Project note

19 September 2016

To	Colin Mazengarb/Ted Rigby Mineral Resources Tasmania/Rienco
From	Bill Cohen **p** +61 3 6245 4513 **e** Bill.Cohen@entura.com.au
Project	South East Tasmania Tsunami Modelling
Project reference	E304976/P510568
Document number	ENTURA-C28B3
Subject	Modelling Summary

Introduction

The objective of the 2015-2016 South East Tasmania Tsunami project has been to take a Tsunami event from a Geoscience Australia database, with a collection of elevation data (bathymetry and land elevation) and roughness surfaces, and update the original Geoscience Australia tsunami modelling of the region. A critical component of this study is the ANUGA hydrodynamic modelling software, developed by Geoscience Australia and the Australian National University.

ANUGA is a very complex software application, is not widely adopted, and has little support and development budget. As such, many of the issues encountered with this project were related to ANUGA, generally for the following reasons:

- Lack of operator experience using the ANUGA software
- ANUGA documentation was not always up to date
- Bugs in the software that were detected during the project

These issues were mitigated as:

- Ted Rigby (Rienco) was made available to the project as a specialist consultant; his understanding of numerical computation and hydraulic modelling in general and ANUGA in particular have been of tremendous support in getting around issues as they were encountered
- Stephen Roberts (ANUGA development team) was very quick to respond to any bugs encountered with the software

The project team wishes to thank and acknowledge the roles that Ted and Stephen played in this project; without their timely support, modelling would not have been possible.

The general modelling issues are described below, as are any issues specific to certain scenarios.

Computing hardware

Entura used one physical workstation and rented several cloud hosted virtual instances (EC2: elastic cloud compute) from Amazon Web Services (AWS).

	Physical workstation	'Compute optimised' AWS	'General purpose' AWS
Label	HP Z400 Workstation	c4.8xlarge	m4.10xlarge
Processor	Intel Xeon 5645 Processor at 2.4 GHz	vCPU (High frequency Intel Xeon E5-2666 v3 (Haswell) processors optimized specifically for EC2)	vCPU (2.4 GHz Intel Xeon® E5-2676 v3 (Haswell) processors)
Memory	16 GB	60 GB	160 GB
Swap space	60 GB	0 GB	0 GB
Cost	Embedded	$2.195 per Hour (USD)	$3.363 per Hour (USD)

The 'general purpose' AWS machine was found to give the most reliable performance. No time was invested into working out how the AWS machines could enable swap space. This is possible, but was not trivial at the time of the project. One concern was that EBS (elastic block storage, the standard disk used by AWS EC2 instances) would charge for large amounts of read/write actions. Swap space has a large amount of read/write actions.

General modelling issues

Swap space

As stated in the install guide, ANUGA is memory-intensive. The desktop workstation used by Entura used the Ubuntu/Linux operating system with 16 GB memory. The initial disk configuration had very little swap space, however. Swap space is hard disk space set aside to be used as memory should the actual physical memory be insufficient. As such, ANUGA would fail before completing model runs. The error messages provided, however, did not point directly to this as an issue. It is possible that the cause of this was due to MPI (Message Passing Interface, the software that ANUGA uses for multi-processing), and not ANUGA itself. A typical message would be:

"mpirun noticed that process rank 0 with PID 3549 on node anuga-workstation exited on signal 9 (Killed)"

The SSD (solid state drive) used by the workstation was repartitioned with 60GB set aside for swap space. This addressed the issue and model runs could then be completed.

Rising hydrograph tail

Some early model runs had a rising hydrograph tail, which was out of character for a tsunami event. It was found that the ANUGA model was double-counting the default stage beyond the timeseries of tsunami boundary. This issue was corrected by setting the initial stage of the tsunami Dirichlet boundary to 0 m instead of the HAT value of 0.8 m.

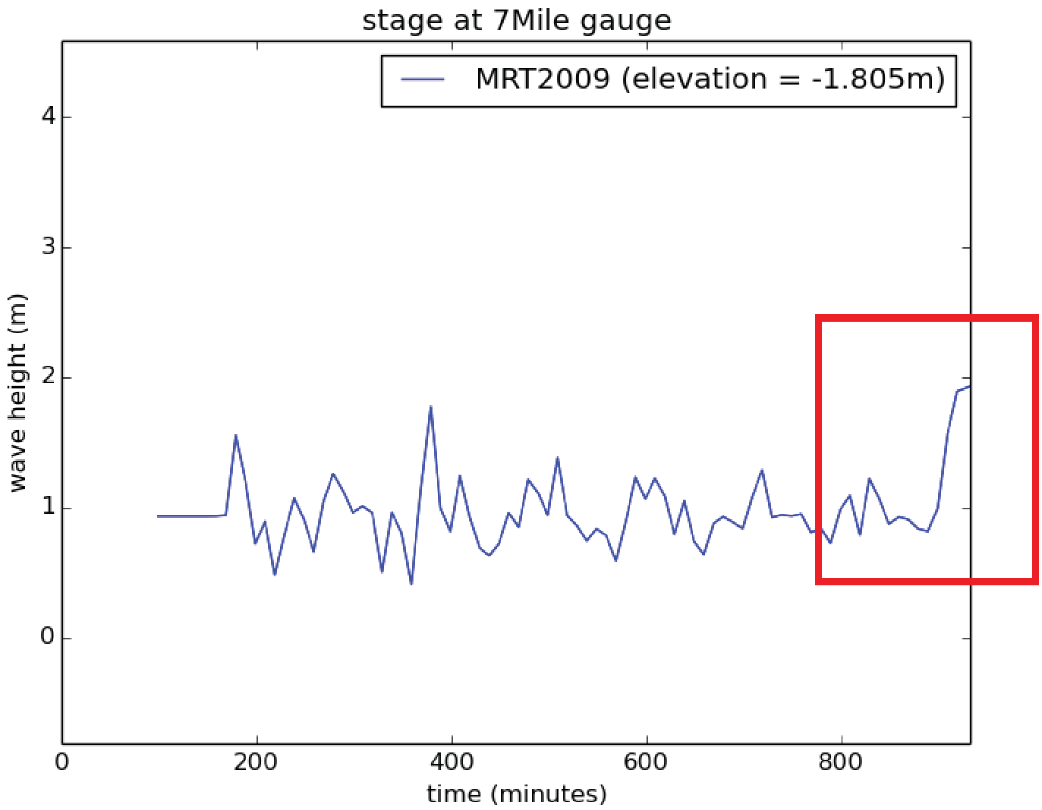

Resolution polygons

ANUGA is quite sensitive to the structure of shapefiles it will accept. Polygon shapefiles cannot have holes in the polygons, these must be infilled. This can be accomplished using GIS software such as ArcGIS or QGIS. Likewise, ANUGA has had difficulties loading polygons with coincident boundaries. Some editing of polygon extents has been required to allow ANUGA to use the polygon extent data.

Run-up issue

Some difficulties were encountered with the calculation of inundation depth in shallow water areas, as described by Ted Rigby (pers. Comms 20 July 2016). Please see Ted's report for more detail on this issue.

General modelling details

Variable mesh resolution was used throughout the model domain. The following resolutions were applied to regions. The extents of these resolution areas are given in the map below. Areas not shaded (ie white) are at the coarse resolution.

Label	Region applied	Mesh resolution	Triangle size
Extra fine	Blackmans Bay	50	10 m x 10 m
Very fine	Study areas	200	20 m x 40 m
Fine	Coastal zone, mean sea level +/- 10 m	1250	50 m x 50 m
Medium	Ocean approaches mean sea level +/- 30 m	20000	200 m x 200 m
Coarse	All other areas	80000	400 m x 400 m

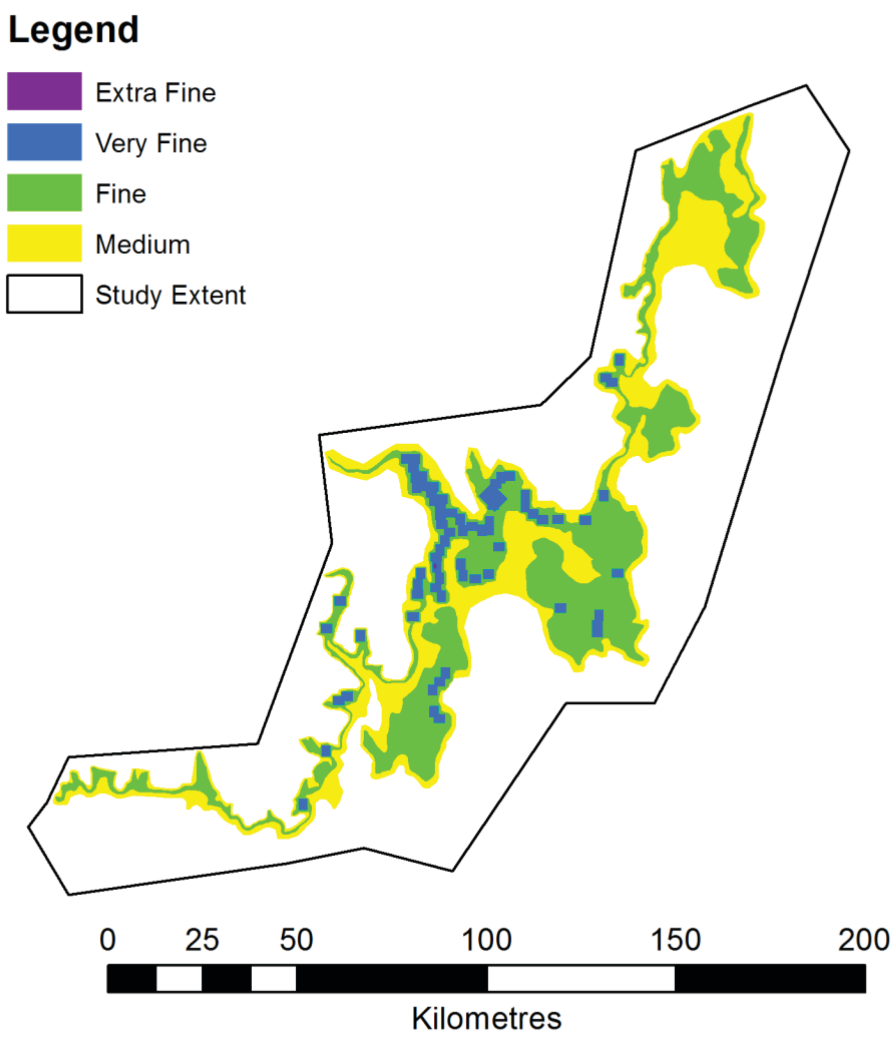

Scenario 1: reinstatement

The objective of scenario 1 has been to reinstate the original model done by Geoscience Australia in 2009 with new data. Potential changes that could have an impact on these results include changes to the ANUGA engine, changes to the elevation data, and the new model resolution extents.

Scenario 2: Core model

The core model has been run for a simulation period of 4 hours. This model ran successfully on an Entura desktop computer. This study used all resolution polygon extents except for the extra fine.

Scenario 3: Navigation Hazard

The Navigation Hazard model took a long time to process on the AWS machine (~80 hours). In addition, a considerable amount of memory was required for the SWW merge process to work. A 'compute optimized' CPU intensive machine with 60 GB of memory failed the final SWW merge step, but a general purpose machine with 160 GB memory was successful. Allocating swap space on the AWS instances was not investigated, but it could prove useful. This study used all resolution polygon extents except for the extra fine.

Timeseries and summary data have been extracted for several series of gauges for the Australian Maritime College. These are given below.

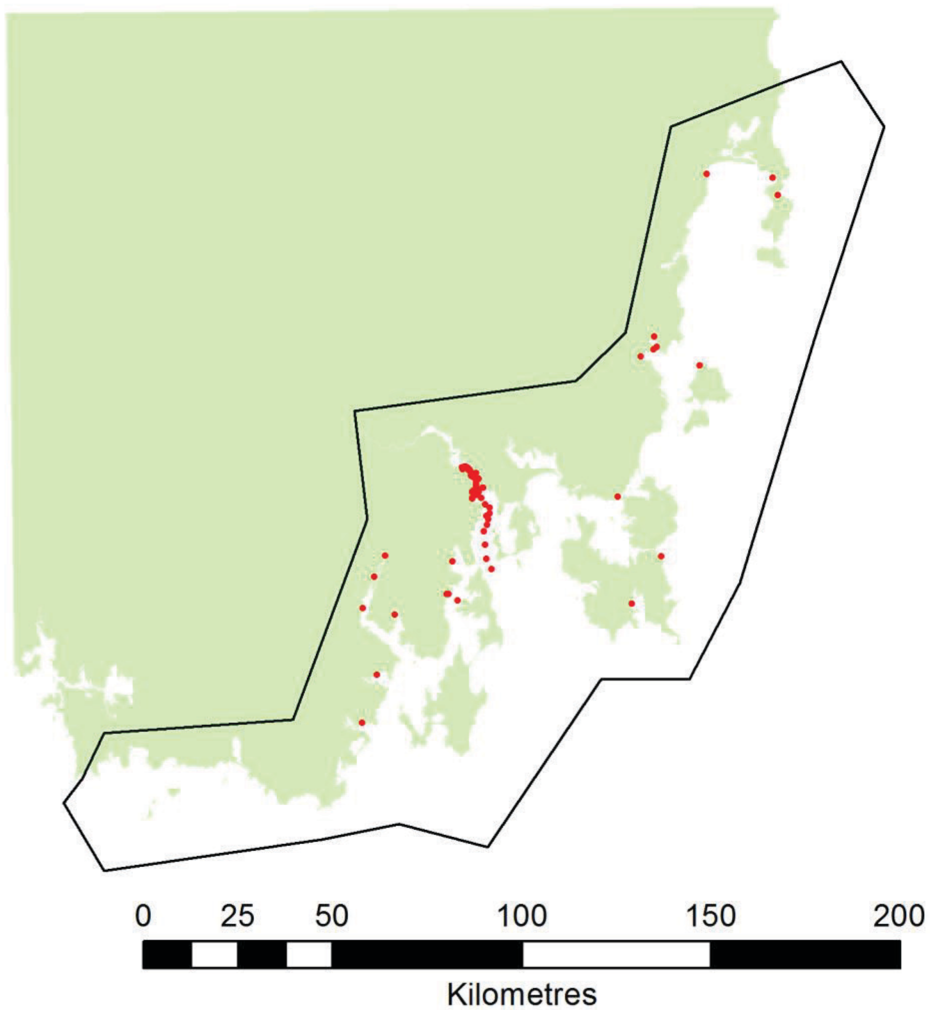

Locations for AMC Gauges 1

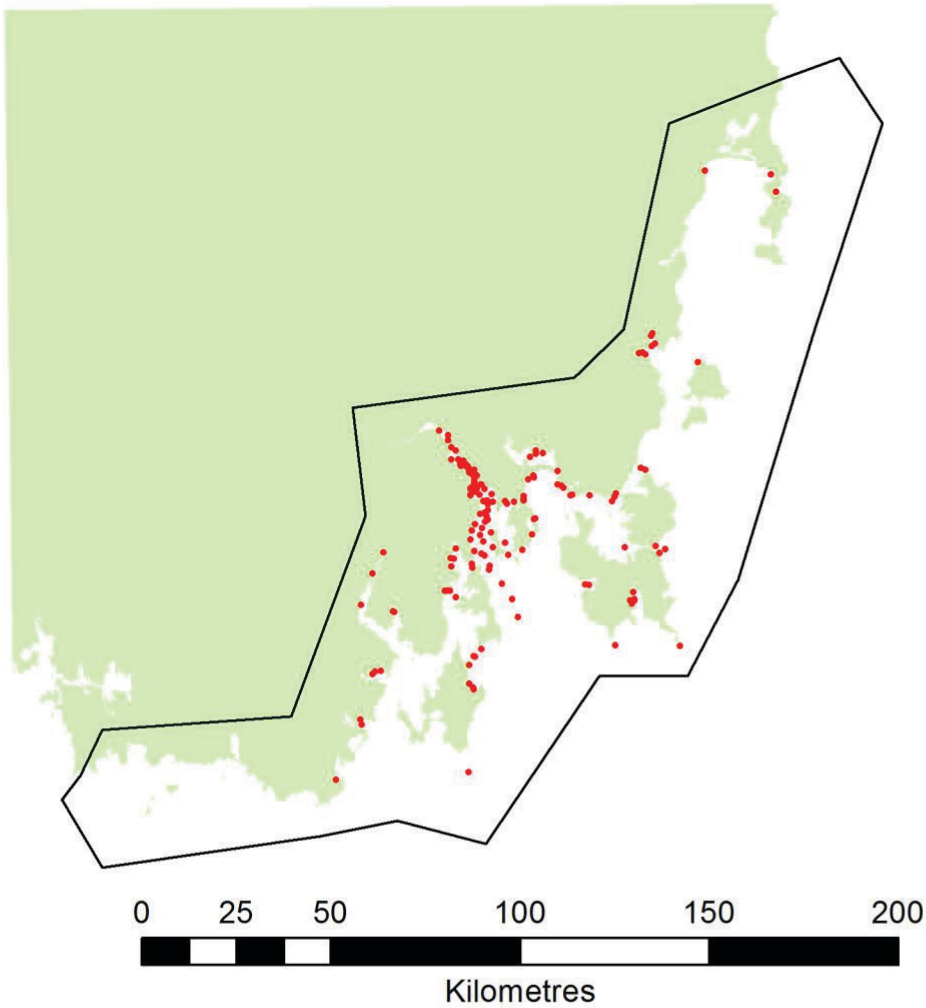

Locations for AMC Gauges 2

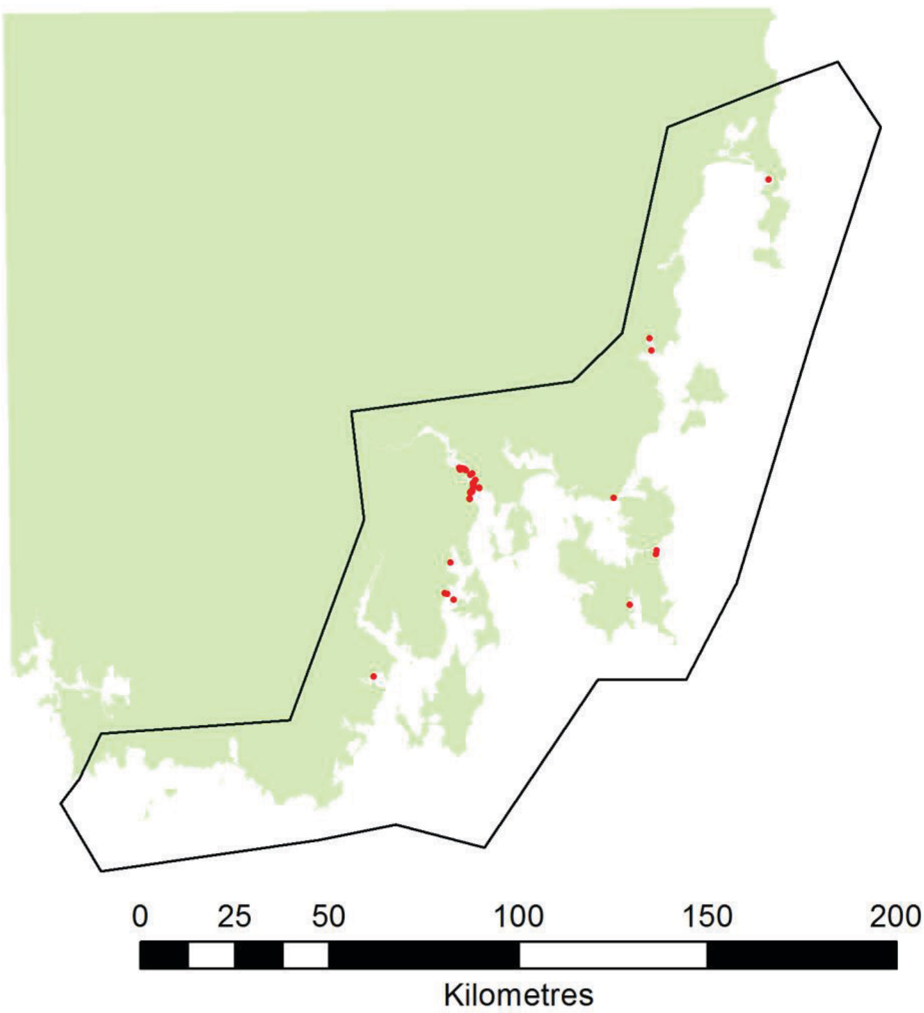

Locations for AMC Gauges 3; these locations are from AMC Gauges 2 that have been moved further offshore due to the ANUGA run-up issue

Scenario 4: Coastal Hazard

The coastal hazard model had a much finer spatial resolution in key analysis locations throughout the model, and extra fine spatial resolution in the Blackman's Bay area. This model was run for a simulation period of 4 hours, which took approximately 50 hours of runtime on an AWS instance, and was expected to take approximately 2 weeks on a desktop computer. This study used all resolution polygon extents.

General recommendations

Generally, it is recommended to keep track of all model configuration files in a version control system, such as git, and use this to track any and all changes to model runs. This allows the tracking of changes to the model and to revert back to previous configurations if required.

The use of cloud services (Amazon Web Services, AWS) has been very beneficial in this study; without it, some model runs would have taken weeks to complete.

Another major challenge of this project has been to handle the volume of data. Not just the size of the datasets, but the number of model configuration options and the number of outputs for each model run. A flexible framework/document for recording all of these details that can work well with the git version control system is recommended. In addition, some restructuring of the Python files could enable greater flexibility and tracking of changes.

References

Entura 2016, ANUGA: Guide to ANUGA on AWS, ENTURA-B9D6D, 25 May 2016