3D Surface Reconstruction of a BGA Connector with Large Specular Highlights by Photometric Stereo

Maxime Vallée* Hideo Saito** Junichi Sugano† Yoshinobu Takizawa‡

We propose to apply the Photometric Stereo method on small objects with a specular surface behavior, such as BGA (Ball Grid Array) connectors. This method is based on a single, fixed camera, and different light sources. From the different images that are obtained in the different light conditions, we first derive the normal vectors map and integrate it to get the 3D shape of the surface. The key issue here is to use as little lights as possible and a camera the less precise as possible while getting a result as precise as possible considering that the surface reflection properties, especially its specular highlights, can make it quite more difficult.

Keywords: photometric stereo, specular, ball grid array connector, 3D shape reconstruction

1. Introduction

The main purpose of computer vision is to describe the real world from the visual information we get from devices (such as cameras) and, as a consequence, one important goal is to reconstruct three-dimensional (3D) shapes from two-dimensional (2D) images. For that purpose, several approaches have been proposed, all having their own assets and drawbacks. It has been proposed to mimic the human binocular vision by using two cameras (stereoscopic vision), or to move the camera in order to get different views of the surface we want to reconstruct (structure from motion)[(1)]. More complex cameras can be used as well to directly get the depth on the image, such as the Microsoft Kinect.

The photometric stereo method relies on the idea that, with a given camera position and light conditions, the measured light intensity only depends on the surface orientation, and some parameters proper to the surface. The local surface orientation having two degrees of freedom, we can obtain it by using several light conditions. This method was first introduced by Woodham[(2, 3)], and then implemented by Silver[(4)].This method has been used for several applications since then[(5, 6, 7, 8)].

Still, the relation between the measured light intensity and the surface orientation is unclear according to the surface properties. Most of the time, it is chosen to focus on particular surfaces, for which this relation is known and easy to formulate. Lambertian surfaces for instance have a diffuse reflection, i.e. the light is reflected everywhere with the same intensity, only depending on the angle between the surface and the incoming light ray. Those are the easiest ones to deal with since the information is worth using everywhere on the image, and the mathematical relationships are simple. Specular surfaces, like the ones of most metals, reflect instead almost all the light in one direction, according to the geometrical optics laws. Given a point on the surface, if the light rays that arrive on it are reflected towards the camera then it can get a very good information about the orientation of the surface, but if they are not (it is the case for the vast majority of the pixels if we use one punctual light source) there is no information available for that pixel. Most materials actually combine those two behaviors. In this paper we will apply the photometric stereo method on a BGA connector, which presents a very specular surface.

Photometric stereo has already been applied to specular surfaces[(9, 10)], but those methods which consist of an estimation of the BRDF (Bidirectional Reflectance Distribution Function) would not work on very wide specular highlights like the ones we have to face.

Section 2 presents the background of our work. Section 3 explains how we can reconstruct the 3D shape of the surface, even with specular reflection, thanks to photometric stereo. Section 4 shows and analyses the experimental results.

Fig. 1. Bottom view of an embedded processor, showing the BGA connections

* Hyper Vision Research Lab,
Graduate School of Science for Open Systems and Environment
Keio University
Tokyo, Japan
E-mail: maxime.vallee@ecl2013.ec-lyon.fr

** E-mail: saito@hvrl.ics.keio.ac.jp

† ViSCO Technlogies
Tokyo, Japan
E-mail: jsugano@visco-tech.com

‡ ViSCO Technologies
Tokyo, Japan
E-mail: ytakizawa@visco-tech.com

2. Background

Precise 3D surface reconstruction has become more important in the industry, for quality control purposes. BGA connectors are being more and more used, especially for connecting processors as shown on figure 1, and as the processors keep getting smaller and smaller the connectors need to stay really precise and efficient. So in this case we want to check that the connecting spheres are at the right position, are all on the same plane (so that they can all touch the connected device) and have the right shape.

Before, this quality control was performed manually. Somebody would just check the measurements, but this is very slow, and the proportion of connectors we could check out of those going out of the factory was too low. Another method would be to actually connect the device and check if it works well but in our case, since a processor is a quite complicated device, it is also quite slow. So the 3D reconstruction seems to be a better solution.

Armand et al.[11] proposed a method inspired by photometric stereo to measure bent pins on electric components (with high specularity) by detecting the specular highlights. We want this time to check on BGA connectors by getting all the 3D shape, and not only the specular information, so we focus on getting the lambertian information.

As we saw in introdution, there are a lot of means to obtain the 3D shape of a surface. First, we could think of Stereoscopic Vision, or Structure from Motion. Those two methods imply getting different views of the surface, and getting the metric information by matching some interest points. The drawbacks of those methods are that they need very identifiable interest points, while the BGR connectors do not have much texture to rely on, which lead to a lot of errors in finding those points. Stereoscopic Vision also has the problem of using two cameras, while we need precise cameras and objectives, i.e. expensive ones, in order to get an image of the connector. In the other hand, Structure from Motion requires estimating the camera position and is thus less precise, while we want to be as precise as possible. Secondly, we could also imagine using depth sensors, but if they are precise enough (unlike the Microsoft Kinect), they get too expensive. Eventually, Structured Light[12] gets similar results to Photometric Stereo, but the projection system is more complex, while any light source can be used for the latter.

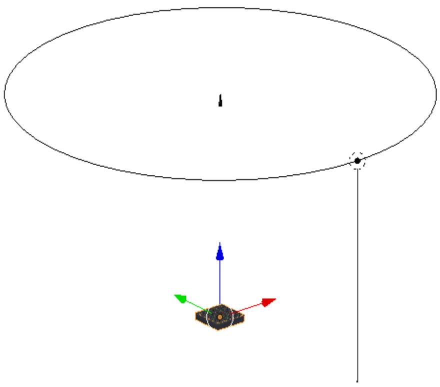

Fig. 2. Relative positions of the camera, connector and lights

3. Normal Vectors Map Computing

All the following algorithms are used on images obtained from simulations. This lets use change easily the reflection properties of the surface to evaluate the algorithm's efficiency.

We decide to use 8 light sources arranged regularly in a circle around the camera, as described figure 2. We get the 8 images like the ones shown on figure 3, on which we will work on.

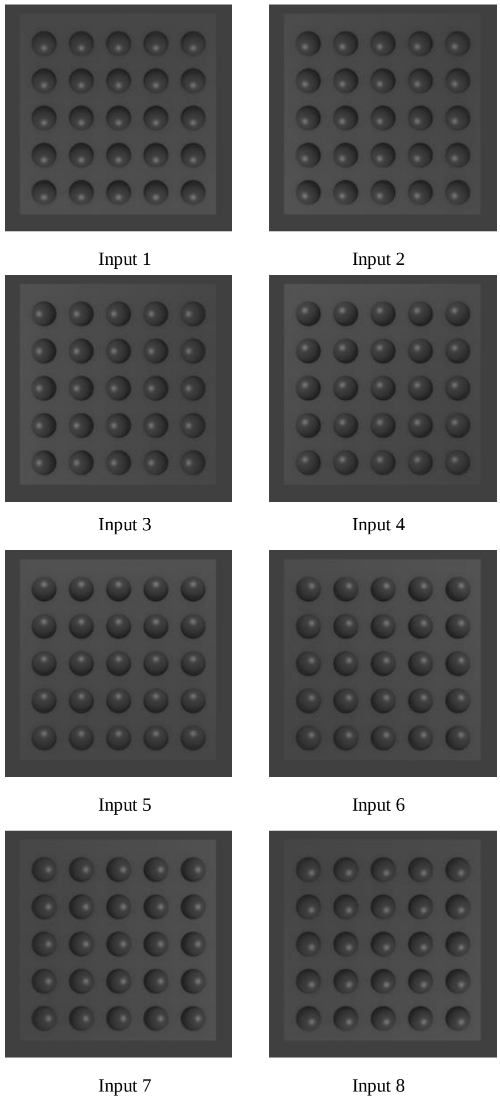

Fig. 3. Input Images

3.1 Lambertian reflection

We first assume the surface is Lambertian (or diffuse), i.e. the light is reflected on the surface equally in every direction, only depending on the angle between the surface normal (vector $\vec{n}$) and the incoming light ray.

$$\mathrm{I} = k_d L.\vec{n}$$

Where I is the vector of the 8 intensities of a given pixel (for the 8 light sources), and L is the 8x3-matrix giving the incoming directions of each light ray. The factor k_d is proper to the surface proprieties. In that case, the normal vector at each point of the surface can be obtained from the intensity at each point for each light condition knowing that we can compute the matrix $\mathrm{A} = (L^T L)^{-1} L^T$ very easily ($L^T L$ is a 3x3 matrix), and get $\vec{n}$

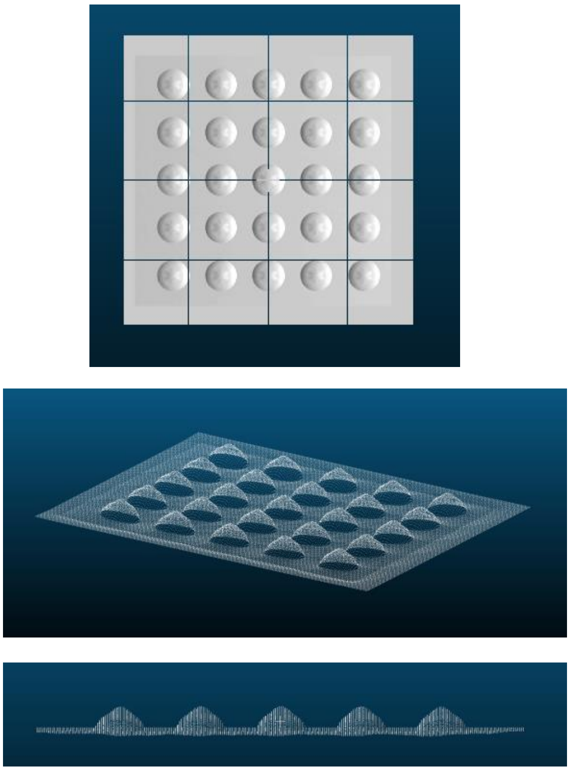

Fig. 4. Naively reconstructed 3D mesh of the surface of a BGA connector (for $I_S = 1$, $I_D = 0.5$, $h = 40$, see the explanation of the reflection parameters section 4)

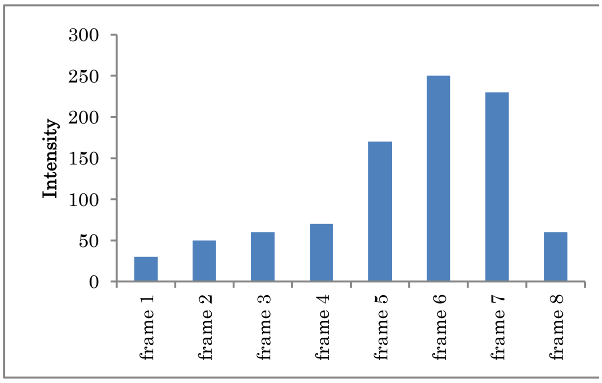

Intensities for one pixel according to the light position (8 frames)

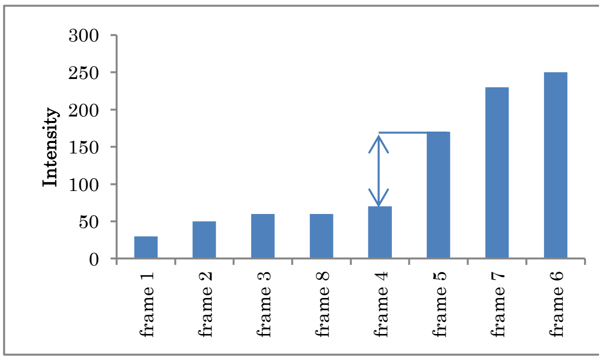

Sorted intensities for the same pixel. The highlighted interval is the first to be higher than the threshold, so the frames 5, 7 and 6 will be ignored

Fig.5. Detection of the specular values for a given pixel

simply by normalizing $k_d\vec{n}$.

$$L^T I = k_d L^T L.\vec{n}$$
$$k_d\vec{n} = (L^T L)^{-1} L^T I$$

3.2 Specular reflection In the case of a BGR connector, and more generally in the case of any metallic surface (and a lot of other surfaces), the light reflection is not only Lambertian, but also specular. That means that the light will be preferably reflected in one direction, according to the laws of geometrical optics. Visually, that creates the clear dots on the spheres on figure 3, and the previous algorithm cannot be naively used. The figure 4 shows the result of the full reconstruction if we naively consider the surface as Lambertian: the spheres become almost cones because when a pixel is caught once in a specular highlight, since the other values are weak in comparison, the slope is only estimated from that value. The algorithm estimates that the slope is orthogonal to the light direction at that time, while the slope is actually closer to the one needed for a mirror reflection from the light source to the camera.

To counter this effect, we treat each pixel before. The figure 5 shows how we eliminate the specular values. We first sort the 8 different intensity values. Then, starting from the darkest one, we check the difference with the next value. When this difference is over a threshold we set before (we set it at 50), and if we have at least 3 intensity values (the minimum needed for the algorithm), we consider that the following values correspond to a specular behavior, and thus ignore them. We use the other values in the algorithm, as if the surface had a simple Lambertian behavior.

3.3 3D shape integration Once we have obtained the normal vector (n_x,n_y,n_z) everywhere on the image, we want to integrate it so that we can get the 3D shape of the surface, i.e. the altitude $z_{i,j}$ of each pixel. For each pixel (i,j) on the image (with the exception of some boundaries):

$$n_x + n_z\left(z_{i+1,j} - z_{i,j}\right) = 0$$

$$n_y + n_z\left(z_{i,j+1} - z_{i,j}\right) = 0$$

We can then build the M matrix, sparse, which verifies (z being the vector built from all the values of $z_{i,j}$):

$$Mz = v$$

For instance, the lower halves of M and v (corresponding to the equations on y) would be as follow (each column and each line corresponds to one pixel (i,j), in lexicographic order. I and J are respectively the number of lines and columns in the image)

$$M_y = \begin{array}{c} \begin{array}{cccccccc} (0,0) & & & (1,0) & & & & \end{array} \\ \left(\begin{array}{cccccccc} -1 & \square & \square & 1 & \square & \square & \square & \square \\ \square & \ddots & \square & \square & \ddots & \square & \square & \square \\ \square & \square & \ddots & \square & \square & \ddots & \square & \square \\ \square & \square & \square & \ddots & \square & \square & \ddots & \square \\ \square & \square & \square & \square & -1 & \square & \square & 1 \\ \hline \square & \square & 1 & \square & \square & -1 & \square & \square \\ \square & \square & \square & \ddots & \square & \square & \ddots & \square \\ \square & \square & \square & \square & 1 & \square & \square & -1 \end{array}\right) \\ \begin{array}{cccccccc} & & (I-1,1) & & (I-1,J) & (I,1) & & (I,J) \end{array} \end{array} \begin{array}{l} (0,0) \\ \\ \\ \\ (I-1,J) \\ (I,1) \\ \\ (I,J) \end{array}$$

$$
v_y = \begin{pmatrix} -\dfrac{n_y(0,0)}{n_z(0,0)} \\ \vdots \\ -\dfrac{n_y(I-1,J)}{n_y(I-1,J)} \\ \dfrac{n_y(I,1)}{n_y(I,1)} \\ \vdots \\ \dfrac{n_y(I,J)}{n_y(I,J)} \end{pmatrix}
$$

This equations system is over determined, so we have to choose a method in order to approximate the altitudes the best way we can. Different algorithms exist, but we decide to use the least squares method of conjugate gradient because of its simplicity and its good results. Since it is better used with a symmetrical matrix, we use it on this equivalent system:

$$M^T M z = M^T v$$

Once this system is solved we obtain the 3D mesh representing the reconstructed surface. The figure 6 shows the output of the algorithm whit the 8 output images shown figure 3.

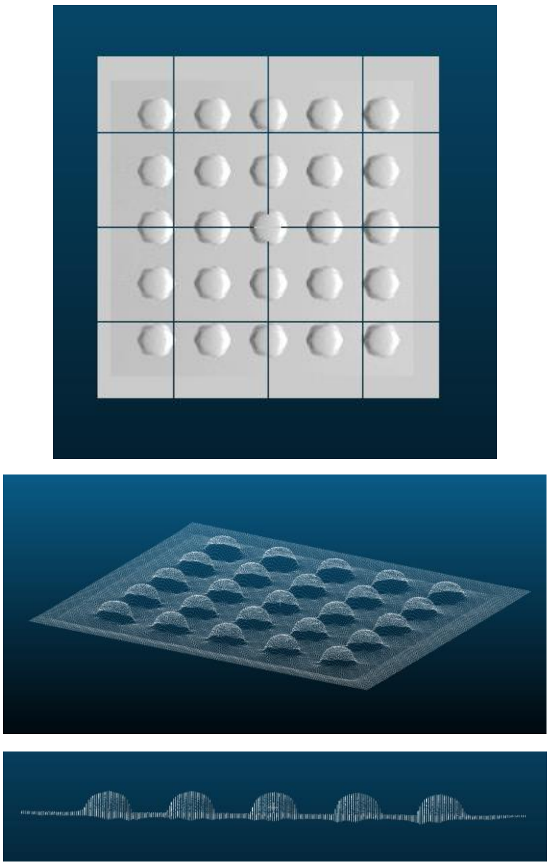

Fig. 6. Reconstructed 3D mesh of the surface of a BGA connector (for $I_S = 1$, $I_D = 0.5$, $h = 40$, see the explanation of the reflection parameters section 4)

4. *Experimental results*

4.1 Results from the simulation For evaluating the accuracy of the method according to the importance of the specular reflection on the surface, we work on a simulated surface. We use the Cook-Torrance model[13] for the specular reflection. This model takes into account many phenomena: it defines the reflectance of any part of the surface, according to the incoming ray direction L, the camera direction V and the normal to the surface N.

$$R_S = \frac{F}{\pi} \frac{DG}{(\vec{N}.\vec{L})(\vec{N}.\vec{V})}$$

Where:

- F is the Fresnel term, who takes into account the light wavelength, because a material does not reflect any wavelength the same.
- G is the geometrical attenuation factor, which describes the self-shadowing due to the micro-facets (the surface is in fact composed of a lot of micro-facets which can face any direction, and not only the mean direction N)
- D represents the distribution of the micro-facets which could have the good normal vector (i.e. equal to H, the bisector of L and V), so that the light is reflected like with a mirror to the camera.

The main term here is D. Several distributions can be used, one of the simplest ones being a Gaussian distribution:

$$\mathrm{D} = e^{-\left(\frac{\alpha}{m}\right)^2}$$

Where α is the angle between N and H (i.e., for instance, the center of the specular highlight is going to be when the conditions for a mirror reflection are fulfilled). The factor m (roughness) is the root mean square of the slopes of the micro-facets. For instance, a very smooth surface will have a roughness of 0. According to this distribution the radius of the specular highlight is going to be proportional to the roughness. The Beckmann distribution is actually used, which is a slightly improved version of the Gaussian distribution, but the factors are the same.

$$\mathrm{D} = \frac{1}{m^2 cos^4 \alpha} e^{-\left(\frac{\tan\alpha}{m}\right)^2}$$

For our simulations, the intensity of the specular reflection (I_S), which is going to take into account the Fresnel and geometrical factors, will vary between 0 (no reflection) and 1 (all the light is reflected) and represent the brightness of the specular highlight. The hardness (h), which is in fact the inverse of the roughness (a hardness of 0 would theoretically mean that the specular highlight is infinite and the higher it gets the smaller the highlight gets), will be fixed to 40 most of the time (to represent better a metal). The diffuse reflection will also play a role, and will be described through its intensity I_D, which corresponds to the brightness of the reflection (between 0 for no reflection and 1 for total reflection).

With that method we can change very precisely the specular and diffuse reflections intensities and the hardness of the surface for the same shape. The figures 7, 8 and 9 show the influence of those three factors on the input images.

This allows us to evaluate our method according to the surface properties. The figures 10, 11 and 12 show the result of our method according to each factor. Since the reconstruction through photometric stereo gives only relative altitudes, we bring all the reconstructed shapes to the same level (by subtracting their minimum altitude) for comparing them.

First, we can notice on those three figures that in any case, the slope estimation is not accurate when it is almost vertical. This is due to the algorithm in itself (the slope becomes infinite so the integration is not possible), but also to the shadowing in those parts.

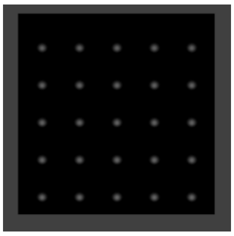

$I_D = 0$

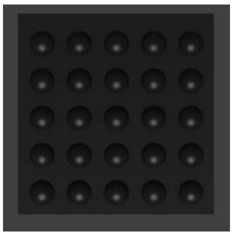

$I_D = 0.1$

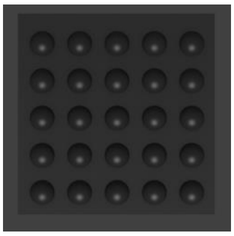

$I_D = 0.2$

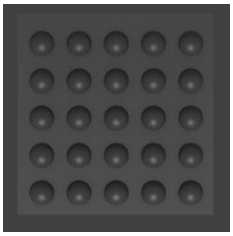

$I_D = 0.5$

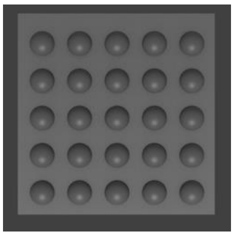

$I_D = 1$

Fig. 7. First input images according to the intensity I_d of the diffuse reflection ($I_s = 1$, h = 40)

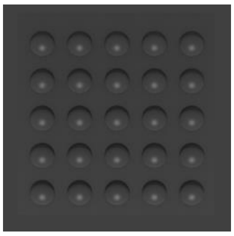

Is = 1

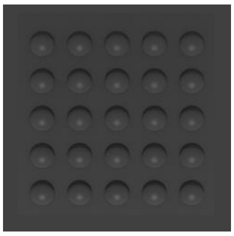

Is = 0.5

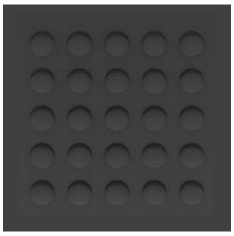

Is = 0

Fig. 8. First input images according to the intensity I_s of the specular reflection (Id=0.5, h=40)

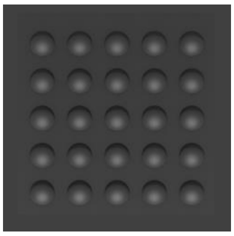

h = 10

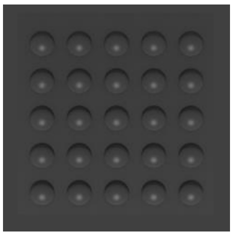

h = 40

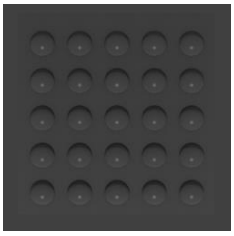

h = 160

Fig. 9. First input images according to the hardness h of the specular reflection ($I_D = 0.5$, $I_S = 1$)

The diffuse reflection is all the information the algorithm gets once the specular highlights are neglected, so it is important to check what amount of diffuse reflection we need to perform the reconstruction. On the figure 10 we can see that the reconstruction stays accurate even with a low diffuse reflection. With diffuse reflection intensity higher that 0.5, the reconstruction is almost perfect while it is still acceptable for 0.05 (we start to notice a small peak at the center of the sphere, because its surroundings are very often in a specular highlight). When the diffuse reflection intensity gets lower, the algorithm does not get enough information anymore to compute a good shape.

The figure 11 shows that the brightness of the specular dots does not influence the result so much: if it is high the specular highlights are going to be filtered by the algorithm, and if it is low the highlights are weak and can be treated as if they were the result of diffuse reflection.

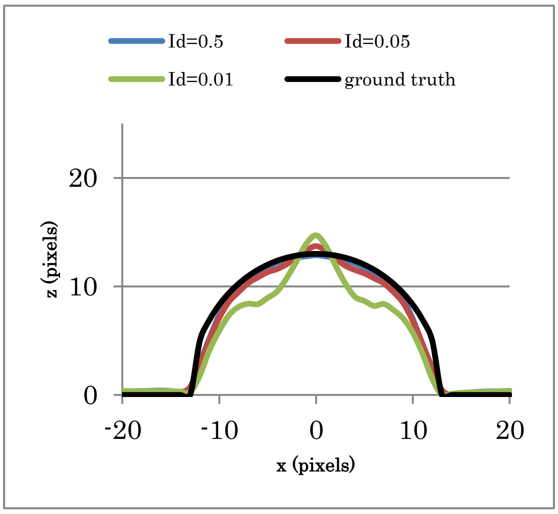

Fig. 10. Profile of a reconstructed BGA sphere according to the intensity Id of the diffuse reflection ($I_S = 1$, h = 40)

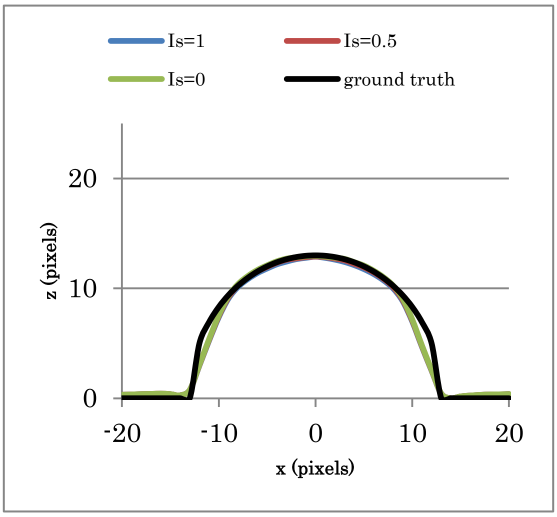

Fig. 11. Profile of a reconstructed BGA sphere according to the intensity Is of the specular reflection (I_D = 0.5, h = 40)

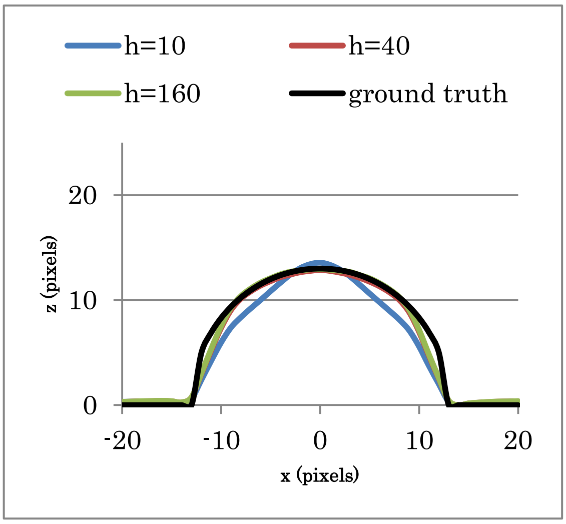

Fig. 12. Profile of a reconstructed BGA sphere according to the hardness h of the surface (I_S = 1, I_D = 0.5)

The size of the specular highlights, represented by the hardness h, has instead a higher impact. As shown on figure 12, if the hardness is too small (i.e. the highlights are too wide), the slope is less accurate. Indeed, if the highlights are wide then a lot of frames will be ignored when we compute the slope of a particular pixel. As a consequence, there will be less possible slope values for the reconstruction of this pixel's area (and the reconstructed 3D shape will be in fact composed of bigger polygons).

4.1 Results from the experiment In order to check the viability of our algorithm for real applications, we set up an experiment, which follows the simulation setup, as shown figure 13. We can see on that picture the 8 lights arranged in a circle around the camera, which is equipped with an objective so that it can focus on the small BGA connector. The synchronization between the lights and the camera is performed by the units on the right, so that it takes less than one second to take all the needed pictures with the different light conditions.

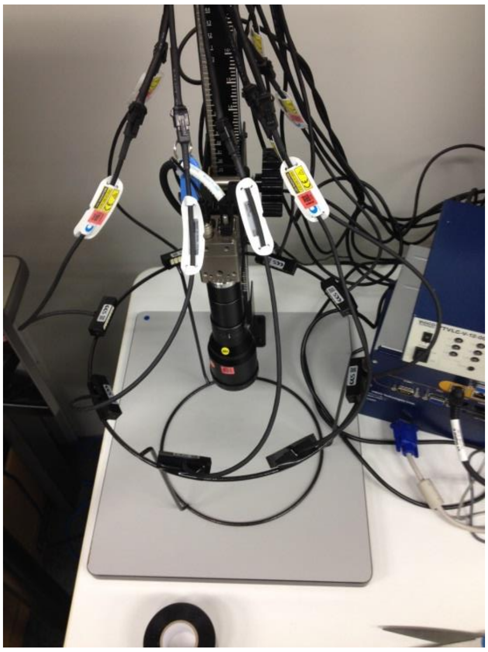

Fig. 13. Experimental setup

We can see on the figure 14 that the quality of the images is, as expected, less optimal than the ones provided by the simulations. Moreover, we seem to be here in a case where the diffuse reflection is very weak, which is where our algorithm started to fail in the simulations, as shown figure 10.

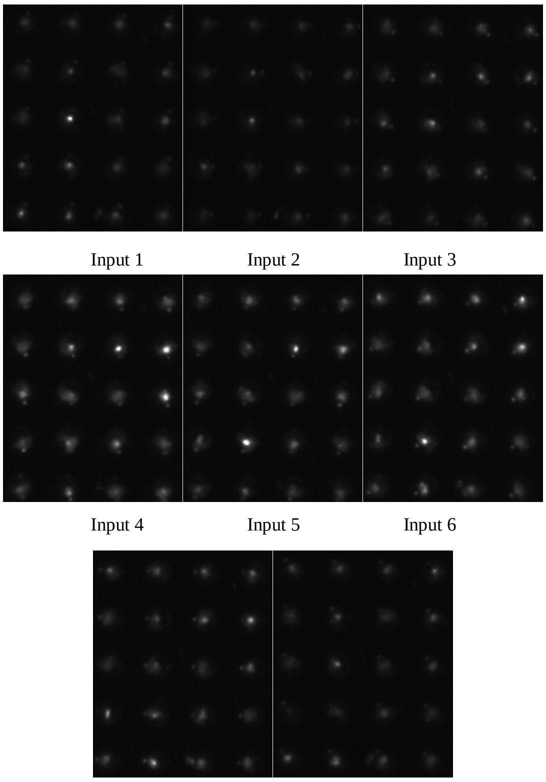

Fig. 14. Input images from the camera

Still, we get some promising results from the algorithm. The figure 15 highlights the 3D mesh obtained from the 8 input images shown figure 14.

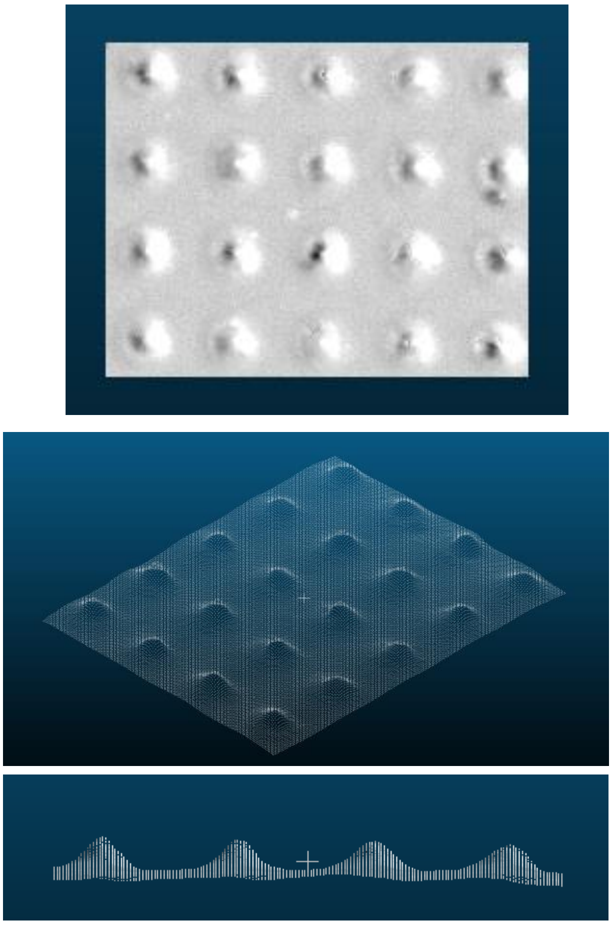

Fig.15. 3D mesh obtained from the real images

The global shape of the mesh is quite good, and the positioning and heights of each ball is well reconstructed. But if we want to improve the individual shape of the spheres (which are not very spherical in our reconstructed mesh), we will need more information around the specular highlights (since the specular reflection is very weak). We should also notice that the camera response function is considered linear in our algorithm, and has yet to be determined to improve the results through a pretreatment of the input images.

5. Conclusion

Our paper shows that the Photometric Stereo method gives good results in the case of a surface with no texture, and is still accurate on a specular surface. This method is better than the others in the sense that it uses less specific — so less expensive — devices, and does not use any point-matching method which would be impaired by the lack of texture on the surface.

It has some drawback though, especially near the parts where the slope is vertical or almost vertical since the integration gets less precise in that case. Also, even though the occlusion problem was not faced in this paper, we know that an occlusion would not only impact on the occluded parts, but on all the reconstructed surface since we need all the information between two pixels to correctly locate them vertically.

For our next research, we will focus on the surfaces showing no diffuse reflection (like on the first image of fig. 6). In that case, we cannot use the Lambertian model and thus need to use only the information obtained through to the specular reflection, so that we would need a totally different algorithm. Eventually, the new algorithm could be used to improve the experimental results of this paper when diffuse and specular reflections are mixed, since until now we do not use the information we get from the specular reflection.

References

(1) Dyer: "Volumetric scene reconstruction from multiple views", Fundations of Image Understanding, pp.469-489 (2001)

(2) Woodham: "Reflectance map techniques for analyzing surface defects in metal castings", Tech. Rep AI-TR-457 (1978)

(3) Woodham: "Photometric method for determining surface orientation from multiple images", Opt. Eng., Vol.19, pp.139-144 (1980)

(4) Silver: "Determining shape and reflectance using multiple images", M. S. thesis (1980)

(5) Ikeuchi: "Determining Surface Orientations of Specular Surfaces by Using the Photometric Stereo Method", Pattern Analysis and Machine Intelligence, Vol.PAMI-3, No.6, pp.661-669 (1981)

(6) Hansson and Johansson: "Topography and reflectance analysis of paper surfaces using a photometric stereo method", Opt. Eng., Vol.39, No.9, pp.2555-2561 (2000)

(7) Ackermann, Langguth, Furhrmann, and Goesele: "Photometric stereo for outdoor webcams", Computer Vision and Pattern Recognition (CVPR), pp.262-269 (2012)

(8) Vogiatis and Hernández: "Self-calibrated, multi-spectral photometric stereo for 3D face capture", International Journal of Computer Vision, Vol.97, No.1, pp.91-103 (2012)

(9) Hin-Shun Chung and Jiaya Jia: "Efficient Photometric Stereo on Glossy Surfaces with Wide Specular Lobes", Computer Vision and Pattern Recognition, pp.1-8 (2008)

(10) Zheng, Ma, Li and Chen: "An Extended Photometric Stereo Algorithm for Recovering Specular Object Shape and Its Reflectance Properties", Computer Vision and Information Systems, Vol. 7, No. 1, pp.1-12 (2010)

(11) Armand, Saito, Sugano, Takizawa: "Detecting and measuring bent pins on electronic components", SPIE Newsroom (2009)

(12) Scharstein and Szeliski: "High-accuracy stereo depth maps using structured light", Computer Vision and Pattern Recognition, Vol.1, pp.195-202 (2003)

(13) Cook and Torrance: "A Reflectance Model for Computer Graphics", ACM transactions on graphics, Vol.1, No.1, pp.7-24 (1982)

Maxime Vallée is currently studying in the Hyper Vision Research Laboratory at Keio University under a double master's degree program with the École Centrale Lyon.

Hideo Saito received his PhD degree in Electrical Engineering from Keio University, Japan, in 1992. Since then, he has been a member of the faculty of Keio University. He is now a Professor in the Department of Information and Computer Science, Keio University. His research interests include computer vision, mixed reality, virtual reality, and 3D video analysis and synthesis.

Junichi Sugano is assistant director in the research department of ViSCO Technologies

Yoshinobu Takizawa is the vice president of ViSCO Technologies.