The 10th International Symposium on Virtual Reality, Archaeology and Cultural Heritage
VAST (2009)
C. Perlingieri and D. Pitzalis (Editors)

Pushing Time-of-Flight scanners to the limit

M. Callieri, P. Cignoni, M. Dellepiane, R. Scopigno

Visual Computing Laboratory, ISTI-CNR, Italy

Abstract

The paper describes a pipeline for 3D scanning acquisition and processing that allow to exploit the utmost precision and quality out of ToF scanners. The proposed approach capitalize on the knowledge of the distribution of the noise to apply sophisticated fairing techniques for cleaning up the data. Leveraging on the very dense sampling of this kind of scanners we show that is possible to attain high accuracy. We present a practical application of the proposed approach for the scanning of a large (5mt) statue with millimetric precision.

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

1. Introduction

3D scanning has become a widely used technology in the cultural heritage field. Many different measuring devices and software tools for data processing are available in the market and in the academic community. While this technology may be considered stable on a quite large variety of "standard" objects, there still exist many complex objects which still require lot of planning, tweaking and specific work, in order to obtain usable and high quality results. It could be said that there is a scanning technology for each object and sampling scale: conoscopy for submillimetric objects, the various kind of triangulation-based for small to human-size objects, phase-shifting TOF for 1..20 mt sized objects, Time-of-Flight for buildings, long range TOF for landscapes. However, world is not so easily confined in categories and not always the best instrument is available; it happens quite often to encounter objects that are borderline between two categories, thus being oversized for a technology and short-sized for another. A typical case is large statues (more than 3-4 meter tall) or very detailed buildings like ruins or heavily sculpted monuments: these artifacts, would require a millimetrical resolution sampling but they are too large for being practically acquired with triangulation scanners, moreover it is considered hard to sample them at high resolution with Time-of-Flight devices due to their characteristics. Newer phase-shifting TOF devices that could fit well of this class of objects are not widespread and sometimes they need significantly longer acquisition time. Two prejudice often arises when considering the Time-of-Flight technology as a possible instrument for this kind of works:

The Time-of-Flight data is too noisy Apart from the lateral resolution one of the characteristics of the TOF devices is that the range accuracy is rather independent of the actual distance of the object. So it is not important if an object is near to the scanner or half a kilometer away the range accuracy is always the same and usually considered too low for objects like large statues.

It is dif cult to build a complete model Given the long acquisition time needed for older TOF devices and the difficulties of framing difficult-to-reach areas (TOF devices are typically bigger, heavier and less orientable w.r.t. triangulation scanners), the models built using these technologies are often fairly incomplete.

The new scanners based on phase shifting or interference Time-of-Flight technology (with their greatly increased precision) would greatly helpin making things easier, but this kind of devices are still very expensive, heavy, sometimes slow and not quite diffused [CSCA05]. The installed base of pure Time-of-Flight scanners is much larger and, while requiring a more careful data processing, we will show that they can be efficiently used to produce high quality result on those borderline artifacts.

The problem of the noise is very disrupting, since the amount of noise in high resolution TOF scan may be large enough to prevents further processing. When working with triangulation range map, the sampling noise is often eliminated in the merging phase, when data acquired from different directions is blended together. In TOF scans the amount of overlap between scans, even if it is sufficient to guarantee a good alignment, may not be enough to provide enough data to perform noise canceling filtering everywhere, so the solution may be filtering each scan. However, to be able to effectively filter a single range map, the data should be dense enough to provide enough redundancy: it is necessary then to do all scans at maximum sampling density. This advice may sound strange, given the additional time it requires, but the large amount of data produced by this procedure, gives us enough redundancy to be able to efficiently and accurately filter the range maps it in order to remove the noise. We will discuss filtering techniques for noise removal in Section 3.

As a further note we remember the importance of a a careful planning of the scanning campaign for having a good coverage without spending ages working on the field. While true for all kind of 3D scanning campaign, this remark is especially important in the case of TOF devices: by considering in advance the difficulties of the scan and by finding the best scanning position it is possible to greatly reduce the time spent on the field and avoid to ”miss" some spots on the object due to a poor choice of acquisition spot. We will show in Section 2.1 how the scanning of this monument has been carefully studied in advance, in order to evaluate the best scanning strategy.

If a correct scanning plan is essential, it is also necessary to be able to process all the gathered data in the most efficient way. A very good set of processing tools (capable to deal with the large amount of data) is really important but, especially in tricky projects like this, it may be even more important to being able to ”bend the pipeline", slightly customizing the processing steps in order to be able to deal with the specific constrains/difficulties/needs of the project. In Section 4 we will detail the process used to generate the complete 3D model of the object, discussing how we tweaked the various steps of alignment, merging and color mapping in order to be able to work on this peculiar dataset.

Figure 1: *Three views of the Neptune fountain. In the Top Left: the Time-Of-Flight scanning device during the acquisition*

2. The Diagnostic Investigation of the Nettuno Statue

The Fountain of Neptune is a fountain in Florence, situated on the Piazza della Signoria, in the north-west corner of the Palazzo Vecchio. The sculptures are a work by Bartolomeo Ammannati (1563-1565) and some assistants (such as Giambologna), and was commissioned in 1565. The central piece of the fountain (see Figure 1) is a five meters tall Neptune statue, placed on a high decorated pedestal. Three small figures (mermen) are between the Neptune s legs, carrying shells that, originally, were water sprinklers. The Neptune statue is friendly called by orentines ”Il Biancone" (the white big man).

This statue is the primary objective of an extensive diagnostic investigation carried out by the Opificio delle Pietre Dure, one of the most important restoration institutes in Italy. Many different analysis will be carried out on the statue (chemical, thermographic, ultrasonic) and many experts will examine the surface of the monument. The aim of the acquisition campaign is to provide a clear overview of the condition of the statue, in order to assess the necessity of a restoration action. The 3D model will be the the base to plan the investigation campaign but also, during the campaign, the substrate on which all the analytic data will be mapped.

It will then be necessary to obtain a 3D model complete enough to be useful as a three-dimensional index for the various analysis, and precise (and realistic) enough to be used as a guide to plan the investigation and to easily locate all the interesting areas. The model should be at a resolution of around 3-5 mm (in order to show enough details), with a precision of around 1-2 mm (to be a precise index) and with surface color (in order to show the very small surface characteristics not acquirable by the laser scanner).

On each corner of the four-sided fountain there is a group of three human-size bronze statues. One of these smaller bronze figures (a satyr) was also included in the acquisition process and diagnostic investigation.

2.1. Planning of the scanning campaign

Due to the artistic (and touristic) importance of the fountain, it was not possible to build a scaffolding around the statue, not even a temporary one. This restriction made impossible the use of triangulation laser scanners, since their range of operativity is quite small, requiring to place the acquisition device very close to the statue. The only possible solution was to use Time-of-Flight devices. The no-scaffolding restriction also affected the possible placements of the TOF devices: all acquisition should have been done from ground level. These conditions posed serious doubts about the possibility of generating a good and complete 3D model; for this reason, the planning of the scanning positions has been carried out with particular care.

Being on ground level, while the statue rested on a platform at around 2.5 meters of height, gave a good opportunity to avoid unsampled areas visible from bottom. We decided that, to have the optimal view angle, the best position would be a couple meters away from the fountain border. Judging from the general shape of the statue, we would expect that the only areas not acquirable from this lower position would be the head top and the upper part of the shoulders. A single elevated viewpoint was available from the terrace of Palazzo Vecchio and we decided to exploit it to cover those difficult-to-reach areas, even though from this point it was only visible a side of the statue. Unfortunately, the buildings on the other sides of the statue were too far to obtain a coverage dense enough to overcome the measurement errors.

During the planning of the scanning campaign, we were particularly worried about the area under Neptune s legs, where three small figures are interweaved, generating a very complex geometry. The groups of statues on the angles of the fountain prevented a diagonal placement; so we decided to use two point of view from each side of the fountain, giving eight equispaced points of view. We believed that eight well distributed viewpoints could be enough to cope with the complexity of the base figures.

2.2. Data Acquisition

The scanning has been performed by using a RIEGL LMS-Z390i, a standard Time-Of-Flight scanner that exploits the time used by the laser light to travel from the scanner to the acquired surface and back to estimate the distance of the sampled point.

The acquisition process, carried out exactly as it was planned, took around four hours; like in most of this cases, almost half of the time was spent in moving the scanner from one position to another. We acquired the statues from eight scan positions from the ground, plus one from the Palazzo Vecchio balcony. One of the ground scans has been done with the device in a perfectly horizontal position, in order to preserve the vertical axis of the statue in the final 3D model.

Knowing the kind of error we could expect from this kind of device, we decided to have as much data as possible. So, we did all the scan at the maximum possible resolution (0.02 degree); even with such detailed sampling, the acquisition time was not that long. Each of the eight ground scans took between 12-18 minutes (the complex is smaller from the side position and larger from front and back), while the scan from the building took a bit less. Each range map contained around 2-2.5 million points (4-5 million triangles), giving a grand total of 19 million points sampled.

A professional photographer, during the following mornings (in order to have a less direct lighting), carried out an extensive photographic coverage of the statue. The photographic campaign produced a very large dataset of nearly 250 photos at 12 MPixels resolution. A subset of the photos taken, has been used to map color on the surface of the 3D model, this will give the diagnostic team a more clear reference to locate areas on the 3D surface.

Scanning the bronze statue was much easier; when moving around the main statue, we acquired the satyr from three positions. The only problem, in this case, was the rear of the statue: it was not possible to acquire it from inside of the tub (since the back was too close to the border) or from the other side of the fountain (too far, and not enough visibility due to the other bronze statue). For this reason we used a triangulation scanner to cover the hard-to-reach areas, combining the different scans during processing.

3. Sampled Data Filtering

The main idea driving the processing of the acquired dataset was the fact that, while the input data is very noisy, we have a good knowledge of the distribution of this noise and we exploit it during acquisition and processing. Infact TOF laser scanners are based on the principle of measuring the time used by the light to travel from the scanner to the acquired surface and back to estimate the distance of the sampled point. This information together with the knowledge of direction of the laser ray, allows to compute the position of the hit point with a given degree of accuracy. One of the characteristic of the TOF devices is that the accuracy of the measure of the distance is rather independent of the actual distance of the object, so for that part of the measure it is

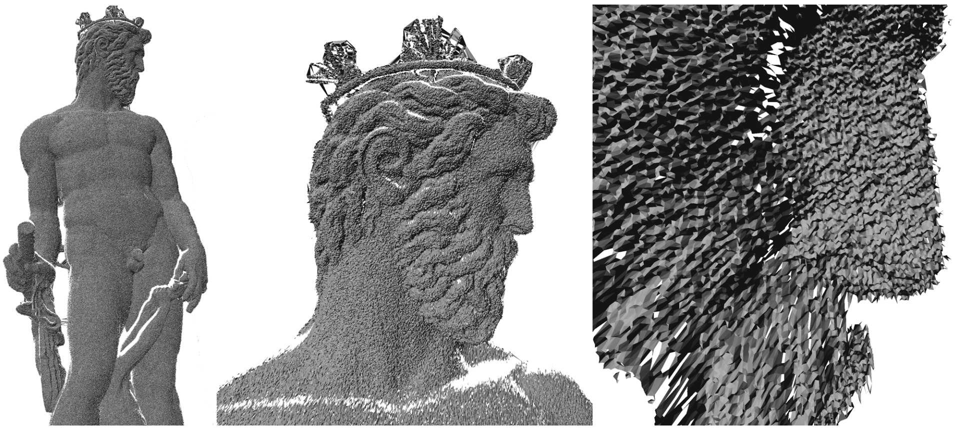

Figure 2: *A range map before filtering (Left), almost all detail is hidden beneath the noise. The two close-up (Center and Right) clearly show how the size of the noise exceed the sampling rate*

not important if an object is near to the scanner or half a kilometer away.

In practice the above considerations means that the distribution of error of the returned positions is not something constant, but it depends on the distance of the acquired points: on close to the scanner points the ray direction estimation is much more accurate than the distance estimation so the position of the measured point has an ellipsoid of uncertainty that is aligned and stretched along the ray direction, like is depicted in figure **??**.

For a number of reasons (including the fact that it is quite simple to produce a very precise and repeatable polar movement system but it is difficult to obtain really good precision in the distantiometer that measures the distance using a calculation based the the speed of light) TOF are designed to operate in the range of tens/hundreds of meters and therefore for that distance range they have reasonably well shaped ellipsoids of uncertainty. As a consequence when used on close ranges, they are able to sample very densely the surface, with a sampling step that is much smaller than the expected error in distance estimation. Common practice suggests that you should tune the acquisition parameters in order to keep the sampling step more or less of the same size of the distance estimation error, but as we show in this paper, oversampling and careful fairing can vastly improve the quality of the acquired meshes.

As expected after acquiring at an high sampling density a brief inspection of the sampled data shows the high level of noise that is present on the range maps. The range map shown in Figure 4 appear quite confuse: getting close to the surface, it is possible to see that this blurring effect is caused by the many ”spikes" that are visible on the local geometry, whose lenght exceeds the sampling rate, do hide the real geometric information.

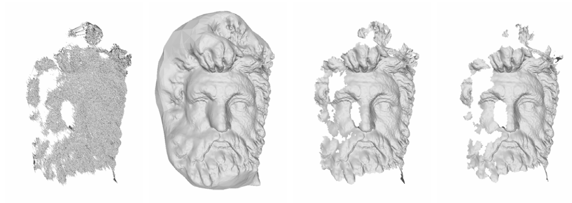

Figure 3: *Using a Poisson based surface reconstruction algorithm can remove most of the scanning noise but creates additional interpolating surfaces that can make difficult the alignment process. By removing the portion of the reconstructed surface at a distance larger than the scanner error, and further eroding the border of surface we are able to get a* safe *clean smooth surface.*

Visualizing the range maps as points, instead of surface, the details of the statue are more visible. It is clear that there is lot of information, but somehow hidden under the noise.

Considering this, it is easy to see how the sampled points are on the "correct" ray, but probably with a range error. This explains also the fact that, on close inspection (See Figure 4), the direction of all the spikes on the range map surface are clearly directed toward the scanner optical center (the origin of the probing rays). This fact is quite important for the good behavior of the advanced smoothing algorithm discussed later, as seen from the acquisition vantage points the mesh has a nice simple topology and therefore there is a correct underlying meshing that can be recovered by simply moving the vertex positions.

3.1. Poisson Filtering

A first idea to clean the sampled data was to use the Poisson surface reconstruction [KBH06] algorithm that is very

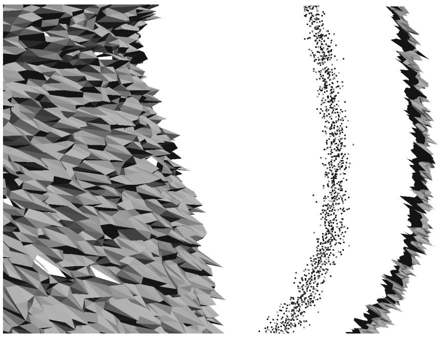

Figure 4: *A detail on the noisy surface show how there are lots of noise "spikes" with a coherent orientation. On Left: a close-up of the surface. Center: a thin slice of the range map point cloud, the size of the noise hides completely the real object surface. On Right: the same thin slice rendered as triangles, making the spikes apparent.*

robust in handling noisy datasets to obtain noise free range maps. In our processing tool, MeshLab [CCR08] we have integrated the original implementation of the standard implementation of the Poisson and by applying it to the raw range maps you can obtain clean surfaces.

The result of this filtering was good, as depicted in Figure 3, but there are some issues, the Poisson surface reconstruction algorithm always tries to build a watertight mesh closing the missing areas of the range map. Even if this feature is very useful for obtaining more complete models (as we will see later), due to the incompleteness of the mesh (after all, it is a single scan), in this case Poisson generated useless filling surfaces in many areas that can disturb the alignment process. On the other hand, the nice smoothness of the correct surface make them very suitable for the alignment process, there fore we automatically purged the reconstructed surface removing all the portions that were farther than a given threshold form the original surface and then we additionally eroded the boundaries of this surface removing a few strips of triangles just to be sure that the during alignment we used only portion of the surfaces corresponding to the original range maps. In this way we got clean range map, that even if they have some loss on the high frequencies, they are significantly smaller and more manageable for an easier alignment.

3.2. Evaluating different filtering strategies

Given the nature of the noise, we leveraged on the vast arsenal of smoothing algorithms that are implemented in MeshLab and tried what worked in the best way for our dataset. Each smoothing algorithm has different characteristics:

- *Laplacian Smoothing*; this is the classical Laplacian smoothing [ABE97] that moves each vertex onto the average position of their 1-ring adjacent vertices.
- *HC Laplacian Smoothing*; this is a improved version of the classical Laplacian smoothing, it exhibits less shrinking effects that the standard approach and has a better behavior about oversmoothing [VMM99].
- *Two Step Smoothing*; this approach, described in [SB04], is based on the idea of operating in two phases: first the smoothing the normals of the faces and then the fitting the vertex position to a surface passing through this new normal field. If the starting surface is not very noisy this approach it is strikingly able to preserve large high frequency features in a very good way (e.g. it can maintain long sharp edges, while smoothing curved surfaces).

We carefully tested these smoothing filter in order to choose the best combination for our particular case. In figure 5 we show the application of the above smoothing filters over a very small (10x10cm) patch of the whole statue. From left to right the depicted meshes represent:

- the original unsmoothed surface; note the high level of noise present in the patch;
- the surface after the application of a Laplacian smoothing. Most of the noise is gone, but also high frequencies have been affected. Moreover the surface suffered a significant shrinking.
- the surface after the application of a HC Laplacian smoothing. Similar to the previous mesh, but with a smaller shrinking effect.
- the surface after a direct application of a Two-Step smoothing. Most of the noise is gone, but the excessive noise of the starting mesh has biased the fitting procedure so the final mesh is affected by a significantly bad triangulation.
- the surface after the application of very moderate HC Laplacian smooth, that reduced only partially the noise, and the subsequent application of the Two-Step smoothing. The combined action of these two filters allow to recover at best the original surface: in this case we succeed to remove all the noise and to preserve large scale high frequency features like the sharp valley and the crest on the right side of the patch.

We want to remember that the proposed method works really well only when the initial data is very dense. This kind of filtering heavily rely on the level of redundancy to be able to eliminate the white noise in the acquisition. As a rule of thumb, it is possible to say that, in order to work properly, the initial sampling rate of the scanning should be at least two times the "desired" level of precision. This may be not covenient, especially when using older (but quite common on the field) Time-of-Flight scanners that are much slower than triangulation devices. However, depending on the project requirements, the result may be worth the extra time.

4. Complete Model Generation

Two problems were still open: how to manage the amount of data to be aligned in a common reference space and how to

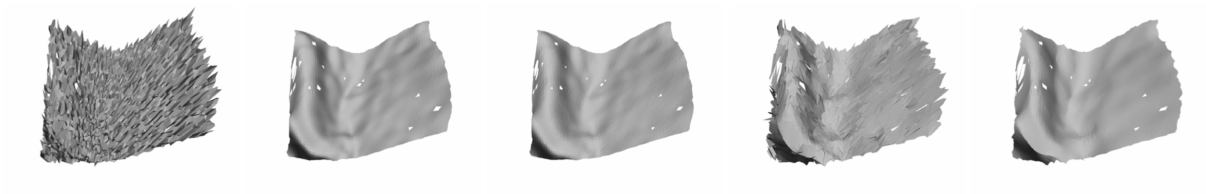

Figure 5: *Different smoothing algorithm on a very small patch (10x10 cm) of the statue. From left to right, original rangemap, Standard Laplacian, HC-Laplacian, Two-Step smoothing, and combined HC-Laplacian and Two-Step smooth.*

deal with the unsampled areas. In most cases, the data coming from a Time-of-Flight scanner is aligned using markers that are placed in the scene at the scanning time. Again, for the non-invasive constrains of the project, it would have been quite difficult to place markers inside the fountain and on the Neptune body. It was then decided to use the standard ICP based approach for the alignment. Then, even if the scanning had provided a very good coverage, some of the recesses between the small figures and the Neptune resulted in unsampled areas. To cope with these missing areas, we decided to use a peculiar hole-filling approach, in order to complete the statue in an effective (but controllable) way, as we will show in the following sections.

4.1. Alignment

The alignment of the range maps was quite difficult for two main reasons:

- the size of each individual map, from 3 to 5 millions triangle each, that posed several memory issues
- the amount of noise, that prevented a good surface-to-surface adherence

To solve those problems we tried to align the Poisson-filtered meshes instead of the original range maps. As we said in the previous section, the mesh generated by the Poisson filter are much smaller (around 600k triangles, making the alignment much faster) and they have a lot less noise, even if some of the high frequency details have been lost. The watertight reconstruction with the false areas is not really a problem, since the expansion ares can easily be identified and deleted, remaining with just original filtered data. By using the Poisson meshes, we were sure that the alignment would have been a problem of much lower computational difficulty. The question was then: if using such a low-res mesh, would the result still be acceptable?

The idea was that, even with a reduced level of detail, the filtered mesh would still contain enough data good for the alignment algorithm. We speculated that, while smoothing out some of the original mesh high-frequency details, the Poisson filtering would preserve the kind of low-frequency details that can be used to select map-to-map correspondences (which are used for the rough alignment) and that prevents the sliding of range maps during the ICP alignment (which is the main problem when dealing with low-resolution, smooth range map). Clearly, the overlapping portion between the shots and the complex topology of the object helped in obtaining a good, rigid alignment. Again, the tight scan coverage that has been decided during the acquisition planning proved to be a wise choice; the extra time required for the scanning was well worth.

4.2. Two-step Merging

Generally, a 3D model built using data from 3D scanning contains lots of small unsampled areas. In literature, there are various algorithms and approach to deal with such incomplete areas and produce closed (sometimes watertight) 3D models. However, in our experience, workers in the cultural heritage field are very conservative and do not want any of this unsampled areas to be filled because doing so would introduce data that has not been measured, thus invalidating the scientific value of the model. In this particular case, however, the main aim of the model was to provide a support for the diagnostic investigations. So, it was more important to provide a surface as complete as possible with respect to have an all-original surface. We decided to aim at a complete reconstruction of the surface, but using a method that could somehow preserve the integrity of the original data. After a first step of volumetric merging, the result was a good but incomplete model: most of the incompleteness were quite small and located, as expected, in the gaps between the three small figures and the leg of the Neptune. Two areas (the left shoulder and the top of the head) were missing due to the impossibility to be acquired from the available scanning position. This model, produced with only the original data, was the base to provide a good quality hole filling to build a complete model to be used in the documentation of the diagnostic project. Again, using the Poisson filter, we were able to process the model from the first merging result to produce a closed surface; this result alone, however, was not what we wanted since, due to the Poisson smoothing, lots of the original detail was lost. So, we used this closed version treating it like an additional starting mesh data, by performing a second merging of the original range maps with this ”filler" geometry. To be sure of using the filler data only when necessary (on holes) and to avoid an excessive smoothing effect, we assigned a very low weight to this additional geometry (1 hundred times lower) with respect to the other range maps.

As shown in Figure 6, the result was very convincing closed surface, that followed the sampled data (where available) and interpolated the small missing areas in a very natural way. However, the hole-filling algorithm could not cope with a severe lack of data on one corner of the base of the statue. The area of the north-west corner of the base was hidden from all the acquisition viewpoints, and there was not enough information in the surroundings area to make the Poisson filling able to ”guess" the correct shape of the missing part. We plan to integrate this area during the diagnostic analysis, when a movable platform will make possible to use triangulation scanner near enough to the statue. We checked all the other parts that have been filled using this two-step approach, and found no significant deviation from the geometry of the actual statue. This check has been carried out during the color mapping phase: photos that covered an interpolated area has been registered to the 3D model using points on the original sampled data. Then, the silhouette and features of the filled areas has been compared to the silhouette and features on the photos; as said, no major discrepancy was found.

Figure 6: *The untextured geometry of the final model, after the two-pass merging that filled missing areas.*

A nice characteristic of this method is that, thanks to the two-step merging we know exactly which parts of the final model are based on the original data and where the filler has been used to close the holes. This information has been stored in the obtained model, in this way combining surface completeness and data integrity.

4.3. Color Mapping

The main goal of the acquisition of the Nettuno statue was to provide a 3D model which could have been extensively used by expert to plan and perform the diagnostic campaign. In order to do so, it should have been possible to restorers and technicians to precisely locate areas of interest on the surface of the statue. To simplify this task we decided to enrich the geometry of the 3D model using the color information coming from the photos. The idea was that color could replace the very small geometric features (small cracks, surface imperfections, crusts) that could not have been captured in any case with a TOF scanner, making those features easily identifiable.
A very detailed photographic campaign (nearly 250 images) was performed on the statue, trying to follow the main requirements needed to obtain a high quality colored model (overlaps between images, very small difference in illumination condition between shots, accurate surface coverage).
A subset of the images (nearly 70) was then registered [FDG*05] and projected [CCCS08] on the 3D model. Some snapshot of the final result are shown in Figure 7 and Fig. 8. The presence of the ”filled" parts did not prevent from obtaining both a good registration and an accurate projection of the color information.

5. Conclusion

In this paper we have presented the results of an acquisition campaign that was performed exploiting at its best capabilities a Time of Flight scanner. We relied on a very high density sampling to comply with the large acquisition error returned by this kind of devices. The high density sampling allowed to apply advanced smoothing algorithm over the single acquired range maps.

With the described procedure we were able to reconstruct the statue with a precision that challenges the the maximum precision that is possible to reach with this device in controlled laboratory settings as declared by the scanner manufacturer, approximatively increasing of almost an order of magnitude the precision that you can get using this devices on the field with standard procedures.

Acknowledgements

The research leading to these results has received funding from the EG 7FP IP “3D-COFORM" (2008-2012, n. 231809). The 3D acquisition of the Nettuno statue was commissioned to ISTI-CNR by the Comune di Firenze; it was executed in cooperation with Studio Titoni (a digital survey company located in Pisa).

References

[ABE97] AMENTA N., BERN M., EPPSTEIN D.: Optimal point placement for mesh smoothing. In *Proceedings of the eighth annual ACM-SIAM symposium on Discrete algorithms* (1997), Society for Industrial and Applied Mathematics Philadelphia, PA, USA, pp. 528–537.

Figure 7: *Snapshots of the colored model. Panels on Right: four close-up of the colored surface, most of the very small features are clearly visible*

[BUSM06] BALZANI M., UCCELLI F., SCOPIGNO R., MONTANI C.: La cattedrale di pisa nella piazza dei miracoli: un rilievo 3d per lÕintegrazione con i sistemi informativi di documentazione storica e di restauro. In *Restauro 2006: Salone dell Arte del Restauro e della Conservazione dei Beni Culturali e Ambientali*. Acropoli srl, 2006. Ferrara, Aprile 2006.

[CCCS08] CALLIERI M., CIGNONI P., CORSINI M., SCOPIGNO R.: Masked photo blending: mapping dense photographic dataset on high-resolution 3d models. *Computer & Graphics 32*, 4 (Aug 2008), 464–473. for the online version: http://dx.doi.org/10.1016/j.cag.2008.05.004.

[CCR08] CIGNONI P., CORSINI M., RANZUGLIA G.: MeshLab: an open-source 3D mesh processing system. *ERCIM News (73)* (2008), 45–46.

[CSCA05] CIOCI A., SPINETTI A., CAROSSO L., ATZENI C.: Digital Three-Dimensional Modeling of Heritage by Frequency-Modulated Laser Radar: the case of DonatelloÕs" David. In *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage* (2005).

[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.

[KBH06] KAZHDAN M., BOLITHO M., HOPPE H.: Poisson surface reconstruction. In *Geometry Processing 2006: Fourth Eurographics Symposium on Geometry Processing, Cagliari, Sardinia, Italy, June 26-28, 2006* (2006), Eurographics, p. 61.

[SB04] SHEN Y., BARNER K.: Fuzzy vector median-based surface smoothing. *IEEE Transactions on Visualization and Computer Graphics 10*, 3 (2004), 252–265.

[VMM99] VOLLMER J., MENCL R., MULLER H.: Improved Laplacian smoothing of noisy surface meshes. In *Computer Graphics Forum* (1999), vol. 18, Blackwell Publishers Ltd, pp. 131–138.

Figure 8: *The final model reconstructed and textured.*