A Distributed Computing Workflow for Modelling Environmental Flows in Complex Terrain

Stuart R. Mead^{1,2,*}, Mahesh Prakash², Christina Magill¹,
Matt Bolger², and Jean-Claude Thouret³

¹ Risk Frontiers, Faculty of Science, Macquarie University, Sydney, Australia

² CSIRO Digital Productivity Flagship, Melbourne, Australia

³ Laboratoire Magmas et Volcans UMR6524 CNRS, IRD and OPGC, University Blaise Pascal,
Clermont-Ferrand, France

Stuart.Mead@mq.edu.au

Abstract. Numerical modelling of extreme environmental flows such as flash floods, avalanches and mudflows can be used to understand fundamental processes, predict outcomes and assess the loss potential of future events. These extreme flows can produce complicated and dynamic free surfaces as a result of interactions with the terrain and built environment. In order to resolve these features that may affect flows, high resolution, accurate terrain models are required. However, terrain models can be difficult and costly to acquire, and often lack detail of important flow steering structures such as bridges or debris. To overcome these issues we have developed a photogrammetry workflow for reconstructing high spatial resolution three dimensional terrain models. The workflow utilises parallel and distributed computing to provide inexpensive terrain models that can then be used in numerical simulations of environmental flows. A section of Quebrada San Lazaro within the city of Arequipa, Peru is used as a case study to demonstrate the construction and usage of the terrain models and applicability of the workflow for a flash flood scenario.

Keywords: Structure-from-Motion, photogrammetry, numerical modelling, rapid mass flow, natural hazards.

1 Introduction

Extreme environmental flows and mass movements such as floods, landslides, avalanches and debris flows pose significant risk to exposed populations and can cause substantial damage to buildings, infrastructure and the environment. Due to safety concerns and difficulty in predicting the occurrence of these events, field measurements are rare and, when available, are generally limited to depth and point velocity measurements (e.g. [1]). As a consequence, computational flow models are commonly employed in an attempt to predict the outcomes of specific events, understand fundamental processes and to provide a greater understanding and delineation of the hazard.

* Corresponding author.

Two-dimensional (2D) depth-averaged numerical models such as the shallow-water (SW) method for fluids or Savage-Hutter method for granular mass movements are widely used to simulate environmental flows. These depth-averaged approaches are suitable for predicting large scale flow features and inundation; however the shallowness assumption can be limiting within complex environments (such as urban areas) where there are large and sudden changes in the terrain gradient. Complex and varying topography causes the flow to have three dimensional features and varying vertical velocity profiles which cannot be captured by depth-averaged approaches. For these circumstances, three dimensional (3D) particle based flow modelling methods, such as smoothed particle hydrodynamics (SPH), may be more suitable as they have the ability to predict and track the motion of objects or debris within the flow, model complex flooding scenarios including interaction with buildings, and predict forces on structures [2]. Use of SPH and other 3D methods has traditionally been limited due to their high computational requirements. However, advances in processing power and parallel computing solutions such as OpenMP and Message Passing Interface (MPI) have created opportunities to simulate larger areas at higher resolutions.

In both depth-averaged 2D and fully 3D flow modelling approaches, topographic information (in the form of terrain models) is required as a primary input and modelled outcomes are highly sensitive to the accuracy of this topographic data [3]. For example, Legleiter, Kyriakidis, McDonald and Nelson [4] examined the effects of uncertain topographic data on a typical 2D depth-integrated (i.e. Shallow Water) method, finding that the method was sensitive to morphological features such as point bars. Small-scale bed features were also found to affect results, albeit at a reduced level as flow depth increased. The applicability and accuracy of simulation predictions is therefore reliant on accurate reproduction of small scale flow-steering features and topographic obstacles in the terrain model.

While advances in computational power and remote sensing methods have generally increased the dimensionality and resolution of terrain models [5], quality and accuracy has not necessarily increased at the same pace. Measurements of the terrain are commonly acquired from satellite or aerial platforms using techniques such as LiDAR, stereo photogrammetry or radar interferometry. This information is then used to create gridded digital terrain models (DTM), a representation of the topography at a single point in time. Frequent data acquisition is critical to ensuring currency of the DTM, as sediment transport processes and human interaction continually modify the terrain and create small scale features that may affect the behavior of flows. The frequency of data acquisition is, however, limited by high computational requirements, costly equipment and lengthy data processing times [5], resulting in an inaccurate representation of the current shape and features of the terrain. In addition, terrain models can be limited by the acquisition platform. High altitude aerial platform methods are not suitable when the area of interest has large amounts of cloud cover, which is often the case before and after weather based flow events. Tall trees, buildings and bridges may also cause occlusion, leading to inadequate representations of key terrain features.

For high resolution modelling, accurate reproduction of small-scale terrain features and the 3D terrain structure is particularly important, although grid based terrain models are limited in their ability to reproduce these 3D features. Reasons for this have been described by Kreylos, Oskin, Cowgill, Gold, Elliott and Kellogg [6] and include:

- Constraints of a rectangular grid system limits the gradient of the slope, and therefore flow direction, to eight cardinal directions.
- Re-projection of steep topographic features, such as dips, overhangs or bridges, onto a horizontal plane degrades the level of detail captured, implicitly reducing the resolution.
- Gridding can have a directional dependence, and features such as channels that do not align with the primary grid direction can be degraded.

Three dimensional terrain representations, such as triangulated irregular networks (TIN) or point based methods [6] present a feasible alternative to reduce the limitations of grid-based terrain models as well as more accurately reproducing terrain features.

Here, we demonstrate a workflow to acquire and process detailed 3D terrain models for use in 3D flow modelling and visualisation. The process utilises and integrates several core technologies to provide a fast, reliable and inexpensive method for terrain generation and numerical modelling that eliminates or reduces many of the common limitations of traditional terrain models. Open source libraries and data structures are used to enhance the flexibility of the process and allow users to extend this methodology. The workflow utilises parallel and distributed computing to allow for rapid turnaround time from terrain acquisition to output of numerical simulation results. The key components of the workflow are (1) acquisition of ground images and low altitude aerial images from a light remote control quadcopter, (2) image feature detection, matching and 3D scene reproduction using Structure-from-Motion (SfM) photogrammetry, (3) point cloud processing, including reprojection and meshing; and (4) numerical modelling using SPH. The low cost and high speed of this method allows for faster acquisition of terrain models, increasing their currency, while the point based approach enables features to be more accurately represented. The application of this workflow is demonstrated as a case study by simulating flash flood scenarios in complex terrain for the city of Arequipa, Peru.

2 Methods

2.1 Workflow Outline and Software Integration

The four stages of the terrain generation and modelling process are shown in figure 1. The first stage, image acquisition, involves capturing images of the terrain from multiple angles, taken either aerially and from ground level. The main objective of this stage is to collect as many images of the area of interest as possible. The Structure-from-Motion (SfM) stage uses these images to generate a point cloud representation of the area using photogrammetry and computer vision methods including feature

identification and matching, camera pose estimation and point cloud reconstruction using multi-view stereo (MVS). The dense point cloud produced by the SfM stage is then processed using algorithms to filter out erroneous points, reduce the size of the point cloud and reconstruct a manifold, watertight surface mesh representing the terrain. In the final numerical modelling stage, the terrain mesh model is converted into an input for simulation and subsequent visualisation.

Fig. 1. Terrain generation and modelling workflow outline

The modelling and terrain generation process, shown in figure 1, is sequential with each step relying on inputs from the previous. Aside from the requirements for inputs, each component of the process is independent from each other. This makes the process well suited for implementation into a workflow tool, where each software component can be developed and modified independently. This is particularly important for image feature detection (step 1 in figure 1), surface reconstruction (step 6) and numerical modelling (steps 7-9), as these are all active areas of research where a range of approaches (e.g. fully 3d or 2d depth-averaged numerical models) could be employed and may be more appropriate depending on user requirements.

We chose to implement this process (referred to hereafter as the modelling workflow) into Workspace, a cross-platform workflow framework with a plug-in architecture [7]. The plug-in architecture exposes data types and operations, enabling the development of new workflows and operations. Other advantages of Workspace for this application are:

- 3D visualisation capabilities allows for integrated model checking, without the need for additional software packages,
- the distributed/parallel execution engine, which allows components of the workflow to be processed on multiple CPU's, both locally and distributed through TCP and cluster based systems, and,
- the engine is built upon permissive free software licenses, allowing for extensions and applications to be developed without third party licensing requirements.

Currently image feature matching, point cloud processing (steps 4-6), input generation (step 7) and visualisation (step 9) are implemented within Workspace. The triangulation, bundle adjustment and multi-view stereo processes are run separately, but

are executed within the workspace from a command line operation. The SPH simulations (step 8) are run externally using the terrain model generated within the workflow.

2.2 Structure-from-Motion (SfM)

Structure-from-Motion (SfM) is a photogrammetry method that reconstructs sparse 3D points and camera viewpoints from image collections of a scene. Dense 3D point clouds can be reconstructed by coupling with MVS methods, which construct dense 3D point clouds from overlapping images and camera viewpoints [8]. The accuracy of SfM-MVS approaches is said to be nearly on par with laser scanners [9].

To obtain camera viewpoints, the SfM process is as follows: (1) within each image, find and compute descriptors for unique features such as building corners (keypoints); (2) find matching keypoints in images taken from different angles; (3) obtain an estimate of camera parameters and location through triangulation of the image matches; and, (4) optimise the matrix of camera parameters and locations (bundle adjustment). For a detailed description of the SfM process, see [10].

In the workflow, our approach differs slightly from the commonly used methods presented by Snavely, Seitz and Szeliski [10] in that we utilise the open source library OpenCV (www.opencv.org) to compute and match features. OpenCV is a BSD-licensed computer vision and machine learning software library with a feature detection and matching framework that allows for a variety of different keypoint extraction and descriptor methods. Currently implemented matching methods are SIFT [11], SURF [12] and BRISK [13]. These provide a range of keypoint descriptors that can be utilised depending on requirements for speed, robustness or commercial purposes (SIFT and SURF are patented methods, freely available for non-commercial use only).

Following the triangulation and bundle adjustment stage of the workflow, the camera pose and sparse point cloud are used to create a dense reconstruction of the scene using MVS. The MVS method used here implements view-clustering to group images in order to reduce memory requirements and reconstruction speed. For details on the view-clustering MVS method used here refer to [14].

2.3 Point Cloud Processing and Surface Generation

Following dense reconstruction, the point cloud is converted into a structure that can be used in the Point Cloud Library (PCL), a large scale open project for point cloud processing [15]. The PCL libraries contains algorithms for filtering, feature and object detection, processing and smoothing, model fitting and surface reconstruction. Currently, the following PCL operations are implemented in Workspace: (1) normal and curvature estimation using OpenMP; (2) statistical processing and filtering; (3) moving least squares smoothing and upsampling (useful for DTM generation); and, (4) iterative closest point registration. These operations, and others within the libraries can be used interchangeably in the workflow to filter, smooth and transform the point clouds depending on output requirements. As a result, the workflow can be easily modified or extended to have uses beyond the creation of surfaces for numerical modelling.

For simulation of environmental flows, a watertight reconstruction of the surface is required, with minimal artificial perturbations or mesh artefacts such as non-manifold geometry. The main challenge with creating surfaces from the SfM-MVS process is handling non-uniform point density, noise and misaligned points; while being scalable to large point clouds. There are several surface reconstruction methods currently available and research in this area is active. In this workflow, the Smoothed Signed Distance Colored (SSD-C) surface reconstruction method of Calakli and Taubin [16] is used. This method was chosen for its speed and durability, however additional methods may be easily substituted if required.

2.4 Numerical Modelling

The reconstructed surface created from the previous steps has a number of potential applications, including visualisation and as an input for numerical modelling. Here, we demonstrate an application for high-resolution numerical modelling within a complicated urban environment using 3D smoothed particle hydrodynamics (SPH). SPH is a versatile, meshless, Lagrangian particle method that has been used to model environmental phenomena such as tsunamis, dam breaks, landslide initiation, lava flows and mudflows [2]. SPH is particularly advantageous for modelling complicated flows within complex environments due to the natural handling of complex topography and free surface flows, and the ability to include additional physics such as non-Newtonian rheology and entrainment of dynamic objects.

In the workflow, the reconstructed surface is converted into a data structure that can be read by the SPH software, which is run externally to the workflow due to the long computational time (in the order of weeks to months) required for high resolution simulations. Once completed, the data can be analysed and visualised within Workspace using separate data analysis workflows.

3 Case Study: Arequipa, Peru

To demonstrate the applicability of this method, we present a case study of the workflow for Arequipa City, Peru. Arequipa's central business district is located approximately 17 km south-west of El Misti, a steep potentially active volcano. The city is exposed to flash floods and lahars (volcanic debris flows) due to climatic conditions and the city's proximity to El Misti. Areas along the main quebradas (ravines) of San Lazaro and Huarangal are at risk from flash floods and lahars during the rainy season or after eruptions [17]. For example, flash floods in February 2011 destroyed 20 houses and caused damage to another 400, mostly along the quebradas and steep slopes. Much of this damage was attributed to heavy rainfall or breakout floods caused by infill obstacles such as makeshift crossings [17]. High resolution terrain models are required to represent the small scale of these features, while the rapidly changing environment requires frequent updating to accurately represent the topography for hazard simulations. This area was therefore chosen as a case study to demonstrate the value of the modelling workflow.

3.1 Image Collection

An approximately 200 m long and 200 m wide area of Quebrada San Lazaro was studied and is shown in figure 2a, looking upstream from a road bridge. The area in question contained the bridge and an upstream bend in the channel, which are thought to create overbank flow and strong 3D flow features that are ideal to test the use of SPH. In total, 2000 images of the area were acquired in September 2013; 600 images were taken from the ground and 1400 aerially. The aerial images were obtained from a DJI Phantom, a small ($35 \times 35 \times 19$ cm), lightweight (take off mass < 1 kg) quadcopter with an attached camera.

The objective of the image collection stage was to acquire as many images as possible from multiple angles to ensure all features were reproduced correctly and that there was minimal occlusion. A quadcopter is ideal in these circumstances, as its light weight and maneuverability means that it can be flown in close proximity to buildings to acquire images from multiple angles. In other areas (e.g. larger regions with less changes to terrain) a higher, more stable, but less maneuverable, aerial platform such as a helium balloon may be more suitable.

3.2 Feature Matching, SfM and MVS

Image features for the case study were detected and matched using SIFT feature descriptors. In an unordered image set, such as the one used in this case study, each image can potentially match each other image, meaning that image matching has $O(n^2)$ complexity. To reduce the time taken to match images, we parallelised this portion of the workflow, enabling matches to be determined on local and external CPU cores using the Workspace parallel execution engine. A speedup of 600% was achieved using parallel processing on 12 CPU cores (at 2.93 GHz), the speedup of the parallel implementation in this case is limited by the file I/O. Following image matching, the triangulation, bundle adjustment and multi-view stereo processes are run separately using VisualSfM [18], which utilises the CMVS-PMVS method of Furukawa and Ponce [14].

3.3 Mesh Generation

The SfM-MVS process created a dense point cloud containing ~8.5 million nodes. An image of the point cloud taken from a similar view angle to figure 2a is shown in figure 2b, where the point cloud model appears to correspond well with the real topography. The density of the raw point cloud varied due to occlusion and the number of images taken of a particular area. The maximum point cloud density was approximately 250 points in a 0.25 m^2 area. However, for this application point densities of this magnitude are much higher than required and the point cloud was reduced by 10% to have a maximum of 1 point within a 1 cm radius (7.6 million points over the entire case study area). The filtered point cloud was then reconstructed into a watertight, manifold mesh using SSD reconstruction, shown in figure 3. For this example, the entire point cloud processing workflow took 10 minutes; however the amount of required memory (~10Gb) may limit the scalability of SSD reconstruction to much larger areas.

Fig. 2. (a) Case study area in Quebrada San Lazaro and, (b) dense point cloud reconstruction from the SfM-MVS

Fig. 3. Watertight surface reconstruction of the case study area using SSD (looking downstream towards bridge)

3.4 SPH Simulation

Multiple SPH simulations examining the effect of flash flood flow rates on inundation, flow structures and velocities were run, using the terrain model generated from the previous steps. Computational time for the simulations increases by the cube of particle spacing and can therefore be reduced significantly through increasing particle spacing and simplification of the terrain. Terrain resolution and particle spacing was chosen here to be the largest resolution size of the smallest features that might affect the flow. Input flow rates of 25, 50 and $150 \text{ m}^3\text{s}^{-1}$ were considered, with a particle spacing of 12.5 cm. In the highest flow rate scenario, this required 2.8 million fluid particles to be simulated, which took 2 weeks of runtime to simulate 150 seconds of inundation on a 12 core, 2.93 GHz processor.

Figure 4 displays the velocities and inundation patterns 40 seconds after the flash flood was initiated for each flow rate. The fluid is shaded by velocity, with red being 10 ms^{-1} and blue being 0 ms^{-1} . The inundated area and velocities increase with flow rate, which is expected given the larger volume of water in the flow. The shape and magnitude of some flow structures are also different for the three flow rates. The effect of small features in the terrain is largely invisible in the highest flow rate scenario, with the flow mainly being guided by the bend in the channel. The effects of smaller features on the flow are noticeable for the 50 and $25 \text{ m}^3\text{s}^{-1}$ scenarios. For example, the small access road into the channel (circled in figure 4) causes hydraulic jumps to form in the lower flow rate scenarios, but the high flow rate scenario is largely unaffected by this feature. This demonstrates the necessity of detailed 3D representations of terrain and modelling, particularly when considering smaller, high frequency flooding scenarios.

Fig. 4. SPH simulations of a flash flood in Quebrada San Lazaro with flow rates of (a) 150, (b) 50 and, (c) 25 m³s⁻¹. White circle highlights small road seen to affect flow structures in lower flow rate simulations.

4 Conclusion

The workflow presented here describes a process for generating 3D terrain models by utilising advances in structure-from-motion, multi-view stereo, point cloud processing and surface generation. The process is inexpensive, rapid, captures fine scale features and accurately represents 3D structures within complex environments. The parallel processing of image matching ensures the workflow is scalable to a large number of images. By reducing the cost and acquisition time, terrain models can be generated at a higher frequency, better capturing transient small-scale features such as flow blockages and obstacles that can affect flow behaviour. Data processing is achieved through the use of a modular workflow engine, which allows for new methods to be implemented as research into particular areas progresses. The flexibility of the workflow engine, in addition to the usage of open source libraries and data structures provides opportunities beyond the use case demonstrated here, as the process and data generated may be used in various visualisation, terrain analysis and risk assessment fields of study.

The use and an application of this workflow was demonstrated using the case study area of Arequipa, Peru, where the resulting terrain model was used to predict the impacts of flash flood events using smoothed particle hydrodynamics. In addition to the presented numerical modelling application, the resulting terrain models can be used in a variety of applications such as visualisation, building classification and vulnerability estimation.

References

1. Manville, V., Cronin, S.J.: Breakout Lahar from New Zealand's Crater Lake. *Eos, Transactions American Geophysical Union* 88, 441–442 (2007)
2. Cleary, P.W., Prakash, M.: Discrete–element modelling and smoothed particle hydrodynamics: potential in the environmental sciences. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* 362, 2003–2030 (2004)
3. Williams, R.D., Brasington, J., Hicks, M., Measures, R., Rennie, C.D., Vericat, D.: Hydraulic validation of two-dimensional simulations of braided river flow with spatially continuous aDcp data. *Water Resources Research* 49, 5183–5205 (2013)
4. Legleiter, C.J., Kyriakidis, P.C., McDonald, R.R., Nelson, J.M.: Effects of uncertain topographic input data on two-dimensional flow modeling in a gravel-bed river. *Water Resources Research* 47, W03518, 3518 (2011)
5. Javernick, L., Brasington, J., Caruso, B.: Modeling the topography of shallow braided rivers using Structure-from-Motion photogrammetry. *Geomorphology* 213, 166–182 (2014)
6. Kreylos, O., Oskin, M., Cowgill, E., Gold, P., Elliott, A., Kellogg, L.: Point-based computing on scanned terrain with LidarViewer. *Geosphere* 9, 546–556 (2013)
7. Workspace (2014), <http://research.csiro.au/workspace>
8. Furukawa, Y., Curless, B., Seitz, S.M., Szeliski, R.: Towards Internet-scale multi-view stereo. In: 2010 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1434–1441 (2010)

9. Seitz, S.M., Curless, B., Diebel, J., Scharstein, D., Szeliski, R.: A Comparison and Evaluation of Multi-View Stereo Reconstruction Algorithms. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), pp. 519–528 (2006)
10. Snavely, N., Seitz, S.M., Szeliski, R.: Photo tourism: exploring photo collections in 3D. *ACM Trans. Graph.* 25, 835–846 (2006)
11. Lowe, D.G.: Object recognition from local scale-invariant features. In: The Proceedings of the Seventh IEEE International Conference on Computer Vision, vol. 1152, pp. 1150–1157 (1999)
12. Bay, H., Tuytelaars, T., Van Gool, L.: SURF: Speeded Up Robust Features. In: Leonardis, A., Bischof, H., Pinz, A. (eds.) *ECCV 2006, Part I. LNCS*, vol. 3951, pp. 404–417. Springer, Heidelberg (2006)
13. Leutenegger, S., Chli, M., Siegwart, R.Y.: BRISK: Binary Robust invariant scalable keypoints. In: 2011 IEEE International Conference on Computer Vision (ICCV), pp. 2548–2555 (2011)
14. Furukawa, Y., Ponce, J.: Accurate, Dense, and Robust Multiview Stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 32, 1362–1376 (2010)
15. Rusu, R.B., Cousins, S.: 3D is here: Point Cloud Library (PCL). In: 2011 IEEE International Conference on Robotics and Automation (ICRA), pp. 1–4 (2011)
16. Calakli, F., Taubin, G.: SSD-C: Smooth Signed Distance Colored Surface Reconstruction. In: Dill, J., Earnshaw, R., Kasik, D., Vince, J., Wong, P.C. (eds.) *Expanding the Frontiers of Visual Analytics and Visualization*, pp. 323–338. Springer, London (2012)
17. Thouret, J.-C., Enjolras, G., Martelli, K., Santoni, O., Luque, J., Nagata, M., Arguedas, A., Macedo, L.: Combining criteria for delineating lahar-and flash-flood-prone hazard and risk zones for the city of Arequipa, Peru. *Natural Hazards and Earth System Sciences* 13, 339–360 (2013)
18. Wu, C.: Towards Linear-Time Incremental Structure from Motion. In: 2013 International Conference on 3D Vision (3DV), pp. 127–134 (2013)