Available online at www.sciencedirect.com

ScienceDirect

Procedia Computer Science 4 (2011) 56–65

Procedia Computer Science

International Conference on Computational Science, ICCS 2011

Towards High-Dimensional Computational Steering of Precomputed Simulation Data using Sparse Grids

Daniel Butnaru*, Dirk Pflüger*, Hans-Joachim Bungartz*

Institut für Informatik, Technische Universität München,
Boltzmannstr. 3, 85748 Garching, Germany

Abstract

With the ever-increasing complexity, accuracy, dimensionality, and size of simulations, a step in the direction of data-intensive scientific discovery becomes necessary. Parameter-dependent simulations are an example of such a data-intensive tasks: The researcher, who is interested in the dependency of the simulation's result on a set of input parameters, changes essential parameters and wants to immediately see the effect of the changes in a visual environment. In this scenario, an interactive exploration is not possible due to the long execution time needed by even a single simulation corresponding to one parameter combination and the overall large number of parameter combinations which could be of interest.

In this paper, we present a method for computational steering with pre-computed data as a particular form of visual scientific exploration. We consider a parametrized simulation as a multi-variate function in several parameters. Using the technique of sparse grids, this makes it possible to sample and compress potentially high-dimensional parameter spaces and to efficiently deliver a combination of simulated and precomputed data to the steering process, thus enabling the user to interactively explore high-dimensional simulation results.

Keywords: Computational Steering, CFD Simulations, Sparse Grids, High Dimensionalities

1. Introduction

In the last decades the computational branch of science has made significant progress in both modeling and in performing accurate simulations of very complex phenomena. A consequence of such a step is the availability of large amounts of data generated by various simulations, and the focus now moves on to how to manage such data and explore it in a convenient way for the researcher. Another unfortunate characteristic inherent to simulations is that a higher accuracy (resolution) of the simulation demands a higher computational effort and thus significantly slows down the exploration process.

In this paper we consider (visual) data exploration due to parameter variation, see Fig. 1. Inside a computational steering environment, a researcher observes the effects of changes of the simulation's main parameters to the simulation results. The goal is, for example, to identify correlated as well as unimportant parameters, or to discover new and

*Corresponding author

Email addresses: `butnaru@in.tum.de` (Daniel Butnaru), `pflueged@in.tum.de` (Dirk Pflüger), `bungartz@in.tum.de` (Hans-Joachim Bungartz)

Selection and/or peer-review under responsibility of Prof. Mitsuhisa Sato and Prof. Satoshi Matsuoka doi:10.1016/j.procs.2011.04.007

unexpected patterns in the data. Due to storage costs and computational effort, it is however not possible to generate and store simulation results for every parameter combination of interest, especially in high-dimensional settings. Just consider that taking only ten distinct values for each of five parameters into account would require to compute and store 10^5 different simulation results. In order to deal with this problem, we introduce a method of sampling and compressing the simulation's parameter space based on sparse grids [1]. By reducing the size of the data, fast exploration of certain multi-dimensional data sets is enhanced, while sufficiently good accuracy of the stored simulation results is preserved.

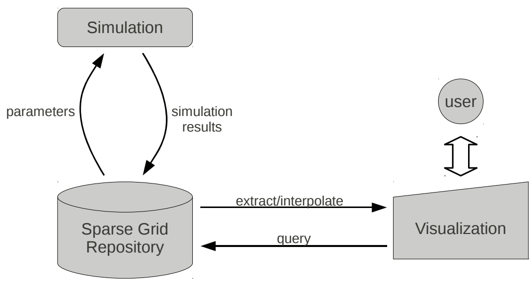

Figure 1: Workflow for computational steering for high-dimensional simulations.

The main idea is to consider the simulation as a function of various simulation parameters. Such a multi-dimensional function can be numerically represented and treated; however, classical discretizations of the parameter space suffer with increasing dimensionality from the so-called curse of dimensionality, the exponential dependency of the effort on the number of dimensions. Sparse grids enable us to mitigate the curse of dimensionality to some extent, allowing to tackle dimensionalities that are of interest in engineering settings, where models depend on a moderate number of variables. Instead of running a simulation for every parameter combination—an unrealistic task in itself—the sparse grid sampling dictates which parameter combinations actually need to be examined and stored. The associated sparse grid interpolation scheme then offers access to all other parameter combinations.

To demonstrate our method, we show results from a simple computational fluid dynamics (CFD) simulation, the so-called driven cavity. In this scenario, the fluid (e.g. water) in a cavity is stimulated by the cavity's moving lid. Location and shape of the emerging vortices mainly depend on the velocity of the lid and the viscosity of the fluid. We assume these two parameters to be continuous within a certain range, and together with time they form a three-dimensional parameter space. We then use a sparse grid to sample and approximate the flow in the considered parameter range, without the need to run new simulations for every combination in between. This enables an efficient real-time visualization of simulation results, well-suited for interactive computational steering.

2. The Sparse Grid Technique

Sparse grids help to overcome the curse of dimensionality to a great extent. Interpolating on a regular grid with a resolution of N grid points in one dimension, they enable one to reduce the number of grid points significantly in d dimensions from $O(N^d)$ to $O(N(\log N)^{d-1})$ while maintaining a similar accuracy as in the full grid case. The only requirement is that the underlying function f has to be sufficiently smooth [1]. Note that it has been shown that even functions that do not meet the smoothness requirements can be successfully dealt with if adaptive refinement is employed [2]; we will address adaptive refinement in Sec. 3.3. The notion sparse grids was coined in 1990 [3] for the solution of high-dimensional partial differential equations, and they have meanwhile been successfully employed in a whole range of applications, ranging from astrophysics and quantum chemistry to data mining and computational finance, see, e.g., [1, 2] and the references cited therein. In the following, we briefly describe sparse grids and the main principles they base upon, a hierarchical representation of the one-dimensional basis and the extension to the d-dimensional setting via a tensor product approach; for further details, see [1, 2] again.

We consider the representation of a piecewise d-linear function $f_N : \Omega \to \Gamma$ for a certain mesh-width $h_n := 2^{-n}$ with some discretization level n. The function $f_N(\underline{x})$ thus maps a set of parameters $\underline{x}$ out of the parameter space Ω to a simulation result Γ. For the parameter space Ω we consider rectangular domains which we scale to $\Omega := [0, 1]^d$.