Transp Porous Med (2015) 108:151–169
DOI 10.1007/s11242-014-0330-7

Field-Scale Simulation of Production from Oceanic Gas Hydrate Deposits

Matthew T. Reagan · George J. Moridis · Jeffery N. Johnson · Lehua Pan · Craig M. Freeman · Lehua Pan · Katie L. Boyle · Noel D. Keen · Jarle Husebo

Received: 1 April 2013 / Accepted: 6 May 2014 / Published online: 3 June 2014

Abstract The quantity of hydrocarbon gases trapped in natural hydrate accumulations is enormous, leading to a significant interest in the evaluation of their potential as an energy source. It has been shown that large volumes of gas can be readily produced at high rates for long times from some types of methane hydrate accumulations by means of depressurization-induced dissociation, and using conventional horizontal or vertical well configurations. However, these resources are currently assessed using simplified or reduced-scale 3D or 2D production simulations. In this study, we use the massively parallel TOUGH+HYDRATE code (pT+H) to assess the production potential of a large, deep ocean hydrate reservoir and develop strategies for effective production. The simulations model a full 3D system of over 38 km^2 extent, examining the productivity of vertical and horizontal wells, single or multiple wells, and explore variations in reservoir properties. Systems of up to 2.5 M gridblocks, running on thousands of supercomputing nodes, are required to simulate such large systems at the highest level of detail. The simulations reveal the challenges inherent in producing from deep, relatively cold systems with extensive water-bearing channels and connectivity to large aquifers, mainly difficulty of achieving depressurization and the problem of enormous water production. Also highlighted are new frontiers in large-scale reservoir simulation of coupled flow, transport, thermodynamics, and phase behavior, including the construction of large meshes and the computational scaling of larger systems.

Keywords Gas hydrates · Methane hydrates · Oceanic hydrates · Gas production

1 Introduction

Gas hydrates are solid crystalline compounds in which gas molecules occupy the lattices of ice-like crystal structures called hosts (Sloan and Koh 2008). They may occur in two distinctly

M. T. Reagan (✉) · G. J. Moridis · J. N. Johnson · L. Pan · C. M. Freeman · L. Pan · K. L. Boyle · N. D. Keen
Lawrence Berkeley National Laboratory, 1 Cyclotron Rd., Berkeley, CA 94720, USA
e-mail: mtreagan@lbl.gov

J. Husebo
Statoil ASA, Research Centre, Sandsliveien 90, 5020 Bergen, Norway

different geographic settings, in the permafrost and in deep ocean sediments, where the necessary conditions of low T and high P exist for their formation and stability. The majority of these naturally occurring hydrates contain CH_4 in overwhelming abundance. Interest in hydrates is enhanced by ever-increasing global energy demand and the environmental desirability of natural gas. Although there has been a limited work mapping and evaluating this resource on a global scale (Moridis et al. 2009), current estimates of in-place volumes vary widely (ranging between 10^{15} and 10^{17} ST m^3), but the consensus is that the worldwide quantity of hydrocarbon gas hydrates is vast (Milkov 2004; Klauda and Sandler 2005; Burwicz et al. 2011). However, not all hydrates are desirable targets for production (Moridis et al. 2011b). Yet even if only a small fraction of the most conservative estimate is recoverable, the sheer size of the resource is so large that it demands evaluation as a potential energy source.

Recent studies (Moridis and Reagan 2007a, b) have indicated that, under certain conditions, gas can be produced from natural hydrate deposits at high rates over long periods using conventional technology. Of the three possible methods of hydrate dissociation (Makogan 1997) for gas production—depressurization, thermal stimulation, and use of inhibitors—depressurization appears to be the most efficient (Moridis et al. 2009). Thermal stimulation has shown promise only as a means of protecting the near-well zone from permeability-lowering secondary hydrate formation (Moridis and Reagan 2007a, b), not as a means of facilitating large-scale hydrate dissociation within the reservoir. The same holds for dissociation via inhibitors—the circulation of warm, saline water in the vicinity of the well may be effective in clearing the near-well zone of ice or secondary hydrate, but it is an inefficient way to create large-scale dissociation in the reservoir (Moridis and Reagan 2007a). A key feature of productive reservoirs is impermeable boundaries, both to allow for the formation of free gas through significant depressurization over a large spatial extent in the reservoir, and to ensure that water production due to flow through the boundaries is manageable (Reagan et al. 2008; Myshakin et al. 2012). Earlier simulation studies have focused on production from vertical wells, but more recent studies (Moridis 2008b; Moridis et al. 2011b) show that horizontal wells can be more productive and easier to manage than vertical wells for confined Class 3 and some Class 2 hydrate reservoirs (Boswell et al. 2011). Simulation work in this area has expanded to heterogeneous 2D simulations based on more complex, data-driven geological models (Myshakin et al. 2012), and to 3D simulations, with model development through the use of seismic data becoming possible as simulation technology improves and larger simulations become possible (Myshakin et al. 2012).

The objective of this study is to simulate a realistic, 3D gas hydrate reservoir, using real geophysical data at the field scale, and also to evaluate a real production target for feasibility and productivity. Previous studies using TOUGH+HYDRATE (Moridis et al. 2009, 2011a, b) have tended to focus on simple 2D modeling or 2D modeling with limited heterogeneity. Through a collaboration with Statoil, we were able to access real data on the geometry and geology of a known oceanic hydrate system that has been considered for commercial exploitation. Due to confidentiality agreements, the exact location of this deposit cannot be published, but the knowledge gained from this simulation work can be generally applied, and used to evaluate other similar oceanic deposits.

The deposit is illustrated in Fig. 1. It is an oceanic system located in the Gulf of Mexico, approximately 7 × 5.5 km and 350 m in thickness. The hydrate-bearing coordinates, represented by a point cloud (brown pixels) in the Figure, are arranged in high-permeability "channels" bounded by slightly lower-permeability "levees" and regions of low-permeability shale. Note, via the slicing plane, the layering of the hydrate-bearing channels. Each of the layers of hydrate may be up to 20 m thick, and several layers extend horizontally for thou-

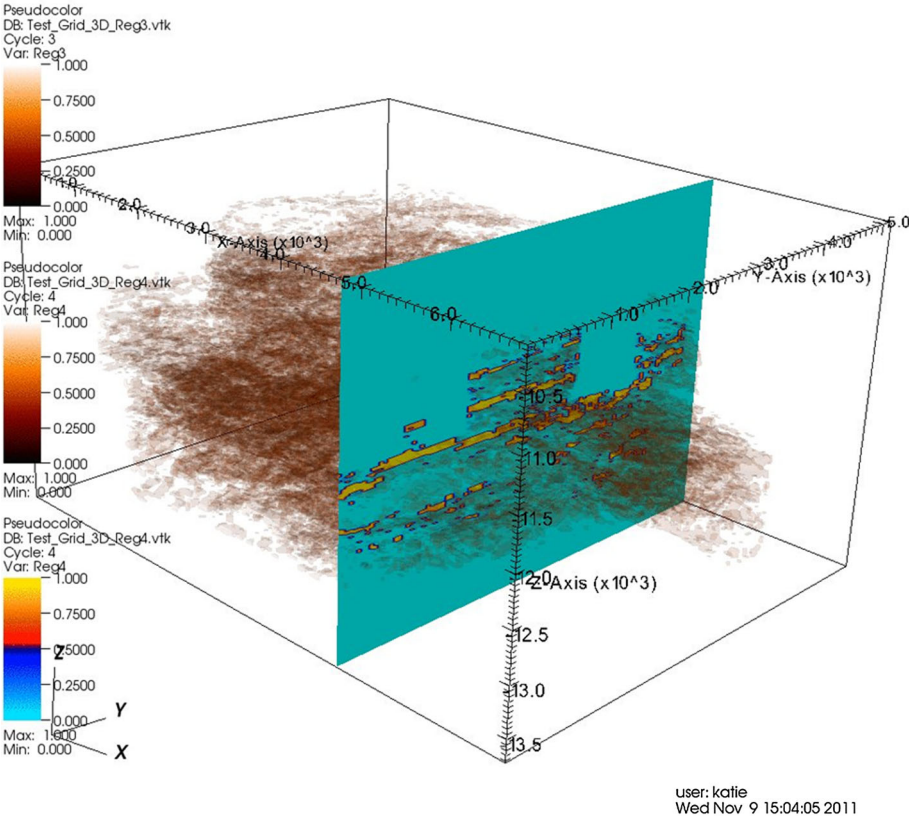

Fig. 1 Illustration of the location of hydrate layers within the reservoir domain. Each *brown* pixel represents a hydrate-containing location within the 3D geophysical dataset. Hydrate locations crossed by the slicing plane are highlighted in *yellow*

sands of meters. The system is likely impermeable at the top and bottom boundaries, but may connect to an aquifer along the $x - z$ face.

2 Methodology

2.1 Mesh Generation

The construction of 3D volume meshes for systems of this scale is well beyond the means of the standard TOUGH+ MeshMaker routine. Therefore, we deploy advanced tools.

For vertical well scenarios, the grid was generated using WinGridder (Pan 2008), an interactive application developed for the Yucca Mountain project. The initial vertical well configuration involves a roughly rectilinear, layer-by-layer mesh of the reservoir, as taken directly from the Statoil dataset, with the discretization matching that of the geophysical data. A cylindrical mesh, with a center at the well, was placed through the region of greatest hydrate accumulation, with the two meshes interpolated at the boundary to conform to TOUGH element-connection rules. The vertical well mesh is illustrated in Fig. 2.

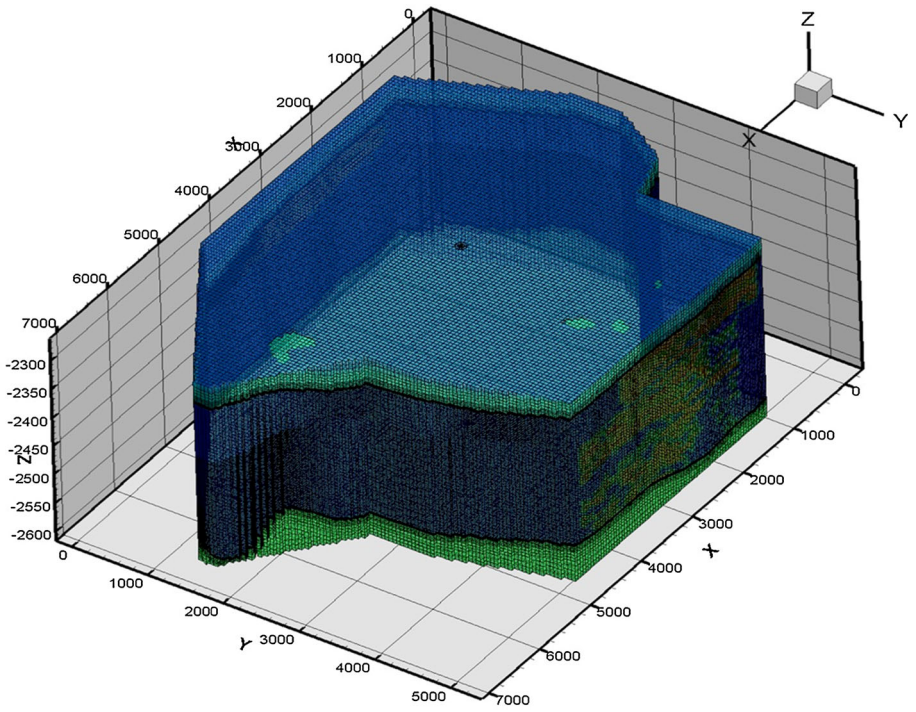

Fig. 2 Illustration of the 3D finite-volume mesh, with trimmed boundaries, for the single vertical well scenario. The intersection of the cylindrical well zone mesh with the top boundary can be seen in the center of the top surface

Note in the figure that the domain has been trimmed, with regions of the mesh that do not represent permeable reservoir rocks removed to promote computational efficiency. A 30 m buffer is placed at the top and bottom of the system, large enough to place the upper and lower fixed boundaries beyond the limit of heat transfer over multi-year production timescales. All subsequent meshes begin with this "reduced" mesh, which contains 1,663,900 elements.

WinGridder does not have the ability to interpolate a horizontal near-well zone into a roughly rectilinear mesh. Therefore, to build the horizontal well system, a combination of custom tools for manipulating the raw list of element centroids and creating new mesh-element configurations (i.e., cylinders) is combined with the meshing toolkit Voro++ (Rycroft 2009). Voro++ is a C++ library capable of generating a fully-3D Voronoi mesh for any valid configuration of cell centers. The Voro++ library keeps track of all relevant cell and interface properties and makes them available for manipulation. This allows the creation of highly flexible and dynamically refined meshes. Voro++ is the heart of a new python-language TOUGH+ meshing toolkit known as MeshVoro, and the horizontal mesh created here in the first application of this code.

For the horizontal well, the geophysical data were analyzed to locate the longest column of continuous horizontal layers of hydrate more than 60 m thick. A region with thickness of roughly 65 m was also located near the core of the system, accommodating a horizontal well 1,900 m in length that has no intersection with non-hydrate-bearing channels or levees. A cylindrical mesh was constructed around the well to a radius of 250 m, and this radial

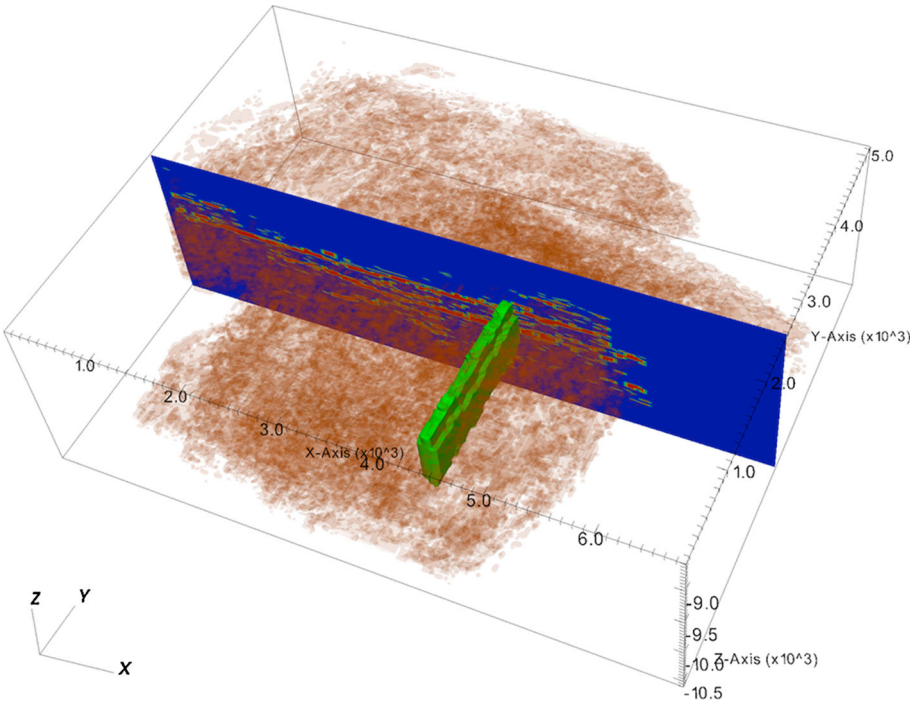

Fig. 3 Diagram of the placement of the horizontal well and surrounding 250 m-radius cylindrical mesh. Hydrate-bearing sediments are plotted as *brown* volumes

mesh was inserted into the rectilinear mesh derived from the geophysical data (Fig. 3). Element properties (i.e., hydrate saturation) were interpolated from the rectilinear data onto the new mixed-geometry mesh. The resulting mesh created by MeshVoro contains 2,264,000 elements, requiring the simultaneous solution of over 9,000,000 equations, making this one of the largest TOUGH simulations ever attempted.

Not surprisingly, meshes of this size are challenging to manipulate. Bespoke perl and python-based scripts were developed to edit the MESH files after generation (as they are too large for most fully-graphical text editors). The system was brought first to hydrostatic equilibrium at the known depth (proprietary), and then to thermal equilibrium using the known geothermal gradient. Due to the heterogeneity of both the mesh and system itself, equilibration through replicating a handful of 1D columns was not possible, and a full 3D representation of the system had to be equilibrated at each step, consuming significant computing time. Hydrate was added to the system according to the distribution indicated from the geophysical data. The system was then allowed to reach full thermal, hydrological, and chemical equilibrium at stated conditions before any simulations were performed.

2.2 Simulators

Due to the computational challenges of simulating systems of this size in 3D, with appropriately fine discretization, we use the MPI-parallel TOUGH+HYDRATE code (pT+H) (Zhang et al. 2008). pT+H contains the same coupled thermal- hydrological- chemical capabilities

Table 1 Reservoir parameters

Region	k (mD)	ϕ	S_H	S_W
0. Boundary layers	0	0.0	0.0	1.0
1. "Levee" deposits 1	100	0.2	0.0	1.0
2. "Levee" deposits 2	100	0.21	0.0	1.0
3. "Levee" deposits 3 (w/hydrate)	100	0.22	0.7	0.3
4. Channel 1 (w/hydrate)	10,000	0.33	0.7	0.3
5. Channel 2	10,000	0.34	0.0	1.0
6. Channel 3	10,000	0.35	0.0	1.0

of serial TOUGH+HYDRATE v1.1 (Moridis et al. 2008a), but can be executed on shared- or distributed-memory clusters using standard OpenMPI libraries (http://www.open-mpi.org/).

The current version of the code typically has been used on small clusters, running on 50–100 nodes for 2D reservoir simulation at medium- to high-resolutions (Moridis et al. 2011a). However, due to the larger mesh (2.46 vs. 1.66 M gridblocks) and the nature of the 3D Voronoi grid (on average, 2X–3X more connections per element), the computational requirements for this simulation are approximately an order of magnitude greater. Initial equilibration was performed using our 220-processor in-house cluster, but for some production cases, pT+H was ported to the Hopper supercomputer, part of NERSC, for multi-day runs using 960 and 1,920 processors.

2.3 Reservoir Properties

Although full details of the system may not be disclosed, the initial set of reservoir properties as derived from proprietary geophysical surveys, is shown in Table 1.

Regions 3 and 4 are the hydrate-bearing media seen in Figs. 1 and 3. Permeabilities as listed are intrinsic permeabilities of the native rock without hydrate. Relative permeability functions governing the effective permeability of hydrate-bearing rocks are described in Table 2. Geophysical data suggested that the presence of a small amount free gas in the reservoir (in the formations "Levee" 2 and Channel 2), but TOUGH+HYDRATE pre-run thermodynamic consistency checks ruled this as unlikely under the stated conditions (and assuming a system at thermal and hydrostatic equilibrium), therefore, the gas was removed from the model. Hydrate saturations could not be estimated from the seismic dataset, and therefore, a uniform hydrate saturation based on limited reported core data has been applied to the hydrate-bearing regions.

For additional reservoir parameters, we used data from one of the few well-characterized, reservoir-grade oceanic hydrate deposits, the Tigershark deposit (Moridis and Reagan 2007a, b). Relative permeability exponents were estimated by fitting to Statoil effective permeability data for the system. To describe the well itself within the simulation domain, flow through the wellbore is represented using the pseudoporous-medium approach of Moridis and Reagan (2007a). For constant-P production, the topmost gridblock in the subdomain representing the wellbore is treated as an internal boundary that was maintained at the constant bottomhole pressure P_w and the contributions of the various gridblocks in contact with the wellbore are determined from the relative mobilities. The pressure specified at the well is selected to be 3.0 MPa (above the quadruple point) to prevent ice formation in the wellbore. Other details about Tigershark and examples of production from such a system, can be found in Moridis and Reagan (2007a, b). Key simulation parameters are summarized in Table 2.

Table 2 Other simulation parameters

Water salinity (mass fraction)	0.035
Reservoir temperature	10.3–16.6 °C
Reservoir pressure	255 bar
Grain density r_R (all formations)	2750 kg/m^3
Constant pressure at the well P_w	3.0×10^6 Pa
Dry thermal conductivity k_{QRD} (all formations)	0.5 W/m/K
Wet thermal conductivity k_{QRW} (all formations)	3.1 W/m/K
Composite thermal conductivity model (Moridis et al. 2008)	$k_{QC} = k_{QRD} + (S_A^{1/2} + S_H^{1/2})(k_{QRW} - k_{QRD}) + f\, S_I k_{QI}$
Capillary pressure model (van Genuchten 1980; Moridis et al. 2008)	$P_{cap} = -P_0 \left[(S^*)^{-1/\lambda} - 1 \right]^{(1-\lambda)}$ $S^* = \frac{(S_A - S_{irA})}{(S_{mxA} - S_{irA})}$
S_{irA}	1
λ	0.77
P_0	10^6 Pa
Relative permeability model (Moridis et al. 2008)	$k_{rA} = (S_A*)^n$ $k_{rG} = (S_G*)^n$ $S_A* = (S_A - S_{irA})/(1 - S_{irA})$ $S_G* = (S_G - S_{irG})/(1 - S_{irA})$ OPM model (Moridis et al. 2008)
n (fitted from data)	4.4292
S_{irG}	0.02
S_{irA}	0.20

3 Results and Discussion

3.1 Production from a Vertical Well

Initial simulation work by Statoil, using the CMG-STARS with gas hydrate add-ons, suggested that the system could be productive if vertical wells were drilled into the lower part of the formation near areas of high hydrate saturation. However, these simulations used a 2D slice of the full system, and assumed the existence of free gas in the initial equilibrium state of the reservoir. Initial TOUGH+HYDRATE equilibration work indicated that such a configuration of gas and hydrate was non-physical. In contrast, we begin by simulating production with fully coupled thermodynamics (including pre-checks), from the 3D, single vertical well shown in Fig. 2.

We simulate three sets of well configurations, with several variations in strategy:

(1a) A well perforated throughout the entire hydrate-bearing zone ("long interval")
(1b) A well perforated only within the topmost hydrate layer ("short interval-top")
(1c) A well perforated only within the bottommost hydrate layer ("short interval-bottom")

(2a) A well perforated throughout the entire hydrate-bearing interval, with the impermeable system boundaries brought inward to simulate the well as part of 500-m-well-spacing pattern (no-flow boundaries 250 m from the well)
(2b) A well perforated only within the topmost hydrate layer, as part of a 500 m pattern
(2c) A well perforated only within the bottommost hydrate layer, as part of a 500 m pattern

We also estimate sensitivity to system permeability by performing four additional simulations based on Case 2b, reducing the permeability of the non-hydrate-bearing levees to:

(3a) $k_{\text{levee}} = 100$ mD,
(3b) $k_{\text{levee}} = 10$ mD,
(3c) $k_{\text{levee}} = 1$ mD, and
(3d) $k_{\text{levee}} = 0.01$ mD.

For each case, we simulate up to 3 years of constant pressure production, using 200–220 processors with pT+H on a dedicated computing cluster. Note that for all cases (including the full system without no-flow boundaries at 250 m), the sub-region within 250 m of the well contains the largest total quantity of hydrate of any such 500-m square sub-region within the deposit, a quantity that is approximately twice that of the median quantity for the set of all 500-m sub-regions. These simulations, particularly Cases 2a–c and 3a–d, are therefore, a good "best-case" example of production from a single vertical well, as we are always attempting to produce from the region of the deposit with the highest total hydrate content.

Figures 4 and 5 track the evolution of hydrate dissociation, and hence the release of gas into the reservoir, for the ten single vertical well Cases. In Fig. 4, we see that hydrate dissociation begins immediately for the case of a long production interval (Case 1a), but quickly subsides

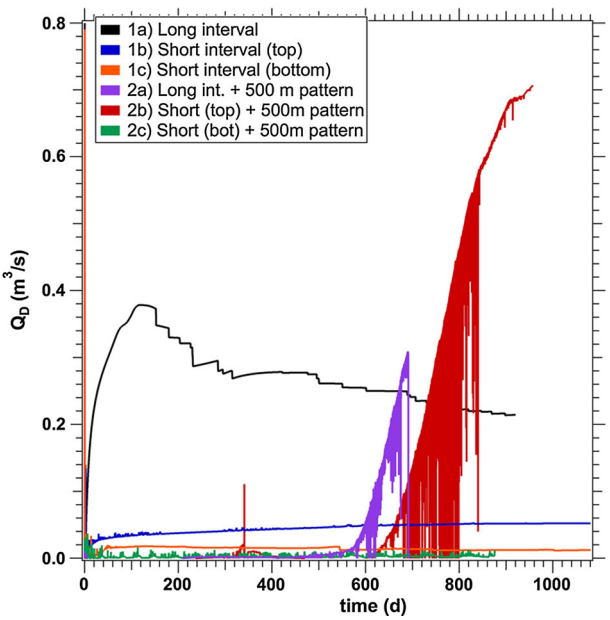

Fig. 4 Rate of hydrate dissociation, Q_D, across the system as a whole, for the vertical well scenarios Cases *1a–c* and *2a–c*

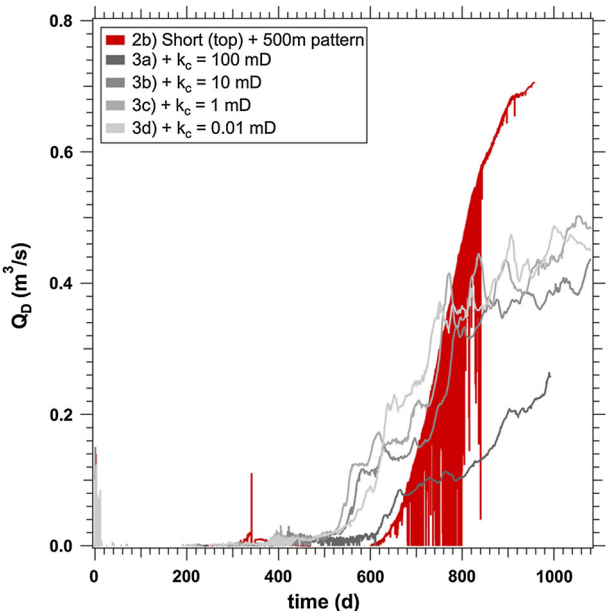

Fig. 5 Rate of hydrate dissociation, Q_D, across the system as a whole, for the vertical well scenarios Case *2b* and *3a–d*

and roughly levels off after 200 d of production. Producing from a shorter interval within the top layer of hydrate (Case 1b), results in far less dissociation, a release rate that plateaus at less than 0.03 m^3/s, and no sign of strong methane release within the 3-year production test window. Case 1c, also a shortened production interval, creates even lower rates of hydrate dissociation.

However, if we are producing as a part of a 500 m pattern, as shown in Case 2a–c, the evolution of dissociation versus time is distinctly different. Comparing Case 2a–1a, we see a 550 d delay in the onset of significant hydrate dissociation in the reservoir. However, the dissociation curve for Case 2a stops abruptly at t = 690 d. This is due to production (and simulation) shutdown due to extensive formation of secondary hydrate in and around the well. Comparing Case 2b–1b (short production interval within the top hydrate lens), we see a similar 600 d delay, followed by the strongest onset of hydrate dissociation of any of the previous cases. Case 2c, however, shows no onset of strong dissociation within the 3-year test period, and release rates worse than Case 1c.

One would expect that no-flow boundaries imposed at distances 250 m from the well along the x and y axes would result in the quickest depressurization and the strongest dissociation of hydrate, but Case 1a shows the most rapid onset—although the dissociation trails off after less than a year. The reason for this will become clear later in the production analysis.

In Fig. 5, we replicate the release curve for Case 2b, and then present Cases 3a–d as a comparison (note that Case 2b was not chosen *a priori* for additional simulation cases, but only after confirming the strong dissociation generated by this particular well configuration). We see strong dissociation/methane release after 1.5 years of depressurization for all of the reduced-permeability levee scenarios, and the reduction in permeability corresponds to earlier initiation of hydrate dissociation. It is somewhat surprising that Case 2b still shows

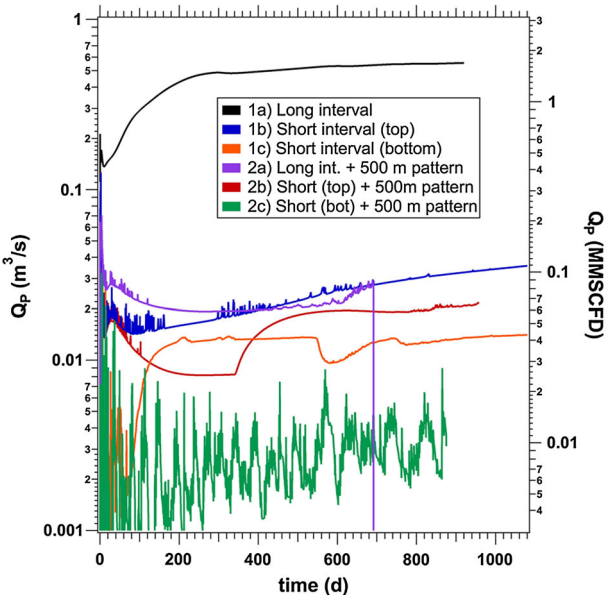

Fig. 6 Rate of methane production at the well, Q_P, for the vertical well scenarios Cases *1a–c* and *2a–c*. Note the log scale

stronger dissociation than any of the reduced-permeability cases in 3a–d, and once again, the explanation will appear later in the analysis.

Figures 6 and 7 show the resulting methane production at the well, including both dissolved and gaseous methane, for the various scenarios. Note that the y-axis, representing production rate, is now a log scale, reflecting the large relative differences in production rates between the cases. In Fig. 6, the full interval, Case 1a, generates the highest rates of total methane production by an order of magnitude (unsurprising at first, due to the much great extent of the perforated interval), with other configurations lagging significantly. However, production levels off quickly at 5 m^3/s (1.7 MMSCFD) in Case 1a, and this method still does not achieve commercially viable rates within 3 year. As Fig. 4 shows no acceleration of hydrate dissociation after 100 d, it is unlikely the rate of methane production will improve at any time for this case. While Case 2b showed promise due to increasing rates of dissociation in Figs. 4 and 5, gas production at the well is lackluster, as are all of Cases 1b–c and 2a–c. Figure 7 again compares Case 2b to 3a–d, with all cases of reduced levee permeability resulting in increased, if not necessarily commercially viable, gas production rates.

The cumulative mass of water removed at the well, for all vertical well cases, is shown in Fig. 8. The difference between Case 1a and all other cases becomes clear—Case 1a requires the removal considerably more water at the well, compared to the next-most productive configurations, to achieve the dissociation and production rates seen in Figs. 4 and 6. All other cases, save the low-production Case 2c, converge upon average water production rates of about 3-30 kg/s, while the full interval would require the removal of a colossal 5×10^{10} kg of water in 2.5 year—a cumulative mass (and a resulting average rate) more than an order of magnitude greater.

Figure 9 shows the amount of water removed (kg) per m^3 of methane produced, reported as cumulative quantities vs. time. Again, we see that massive amounts of water are moved to

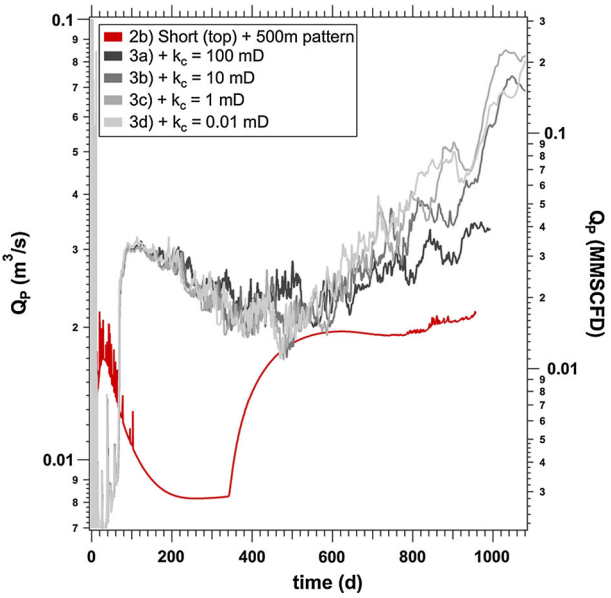

Fig. 7 Rate of methane production at the well, Q_P, for the vertical well scenarios Case *2b* and *3a–d*. Note the log scale

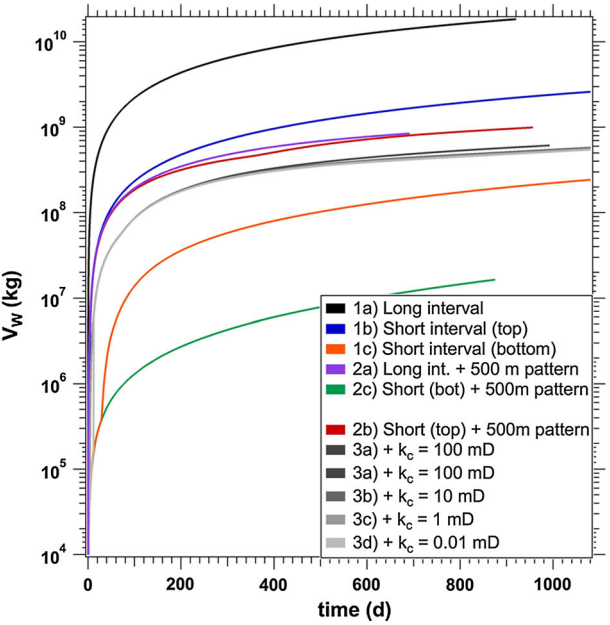

Fig. 8 Cumulative mass of water produced at the well, V_W, for all vertical well cases

achieve even the lowest rates of dissociation and production, even for cases where the communicating water-filled levees have extremely reduced permeability, although water production decreases for all cases, as is typical in production from hydrates. Case 1b, which seemed

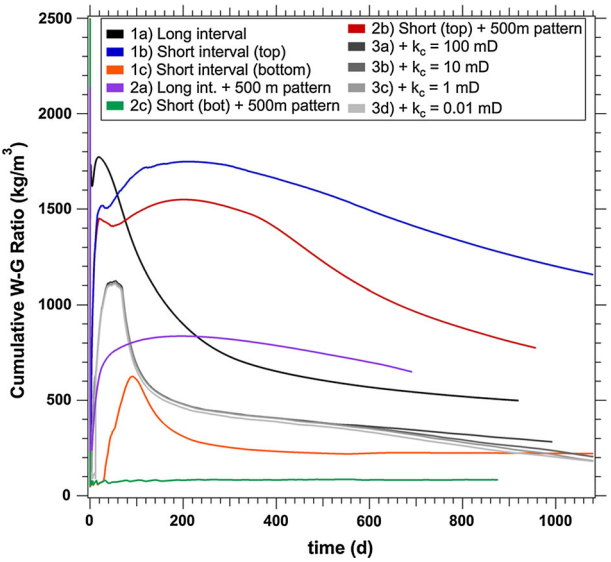

Fig. 9 Ratio of total water production to total methane production (cumulative at time t) for all vertical well scenarios

to have promising production rates, is among the worst in terms of water-gas ratio. For the best-case scenarios (low-permeability Cases 3a–d as well as the low-production Case 1c), 200 kg of water is removed per m^3 of methane after 3 years of production, a value that exceeds the solubility of methane in water under reservoir conditions. This suggests that mobile free gas is not being produced in significant quantities via depressurization, and that the bulk of produced methane results from dissolution of methane directly from solid hydrate via contact with undersaturated water, and the transport of that methane to the well via the aqueous phase.

Figure 10, showing only production of methane in the gas phase excluding aqueous transport, confirms this problem, particular when compared to total production (gas + aqueous phase) in Fig. 7. Only wells in a 500 m pattern produce the bulk of their produced methane in the gas phase. For Cases 1a, b, the production of free gas is orders of magnitude less than gas transport in the aqueous phase. It is clear that the presence of large amount of water in this hydrate reservoir hampers effective depressurization and dissociation. This effect has been seen clearly before, where simulations of Class 2 systems strongly suggest that impermeable boundaries are crucial (Moridis et al. 2009). These effects may also explain the poor performance of Cases 1c and 2c, which is particularly surprising given that deeper, warmer deposits are nearly always easier to dissociate via depressurization due to proximity to the hydrate stability boundary. However, here the lowest hydrate lens are actually relatively small in extent, bounded closely by impermeable layers, with a likely connection to a nearby aquifer. Thus, when producing from the full system, the dissociation/production behavior, we see the effect of free water flow disrupting depressurization, and when producing with close-in no-flow boundaries, we see simply a very small hydrate lens, surrounded by relatively impermeable rock, that is, quickly exhausted.

Three-dimensional visualization of such large systems is a challenging problem. This is partly due to the sheer size of the datasets, but also due to the lack of face/edge data to define the volume elements. The finite-volume meshes used by the TOUGH family of codes

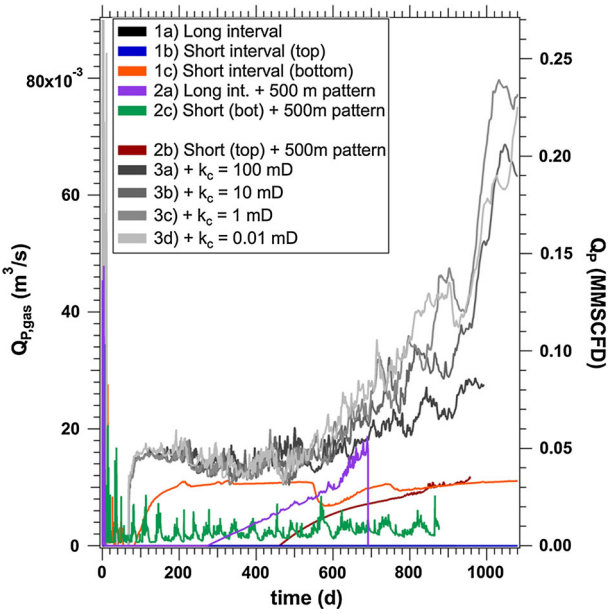

Fig. 10 Rate of methane production at the well in the gas phase only, $Q_{\mathrm{P,gas}}$, for all vertical well scenarios

allow for unmatched flexibility in defining the system (for example, the mix of rectilinear and unstructured grids used here), but also require extensive post-processing to generate visualization meshes. Interpolating the data onto a regular 3D grid is the simplest option, but the huge range of length scales in this problem (cm to tens of meters) results in an unreasonably large visualization problem of tens or hundreds of millions of volume elements. Proper tetrahedralization of an unstructured TOUGH finite-volume mesh is still an active research problem. For this work, we, therefore, use our knowledge of the system to perform limited 2D visualization of the near-well zone for Case 2b (production from the top hydrate lens under the assumption of a 500-m well pattern), using the open visualization package VisIt (2013). We choose to examine this Case in particular in order to understand why a higher interval (presumably further from the hydrate stability limit) exhibits the highest dissociation rates. We select an $x - z$ slicing plain at the y-coordinate of the well, also aligning the slice with a radian of the cylindrical near-well mesh to avoid the necessity to first render 3D volumes before slicing. We then linearly interpolate (using built-in VisIt functions) the data onto a uniform grid to generate color plots of various simulated properties.

Color plots of pressure and hydrate saturation (S_H) at the beginning and end of production for Case 2b are shown in Figs. 11 and 12. The production interval is indicated by the black line segment. In Fig. 11, we see the evolution of reservoir depressurization, with extensive depressurization seen after only 1 day of production along the highly permeable layer adjacent to the production interval. Nearly three years later, at $t = 960$ d, the depressurized zone has expanded to multiple hydrate and water layers, indicating communication between layers and effectiveness of depressurization.

However, this depressurization is not achieving extensive hydrate dissociation, as seen in Fig. 12. Looking at the region of the second panel ($t = 960$ d) labeled A and comparing it to the same location at $t = 1$ d, limited hydrate dissociation is apparent, but not occurring at a

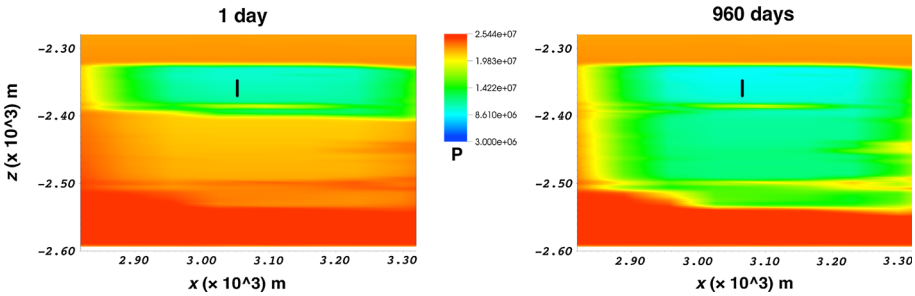

Fig. 11 Color plot of pressure, P, for a 2D ($X - Z$) slice through the domain at the well (production interval shown as *black line*) at the beginning and end of production

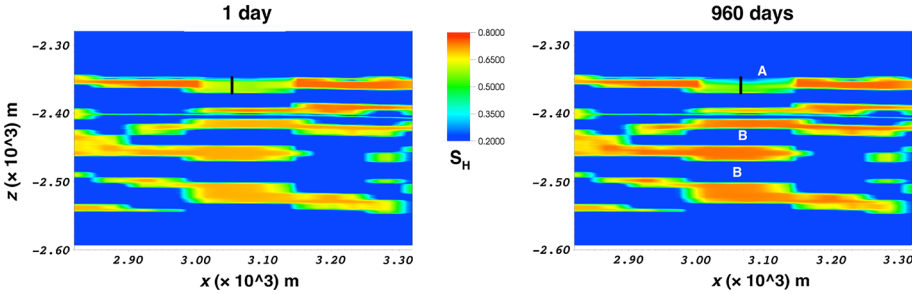

Fig. 12 Color plot of hydrate saturation, S_{H}, for a 2D ($X - Z$) slice through the domain at the well (production interval shown as *black line*) at the beginning and end of production

large scale (note that this plot spans the full 500 m width of the reduced system). Performing a similar visual comparison at the points labeled B in the second panel reveal more surprising behavior. The darkening of the color, indication increased hydrate saturation, shows that additional hydrate is *forming* at some distance from the well. At the estimate temperature and pressure conditions of this system, (1) depressurization alone is not enough to move a large region of hydrate below the stability curve, and (2) methane-bearing water is being pulled into hydrate-bearing zones, and methane in excess of the saturation limit at the local $T - P$ conditions is forming additional hydrate. As Case 2b and its derivatives show the most *favorable* dissociation behavior of the studied scenarios, it is clear that the production difficulties may be insurmountable.

3.2 Production from a Horizontal Well

For many Class 2 and Class 3 hydrate deposits, it has been demonstrated that horizontal wells, properly located, can be a far more effective production strategy (Moridis 2008b). Therefore, we simulate production from the horizontal well described in Sect. 2.1, with all reservoir parameters otherwise unchanged from the vertical well cases.

Evolution of the rate of hydrate dissociation/methane release (Q_R) and the rate of methane production at the well (Q_{P}) are provided in Fig. 13. We see that release exceeds production from the beginning, which is necessary to generate free gas in the reservoir for later production (Moridis and Reagan 2007a, b). However, after less than 5 months of production both release and production cease to increase with time, while water production (Q_{W}) remains roughly constant, at rates comparable to the best-case vertical well scenario (with greater

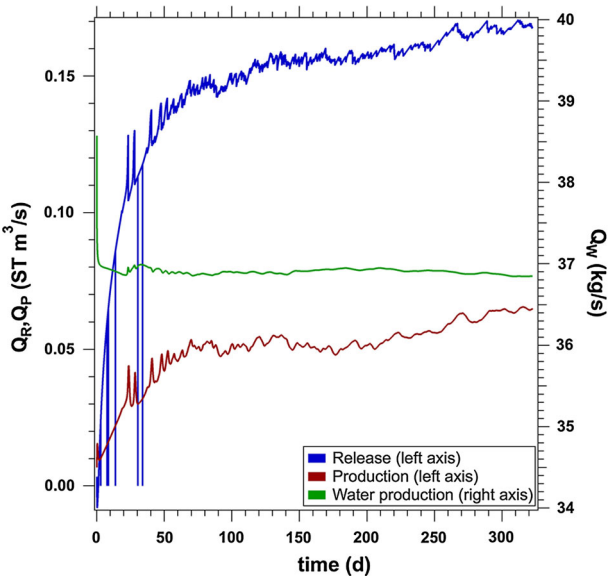

Fig. 13 Rate of hydrate dissociation and methane release (Q_R), rate of methane production at the well (Q_P), and rate of water production at the well (Q_W, right axis) for the horizontal well scenario

access to a larger volume of the reservoir via a single well). This suggests that once again, effective depressurization is being hindered by the inflow of large quantities of water from the surrounding formation. As economically viable production rates for offshore wells need to reach orders of millions of cubic feet per day (0.12 – 1.2 m^3/s), this system is evolving unsustainably.

Cumulative quantities of methane, released and produced, are shown in Fig. 14. More significant, however, is the water-to-gas ratio, shown in green. While this ratio drops rapidly as production proceeds, the curve is heading toward asymptotic behavior at the values of 700–800 kg per m^3, which suggests again that transport of methane is entirely in the aqueous phase, and even then, that the quantity of water produced greatly exceeds that quantity of water needed to dissolve and transport the methane seen at the well.

Figure 15 provides two color plots, at earlier and late times, of a close-in 2D slice of the horizontal well system oriented along the plane of the horizontal well. These plots have been generated by bespoke python-based visualization scripts using matplotlib (Hunter 2007) for rendering (rather than Visit), and are interpolated onto a uniform rectilinear mesh using a nearest-neighbor interpolation scheme to generate sharply defined voxels. Examination of the color plots indicates that there has been extensive early hydrate dissociation focused at the heel (left) and toe of the wellbore, as well as disappearance of hydrate at several locations around the wellbore. The removal of hydrate at these locations is allowing relatively unrestricted flow of water from the surrounding aquifer, hindering depressurization and preventing further dissociation of hydrate. Thus, the formation of free gas is greatly hindered after only a few months of production, and the problems of high water production and poor gas production seen in the vertical well cases reappear. Note also the darkening of many of the gridblocks within the near-well region as other gridblocks show dissociation. Methane-bearing water is being drawn into the hydrate formation from surrounding water-bearing zones, and hydrate saturations have increased slightly due to hydrate formation.

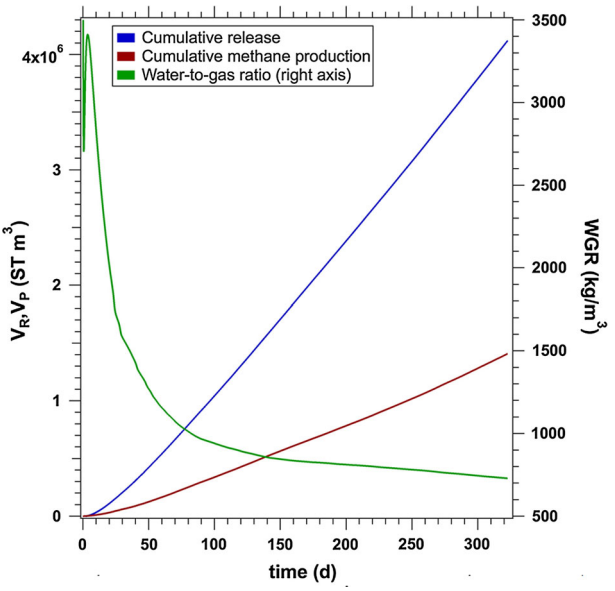

Fig. 14 Cumulative methane release (V_R), cumulative methane production (V_P), and ratio of total water production to total methane production (V_W, right axis), for the horizontal well scenario

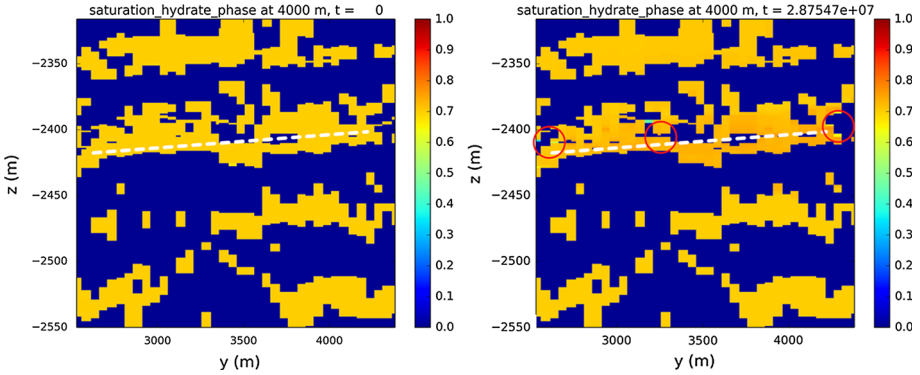

Fig. 15 Color plot of hydrate saturation, interpolated onto a rectilinear visualization mesh, for a $Y - Z$ slice of the horizontal well system at $x = 4000$ m at $t = 1$ d and $t = 330$ d. The rough location and extent of the horizontal well is indicated by the *white dashed line*, with the heel of the well at the left. Locations of significant hydrate dissociation are *circled*

4 Conclusions

This study has been the first supercomputer-driven 3D simulation of production in a large oceanic hydrate reservoir, with the deployment of thousands of processors allowing both individually large simulations and multiple simulations performed in parallel to evaluate a range of production strategies. While the simulations are the largest to date, the insights found here strongly relate to earlier studies on the productivity and non-productivity of various configurations of hydrates.

Earlier work has strongly suggested that impermeable boundaries (i.e., impermeable enough to allow strong depressurization of the reservoir through conventional technology) are the key to efficient and economical production (Moridis and Reagan 2007a, b; Moridis et al. 2009, 2011b). While this reservoir (as modeled) is indeed bounded on a large scale (the possible connection to a larger aquifer was not simulated), the layers of high-saturation hydrate are adjacent to, and effectively surrounded by, levees and channels of relatively high-permeability, water-saturated media. Even when those levees are reduced in permeability, creating bounded, Class 2-like systems, the sheer quantity of available water makes depressurization difficult on either a reservoir scale or a local scale, even when assuming a regular pattern of wells and short perforation intervals isolated within individual hydrate-water layers. Examination of the simulation output also indicates that the achieved pressure drops only move the system to the edge of the hydrate stability curve at reservoir temperatures, and therefore, liquid-gas-hydrate equilibrium is achieved at each layer, but large quantities of free gas do not appear, and the system is not driven fully into the zone of liquid-gas coexistence. As a result, we see (1) poor rates of hydrate dissociation, (2) poor rates of methane production, and (3) high water production in both absolute terms and relative to methane production.

Previous work found that horizontal well, properly placed, can mitigate some of these problems as long as the horizontal well stays within the hydrate layer, and does not connect to surrounding water-filled media (Moridis 2008b). Our results using this configuration within the 3D reservoir domain show that large quantities of water are still able to reach the well early in the depressurization process, reducing depressurization effectiveness. As a result, it is becoming clear that systems with large quantities of water in communication with the hydrate reservoir, even if not connection to aquifers beyond reservoir boundaries, are particularly challenging production targets. Permeability of the bounding media, as well as distance from the hydrate stability boundary and configuration of the hydrate into large contiguous masses, need to be considered in reservoir evaluation.

It is somewhat disappointing that simulations of this size and scale (and with the corresponding effort) result in "lackluster" results—in terms of active hydrate dissociation, gas formation, and production rates. However, characterization of these systems in the field is challenging, and the sensitivity studies suggest that a system a few degrees warmer, with lower-permeability boundaries, or perhaps with thicker and more extensive hydrate layers would perhaps be significantly more productive. Issues such as these should be investigated in upcoming field tests. It is also important to note that these are the largest TOUGH simulations to date, with the simultaneous solution of 9,000,000 equations at each step, plus fully described multiphase hydrate thermodynamics. The simplified geometries and reduced extent often used to study hydrate-bearing systems may in fact be limiting our understanding of real hydrate reservoirs at the scales and geometries seen in the field. Therefore, systems of increasing size and heterogeneity need to be investigated. It is also important to highlight the role of MPI-based, massively parallel simulation codes in this process. Faster individual processors are no longer driving computational performance gains, and multithreading across many cores is not sustainable in terms of providing the necessary resources for ever-larger simulations (Sutter 2005). The development of distributed-memory parallel codes for reservoir simulation (a process that is more complex than simply "adding parallelism" to an existing simulator) is likely to increase in urgency in the near future.

The simulations also show that mesh complexity becomes critical for large and heterogeneous systems, as the "layered" and mainly rectilinear and cylindrical vertical well mesh (1.66 M elements) can be realistically run on a small departmental cluster (250 processors), while the more complex Voronoi grid used in the horizontal well simulations (2.46 M ele-

ments, plus more connections per gridblock) required over a million processor-hours of time to complete the assessment. The development of user-friendly meshing tools and perhaps active mesh refinement and de-refinement may be needed to move TOUGH simulations fully into the massively parallel universe. Finally, the traditional TOUGH family visualization methods, mainly via standard outputs or ASCII-based, Tecplot-formatted files, are insufficient to handle datasets of this size. Preliminary visualization of pT+H using Visit (https://wci.llnl.gov/codes/visit/) has been moderately successful, however, no suitable tools for point-and-click visualization of arbitrary polyhedra are available. Custom visualization tools are still necessary for particularly complex geometries. Key hurdles include finding and implementing modern data formats, and converting the TOUGH finite-volume meshes into forms more compatibility with high-performance visualization software.

Acknowledgments This research was funded by Statoil ASA. This research used resources of the National Energy Research Scientific Computing Center, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

Boswell, R., Moridis, G.J., Reagan, M.T., Collett, T.S.: Gas hydrate accumulation types and their application to numerical simulation. In: Proceedings of 7th International Conference on Gas Hydrates, Edinburgh, Scotland, UK, 17–21 July 2011

Burwicz, E.B., Rupke, L.H., Wallmann, K.: Estimation of the global amount of submarine gas hydrates formed via microbial methane formation based on numerical reaction-transport modeling and a novel parameterization of Holocene sedimentation. Geochim. et Cosmochim. Acta **75**(16), 4562–4576 (2011)

Hunter, J.D.: Matplotlib: a 2D graphics environment. Comput. Sci. Eng. **9**(3), 90–95 (2007)

Klauda, J.B., Sandler, S.I.: Global distribution of methane hydrate in ocean sediment. Energy Fuels **19**, 459 (2005)

Makogon, Y.F.: Hydrates of Hydrocarbons. Penn Well Publishing Co., Tulsa (1997)

Milkov, A.V.: Global estimates of hydrate-bound gas in marine sediments: how much is really out there? Earth Sci. Rev. **66**, 183 (2004)

Moridis, G.J., Reagan, M.T.: Gas Production From Oceanic Class 2 Hydrate Accumulations, OTC 18866. In: Proceedings of 2007 Offshore Technology Conference, Houston, Texas, U.S.A., 30 April–3 May 2007a

Moridis, G.J., Reagan, M.T.: Strategies for gas production from oceanic class 3 hydrate accumulations, OTC 18865. In: Proceedings of 2007 Offshore Technology Conference, Houston, Texas, 30 April–3 May 2007b

Moridis, G.J., Kowalsky, M.B., Pruess, K.: TOUGH+HYDRATE v1.0 User's Manual: a code for the simulation of System behavior in hydrate-bearing geologic media, Report LBNL-0149E. Lawrence Berkeley National Laboratory, Berkeley, CA 2008a

Moridis, G.J., Reagan, M.T., Zhang, K.: The use of horizontal wells in gas production from hydrate accumulations. In: Proceedings of 6th International Conference on Gas Hydrates, Vancouver, BC, July 6–10 2008b

Moridis, G.J., Collett, T.S., Boswell, R., Kurihara, M., Reagan, M.T., Koh, C., Sloan, E.D.: Toward production from gas hydrates: current status, assessment of resources, and simulation-based evaluation of technology and potential. SPE J. **12**(5), 745–771 (2009)

Moridis, G.J., Reagan, M.T., Boyle, K.L., Zhang, K.: Evaluation of the gas production potential of some particularly challenging types of oceanic hydrate deposits. Transp. Porous Media **90**, 269–299 (2011a). doi:10.1007/s11242-011-9762-5

Moridis, G.J., Collett, T.S., Pooladi-Darwish, M., Hancock, S., Santamarina, C., Boswell, R., Kneafsey, T., Rutqvist, J., Kowalsky, M.J., Reagan, M.T., Sloan, E.D., Sum, A.K., Koh, C.: Challenges, uncertainties and issues facing gas production from hydrate deposits in geologic systems. SPE Res. Eval. Eng. **14**(1), 76–112 (2011b)

Pan, L.: User's Information for WinGridder V3.0., LBNL Report 273E. Lawrence Berkeley National Laboratory, Berkeley, CA (2008)

Reagan, M.T., Moridis, G.J., Zhang, K.: Sensitivity analysis of gas production from class 2 and class 3 hydrate deposits. LBNL-01657E, OTC 19554, In: Proceedings of 2008 Offshore Technology Conference, Houston, Texas, USA, 5–8 May 2008

Rycroft, C.H.: Voro++: a three-dimensional voronoi cell library in C++. Chaos **19**, 041111 (2009)
Sloan, E.D., Koh, C.: Clathrate Hydrates of Natural Gases, 3rd edn. Taylor and Francis Inc, Boca Raton (2008)
Sutter, H.: The free lunch is over: a fundamental turn toward concurrency in software. Dr. Dobbs J. **30**, 3 (2005)
van Genuchten, : A closed-form equation for predicting the hydraulic conductivity of unsaturated soils. Soil Sci. Soc. **44**, 892 (1980)
Visit https://wci.llnl.gov/codes/visit/ (2013)
Zhang, K., Moridis, G.J., Wu, Y.S., Pruess, K.: A domain decomposition approach for large-scale simulations of flow processes in hydrate-bearing geologic media. In: Proceedings of 6th International Conference on Gas Hydrates, Vancouver, BC, July 6–10 2008