Open Ocean • Here be Dragons 2018

Project Prometheus Report

Allan Adams¹ Jacob Bernstein¹

¹MIT Future Ocean Lab

Published on: Oct 06, 2021

DOI: <https://doi.org/10.21428/a680be9a.207c346e/ec788172>

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

The Prometheus project aims to change the way underwater caves are mapped, documented, and shared with the public. At present, underwater caves are typically mapped with taut lines, waterproof slates, and compasses; our goal is to create a diver-carried cave-mapping system that can create a 3d map of an underwater cave in real time, for hours on end, and at a cost comparable to a basic dive camera. This challenge has led us into technologically uncharted waters -- traditional approaches (LIDAR) proved too costly, and our chosen approach (ToF imaging) landed us on the bleeding edge of high-speed optical electronics. To date we have built two prototypes and fielded them on two expeditions -- this has given us valuable data and feedback. As of this writing, we have demonstrated proof-of-principle results for the core electronics and optics; key next steps include a systematic calibration strategy and a robust software pipeline, both of which are in early stages of development.

Introduction

Our initial goal was to develop a low-cost underwater LIDAR system similar to those used in autonomous cars, but using blue lasers (which penetrate efficiently in water) rather than industry-standard infrared lasers (which do not). As we worked through a design study for an underwater LIDAR scanner, however, we became aware of the possibility of creating a smaller, more efficient, more robust instrument at significantly lower cost by using Time-of-Flight (ToF) cameras. These cameras use modified CMOS image sensors and modulated light sources to create images where each pixel encodes the distance to an object, rather than its color. They can be manufactured compactly and cheaply enough to embed in a smartphone - or a phone-sized pressure housing. Affordable, adaptable underwater ToF cameras would be valuable for many applications, from documenting archaeological sites to monitoring aquatic life to augmenting the navigational capabilities of ROVs and UAVs.

Adopting ToF cameras for the Prometheus project posed increased risks along with increased opportunities. Traditional LIDAR systems have been demonstrated many times underwater; ToF cameras have not. To estimate the risk:reward ratio, we ran an initial design study modeling the performance of a ToF camera underwater and comparing it to LIDAR, ensuring we could find a combination of image sensor, lens, and LEDs that would satisfy the resolution and operational requirements of the cave mapping system. Our next critical test was developing the electronics which modulated the LEDs, providing sufficient optical power and switching speed to support the maximum resolution of the camera. We used the theory of operation of the ToF imaging chip to prove that our optical switching speed met the operational requirements of the camera (after several design iterations and the invention of a novel LED-driving circuit, including chips so new that they are just becoming publicly available). With our chosen camera components and custom electronics, we were able to take pictures and inspect the point clouds, an important milestone.

From there, the work expanded on parallel fronts, developing a prototype instrument system to enable camera testing, and running camera tests to understand and improve the prototype instrument. We designed a prototype to be tested in the field, with all electronics and mechanical systems designed to fit within a custom pressure

housing. We eventually took two prototypes to the field and learned many lessons about our design in the process. As the prototype matured, we were also able to more fully test the camera's performance.

While our estimates for camera resolution were sound, achieving a comparable accuracy requires calibrating against a number of variables which can be very hard to control or even measure for example the temperature distribution along the imaging sensor itself. While this has significantly lengthened our development process, it has not changed our expectations or our confidence in this strategy.

Design Study

Our technical objective at the project outset was to develop an instrument that satisfied specific performance requirements - capturing 3D scans of cave surfaces with centimeter scale voxels over 10m distances, on dives lasting many hours, borne by a moving diver who could not devote too much attention to the instrument operation, nor be overly impeded by its size. We began the project with a design study to identify a development strategy that could meet the requirements. This survey was conducted in two phases - first, we surveyed the space of candidate technologies, and then we drilled into technical details of our top choices.

We started by looking at LIDAR implementations of 3D mapping. LIDAR is a popular, well-supported technology, works well in water (at the length scale of our interest, when using blue light), and is straightforward to analyze and implement. Essentially, LIDAR is a form of optical distance measurement: the instrument scans a pulsed laser over a surface and measures the time it takes each pulse of light to return (called 'time of flight'). The distance resolution is determined by the precision of the instrument's clock, and the lateral resolution is determined by the raster distance between pulses; both are typically mm scale. The series of distance measurements is called a point cloud, which can be represented in 3D as a sampling of the surface position of whatever objects the instrument is scanning. LIDAR scanners are regularly used to map dry caves with mm scale accuracy, and represent a promising approach for underwater mapping as well.

This strategy becomes much more challenging when the camera is held by a moving diver (or AUV). If a LIDAR scanner moves while recording, each point in the point cloud records a distance from where the scanner was when that pulse was emitted, and relative to the direction the scanner was pointing at that moment. The true point cloud thus can't be reconstructed without knowing precisely how the instrument (and diver!) moved and rotated between each measurement. The precision of the final map is thus limited by the precision with which the 3d trajectory of the diver can be known -- which is obviously very difficult in an unknown underwater cave! This measurement, even performed with an expensive navigation system, tends to be orders of magnitude noisier than the native resolution of the LIDAR scanner, and so surface maps made with a moving scanner are much rougher than maps made with a stationary one. Some of this noise may be mitigated by advanced reconstruction algorithms, known collectively as Simultaneous Localization and Mapping (SLAM), if the same surface is scanned multiple times. Thus, in applications where accurate, moving maps are required, such as planning systems for self-driving cars, speed of scanning is critical.

Since LIDAR modules for self-driving cars constitute a multi-billion dollar industry, there are already fantastically-engineered LIDAR modules available today that tackle this challenge -- miniature, rugged, and capable of generating real-time 3D movies. These might be a great solution for underwater cave mapping, modulo their cost in dollars and Joules, except for one thing - they all use infrared light, which is invisible to humans but rapidly absorbed by water and useless within a few centimeters. Off-the-shelf LIDAR systems would thus be essentially useless to us, and we would have been foolhardy to try to reengineer our own.

Fortunately, self-driving cars are not the only industry in need of optical 3D scanning, and a variety of technologies have been commercially developed to fill various product niches. We searched through these for a solution that would match our needs for size, robustness, acquisition speed, and the ability to use blue light. We found that Time of Flight (ToF) cameras were a perfect match.

ToF cameras perform their time of flight measurement in a more subtle manner than conventional LIDAR systems. Each measurement is made not with a single pulse of a laser beam but with an oscillating light source flashing on and off at frequencies of order 10 MHz for thousands of cycles at a time. This oscillating light acts like a clock: it beams out from the light source, bounces off an object, and returns ever-so-slightly out of phase with the outgoing clock signal. The camera utilizes a CMOS imaging chip which measures both how much light returns (brightness) and the phase delay of the returning signal -- that is, the time of flight.

The upshot is that a ToF camera measures a distance for each pixel in a scene nearly simultaneously using a steady but rapidly-flashing light source. This greatly reduces the mechanical and optical complexity of the system and essentially eliminates the need for precisely measuring the motion of the camera and diver -- an enormous simplification and cost-savings compared to a LIDAR approach. Meanwhile, the light source no longer needs to be a pulsed IR laser -- instead we can work with high-efficiency blue LEDs which penetrate very efficiently through water. This first stage of our design study thus left us optimistic that ToF cameras would provide the technical foundation for the cave mapping system.

To determine whether this strategy could meet the stated performance requirements, we developed a simple model for ToF camera resolution which depends on different variables from LIDAR: while the lateral resolution is set by the camera lens focal length and pixel density, the depth resolution is a function of the lens focal length and aperture, the modulation clock frequency, and the total light exposure -- which is itself dependent on the exposure time, the number, power, and wavelength of the LEDs, and the distance to the measured object, which determines the attenuation from both geometric spread and absorption.

To keep track of all these variables, we developed a software library to model the performance of an underwater ToF camera (shown in **Fig. 1**), exploring different combinations of lenses, lighting, and exposure parameters, to show that, in theory, this approach could satisfy the resolution and operational requirements for cave mapping. Our results were in some ways disconcerting - a complete 360 degree field of view requires 8-12 individual sensors, 100-200 LEDs, and active exposure compensation to capture measurements in optimal

dynamic range - but overall, the design would still win significantly in size, power consumption, and price compared to a scanning LIDAR system.

Figure 1. Screenshot from ToF camera design study. The Input Parameters (left column) allow the user to change the ToF image sensor, camera lens, LED, and exposure parameters. The Design Output (center column) displays various figures of interest for camera resolution, power consumption, and data generation, while the Graphs (right column) plot the lateral resolution (top graph) and brightness (bottom graph) vs distance to target. The horizontal line in the bottom graph shows the image brightness which achieves the best distance resolution.

Hardware Development

We began to outline the system of cameras, lights, optics, sensors, electronics, and computation that would be needed, and sketch designs for the instrument. Our models helped us anticipate the challenges that lay ahead. The infrastructure to manage the operation and data output of the array of cameras would require sophisticated design, as would packing all of this hardware into such a small instrument. Yet the full system requirements could only be discovered by collecting data with ToF cameras to discover how to optimize their performance in this novel application. Accordingly, we built two prototypes, the first a simple platform to perform basic tests with ToF cameras, the second a fully-featured system to pilot solutions to key integration challenges and expand our experimental capabilities.

Our first prototype helped to identify a key engineering challenge to optimizing the ToF camera performance. In assembling the core components of the ToF camera, we quickly discovered that the electronics to control the LEDs would need to be custom designed, and that their performance played a critical role in achieving the camera's maximum resolution. We would more gradually come to appreciate the deep challenge to optimizing this circuitry, as we went through a number of iterations in its design.

Illumination Control

The LED control electronics have a simple-to-state function: steadily flash a bank of high-power LEDs on and off ten million times a second. Doing so, however, is not so simple. The literature suggested that a basic shunt

switching circuit topology would work well up to 100 MHz; when we tested such a circuit, however, it struggled to work at one tenth the speed, visibly dimming when we increased the switching frequency above a few MHz. We measured the optical rise time to be 50 ns, so slow that the LEDs never reached full power when switched at 10 MHz. Why was the rise time was so slow, how could we speed it up, and, crucially, how fast did it need to be?

To answer the last question, we analyzed the phase delay measurement algorithm performed by the camera to understand the impact of illumination waveform on measurement resolution. Ideally, the light output would look like a square wave, on at full power for half the period, and off completely for the other half; this output was assumed in order to simplify the models of the design study. We found that when the output waveform is not ideal, the distance resolution becomes distance dependent. Essentially, time along the output waveform can be mapped onto a distance from the camera, and the resolution at that distance depends on the amplitude of the waveform there. Slow transitions create regions of poor resolution; the slower the transition, the larger the region. While faster transitions are always better, we determined that output rise times of 10 ns would be sufficient. (If necessary, one can shift the location of the transition regions by delaying the LED clock, and take multiple images with different delays to cover the full space at maximum resolution).

We expected to find that our optical rise time was slow due to slow current switching. The parasitic inductance of the wiring to the LEDs acts as a low-pass filter - one which the shunt switching topology could only partially mitigate, as shown in **Fig. 2A** (in the next iteration, we made flexible circuit boards for the LEDs to minimize this inductance). Indeed, we found that the rise time of the voltage across the LEDs was about 20 ns, roughly ten times slower than the reference circuit. Still, the rise time of the optical output was significantly slower. We knew that LEDs have an intrinsic capacitance that acts as a low-pass filter between power input and light output, but we had not expected the filtering to be so strong. However, this capacitance scales with LED output capacity, and we were using the most powerful blue LEDs available. So, to drive these LEDs faster than 10 ns, we would need to modify the controller circuitry.

Figure 2 - LED Control Circuit Current Shunt Topology. The LED controller utilizes a constant current source to power the LEDs, and a transistor in parallel with the LEDs to turn them off. The current source is buffered by a high speed inductor, which mitigates the filtering effects of wiring inductance when the current switches between the LEDs and the transistor. The circuit is very efficient; when the LEDs are switched off, the power consumption is very low, due to the high conductivity of the transistor. The original circuit **(A)** suffered from slow light output switching, due to the intrinsic capacitance of the power LEDs. A passive current filter **(B)** significantly decreased the rise time, at the cost of increased power consumption through the resistor. Our novel controller topology **(C)** utilized a second current source and switching transistor (right side) to provide even faster rise times with higher efficiency. In all cases, the driver was tested with a series of 6-8 high-power blue LEDs. The original current source was set to 0.5 A. When the filter was added, the current source was increased to 1 A. In the two-source topology, the main current was set back to 0.5 A, and the peaking current provided an additional 1 A.

The general strategy for equalizing the low-pass filtering of the light output is to high-pass filter the input signal. For our first adjustment, we added an inductor-resistor current filter in parallel with the LEDs (**Fig. 2B**).

By tuning the value of the two passive components, we achieved a 16 ns rise time. The improvement was significant, but increased speed came at the cost of increased power consumption, as all of the current routed through the filter was wasted. A different circuit topology would be necessary to increase both the speed and efficiency of the LED controller.

Another way to think about equalizing the light output is that, for a given steady-state LED current, there will be a constant charge built up in the LED capacitance; the LED controller should use as large a current as possible to fill that capacitance, and then switch to the desired steady state current. Our circuit already contained one efficient constant-current source; we simply copied it to provide the initial boost of current (**Fig. 2C**). The size of this current was set by the maximum transient current of the LEDs; by tuning its duration, we achieved sub-10 ns rise times. However, the output waveform did not remain steady after the initial rise; it experienced a dip when the peaking current turned off, and damped oscillations thereafter. This new circuit topology was proving promising, but still needed tuning.

The third iteration of our controller circuit replaced key components to address the source of the output oscillations. We found that the dips in output power when the peak current turned off were due to a brief shunting of current from the main source into the peak shunt transistor. This current path was blocked by a Schottky diode (**Fig. 2C**), but during the peak current shutoff, the parasitic capacitance of the diode allowed current to flow in the wrong direction. Similarly, the oscillations were due to the drain-source capacitance of the main transistor resonating with the wiring inductance to the LEDs. While some amount of capacitance is inevitable in any rectifier, this capacitance is much lower in some advanced semiconductor materials. By substituting the silicon components with newly released upgrades (a silicon carbide Schottky diode and gallium nitride MOSFET) we were able to substantially stabilize the light output. We settled on this solution as an optimal balance of performance, cost, and complexity.

This novel circuit serves to substantially expand the application space of ToF cameras. The optical power required by a ToF camera scales with the required resolution and the volume being measured. Standard applications tend to either focus on small volumes (e.g. for industrial automation) or rough measurements (e.g. for gesture recognition). The small LEDs or lasers used in these applications switch fast enough that equalization is unnecessary. Where work on fast-pulsed LEDs has been done in other applications (e.g. for optical communication), the use case was not power-limited, and so our circuit topology has, to our knowledge, never been reported in scientific literature.

Figure 3 - Design and Performance of LED Controller Iterations. Original circuit (A), modified with passive current filter (B), and novel circuit (C), with improved components (D).

Top image in each box shows a complete controller circuit board, and the bottom image shows test results. Note that the time window varies across test results; the original circuit (A) was tested at 2 MHz, but all other circuits were tested at 10 MHz. The measured rise/fall times for each circuit were (A) 48 ns / 24 ns (Green trace), (B) 16 ns / 8 ns (Green trace), (C) 7 ns / 16 ns (Purple trace), and (D) 8 ns / 10 ns (Purple trace). Additional traces show the input clock (square wave in all results) and LED voltage (Yellow trace in (A) and (C), blue trace in (D)).

System Design

Our second prototype included all of the systems necessary for self-contained field testing. Built in an 8 inch diameter, 12 inch long custom pressure housing, it featured two ToF cameras with overlapping 45° x 30° fields of view. It included on-board processing, auxiliary sensing, power storage and distribution, and a basic user interface, which displayed images and user settings, which the user could control with buttons mounted on the end of the housing. We developed software to take fine-grained control of camera exposure settings and capture routines, essential to conducting calibration experiments with the camera in the lab and optimizing image quality in the field.

Figure 4 - Prototype testing in a pool.

This prototype gave us a platform to work on a number of system integration challenges the instrument needs solved. Examples seen in **Fig. 5** include:

- fixing the geometry of the LEDs to create an even field of illumination, as well as providing heat sinking for the LEDs and controllers
- designing the housing to eliminate any internal light leakage pathways between LEDs and cameras
- synchronizing the camera and LED controller clocks so that they can operate simultaneously

Figure 5 - Prototype chassis during assembly. *LEDs on bottom. Internal baffle. Cameras up top*

Calibration

As with any sensor, calibration of the ToF camera is essential to extracting accurate distance measurements from its phase delay measurements. To achieve cm scale accuracy, the camera must calculate the time of flight with <100 ps error, corresponding to a phase estimate accurate within 1 part in 10,000. The calibration must account for the geometry of the cameras and LEDs, the precise illumination waveform generated by the LED controller, and pixel-by-pixel variability of sensitivity and signal propagation delays in the front-end electronics.

The influence of all these variables on the phase measurement is too complex to effectively model, so calibration must be performed with a complete camera system, in the configuration it is going to be operated. We have developed calibration routines for lab measurements in confined volumes, but the instrument will need an efficient method for underwater calibration through the complete measurement volume.

Figure 6 - A particular challenge of the calibration is that signal propagation speeds on the image sensor are very sensitive to temperature, and the power consumption of the image sensor produces significant heat during operation. With changing imaging routines, the sensor temperature can change >5 C/s, causing time delays equating to ~ 10 cm distance errors. It is possible to compensate for these delays if the temperature is known, but the heating of the sensor is nonuniform, and the sensor only provides four points of temperature measurement, so the temperature gradients across the chip are fundamentally unmeasurable. Thus, the temperature of the sensor must be controlled in order for calibrations to be valid.

Field Testing

Mexico - August, 2018

Our first trip to the field was motivated by a unique opportunity for the MIT team to meet Corey Jaskolski in August of 2018 while he was exploring a recently discovered cenote outside of Chichen Itza. When the trip was planned, we had completed the design study and the first iteration of the LED driver, and we were using a demo system camera from a ToF sensor company to run basic imaging tests. We quickly built a minimum viable field-capable prototype (shown on the site of the cenote in **Fig. 6**). Its components were split between two pressure housings, one containing the camera, batteries, and an ethernet buffer, the other, the LEDs and controller. The camera was controlled by an external computer, tethered by a long ethernet cable.

To operate the camera in the cenote, one member of the team remained on the surface while Corey was lowered through a narrow vent in the earth to the surface of the water, 80 feet fellow. Communicating over acoustic underwater telephones, we took images of artifacts submerged on an underwater shelf, and took videos of passes along the cave wall. Unfortunately, the image quality was poor; although we tried to minimize direct light leakage between the LEDs and camera by enclosing them in separate housings, the clear acrylic walls left leakage paths open. Yet in many ways the trip was a success, giving the MIT team a chance work directly with Corey and gain experience in conducting a field deployment.

Figure 7 - The first prototype, on the Yucatan Peninsula. The prototype sits on the yellow case, and MIT team member Jake Bernstein works on connecting the ethernet tether to the surface camera control center.

Channel Islands - July, 2019

Our second field trip was an opportunity to test the second prototype on an expedition with Kenny Broad exploring sea caves off the coast of California. The MIT team gained more valuable experience preparing for deployments and debugging hardware in the field. Fixing a buggy bit of embedded software delayed their arrival to the expedition, and patching a leak in the pressure housing scratched a day of dives (the leak was discovered when the pressure housing, which had been successfully tested in the school pool, was reassembled for deployment). As a result, the prototype was only brought on one dive, with Kenny exploring the sea bottom in the vicinity of the dive boat. The prototype captured a stream of images at a constant frame rate and shutter speed for roughly ten minutes.

These images help demonstrate the challenges and potential of the ToF camera. The patchwork areas of the rocks indicate that local areas with small variation in depth can result in large variations in measurement quality. Regardless of why that happened here - perhaps the rock was rough there, or perhaps parts of the rock were covered in algae that absorb blue light - the fact of these variations has important implications for the operation of the camera. Since the camera achieves optimal resolution in a narrow band of exposure, and the

brightness of surfaces across the scene will vary greatly, the camera will need to take a series of images of each scene, sweeping through a range of shutter speeds to properly expose the entirety.

Next Steps

At this point the core technology is well developed. The key challenges that remain to tackle are calibration issues (including thermal compensation across the imaging sensor), the software pipeline (stitching the point clouds in real time at low power, techniques in which our collaborators have significant expertise), and a third hardware revision focusing on cave-specific deployments. We also need to rethink the user interface to make the system more usable for divers in the field.

As must be clear by now, this problem proved far, far more complex than our initial proposal estimated, and we have spent far, far more resources on the project than initially expected. By the same token, this project has taught us far more than we would have guessed, and has opened doors on underwater imaging that we would otherwise never have reached for. We are excited to see where Prometheus leads us next.

Acknowledgements

Project Prometheus was funded by the National Geographic Society.