

Proof of concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources

Stefan J. Kollet,¹ Reed M. Maxwell,² Carol S. Woodward,³ Steve Smith,⁴ Jan Vanderborght,⁵ Harry Vereecken,⁵ and Clemens Simmer¹

Received 2 October 2009; revised 19 January 2010; accepted 15 February 2010; published 20 April 2010.

[1] We present the results of a unique, parallel scaling study using a 3-D variably saturated flow problem including land surface processes that ranges from a single processor to a maximum number of 16,384 processors. In the applied finite difference framework and for a fixed problem size per processor, this results in a maximum number of approximately 8×10^9 grid cells (unknowns). Detailed timing information shows that the applied simulation platform ParFlow exhibits excellent parallel efficiency. This study demonstrates that regional scale hydrologic simulations on the order of 10^3 km^2 are feasible at hydrologic resolution ($\sim 10^0\text{--}10^1 \text{ m}$ laterally, $10^{-2}\text{--}10^{-1} \text{ m}$ vertically) with reasonable computation times, which has been previously assumed to be an intractable computational problem.

Citation: Kollet, S. J., R. M. Maxwell, C. S. Woodward, S. Smith, J. Vanderborght, H. Vereecken, and C. Simmer (2010), Proof of concept of regional scale hydrologic simulations at hydrologic resolution utilizing massively parallel computer resources, *Water Resour. Res.*, 46, W04201, doi:10.1029/2009WR008730.

1. Introduction

[2] Since *Freeze and Harlan* [1969] proposed a blueprint for an integrated hydrologic model, a number of simulation platforms have been developed in that direction [e.g., *Jones et al.*, 2006; *Kollet and Maxwell*, 2006, 2008a; *Panday and Huyakorn*, 2004; *Qu and Duffy*, 2007; *VanderKwaak and Loague*, 2001]. These platforms have been used to tackle research questions already put forth by *Freeze and Harlan* [1969] including continuity in the variably saturated flow problem, the influence of vegetation on the terrestrial water balance, and feedbacks between the subsurface and weather generating processes of the lower atmosphere. The latter has gained recent attention in both the hydrologic and atmospheric community, which led to model development efforts ranging from simple operational incorporation of subsurface hydrodynamics in atmospheric models [e.g., *Niu et al.*, 2007; *Seuffert et al.*, 2002] to full coupling of three-dimensional subsurface hydrodynamics with land surface and atmospheric models [*Maxwell et al.*, 2007].

[3] In coupling hydrologic and atmospheric models, the obvious disparity in spatial scales and resolutions that are applied in these models is still an unresolved issue. Three-dimensional physics-based hydrologic models commonly

are applied at relatively small spatial scales of $10^{-1}\text{--}10^2 \text{ km}^2$ and high resolution of $10^0\text{--}10^2 \text{ m}$. On the other hand, atmospheric models are commonly applied at the regional to global scale with spatial resolutions ranging from $10^0\text{--}10^2 \text{ km}$. Thus, physics-based hydrologic models basically constitute a single grid cell in a global circulation model. It is important to note that these different scales and resolutions are not arbitrary but commensurate with the major processes that need to be accurately represented in these models. Two major questions arise that require careful scientific and technical consideration: (1) can atmospheric models be applied and coupled at hydrologic resolution (here defined as $10^0\text{--}10^1 \text{ m}$ horizontally and $10^{-2}\text{--}10^{-1} \text{ m}$ vertically) and (2) can hydrologic models be applied at no less than regional scales on the order of 10^3 km^2 while still maintaining hydrologic resolution?

[4] The first question is related to the representation of turbulence and the lower boundary condition in atmospheric models, i.e., the land surface, and constitutes a science question that has recently received considerable attention but remains unresolved [*Huang et al.*, 2009; *Patton et al.*, 2005]. The second question implies the lack of adequate scaling laws in order to upscale the continuity equations of variably saturated subsurface flow (e.g., Richards' equation) to atmospheric resolution. Thus, it is postulated that hydrologic models must be applied at hydrologic resolution over regional scales with acceptable compute times, which requires very large computational resources and has until now appeared impossible.

[5] The benefit of developing and applying highly resolved hydrologic models over regional scales is not only motivated by the potential of coupling with atmospheric models in a physically consistent fashion. These types of models may be useful in serving as virtual laboratories or realities, a term that was quite recently coined in the liter-

¹Meteorological Institute, Bonn University, Bonn, Germany.

²Department of Geology and Geological Engineering and International Groundwater Modeling Center, Colorado School of Mines, Golden, Colorado, USA.

³Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, California, USA.

⁴Computing Applications and Research Department, Lawrence Livermore National Laboratory, Livermore, California, USA.

⁵Agrosphere, ICG-4, Forschungszentrum Jülich, Jülich, Germany.

ature by, e.g., *Weiler and McDonnell* [2006] and *Wood et al.* [2005]. The rational beyond the establishment of virtual laboratories is that experimental studies alone cannot completely solve important scientific problems related to, for example, upscaling of fluxes, variables and parameters, because there will arguably never be enough measured data of appropriate uncertainty across multiple space and time scales. Here high-resolution, regional scale models constituting virtual laboratories, which explicitly resolve the different variances over orders of magnitude, provide a means of exploring these problems and testing upscaling techniques in a formalized fashion. Of course virtual laboratories must always be informed and tested with available measured data. In turn, they constitute a useful tool in experimental design to develop optimized monitoring networks and schedules [*Wood et al.*, 2005]. Thus, a constructive reciprocal connection can be established between experimental studies and virtual laboratories, i.e., large-scale, high-resolution physics-based models.

[6] In this study, we demonstrate that large-scale hydrologic simulations are feasible at hydrologic resolution (10^0 – 10^1 m horizontally and 10^{-2} – 10^{-1} m vertically) and that there is great potential in the use of virtual laboratories. These simulations open new ways of answering important scientific questions related to, for example, upscaling of hydrologic processes in coupled soil-vegetation-atmosphere systems and the analysis of surface-subsurface interactions as well as two-way feedback mechanisms in soil-vegetation-atmosphere systems. We utilized the parallel variably saturated groundwater flow model ParFlow [*Jones and Woodward*, 2001; *Kollet and Maxwell*, 2006] that has been explicitly designed for massively parallel computer environments. Here ParFlow was applied in coupled mode with the land surface model CLM (Common Land Model) [*Dai et al.*, 2001] to account for land surface processes and their interactions with the subsurface. The computational challenge we posed is solving a realistic problem on the order of almost 10^{10} computational cells (unknowns) in a tractable ratio of wall clock time to simulation time, which ultimately affords the simulation of yearly time series. This was done in the framework of a parallel scaling study by successively increasing the problem and processor size in order to evaluate ParFlow's parallel efficiency. In addition, an illustrative numerical experiment was performed that demonstrates the usefulness of the proposed approach.

2. Simulation Platform ParFlow

[7] In this study, the coupled model ParFlow was used to simulate the interactions between land surface processes and variably saturated flow in a heterogeneous subsurface. The core of the integrated watershed simulation platform consists of ParFlow [*Jones and Woodward*, 2001; *Kollet and Maxwell*, 2006], a parallel, three-dimensional, variably saturated groundwater flow code with integrated overland flow that is especially suitable for large-scale, high-resolution flow problems. ParFlow's development has been ongoing for more than 10 years and has resulted in some of the most advanced numerical solvers and multigrid preconditioners for massively parallel computer environments that are available today [*Falgout et al.*, 2006; *Falgout*, 2008].

[8] An additional advantage of ParFlow is the use of a sophisticated octree space partitioning algorithm to depict

complex structures in three-dimensional space, such as topography, different hydrologic facies, and watershed boundaries. In order to generate stochastic models of the hydraulic property distribution of the subsurface, two parallel, correlated, random field simulators have been developed as an integral part of ParFlow [*Tompson et al.*, 1989]. The ParFlow platform also incorporates physical processes that are related to the energy and mass balance at the land surface. This has been done by integrating a land surface model, namely, the Common Land Model [*Dai et al.*, 2003], into ParFlow [*Kollet and Maxwell*, 2008a; *Maxwell and Miller*, 2005]. More specifically, the 3-D variably saturated groundwater flow formulation of ParFlow replaces the 1-D subsurface hydrology of CLM. This also allows to accurately represent a free water table and lateral flow. Thus, ParFlow simulates the 3-D transient moisture redistribution in the subsurface including nonlinear sources and sinks from land surface fluxes (e.g., evapotranspiration) that are calculated by CLM. Topography is derived from digital elevation models and approximated in ParFlow's finite difference framework. Since various land surface fluxes depend on the moisture and energy state of the subsurface, a two-way nonlinear feedback arises between the subsurface and the land surface that is accounted for in the resulting coupled model. The incorporation of the different components in a single numerical framework enables large-scale, high-resolution, integrated watershed simulations that can be used to establish virtual laboratories.

[9] It is important to note that CLM estimates the different land surface fluxes based on similarity approaches that approximate 1-D vertically the turbulent mass, energy, and momentum transport above a rough surface. This means that lateral fluxes are not incorporated explicitly, and that no horizontal scale and resolution are assigned a priori in the application of these approaches. This makes CLM applicable to point measurements as well as global scales, which has been pointed out previously by *Dai et al.* [2003] as a major advantage of their approach. The validity of the applied similarity approaches and thus, the kernel of CLM, mainly depends on the roughness characteristics and heterogeneity of the land surface, and has to be evaluated on a case-by-case basis.

3. Numerical Experimental Setup

[10] In order to evaluate the parallel performance of ParFlow, a scaling study was performed on the JUGENE massively parallel supercomputer (<http://www.fz-juelich.de/jsc/jugene>). JUGENE is an IBM Blue-Gene supercomputer with a total of 294,912 processors and 144TB of memory capable of 0.825 PetaFLOPS (floating point operations per second) that is currently ranked the fourth fastest supercomputer in the world (<http://www.top500.org/lists/2009/11/>). In this scaling study we employed the parallel efficiency, E , following *Jones and Woodward* [2001] and *Kollet and Maxwell* [2006] as

$$E(n,p) = T(n,1)/T(pn,p) \quad (1)$$

where T is the run time as a function of the problem size, n , which is distributed across a number of processors, p . For the case of a perfectly efficient parallel simulator, $E(n,p) = 1$, doubling the problem size and the number of processors




Figure 1a. Semilogarithmic plot of scaled parallel efficiency, E , versus the number of processors, p , of the simulation platform ParFlow.

will result in the same wall clock run time. The problem size per processor (here termed the *unit problem size*) was defined to optimally exploit the relatively small per-processor memory size. This unit problem size was fixed at $45 \times 45 \times 240$ grid cells in the x , y , and z directions, respectively, for a total of 486,000 compute cells. The total problem size was increased by distributing the unit problem over a geometrically increasing number of processors: 512, 1,024, 2,048, 4,096 (one complete JUGENE rack) and 16,384 (four complete JUGENE racks).

[11] The physical problem simulated in this experiment was a fully coupled subsurface–land surface domain. The problem was published previously by Kollet [2009] and consists of a 3-D heterogeneous subsurface with $\Delta x = \Delta y = 1$ m, $\Delta z = 0.025$ m and $n_x = n_y = 45$, $n_z = 240$ as aforementioned. The water table was implemented as a constant head boundary condition at the bottom of the domain with an unsaturated zone above, extending 6 m toward the land surface. The heterogeneous subsurface was simulated as a spatially uncorrelated, log-transformed Gaussian random field of the saturated hydraulic conductivity with a variance ranging over 1 order of magnitude. The vegetation cover was grass using default CLM parameters. The atmospheric time series used to drive the model was the water year 1999 data set previously used by Kollet [2009]. The time step was fixed at 1 h and the total simulation time was 240 h.

4. Results and Discussion

[12] In order to comprehensively interrogate the scaling behavior of ParFlow, timing information was collected for the entire simulation platform (Figure 1a) and the different components (Figure 1b) that are the Solver Setup; the Solver; different objects of the Solver, such as the nonlinear function evaluation (NL_F_EVAL), the nonlinear Newton-Krylov solver KINSOL, the preconditioner MGSEMI; and also the land surface module CLM. Figure 2 shows the relative compute times of the different objects defined as absolute compute times of the different objects scaled by the total wall clock time.

[13] The total wall clock time ranged from some 105 to 182 min for problem sizes ranging from 486,000 cells ($p = 1$) to 7,962,624,000 cells ($p = 16,384$), respectively. This leads to a parallel efficiency that decreases continuously with increasing p and results in $E = 0.58$ for the largest processor number, $p = 16,384$. With increasing problem size, the decrease in E is due to the increase in communication overhead that is required to exchange information at the edges of the computational domains of adjacent processors; nonscalability in the numerical methods; and complete global reduction operations. Nevertheless these E values show that the overall parallel performance of ParFlow is excellent for problem sizes up to one rack of processors ($p = 4,096$), and good for up to four racks ($p = 16,384$).

[14] Figure 1a shows a stronger performance decrease when the problem size is increased from $p = 2,048$ ($E = 0.85$) to one complete rack of processors ($p = 4,096$, $E = 0.75$) and to four racks of processors ($p = 16,384$, $E = 0.58$) which will be explained in the following using Figures 1b and 2.

[15] It is important to note that the total wall clock time of the simulations is relatively small (<182 min). Thus, the time spent in simulation initialization and the setup of the solver infrastructure (Solver Setup) requires an increasingly larger proportion of the total simulation time as the number of processors increases, because the information has to be distributed over a very large amount of processors which requires a considerable amount of interprocessor communication. This is illustrated by an increase of relative compute time from 0.03 to 0.1 for Solver Setup for an increase in p from 4,096 to 16,384 shown in Figure 2. Thus, this negative impact of Solver Setup on the overall parallel performance is small for small problem sizes but becomes amplified with larger numbers of processors due to the small total wall clock times. However, for larger simulation times (i.e., larger wall clock times), the negative impact of Solver Setup on E will decrease considerably, because the setup occurs once at the beginning of the simulation. Thus, E will improve and approach $E \sim 0.65$, which is close to that of




Figure 1b. Semilogarithmic plot of scaled parallel efficiency, E , versus the number of processors, p , for different components of the simulation platform ParFlow including the integrated land surface model CLM.

Solver (i.e., time spent in the solution infrastructure and algorithm alone, excluding communication overhead in the process of initialization) for increasing simulation times.

[16] The above rationale is corroborated by the relative compute times of Solver in Figure 2. Solver, which includes the components CLM, MGSEmi, NL_F_Eval and Kinsol, naturally contributes the largest amount of compute time to the total wall clock time and scales with $E = 0.65$ for the largest processor number, $p = 16,384$. In Solver, the by far largest compute time contribution stems from the nonlinear algebraic solver Kinsol, which includes the nonlinear function evaluation NL_F_Eval. Both components show E values above 0.7 for $p = 16,384$ and, thus, determine the

overall scaling behavior of ParFlow, when the relative contribution of Solver Setup decreases (i.e., when the total simulation time increases). Note that these E values are similar to that observed in the original scaling studies of Kinsol within the ParFlow system [Woodward, 1998].

[17] Simultaneous inspection of Figures 1a, 1b, and 2 allows diagnosis of more aspects of the performance of ParFlow. The scaling of each component taken as a function of overall contribution to simulation time provides great insight into not only which components scale well, but whether they constitute a significant fraction of wall clock time. For example, E values for CLM decrease considerably as the number of processors increases, although there is very

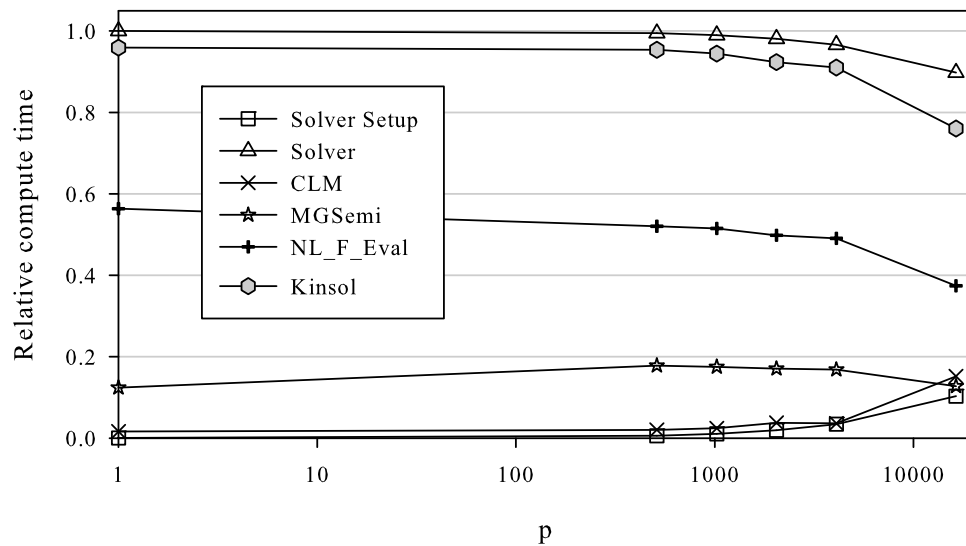


Figure 2. Semilogarithmic plot of relative compute times for different components versus the number of processors, p . Timing information is shown for different components of the ParFlow simulation platform. Compute times of the different components were normalized by compute times for the unit problem size on one processor.




Figure 3. Snapshot of evapotranspiration, ET (W m^{-2}), calculated for a hypothetical, heterogeneous subsurface. The heterogeneity was approximated as a 3-D correlated Gaussian random field in the saturated hydraulic conductivity, K_{sat} , using highly anisotropic correlation lengths of $\lambda_x = 10$ m, $\lambda_y = 1000$ m, and $\lambda_z = 0.1$ m.

little communication overhead in CLM, because the land surface energy and mass balances are calculated for each column individually without lateral transport (i.e., no inter-processor communication is required). The reason for this scaling behavior is that the timing information of CLM includes the initialization of the land surface module, which is performed separately from ParFlow at the beginning of each simulation. This again results in considerable communication overhead, because a large amount of input data needs to be distributed over a very large amount of processors. Thus, better performance can again be expected for increasing simulation times, which will reduce the relative contribution of the initialization at the beginning of each simulation. Additional output overhead from CLM might be another reason, though limited output was specified in the performed simulations.

5. Illustrative Numerical Experiment

[18] In order to illustrate the strength of the proposed approach, a realistic hypothetical example was simulated focusing on the influence of subsurface heterogeneity on evapotranspiration, ET , at the land surface [Kollet, 2009]. In the numerical experimental setup, the difference from the parallel scaling study described above was the representation of the heterogeneity in the saturated hydraulic conductivity, K_{sat} . While uncorrelated Gaussian fields were used in the parallel scaling study, correlated Gaussian fields were generated using highly anisotropic correlation lengths in the x , y , and z directions of $\lambda_x = 10$ m, $\lambda_y = 1000$ m, $\lambda_z = 0.1$ m, respectively. Thus, in order to accurately simulate the influence of heterogeneity in K_{sat} and resulting variability in ET at all scales (ranging from 10^0 to 10^3 m in the horizontal directions), high-resolution, large-scale simulations were required. Due to computational time and memory constraints, these results could not be obtained using commonly applied, nonparallel simulation platforms. Here, these computations were carried out by distributing the aforementioned unit problem size per processor over a total of 16,384 processors (the maximum number used also in the scaling study) resulting in 7,962,624,000 cells in the finite

difference framework of ParFlow. This led to a 3-D computational domain of 6 m thickness and an area of some 33.18 km^2 at a resolution of 1 m and 0.025 m in the horizontal and vertical directions, respectively.

[19] Note that although the simulated problem is hypothetical at this point, it nevertheless constitutes a realistic test case including realistic hydrologic boundary conditions and atmospheric forcing, which is relevant to many science questions dealing with the influence of subsurface hydrodynamics on the mass and energy balance at the land surface. In the presented example, the application of the land surface model CLM at high spatial resolution is warranted, because the land surface roughness characteristics are homogeneous at all spatial scales in the simulations. Under which conditions similarity breaks down is subject of current research and must be evaluated on a case-by-case basis as aforementioned [Huang *et al.*, 2009].

[20] Figure 3 shows a snapshot of the calculated ET distribution for grassland at the land surface due to the heterogeneous subsurface. The ET distribution over the entire domain (Figure 3, left) exhibits clearly the influence of the anisotropy in λ_x and λ_y , which displays narrow, elongated features on the order of kilometers in the y direction. Zooming in on an area of 150×150 m (Figure 3, right) again shows the characteristic elongated features, but additionally resolves explicitly the simulated point variability at meter resolution that is mainly influenced by local, discontinuous subsurface heterogeneity in both the horizontal and vertical directions. Thus, for the first time, simulations resolve the variance in the mass and energy fluxes at the subgrid scale of commonly applied, nonparallel simulation platforms. Since this variability usually needs to be parameterized or interpolated using theoretical approaches in nonparallel simulation platforms, benchmark ParFlow simulations, such as the one presented here, can be useful in the development, testing and verification of these approaches. Additionally, two-way feedbacks in the coupled subsurface–land surface system can be studied at spatial scales ranging over orders of magnitude, which allows developing spatial scaling laws depending on the predominant processes acting

at different spatial scales. Thus, the presented example nicely reflects the basic ideas of virtual laboratories.

6. Implications for Regional Scale Modeling at Hydrologic Resolution

[21] Here we demonstrated excellent parallel scaling of the simulation platform ParFlow to 4,096 processors and good parallel scaling to 16,384 processors for problem sizes up to eight billion compute cells. These parallel scaling experiments also demonstrate a very tractable wall clock to simulation time ratio: some 200 min of wall clock for 240 h of simulation time. These two factors afford, for the first time, catchment-scale simulations ($\sim 10^3$ km²) at the hydrologic resolution (10^0 – 10^1 m resolution laterally and 10^{-1} – 10^{-2} m vertically). For example, assuming a lateral resolution of 10 m with $n_x = 10^4$, $n_y = 10^3$, and a vertical resolution of 0.1 m with $n_z = 10^3$, 10^{10} elements may be simulated in a 1000 km² catchment with a subsurface of 100 m depth! Applying an atmospheric time series in 1 h time step will result in about 4 days of compute time for 1 year of simulation time. Thus, it is now possible to perform large-scale, multiyear numerical experiments at very high spatial and temporal resolution. This unique capability allows us to directly address many important hydrologic scientific questions related to interactions of the subsurface–land surface system, and the scaling behavior of hydrologic processes and parameters. These simulations are currently in preparation in order to study the influence of topography on the scaling behavior on subsurface hydrodynamics at the catchment scale [Kollet and Maxwell, 2008b].

[22] The hypothetical dimensions (n_x , n_y , n_z) provided above also suggest that variability in hydraulic properties, such as hydraulic conductivity, can be resolved over 3–4 orders of magnitude in three spatial dimensions. This presents new possibilities in the observation of the aforementioned spatial scaling behavior and the derivation of accurate spatial scaling laws. On the other hand, since variability can be resolved over many orders of magnitude it will be possible to provide benchmark simulation results that can be used in the validation of novel theoretical approaches of flow and transport in heterogeneous porous media. This was demonstrated using an example simulation of the influence of spatially correlated subsurface heterogeneity on evapotranspiration at the land surface. The example simulation also reflects very well the idea and potential of virtual laboratories that may be used in combination with field data to address the aforementioned scientific questions in formalized fashion.

[23] At this point the parallel scaling study is limited to physics-based simulations of the interactions of the subsurface and the land surface using ParFlow coupled with the land surface model CLM. In the future, additional scaling studies will be performed using ParFlow coupled with the climate model ARPS (Advanced Regional Prediction System) described by Maxwell et al. [2007] to investigate two-way feedbacks of subsurface hydrodynamics with weather generating processes of the atmosphere.

[24] Finally the study shows that while hydrologic sciences may have been slow in comparison to, e.g., atmospheric sciences to adopt parallel programming and computing as a standard tool to tackle important scientific questions, this gap may soon be closed. Therefore, we strongly rec-

ommend increased development efforts and financial support of parallel simulation platforms in the hydrologic science community.

[25] **Acknowledgments.** The financial support by the SFB/TR 32 “Pattern in Soil-Vegetation-Atmosphere Systems: Monitoring, Modeling, and Data Assimilation” funded by the Deutsche Forschungsgemeinschaft (DFG) is gratefully acknowledged. We would also like to thank the John von Neumann Institute for Computing of the Forschungszentrum Jülich and project JICG42, “Inverse Modeling of Terrestrial Systems,” for providing the required compute time on JUGENE. Portions of this work were performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. We also would like to thank Praveen Kumar, John Selker, Eric Wood, Dennis Lettenmaier, and one anonymous reviewer for their constructive comments and suggestions that greatly improved the quality of the manuscript.

References

- Dai, X. P., X. Zeng, and C. D. Dickinson (2001), The Common Land Model (CLM): Technical documentation and user's guide.
- Dai, Y. J., et al. (2003), The Common Land Model, *Bull. Am. Meteorol. Soc.*, 84(8), 1013–1023, doi:10.1175/BAMS-84-8-1013.
- Falgout, R. D. (2008), Multigrid methods, *Numer. Linear Algebra Appl.*, 15(2–3), 85–87, doi:10.1002/nla.586.
- Falgout, R. D., J. E. Jones, and U. M. Yang (2006), The design and implementation of hypre, a library of parallel high performance preconditioners, in *Numerical Solution of Partial Differential Equations on Parallel Computers, Lecture Notes Comput. Sci. Eng.*, vol. 51, edited by M. Bruaset and A. Tveito, pp. 267–294, Springer, Berlin.
- Freeze, R. A., and R. L. Harlan (1969), Blueprint for a physically based, digitally simulated hydrologic response model, *J. Hydrol.*, 9, 237–258, doi:10.1016/0022-1694(69)90020-1.
- Huang, J. P., X. H. Lee, and E. G. Patton (2009), Dissimilarity of scalar transport in the convective boundary layer in inhomogeneous landscapes, *Boundary Layer Meteorol.*, 130(3), 327–345, doi:10.1007/s10546-009-9356-8.
- Jones, J. E., and C. S. Woodward (2001), Newton-Krylov-multigrid solvers for large-scale, highly heterogeneous, variably saturated flow problems, *Adv. Water Resour.*, 24(7), 763–774, doi:10.1016/S0309-1708(00)00075-0.
- Jones, J. P., E. A. Sudicky, A. E. Brookfield, and Y. J. Park (2006), An assessment of the tracer-based approach to quantifying groundwater contributions to streamflow, *Water Resour. Res.*, 42, W02407, doi:10.1029/2005WR004130.
- Kollet, S. J. (2009), Influence of soil heterogeneity on evapotranspiration under shallow water table conditions: Transient, stochastic simulations, *Environ. Res. Lett.*, 4(3), 035007, doi:10.1088/1748-9326/4/3/035007.
- Kollet, S. J., and R. M. Maxwell (2006), Integrated surface-groundwater flow modeling: A free-surface overland flow boundary condition in a parallel groundwater flow model, *Adv. Water Resour.*, 29(7), 945–958, doi:10.1016/j.advwatres.2005.08.006.
- Kollet, S. J., and R. M. Maxwell (2008a), Capturing the influence of groundwater dynamics on land surface processes using an integrated, distributed watershed model, *Water Resour. Res.*, 44, W02402, doi:10.1029/2007WR006004.
- Kollet, S. J., and R. M. Maxwell (2008b), Demonstrating fractal scaling of baseflow residence time distributions using a fully coupled groundwater and land surface model, *Geophys. Res. Lett.*, 35, L07402, doi:10.1029/2008GL033215.
- Maxwell, R. M., and N. L. Miller (2005), Development of a coupled land surface and groundwater model, *J. Hydrometeorol.*, 6(3), 233–247, doi:10.1175/JHM422.1.
- Maxwell, R. M., F. K. Chow, and S. J. Kollet (2007), The groundwater-land-surface-atmosphere connection: Soil moisture effects on the atmospheric boundary layer in fully coupled simulations, *Adv. Water Resour.*, 30(12), 2447–2466, doi:10.1016/j.advwatres.2007.05.018.
- Niu, G. Y., Z. L. Yang, R. E. Dickinson, L. E. Gulden, and H. Su (2007), Development of a simple groundwater model for use in climate models and evaluation with Gravity Recovery and Climate Experiment data, *J. Geophys. Res.*, 112, D07103, doi:10.1029/2006JD007522.
- Panday, S., and P. S. Huyakorn (2004), A fully coupled physically based spatially distributed model for evaluating surface/subsurface flow, *Adv. Water Resour.*, 27(4), 361–382, doi:10.1016/j.advwatres.2004.02.016.

- Patton, E. G., P. P. Sullivan, and C. H. Moeng (2005), The influence of idealized heterogeneity on wet and dry planetary boundary layers coupled to the land surface, *J. Atmos. Sci.*, *62*(7), 2078–2097, doi:10.1175/JAS3465.1.
- Qu, Y., and C. J. Duffy (2007), A semidiscrete finite volume formulation for multiprocess watershed simulation, *Water Resour. Res.*, *43*, W08419, doi:10.1029/2006WR005752.
- Seuffert, G., P. Gross, C. Simmer, and E. F. Wood (2002), The influence of hydrologic modeling on the predicted local weather: Two-way coupling of a mesoscale weather prediction model and a land surface hydrologic model, *J. Hydrometeorol.*, *3*(5), 505–523, doi:10.1175/1525-7541(2002)003<0505:TIOHMO>2.0.CO;2.
- Tompson, A. F. B., R. Ababou, and L. W. Gelhar (1989), Implementation of the three-dimensional turning bands random field generator, *Water Resour. Res.*, *25*(10), 2227–2243, doi:10.1029/WR025i010p02227.
- VanderKwaak, J. E., and K. Loague (2001), Hydrologic-response simulations for the R-5 catchment with a comprehensive physics-based model, *Water Resour. Res.*, *37*(4), 999–1013, doi:10.1029/2000WR900272.
- Weiler, M., and J. R. J. McDonnell (2006), Testing nutrient flushing hypotheses at the hillslope scale: A virtual experiment approach, *J. Hydrol.*, *319*, 339–356, doi:10.1016/j.jhydrol.2005.06.040.
- Wood, E. F., J. Bolle, P. Bogaart, and P. A. Troch (2005), The need for a virtual hydrologic laboratory for PUB, *IAHS Publ.*, *301*, 189–203.
- Woodward, C. S. (1998), A Newton-Krylov-multigrid solver for variably saturated flow problems, in *Proceedings of the Twelfth International Conference on Computational Methods in Water Resources*, vol. 2, edited by C. A. Brebbia et al., pp. 609–616, Comput. Mech. Publ., Southampton, U. K.
-
- S. J. Kollet and C. Simmer, Meteorological Institute, Bonn University, Auf dem Huegel 20, Bonn, D-53121, Germany. (stefan.kollet@uni-bonn.de; csimmer@uni-bonn.de)
- R. M. Maxwell, Department of Geology and Geological Engineering, 1500 Illinois St., Colorado School of Mines, Golden, CO 80401, USA. (rmaxwell@mines.edu)
- S. Smith, Computing Applications and Research Department, Lawrence Livermore National Laboratory, Livermore, CA 94551, USA. (smith84@llnl.gov)
- J. Vanderborgh and H. Vereecken, Agrosphere, ICG-4, Forschungszentrum Jülich, Jülich, D-52425, Germany. (j.vanderborgh@fz-juelich.de; h.vereecken@fz-juelich.de)
- C. S. Woodward, Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, CA 94551, USA. (cswoodward@llnl.gov)