EUROGRAPHICS 2001 / A. Chalmers and T.-M. Rhyne
(Guest Editors)

Volume 20 (2001), Number 3

A low cost 3D scanner based on structured light

C. Rocchini, P. Cignoni, C. Montani, P. Pingi and R. Scopigno[†]

Istituto di Scienza e Tecnologie dell'Informazione (ISTI)[‡] – Consiglio Nazionale delle Ricerche, C.N.R., Pisa, Italy

Abstract

Automatic 3D acquisition devices (often called 3D scanners) allow to build highly accurate models of real 3D objects in a cost- and time-effective manner. We have experimented this technology in a particular application context: the acquisition of Cultural Heritage artefacts. Specific needs of this domain are: medium-high accuracy, easy of use, affordable cost of the scanning device, self-registered acquisition of shape and color data, and finally operational safety for both the operator and the scanned artefacts. According to these requirements, we designed a low-cost 3D scanner based on structured light which adopts a new, versatile colored stripe pattern approach. We present the scanner architecture, the software technologies adopted, and the first results of its use in a project regarding the 3D acquisition of an archeological statue.

1. Introduction

Classical 3D modeling tools result inadequate to model the shape of the art works of interest in Cultural Heritage applications. This is both due to the shape complexity of most artefacts (e.g. sculptures) and also to the accuracy requested. The 3D model in many cases should not only look visually similar to the real object, but should also be very accurate, from a geometrical point of view. This to allow a number of important uses, such as the construction of 3D catalogues, the automatic reproduction of copies, the use of 3D models in the context of restoration plans, etc. 3D scanning technology has been adopted in a number of recent projects in the framework of Cultural Heritage. Just to give some examples, we may cite the Digital Michelangelo Project of the Stanford University [15], the acquisition of a Michelangelo's Pietá by a team of the IBM T.J. Watson Research Center [19], or the acquisition of a section of the Coliseum in Rome [11]. 3D scanning technology evolved in the last few years, but unfortunately not as fast as expected. One critical problem is still the high cost of high quality scanning devices. A good scanner, able to produce data at the accuracy required by the class of applications we are dealing with, often costs more than 100K US$, and therefore some recent scanning projects required very high budgets [15]. This could discourage the use of 3D scanning in many Cultural Heritage institutions. Therefore, we investigated an alternative design for a low cost and medium quality optical scanner. Our scanner has been designed to fulfill the following goals:

- use only consumer technology, to ensure affordable hardware cost and a very fast technological advance;
- support sufficient accuracy and resolution, i.e. situating midway between commercial low cost laser scanners and high quality ones;
- ensure easy of use and flexibility.

The scanner has been designed around two very common electronic devices: a video projector and a digital still camera. A video projector is used to project structured light patterns on the object to be scanned. The digital still camera is used to acquire images of the object under structured lighting. Both devices are driven by a software tool running on a standard PC, which produces the series of patterns projected by the emitter and drives the camera. Photos are taken, to acquire: images of the distorted patterns (from which the geometry is reconstructed), and

[†] Email: {rocchini, cignoni, montani, pingi}@iei.pi.cnr.it, roberto.scopigno@cnuce.cnr.it

[‡] Via G. Moruzzi 1, 56124 Pisa, ITALY

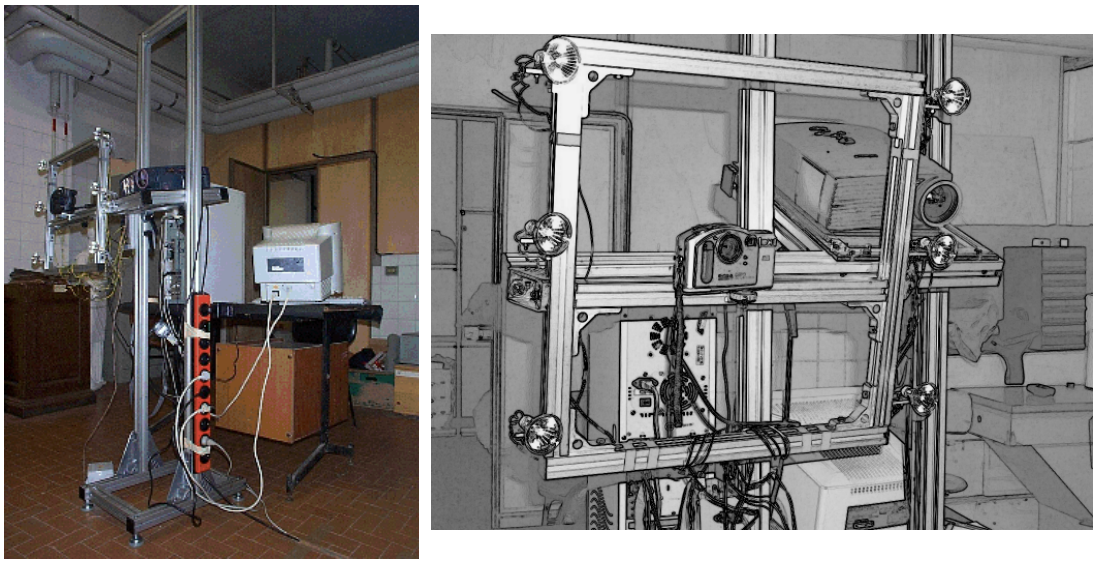

Figure 1: *The low cost 3D scanner developed at CNR, based on structured light and consumer electronic technology.*

images of the object under different illumination conditions (from which the illumination-invariant color, or *albedo*, of the object surface is reconstructed). The color images by definition are self-registered with the geometry (the range maps). The global system, depicted in Figure 1, is able to produce: a range map of each selected surface portion, with sample density of 0.7 mm in Y axis and 1.4 mm in the X axis (when a 1024*768 video projector is used, located approximately at 1.3 meters from the artifact); an illumination-invariant color texture (aligned with the corresponding range map) which represents the pictorial detail present on the acquired surface section. A limitation of our system, common to all methods which project a series of patterns, is that the scene should be static within the acquisition period.

The paper is organized as follows. A brief overview and characterization of 3D scanning technologies is presented in Section 2. The architecture of the proposed scanner is described in detail in Section 3. The device has been tested and evaluated in the framework of a complex 3D scanning session; the results are reported in Section 4. Finally, concluding remarks and possible extensions are presented in Section 5.

2. Previous work

Range scanning technology has evolved in a considerable manner in the last few years. A complete overview of the field, covering both hardware and software issues, is available in a couple of recent papers [2, 10]. Many different systems have been proposed; a common characterization subdivides them into *contact* and *non-contact* devices. An important subclass of the latter (at least, for the applications we are interested in) is the one based on the adoption of *optical* technology, and it can be further subdivided into *active* and *passive* approaches (see Figure 2). We give in the following a brief characterization of *optical* devices based on the *active* approach.

Active optical devices are based on an *emitter*, which produces some sort of structured illumination on the object to be scanned, and a *sensor*, which is typically a CCD camera and acquires images of the distorted pattern reflected by the object surface. In most cases the depth information is reconstructed by triangulation (see for example [3] and Figure 3), given the known relative positions of the emitter-sensor pair. The emitter can produce coherent light (e.g. a laser-beam) or incoherent light; in both cases, a given light pattern (point-wise, stripe-wise or a more complex pattern) is projected on the object surface. Different technologies have been adopted to produce the structured light pattern: laser emitters, custom white light projectors (which can filter light by means of a glass slide with a stripe pattern engraved via photolithography), low cost photographic slide projectors, and finally digital video projectors. One advantage of the latter is the possibility to choose from a virtually infinite number of different patterns. The pattern used should satisfy some characteristics: it should not be drastically altered by the reflective characteristics of the object surface (e.g. a colored light beam can change color depending on the reflectance of the ob-

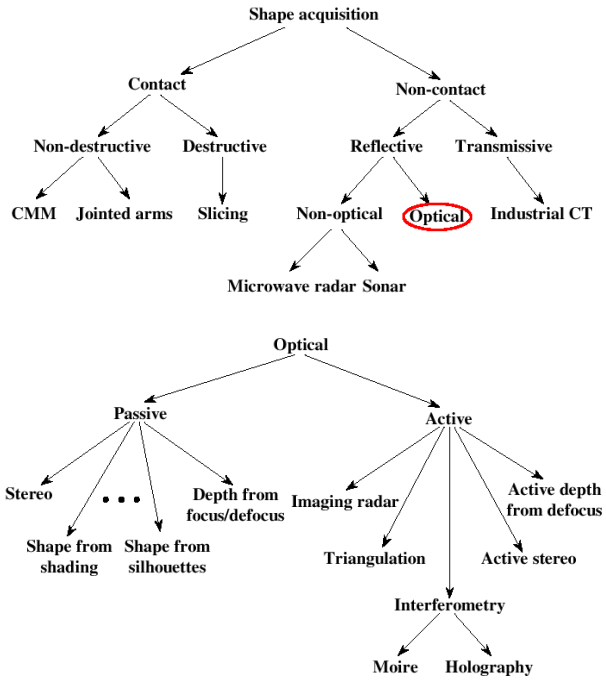

Figure 2: *A taxonomy of the systems for the automatic or semi-automatic acquisition of the shape of real 3D objects, according to* [10].

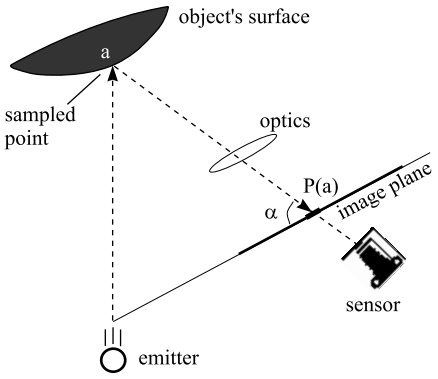

Figure 3: *A scheme of a typical optical scanner, where the 3D positions of the sampled points are computed by triangulation given the sampled point projection $P(a)$ on the sensor plane and the known relative position/orientation of the* emitter *and the* sensor.

ject surface); the detection of its constituent features (e.g. points or lines) from the reflected image should be easy and accurate; the reconstruction of the indexing of the features contained in the pattern (e.g. re-indexing the reflected stripes with respect to the original order in a stripe-based pattern produced by the emitter) should also be easy and accurate.

If we know exactly the position of the emitter w.r.t. the sensor (or viceversa), pairing the sampled camera pixels with the corresponding location on the light pattern is sufficient to reconstruct the spatial location of the corresponding point on the object surface. The reconstruction method is in general directly related with the characteristics of the pattern adopted. The simplest solution is to shoot in each instant of time a single light plane (e.g. the solution adopted by most laser-based devices), which draws a single profile on the scanned surface. Other solutions tend to accelerate the ranging process by projecting multiple beams or planes of light (e.g. series of stripes). However, the reconstruction of a multiple feature pattern is slightly more complex than the reconstruction of a single feature pattern. The sequence in which these features are reflected is not necessarily the same as that projected, because it depends on the local scanned surface topology and curvature. Therefore, a possible ambiguity arises in the identification of complex patterns, e.g. multi-stripes patterns. Some methods work only under the strong assumption of *spatial coherence* of the object surface. In the case of methods based on the projection of a single B/W multi-stripe pattern [20], a correct reconstruction of the stripe order from the reflected image is possible only if the observed surface section does not contain self-occlusions or disconnected components. This is due to the different incident angle of the emitter and the sensor: some of the reflected stripes can be occluded and are missing in the acquired CCD image. In this case the disappeared stripes can originate a wrong stripe re-indexing. This problem is often prevented by the use of either color-encoded stripes [4, 7] or set of different patterns projected sequentially on a static object [16, 20]; the latter approach is called *spatio-temporal* modulation of the illumination, and an example of a possible solution is presented in Figure 4. Several methods have been proposed in literature, and a few of them are discussed in the following Section 3.1.

3. A low cost 3D scanning tool

In the design of our low-cost scanner, we chose to implement the *emitter* unit by using a standard video projector. The reason was the flexibility of this device (which allows to experiment any type of light pattern) and its wide availability. A similar choice was also done in [6]. The *sensor* can be either a custom device, a standard digital still camera or a video camera. In our project the requirements of the sensor device are:

- it must support high quality color capture (i.e. ac-

Figure 4: *An example of a B/W binary pattern of stripes projected on the Minerva case study.*

quisition of high dynamic range) and possibly it should not perform lossy compression of the images captured;
- the resolution of the image should be as high as possible (at least, not less than the resolution of the video projector);
- it should be driven by computer, possibly by giving access to the exposition and shutter time parameters.

We chose to use a digital still camera because we gave priority to the resolution and image quality of the sensor device. Off-the-shelf digital cameras offer much higher resolution than standard consumer digital video cameras. The disadvantage is that image acquisition (and transmission to the computer) is in general slower. Finally, a personal computer generates the multiple patters, drives the digital camera (commands to acquire each image and to transmit it to the PC), and carries out the shape and color reconstruction task.

3.1. Structured light: evaluation of different patterns

One advantage of using a video projector to produce a structured light pattern is that we can easily experiment and evaluate different patterns. Some classical range finders use a single multi-stripe pattern [20]. As we pointed out previously, the disadvantage of this approach is that if the object has a complex topology and self-occlusions in the line of sight direction, then we can have ambiguous stripe configurations, and stripe indexing could become not easy. The term *stripe indexing* indicates the reconstruction of the ordinal placement, i.e. a match between each detected stripe in the image plane and the projected stripe which caused that reflected pattern. Some methods have covered this problem by encoding the stripes using multiple colors [4]. However, if indexing depends heavily on color matching, stripe detection error can be easily produced by surfaces with non-neutral color (e.g. non-uniformly painted surfaces). Other approaches make use of multiple sensors (e.g. acquisition of stereo pairs [7] or multiple images).

The indexing problem can be solved by using a series of images, produced in a number of steps by binary subdivision of a binary pattern [16, 23, 20], or by using intensity-modulated gray-level images [5, 12, 13]. The first implementation of our system used a standard binary pattern, composed of black and white stripes which were subdivided recursively until the projector resolution was reached. An example of this series of images is shown in Figure 4. A Gray code is associated to each light plane, and each illumination pattern represents one of the bits of the Gray code labels (the number of projected patterns is therefore equal to the number of bits needed to index the light stripes at the requested resolution). Stripe recognition is thus implemented using a progressive approach: the new stripes detected in image k are recognized and indexed on the base of the knowledge acquired while reconstructing images $[1, .. k-1]$. This approach needs a very precise alignment of the patterns [4] which made its practical implementation difficult in the past but very easy nowadays because video projector can support a precise alignment of the projected frames.

An improved solution has been implemented and tested in the current version of our scanner, based on the use of colored patterns. For each step we divide each region in two subregions, drawing a one-pixel-wide green line in the middle of the region and assigning red or blue colors to the points on the left and on the right of this dividing line (an example of the recursive process is in Figure 5*). The adoption of this new hybrid approach, which uses colored bands and lines, has some advantages. The colored bands allow to reconstruct easily the Gray code and the indexing of the single vertical lines (the green ones). Shape reconstruction is operated by extracting from the image the center line of a thin green line (which is analogous to the approach used in laser scanning), instead of detecting the discontinuity between two regions of different color. The latter solution has the disadvantage of being less precise due to the widening effect that a brighter light band has towards a darker one: the CCD response shifts the discontinuity border towards the darker band. Conversely, a bright green line can be slightly widened by the CCD, but its medial axis remains correct if the brightness of the two surrounding bands is approximately the same.

3.2. Scanner calibration

The initialization phase concerns the *calibration* of the scanning unit with respect to the object to be acquired, i.e. with respect to a given acquisition volume. It has to be performed only when the respective loca-

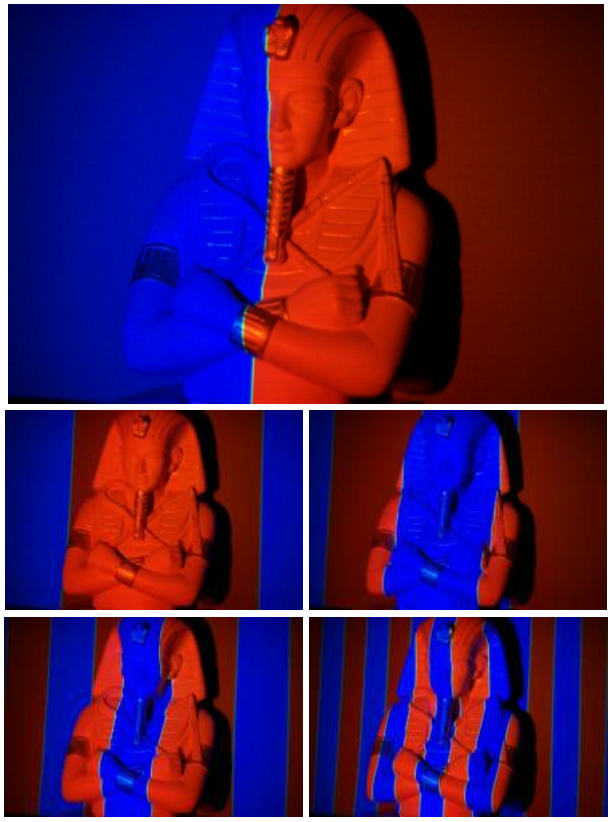

Figure 5: ** The set of RGB patterns adopted, produced by recursive subdivision.*

tion of the scanning unit and the object vary. System calibration entails:

- **Camera intrinsic calibration**. It is done once, after fixing: focus, zoom and exposition (diaframma) of the camera. It detects the intrinsic parameters of the camera: focus distance, perspective focus, radial distortion;
- **Projector intrinsic calibration**. It is done once, after fixing projector's focus and zoom. It detects the 3D parameters of the set of planes generated by each projector's vertical line of pixels;
- **Overall system calibration**. It is a by product of the previous steps, and returns the roto-translation of the projector with respect to the camera.

The *camera calibration* is performed by acquiring a *calibration panel* (see Figure 6) which is put in a set of known position in the scene such that the entire acquisition volume is swept by the planar calibration

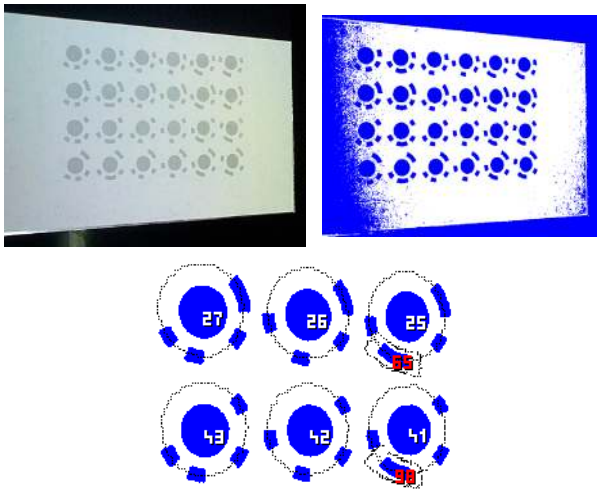

Figure 6: *An image of the calibration panel is on the left; the calibration pattern after filtering and quantization is on the right, and finally the recognized marker positions are in the bottom image.*

panel[†]. The calibration panel contains a set of markers, with known relative position, and it is lighted with a uniform white light. The markers are detected from the acquired image as follows: the input image is filtered and quantized in an adaptive manner; the markers (circle + bar code) are detected; for each marker we compute the circle baricenter and the principal axes of the projected elliptical shape (with sub-pixel accuracy); finally, each marker is indexed by decoding the associated bar code.

Therefore, for each marker center point we have the reconstructed image coordinates (u, v) and the associated known coordinates in the probe plane. Given a small set of these corresponding pairs, the *camera calibration* can be computed using the approach proposed by Tsai [21, 22].

The *projector calibration* is performed by acquiring the same calibration panel lighted with our structured light patterns and placed in two *unknown* different positions in the scanning volume[‡]. First, the 3D location

[†] Please note that this calibration step has to be performed only once for a given selected camera. This calibration is generally performed in the lab, using a track gantry which allows to sweep the calibration probe precisely in the space, passing through a discrete set of known positions.

[‡] In contrast with the first calibration step, projector calibration is usually performed on site, just before the actual acquisition starts. In this case the calibration panel is positioned in the space by hand, in the approximate location of the surface that has to be acquired.

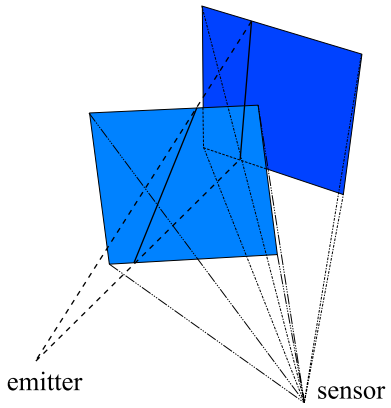

Figure 7: *The projector calibration is reconstructed by acquiring two images of the calibration panel (positioned in two non-planar locations); each light plane coordinates can be computed as the ones defining a plane passing through the corresponding pair of stripe medial axes detected in the two images.*

of the probe plane is computed automatically by detecting the markers' position and deformation in the acquired images. Second, all the light stripes projected on the probe plane are detected, and the corresponding line equations are computed in object space. The probe panel is then moved in a second non-planar location and the procedure is repeated (see Figure 7). For each vertical light stripe, we reconstruct the coefficient of the two corresponding lines on the two probe planes. The equation of the plane in 3D space which defines each vertical light stripe can be easily reconstructed by computing the coefficient of the plane which passes through these two corresponding lines and the emitter focus point.

Slight numerical inaccuracies can be introduced in the evaluation of the planes equations, due to the image-based reconstruction of the stripe central axes and to insufficient numerical accuracy. But we have some knowledge of the emitter-sensor system that can be used to improve the accuracy of the calibration phase: by definition, all light planes are regularly distributed in space and are projected from the same point; we can also assume a sufficiently good optical response of the projector. Following these conditions, the potential numerical inaccuracies can be removed by applying a global fitting phase on all the planes produced by the initial calibration procedure.

3.3. Stripe detection and range map reconstruction

After the calibration of the system, range map acquisition can start. Each set of images of the reflected patterns is then processed to reconstruct the corresponding range map. For each set of k images we perform the following steps:

- FOR EACH image i in the set DO
  - detect and purge pixels under shadow;
  - detect and index properly the green stripes in the images (per-pixel detection) and compute the medial axis of each stripe (sub-pixel accuracy);
  - add these reconstructed medial axes to the ones detected in images $[1 \,..\, i-1]$;
- FOR EACH indexed stripe j DO
  - FOR EACH image point p on the stripe j medial axis (with sub-pixel accuracy), compute the 3D point corresponding to the intersection of the line-of-sight ray associated to the point p and the structured light plane corresponding to stripe j (triangulation kernel);
- remesh the resulting 3D point cloud to produce a range surface.

Let us describe in more detail the steps above.

Pixels under shadow. Some pixels in the image correspond to shadow regions (see Figure 8*). Due to potential self-occlusion and curvature of the real objects, there could be surface regions which are visible by the sensor and hidden to the emitter. The presence of shadow regions depends on the different optical axes of the projector and the camera (approximately, they intersect with a 20 degrees angle). These pixels are detected (marked in red in the image on the right in Figure 8*) and not further considered.

Stripes detection. The images are filtered to remove possible noise peaks, and to avoid to classify a noise peak as a light band peak. For each raw image, the light stripe center is computer with sub-pixel accuracy [14]. A crucial point in the acquisition of the images of the distorted pattern is to use the right exposition, to prevent the production of over-saturated images. Because shape acquisition accuracy depends primarily on the accuracy in the reconstruction of the stripe medial axis, the saturation of the images is a crucial issue. We show with a practical example in Figure 9 the problem that can be originated by an excessive saturation. The saturation of the cropped section of the image (see the small clip on the left) makes the green line too much saturated, as it is shown in the plot of the green channel presented in the figure; the intensity slope is too abrupt, and this causes a loss

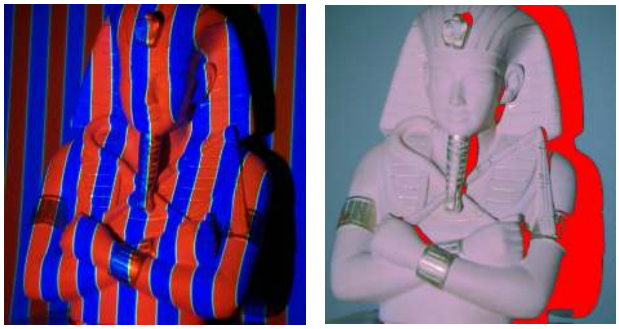

Figure 8: * *Detection and elimination of pixels under shadow (the ones coded in red in the image on the right).*

of accuracy in the reconstruction of the stripe medial axes. This can affect the quality of the 3D mesh (see the gray-shaded mesh portion in the center) where some aliasing is evident (the surface presents vertical ribs, due to a not sufficiently accurate reconstruction of some vertical set of sampling points). Conversely, a less saturated image (e.g. the plot on the right in Figure 9) makes the slope of the green channel plot less steepy and allows a more precise computation of the center axe of the stripe (with sub-pixel precision).

3D points computation. For each point s, reconstructed with sub-pixel accuracy on the medial axis of a reflected stripe, its coordinate in the 3D space can be easily computed. We know the equation of the corresponding light stripe plane and therefore we only have to evaluate a ray-plane intersection.

Range map meshing. To build a range map of the sample points detected we adopted an optimized triangulation approach recently proposed in [1], which locally minimizes the Gaussian curvature.

3.4. Surface color acquisition

The system allows to acquire also the "color" attribute of the object's surface. In particular, we are able to acquire either the texture color observed on the object surface, which strongly depends on the lighting conditions, or the illumination-invariant surface reflectance properties (also called *albedo*), computed by removing highlights and diffuse shading. The latter allows to produce higher-quality 3D models and follows an approach originally proposed in [18]: multiple images are taken from the same viewpoint and under different (and controlled) lighting conditions, to allow easy removal of the illumination effects (light shading, highlights and shadows). The relative positions of the lights and the camera are fixed (see the spot light gantry which sorrounds the camera in Figure 1), and are computed by an automatic image-based calibration procedure executed once in a pre-processing phase. The object is assumed to have a surface with *non-uniform* reflectance properties, and to behave as a Lambertian surface. Thus, we compute the object surface reflectance properties only partially: we do not produce a bidirectional reflectance distribution function (BRDF) relative to the sampled object surface [24], but determine only the diffuse reflection coefficients (i.e. the RGB albedo) for each observed surface point by inverting the illumination model. The result of this phase is a set of so called *un-shaded* images, one for each range map. This approach has been developed and experimented in a previous project regarding color acquisition and mapping on existing 3D digital models [8], and it was included in the 3D scanner.

4. Results

The performances and features of our scanner can be summarized as follows. *Range map size* depends on the resolution of the emitter and sensor device. In our case (due to the resolution of our camera, which is approximately the same size of the projector), images where green stripes pixels are interleaved by only one color-band pixel do not ensure a sufficient quality in the reconstruction process. Therefore, the hierarchical pattern projection process stops after 9 passes, producing a total of 512 green stripes (with around 3 pixels in between each pair of green lines). On each stripe we detect a number of points dependent on the resolution of the CCD camera (in our case, 1440 pixels in the horizontal axis and 960 in the vertical one), and on the width of the pattern section which actually is projected on the object surface (some stripes sections can project on the background). The resulting number of sampled 3D points stored each range image is therefore (at most) 960*512.

The *acquisition rate* depends on the speed of acquisition/transmission of the digital camera. This is the bottleneck in our system, which uses a Kodak 290 camera with an USB link: saving and transferring an image to the PC takes around 20 seconds. But this will improve in the near future with the introduction of faster interconnecting links.

The *working volume* depends on the focusing capabilities of the video projector; usually, we work at a distance of approximately 0.8 - 1.5 meters. At a distance of around 1.3 meters sample spacing on the object surface is approximately 0.7 mm in Y axis and 1.4 mm in the X axis.

The accuracy of the scanner depends on a number of factors: the optical quality of the video projector (which depends mainly on the lenses); the quality of the digital camera (i.e. resolution and color fidelity; lens radial distortion is detected and corrected in the

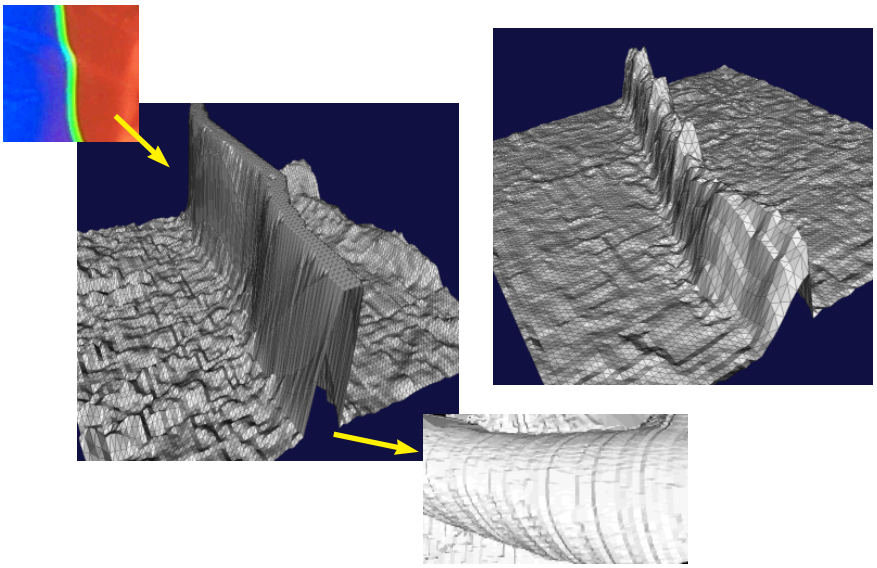

Figure 9: *The figure illustrates the problem that arise when images are too much saturated (see the text for a detailed description of the topic).*

calibration phase); the emitter-sensor angle (a small angle decreases precision, but also reduces the pixels under shadow; we generally use a 20 degree angle); finally, accuracy depends on the reflective properties of the scanned surface (e.g. acquiring highly specular surfaces is rather hard). Therefore, evaluating empirically the accuracy is not easy. We run a simple test based on the acquisition of a planar ceramic tile, and measured the mean distance of the sampled points from a fitting plane, getting a result of 0.3 mm.

The scanner has been used at the Restoration Laboratory of the Soprintendenza Archeologica Toscana in the framework of the restoration of the Minerva, a bronze statue of the Museo Archeologico at Florence (see Figure 10). As it is usual in 3D scanning, complex objects are modeled via the acquisition of a set of partially overlapping range scans. The whole Minerva statue (170 cm. high) has been scanned in October 2000 by taking 146 range maps. For each range map we take: one image of the object under white light plus 9 images to get all the stripe patterns; 6 images for the acquisition of the color (six different lighting conditions). In total we have 16 images for each range map (shape + color). The number of total images acquired to scan the Minerva, including some test ones, was therefore around 2,500. On-site scanning required four full-time days.
Only a few of these images were processed on-site (to check the correctness of the calibration and of the overall setting), while the others were processed in our laboratory in the following weeks. Once the corresponding range maps were reconstructed, our proprietary software has been used to register and merge all the range maps in a single triangulated mesh. A complete digital model of the Minerva has been produced, reconstructed from the range maps using a volumetric reconstruction approach [17] with inter-voxel distance of 0.57 mm. The mesh reconstructed is rather big: 26M faces. To improve usability, the original mesh was simplified using an external memory simplifier, based on edge collapse [9]. The results presented in Figures 11

Figure 10: *The Minerva statue; note the corrosion of the bronze surface (pre-restoration status).*

Figure 11: *A digital model of the Minerva (simplified down to 1,034,029 faces).*

Figure 12: *The Minerva head section (simplified down to 1,094,368 faces).*

and 12 are simplified models of the statue (around 1M faces) and of the head section (around 1.1M faces).

Limitations The scanner developed satisfies most of the original goals, but it has also some limitations. Being based on a multi-pattern acquisition approach, it can acquire only static objects, because the object should be static within the acquisition period (which is in the order of 6 minutes according to current low cost off-the-shelf technology). Fortunately, this is not a problem in the case of Cultural Heritage applications. The restricted depth of field of the emitter (the video projector) limits the depth of the acquisition volume that can be acquired in a single range map. But in the case of complex acquisitions, we produce many range maps and usually any surface region sampled with a range map spans a thin depth interval; therefore, this potential limitation did not show up in our scanning experience. Moreover, video projectors cannot focus at short distance (in our case, approximately 80 cm.), and this limits the sampling density on the surface. A higher sampling density could be obtained by the adoption of custom lenses.

5. Conclusions

We have presented a 3D scanner based on structured light that can be implemented using only consumer technology (a video projector and a digital camera). It adopts an innovative set patterns produced by recursive subdivision, which mix thin stripes (from which the shape is reconstructed) and colored bands (used to re-index the stripes). An important feature of the scanner is its capability of acquiring self-registered and high-resolution shape and color data. The scanner has been implemented and tested on a real, complex artwork: the Minerva of Arezzo. The scanner produced good results, even on such a complex testbed (characterized by a complex shape, with extensive corrosion and bumpyness of the statue surface, and the complex reflectivity properties of the bronze alley). We are

working to some extensions of the system, mainly oriented to improving color acquisition in the case of non-Lambertian surfaces (such as the bronze alley).

Acknowledgements This project has been developed in cooperation with the Centro di Restauro (Restoration Laboratory) of the Soprintendenza Archeologica Toscana in Florence and with the Istituto Nazionale di Ottica. In particular we would like to thank M. Cygielman, R. Giachetti, M. Giovanelli, G. Gori and L. Pezzati for the fruitful collaboration. We acknowledge the financial support of the Progetto Finalizzato CNR "Beni Culturali" and of the Tuscany Region's RIS+ Project.

References

1. L. Alboul, G. Kloosterman, C.R. Traas, and R.M. van Damme, *Best data-dependent triangulations*, Tech. Report TR-1487-99, University of Twente, 1999.
2. F. Bernardini and H. E. Rushmeier, *3D Model Acquisition*, Eurographics 2000, State of the Art Reports Proceedings, Eurographics Association, August 24–25 2000, pp. 41–62.
3. P. J. Besl, *Active, optical range imaging sensors*, Machine Vision and Applications **1** (1988), 127–152.
4. K. L. Boyer and A. C. Kak, *Color-encoded structured light for rapid active ranging*, IEEE Transactions on Pattern Analysis and Machine Intelligence **9** (1987), no. 1, 14–28.
5. Brian Carrihill and Robert Hummel, *Experiments with the intensity ratio depth sensor*, Computer Vision, Graphics, and Image Processing **32** (1985), no. 3, 337–358.
6. D. Caspi, N. Kiryati, and J. Shamir, *Range imaging with adaptive color structured light*, IEEE Transactions on Pattern Analysis and Machine Intelligence **20** (1998), no. 5, 470–480.
7. C.S. Chen, Y.P. Hung, C.C. Chiang, and J.L. Wu, *Range data acquisition using color structured lighting and stereo vision*, Image and Vision Computing **15** (1997), 445–456.
8. P. Cignoni, C. Montani, C. Rocchini, and R. Scopigno, *Multiple textures stitching and blending on 3d objects*, Rendering Techniques '99, Springer-Verlag Wien, 1999, pp. 119–130.
9. ———, *External memory management and simplification of huge meshes*, Tech. Report B4-01, I.E.I. – C.N.R., Pisa, Italy, Gennaio 2001.
10. Brian Curless and Steven Seitz, *3D Photography*, ACM Siggraph '00 Course Notes, Course No. 19, August 24th 2000.
11. M. Gaiani, M. Balzani, and F. Uccelli, *Reshaping the Coliseum in Rome*, Computer Graphics Forum (Eurographics 2000) (M. Gross and F. R. A. Hopgood, eds.), vol. 19(3), 2000, pp. 369–378.
12. E. Horn and N. Kiryati, *Toward optimal structured light patterns*, Image and Vision Computing **17** (1999), no. 5, 87–97.
13. S. Inokuchi, K. Sato, and F. Matsuda, *Range-imaging system for 3-D object recognition*, Seventh International Conference on Pattern Recognition (Montreal, Canada, July 30-August 2, 1984), IEEE, 1984, pp. 806–808.
14. R. Jain, R. Kasturi, and B.G. Schunck, *Machine vision*, Mc Graw Hill, 1995.
15. M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, et al., *The Digital Michelangelo Project: 3D scanning of large statues*, Comp. Graph. Proc., Annual Conf. Series (Siggraph '00), ACM SIGGRAPH, Addison Wesley, July 24-28 2000, pp. 131–144.
16. J. L. Posdamer and M. D. Altschuler, *Surface measurement by space-encoded projected beam systems*, Computer Graphics and Image Processing **18** (1982), no. 1, 1–17.
17. C. Rocchini, P. Cignoni, C. Montani, and R. Scopigno, *The Marching Intersections algorithm for merging range images*, Technical Report B4-61-00, I.E.I. - C.N.R., Pisa, Italy, June 2000.
18. H. Rushmeier and F. Bernardini, *Computing consistent normals and colors from photometric data*, Proc. Second Int. Conf. on 3D Digital Imaging and Modeling (Ottawa, Canada), 1999, pp. 99–108.
19. H. Rushmeier, F. Bernardini, J. Mittleman, and G. Taubin, *Acquiring input for rendering at appropriate levels of detail: digitizing a pietá*, Rendering Techniques '98 (G. Drettakis and N. Max, eds.), Springer Wien, 1998, pp. 81–92.
20. K. Sato and S. Inokuchi, *Three-dimensional surface measurement by space encoding range imaging*, J. Robotic Systems **2** (1985), 27–39.
21. R. Tsai, *A versatile camera calibration technique for high accuracy 3D machine vision metrology using off-the-shelf TV cameras and lenses*, IEEE Journal of Robotics and Automation **RA-3** (1987), no. 4.
22. Reg Willson, *Tsai camera calibration software*, More info on: http://www.cs.cmu.edu/~rgw/TsaiDesc.html, 1995.
23. H. S. Yang, K. L. Boyer, and A. C. Kak, *Range data extraction and interpretation by structural light*, The First Conference on Artificial Intelligence Applications (Denver, CO, December 5-7, 1984), IEEE, 1984, pp. 199–205.
24. Y. Yu and J. Malik, *Recovering photometric properties of architectural scenes from photographs*, SIGGRAPH 98 Conference Proceedings (Michael Cohen, ed.), Annual Conference Series, ACM SIGGRAPH, Addison Wesley, July 1998, pp. 207–218.