



Performance of Visual Computing Algorithms

Valentin Bruder

Visual computing is a field of computer science that deals with handling images and 3D models. This includes image generation as used in movies, computer games, or VR/AR, but also image processing and computer vision that can be used for autonomous driving. Due to the nature of image data to be relatively large in terms of memory usage, especially for data changing over time, performance is crucial in the field of visual computing. Our research is focused on assessing the performance of different algorithms, model it and ultimately predict performance requirements.

Knowing the Performance

Who doesn't know it, you are streaming video or playing a video game and suddenly the image freezes for a few seconds. There are many different reasons for this to happen: an unstable network connection, someone else straining the network, a problem on the provider side, or insufficient resources on your device. Some streaming services adapt the image quality based on the available resources in order to provide an interruption-free playback. But streaming videos is a comparatively linear application in terms of performance requirements. There is a known amount of data that is passed from a server to the client. If you consider an interactive application like a video game or a visualization, knowing the performance beforehand has much more impact. Correct calculations lead to stable behavior of the application.

Challenges

One approach to avoid freezes or lags in interactive visual computing applications is to lower the overall rendering quality when needed. However, this approach has several drawbacks. It is important to have a good estimate on how much compute power you will need in maximum. While a too optimistic setting can still lead to lags, a too conservative and inflexible adjustment leaves resources unused for computationally cheaper scenes or images and unnecessarily leads to a lower quality of the renderings.

Algorithms in the field of Visual Computing pose different challenges due to the nature of the image data that is being processed or rendered. The subfield of scientific visualization for example is concerned with efficiently rendering simulation data from different fields such as physics, biology or metrology (fig. 2 and 3). It is not uncommon for such simulation




Fig. 2: Simulation of a Supernova. Performance modeling and prediction made a fluent interaction with a high quality application possible.
Data source: John M. Blondin; rendering: Valentin Bruder.

data to be multi-dimensional (3D space, time, and multiple physical parameters such as energy, density etc.) and of a high resolution resulting in a huge amount of data.

Processing all this data in real time for an exploratory analysis is a big challenge. In addition, there are many influencing factors for an interactive applications such as viewing angle, quality parameters, screen resolution, and others that impact on the performance. Also, modern computing devices often have multiple different parallel hardware hierarchies which need to be taken into account for optimal exploitation of the available resources. The heterogeneity of those devices is a challenge, too. A smartphone has substantially different properties and requirements compared to a high performance computer cluster containing thousands of processing cores.

What is our research for?

There are many application scenarios where you can greatly benefit from performance modeling and prediction. We address those in our research.

One scenario is infrastructure planning. We developed a machine learning approach to predict the performance of a new or upgraded cluster. This helps in saving money and resources.

Furthermore, we developed techniques to improve the image quality during runtime based on available resources. We predict the cost of the upcoming frame and use that information to adjust

»Knowing the performance of a Visual Computing algorithm on specific hardware beforehand is an efficient way to save resources«




Fig. 3: Inside a Supernova. A huge amount of data had to be processed for this interactive visualization.
Data source: John M. Blondin; rendering: Valentin Bruder.

the quality of the next image. This guarantees fluent interaction while showing the best quality given the constraints. Similarly, we use the prediction to perform load balancing between for example two graphic cards in a desktop workstation.

The term performance does not only include runtime performance, i.e., how long do we need to calculate an image, but also memory and energy consumptions. We investigated this in the context of rendering on mobile devices where it is highly relevant.

Summary

Ever growing data sizes pose challenges to efficiently compute Visual Computing algorithms. Quantifying the performance of those algorithms is one way to address several of these challenges and optimize execution on different devices.

About the Author

Valentin Bruder is Doctoral Researcher in Project A02 Quantifying Visual Computing Systems. He does research on volume rendering and hardware performance.
valentin.bruder@visus.uni-stuttgart.de

