GPGPU Acceleration of Environmental and Movement Datasets

Daniel Bird
Daniel.Bird@uea.ac.uk
University Of East Anglia
Norwich

Stephen Laycock
S.Laycock@uea.ac.uk
University Of East Anglia
Norwich

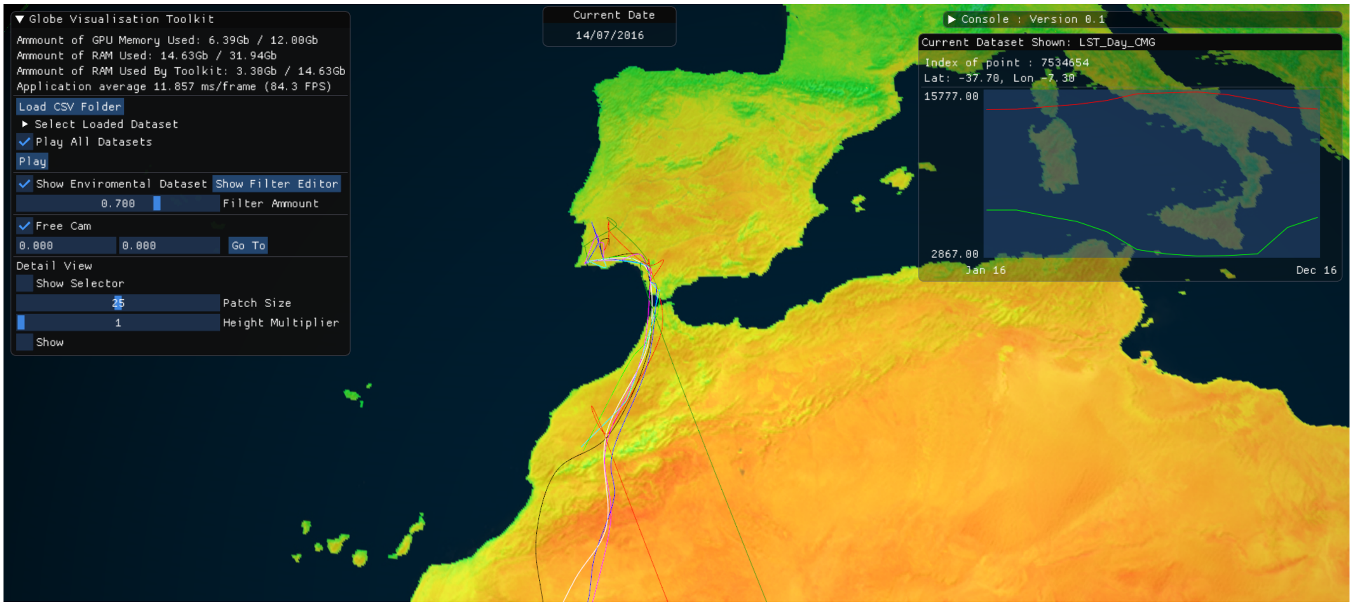

Figure 1: Visualisation of the MODIS Land Surface Temperature dataset with the migration routes of multiple white storks.

CCS CONCEPTS

• **Human-centered computing** → **Visualization toolkits**; • **Applied computing** → Earth and atmospheric sciences.

KEYWORDS

Movement Ecology, GPGPU, Visualisation Toolkit

ACM Reference Format:
Daniel Bird and Stephen Laycock. 2019. GPGPU Acceleration of Environmental and Movement Datasets. In *Proceedings of SIGGRAPH '19 Posters*. ACM, New York, NY, USA, 2 pages. https://doi.org/10.1145/3306214.3338584

1 INTRODUCTION

Due to the increased availability and accuracy of GPS sensors, the field of movement ecology has been able to benefit from larger datasets of movement data. As miniaturisation and the efficiency of electronic components have improved, additional sensors have been coupled with GPS tracking to enable features related to the animal's state at a given position to be recorded. This capability is especially relevant to understand how environmental conditions may affect movement.

Surveys that have been conducted in the field of visualisation for movement ecology have highlighted that tool-sets investigating movement within context are important [Andrienko and Andrienko 2013; Demšar et al. 2015]. Investigating movement within context allows for possible inference on why and where movement takes place [Nathan et al. 2008]. Due to the exceptionally large size of global datasets, visualisations are often calculated for a local area, which is viable for studies that only remain within a local area, but for large scale movement, such as migratory movement, it is more difficult. For example, [Gilbert et al. 2016] studied the migratory pattern of white storks, that migrate from Portugal to Morocco, to observe the effects of environmental factors, such as vegetation and land use.

2 OUR APPROACH

In order for global datasets, such as the MODIS/Terra Land Surface Temperature and Emissivity dataset containing over 25 million points, to be visualised, our approach applies GPGPU acceleration to the colour calculation of these data-points, allowing for a colour map to be applied to the entirety of the dataset at interactive speeds. This form of visualisation lends itself exceptionally well to the parallel processing nature of GPGPUs, with every data-point being processed independently. The result of this can then be displayed

SIGGRAPH '19 Posters, July 28 - August 01, 2019, Los Angeles, CA, USA

ACM ISBN 978-1-4503-6314-3/19/07.
https://doi.org/10.1145/3306214.3338584

upon a fully rendered 3D globe. This allows for multiple simultaneous studies, or studies that encompass a large area or distance, such as migration, to be visualised.

2.1 Global View

First, a given number of environmental datasets are loaded into the dataset series with its corresponding date. Then, whenever the visualisation reaches a point where an update to the colour map is needed, the dataset details, such as resolution and data type, are read from the dataset series. Here, if the resolution of the datasets have changed between time stamps the memory allocated to the output texture is dynamically changed to match the resolution of the loaded dataset, allowing for minimal use of GPU resources.

Once the globe is at the correct resolution the colour map kernel is called. The datasets assigned to the given time-stamp, or closest if there is no dataset assigned to the time-stamp, is then loaded onto the GPU. This approach allows for multiple datasets to be uploaded onto the often larger capacity system memory, only loading the dataset onto the often smaller graphics memory when needed. This facilitates the use of lower end GPUs that may not have much memory on the condition that it can store one month of data. A CUDA kernel is then used to apply this colour map to every point within the dataset. To reduce the amount of GPU memory used in storing the results, the colour mapping is saved to texture memory. This allows for the resolution of the globe itself to be far lower than the resolution of the dataset currently shown, increasing performance. This use of texture memory allows for interpolation of the dataset to utilise the texture filtering capability of the GPU. This also allows the user to switch between nearest neighbour or bi-linear interpolation for visualisation of the colour map.

Currently, calculation of a 7200 by 3600 point colour map requires 2.59 milliseconds using a NVIDIA Titan Xp. While this card has a 12GB memory capacity, the calculation only requires 1.39GBs plus the size of one time stamp of the dataset of GPU memory, this allows for systems with lower specifications still being able to create real-time visualisations. Using this approach allows for significantly reduced calculation times, allowing for the rapid switching of these datasets, enabling real-time visualisation of large datasets over time.

To facilitate the exploration of loaded datasets, a ray may be cast from the camera location to the globe with the current colour map displayed. The index of the selected data-point is then calculated from the point that collides with the ray. A year of data at the selected index is then displayed. This allows for the consideration of not only the current value of every loaded dataset to be taken into account for the current position, but also changes in the environmental context.

2.2 Detail View

Due to the exceptionally large size of some datasets, such as the 233 million point ETOPO1 Global Relief Model, it would be unrealistic to apply the same method and achieve a real-time visualisation. To allow for this, the detail view is used. This takes a small section of the currently displayed dataset and uses the texture output from the previous step to map the higher resolution geometry to the colour map. The detail patch is generated using a CUDA kernel that recalculates the geometry using the higher resolution digital elevation model, allowing for the use of exceptionally large environmental datasets, such as height data.

These methods, when paired with the ability of to load and playback tag data, allow for both the investigation and the visualisation of the data, while simultaneously visualising the underlying environmental conditions on a global scale. This allows for multiple or large-scale studies across the world to be investigated with environmental context.

ACKNOWLEDGMENTS

We gratefully acknowledge the support of NVIDIA Corporation with the donation of the Titan Xp GPU used for this research. We would also like to thank Dr Aldina Franco for supplying the GPS data for the White Storks [Gilbert et al. 2016] and the discussions on the research.

REFERENCES

Natalia Andrienko and Gennady Andrienko. 2013. Visual analytics of movement: An overview of methods, tools and procedures. , 3–24 pages. https://doi.org/10.1177/1473871612457601

Urška Demšar, Kevin Buchin, Francesca Cagnacci, Kamran Safi, Bettina Speckmann, Nico de Weghe, Daniel Weiskopf, and Robert Weibel. 2015. Analysis and visualisation of movement: an interdisciplinary review. *Movement Ecology* 3, 1 (mar 2015), 5. https://doi.org/10.1186/s40462-015-0032-y

Nathalie I. Gilbert, Ricardo A. Correia, João Paulo Silva, Carlos Pacheco, Inês Catry, Philip W. Atkinson, Jenny A. Gill, and Aldina M. A. Franco. 2016. Are white storks addicted to junk food? Impacts of landfill use on the movement and behaviour of resident white storks (Ciconia ciconia) from a partially migratory population. *Movement Ecology* 4, 1 (dec 2016), 7. https://doi.org/10.1186/s40462-016-0070-0

Ran Nathan, Wayne M Getz, Eloy Revilla, Marcel Holyoak, Ronen Kadmon, David Saltz, and Peter E Smouse. 2008. A movement ecology paradigm for unifying organismal movement research. *Proceedings of the National Academy of Sciences of the United States of America* 105, 49 (dec 2008), 19052–9. https://doi.org/10.1073/pnas.0800375105