3D Reconstructions for the Humanities and Engineering: Tomography and Digital Photogrammetry

George Bevan
Associate Professor, Classics, Art Conservation/Art History, and Geological Sciences and Engineering

3D reconstructions of either volumes or surfaces by Tomography or Digital Photogrammetry are very computationally demanding. Both techniques have seen increased use within the humanities by archaeologists, art historians and conservators, as well as in industry in areas such as Precision Agriculture, Civil Engineering, Mining, and Geological Engineering. Compute times for large data-sets can be measured in days, not hours, particularly those derived from the latest high-resolution sensors. New tomographic reconstruction techniques such as the Simultaneous Algebraic Reconstruction Technique (SART), an algorithm well suited for limited-angle tomography, a technique useful for large, thin, objects like paintings, are extremely demanding but can be easily parallelized with existing, commercial packages like Octopus (https://octopusimaging.eu/). SART is also important for other objects that suffer from photon or neutron-starvation artifacts in standard cone-beam reconstruction, usually seen as star-like streaking across the reconstructed slices that obscure important features. For instance, pacemakers and hip-implants often cause star artifacts that obscure soft tissue scans looking for tumors. This same algorithm can be used for looking at layers of paint, or the interior of statues.[1]

Digital Photogrammetry has been proposed as the best option for recording endangered cultural heritage, but the latest algorithm used to process the sub-optimal imagery captured by amateur photographers, Semi-Global Matching, is computationally expensive. Normally compute-times are reduced by the use of the GPU, and by down-sampling the original imagery so that instead of, say, 36 mega-pixel images being used for dense-stereo matching, 12 mega-pixel images are used instead. This image resampling not only increases the ground pixel size by a factor of four-- ground-pixel size, along with stereo base, is the key parameter in photogrammetric accuracy -- but also leads to undesirably quantizing effects, usually observed as symmetrical ripples on the models, particularly on sharp edges. With commercial software like Agisoft Photoscan, hardly any researchers or industry professionals use full-resolution images for their reconstructions and so lose important data in their 3D models (http://www.agisoft.com/). Recently, Photoscan has offered a new, distributed computing option that could significantly reduce the compute-times for dense stereo matching at full-resolution, as well as the meshing and texture-mapping of the surface of large photogrammetry projects (>500 images). The details are contained in a new 2015 white-paper.[2]

A workflow that depends on desktop or even workstation-class machines to do the heavy lifting is cost ineffective and does not scale particularly well. Assigning expensive workstations to weeks of computing means that they cannot be used for other interactive phases of 3D analysis. While distributed computing and GPU computing has been used to improve reconstruction times, the use of multiple workstations connected by ethernet and

[1] http://www.ncbi.nlm.nih.gov/pmc/articles/PMC2780629/
[2] http://www.agisoft.com/pdf/photoscan_network_processing.pdf

equipped with many GPU cores is cost-prohibitive for most researchers in the humanities. What is needed is a cloud-based service for 3D reconstruction, not unlike the established practice of “render farms” in the 3D graphics industry that can be used when needed and then released to other users when no longer needed. The need for a scalable, cost-effective solution is felt especially keenly in the humanities where funding constraints seldom allow for regular updates of equipment, or even large hardware purchases.

The advent of Cloud Computing tools like hardware-accelerated virtualization using OpenStack and large shared-memory systems like the SGI Ultra-Violet 300 in the last two years, means that the cost-effective scaling-up of Commercial Off-the-Shelf (COTS) software like Octopus and Photoscan, both of which are well known and respected by researchers, is now possible.[3] In the future researchers such as myself would like to see shared memory systems like the UV 300 running OpenStack, along with dedicated GPUs, to run this commercial code. We would propose creating many virtual workstations in memory and linking them through existing distributed computing protocols over virtualized ethernet to realize what we hope would be a speed-up in reconstruction by at least an order of magnitude. While custom HPC code to replace the COTS packages may be desirable in the future, most end-users would be willing to see a marginal decline in speed so that they can retain familiar user-interfaces and predictable application behaviour.

We believe that Compute Canada is well positioned to offer these services to humanities scholars in the coming years, and would like to see these applications taken into account in the next round of capital expenditures.

[3] http://blog.sgi.com/sgi-uv-300-ultimate-virtualization-platform/
http://www.nextplatform.com/2015/04/03/openstack-and-kvm-give-sgi-big-iron-os-options/