Computers and Geosciences 119 (2018) 39–48

Contents lists available at ScienceDirect

Computers and Geosciences

journal homepage: www.elsevier.com/locate/cageo

Research paper

Near real-time coastal flood inundation simulation with uncertainty analysis and GPU acceleration in a web environment

Janne Kovanen[*,1], Juha Oksanen[2], Tapani Sarjakoski[3]

Finnish Geospatial Research Institute, National Land Survey of Finland, P.O. Box 84, FI-00521, HELSINKI, Finland

A R T I C L E I N F O

Keywords:
Inundation forecasting
Flood risk management
Spatial analysis
Digital elevation model
Graphics processing unit
Uncertainty

A B S T R A C T

A proof of concept is presented on how to produce uncertainty-aware near real-time coastal flood inundation Web maps from water-level observations and predictions, which have been computed for tide gauge sites and made publicly accessible. The stochastic inundation simulation takes into account several sources of uncertainty, which have until now not been employed in either bathtub models or hydrodynamic models. The simulation is based on the Monte Carlo method. The feasibility of the proposed approach is demonstrated by an implementation using general-purpose computing on graphics processing units. The outcome of the research is that the current technologies enable the building of a novel system that connects to official data sources and that takes into account sources of uncertainty whose inclusion in the past has been avoided by being either weakly known or computationally too expensive.

1. Introduction

In the European Union, the Floods Directive (The European Parliament and the Council of the European Union, 2007b) is put into effect, which requires the member states to assess the flood risk in their coastal zones in the form of flood hazard and risk maps. The maps are static in nature and reflect the likelihood of various risk levels. A flood hazard map visualizes at least the extent and depth of the flood (Fig. 1), while a risk map incorporates the number of inhabitants, type of economic activities, and other critical objects and infrastructure of the particular scenario. Still, one of the major deficiencies in flood mapping is the assumption of error-free data.

At the same time as flood risk assessment has gained attention (de Moel et al., 2009), the open data movement has gained popularity and technological advances made it possible to share the data. Consequently, some national agencies measuring and forecasting sea levels are making their data public; for example, the Finnish Meteorological Institute (FMI) opened its first digital data sets for the general public in 2013. The data is published via standardized web interfaces and data schemas. This allows anyone to derive new services and data products from it.

In this paper, we go beyond state-of-the-art of projected flood hazard maps by presenting a proof of concept on how to produce near real-time coastal flood inundation maps that visualize the impact of uncertainties in a number of input data sources. As the source data for computation, we use water-level forecasts for tide gauges and a digital elevation model, both of which are loaded from publicly-accessible web interfaces. Using a Monte Carlo (MC) -based method, we take into account several sources of uncertainty, some of which are recognized now for the first time in the inundation context. Finally, the outcome is published using open interfaces and an interactive web map.

The organization of the paper is the following. This Section continues with a short introduction to inundation models and describes what methods have earlier been used to handle uncertainty. In Section 2, we describe the used model. In Section 3, we describe the data, how its uncertainty is considered, and the related existing web interfaces. In Section 4, we present the system architecture and explain the choices made to speed-up the modelling, like partitioning the data and the use of graphics processing units for general-purpose computing. In Section 5, we describe our test area and hardware used to validate the feasibility of the concept. Finally, in Section 6, we discuss the limitations and further potential of the presented system.

1.1. Inundation models

Coastal flood hazard maps are made mainly with hydrodynamic

* Corresponding author.
E-mail addresses: janne.kovanen@iki.fi (J. Kovanen), juha.oksanen@nls.fi (J. Oksanen), tapani.sarjakoski@nls.fi (T. Sarjakoski).

[1] Kovanen contributed in the development of the idea, carried out the implementation and wrote the manuscript.
[2] Oksanen conceived the original idea, contributed to its development, provided the source code basis for process convolution and critically revised the manuscript.
[3] Sarjakoski contributed in the development of the idea and supervised the work.

https://doi.org/10.1016/j.cageo.2018.06.007
Received 15 April 2016; Received in revised form 1 June 2018; Accepted 14 June 2018
Available online 19 June 2018
0098-3004/

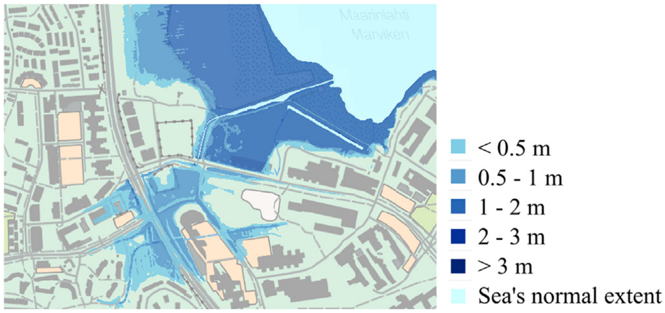

Fig. 1. A view of the Flood Map Service of the Finnish Environment Institute. The water depth is non-linearly colour-coded with shades of blue. (For interpretation of the references to colour in this figure legend, the reader is referred to the Web version of this article.)

models. A technical review on inundation models is given, for example, by Néelz and Pender (2009) and Di Baldassarre (2012). The models can be categorized according to their dimension. In case of coastal floods, computationally cheap 1D models are rarely used. Among other things, they would have to be applied to beach profiles or form a network of channels through which the inundation advances inland. On the other hand, 3D models have a high computational cost. For them to be feasible, a compromise needs to be made between the cell size, domain size and complexity of the closure scheme (Woodhead et al., 2007). Less accurate 2D models are more practical in producing water levels and flow velocities that vary in the horizontal directions. However, even 2D models with a reasonable resolution may be impractical from the runtime viewpoint to be applied over extensive areas or in iterative analyses that do not apply parallel computing techniques (Néelz and Pender, 2010).

Besides the aforementioned categories, inundation may be determined by hybrid models and a branch of techniques that is unbound by fluid hydraulics. The latter include bathtub methods, which are also referred to as flat water, single-value surface and equilibrium methods. They compare a projected water-level surface with the terrain's elevation. Inundation occurs when the water level exceeds the terrain's elevation. In tandem, the elevation difference can be captured in a flood depth map. The water level may be based on a given return period or sea level rise scenario (Coveney and Fotheringham, 2011; Gesch, 2009; Poulter and Halpin, 2008), but it may be as well interpolated from tide gauge data or a 1D flood model applied to cross sections perpendicular to a river's flow (Werner, 2001).

The feasibility of bathtub models has been demonstrated by Bates and De Roo (2000), who fitted a water-level surface to recorded tide gauge levels. The model was placed side by side with the solutions of two 2D models. The comparison was based on river flood extents, which were interpreted from aerial and SAR imagery. Bates and De Roo came to the conclusion that, at least in their study, even if the 2D inundation models outperformed the bathtub model, the latter performed well.

1.2. Uncertainty-aware modelling

Coastal floods are invariably uncertain as a consequence of the uncertainties related to the projections of the climate, the complex nature of coastal processes, and DEMs (Cowell and Zeng, 2003). This also applies to bathtub models, where the errors of input data together with all uncertainties in the model, in general, propagate to the inundation maps. While deterministic models draw a crisp extent boundary, probabilistic models define the extent as a function of probability or for certain confidence levels. In the probabilistic models, methods range from the use of Artificial Neural Networks (Shrestha et al., 2009) to the use of fuzzy sets where membership vectors define the possibility of an element being inundated (Pappenberger et al., 2007). However, parameter uncertainty has been mostly quantified with an informally Bayesian method – the Generalized Likelihood Uncertainty Estimation (GLUE) (Beven and Binley, 2014). A likely reason for GLUE's success is its conceptual simplicity and applicability to nonlinear systems (Stedinger et al., 2008) but as embarrassingly parallel it also suits distributed systems (Vrugt et al., 2009).

Relating to stochastic bathtub models, Gesch (2009) adds to the DEM a coefficient called linear error at the 95% confidence level, which describes the vertical accuracy of the DEM. The coefficient is defined by 1.96 × root-mean-square error (RMSE), as the data is assumed to be normally distributed and unbiased. A successor of the model, published by Schmid et al. (2014), also takes into account the geographically varying uncertainty to convert the water levels to a geodetic height system, which in their model can be assessed to be within 5–23 cm (NOAA, 2015). Unlike the other bathtub models, the approach presented by Zerger et al. (2002) creates, in Monte Carlo-style, perturbed instances of a DEM. The random noise of the model is based on the RMSE of the data used to compute the DEM.

2. The stochastic bathtub model

In our stochastic bathtub model, we take into account the uncertainties of DEM and tide gauge predictions. The stochastic model is implemented in the MC methods fashion where every iteration produces equally probable instances of terrain (Fig. 2) and water-level (Fig. 3). At the end of each MC run, a cell-wise comparison is made between the two (Fig. 4). The final probability after n runs is computed with the following equation

$$P = \frac{1}{2n}\sum_{i=1}^{n}(sgn(W_i - D_i) + 1), \quad sgn(x) = \begin{cases} 1 & \text{if } x > 0 \\ -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \end{cases} \tag{1}$$

where W_i and D_i correspond to water surface and terrain elevation instances, respectively. In the process of computing the inundation based on observations, in order to correct earlier forecasts, we only consider the uncertainty of the terrain elevation.

In the end, we enforce hydrological connectivity by removing inundated cells not connected to the sea directly or via other cells. We apply eight-connectivity, as was done in earlier bathtub models (Gesch, 2009; Poulter and Halpin, 2008). This raises the significance of topographic features, like embankments, and helps to avoid erroneous conclusions. Once connectivity to the sea exists, the surface will still contain those rivers, lakes, and reservoirs that have been valid pathways for the flood to proceed inland. We remove them by masking in a manner similar to that of Rowley et al. (2007) and Gesch (2009).

2.1. Terrain elevation instance

According to the First Law of Geography (Tobler, 1970) geographically proximal objects tend to have correlated values of an attribute. This phenomenon, termed ”spatial autocorrelation”, also applies to the errors related to the values. In their study, Zerger et al. (2002) used a spatially autoregressive process presented by Hunter and Goodchild (1997) to estimate the spatial autocorrelation of the DEM. In this study, however, we use a technique known as process convolution (Higdon, 1998; Thiébaux and Pedder, 1987) to model the spatial dependence between DEM errors (Oksanen and Sarjakoski, 2005) because in our case the correlogram model is known and the technique is computationally fast with small convolution kernels. We use the correlogram $\rho(h)$, computed from the practical range r_p, to create an isotropic smoothing kernel $K(h)$ with radius r_{kernel}

$$\rho(h) = e^{-\frac{r}{a}} \tag{2}$$

$$a \approx \frac{r_p}{3} \tag{3}$$