Compact high-resolution 3D real-time imaging for robotic vision with sub-mm accuracy

Thielemann, Jens T.
SINTEF AS
Oslo, Norway
jtt@sintef.no

Thorstensen, Jostein
SINTEF AS
Oslo, Norway
jostein.thorstensen@sintef.no

Dubanchet, Vincent
Thales Alenia Space
Cannes, France
vincent.dubanchet@thalesaleniaspace.com

Kirkhus, Trine
SINTEF AS
Oslo, Norway
trk@sintef.no

Hjelmervik, Jon
SINTEF AS
Oslo, Norway
jon.m.hjelmervik@sintef.no

Torben, Jakob
SINTEF AS
Oslo, Norway
jakob.torben@sintef.no

Ohl, Julia
Thales Alenia Space
Cannes, France
julia.ohl@thalesaleniaspace.com

***Abstract*— We have built and characterized a compact, simple, and flexible 3D camera based on interferometric fringe projection and stereo reconstruction. The camera uses multi-frame active stereo as basis for 3D reconstruction, providing full-field 3D images. We demonstrate real time 3D reconstruction in direct sunlight, with 3D reconstruction time < 100 milliseconds. Total 3D measurement noise is 0.3 mm at 0.6 m measured at per pixel-level. When employed for object localization, precision is 20 µm and 0.3 mrad at 500 mm distance. The system provides reliable object localization for relative motion rates up to at least 50 mm/second. Compliance with a wide set of flight materials and space robotics requirements has been investigated.**

Keywords—3D camera, object localization, in-orbit servicing, manufacturing, and construction

I. INTRODUCTION

There is a wide range of optical 3D sensing modalities available, that each offer specific trade-offs between data density, accuracy and working range (amongst other concerns). Structured Light and Active Stereo (multi-frame) systems can deliver high accuracy 3D data at short range, through projection of multiple patterns. This makes them ideal cameras for robot operations like inspection, manipulation or interaction with objects of both known and unknown shapes, e.g., for in-orbit servicing or planetary exploration.

Traditionally, structured light systems have tended to be bulkier and more complex than their Assisted Stereo (single frame) counterparts and use in sunlight has traditionally been difficult due to the relatively weak projected signal.

We present a novel 3D camera (Figure 1) design based on a stereo camera setup with an interferometric projector making the system simple and compact, while at the same time powerful enough to provide robust data even in sunlight.

The 3D-camera is designed for space-based applications. As opposed to ground-based, often indoor applications, such a use case comes with its own set of design constraints. A camera for space needs to be compact, robust, low-power and must be functional in direct sunlight. At the same time, the camera must deliver the excellent 3D data quality expected from state-of-the-art Structured Light 3D cameras, i.e. dense point clouds with per pixel noise on the order of 0.1 mm.

The work presented is an expansion of previous work [1], where we can now demonstrate real time 3D capture and sunlight resistance. We have also extended our results to not only report on the quality of the point clouds, but also how well the camera performs for object localization of objects in the scene. Testing at Thales Alenia Space's ROBY facilities indicates micro-meter level accuracy in object localization, significant motion tolerance (at least 50 mm/s), and compliance with many relevant space materials.

At the same time, no single 3D system can cover all use cases. For e.g. in-orbit servicing, we would expect 2-3 different systems to be necessary to cover the entire working range from initial localization to actual grasping and manipulation by a robotic end-effector. Table I presents a summarized trade-off of different technologies, and briefly outlines their relative strengths and weaknesses. As one can see, the presented system is particularly useful for close-range interactions, e.g. for the final phases of approach, or during the actual servicing of in-orbit satellites.

The system is not dependent on fiducial markers, making it capable of localizing both compliant and non-compliant objects, as well as providing high-resolution 3D maps for e.g. integrity monitoring. In addition, without the projector it can also be employed as a regular unassisted stereo camera, enabling longer range or lower power measurement modalities.

For covering the entire localization chain, our system could fit well together with e.g. a Flash LIDAR system that allows object localization far beyond the working range of our system. Together these systems can provide robust localization of both cooperative and non-cooperative clients.

The remaining paper is structured as follows: Chapter II and III outline the hardware and software implementation. Chapter IV and V describe our results in terms of testing for accuracy and in-orbit applications. Chapter VI concludes and summarizes the paper.

Figure 1: Actual camera. Pen placed in front as size reference.

TABLE I. INDICATIVE TRADE-OFF TABLE BETWEEN OPTICAL 3D LOCALIZATION TECHNOLOGIES

Technology	Working range	Accuracy	Data density	Sunlight tolerance	Power	Size	Computational load
Structured light (ours)	0.5 m -1.5 m	High	Dense	Good	Medium	Medium	High
Flash LIDAR [2]	2 m - 4 km	Low	Dense	Good	High	Large	Medium
SPAD Flash LIDAR [3], [4]	1 m -300 m	Low	Dense	Bad	Medium	Large	Medium
Scanning LIDAR[5]	1 m – 1.5 km	Medium	Dense	Good	High	Very large	Low
Stereo (unassisted) [6]	0.5 m -1.5 m	Medium	Sparse	Very good	Medium	Medium	High
2D cameras with markers	0.5 m -1.5 m	Medium	Very sparse	Good	Low	Small	Medium

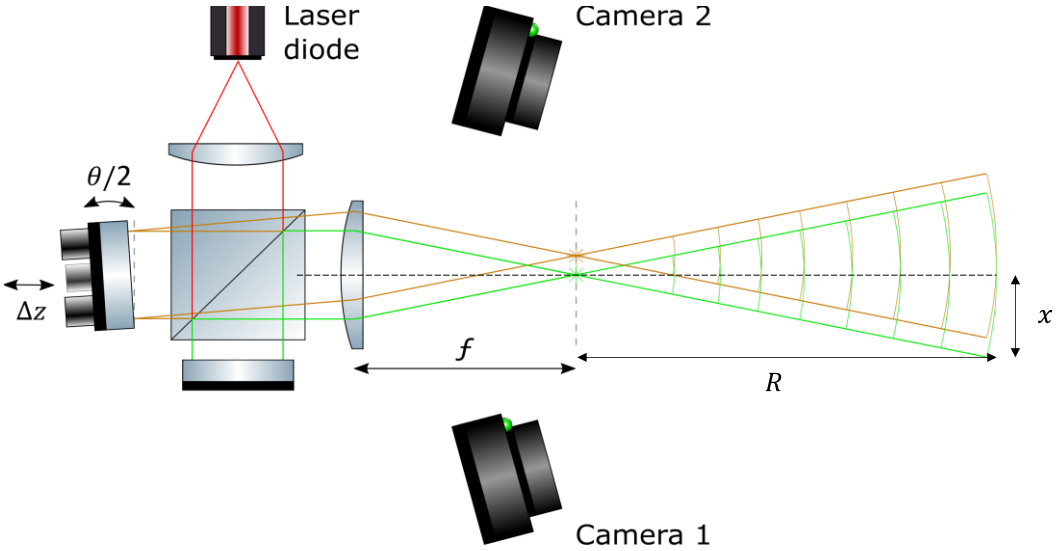

Figure 2: Schematic layout of the camera/interferometric projector.

II. OPTICAL CONCEPT

As noted above, a 3D camera for space-based applications comes with a quite particular set of design constraints, and one cannot simply re-use a system architecture that would work in an in-door setting. A camera for space needs to be compact, robust, low-power and must be functional in direct sunlight. At the same time, the camera must deliver the excellent 3D data quality expected from state-of-the-art Structured Light 3D cameras, i.e. dense point clouds with per pixel measurement noise (precision) on the order of 0.1 mm.

Starting with the requirement of low-power operation in sunlight, it becomes apparent that laser illumination combined with spectral filtering is required. Compactness and robustness preclude the use of DMD based projection. Instead, we have developed an Active Stereo 3D camera with a projector based on a Michelson interferometer and a movable micro-mirror using laser illumination. This projector design can be made exceedingly compact, while the use of laser illumination combined with sunlight filtering allows us to project patterns that can outshine the sunlight background, even with moderate power consumption.

The optical layout of the camera is shown schematically in Figure 2, and described in detail in [1]. In short, a collimated laser beam is sent through a Michelson interferometer where one arm carries a MEMS micro-mirror which is movable in tip, tilt and piston motion. After recombination, the two beams are focused with a focusing lens, creating two mutually coherent virtual point sources.

At a distance $R = \sqrt{x^2 + y^2 + z^2}$, the intensity distribution, I, of a source with wavelength λ is given by:

$$I = A\left(1 + \cos\left(\frac{2\pi x f}{\lambda R}\tan\theta + \frac{2\pi\Delta z}{\lambda}\right)\right)$$

We observe that we can project sinusoidally varying patterns with adjustable spatial frequency and phase. Such a set of patterns is sufficient for high-quality 3D reconstruction, provided that a suitable set of frequencies and phases can be chosen [7]. The relative position of the two virtual point sources, and thereby the intensity profile on the scene, is given by Δz and θ, i.e. the displacement of the movable mirror. We have developed a piezo-actuated MEMS micro-mirror capable of re-positioning in 2 ms, while maintaining a static position within 15 nm or 5 µrad during camera exposures [8]. The maximum tilt angle of the mirror (6 mrad) defines the range of spatial frequencies that can be projected.

While our projector design is simple and compact, the MEMS mirror is a novel component whose robustness must be verified prior to use in space. It has therefore been tested for radiation compliance to 100 kRad, TVAC cycling (-40°C - +70°C) and vibration (20g), with no signs of degradation, making it a highly suitable candidate for space applications [9]. The tip-tilt and piston range of the mirror gives us access to 2π phase modulation, and a frequency span Δf of around 20 fringes visible within our acquired images.

The projected pattern is observed by two cameras acquiring synchronized image pairs, operating with >100 Hz image rate and a resolution of 500 x 500 pixels. We use 15 image pairs for 3D reconstruction, moving the mirror to different θ and Δz between each image pair. With a piezo-based system, we expect slow drift in θ and Δz over time, and we select a 3D reconstruction method which is robust to such drift, as will be discussed in section III. The 15 images are captured in 150 ms, limited only by our current camera driver. We automatically adjust the exposure time for the cameras to avoid saturation, with the same exposure time used for all images within a single capture.

To ensure that the projected pattern is visible also in direct sunlight, the cameras are equipped with narrow-band pass sunlight filters with a pass width of 4 nm centered at the laser wavelength. To utilize such a narrow passband, a number of design choices are made. Firstly, the filter is placed between the lens and the sensor, where the angular distribution of the incoming light is the smallest. Secondly, the lens is designed for chief ray angle normal to the sensor across the sensor. Thirdly, the f-number of the lens is chosen to f/5.6, further reducing the extent of the incoming light cone. This relatively high f-number has the added advantage of a relatively large depth of focus, ensuring a wide operating range for the camera. Finally, the laser itself is temperature (and thus wavelength) stabilized using a built-in TEC element and operates at 810 nm and 800 mW optical power. In total, the 3D camera consumes around 10 W (excluding 3D reconstruction) for the laser (5 W) and the two cameras (2×2.5 W).

III. 3D Reconstruction Algorithm

Our 3D reconstruction algorithm is designed with a basis in stereo algorithms. A full description of the algorithm is given in [1], we provide an abridged description here.

Our basis in stereo algorithms means that we are interested in first establishing a subpixel disparity estimate between camera 1 and 2, and subsequently performing a homographic transform to convert the disparity estimates into a 3D point cloud.

The system works by illuminating the scene with patterns of selected spatial frequencies and phase. The goal is that every object point receives its unique temporal "signature" that makes it possible to determine for a pixel A in camera 1 which is the corresponding pixel B in camera 2. As we have a stereo system, we can use the epipolar constraint [10] to constrain the search in the vertical direction, meaning that the projected codes only need to be unique in the horizontal direction. Unlike traditional stereo, the use of temporally changing patterns that cover the whole field-of-view allows us to capture detailed, full 3D data of surfaces that do not have any structure or surface texture, providing us with one independent distance measurement per camera pixel.

The use of sinusoidal patterns would normally hint towards the use of atan based reconstruction algorithms for matching [11], possibly combined with number theoretic or other approaches [12] for resolving the phase ambiguities. However, as the mirror cannot guarantee an exact phase shift between patterns (confer section II) and the phase will drift over time, we have selected to use a conservative approach relying on zero-normalized cross-correlation (ZNCC) instead, noting that it in other works outperforms many other decoders [13]. In addition, ZNCC can easily be made robust towards differences in signal amplitude between the two cameras, unlike other disparity metrics like sum of absolute differences.

We perform a full search for the best match within the disparity window/working range of the camera. This ensures that we will find the best match irrespective of any sudden discontinuities in the scene structure, at the expense of a larger computational load.

We calculate depth data directly from stereo disparity estimates by using a homographic transform after median filtering the disparity estimates. The homographic transform is established on basis on the Zhang's camera calibration method [14].

A. Pattern / frequency selection for projection

Projecting a repeating pattern, such as a sine wave, will lead to ambiguities in the 3D reconstruction, as the signal will look the same in several positions in the image. We need to project sine waves with a carefully selected set of frequencies to resolve these ambiguities, whilst simultaneously managing laser speckle, slightly unreliable frequency/phase information, and partially specular objects. By selecting the right set of patterns, we can enable the disparity estimation algorithm to reduce the impact of these artifacts.

We use the following strategies to select the frequency and phase of our patterns:

1. We include the highest frequency the interferometer can project into the set, as we previously have found this to significantly improve accuracy [15].
2. We project 3 phases $\{0°, 120°, 240°\}$ to ensure a constant amplitude for the signal. This allows us to robustly compensate for object albedo.
3. Beyond this, we choose the set of frequencies that provide the best ratio of main lobe to side lobe height for the working range of our system for the cross-correlation matrix of the selected codes.

As we project a limited number of patterns, we have found it possible to perform an exhaustive search of relevant frequencies to optimize the system according to its working range, as detailed in our previous publication [1].

B. GPU implementation

The algorithm operates almost exclusively on individual pixels, meaning that it is very well suited for modern parallel architectures. In short, our GPU implementation of the algorithm can be broken down into the following steps:

1. Stereo rectification, removing lens distortion. This can be done simultaneously with the reception of the images from the cameras and upload to the GPU. The step is realized through a look-up table.
2. Image stack normalization, removing ambient light and object albedo. This is performed per pixel right after reception. At the same time, data is reshuffled from a space-oriented to a time-oriented data layout for improving the speed of later memory accesses.
3. Cross-correlation between images to produce a disparity map along the epipolar lines of the images.
4. 3D reconstruction by multiplying the resulting disparity map with the homomorphic matrix provided by the camera calibration.

We chose to implement the algorithm in CUDA, using the Python Numba wrapper, instead of a higher-level GPU interface like Halide [16]. Our consideration was that this provided the best balance between implementation effort, flexibility, and performance for this application.

Our GPU implementation follows closely the original Matlab® implementation used in our previous reporting of results [1]. Small modifications have been made to adapt the algorithm for efficient implementation, like changing 2D noise filters to 1D filters. These changes have had no real impact on 3D performance but reduced 3D reconstruction time to less than 100 ms on all tested GPUs, meaning that reconstruction is real-time. Our algorithm uses roughly 30 GFLOP per fully reconstructed 3D image.

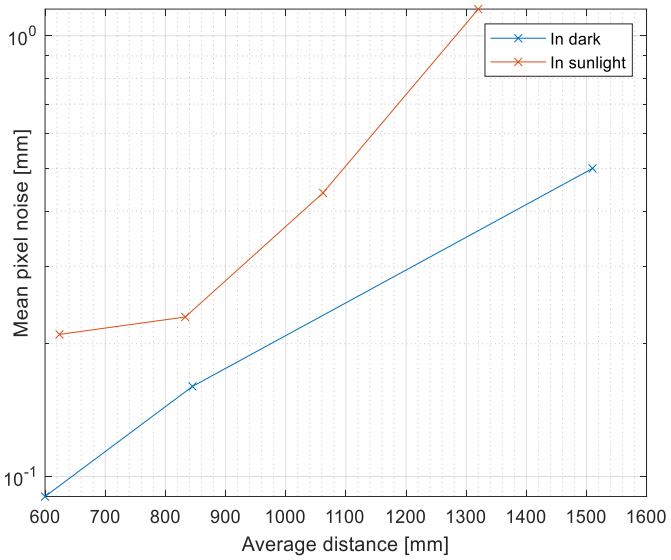

Figure 3: 3D point precision in dark (blue) and in sunlight (red) as function of distance to object.

C. 6DOF Matching Algorithm

In our experiments, we have quantified not only the per-point accuracy of our 3D point cloud, but also the total registration error after registering a known 3D shape to our point cloud. Registration in this context means recovering the position and orientation of the object in the scene coordinate frame.

We used Iterative Closest Points [17] to perform the actual registration. We used point-to-plane as our distance metric during the iterations. The algorithm was initialized manually once per approach trajectory (section V) for the first image only. For subsequent images, the previous position was used to initialize the registration. The point cloud was subsampled prior to registration, and points far away from the initial position were removed prior to registration.

IV. Results

To verify the 3D data quality from the camera, we perform a set of tests, each intended to probe different aspects of the camera performance. The tests presented in this chapter are performed by imaging a planar surface, to isolate the effect of the camera (as opposed to effects introduced e.g. by point cloud matching algorithms). Firstly, we perform imaging of compliant targets at rest and calculate the point precision of the camera. This represents a best-case scenario and is the limiting performance of the camera. Secondly, we introduce sunlight to our scene. As argued above, a 3D camera for space needs to provide data also in direct sunlight. Sunlight will reduce the observed contrast of our projected patterns and will thus degrade our precision and reduce the useful range of the system. Thirdly, we introduce relative motion between camera and scene. Motion during capture will make our patterns shift during acquisition, again degrading the 3D data quality.

A. Point precision at rest

We capture repeated 3D images on planar surfaces both indoor and in sunlight and calculate pr pixel standard deviation as function of distance. The results are shown in Figure 3.

In darkness, we observe as expected a standard deviation proportional to distance squared [15]. At 60 cm distance, the pr pixel standard deviation (precision) is <100 µm. In sunlight,

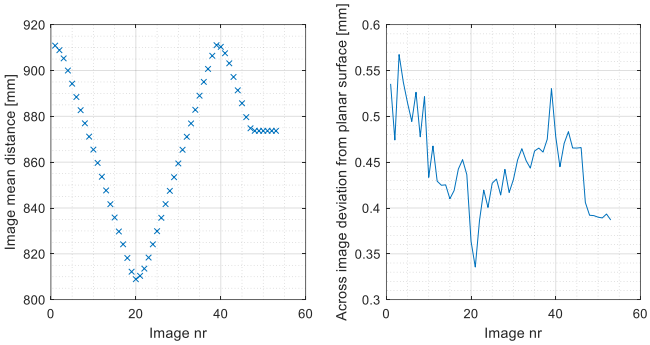

Figure 4: Data from tests during motion at 4 cm/s. X-axis: 3D image index. Left: Average distance to imaged surface. Right: Deviation from planar surface. We observe a slight increase in measurement noise during motion.

the background signal adds to the measurement standard deviation, with increasing influence towards longer distances. Nevertheless, we demonstrate successful 3D reconstruction at up to 1.3 m in direct sunlight. Improved exposure parameters will extend useful range in direct sunlight.

B. Relative motion

Further, we perform the same imaging of a planar surface in an indoor setting, but this time with the camera moving perpendicular to the planar surface. This test probes the data quality as function of relative velocity.

Figure 4 shows an example data set, where the camera is moved at 4 cm/s. The left figure shows the recorded average distance, while the right figure shows the average deviation from planar surface across the image. Assuming that the imaged surface is indeed planar, this deviation will give an indication of data quality deterioration as velocity increases.

We observe that the deviation from planar surface with the camera at rest is approx. 400 µm, while the deviation during motion is approx. 450 µm, a relatively modest deterioration. These tests indicate that the camera can handle relative motion rates up to 4 cm/s. At motion rates of 8 cm/s, the deviation increases above 1500 µm, indicating that 3D reconstruction is severely affected at such high motion rates.

V. Tests towards in-orbit servicing

Having verified the basic functionality of the camera under somewhat idealized conditions, we now seek to test the camera in conditions closer to an in-orbit servicing scenario. In the previous chapter, we assessed the point precision of the camera on compliant surfaces. In this chapter, we present testing of precision as well as trueness. Further, we are well aware of the fact that the surface under consideration will strongly affect the measurement result. Traditionally, black and metallic / specular surfaces are known to be challenging for 3D imaging systems. However, in a satellite scenario, these are typically the surfaces that will be encountered.

For testing in relevant settings, we use the Thales Alenia Space - France Robotic Orbital Facility (ROBY) test bench in Cannes (Figure 6). The ROBY test bench consists of two robots, the “Sensor Robot” (SEN) where the 3D camera is mounted, and the “Mockup robot” (MKP) where the 3D target to be imaged is mounted. In an in-orbit situation, the SEN robot would equal the “servicer” whereas the MKP would equal the “client”. Each of the two robots can move independently in 6 degrees of freedom (DOF). In addition, one of them is mounted on a 10m rail for longer range motion. The

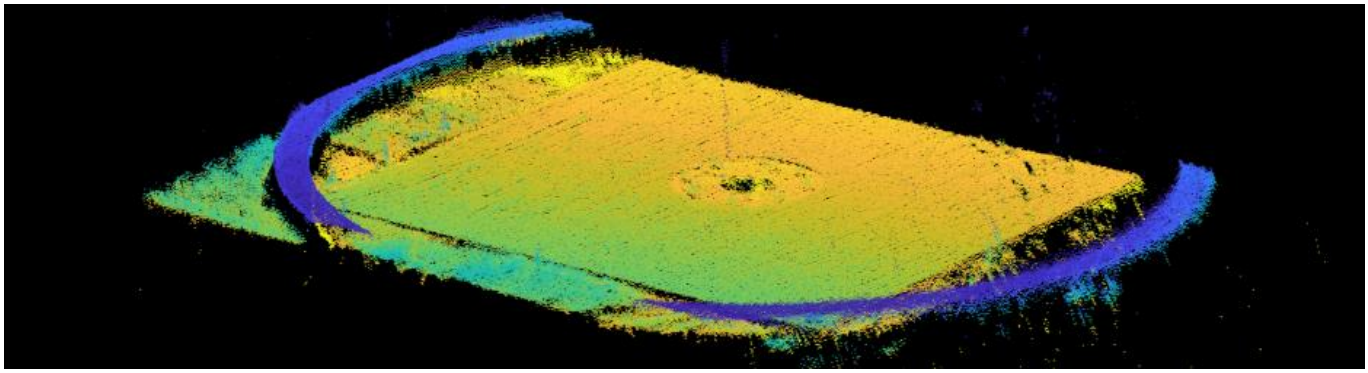

Figure 5: Launch adapter ring. Distance is color-coded to emphasize that the LAR is clearly standing out in the data.

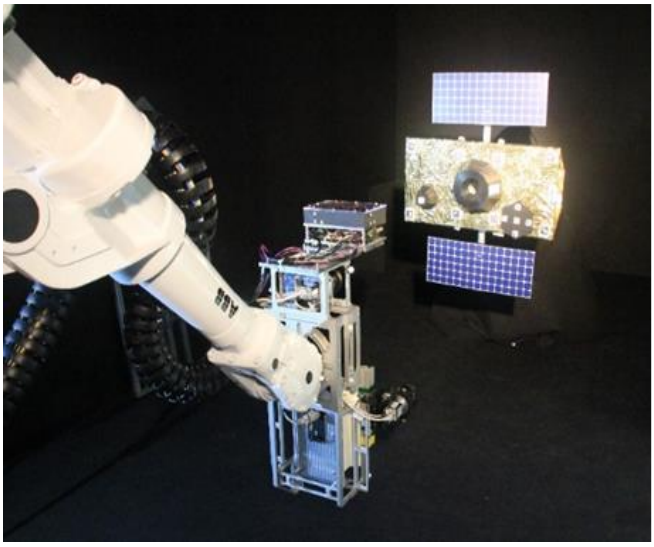

Figure 6: Test of camera at robotic test bench at Thales Alenia Space on a satellite mockup.

ROBY facility also has access to satellite mockups and a set of flight material samples.

For a visual impression of the 3D data quality, we image a satellite mockup. Figure 5 shows a segment of the mockup with the details of the Launch Adapter Ring (LAR) clearly visible. Figure 7 shows further details of the LAR, as well as an alignment pyramid marker on a background of a slightly wavy black MLI.

A. 6DOF precision and accuracy

For quantitative assessment of 3D data quality, we mount a brushed metal cube on the MKP. This represents a simple target, both in terms of imaging and subsequent point cloud matching. We perform a series of approaches with the SEN, at varying relative motion rates. We perform real-time 3D-acquisition and reconstruction of the scene, followed by offline 6DOF matching as outlined in section III.C.

The recorded 6DOF figures are matched with those reported by the robotic positioning system. An example result is shown in Figure 8. The rotation coordinates are kept fixed during this trajectory, and we present XYZ-coordinates from 3D camera and from robotic positioning. At distances between 450 and 1300 mm, we find that the coordinates correspond to within 1 % of the measured distances, even at motion rates up to 7 cm/s. This is indicative of excellent trueness (accuracy), while remaining aligned with the usual mission requirements for the navigation chain.

Much of the discrepancy can be explained by a <100 ms timing offset between 3D camera and robot, and by <25 ms timestamping error / jitter. Both of these effects are most prominent at high velocities. With stationary robots, we calculate the standard deviation in 6DOF estimates, finding that 6DOF precision is better than 50 µm and 1 mrad, even up to distances of 1.3 meters.

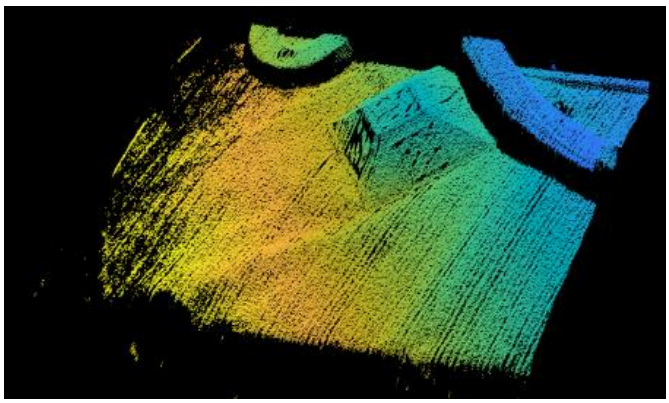

Figure 7: Detail of launch adapter ring with alignment pyramid marker.

The precision and accuracy obtained with our 3D camera are equal to or better than the positioning accuracy of the robotic test bench ground truth. As such, we cannot with confidence state whether the error lies mainly at the camera or at the robotic positioning. Previous tests from the ROBY facility using 2D cameras and markers show typical precision and accuracy of such a traditional image processing chain around 1 cm / 1 degree, i.e. approx. 20 times higher measurement noise. The advantages of an Active Stereo 3D camera are evident for the cases where detailed 3D / 6DOF data is needed.

B. Surface/materials testing

To assess the expected data quality on realistic targets, we acquire 3D data on a set of flight material samples. Quite expectedly, we find that the data quality depends on the surface under investigation, ranging from simple surfaces such as brushed metals, plastics and composites giving good 3D quality in all cases, to mirroring surfaces such as Optical Solar Reflectors which do not provide 3D data using measurement methods with point-like illumination.

Table II shows a simplified compliance table. Such a table is useful to understand the expected data quality, and to select a suitable imaging geometry. In particular, avoiding imaging perpendicular to specular surfaces can significantly increase data quality and coverage of the scene.

TABLE II. FLIGHT MATERIAL COMPLIANCE.

Materials	Compliance
Brushed metal surfaces **Plastic and 3D printed surfaces** **Composite materials**	Tolerant at all angles
MLI surfaces **Polished metal surfaces** **Glossy paint**	Tolerant at surface normals >10 degrees away from camera z-axis.
Highly specular surfaces (solar panels etc).	Limited compliance. Best effort at surface normals >10 degrees away from camera z-axis.
Mirroring surfaces (Optical Solar Reflectors etc.)	Non-compliance.

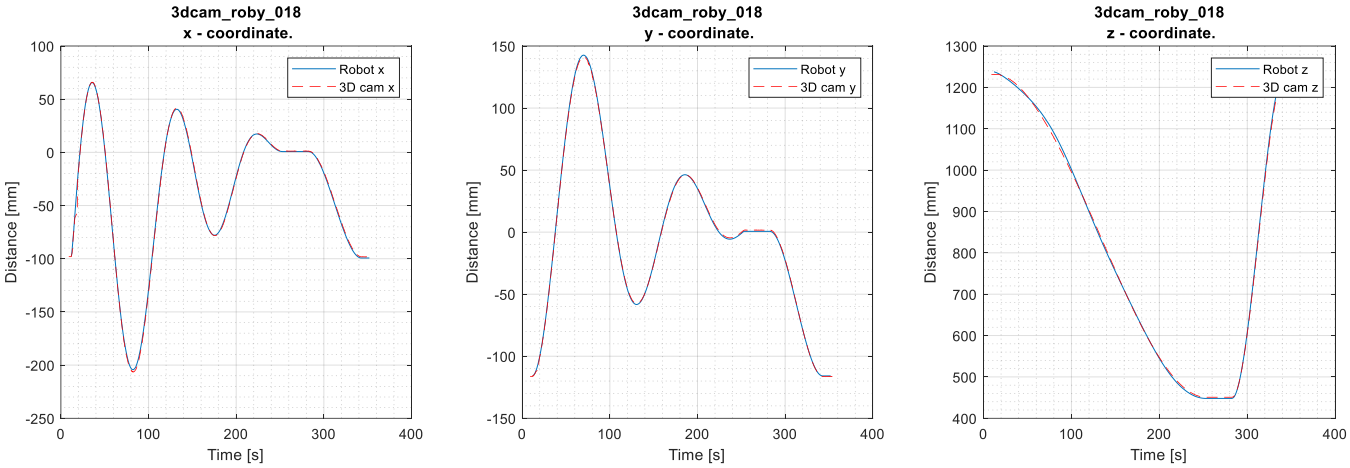

Figure 8: Tests towards motion, XYZ coordinates. Robot position (ground truth) and recovered camera position plotted on top of each other.

VI. SUMMARY AND OUTLOOK

We have presented a compact, robust and low-power 3D camera that in tests shows good 3D data quality, motion tolerance, sunlight resistance and compatibility with most relevant space materials.

Table III shows the specification of our current prototype, and we see clear possibilities for optimizing the design to achieve a further significant reduction in size and mass. Figure 9 shows a miniaturized version of the 3D camera, where the projector has been reduced in size to 20 x 30 x 50 mm, limited only by the need for manual assembly. This is small enough to fit between the two cameras, significantly reducing the size of the 3D camera.

The MEMS micro-mirror has been radiation, TVAC and vibration tested with no degradation in performance. The camera's capability of reliably imaging both cooperative and non-cooperative clients broadens its usage and provides robustness for interaction. The 3D reconstruction algorithm is well-suited for implementation on parallel architectures like GPU's and FPGA's.

For implementation into an actual space camera, it will be natural to both embed the 6DOF estimation into the camera itself, plus moving the 3D reconstruction algorithm from GPU to a space-compliant FPGA. Such space-compliant FPGAs with significant computational capacity have recently been launched (e.g. Xilinx™ Versal), which could enable a power-efficient, high-performance system.

For future space applications like in-orbit servicing and planetary exploration, the 3D camera can fit very well in a 3D sensing portfolio, providing high quality 3D data for shorter working ranges. This can enable robust, efficient, and safe inspection and manipulation for space robotics.

TABLE III. CAMERA SPECIFICATION FOR CURRENT PROTOTYPE AND OPTIMIZED VERSION

	Current prototype	Optimized version
Resolution	500x500 pix	500x500 pix
Field-of-view	15x15 deg	15x15 deg
Capture time	150-450 ms	75-200 ms
Compute time	< 100 ms	< 100 ms
Size, mass	175 x 125 x 88 mm	145 x 80 x 50 mm
Mass	2 kg	0.4 kg
3D point precision	0.1 mm at 0.6 m	0.1 mm at 0.6 m
3D trueness	<1 % of distance	<1 % of distance
6DOF precision	<50 µm, <1 mrad	<50 µm, <1 mrad
Power incl. stereo reconstruction	30 W @ 10 Hz	10 W @ 10 Hz

ACKNNOWLEDGEMENTS

The authors thank Fanny Keller and Volker Kirschner, European Space Agency, for a very constructive dialogue throughout the project.

The camera has been developed under a contract with the European Space Agency, contract 4000129212-ECR001.

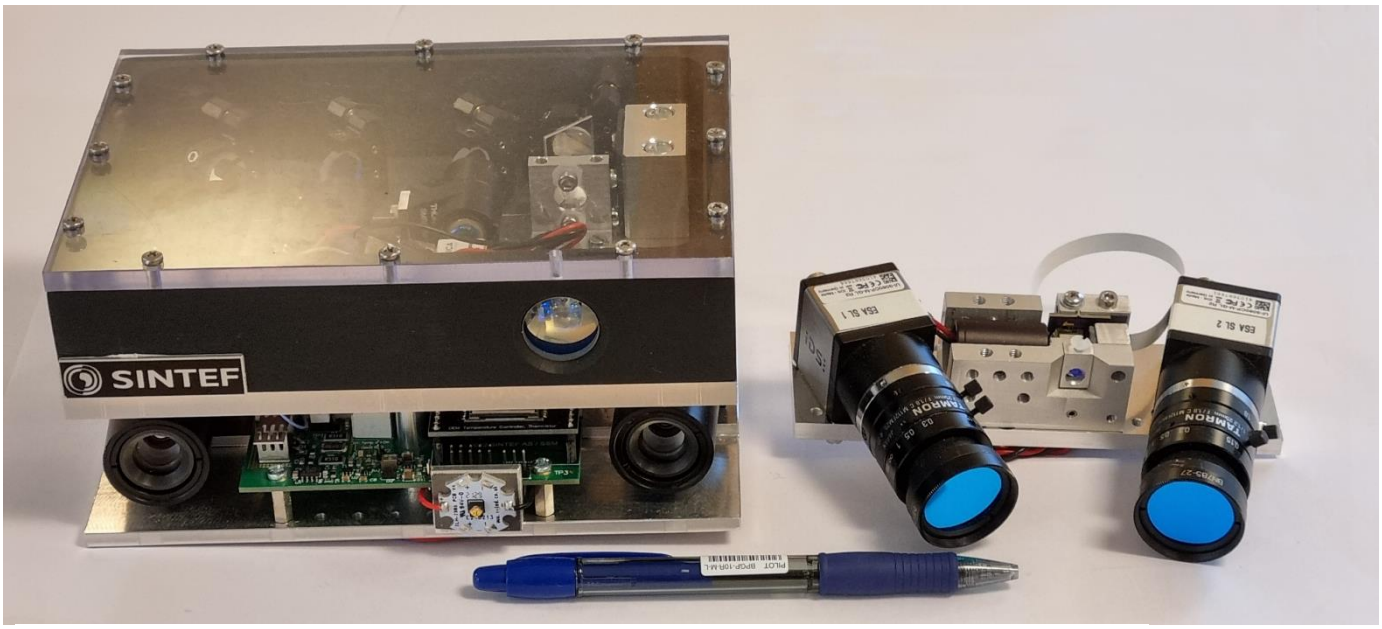

Figure 9: Current prototype (left) and miniaturized projector (right). Pen included as size reference.

REFERENCES

[1] J. B. Thorstensen, J. T. Thielemann, P. Risholm, J. Gjessing, R. P. Dahl-Hansen, and J. Tschudi, "High-quality dense 3D point clouds with active stereo and a miniaturizable interferometric pattern projector," *Opt. Express*, vol. 29, no. 25, pp. 41081–41097, 2021.

[2] K. H. Haugholt, A. H. Hansen, P. Risholm, J. T. Thielemann, and G. Tzeremes, "Development of A Flash-Lidar Elegant Breadboard Model for Rendezvous Applications," in *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*, Waikoloa, HI, USA: IEEE, Sep. 2020, pp. 3483–3486. doi: 10.1109/IGARSS39084.2020.9324046.

[3] A. Pollini, C. Pache, and J. Haesler, "CSEM Space Lidars for Imaging and Rangefinding," in *IGARSS 2018 - 2018 IEEE International Geoscience and Remote Sensing Symposium*, Jul. 2018, pp. 1849–1852. doi: 10.1109/IGARSS.2018.8519241.

[4] G. S. Aglietti *et al.*, "RemoveDEBRIS: An in-orbit demonstration of technologies for the removal of space debris," *Aeronaut. J.*, vol. 124, no. 1271, pp. 1–23, Jan. 2020, doi: 10.1017/aer.2019.136.

[5] F. M. Kolb *et al.*, "THE LIRIS-2 3D IMAGING LIDAR ON ATV-5," presented at the ASTRA, 2015.

[6] J. F. Bell *et al.*, "The Mars 2020 Perseverance Rover Mast Camera Zoom (Mastcam-Z) Multispectral, Stereoscopic Imaging Investigation," *Space Sci. Rev.*, vol. 217, no. 1, p. 24, Feb. 2021, doi: 10.1007/s11214-020-00755-x.

[7] C. Zuo, S. Feng, L. Huang, T. Tao, W. Yin, and Q. Chen, "Phase shifting algorithms for fringe projection profilometry: A review," *Opt. Lasers Eng.*, vol. 109, pp. 23–59, Oct. 2018, doi: 10.1016/j.optlaseng.2018.04.019.

[8] T. Bakke, A. Vogl, O. Żero, F. Tyholdt, I.-R. Johansen, and D. Wang, "A novel ultra-planar, long-stroke and low-voltage piezoelectric micromirror," *J. Micromechanics Microengineering*, vol. 20, no. 6, p. 064010, Jun. 2010, doi: 10.1088/0960-1317/20/6/064010.

[9] R. Dahl-Hansen, J. Gjessing, P. Mardilovich, C. Fragkiadakis, and J. Thorstensen, "Reliable Pb(Zr,Ti)O3-based thin film piezoelectric micromirrors for space-applications," *Appl. Phys. Lett.*, vol. 121, no. 13, p. 132901, Sep. 2022, doi: 10.1063/5.0106933.

[10] Y. Liu, N. Pears, P. L. Rosin, and P. Huber, *3D Imaging, Analysis and Applications*. 2020. doi: 10.1007/978-3-030-44070-1.

[11] J. Thorstensen and J. Thielemann, "Adaptive multi-frequency phase stepping for optimal 3D depth reconstruction," in *Imaging and Applied Optics 2018 (3D, AO, AIO, COSI, DH, IS, LACSEA, LS&C, MATH, pcAOP), OSA Technical Digest*, 2018, p. JTu4A.15. doi: 10.1364/3D.2018.JTu4A.15.

[12] C. Zuo, L. Huang, M. Zhang, Q. Chen, and A. Asundi, "Temporal phase unwrapping algorithms for fringe projection profilometry: A comparative review," *Opt. Lasers Eng.*, vol. 85, pp. 84–103, Oct. 2016, doi: 10.1016/J.OPTLASENG.2016.04.022.

[13] P. Mirdehghan, W. Chen, and K. N. Kutulakos, "Optimal Structured Light a la Carte," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2018. doi: 10.1109/CVPR.2018.00654.

[14] Z. Zhang, "A flexible new technique for camera calibration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 22, no. 11, 2000, doi: 10.1109/34.888718.

[15] G. Bouquet, J. Thorstensen, K. A. H. Bakke, and P. Risholm, "Design tool for TOF and SL based 3D cameras," *Opt. Express*, vol. 25, no. 22, pp. 27758–27769, Oct. 2017, doi: 10.1364/OE.25.027758.

[16] J. Ragan-Kelley *et al.*, "Halide: decoupling algorithms from schedules for high-performance image processing," *Commun. ACM*, vol. 61, no. 1, pp. 106–115, Dec. 2017, doi: 10.1145/3150211.

[17] S. Rusinkiewicz and M. Levoy, "Efficient variants of the ICP algorithm," in *Proceedings Third International Conference on 3-D Digital Imaging and Modeling*, Quebec City, Que., Canada: IEEE Comput. Soc, 2001, pp. 145–152. doi: 10.1109/IM.2001.924423.