https://doi.org/10.2352/ISSN.2470-1173.2020.7.ISS-227

TunnelCAM- A HDR Spherical Camera Array for Structural Integrity Assessments of Dam Interiors

Dominique E. Meyer [1], ***Eric Lo*** [1], ***Jonathan Klingspon*** [1], ***Anton Netchaev*** [2], ***Charles Ellison*** [2] ***and Falko Kuester*** [1]
[1]Dronelab, University of California San Diego; La Jolla, CA
[2]U.S. Army Corps of Engineers, Engineer Research and Development Center

Abstract

The United States of America has an estimate of 84,000 dams of which approximately 15,500 are rated as high-risk as of 2016. Recurrent geological and structural health changes require dam assets to be subject to continuous structural monitoring, assessment and restoration. The objective of the developed system is targeted at evaluating the feasibility for standardization in remote, digital inspections of the outflow works of such assets to replace human visual inspections.

This work proposes both a mobile inspection platform and an image processing pipeline to reconstruct 3D models of the outflow tunnel and gates of dams for structural defect identification. We begin by presenting the imaging system with consideration to lighting conditions and acquisition strategies. We then propose and formulate global optimization constraints that optimize system poses and geometric estimates of the environment. Following that, we present a RANSAC frame-work that fits geometric cylinder primitives for texture projection and geometric deviation, as well as an interactive annotation frame-work for 3D anomaly marking. Results of the system and processing are demonstrated at the Blue Mountain Dam, Arkansas and the F.E. Walter Dam, Pennsylvania.

Problem Statement

Dam safety relates to the risk a dam imposes in the case of structural failure to the environment in the path of flooding. A dam can be rated ”safe”, ”unsafe” and ”emergency unsafe” (critical condition) according to a structural evaluation. The status ”unsafe” is generally achieved when the dam is declared inadequate to handle a Probable Maximum Flood (PMF)[13], or when structural failure is likely during geological and seismic activity [4]. A set of dams have been subject to damage due to excessive loading, geophysical phenomena or have suffered aging with unknown effects on integrity [4]. A primary failure point are the flood gates which are used to control the water flow. For dams that have water outlet channels that need to be inspected or the gates themselves, inspectors are put at direct risk of the holding capability of the gates. Materials used in construction of dam outlet works are vulnerable to corrosion and fatigue issues that may compromise the structural integrity of the system and put personnel entering the conduit/tunnel at risk. US Army Corps of Engineers (USACE) maintains many of these aging structures in need of an updated inspection, but placing personnel, including inspection personnel, inside dam outlet works of a facility in an unknown or deteriorated condition presents unacceptable risk to personnel and does not comply with current safety requirements. This motivates the development of remote imaging systems that are deployed on remotely tethered and controlled rovers or on fully autonomous systems.

Informal and formal inspections are encouraged by Federal agencies to note change in visible and measurable features [5]. Visual inspection of accessible dam interiors contribute to a ma-

Figure 1: TunnelCAM deployed to inspect the water evacuation tunnel at the Francis E. Walter Dam, Philadelphia, USA

jor component to asset assessment and is conducted by trained individuals. Thorough documentation strategies are encouraged by federal agencies [5], however it is common to only take low-resolution focused on defects and rely on hand notes for anomaly location. Research and Development efforts are evaluating high-resolution 3D structure documentation for standardization in inspections and temporal data analysis.

We summarize the project into a set of hypotheses that drive both methodology and system.

- It is possible to complete remote digital inspections of dam outflows
- Data reconstruction can achieve human equivalent acuity
- Geometric deviation is possible to be applied to outflow tunnel evaluation
- Digital models can serve as Building Information Model (BIM) records
- The imaging system requires self-illumination and HDR
- Geometric alignment & accuracy can be achieved with only exterior GCPs
- We can constrain inter-camera pose to accelerate and improve reconstruction
- We can initialize pose using VIO from the Realsense T265 camera

The development of the TunnelCAM imaging system is a result of collaborative efforts between the University of California San Diego Center of Advanced Imaging and the the U.S. Army Corps of Engineers, Engineer Research and Development Center to evaluate diagnostic imaging of dam interiors.

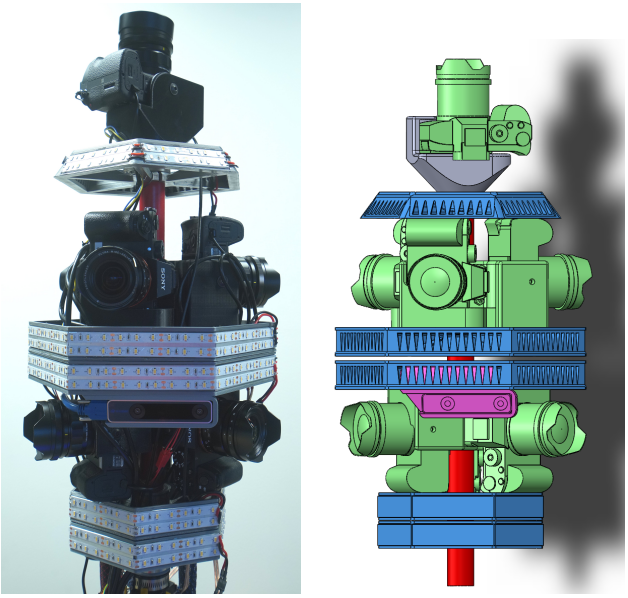

Figure 2: TunnelCAM system, left: as-built, right: CAD model

TunnelCAM System

In this section we will briefly describe the TunnelCAM system as it was built. The motivation behind this design is to facilitate a user friendly integration of consumer grade mirror-less cameras while maintaining synchronization of the sensors to satisfy the timing constraints imposed by the post processing stages of the workflow. We highlight the engineering details in support of 3 assumptions we define for the reconstruction algorithm: 1) it is reasonable to assume sensor synchronization to $<$1 millisecond such that all cameras have been subject to the same rigid-body transformations τ between subsequent captures; 2) intrinsic parameters remain constant for every respective camera over capture runs and 3) the relative pose between cameras remains constant for all captures due to the rigid mechanical assembly of the system.

During the reconstruction stage, the captured image data is processed through a structure from motion pipeline that allows us to estimate the intrinsic & extrinsic parameters for every camera at every location, as well as the geometry of the scene itself. While sparse scene estimation is done using a sparse subset of features, we employ dense methods to achieve model resolutions required for geometric and texture based inspection.

System Description

The present TunnelCam system is comprised of seven Sony A7R II mirror-less digital cameras [2], arranged in two offset rows of three cameras and one upward-facing camera, as shown in green in Fig. 2; a configuration that aptly accommodates for spherical panoramic image capture while keeping the mass of the system as close to the center as possible. The back illuminated CMOS sensors provide a maximum image resolution of 42.2 Megapixels at an aspect ratio of 3:2. Each camera is fitted with a Voigtlander Ultra Wide Heliar 12mm Aspherical III lens, rendering rectilinear images with a diagonal Field of View(FoV) of 121°. This sensor-lens combination provides a visual acuity of 3.77 millimeters per pixel at a focal distance of 10 meters, exceeding human visual acuity [16]. The choice of a rectilinear lens was to minimize radial distortion compared to a fisheye lens known to cause issues when not properly modeled in intrinsic rectification [11]. The image sensor of the a7R II provided improves signal to noise ratios for amplification of electrical signals from the pixels. This is a key characteristic when acquiring low-light images, especially when exposure times need to be kept at $<$1/100th of a second for minimizing motion blur [17]. Driving the choice of consumer grade cameras was the availability of sensor shifting technology provided by Sony. A 5-axis stabilization allows for the sensor to be shifted such as to compensate for motion during exposure time. While this is of great importance for low-light capture where the integration time is large enough to affect blurring. This does however open questions about how reasonable it is to uphold assumption 2 as individual

Five double rows of PWM dimmable LED rings are used on the system to provide a total luminous flux of approximately 32760 lumens, as shown as the blue features on Fig. 2. For the cameras pointed directly at the tunnel walls, the illumination is sufficient to completely illuminate the field of view to within 2 stops, so it is reasonable to push the exposure in the darker areas while retaining detail. For the off-axis cameras, exposure bracketing for longer exposures can compensate for the decreased illumination due to increased distance further down the tunnel. Due to the reflective nature of the wet tunnel, these off axis viewpoints are also important to minimize direct specular reflection when the LEDs reflect directly into the cameras shooting normal to the tunnel surface. We can approximate the luminosity distribution in the inside of the tunnel by assuming the LED on the TunnelCAM to

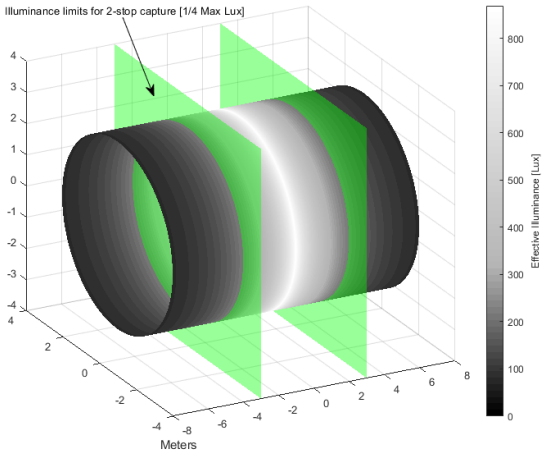

Figure 3: Simulated illuminance of a tunnel with 7 meter diameter and central point light source of 32760 lumens

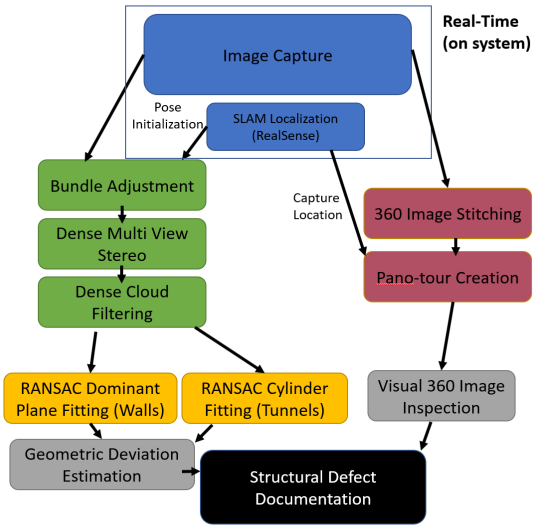

Figure 4: TunnelCAM Data Flow for Capture and Processing

be represented as a single point source uniformly radiating light. Fig. 3 shows this distribution along with a set of planes which indicate the fall-off limit equal to 2-stops. These 2 stops are equivalent to a quarter of the luminance from the inside of the tunnel. We consider this to be the area which is sufficiently illuminated to correctly fill the dynamic range of the sensors.

Finally, a single Intel RealSense T265 camera [3] is positioned on one side of TunnelCam as seen in pink in Fig. 2, enabling ego-motion estimation by means of stereo visual odometry. The Realsense T265 camera uses an internal Simultaneous Localization and Mapping (SLAM) algorithm with hardware acceleration to estimate the pose of the system [3]. This pose estimate is used purely to initialize our optimization and may be omitted in future revisions. Synchronization of the cameras is achieved through the cameras' USB PTP interface connected to a small-form-factor Intel Next Unit of Computing (NUC). Necessary capture settings, such as bracketing HDR modes and exposures, are managed by a custom backend based on libgphoto2. The NUC is also responsible to aggregate and tag the pose estimation from the RealSense camera with the respective captures of the Sony A7Rii cameras. Captures are triggered by the GPIO output of an Arduino Nano microcontroller, and can be induced by physical trigger remote input or serial signals received from the NUC.

Acquisition Methodology

Acquisition runs of the TunnelCAM involve hand-carrying as in Fig. 1 or rover mounting the system in the dam interiors. The camera rig is first exposure set and LED power adjusted at the dam tunnel entrance to ensure proper illumination of the surrounding geometry through a live camera feed. Once initialized adequately, the synchronized capture run of both the Realsense and Sony cameras is begun through the NUC. The high-resolution Sony sensors are manually triggered every time the system is moved by a desired distance. This trigger commands the sensors to complete 3,5 or 7 exposures with bracketing (HDR) while simultaneously instructing the NUC to register the pose from the RealSense camera with that trigger. This is repeated for the extent of the scene with an approximate translation of 20 centimeters between captures. The high-resolution Sony frames result in being captured at 5-10 sec intervals while the RealSense camera runs uninterrupted at 60 Hertz. Dam interior inspections are normally limited to sub-hour periods to reduce stress build-up of accumulating water on the reservoir side [4], which puts pressure for rapid imaging of long interior sections. Changes in geometry may require the system operator to adjust LED power and exposure times during the course of the acquisition.

Processing Pipeline

Data is stored locally on the system while the acquisition is running however no full 3D reconstruction is run in real-time. After the acquisition, we can segment the processing pipeline in two main parts: a multi-view dense reconstruction and a naive panorama stitching pipeline. Fig. 4 highlights the use of the pose estimates derived by the RealSense camera for both the initialization in the bundle adjustment as well as the approximate position used to generate the 360 photo-sphere tours. Panorama image stitches generally assume perfect rotation about the focal point of the cameras [18], which does not hold true for this system due to the physical separation of the cameras. As such, we expose this method of processing and viewing the data to collaborators for ease of generation but emphasize the benefits of full multi-view reconstructions. This section will discuss the constrained localization method leveraged with initialized pose parameters and the densification step used to generate dense 3D point clouds.

Preprocessing

A set of image pre-processing tasks are completed prior to the multi-view reconstruction part of the pipeline. Imaging in dark environments with a lighting source introduces the issue of deciding what exposure to choose such as to optimally represent the lighting difference across the sensor with a finite number of bits. The sensors used achieve 14-bit RAW bit depth [2]. The use of multiple exposures enable the use of tone mapping [8] to reduce the multiple frames (3,5 or 7) of 14-bits into a single frame with darker areas stretched for increased luminosity and brighter areas darkened to prevent saturation. Tomasi et al [14] proposed a bilateral filtering algorithm which is what we use to merge the

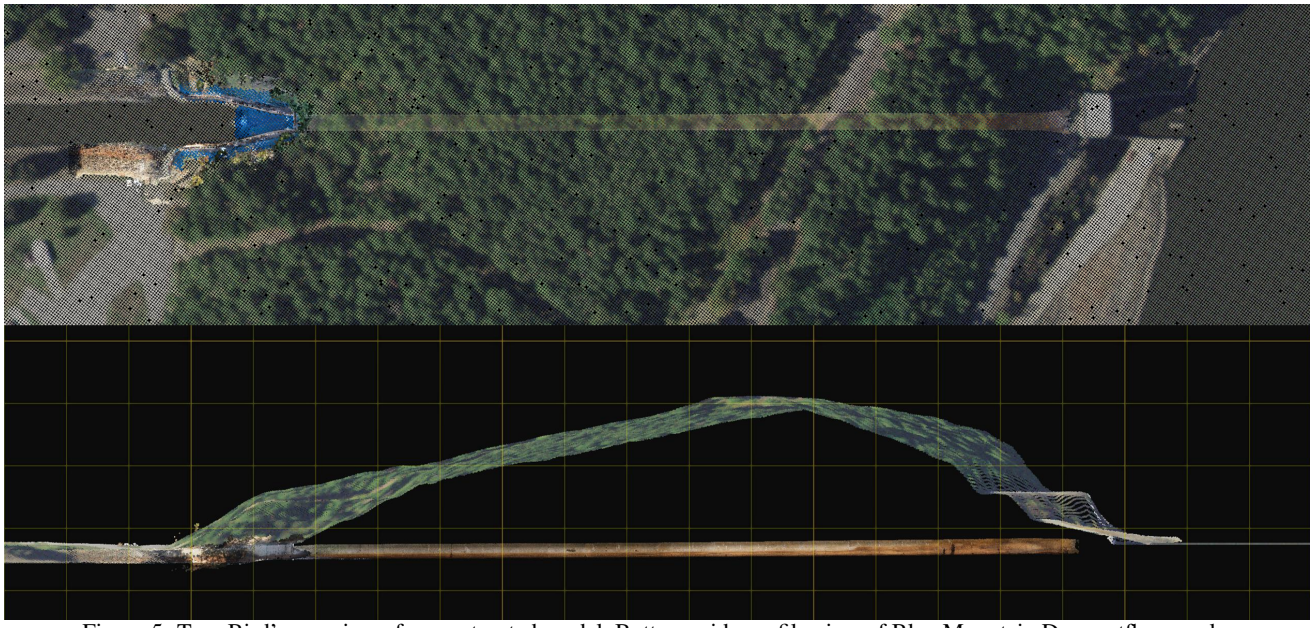

Figure 5: Top: Bird's eye view of reconstructed model, Bottom: side profile view of Blue Mountain Dam outflow works

multi-exposure frame sets. We further compress the images to 12-bit for ease of processing.

Photogrammetric Reconstruction

Photogrammetry is known for the generalized problem of simultaneously solving for the 3D location of a set of points which are observed from N cameras where each camera can be characterized with camera matrix P [6]. SLAM solves for sensor pose and environment geometry [12]. SLAM algorithms are generally targeted at applications where image and IMU data is sequentially acquired and where results are expected in real-time or near real-time on the acquisition platform. Similar non-linear optimization techniques are exploited in sequential bundles (windowed Bundle Adjustment) [10] when compared to Photogrammetry methods which leverage optimizations over the full set of bundles.

Acquisition of the data is completed at intervals that are temporally constant, but irregular motion yields in slightly different rigid body transforms between captures. The TunnelCAM centering within the cylindrical tunnels, the spacing between captures, and the rotation of the system are all subject to random human error. We hence have to derive the the 6- dof transformation τ_j between subsequent captures $\{j = 0, 1, ..., k\}$ in the same the world coordinate system. Since the cameras are mechanically constrained we assume rigid body motion of the system and consider the relative camera extrinsic parameters R_i, t_i to be constant for cameras $\{i = 1, ..., 7\}$.

From a set of observed image features x^i_j with corresponding 3D points X_j viewed from the set of cameras with camera matrices P^i over a set of captures, we minimize the re-projection error in 1 using a Levenberg Marquardt algorithm [9].

$$\sum_{i,j} ||x^i_j - (M^i X_j + t^i)||^2 \tag{1}$$

Initial pose estimation can be obtained from onboard localization systems such as the Realsense camera that computes VIO, it is also possible to compute it from general photogrammetry through software packages like Agisoft Metashape [1]. Differences in pose estimate performance will be later discussed in the results section.

The goal is to achieve geometric and photometric reconstruction of the outflow works at resolutions that are equivalent or greater than human visual acuity. As the camera system resolves at 3.77mm/pixel resolution at 10m, slightly better than human acuity, it is necessary that reconstruction outputs have similar number of world points as image points (unique pixel count over overlapping areas) $X_j \sim x^i_j$. We leverage the Semi-Global Matching (SGM) algorithm [7] on image pairs for the dense reconstruction of the tunnel.For every camera we combine the frame with the frame from the previous capture to form a stereo pair. This pair is rectified to obey epipolar constraints between and the densification algorithm is run. The subsequent depths for every respective pixel are then projected with the intrinsic and extrinsic camera parameters to 3D world points.

Primitive Fitting

The resulting reconstruction is in the form of a dense point cloud where each point has a set of parameters - $X_j \in \{\{x, y, z\}, \{r, g, b\}, \{n_x, n_y, n_z\}, c\}$, which are location, color, point normals and point confidence respectively. The point confidence is a value estimated from the number of image pairs that contributed to the estimation of that point. We leverage the MLESAC algorithm proposed by Torr et al. [15] to estimate a cylinder to the dense cloud of the tunnel. The cylinder axis of the tunnel is initialized by Principal Component Analysis (PCA) of the set of dense points. Improved fitting is achieved through prior thresholding of points that have c below a certain user-defined value. This generally eliminates distant point estimates from poses that deviate further from the tunnel geometry and can be considered as noise in the dense reconstruction. The resulting cylinder equation is used to calculate a point-to-cylinder distance for every point in the model, that serves as our deviation estimate. All points are then remapped into cylindrical coordinates to form a 2D cylindri-

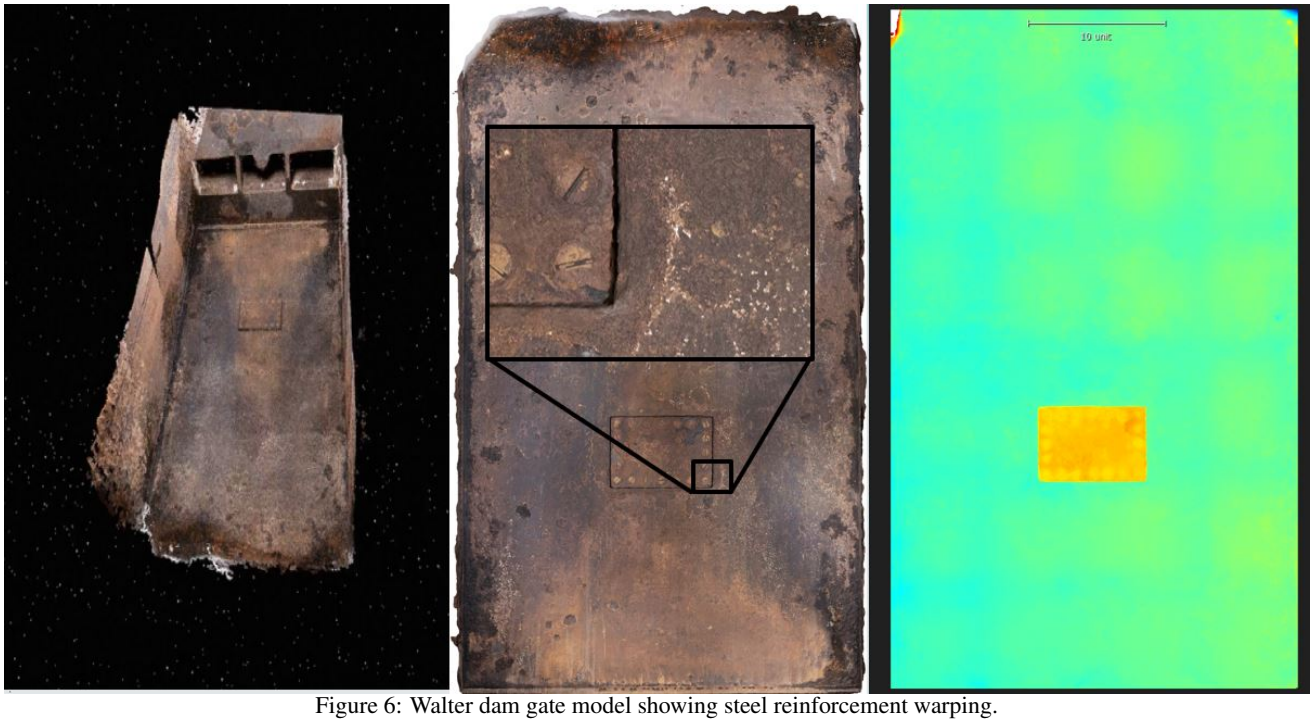

Figure 6: Walter dam gate model showing steel reinforcement warping.

cal image of the tunnel, with RGB and distance values for every point.

Results

This section provides results from the TunnelCAM system and proposed processing pipeline deployed at the F.E. Walter Dam in Philadelphia, USA and the Blue Mountain Dam in Arkansas, USA. The first served dam tests served primarily as an operation test, and provided gate imaging results, while the Blue Mountain dam allowed for imaging of the whole outlet works. A total of 7504 images were captured across the 7 cameras during the imaging of the 300m Blue Mountain Dam tunnel. The captures were set with a regular acquisition time interval of 5 sec, while the system operators moved the system between subsequent captures. The system was acquiring data starting at the tunnel entrance and during it's path in and out. At the end of the tunnel were 3 steel gates that partially leaked water and prevented close-up imaging. During acquisition, a live-stream served for a remote operator to adjust exposure timing on the sensors and ensure that all captures were synchronized.

The Blue Mountain Dam outflow works were also surveyed in with a Faro ground LiDAR that was georeferenced with ground control points located outside the tunnel. The photogrammetry model derived from the TunnelCAM was evaluated to a known set of points in the LiDAR model and the results are shown in Figure 7. It can be seen that the geometric consistency was on the order of 0.17m throughout the length of the tunnel. This reflects global accuracy of the model rather than the models local consistency that shows geometric resolutions to a few mm.

The data processing pipeline used a set of image masks that were derived manually from the system after it was mounted on the boat vessel. This was to prevent the operators from being matched between frames and reduce mis-alignment instances. The Realsense localization of the TunnelCAM system failed, as seen in Figure 8, forcing the image processing without initial pose estimates. The photogrammetry alignment was able to localize all cameras with an average re-projection error of 2.31 pix, a spatial resolution of 10mm/px at 3m, and a total dense point cloud point count of 402 million points.

Label	X error (m)	Y error (m)	Z error (m)	Total (m)	Image (pix)
F	-0.0860723	0.131696	-0.0217955	0.158831	2.672 (721)
G	-0.117641	0.207652	-0.0153825	0.239155	1.868 (764)
H	-0.172667	0.20575	-0.0118351	0.268862	2.523 (50)
E	-0.0913591	0.178404	-0.0113083	0.200755	1.652 (4)
Q	0.106497	-0.105622	-0.00954919	0.150296	1.807 (706)
P	0.0760174	-0.0878003	0.00645889	0.116315	1.417 (781)
R	-0.033518	0.113241	-0.00928697	0.118462	13.549 (12)
J	-0.0588957	-0.0685505	-0.0177888	0.0921103	6.180 (155)
I	-0.0941688	0.0315672	-0.0195749	0.10123	4.470 (205)
K	-0.0677802	-0.157664	0.0824398	0.19039	10.705 (4)
L	0.157028	-0.152691	0.0261215	0.220578	0.877 (3)
M	0.0781087	-0.0310748	0.0103099	0.084693	0.741 (4)
N	0.155806	-0.128649	-0.000215355	0.202054	1.379 (2)
O	0.148654	-0.136256	-0.00859926	0.201836	1.742 (3)
Total	**0.110749**	**0.135294**	**0.0261279**	**0.176784**	**2.691**

Table 11. Control points.
X - Easting, Y - Northing, Z - Altitude.

Figure 7: Accuracy results from LiDAR reference.

The cylinder fitting for the dense point cloud of the tunnel section was relatively

The gate reconstruction was completed by moving the camera array around the gate without specific pattern, but by maximizing coverage. This guaranteed that every point was ovserved by at least 10 camera poses. Geometric and photometric recon-

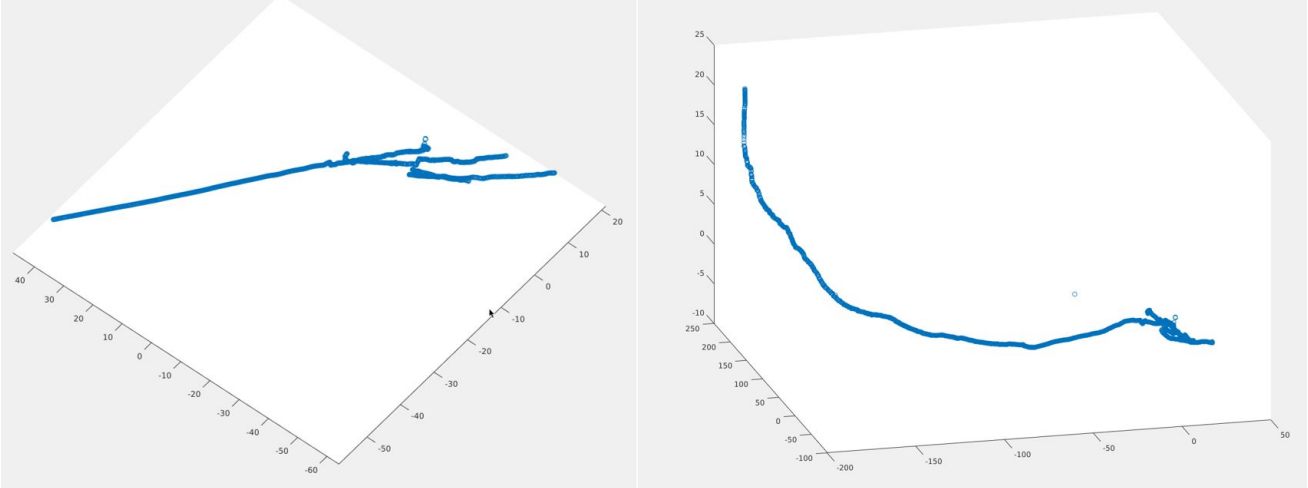

Figure 8: Mis-localized TunnelCAM system from Realsense localization unit.

struction results are seen in Figure 6. The plane deviation model clearly identified the grid pattern behind the gate that was subject to warping during the casting process.

Discussion and Conclusion

The results confirm that it is possible to achieve human equivalent inspection of outflow works through remote data capture using a system like TunnelCAM. While spatial resolutions of the reconstruction supports texture and local geometric consistency, the global alignment accuracy of 0.2m is indicative that the model is subject to warping and drifting. One potential issue is the lack of intrinsic camera parameter consistency due to the optical image stabilization. Furthermore, mis-matches in the point features could also have contributed to reduced alignment and reconstruction performance. Pose initialization can also be improved with better low-light sensors and algorithms that are more robust. This could also further increase the high-resolution reconstruction accuracy, and reduce alignment time. The primitive fitting algorithm showed larger deviations which were likely due to the tunnel not being a perfect cylinder. We foresee better deviation scaling if parametric surfaces would be used. Future work will also investigate alternative imaging techniques such as multi-spectral and X-ray transmission imaging for internal failure analysis.

To conclude, this work has addressed the automation of diagnostic imaging in dam outflow works for structural inspection and assessments. We have introduced a new self-illuminating camera array with 7 cameras to cover a full viewing-circle at a total resolution of 294 Megapixels. The array serves for handheld and rover mounted data collection inside dam tunnels. A novel processing pipeline is formulated that leverages the array to quickly derive inspection grade geometric and photometric cylindrical plans. We generate a dense 3D point cloud which allow for geometric analysis of the scene. We provide a complete data collection and processing results from the Blue Mountain dam as well as gate models from the the F.W. Walter dam in the month of October. The results confirm the feasibility to complete remote inspections, and that digital records can provide sufficient detail for digital documentation (BIMs) of such dam outflow works.

Acknowledgments

This publication is based on work supported by the US Army Corps of Engineers under research Cooperative Agreement W912HZ-17-2-0024, NIST Award #70NANB17H211, as well as NSF award #CNS-1338192, MRI: Development of Advanced Visualization Instrumentation for the Collaborative Exploration of Big Data. Additional support was provided by the Kinsella Expedition Fund, the Qualcomm Institute at UC San Diego and the World Cultural Heritage Society. We acknowledge the support from Luca De Vivo in deployment assistance, Danylo Drohobytsky for mechanical manufacturing and Vid Petrovic for the Point Cloud visualization framework. We thank all collaborators in Calit2, UC San Diego, as well as all other contributors to ideas, suggestions and comments of this work. Opinions, findings, and conclusions from this study are those of the authors and do not necessarily reflect the opinions of the research sponsors.

References

[1] LLC Agisoft. Agisoft metashape user manual, professional edition, version 1.5. *Agisoft LLC, St. Petersburg, Russia, from https://www. agisoft. com/pdf/metashape-pro_1_5_en. pdf, accessed June*, 2:2019, 2018.

[2] Anon. Sony a7r ii with back-illuminated full-frame image sensor.

[3] Anon. Tracking camera t265 – intel realsense depth and tracking cameras.

[4] Bruce Brand, David Dollar, Husein Hasan, Luis Hernandex, Larry Nuss, Rex Rowell, and William Wallace. Selecting Analytic Tools for Concrete Dams Address Key Events Along Potential Failure Mode Paths. Technical report, Federal Emergency Management Agency, 07 2014.

[5] Misc. FEMA. Federal Guidelines for Dam Safety. Technical report, Federal Emergency Management Agency, 04 2004.

[6] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.

[7] Heiko Hirschmüller. Semi-global matching-motivation, developments and applications. In *Photogrammetric Week*, volume 11, pages 173–184. Wichmann Verlag Heidelberg, Germany, 2011.

[8] Jiangtao Kuang, Hiroshi Yamaguchi, Changmeng Liu, Garrett M Johnson, and Mark D Fairchild. Evaluating hdr rendering algorithms. *ACM Transactions on Applied Perception (TAP)*, 4(2):9,

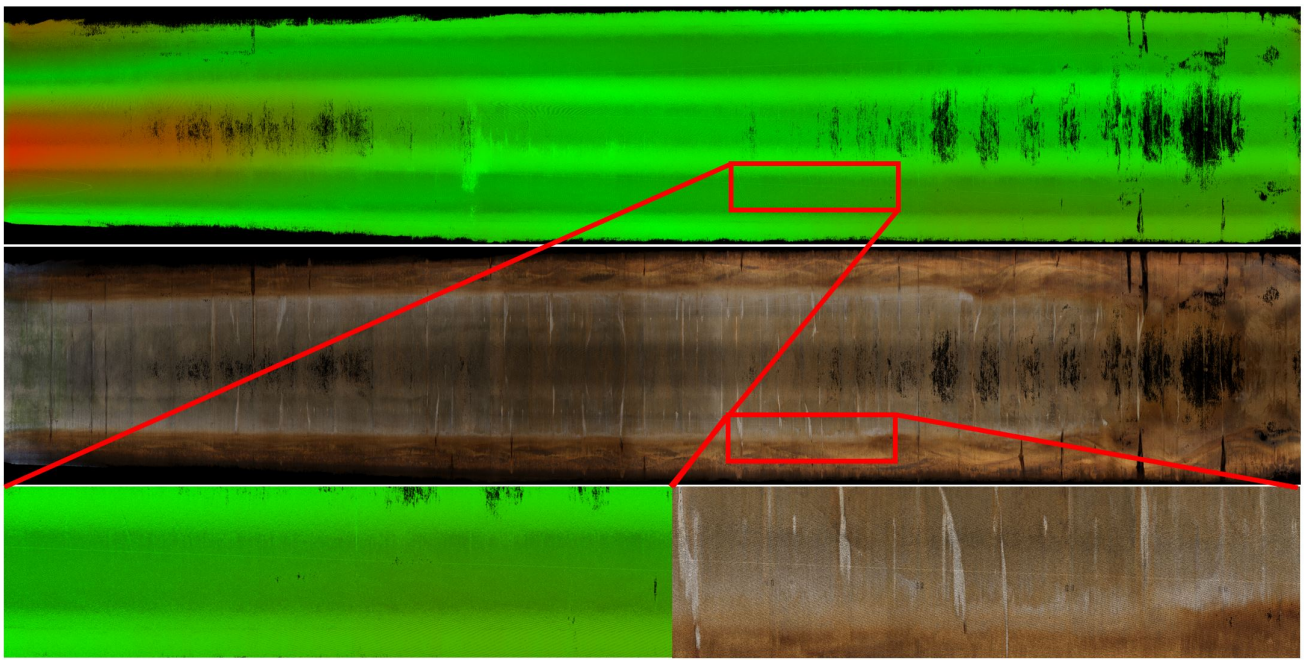

Figure 9: Cylindrical model representation of Blue Mountain Dam Tunnel. Deviation is in the range of 0-5cm.

2007.

[9] K Levenberg. Math. 2, 164 (1944). *Google Scholar DW Marquardt, SIAM J. Appl. Math*, 11:431, 1963.

[10] Julien Michot, Adrien Bartoli, and François Gaspard. Bi-objective bundle adjustment with application to multi-sensor slam. *3DPVT'10*, 3025, 2010.

[11] Luca Perfetti, Carlo Polari, Francesco Fassi, et al. Fisheye photogrammetry: tests and methodologies for the survey of narrow spaces. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 42(W3):573–580, 2017.

[12] Joao Machado Santos, David Portugal, and Rui P Rocha. An evaluation of 2d slam techniques available in robot operating system. In *2013 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 1–6. IEEE, 2013.

[13] Wilson H Tang and Ben Chie Yen. Dam safety inspection scheduling. *Journal of Hydraulic Engineering*, 117(2):214–229, 1991.

[14] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *Iccv*, volume 98, page 2, 1998.

[15] Philip HS Torr and Andrew Zisserman. Mlesac: A new robust estimator with application to estimating image geometry. *Computer vision and image understanding*, 78(1):138–156, 2000.

[16] CO Universal. Visual acuity measurement standard. *Visual Functions Committee*, 1984.

[17] Yitzhak Yitzhaky and Norman S Kopeika. Identification of blur parameters from motion blurred images. *Graphical models and image processing*, 59(5):310–320, 1997.

[18] Fan Zhang and Feng Liu. Parallax-tolerant image stitching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3262–3269, 2014.

Author Biography

Dominique E. Meyer is a PhD student at the University of California San Diego, where he previously received his B.S. in Physics in 2017 and his M.S. in Computer Science Engineering in 2019. He is working with Professor Kuester at the Dronelab with a research focuses on the development of real-time 3D imaging systems for robotics and 3D mapping.

Eric Lo is a Development Engineer at the Qualcomm Institute at UC San Diego. He received his B.S. in 2015 in the department of Electrical and Computer Engineering and the UC San Diego. His interest is in developing and optimizing autonomous robotic platforms for high resolution, large scale data acquisition, from centimeter level aerial mapping to micron level artifact scanning.

Jonathan Klingspon is an undergraduate student in the joint program of Mathematics and Computer Science at the University of California San Diego. His focus at the Center for Advanced Imaging and the dronelab is low-level programming for embedded camera systems. He has independently explored a range of subjects including computer vision, robotics control, and operating systems.

Falko Kuester received a M.S.E. in Mechanical Engineering and Computer Science and Engineering from the University of Michigan, Ann Arbor, in 1994 and 1995 respectively and a Ph.D. in Computer Science from the UC Davis in 2001. He is a professor at UC San Diego, director of the UC San Diego Center DroneLab as well as the Cultural Heritage Engineering Initiative, advancing research in robotics, remote imaging, large-scale visual analytics and virtual reality.

JOIN US AT THE NEXT EI!

IS&T International Symposium on

Electronic Imaging

SCIENCE AND TECHNOLOGY

Imaging across applications . . . Where industry and academia meet!

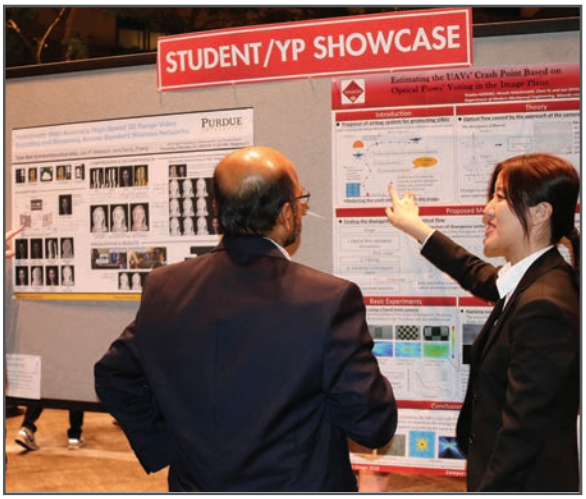

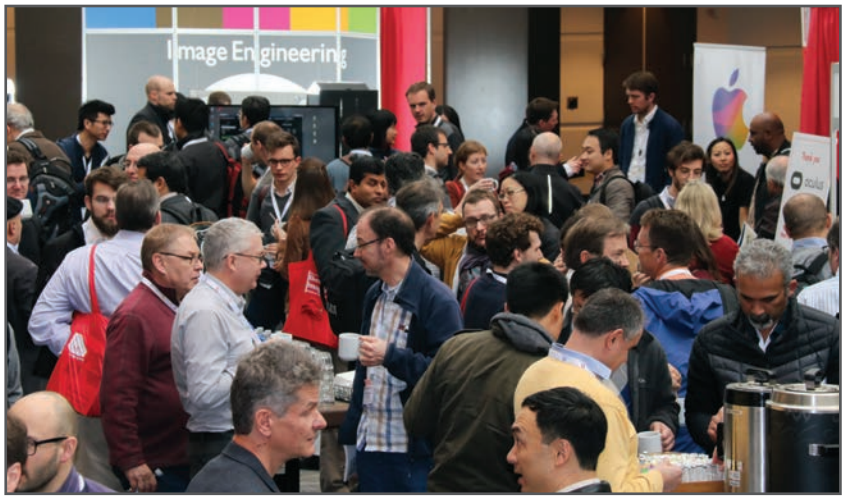

• SHORT COURSES • EXHIBITS • DEMONSTRATION SESSION • PLENARY TALKS •
• INTERACTIVE PAPER SESSION • SPECIAL EVENTS • TECHNICAL SESSIONS •

www.electronicimaging.org

IS&T
imaging.org