeff was able to use a PIC microcontroller to supply a pulse-width signal to the Sony camera in order to manually increase its exposure time and trigger it. This approach was extremely successful and numerous images were taken outside the college looking at the sky. The longest exposure time of 250 ms was not needed; the camera was saturating at much less. A reduction in intensity was clearly seen when a gas cell containing SO2 was placed in front of the camera. This drastic increase in sensitivity makes the camera capable of detecting SO2 in increments – which is a huge advancement in the project.

Fourth Week Of March (22nd – 28th)

Since the exposure time of the camera was increased to a maximum of 250 ms, the video output no longer adhered to composite video format standards. Therefore, a LabVIEW program was made to only grab good frames, and discard the random noise and bad frames between them. This made the video output smooth and watchable as opposed to looking like random noise.

Major progress on the UAV was also made, mainly the installation of a microcontroller system to interface with the onboard computer. A PIC was used to send and receive serial data to the computer, take servo positions and control 4 servos. The PIC would also select which input source to use for the servos - either the manual radio receiver, the onboard computer, or the backup autopilot.

End Of March / Beginning of April (29th – 4th)

Further progress on the LabVIEW software interface was made in this week. The interface combined the servo-controlled gas cell switcher, the false-colour algorithm, and the code to only grab good frames into one program. With this software the camera could be taken outside and evaluated. Many images were captured and the correlation algorithm was applied to some of them. A weak result was obtained, indicating the system needed to be optimized. Jeff analyzed the recorded data and began to work out what needed to be changed.

First Week Of April (5th – 11th)

The UAV aspect of the project rapidly advanced, given the nicer weather. An onscreen display was installed in the UAV which overlays GPS data onto the video feed before it is transmitted to the base station. With this setup the pilot could fly the plane using the video feed, and have course, speed, altitude, and plenty of other useful information at hand to avoid getting lost or running out of battery capacity.

A teleconference with Bill of Resonance was held, and various optimization techniques were discussed for the camera, including pixel averaging, aspheric lenses, and using two band-pass filters to look at different SO2 absorption regions.

Second Week Of April (12^{th} – 18^{th})

Nearing the end of the semester, the SO_2 camera was put on hold to focus on exams. Jeff and Mike were both hired full-time by Resonance where the project will be completed in the near future.

As for the UAV, radio modems were installed to allow real-time telemetry data to stream back from the aircraft to the base station. Using Google Earth as a moving map system, the UAV could be tracked and its location known in real-time, at all times.

Current State of Project: Objectives, Status and Progress

Put simply, the following list contains the objectives associated with this project that were created in September. Each point will be reviewed and evaluated to ascertain the current state of the project.:

- Produce a hand-held, portable prototype camera (under 750 g);

  This point will be discussed in the conclusion.

- Remove the window of a standard CCD camera to extend the UV sensitivity and research the possibility of using a window coated with Lumogen to extend the UV sensitivity of a standard CCD camera;

  The Pulnix camera originally experimented with did not have any sprt of window covering the CCD. The Sony camera had a glass window covering the CCD, but it was not attempted to remove it, since comparable sensitivity was seen with the Sony as the Pulnix. A Lumogen coating was applied to glass slides to evaluate how well it would fluoresce. The lumogen was mixed with a UV curable epoxy, and spin coating was used to uniformly distribute it on the slides. It was found that the lumogen layer did not emit light in the right direction, and therefore was not effective in converting the 310 nm light into 600 nm light that the CCD was much more sensitive to. Increasing the exposure time of the camera was the solution to obtaining UV sensitivity.

- Produce a prototype capable of capturing an image in the UV (~300 nm range) and analyzing the intensity of the image with and without an SO2 gas cell to determine the concentration of SO2 across the entire image (gas correlation technique) [1];

  The gas correlation technique was used to evaluate the ability of the Sony camera to produce a UV-sensitive, fairly high resolution image of a gas cell. A strong difference was seen in the intensity of light through the SO2 gas cell compared to the background illumination. This proved that the camera could indeed be used to detect SO2, however its sensitivity to SO2 levels was still unknown.

- Produce a prototype camera with at least 1024x768 resolution;

  The Sony camera chosen has a resolution of 720x480. Monochrome cameras at higher resolutions are more costly, but not out of the question. Production versions of the camera would possibly achieve the 1024x768 resolution parameter, but this is not critical.

- Produce a camera capable of remote sensing of gases many kilometres away;

Since the camera uses the same gas correlation technique as well established spectrometer-based systems, this goal has also been achieved. It is only a question of how sensitive the camera can be made.

- Sensitivity of instrument on the order of 100 ppm

No concrete sensitivity measurements have been performed on the camera yet. It is estimated that with the current UV sensitivity and some pixel averaging done in :LabVIEW, sensitivity of 100 ppm is obtainable.

- Install the prototype system into a UAV (unmanned air vehicle). Resonance, the industry partner, has a custom built UAV for this application.

The UAV is ready to fly and carry the camera to a remote location. The only limiting factor is the weight of the camera, since the maximum payload capacity of the UAV is about 500 grams. With a micro servo serving as the cell switcher, and a minimal framegrabber, the camera system falls within these weight limits and can be flown without any trouble on the UAV.

Current State of Project

At this point the SO2 camera is a work in progress. It has been demonstrated that the camera can detect SO2, however its sensitivity has yet to be evaluated. The project and research will continue at Resonance under supervision of Dr. Morrow, with most of the work done by Mike Taras and Jeff Brown. The low-cost factor of the camera has been retained, and a polished, finished system would cost less than $5000 Cdn in parts. More research needs to be done in regards to the sensitivity of the camera and how to increase it effectively.

References:

[1] W. H. Morrow, "Filter Correlation Spectrometers For Remote Sensing Of Tropospheric Gases," Ph.D. Dissertation, York University, Toronto, ON, Canada, 2002.

[2] H. Deeg, Z. Ninkov, "Characterization of a large-format charge-coupled device", Optical Engineering 34N1 (1995) 43-49

[3] Ontario Ministry of the Environment, Sulphur Dioxide (SO2) http://www.airqualityontario.com/science/pollutants/sulphur.cfm