The 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage
VAST (2006)
M. Ioannides, D. Arnold, F. Niccolucci, K. Mania (Editors)

High Quality PTM Acquisition: Reflection Transformation Imaging for Large Objects

Matteo Dellepiane, Massimiliano Corsini, Marco Callieri, Roberto Scopigno

Visual Computing Lab, ISTI - Area della ricerca CNR
Via G. Moruzzi 1, 56124, Pisa, Italy
Email: `{dellepiane, corsini, callieri, scopigno}@isti.cnr.it`

Abstract

Reflection Transformation Imaging has proved to be a powerful method to acquire and represent the 3D reflectance properties of an object, displaying them as a 2D image. Recently, Polynomial Texture Maps (PTM), which are relightable images created from a set of photos of the object taken under several different lighting conditions, have been used in Cultural Heritage field to document and virtually inspect several sets of small objects, such as cuneiform tablets and coins. In this paper we explore the possibility of producing high quality PTM of medium or large size objects. The aim is to analyze the acquisition pipeline, resolving all the issues related to the size of the object, and the conditions of acquisition. We will discuss issues regarding acquisition planning and data gathering. We also present a new tool to interactively browse high resolution PTMs. Moreover, we perform some quality assessment considerations, in order to study the degradation of quality of the PTMs respect to the number and position of lights used to acquire the PTM. The results of our acquisition system are presented with some examples of PTMs of large artifacts like a sarcophagus of 2.4 × 1 m size. PTM can be a good alternative to 3D scanning for capturing and representing certain class of objects, like bas-relieves, having lower costs in terms of acquisition equipment and data processing time.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation I.3.4 [Computer Graphics]: Graphics Utilities I.3.6 [Computer Graphics]: Methodology and Techniques

1. Introduction

Reflection Transformation Imaging has proved to be a very interesting and powerful method to acquire and represent the 3D reflectance properties of an object, displaying them on a 2D image. Typically, this is done by calculating an approximation of the reflectance function of the objects' surface for each pixel of the final image. This approximation is calculated starting from a set of photos of the object, each one taken under controlled illumination. One of the most popular techniques for reflection transformation imaging is Polynomial Texture Mapping (PTM) [MGW01]. Several applications of this technique have been proposed, mainly in the field of Cultural Heritage. The use of contrast enhancement mechanisms (specular enhancement, diffuse gain) proved to be very useful not only in term of documentation, but also in terms of analysis of the surface features. Several features which were not visible during physical inspection were discovered by this representation. Nevertheless, until now the method has been applied only to small sized objects. In this paper we explore the possibility of producing high quality PTMs of medium-large objects, from 60-70 cm of width by 50-60 cm of height up to 2m of width by 1m of height. PTMs may be an interesting alternative to the other well-known 3D acquisition techniques, like 3D scanning, which are "expensive" in terms of equipment, acquisition time and post-processing. Moreover, visualization of detailed 3D models is problematic in online environment while PTMs can be a good representation to make data more accessible to the public. Finally, PTMs seem to be a better solution for the visualization of certain objects, like bas-relieves, where the information provided by re-illumination is more important than the one provided by geometry. For all of these reasons our goal is to define a fast, flexible, low-cost system for PTM acquisition of large objects, dealing with all the issues related to the adaptation of the usual "PTM acquisition pipeline" to bigger spaces. Therefore, after an overview of related work in Section 2, in Section 3 we analyze the method of acquisition in all its aspects, from acquisition planning to data processing. In Section 4 we present our solution for browsing of high resolution PTMs, and in Section 5 we propose a preliminary analysis of quality in comparison with 3D scanning, and a study on quality degradation relatively to the number of shots used to calculate the illumination function. Some examples of objects acquired with our system are presented in Section 6. Finally, conclusions and future work are outlined in Section 7.

2. Related work

The appearance of a surface normally varies under different lighting conditions. Traditional CG techniques that involve 3D geometry modeling, specification of reflectance characteristics and lightning conditions, and global illumination rendering methods are used to capture such appearance [DRWP04, Goe04]. But modeling the reflectance function of an object can be very time consuming and computationally intensive. That is why Image-based Relighting (IBRL) can be a fast, alternative method to reproduce photo-realistic lighting effects, like subsurface scattering, interreflection, shadowing and refraction [CC06]. Reflectance imaging is a photographic process that captures views of a surface under varying lighting conditions. For a static object and camera, per-pixel reflectance functions can easily be captured and modeled. Polynomial Texture Maps is a simple but effective technique proposed by Malzbender et al. [MGW01] in 2001. For each pixel, the reflectance function is approximated by a biquadratic polynomial in the following way:

$$\begin{aligned} L(u,v,l_u,l_v) &= a_0(u,v)l_u^2 + a_1(u,v)l_v^2 + \\ &+ a_2(u,v)l_u l_v + a_3(u,v)l_u + \\ &+ a_4(u,v)l_v + a_5(u,v) \end{aligned} \tag{1}$$

where (l_u, l_v) is the direction of the incident light and (u,v) are the pixel coordinates. Hence, each pixel of a PTM image is composed by the RGB values and the six coefficients of the model function. Coefficients are calculated starting from a set of photos taken from a fixed point of view, with different light positions. In order to estimate the coefficients $(a_0, \ldots, a_5)$ the light positions have to be known. Several Cultural Heritage projects used PTM for the inspection of artifacts. The first one was imaging of cuneiform epigraphy [MGWZ00]. HP Labs PTM viewer [MGW01] was used to inspect clay cuneiform tablets under different (optimal) light conditions. Reflection transformation tools were used also in Paleontology, to provide noticeable improvement in imaging of low color contrast, high relief fossils [HBMG02]. The application of PTM method on ancient stone tools revealed fine details of concoidal knapping fractures, use scarring and stone grain [Mud04]. A joint work done by National Gallery and Tate Gallery of London showed that PTM under specular enhancement provided additional information about the surface textures of oil paintings [PSM05]. Cuneiform tablets were analyzed using both 2D (PTM) and 3D (structured light scanner) information. The PTMs were texture mapped on the model, and a special 3D viewer was created [Mud04]. Recently, the application of PTMs and scanning techniques on a large numismatic collection permitted the creation of a ”virtual exhibition” [MVSL05]. Moreover, the use of specular enhancement and diffuse gain produced an improvement in data discernment. These brief overview covers the most important applications of PTM in Cultural Heritage; but Polynomial Texture Map can have even other applications, such as investigative work [Mor03] or actor performances [WGT*05].

3. Method

The usual PTM acquisition pipeline had to be re-designed, since the acquisition of object of medium to large size presents several specific issues. In fact, typically PTM are acquired by positioning the object of interest inside a light dome of fixed size (see Section 3.2). This permits to automatically change the light direction during photos acquisition, but limits the flexibility of the overall system. Since the size of our objects is too big to create a fixed dome, we decided to deal with a “virtual” light dome as explained in the next sections. In particular, we divided the whole process in three steps. First of all we considered the physical acquisition. The size of the objects, and the fact that in most cases they cannot be moved from their place led us to the necessity to think about a new kind of non-fixed acquisition system. The proposed system is very simple, but to fasten the acquisition, the planning had to be considered a critical aspect. Acquisition of the photo set was very useful, and some work was done also for the data processing. We present our experience about each step in the next 3 subsections.

3.1. Acquisition planning

Selecting the correct lighting point is an important step in the PTM acquisition of large objects; given the size and position (in the majority of cases, on a wall) of an object, in general we do not have the possibility to use a physical dome to illuminate the object. Instead, we will have to manually place the light in different positions, forming a “virtual” illumination dome. The size of this illumination dome and its light distribution will depend on the size of the target object and on the number of light directions we want to use to sample the reflectance function of the object. To simplify the light placements we developed a software tool, called *PTM Planner*. With this tool it is possible to define the properties of the lighting dome, to visually check its correctness and to automatically generate the coordinates for the light placements. The tool usage is quite simple; the scene setup is generated as the user inputs the size of the object to be acquired, its height from the ground and the distance of the camera. Objects in the scene are scaled according to user specifications; camera is pointed towards the center of the object. Next step is the definition of the acquisition pattern. The array of light can be generated by choosing the light distance and two angles (vertical and horizontal step). The tool can automatically exclude the light positions that are near to the "wall" (there can be problems in positioning the light source in such position) and that are aligned with the camera axis (light will be shadowed by camera or will occlude the camera).

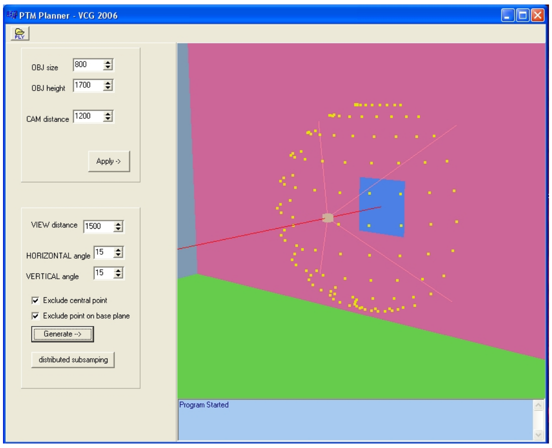

OBJ size 800
OBJ height 1700
CAM distance 1200

VIEW distance 1500
HORIZONTAL angle 15
VERTICAL angle 15
☑ Exclude central point
☑ Exclude point on base plane

Program Started

Figure 1: *PTM Planner tool.*

The points are generated using a parallel-meridian grid as showed in Figure 1. This method does not guarantee a uniform distribution over the sphere but, as we will show describing the acquisition procedure (Section 3.2), having a series of light position at the same height will result in a much faster acquisition. For some parameters, it can happen that some points are generated below the floor; those points are automatically excluded. The user can also manually *turn off* (by clicking on the 3D view) the light positions that will probably be impossible to be used due to occlusions. A more advanced possibility could be to automatically detect light position that are impossible to be used due to obstruction caused by the camera tripod (in cases where the camera is between the light and the object). Moreover, a complete description of the room geometry would be required to perform an exhaustive check of the correctness of each light position. However, we believe this level of automatic checking would be too cumbersome, giving a minimal advantage to the user, who is able to evaluate much more efficiently the presence of occluders. Finally, given a complete dome, the program can perform a light pruning following the "distributed" scheme (described in Section 5). This scheme, by generating a more uniform distribution, greatly reduces the number of required light positions while not influencing excessively the PTM quality. When the light setup has been completed, the PTMPlanner tool can save a written description of the scene. The saved data consists in:

- **Scene Data**: a review of the chosen parameters (dimension, angles), total number of photos and instructions for acquisition.
- **Acquisition Plan**: the description of all the points where the light should be positioned during acquisition. The points are described row by row, using cylindrical coordinates: as will be explained in Section 3.2, using an angle template on the floor it is very easy to place the light in the correct position by measuring its distance from the center and its height.
- **PTM Script**: the software that builds the PTM from the photos need to know the light position used for the photos; the PTMPlanner saves the position data in the required format. User will only have to manually change the filename of the images.

Even though PTMPlanner is a quite simple software, it greatly helped us in speeding up the acquisition process both during planning, by giving visual feedback and instant parameters editing, and during acquisition, by providing step-by-step instructions on light placement.

3.2. Acquisition

Several experimental devices has been created to acquire PTMs. Two of them are shown in Figure 2.

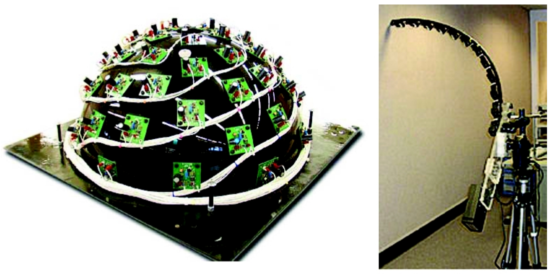
(a) (b)

Figure 2: *(a) PTM acquisition dome; (b) PTM acquisition arc (Photos: HP Labs)*

The object in Figure 2(a) is suitable for sampling small objects (nearly 15 cm.). It is a 90 cm diameter black plastic hemisphere, with fifty evenly distributed strobe lights mounted such that they illuminate the hemispheric dome's interior. The digital camera is positioned at the top of the hemisphere and photographs the PTM subject through a view port cut in the dome. In Figure 2(b) a device designed for larger objects is shown. A 90 degree arc 1.50m in diameter is mounted with 12 strobe lights facing towards the center of the arc. One end of the arc is connected to a circular bearing race in the shape of a doughnut. This allows the arc to spin in a 360 degree circle around the bearing race.

These two experimental devices work very well but they are not suitable for our target. Various reasons support this statement:

- The diameter of the "hemisphere" formed by all the light positions depends on the size of the object, since for each photo the light must completely cover the target. For the object shown in Section 5, the minimum diameter for the “virtual” dome was 3 m.
- The lights which compose the arc in the machine in Figure 2(b) can’t completely illuminate a large sized object. We need heavier and more powerful types of light.
- In most cases, the target object cannot be moved from its place, so we have to deal with the fact that it’s not always possible to exploit all the light positions, due for example to the height from the ground.

Figure 3: *Our acquisition setup.*

Following these remarks, we had to think to a simple and cheap acquisition equipment. Our solution is shown in Figure 3.

Since it was not feasible to use a big number of lights, we decided to use only one, and to change its position for every photo of the set. The time needed to position the light was minimized by an accurate planning, as shown in the previous Section, and by some references placed on the floor. We fastened the acquisition using a printed scheme of the angle directions (it helped in placing the references on the floor very quickly), and a plumb line attached to the light in order to facilitate the positioning. Our acquisition equipment was composed of an 8MPixel Canon Digital Camera, a 1000W halogen floodlight, a tripod and a boom stand. The fact that we used only one light explains also the parallel-meridian placement of lights: with these configuration we needed to set the height and direction of the light only once for each level of height. The acquisition can be summarized in this way:

- Take the measures of the object, find the center of it and its height from the ground.
- Using these data, generate the “virtual dome” and choose the positions of all the lights.
- Position the digital camera on the tripod. Measure aperture and shutter speed under the illumination of the central light. Keep these values fixed for all the photos, in order to have a constant exposure.
- With the help of the output of PTM planner, put the reference marks related to each light.
- For each level of height, set the height and the direction of the light, then put it on each reference mark related to the level, and take the photo.

Following this approach and the results of the quality assessment (see Section 5), we were able to acquire several PTMs of an object in a short time (see Section 6). Other big advantages of this equipment are that it is quite cheap (nearly 1000 Euros in total) and easily transportable.

3.3. Data processing

In order to calculate a precise illumination function, a critical factor is that the digital camera must not move from one photo to the other. Even a misalignment of a few pixel can produce a bad result, with visible aliasing. In our experimental acquisition set it’s almost impossible not to have small movements of the camera. This led to the necessity of aligning the set of photos before building the PTM. We performed the alignment automatically using a freeware tool for panoramic images. This was the only data processing we had to add to the usual PTM construction pipeline: there was no need of any image to image calibration, since all the photos had the same exposure. The PTMs were created and visualized with the software we developed. Section 4 will provide a detailed description of how it works.

4. Remote Fruition

In this section we give a brief description of another set of tools that we developed to efficiently browse PTMs of very high resolution.

As previously stated, PTMs are a good multimedia representation of artifacts since the interaction with light in general satisfies the user and its size is typically lower that of a 3D models with fine details. Nevertheless, when the resolution of PTMs is considerable high, such as thousands by thousands pixels, the size of the PTM make it necessary to find a specific solution for browsing it. In fact, a non-compressed PTM of 4000 × 3000 pixels occupies about 70MB of spaces making a download of it quite tedious. For this reason we decided to develop a specific tool to browse a multi-resolution version of huge PTMs. In this way an Internet user can interact with a low-resolution version of the PTM, while medium and high resolution portions of the image are transmitted when needed, using a progressive transmission approach.

To be more specific, two tools has been developed for this purpose; a first tool, called *HPTM Builder* which gets in input an high-resolution PTM and decomposes it properly into sub-ptms; and the effective browser, a Java applet called *HPTM Browser* (Huge PTM Browser) which handles the stored sub-ptms and browses them following a multi-resolution scheme. Both these tools are based on the Java PTM Library developed by Clifford Lyon (`http://ptmviewer.dev.java.net`).

4.1. HPTM Builder

First of all, a high resolution PTM is decomposed in several sub-ptms in a proprietary format and stored on the server. This sub-ptms represent the leafs of a quadtree, which support the progressive transmission of the level of details of the PTM. The format of these patches is slightly different from the standard format defined by HP. The HP JPEG PTM format treats the coefficients of the pixel as separated images and compress each plane coefficient using standard JPEG/JFIF compression. Before compression the range of values of each plane coefficient are reduced by a scale and bias factor. Our format differs from this one in the compression algorithms used: in place of JPEG/JFIF compression we use lossless JPEG2000 compression. We have not studied in deep the consequence of using lossy JPEG2000 compression. This aspect is left as an interesting work for future investigations.

During the decomposition each sub-ptms is properly indexed to be retrieved efficiently by the PTM browser in a multi-resolution way. In particular, chosen the number of resolution levels of the quadtree the corresponding indexed sub-ptms are generated. The indexing scheme is done accordingly to the Z-order space filling curve. This curve is particularly suitable for this purposes since it is easy to convert from the d indices of a d-dimensional matrix to the 1D index along the curve. For further information about this topic see [PF01, BKV99] In other words it is simple to navigate through the quadtree without storing it explicitly by using a matrix, filled using the Z-order filling curve, that we call Z matrix. For example, two levels of resolution produce 16 sub-ptms indexed by the following Z matrix:

$$Z_2 = \begin{bmatrix} 0 & 1 & 4 & 5 \\ 2 & 3 & 6 & 7 \\ 8 & 9 & 12 & 13 \\ 10 & 11 & 14 & 15 \end{bmatrix} \tag{2}$$

where $0, 4, 8$ and 12 are the node at the coarse level of resolution, $1, 2$ and 3 are the child of the node 0, $5, 6$ and 7 are the child of node 4, and so on. It is possible to exploit the property of this matrix to efficiently retrieve the patches at a given level of resolution. For example, as it is possible to notice, the indices of all the patches at a level of resolution l differs from $\Delta = 4^l$.

4.2. HPTM Browser

The *HPTM Browser* is a Java applet composed by two panels, a view panel where the user can interact to change the light and explore the image, and a navigation panel that shows the part of the PTM currently under viewing with respect to the whole PTM. During the interaction a thread called *update thread* loads from the server (using the HTTP protocol) the sub-ptms increasing progressively the level of details of the part of the PTM under viewing. The update thread works with a resolution table T that stores the information about the patches loaded. In fact, all the patches have the same resolution, but when a parent node is loaded it is interpolated to cover all the part of the PTM at the target level of resolution. When a level of resolution is complete the thread starts to load the sub-ptms corresponding to the immediate child patches. The process ends when all the leafs are loaded.

It is important to underline that to obtain optimal performance the update thread processes only the set of patches currently under viewing. The set of patches under viewing is calculated whenever the user pan or zoom the PTM. In this way the browser is capable to visualize very quickly a PTM at full resolution after a big zoom operation. Obviously, obtaining the whole PTM at full resolution requires a long time since all the sub-ptms have to be download.

5. Quality assessment

Besides trying to create high resolution PTMs of large artifacts, we wanted also to consider some issues regarding quality assessment. In particular, we performed some quality evaluation both with respect to the “best” representation of the object and the number and position of lights considered for acquisition. In order to perform this quality evaluation, we considered a 70 by 80 cm section of the XIVth Century Tomb of Archbishop Giovanni Masotti as a case study. We performed a very accurate PTM acquisition, using a large number of lights (105 light positions, 11 angles and 11 height levels). We acquired the same object also with a triangulation Scanner (Minolta 910i). We consider the 3D scanned model as a “ground truth”. For larger objects 3D scanning represents a very reliable technique, in terms of accuracy [BR02], even if recently it has been proposed to improve its field of normals using a combination with photometric stereo [NRDR05]. Following the pipeline described in Section 3, we created a PTM using all the 105 photos. We also generate an high-precision 3D model (nearly 2.4 millions of faces, $\frac{1}{3}$ of millimeter of sampling resolution) from a set of 68 range maps. Our first comparison was between these two representations, in order to estimate the quality of the normals calculated from the PTM data. We aligned the 3D scan model to the PTM [FDG*05] and we calculated the normals of both the model and the PTM. In Figure 4 a comparison of the normal maps is shown. The variation of the normals in the PTM is smoother than in the corresponding 3D scan, but their values are coherent. This test demonstrates that, even though PTM provides an approximation of the objects’ geometry, the obtained data are reliable. It also demonstrates that our setup doesn’t introduce errors in the representation. The other analysis was related to the degradation of PTM quality respect to the number and position of lights. For this purpose, we created four PTMs starting from subsets of

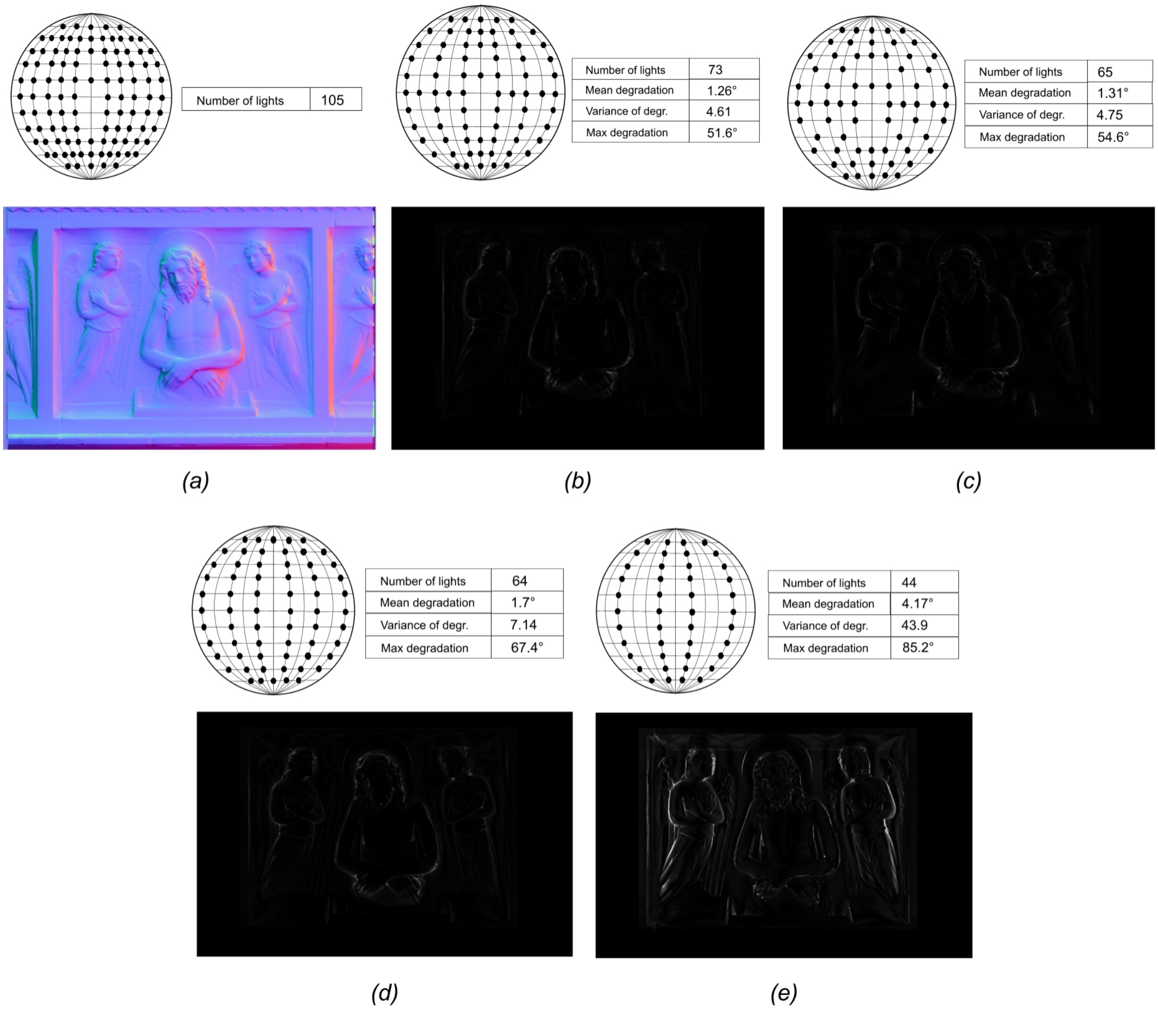

Figure 5: *Quality degradation: (a) Best quality PTM (normal map) (b-e) Maps of the differences in dihedral angle of normals. The sphere shows the lights placement.*

the original lights. Then we made a comparison between the normal maps of the "best" PTM (the one with 105 lights) and the "subsampled" ones. The comparison was made calculating the difference in dihedral angle between the normals of each pixel. In Figure 5 we show the analysis of the difference between the best PTM and four possible subsets. In terms of number of lights, we can observe that we can considerably reduce the number of lights without having an excessive degradation of quality. For example, we can reduce the number of photos up to 65 (see Figure 5(c) and 5(d)) and we will have a PTM where mean value and variance (nearly 1.5 and 6 respectively) of the overall degradation are still satisfying. As regards the different placement of lights, we can observe the case of Figure 5(c) and 5(d). Even though we have almost the same number of lights, a more "distributed" position of the lights brings to lower mean degradation and peak error. Considering these facts, we can conclude that a set of 60-70 properly distributed photos can produce a high-quality PTM. The results of this quality assessment were used to reduce the acquisition time for the objects shown in the next section.

6. Results

Several objects have been acquired with the developed system. In this Section we present 3 examples of the PTMs we produced; all this material is available for real-time exploration at `http://get-me.to/ptm` . The first example is one face of a small (30 × 30 × 20 cm.) Medieval Capital from the Museum of S.Matteo in Pisa. With the help of a professional photographer, we created a set of 36 high resolution (5440 × 4080) photos. In this case, we did not

Figure 4: *Comparison between the normal maps of the 3D scanning and the PTM: full model and particular.*

use the equipment described in Section 3.2, but a 20 MPixel Monorail View Camera and a professional flash light. We produced a very detailed horizontal PTM of the Capital: a snapshot is shown in Figure 6. The acquisition time for this object was nearly 1 hour.

Figure 6: *The Museum of San Matteo Capital.*

The second example was already considered in Section 5: it was a part (70 × 80 cm.) of the XIVth Century Tomb of the Archbishop Giovanni Scherlatti, by Nino Pisano (Museum of the Opera Primaziale in Pisa). We performed a very detailed acquisition (105 light positions, image resolution 3496 × 2280), which lasted about 3.5 hours. In this case we used the acquisition system described in Section 3.2. We were able to produce a very detailed PTM: in Fig 7 we show a snapshot. The third test object was a II Century A.D. Roman Sarcophagus, representing the Phedra and Hyppolitus Myth. This artifact is in the Camposanto Monumentale of Pisa. We chose this particular example for two main reasons: first of all we wanted to make use of the results of the quality assessment in order to perform a detailed acquisition with a lower number of lights (and consequently a shorter time). The second reason was that the Sarcophagus is situated outdoors, so we wanted to experiment if the proposed acquisition system could produce a good PTM also when the ambient light is considerable high with respect to the light equipment we used.

Figure 7: *Placed Christ from the Museum of Opera.*

Figure 8: *Roman Sarcophagus from Camposanto Monumentale.*

After 2.5 hours of work we were able to produce a complete PTM of a 90 × 60 cm portion (66 photos, resolution 3496 × 2280, see Fig. 8), and two horizontal PTMs (10 photos each) of the two halves of the Sarcophagus As can also be seen in the web site, the results was more than satisfying, considering the non ideal condition of lighting.

7. Conclusions and Future Work

In this article we extended the use of Polynomial Texture Mapping to large objects, and we demonstrated that it is possible to produce very detailed representations with a low-cost and simple acquisition system. In order to achieve this result, we had to rethink the PTM acquisition pipeline. Due to the impossibility to utilize the usual fixed acquisition dome, acquisition planning became very important. We defined also a quite simple acquisition system, and we analyzed the data processing step. After an accurate planning the acquisition can be very fast, and the equipment needed is cheap and easy-to-use. In order to preserve the high detail and the interactivity of exploration, we developed new software to progressively browse PTMs. Moreover, we studied some issues about quality assessment of PTMs. This studies gave us useful suggestions to perform the acquisition more quickly, without losing quality in the final result. All the examples gave satisfying results, and showed us that PTM can definitely be an alternative method for documenting and communicating Cultural Heritage information also for large size objects. In particular, the variation of illumination can be the best way to inspect objects such as bas-relieves or paintings.

Some future work could be useful to improve the technique. We can exploit an automatic system to estimate the light direction starting from the photo set in order to make the acquisition process more fast and accurate. Another useful feature could be the removal of the ambient lighting contribution, which lowers the quality of representation especially for outdoors objects.

Acknowledgements

We would like to thank the museums of San Matteo and Opera Primaziale in Pisa for their support. We would also thank Mark Mudge for his suggestions and enthusiasm about PTM technology. We also thank Clifford Lyon for share with us and with the Java community his Java PTM library. This work was partially funded by Epoch: European Research Network on Excellence in Processing Open Cultural Heritage (EPOCH) and Cultural Heritage Department of Italian National Research Council.

References

[BKV99] BALMELLI L., KOVACEVIC J., VETTERLI M.: Quadtrees for embedded surface visualization: Constraints and efficient data structures. In *Proceedings of International Conference on Image Processing 1999* (October 1999), vol. 2, pp. 487–491.

[BR02] BERNARDINI F., RUSHMEIER H. E.: The 3d model acquisition pipeline. *Comput. Graph. Forum 21*, 2 (2002), 149–172.

[CC06] CHOUDHURY B., CHANDRAN S.: A survey of image-based relighting techniques. In *Proceedings of the 2nd International Conference on Computer Graphics Theory and Applications* (March 2006).

[DRWP04] DEBEVEC P., REINHARD E., WARD G., PATTANAIK S. (Eds.): *High Dynamic Range Imaging, (ACM SIGGRAPH Course Notes No.13)*. ACM Press, 2004.

[FDG*05] FRANKEN T., DELLEPIANE M., GANOVELLI F., CIGNONI P., MONTANI C., SCOPIGNO R.: Minimizing user intervention in registering 2d images to 3d models. *The Visual Computer 21*, 8-10 (sep 2005), 619–628. Special Issues for Pacific Graphics 2005.

[Goe04] GOESELE M.: *New acquisition techniques for real objects and light sources in computer graphics*. PhD thesis, Universität Saarbrücken, 2004.

[HBMG02] HAMMER O., BENGTSON S., MALZBENDER T., GELB D.: Imaging fossils using reflectance transformation and interactive manipulation of virtual light sources. In *Paleontologia Electronica* (2002), no. 1.

[MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 2001), ACM Press, pp. 519–528.

[MGWZ00] MALZBENDER T., GELB D., WOLTERS H., ZUCKERMAN B.: *Enhancement of Shape Perception by Surface Reflectance Transformation*. Tech. Rep. HPL-2000-38R1, Hewlett-Packard Laboratories, Palo Alto, California, 2000.

[Mor03] MORTON M.: Foresinc control examiner for the san francisco police criminalis laboratory, presentation to the criminalistics section of the american academy of foresinc sciences, 2003.

[Mud04] MUDGE M.: Siggraph 2004 conference presentations, web graphics/special sessions/ panels, cultural heritage and computer graphics panel, 2004.

[MVSL05] MUDGE M., VOUTAZ J.-P., SCHROER C., LUM M.: Reflection transformation imaging and virtual representations of coins from the hospice of the grand st. bernard. In *Proceedings of 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST2005)* (November 2005), Mudge M., Ryan N., Scopigno R., (Eds.), Eurographics Association, pp. 29–39.

[NRDR05] NEHAB D., RUSINKIEWICZ S., DAVIS J., RAMAMOORTHI R.: Efficiently combining positions and normals for precise 3D geometry. *ACM Transactions on Graphics (Proc. SIGGRAPH) 24*, 3 (Aug. 2005).

[PF01] PASCUCCI V., FRANK R. J.: Global static indexing for real-time exploration of very large regular grids. In *on-line proceeding* (November 2001).

[PSM05] PADFIELD J., SAUNDERS D., MALZBENDER T.: Polynomial texture mapping: a new tool for examining the surface of paintings. ICOM Committee for Conservation, 2005.

[WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph. 24*, 3 (2005), 756–764.