Results and experiences using grid-based optimization in groundwater management

O. Arndt[1], U. Junghans[1], S. Kaden[1], P. Schätzl[1], F. Thilo[2], A. Gädeke[3]

[1] DHI-WASY GmbH, Member of DHI Group, Waltersdorfer Str. 105, D-12526 Berlin, Germany
[2] Information Systems Institute, University of Siegen, Hölderlinstr. 3, D-57068 Siegen, Germany
[3] Lehrstuhl Hydrologie und Wasserwirtschaft, Konrad-Wachsmann-Allee 6, D-03046 Cottbus, Germany

Abstract A common technique for groundwater management modeling is the Finite-Element method. Optimizing these complex models often requires many hundreds of model simulations which implies an impracticable runtime of the optimization process even on modern computation resources. A grid infrastructure was used to calculate multiple instances of the simulation task in parallel. Results are presented that show the influence of control parameters in a location optimization task.

Keywords simulation, groundwater, grid, optimization, Finite-Element method

Introduction

In the field of groundwater management, increasingly complex scenarios have to be considered, not only from the technical point of view but also with respect to economical aspects. Therefore, large-scale high-resolution groundwater models – often based on widely used Finite-Element (FE) methods – are developed to model environmental systems such as subsurface flow and transport processes in groundwater. In addition to traditional groundwater-management modeling tasks there is an increasing need for modeling under economical aspects like optimal placement of facilities or reduction of energy costs for the operation of pumping stations. Such optimization tasks often involve hundreds or thousands of time-consuming individual simulations which can hardly be realized using traditional computing infrastructure. New techniques are needed to solve complex groundwater modeling tasks on parallel computing resources.

In the past, application of parallel computing in groundwater modeling was rare. If it took place, networks of workstations or compute clusters were the architecture of choice. The main disadvantage of these cluster-based solutions is that in general the scalability is limited by the size of the cluster and not by the problem size, i.e. the full potential of parallelism cannot be exploited. Also the installation and maintenance of clusters can become very costly. An appropriate infrastructure is required that allocates resources on demand. However, there is only a limited availability of such clusters for practical engineering tasks. Grid technology (Foster and Kesselman 2004) is a promising approach to implement a distributed resource infrastructure which is not limited to single clusters, sites or administrative boundaries.

In this paper, the results of an approach are presented which combines the simulation core of the FEFLOW® (Diersch 2007) groundwater modeling and simulation application and the OpTiX suite (Barth et al 2000) as grid-services in a Globus® Toolkit (globus.org 2008) grid environment to enable model-based optimization. The presentation focuses on a well location optimization problem in the Lower Left Rhine coal mining area.

Grid-based architecture

In grid computing, resources from different physical and administrative domains – such as high-performance computing centers or research institutions – are used to solve a single computing task that requires a large number of processing cycles or a huge amount of data. Since version 4 of the Globus Toolkit (GT4) this can be done by implementing and using grid services based on the software engineering concept of Service Orient Architectures (SOA; Singh and Huhns 2005).

To solve complex groundwater modeling tasks, several services were developed and deployed to resources in the grid environment. This includes a simulation service based on the simulation core of the FEFLOW® groundwater modeling and simulation application, an optimization service which utilizes the OpTiX suite as a backend, and a client component integrated into the FEFLOW user interface. A complete service list and a more detailed description were presented in a previous paper (Arndt et al. 2008).

These services were orchestrated to perform the necessary workflow and compute the required simulations on distributed grid resources provided by particular experts, e.g., the optimization service was provided by an optimization expert while computation services were offered by research institutions.

Practical application

In many areas man-made modifications caused dramatic environmental consequences; such as changes in the groundwater table due to water gates or mining activities. To act against a groundwater rise, often pumping stations must operate to protect settlement areas and parks or to retain the course of rivers. In general such pumping stations are built as the need arises, and thus they are often located and operated in a suboptimal manner. Rising energy costs make it imperative to reduce operational costs, e.g., by implementing optimal extraction strategies and by optimizing location placement in case of pumping-station substitution. For this, FE-based optimization is one of the most promising approaches. But even for relatively small optimization tasks, the large number of required simulations prohibits the use of sequential and often also cluster-based optimization of groundwater models. Using the grid-based architecture it was possible to demonstrate the feasibility of this optimization approach by reducing the operational dewatering costs using an optimized pumping station placement.

The working model is based on an existing groundwater model representing the quaternary aquifer between river Rhine and river Niers in the Lower Left Rhine Area with a total area of more than 1000 km². Wide parts of this aquifer are influenced by ground subsidence caused by coal mining activities. During decades of coal mining surface and subsurface dewatering structures were built up as the need arose. As all mining activities in this area will be given up in the near future the no further changes to the underground structure are to be expected which allows a long-term planning of the dewatering system. A test area of 20 km² was chosen to prove the general suitability of the selected optimizing algorithm for location optimization. The test area is very heterogeneous in both land use and hydrogeology which has to be considered during the optimization process.

Within the test area there are several pumping stations at different locations with a total rate of about 27000 m³/d of water. The objective of the optimization task was to minimize the overall pumping rate by keeping a maximum groundwater table within the test area.

Pumping station types are defined by allocating delivery rates. These pumping station types form the decision variables along with the real-world coordinates of the locations and the corresponding extraction rate. The minimum allowed depth to the water table depending on land use was used as constraint that had to be satisfied everywhere. Generated from a digital terrain model (DTM) and a land-use map, and the constraint was stored in the FEFLOW groundwater definition file. But pumping station placement was restricted. The model area was divided in a "go area" where pumping station can be placed, and "no-go zones" unsuitable for pumping stations due to low hydraulic conductivity or other reasons. Inside the go area there were 17 pumping stations represented by 32 single abstraction wells with a total rate of about 17000 m³/d; some stations with an additional amount of 10000 m³/d were not included in the optimization process. The optimization process was started without an initial solution, thus those 17 stations were removed from the model. In addition, the optimization task was configured to place up to 20 abstraction wells with a maximum rate of 5000 m³ per day and well within the go area.

The optimization process was started using 34 computation resources. Each simulation required one computation resource for about 1 hour, hence up to 34 simulations could be calculated in parallel. During the optimization task about 22000 simulations were computed so the optimization took approximately 26 days.

The best solution which was achieved by the optimization had a total extraction rate of 15795 m³/d distributed over 7 well locations, which translates to a reduction of approximately 7%. Perhaps more important than the reduction were two other conclusions: First, the grid-based optimization principally works even for optimizing of locations. Second, obtaining a valid optimization result involving rates and real-world location coordinates requires an enormous number of simulations even for a relatively small model area.

Based on the optimization result several scenarios were analyzed. In a first test case the optimization was restarted with the previous setting but also with the previously achieved optimization result as initial solution. After roughly 9000 additional simulations, the optimization was able to reduce the extraction rate by a further 1.4%. It was noticeable that most of the newly placed wells were located at the same positions as previous wells and also that just 8 of the 20 possible wells were active, i.e., extracting water.

With regard to the small number of active wells, an additional optimization was started. The optimization was configured exactly as the test case before but the number of allowed abstraction wells was limited to 8. The expectation was to achieve better results by reducing the number of decision variables of the optimization algorithm, but the optimization failed. After more than 4000 simulations, the optimization was not able to find an adequate number of valid results to set up the initial reference set which is required by the employed optimization algorithm. For comparison, the previous optimization with 20 possible locations was able to build the reference set after less than 200 simulations.

After the failed optimization attempt, another optimization was started with the number of allowed wells increased to 11. The intent was again to reduce the number of decision variables but to increase the chance to generate feasible solutions and so to build a valid initial reference set. Indeed, the optimization was able to build the initial reference set after about 1900 simulations. After another 12500 simulations, the optimization was able to reduce the extraction rate by 9.14% compared to the initial solution.

Analyzing the previous test cases, it can be summarized that a reduction of decision variables (wells to be placed) can result in a better optimization solution as long as there remain sufficient degrees of freedom to find an appropriate number of valid solutions to build up the initial reference set. It seems reasonable to assume that another reduction in the number of decision variables (e.g., the maximum extraction rate or the go area) can also improve the optimization result.

Based on the previous optimization results and a more detailed analysis of the model simulations, it was decided to exclude a small part of the area from the go area. This was due to the fact that there were some nodes that strongly violated constraints and therefore it was necessary in the initial solution to manually place abstraction wells in that area to avoid constraint violations.

Within the new go area there were 30 abstraction wells with a total rate of 12443 m^3/d. Based on the new configuration, the optimization was restarted with a maximum number of 10 wells to be placed and a maximum extraction rate of 5000 m^3/d per well. Unfortunately, the optimization run stopped after about 2500 simulations due to technical problems. Nevertheless, it was able to reduce the total extraction rate by approximately 3% within that time.

Afterwards, the optimization was repeated with 10 wells to be placed but also with a reduced maximum extraction rate of 3000 m^3/d per well. Using this configuration, the optimization was able to reduce the total extraction rate by more than 19.5 % after approximately 5500 simulation runs. The results of the test cases with the reduced go area are summarized in table 1.

Conclusions and future work

In this paper, results and experiences of a FEM-based optimization using grid technologies were presented. Using a grid infrastructure based on the FEFLOW core simulation system and the OpTiX suite, it was possible to process several test cases and to demonstrate the feasibility of the approach even for location optimization. The test cases gave promising results. It was possible to reduce the overall extraction rates of these test cases using the optimization compared to the initial configuration and it was also possible to eliminate unused wells in that area. It can also be expected that there is a range to improve the results by modifying the optimization settings.

Nevertheless, there is still some work to do. The next step will be to improve the optimization settings, and thus to demonstrate the improvement which can be obtained by using FEM-based

Table 1 *Optimization results for the LINEG test area with a reduced go area*

Number of Wells to be placed	Extraction Rate Interval	Former Extraction Rate	Optimized Extraction Rate	Reduction [%]	Required Simulations
10	0 – 5000 [m³/d]	12443 [m³/d]	12064 [m³/d]	3.0 %	2506
10	0 – 3000 [m³/d]	12443 [m³/d]	9999 [m³/d]	19.6 %	5515

optimization. That includes comparative studies of different algorithms and models as well as scalability analysis using additional resources and larger model areas. Cases where technical problems lead to a premature end of the simulations must be handled in a user-friendly manner so that optimization runs can be continued at any point.

Acknowledgements

This research was supported within German Grid Initiative (D-Grid) by the German Federal Ministry of Education and Research (BMBF) under the contract 01AK806E and within BEinGRID Business Experiments in GRID by the European Commission under the contract number 034702.

References

Ahn, C.W.: "Advances in Evolutionary Algorithms: Theory, Design and Practice", Springer, 2006.
Arndt, O., Junghans, U., Kaden, S., Schätzl. P., Thilo, F.: "Grid-based optimization in groundwater management", in "Proceedings of the 10th IMWA Congress 2008 – Mine Water and the environment", edited by Rapantova, N. and Hrkal, Z., pp. 413 – 416. Ostrava, Czech Republic. 2008.
Barth, T., Freisleben, B., Grauer, M., and Thilo, F.: "Distributed Solution of Simulation-Based Optimization Problems on Networks of Workstations", Journal on Computer Science and Systems, pp. 94—105, 2000.
H.-J. Diersch: "FEFLOW 5.3 – Finite Element Subsurface Flow & Transportation System", User's Manual, WASY GmbH, Berlin, 2007.
Foster, I., and Kesselman, C. (Eds.): "The Grid 2: Blueprint for a New Computing Infrastructure", Morgan Kaufmann, 2004.
globus.org: The Globus Alliance: http://www.globus.org (visited 2008-03-12)
Hough, P., Kolda, T. G., and Torczon, V.: "Asynchronous Parallel Pattern Search for Nonlinear Optimization", SIAM Journal of Scientific Computing 23(1), 134—156 (2001).
Kennedy, J., Eberhart, R., and Shi, Y.: "Swarm Intelligence", Morgan Kaufmann Publishers, 2001
Marti, R., Laguna, M., and Glover, F.: "Principles of Scatter Search", European Journal of Operational Research, Vol. 169, Issue 2, pp. 359—372, 2006.
Singh, M., P.; Huhns, M., N.: "Service-Oriented Computing Semantics", Processes, Agents, Wiley 2005.
W3C.org: World Wide Web consortium: http://www.w3.org (visited 2008-03-13)