

# Interactive Data Exploration for High-Performance Fluid Flow Computations Through Porous Media

Nevena Perović, Jérôme Frisch, Ralf-Peter Mundani, and Ernst Rank  
 Chair for Computation in Engineering  
 Technische Universität München, Germany  
 Email: nevena.perovic@tum.de

**Abstract**—Huge data advent in high-performance computing (HPC) applications such as fluid flow simulations usually hinders the interactive processing and exploration of simulation results. Such an interactive data exploration not only allows scientist to ‘play’ with their data but also to visualise huge (distributed) data sets in both an efficient and easy way. Therefore, we propose an HPC data exploration service based on a sliding window concept, that enables researches to access remote data (available on a supercomputer or cluster) during simulation runtime without exceeding any bandwidth limitations between the HPC back-end and the user front-end.

**Keywords**—high-performance computing, interactive data exploration, adaptive data structure, multi-grid-like solver concept, flow through porous media

## I. INTRODUCTION AND MOTIVATION

In its 1987 report [1], the U.S. National Science Foundation stated *computational steering* as valuable scientific discovery. Computational steering is the interaction from a visual front-end with a running simulation (back-end) in order to modify properties such as geometry or boundary conditions for immediate feedback, i. e. allowing users to intervene with an application during runtime and explore any effect caused by their changes [2]. According to [3], “*steering enhances productivity by greatly reducing the time between changes to model parameters and the viewing of the results*”. Even there has been a lot of research on steering within the last two decades, the combination of interaction and high-performance computing (HPC) is still a challenging endeavour.

Typically, massive parallel simulation codes lead to huge data advent that cannot be interactively processed as necessary for a steering approach – in our case the computed results of a fluid flow simulation through porous media. Here, one major problem is the link between front- and back-end, i. e. the network with its bandwidth and latency restrictions. A proper paraphrase, to be found in [4], states that *arithmetic is cheap, latency is physics, and bandwidth is money*. Hence, to overcome this bottleneck special services are necessary in order to leverage steering for demand-driven interactive HPC data exploration.

Therefore, we have designed a so-called *sliding window concept* [5] that allows users to ‘navigate’ through data and retrieve computation results with respect to varying regions of interest and resolutions. Key feature of that concept is to keep the total amount of data to be transferred between back- and

front-end constant, hence any bandwidth limitations are not exceeded at all times. Based on the sliding window, we have developed an HPC service for interactive data exploration, that was tested for parallel CFD applications. We will demonstrate the applicability of this service for micro-scale simulations of flows through porous media and highlight the benefits for researchers in order to gain insight.

The remainder of this paper is as follows. First we will present the underlying physical fluid flow phenomena and their modelling in order to perform a high performance computation. We will then introduce the proposed concept in order to integrate services into the previously described parallel fluid flow computations and present a complex engineering example of flows through porous media. Finally, the paper will close with a short summary and outlook.

## II. HIGH PERFORMANCE COMPUTING CONCEPT FOR COMPUTATIONAL FLUID DYNAMICS SIMULATIONS

The computational fluid dynamics (CFD) simulations are based on an incompressible, isothermal, single-phase Newtonian fluid flow described by the conservation of mass equation written in vector form as

$$\nabla \cdot \vec{u} = 0 \quad (1)$$

and three conservation of momentum equations

$$\frac{\partial u_i}{\partial t} + \nabla \cdot (u_i \vec{u}) = \nabla \cdot (\nu \nabla u_i) - \frac{1}{\rho} \nabla \cdot (p \vec{e}_i) + b_i \quad (2)$$

for  $i \in \{x, y, z\}$ , where  $\vec{u} = (u_x \ u_y \ u_z)^T$  denotes the velocity field in [m/s] in the three spatial dimensions  $x, y, z$ ,  $t$  the current time of the simulation in [s],  $\nu$  the kinematic viscosity in [m<sup>2</sup>/s],  $\rho$  the density in [kg/m<sup>3</sup>],  $p$  the pressure in [Pa],  $\vec{e}_i$  the unity vector in the direction  $i$ , and  $b_i$  sums up external volume forces such as acceleration due to gravity in [m/s<sup>2</sup>].

For the complete historical derivation of the full Navier-Stokes equations (1) and (2) the interested reader is referred to standard literature such as [6, 7, 8].

In order to solve the Navier-Stokes equations, a fractional step approach (also called projection method) proposed by Chorin [9] is applied. An intermediate velocity field is computed while the pressure field is neglected at first. Using this intermediate velocity field, a pressure Poisson equation is established and has to be solved in order to correct and update

the intermediate velocity field and guarantee a convergence free velocity field at the next time step.

The numerical discretisation is based on a finite volume discretisation in space and a second order explicit Adams-Bashforth method in time [10]. While applying Chorin’s projection method, the velocity terms are treated explicitly and the pressure terms are treated implicitly.

The backbone of the numerical discretisation and the basis for a parallel computation is formed by a highly adaptive data structure described in detail in [11]. As a matter of completeness however, the principal points will be introduced briefly.

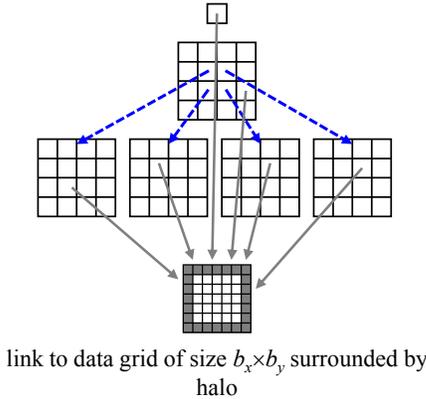
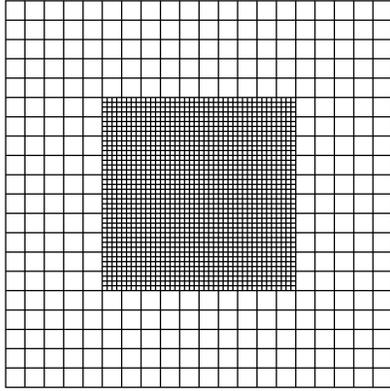


Figure 1. Top part: non-overlapping adaptive grid structure. Bottom part: hierarchical construction of the data structure out of logical grids and data grids.

The data structure is constructed out of non-overlapping hierarchical block-structured orthogonal Cartesian grids and is depicted in figure 1. In the top part, the total agglomerated computational grid can be seen, which will be distributed across different processes for performing a parallel computation in a distributed memory sense. In the bottom part, the hierarchical construction is shown. In general, the data structure comprises two main parts: on the one hand, logical grids manage all hierarchical topological information such as links to parents and children as well as geometrical information such as bounding boxes. On the other hand, data grids containing actual variables in order to perform numerical computations such as velocities, pressures, temperatures etc. are linked to

the previously mentioned logical management grids.

The hierarchical grid construction starts with a logical root grid at depth 0 which can be refined  $r_x^t \times r_y^t \times r_z^t$  times. In case of the example shown in figure 1, four logical grids are refined again  $r_x^s \times r_y^s \times r_z^s$  times which can be repeated recursively until the desired depth has been reached. The refinement on the root level and on all subsequent lower levels can be selected arbitrarily in order to account for a non-cubic initial domain geometry.

Every logical grid is linked to one data grid containing all the variables which are necessary during the numerical simulation. A data grid is oblivious to topological and geometrical information and relies on the link to the logical grid management structure. A data grid can also be called block and is an orthogonal Cartesian grid with the size of  $s_x \times s_y \times s_z$ . Every data grid is identical and  $s_i$  must be divisible without remainder by  $r_i^s$  and by  $r_i^t$  in order to avoid non-conforming setups.

The data grids are padded by one layer of ghost cells in all directions in order to perform a domain decomposition following a Schwarz approach [12]. Hence, the data grids do not overlap, but the ghost cells of the respective grids are overlapping. The solution process is repeated in an iterative manner containing two different phases: an update (or communication) phase and a computation phase. During the update phase, boundary values are exchanged and updated in the corresponding neighbouring data grid’s ghost cells. This process is repeated until the solution has converged.

In case of a serial computation, the order between the update phase and the computation phase is given implicitly by the ordering of the lines of code in the program. In a parallel case using a distributed memory approach no global time is present, and the synchronisation has to be implemented using explicit message exchanges. The order of the synchronisation calls is crucial, as a wrong synchronisation procedure would result in undefined data states. The communication is handled by explicitly exchanging messages using the message passing interface (MPI). The communication structure is divided into three phases (bottom-up, horizontal, and top-down) where a mixture of blocking and non-blocking communication calls ensures a correct predefined ordering. Details about the communication structures can be found in [13].

The implicit pressure Poisson equation is solved immediately on the grid itself rather than assembling one huge system matrix. As solution procedure a multi-grid-like solver integrated into the data structure itself is applied. It uses the previously mentioned communication structures for its solving process and is able to treat large data amounts. Details about performance results are published in [11].

In addition to the computation processes containing the distributed grids, special dedicated processes are running concurrently containing a topological and geometrical repository of all data grids over all processes without containing any actual data values such as pressures or velocities. These processes are called *neighbourhood servers* (NBH) as their primary task is to answer queries such as ‘*who is my neighbour on the*

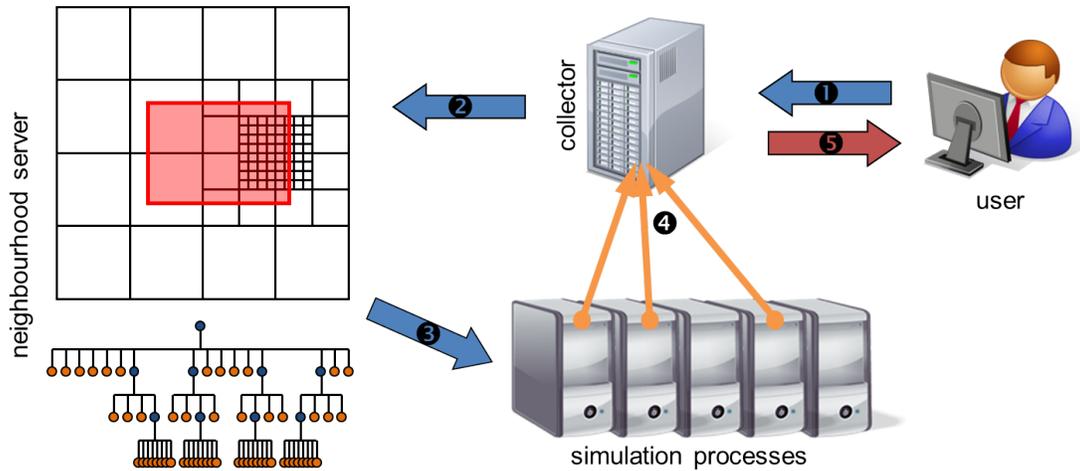


Figure 2. Service Based Concept – User posts requests to collector (1) which are forwarded to the neighbourhood server (2). NBH identifies the simulation processes holding data grids and orders a visualisation (3). The data is then sent to the collector (4), which compresses them and forwards them to the user (5).

*east side?*'. Therefore, the NBH servers have a global view over the complete domain. A unique grid identification (UID) consisting of a 64 bit integer ensures a specific addressing over different domains and different grids containing the encoded rank of the process as well as the grid ID on that specific process. If only one neighbourhood server is used, every computing process has to communicate with the NBH server for querying information about neighbouring grids. This will lead at some point to a bottleneck. The paper [14] describes remedies using multiple NBH servers and special synchronisation techniques in order to overcome the bottleneck problem.

In order to get a good parallel performance, a load distribution has to be performed in order to distribute the amount of grids and, thus, the workload fairly over all available processes. To this extent, space filling curves as introduced in [15] can be used. The presented concept applies a Lebesgue curve [16] (also called z-curve due to its shape) as its linearisation can be described and implemented easily by a Morton ordering [17].

Finally, Frisch2014SCPE shows performance measurements for the parallel computation of uniform as well as adaptive grid setups using all of the aforementioned concepts on the Blue Gene/P installed at Universitatea de Vest din Timișoara (UVT) in Romania.

### III. SERVICE-BASED INTERACTION CONCEPT

Service-based interaction concepts get more and more important in nowadays computation environments. For the visualisation, the paper [5] introduces a so called *sliding window concept* in order to overcome typical bandwidth problems while visualising large domains. Typically, CFD simulations generate giga bytes of data per time step which has to be treated and visualised somehow. The typical state of the art procedure is to dump data for each time step to the supercomputer's network file system (or hard disk), transfer it

completely to a front-end visualisation machine, store it there, and explore and evaluate it in a visualisation environment in order to gain valuable scientific information, as 'the purpose of computing is insight, not numbers'<sup>1</sup>. Unfortunately, this procedure puts a lot of strain on the memory resources as well as the network bandwidth. A sample computation which was done with the above mentioned program using 80 billion cells on SuperMUC (one of Germany's national supercomputers) had a memory footprint of over 28 tera byte per time step. Even if this information can be stored on disk, no visualisation program will be able to open and treat the contents interactively.

Thus, a different approach was chosen. The user selects a window (i. e. region) of interest and together with a given network bandwidth this request is sent to the simulation program. The program then selects the relevant parts by intersecting the window with the computational domain and checks how many information it can put inside the stream until the given bandwidth is exceeded and sends the results back to the user who can visualise it on-the-fly. Using this approach, a user can either have a look at the complete domain in a coarse resolution or zoom into a particular part and see it with all its details. The numerical computation is not influenced by the selection of the view itself; it is always performed on the finest possible resolution given by the data structure and the previous problem definition.

Figure 2 illustrates this request and selection service. A user is requesting to view a certain area with a certain resolution which his network connection can support without problem and sends this request to a dedicated collector server (arrow 1). The collector contacts the neighbourhood server (arrow 2) and transmits the required information.

Once the NBH server has identified the processes which

<sup>1</sup>written by Richard Wesley Hamming in the 1962 preface to *Numerical Methods for Scientists and Engineers*.

contain parts of the domain the user requested, it contacts them (arrow 3) and orders these processes to send the required information to the collector (arrow 4). The collector receives the information, compresses them in a serialised stream and sends them to the user (arrow 5). In case of a slow communication between collector and user, the simulation processes are not blocked as they continue computing as soon as they have sent their information package to the collector.

Thus, the collector server represents an intermediate step separating the user from the computation processes in order not to disturb the simulation itself. Furthermore, it technically decouples the computation processes' MPI communication from the user's socket based communication and acts as a collector and gateway to the outside world. Hence, the complete simulation back-end can run on a supercomputer and the user connects at an arbitrary time to the collector. It is not necessary to be connected at the starting time of the simulation. Furthermore, a user can connect or disconnect to or from the same simulation more than once. Thus, a user can check the progress of the simulation at multiple points in times while it is still running without the necessity of interrupting the simulation or predefining the visualisation steps a-priori. This visualisation based service showed very promising results so far and will be used for porous media flows next.

Instead of ordering now visualisation tasks for interactive data explorations, different kinds of requests can be send to the collector and interpreted accordingly, such as changing of boundary conditions, or refinement of grids. The complete simulation process can be driven by user requests submitted to a back-end simulation framework running on a supercomputer and acting as a service handler.

#### IV. APPLICATION EXAMPLE: GROUND WATER FLOW THROUGH POROUS MEDIA

The following section of the paper will describe a complex engineering application example treating a flow through a porous media using the above mentioned program concepts and services.

In order to describe a fluid flow through an arbitrary porous media, scientists rely on mathematically established and experimentally proven concepts established by Henry Darcy, a French hydraulic engineer of the 19th century. According to his formulation, the ground water flow is directly proportional to the hydraulic gradient over a measured domain  $i = \frac{\Delta p}{\Delta L}$  and a hydraulic permeability  $k$  of earth layers through which the flow circulates.

The hydraulic permeability  $k$  carries within itself a vast amount of information about the physical properties of an investigated medium such as the consistency and distribution of soil particles, the order of soil layers assuming an existence of non-homogeneous scenarios, and the type of soil (in a comparison to already experimentally determined values for several hundred types of soil). Due to a high impact of  $k$  on the accuracy of other hydraulic parameters in a wide range of engineering applications, a lot of effort is put into a precise definition, measurement, and interpretation of the permeability.

The detailed description of experimental apparatuses which can be used to measure hydraulic conductivity  $K$  and, thus, indirectly a permeability  $k$  will not be presented here, as the relation between  $k$  [ $\text{m}^2$ ] and  $K$  [ $\text{m/s}$ ] is already well established and consists only of known properties of a fluid, namely, viscosity, density, and acceleration due to gravitational forces. Furthermore, the settings of physical experiments are highly dependable on the specific goal of examiners such as global values of permeability in one phase flow or particular values for every single flow-phase and for every given soil layer. Due to the complexity of the measurement, we will not attempt to incorporate the complete procedure here, but the interested reader is referred to already existing literature such as [18].

In recent years, the ever increasing performance of computers provided scientists many possibilities to conduct very fine calculations on rather small domains. Hence, current researches often address the coupling of the Navier-Stokes equations (micro scale) with the Darcy equations (macro scale) in order to bridge scales and to use micro scale results within macro scale computations. Any call to the micro scale (as small part of the macro scale domain) requires a full 3D run of the simulation code, which is costly in terms of time and, thus, should be executed only when necessary, i. e. demand-driven due to user requests.

Rather than making general conclusions about advantages and disadvantages that arise from the common concept of these two models, we have intentionally chosen the permeability variable as an obvious link between them and based our research on how this value can be evaluated in one model and then without restriction used in another one. The complete study is conducted with respect to the following assumptions:

- a) the granular pore geometry is highly irregular, without a specific pattern that repeats itself throughout the domain,
- b) the irregularity of the domain is defined by a realistic granulometric sand curve,
- c) a laminar flow regime is used for both models by fixing the Reynolds number to a low value,
- d) the fluid is treated as incompressible,
- e) the dominant flow direction is horizontal, i. e. parallel to the  $x$ -direction of the Cartesian coordinate system,
- f) the influence of the gravity force as a body force is negligible compared to the viscous and pressure terms, but will be taken into consideration, in order to be consistent in the formulation.

##### A. Governing Equations: Darcy Law

In his experimental work Henry Darcy established a very important relationship among several hydraulic parameters which became the basis for vast amount of subsequent theoretical studies. Hence, the water seepage through naturally settled ground layers  $Q$  [ $\text{m}^3/\text{s}$ ] is directly proportional to the hydrostatic pressure difference and the permeability, and inversely proportional to the length of the sample and to the kinematic viscosity of the fluid. The Darcy equation has the following form:

$$\frac{Q}{A} = -\frac{k}{\mu} \cdot \frac{\Delta H_p}{L}, \quad (3)$$

$$K = \frac{k\rho g}{\mu}, \quad \text{and} \quad (4)$$

$$U_{darcy} = -K \frac{\Delta P}{L}, \quad (5)$$

where the ratio  $Q/A$  is often referred to as Darcy flux, which should not be confused with the fluid velocity  $\vec{u}$  [m/s].  $k$  denotes the permeability of the chosen medium,  $\mu$  the dynamic viscosity of the fluid [kg/(m·s)],  $K$  the hydraulic conductivity [m/s], and  $\Delta H_p$  [Pa] the pressure drop on the length  $L$  [m] of the sand bed.

As the calculation and the visualisation of results are essentially conducted in all three directions, the corresponding vector components of the velocity are linked to the respective values of the pressure drop, permeability values [ $k_x, k_y, k_z$ ], and length differences  $\Delta x, \Delta y$ , and  $\Delta z$ .

### B. Governing Equations: Navier-Stokes Equations

The Darcy Law is derived from physical Navier-Stokes equations introduced in section II which describe the fluid flow on a micro scale.

### C. Simulation Setup

In order to be able to evaluate the obtained results, the same pore geometry is used in a set of different scenarios, in which the value of porosity is defined as constant for a single scenario. Such restrictions led to unambiguous calculation settings, defined in figure 3, which are then suitable for subsequent analyses.

Furthermore, in this chapter two cases based on different geometries will be presented, with calculated values of permeability, challenges that we meet, description of ongoing studies as well as how the previously introduced service-based exploration concept can be used throughout the process, in order to analyse our data sets.

### D. Generation of the Domain

The generated testing domains can be split into two groups:

1. with random, irregular geometry, that can be seen in figures 4 and 6 and
2. with regular geometry setup, which is thoughtfully tested, but not presented here, as the first irregular case offers a far better basis for exploration of the introduced sliding window concept.

Whereas all regular domains are generated analytically, taking into consideration only geometrical parameters of inserted spheres, irregular domains are far more complicated and generated as a rectangular shape with predefined dimensions, in which all spheres with different diameters representing soil grains according to the granulometric distribution (cf. figure 3) are placed using random distribution functions. Geometric intersection tests are performed in order to prevent

Sand Grain Diameter [mm]	Phi Scale	Percentage	Number of spheres
3	-1.58	0.013%	1
2	-1.00	0.050%	4
1	0.00	0.502%	40
0.75	0.42	1.442%	115
0.5	1.00	6.070%	484
0.376	1.41	9.582%	764
0.25	2.00	19.403%	1547
0.188	2.41	30.666%	2445
0.126	2.99	25.875%	2063
0.094	3.41	6.133%	489
0.062	4.01	0.213%	17
0.046	4.44	0.050%	4

Figure 3. Granulometric curve along with calculated number of spheres, extracted from the regular domain in order to form a fluid flow domain

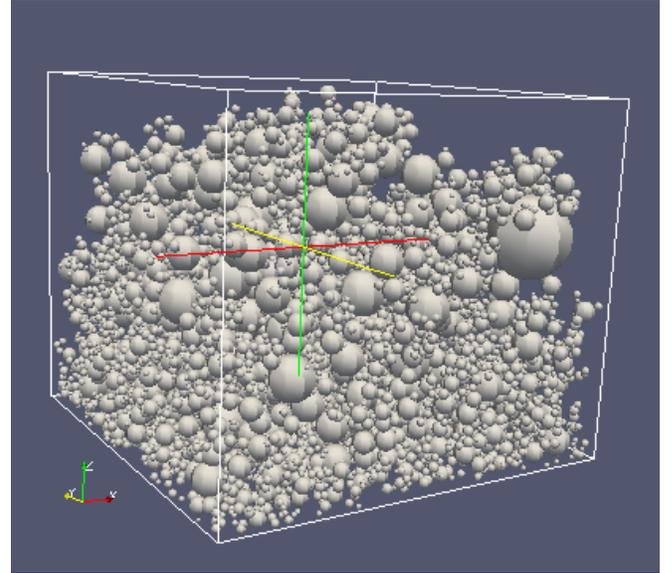


Figure 4. Complex geometry generated using random distribution function (stage 1). Inversion of this geometry forms a fluid flow domain (stage 2)

the overlapping of spheres, such that the generated sample has a physical validity.

Once the sample geometry is generated and described using a boundary representation method, a block-mesh generator based on an octree generation as described in [19] is creating a 3D volume discretisation of the domain by setting the corresponding boundary conditions on the data grids.

Due to agglomerations of ‘artificial’ sand grains near the walls, the results for the relevant parameters such as velocity, pressure, and permeability values were very disperse. In order to reduce local boundary influences, subsequent analyses are conducted on sub-domain samples closer to the centre of the main domain.

The amount of solid particles goes from only 35% up to densely packed 65% of the total volume. As the inflow velocity

can have a significant influence on the velocity distribution throughout the domain – leading to possible numerical instabilities – a larger simulation domain was intentionally set in  $x$ -direction, where a fixed boundary condition value was applied.

The implemented boundary conditions are defined as follows:

- at the left-most boundary – Dirichlet boundary conditions for the inflow velocity with a fixed value of  $u_{in,x} = 1 \text{ m/s}$ ,  $u_{in,y} = u_{in,z} = 0$  and homogeneous Neumann boundary conditions for the pressure,
- at the right-most boundary – free outflow boundary conditions with homogeneous Neumann boundary conditions for the velocity and Dirichlet conditions with a pressure value set to zero,
- at all other walls of the domain – No slip boundary conditions with Dirichlet boundary conditions fixed at zero for the velocities (i.e. the velocity of the fluid is at rest in the vicinity of the wall) and Neumann boundary conditions for the pressure, and
- at the surface of sand grains – No slip boundary condition, as all spheres are treated as solid (similar to the walls), with fixed position over the space and non-deformable over the time particles.

Figure 5 highlights the benefits of an interactive data exploration using the sliding window concept for a complex flow through a porous media. The simulation is always executed on the deepest level of refinement (in this case three levels) using around 6 million cells. On the top level (top part of figure 5) only a few cells are visualised, but they can give a good impression of the overall flow behaviour using a very limited amount of the available, i.e. computed information.

Moreover, the smallest details would not be distinguishable any more as the differences in scales are too big. If more details should be visualised, only one small part of the geometry is visualised in more details (middle part of figure 5). Here, more details can be analysed around a zoomed area visualising more or less the same amount of cells and, thus, keeping the transmitted bandwidth the same as for the top level visualisation. The initial geometry was overlaid in all of the examples in order to get an idea of the zoomed region. Even more details can be seen on a smaller region in the lower part of figure 5.

### E. Further analyses

Taking into account the previously described concept and all benefits that arise from its implementation, different parameters' analyses can be carried out already after first steps of the simulation process, giving us an overall trend of the physical validity of the conducted experiments. One of these examples is depicted in the figures 6, 7, and 8, where the fluid flow simulation through the generated porous medium is presented with depicted sensitivity analyses of permeability value, previously denoted in subsection IV-A. Having on mind the complexity of the generated geometry

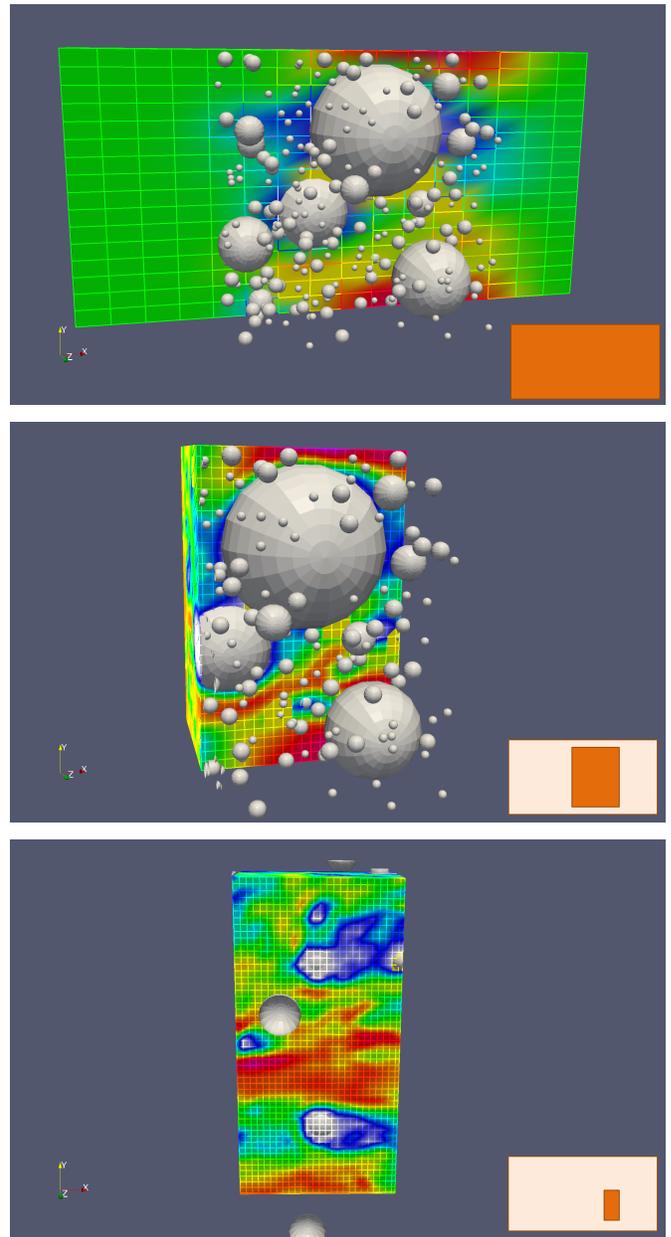


Figure 5. Interactive multi-level domain exploration using the sliding window concept. Displayed are velocity magnitudes on a selected region of interest. In the top picture a ‘global’ view is selected giving a global impression. In the middle and the lower picture, more and more details can be seen while zooming in.

and the complexity of the data structure that supports such cases, analysing the complete data set in real time would be a demanding task, if at all possible to transmit such huge amounts of data over the underlying network, without having side effects. Furthermore, for the most such analyses, the whole data set is not really necessary for studying a specific parameter in a specific part of the domain under particular circumstances. This concept allows us to conduct memory and time consumption saving tests without any loss of data that could influence the results itself.

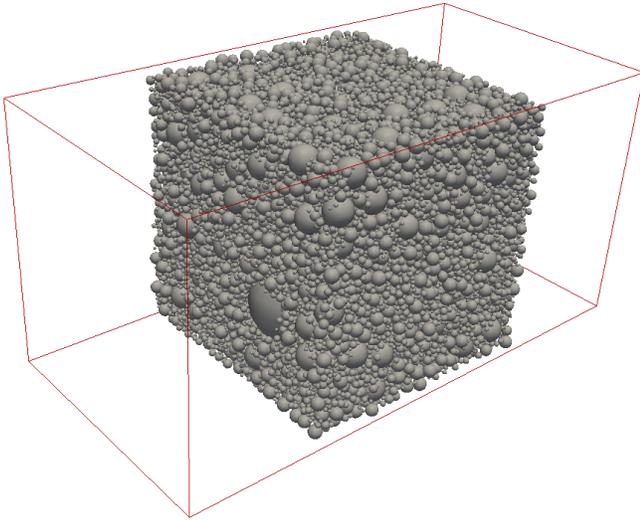


Figure 6. Calculation domain randomly filled up with spheres of different size, with predefined volumetric contribution of all fractions to the total volume of solid phase

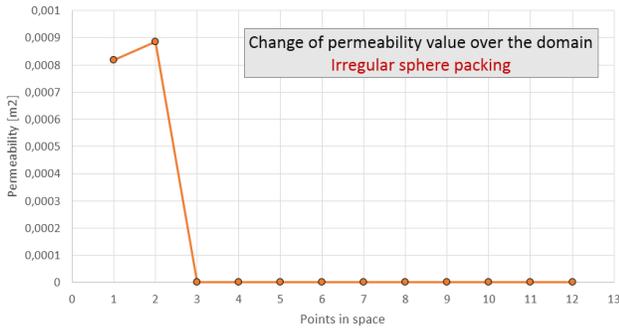


Figure 7. Calculated permeability value for 12 arbitrarily chosen points throughout irregularly (randomly) set domain. Two first points are taken from the part of the domain, where no spheres are positioned

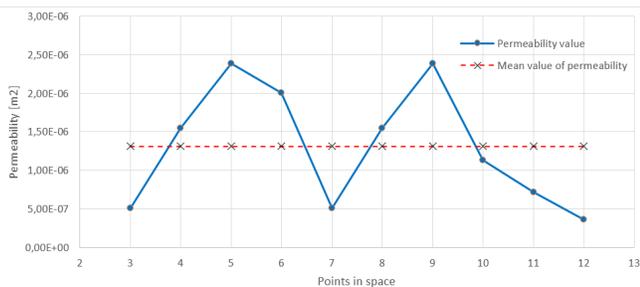


Figure 8. Calculated and mean value of permeability for 10 points solely from part of the domain filled with spheres – irregular sphere packing

In figure 6 randomly generated highly irregular geometry, which was investigated in this particular study is shown. On the deepest level of our data structure (resulting to the resolution of  $320 \times 160 \times 160$  cells) more than 8 million cells are taken into calculation. Analysing parameters at this level could be cumbersome in the means of transferring calculated data for

a particular desired time step. For that reason, the visualisation of the whole domain is done on the coarsest level, in order to check all set boundary conditions and possible conflicts that could be already there effectively spotted. On the next finer level, also depicted in figure 5 (middle plot), the general trend of permeability value is visualised and consequently shown in figure 7, whereas on the deepest, in this case third level of refinement the detailed analyses are conducted, where necessary plots of spatial distribution of explored values are produced. One similar plot is shown in figure 8.

Closely observing enclosed diagrams, it is obvious that some irregularities in behavior of examined parameter are not to be detected on the coarser level, show in figure 7. If there were no possibility to explore mentioned parameter on a finer level, the shown behavior would lead to the wrong conclusion that our permeability value is constant throughout domain and representative for every spatial distribution of inserted spheres. Nevertheless, enclosed finer representation, depicted in figure 8 highlights the spatial fluctuation of permeability (drawn with blue line) comparing to calculated mean value (drawn with red dot line). Such aberration can be explained with fact that the domain was highly irregular with no patterns to be repeated throughout it and therefore the values of permeability are expected to be highly diverse. Ongoing research goes into direction of defining such relationship between aforementioned value and sieve curve (cf. figure 3) on which basis the complex geometry was generated. Such defined permeability value should be completely independent of spatial distribution of solid bodies over the domain. Already gathered results are promising.

Shortly explained setting of conducted analyses is just an example of application of our concept in diverse engineering problems. Nevertheless, it shows clearly, that such analyses are highly feasible for wide range of physical phenomena and their fundamental parameters, providing a powerful tool for experts to either establish further new relationships among variables of interests or explore new behavior by interactively changing the set of input values on-the-fly, having immediate response and logically comprehensible feedback.

Along with the detailed explanation of the implemented high performance computing routines, some first obtained results are presented in this section, with intention to validate (in future) our finite volume concept for porous media flow in comparative analyses to published results based on finite elements solvers as well as to data obtained from – specifically for this purpose – conducted experiments.

## V. CONCLUSION

In order to tackle the huge data advent arising from massive parallel computations, a sliding window concept based on services was introduced in this paper, enabling a user to select the computational domain as a whole or highlight a specific detail and visualise it interactively.

This paper presented the basic concepts and data structures of the underlying computational domain as well as the necessary data exchange and load balancing strategies in order to

execute a simulation on a massive parallel computing systems. Hence, this paper combined the concepts of the sliding window visualisation with HPC computing in order to simulate and interactively explore a complex engineering scenario such as a flow through porous media. The HPC service based approach shall be investigated further in future in order to obtain benefits from on-demand computing and adaptive mesh refinement for instance.

## VI. ACKNOWLEDGEMENT

This publication is partially based on work supported by Award No. UK-c0020, made by King Abdullah University of Science and Technology (KAUST). Furthermore, the authors would like to cordially thank for the support and usage of the Blue Gene/P at Universitatea de Vest din Timișoara (UVT) in Romania.

## REFERENCES

- [1] B.H. McCormick, T.A. DeFanti, and M.D. Brown, “Visualization in scientific computing,” *Computer Graphics*, vol. 21, no. 6, 1987.
- [2] J.D. Mulder, J.J. van Wijk, and R. van Liere, “A survey of computational steering environments,” *Future Generation Computer Systems*, vol. 15, no. 1, pp. 119–129, 1999.
- [3] R. Marshall, J. Kempf, S. Dyer, and C.C. Yen, “Visualization methods and simulation steering for a 3D turbulence model of Lake Erie,” *ACM SIGGRAPH Computer Graphics*, vol. 24, no. 2, pp. 89–97, 1990.
- [4] M. Hoemmen, *Communication-avoiding Krylov Subspace Methods*, Ph.D. thesis, University of California at Berkeley, 2010.
- [5] R.-P. Mundani, J. Frisch, and E. Rank, “Towards interactive HPC: Sliding window data transfer,” in *Proc. of the 3rd Int. Conference on Parallel, Distributed, Grid and Cloud Computing for Engineering*. 2013, Civil-Comp Press.
- [6] G.K. Batchelor, *An Introduction to Fluid Dynamics*, Cambridge University Press, 2000.
- [7] J.H. Ferziger and M. Perić, *Computational Methods for Fluid Dynamics*, Springer, 3rd rev. edition, 2002.
- [8] C. Hirsch, *Numerical Computation of Internal and External Flows, Volume 1*, Butterworth–Heinemann, 2nd edition, 2007.
- [9] A.J. Chorin, “Numerical solution of the Navier-Stokes equations,” *Mathematics of Computation*, vol. 22, no. 104, pp. 745–762, 1968.
- [10] H. R. Schwarz and N. Köckler, *Numerische Mathematik*, Vieweg + Teubner, 8th rev. edition, 2011.
- [11] J. Frisch, R.-P. Mundani, and E. Rank, “Adaptive multi-grid methods for parallel CFD applications,” *Scalable Computing: Practice and Experience*, vol. 15, no. 1, pp. 33–48, 2014.
- [12] H. A. Schwarz, “Ueber einen Grenzübergang durch alternierendes Verfahren,” *Vierteljahrsschrift der Naturforschenden Gesellschaft in Zürich*, vol. 15, pp. 272–286, 1870.
- [13] J. Frisch, R.-P. Mundani, and E. Rank, “Communication schemes of a parallel fluid solver for multi-scale environmental simulations,” in *Proc. of the 13th Int. Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)*. 2011, pp. 391–397, IEEE Computer Society.
- [14] J. Frisch, R.-P. Mundani, and E. Rank, “Resolving neighbourhood relations in a parallel fluid dynamic solver,” in *Proc. of the 11th Int. Symposium on Parallel and Distributed Computing*. 2012, pp. 267–273, IEEE Computer Society.
- [15] M. Bader, *Space-Filling Curves – An Introduction with Applications in Scientific Computing*, Springer-Verlag, 2013.
- [16] H.L. Lebesgue, *Leçons sur l’intégration et la recherche des fonctions primitives*, Gauthier-Villars, 1904.
- [17] G.M. Morton, “A computer oriented geodetic data base and a new technique in file sequencing,” Tech. Rep., IBM Ltd., Ottawa, Ontario, Canada, 1966.
- [18] “Method 9100 – Saturated hydraulic conductivity, saturated leachate conductivity, and intrinsic permeability,” Tech. Rep., United States Environmental Protection Agency EPA, 1986.
- [19] R.-P. Mundani, *Hierarchische Geometriemodelle zur Einbettung verteilter Simulationsaufgaben*, Ph.D. thesis, Universität Stuttgart, 2006.