

THE UNIVERSITY OF CALGARY

Exploring Immersive Virtual Environments for Well Placement Optimization in Reservoir Models

by

Roberta Cabral Ramos Mota

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF MASTER OF SCIENCE

GRADUATE PROGRAM IN COMPUTER SCIENCE

CALGARY, ALBERTA

JANUARY, 2017

© Roberta Cabral Ramos Mota 2017

ABSTRACT

Immersive virtual environments have been considered promising mediums to attend specific demands from the oil and gas industry. In this thesis, we explore immersive technologies' benefits for the execution of tasks associated with well placement optimization. We present a) an analytical method to perform static connectivity analysis as a proxy for flow simulation, b) an application to support well optimization using our method, and c) an exploration of our application in three immersive environments – a CAVE with a tracked gamepad; a HMD with a tracked gamepad; and a HMD with a leap motion controller – in the search for visualization and interaction techniques that facilitate well placement studies. Based on primary study conducted with reservoir engineers, we provide an examination of the usefulness of our application. We also discuss our findings considering engineers' preferences as well as the suitability of the different immersive environments for designing and assessing well placement scenarios.

PUBLICATIONS

Some of the materials, ideas, and figures in this thesis have previously appeared in the following publications:

- R. C. R. Mota, S. Cartwright, E. Sharlin, H. Hamdi, M. C. Sousa and Z. Chen, “Exploring Immersive Interfaces for Well Placement Optimization in Reservoir Models”, Proceedings of the ACM Symposium on Spatial User Interaction (SUI'16). 2016. Tokyo, Japan.
- R. C. R. Mota, H. Hamdi, M. C. Sousa, E. Sharlin and Z.X. Chen, “A Visual Framework for Reservoir Connectivity Analysis”, Proceedings of the 78th European Association of Geoscientists & Engineers Conference and Exhibition. 2016. Vienna, Austria.

ACKNOWLEDGEMENTS

During the last two years, I was part of three research groups and each of them helped me choose the theme of this thesis and carry on this research. I thank Dr. Mario Costa Sousa for introducing me a new domain in which I could make use of Computer Science and, specifically, of Virtual Reality. Even though working on oil and gas was challenging at first, the journey has been rewarding. I also extend the thanks to the Illustrares Research Group; in particular, to Julio Daniel, Clarissa Codá, Allan Rocha, and Felipe Carvalho. Hamidreza Hamdi, thanks for all the long talks and all the feedback.

I also thank Dr. Zhangxing (John) Chen for the opportunity to be part of the Reservoir Simulation Research Group and to work at the Collaboration Centre, where I spent most of my time; having access to all the technology it offers was essential to advance this project. I take this opportunity to thank the Collaboration Centre's staff, particularly Stephen Cartwright and Emil Selin. To Dr. Ehud Sharlin, thank you for being so welcoming, for the insightful meetings, and for supervising me in projects on areas that I find to be incredible: HCI and HRI. Finally, I thank uTouch Research Group's colleagues Sowmya Somanath, Nico Li, and Ahmed Mostafa.

Um agradecimento com saudade a Fernando Mota, Fátima Guedes, Cristina Guedes e Renan Mota. Especialmente aos meus pais, agradeço por todo o esforço investido na minha educação, pelos incentivos constantes e pelo suporte nas horas difíceis. Agradeço de coração a Francisco, Rinalda e Marília Ramos por todo o apoio, por me sentir família e por tudo que isso representa. Por fim, agradeço a Mariana Mota por acreditar em mim e por sempre me incentivar a buscar mais; por seu sorriso, seu abraço e seu olhar diariamente me lembrarem que, apesar dos desafios inerentes à vida, nela eu encontro a quietude de um amor resolutivo.

CONTENTS

| | |
|--|-----------|
| 1. INTRODUCTION | 1 |
| 1.1. IMMERSIVE VIRTUAL ENVIRONMENTS | 4 |
| 1.2. STATIC CONNECTIVITY ANALYSIS..... | 6 |
| 1.3. GOALS..... | 7 |
| 1.4. METHODOLOGY | 8 |
| 1.5. CONTRIBUTIONS | 10 |
| 1.6. THESIS OVERVIEW..... | 10 |
| 2. IMMERSIVE VIRTUAL ENVIRONMENTS | 13 |
| 2.1. THE REALITIES | 14 |
| 2.2. THE VIRTUAL REALITY SYSTEM | 15 |
| 2.3. THE INTERFACE TO THE VIRTUAL WORLD – INPUT..... | 16 |
| 2.3.1. HAND INPUT DEVICES..... | 17 |
| 2.3.2. NON-HAND INPUT DEVICES | 20 |
| 2.4. THE INTERFACE TO THE VIRTUAL WORLD – OUTPUT | 24 |
| 2.4.1. VISUAL DISPLAYS | 24 |
| 2.4.2. AURAL DISPLAYS..... | 27 |
| 2.4.3. HAPTICS DISPLAYS | 29 |
| 2.5. THE INTERFACES TO THE VIRTUAL WORLD – IMMERSION..... | 32 |
| 2.5.1. PRESENCE..... | 33 |
| 2.5.2. IMMERSION..... | 34 |
| 2.5.3. BENEFITS OF IMMERSION..... | 35 |
| 3. OIL AND GAS EXPLORATION, DEVELOPMENT, AND PRODUCTION | 42 |
| 3.1. OIL AND GAS E, D&P STAGES | 43 |
| 3.1.1. EXPLORATION STAGE..... | 44 |
| 3.1.2. DEVELOPMENT STAGE | 44 |
| 3.1.3. PRODUCTION STAGE..... | 52 |
| 3.2. WELL PLACEMENT OPTIMIZATION | 55 |
| 3.2.1. USING DYNAMIC RESERVOIR SIMULATION | 56 |
| 3.2.2. USING STATIC CONNECTIVITY ANALYSIS | 58 |

| | | |
|-----------|--|------------|
| 3.2.3. | USING OUR STATIC CONNECTIVITY ANALYSIS | 73 |
| 3.2.4. | EVALUATING LOCAL CONNECTIVITY METRICS..... | 84 |
| 4. | IVES IN OIL AND GAS E, D&P | 94 |
| 4.1. | FOR INCREASED INFORMATION SPACE | 95 |
| 4.1.1. | FOR DATA INTEGRATION | 95 |
| 4.1.2. | FOR COLLABORATION..... | 98 |
| 4.2. | FOR INCREASED SPATIAL UNDERSTANDING | 102 |
| 4.3. | FOR MORE NATURAL INTERACTIONS..... | 104 |
| 4.4. | FOR INCREASED SENSE OF PRESENCE..... | 109 |
| 4.5. | FOR WELL PLACEMENT OPTIMIZATION | 112 |
| 5. | IVES FOR WELL PLACEMENT OPTIMIZATION | 123 |
| 5.1. | HARDWARE SPECIFICATIONS..... | 124 |
| 5.2. | SOFTWARE SPECIFICATIONS..... | 125 |
| 5.3. | INTERACTIONS USING GAMEPAD | 125 |
| 5.3.1. | GENERIC | 125 |
| 5.3.2. | DOMAIN-SPECIFIC | 130 |
| 5.4. | INTERACTIONS USING LEAP MOTION..... | 149 |
| 5.4.1. | GENERIC | 149 |
| 5.4.2. | DOMAIN-SPECIFIC | 158 |
| 6. | IVES FOR WELL PLACEMENT OPTIMIZATION: AN EVALUATION | 162 |
| 6.1. | EVALUATION SESSION | 163 |
| 6.2. | ON THE APPLICATION..... | 163 |
| 6.3. | ON THE IMMERSIVE INTERFACES | 164 |
| 6.4. | PARTICIPANTS | 165 |
| 6.5. | DATASET | 166 |
| 6.6. | RESULTS..... | 166 |
| 6.6.1. | ON THE APPLICATION..... | 167 |
| 6.6.2. | ON THE IMMERSIVE INTERFACES | 171 |
| 6.7. | DISCUSSION..... | 173 |
| 6.7.1. | ON THE APPLICATION..... | 173 |
| 6.7.2. | ON THE IMMERSIVE INTERFACES | 174 |
| 6.8. | FUTURE WORK..... | 176 |

| | | |
|-------------------|--|------------|
| 6.8.1. | ON THE STATIC CONNECTIVITY ANALYSIS..... | 176 |
| 6.8.2. | ON THE APPLICATION..... | 178 |
| 6.8.3. | ON THE IMMERSIVE INTERFACES | 183 |
| 7. | CONCLUSION..... | 184 |
| 7.1. | CLOSING REMARKS..... | 187 |
| APPENDIX A | | 189 |

LIST OF FIGURES

Figure 1: The same reservoir model being displayed in two immersive virtual environments: a CAVE and a regular desktop system. 1

Figure 2: Examples of novel visualization techniques in oil and gas domain. From left to right: decal-maps for multidimensional data analysis, SBM for 3D geological modeling, and cutaway views for investigating internal reservoir cells. 5

Figure 3: Birdly [13] is a full-body simulator ride that uses multi-hinged flags, head tracking, and motion feedback to evoke the sense of flying. In addition, there is a fan in front of the user to simulate headwind which gives a sense of flying speed. 13

Figure 4: The virtuality continuum – adapted from [14]. 14

Figure 5: A virtual reality system comprises an input interface, a virtual world, and an output interface – adapted from [15]. 15

Figure 6: Examples of world-grounded devices developed in the early 1990’s. (a) A person uses a real bulldozer steering wheel on the left that meets the virtual steering wheel on the right; (b) Disney Imagineering designed a steering mechanism to give guests the illusion of flying a magic carpet as in the film Aladdin; (c) surgeon practices a surgical procedure using an artificial cavity together with the correct instrument set coupled with force-feedback devices; (d) a researcher uses a force-feedback device to control the tip of a scanning tunneling microscope probe, exploring the surface of a sample of atoms. 17

Figure 7: Examples of hand-held controllers. Oculus Touch [18], Sony PS3 Move [19], THE VOID Nova IV Gun [20], and WorldViz Wand [21] are tracked controllers – (a), (b), (c), and (d), respectively. MOGA Pro Controller [17] is a non-tracked gamepad – (e). 18

Figure 8: CyberGlove (a) and Myo (b) are hand-worn input devices. The leap motion (c) is a bare-hand input device. 20

Figure 9: Oculus Rift [28] uses sensor fusion and Constellation tracking to provide head-tracking input – (c). FOVE [29] uses infrared sensors for eye tracking – (a, b, d). By using eye-tracking, the FOVE system allows two innovative features: eye contact monitoring and foveated rendering. Foveated rendering refers to loading high-resolution textures where users are actually looking, and optimizing out-of-focus areas with lower-resolution or fewer effects. 21

Figure 10: PriorVR suit [32] places inertial sensors onto the user’s body, while Microsoft Kinect uses optical-based sensors placed on an add-on peripheral to provide full-body tracking input. – (a, b) and (c), respectively. Cyberith Virtualizer [30] is a treadmill that uses motion sensors in its base plate, pillars, and ring to provide foot tracking input – (d)..... 23

| | |
|---|----|
| Figure 11: The EVL CAVE (a) and ImmersaDesk (d) are projection-based visual displays, while the ActiveWall (b) and ActiveDesk (c) are monitor-based displays designed by Virtualis [39]. | 25 |
| Figure 12: On the top row, HTC Vive [40], Samsung Gear VR [41], and Google Cardboard [42] are non-see-through HMDs. On the bottom row, Vuzix M100 [43] and NVisor ST50 [44] are OST HMDs, while Visette45 [45] is a VST HMD. | 27 |
| Figure 13: Types of aural displays: standalone and built-in headphones (left, middle) and loudspeakers (right). | 28 |
| Figure 14: The 'pit room' using ledges as passive haptic feedback. | 30 |
| Figure 15: Dexmo (a) provides force feedback, while Null Space (b) and KOR-FX (c) suits simulate tactile feedback. | 31 |
| Figure 16: Widerun is compatible with any kind of bike. It gets inputs both from the game software and from the bike trainer to regulate back the electrical signals accordingly to the virtual experience the user is having, like steering degree, speed magnitude and ground resistance. Widerun is compatible with Oculus Rift and Samsung Gear VR. | 32 |
| Figure 17: The continuum extended into a two-dimensional plane of virtuality and immersion. | 35 |
| Figure 18: Some application domains rely on the realistic experience that VR systems may provide to the user, such as military and medical training (top) and phobia therapy (bottom). | 37 |
| Figure 19: Climate data running in TESSIN VISLab – images from [70]. | 39 |
| Figure 20: Illustrations of spatial interactions for spatial construction – images from [72]. | 40 |
| Figure 21: Oil and gas exploration, development, and production stages along with the various disciplines and activities involved. | 42 |
| Figure 22: Seismic studies and exploration wells are some of the efforts performed in the exploration stage. | 44 |
| Figure 23: The development stage includes procedures such as geological and simulation modelling | 45 |
| Figure 24: Sequential steps of a conventional workflow of geological reservoir modelling. | 46 |
| Figure 25: Three geological realizations displaying the distribution of facies property. Even though there are significant differences between the realizations, parts of the reservoir where real well log data is available are honoured. | 48 |
| Figure 26: Upscaling is a process that converts a fine-scale model into a coarse-scale one. | 50 |
| Figure 27: History matching procedures and drilling of development wells are some of the efforts performed in the production stage. | 52 |

| | |
|--|----|
| Figure 28: A common production profile of an oil field development. | 53 |
| Figure 29: Reservoir recovery techniques: primary (a), secondary with water injection (b), and tertiary with steam injection (c)..... | 54 |
| Figure 30: Illustrations of the most common trajectories of wells: conventional wells, which are vertical or slightly deviated, and horizontal wells, which have inclination greater than 80-85 degrees. Horizontal wells can be further classified into multilateral, MRC, fishbone, and extended reach wells – figure redrawn from [73]. | 56 |
| Figure 31: Fluid flow simulation results for water saturation: (a) 80 days, (b) 200 days, (c) 400 days, (d) 500 days - images from [89]. | 56 |
| Figure 32: On the left, figurative geobodies representing individual channel deposits (yellow) and an amalgamated channel deposit (green). Geobody connectivity is defined as the proportion of the volume of the biggest geobody (green) to the sum of all geobodies (green and yellow). On the right: reservoir-to-well connectivity is calculated as the proportion of volume of geobodies connected to one or more wells to the sum of all geobodies - figures redrawn from [112]. | 60 |
| Figure 33: Three examples of connected clusters colored in red for three different pc values. For some specific value of pc, the percolation cluster spans across the entire section – figure from [103]. | 61 |
| Figure 34: Connectivity as a function of time refers to, at any given time, fluid may be actively flowing only in fractions of channel deposits. Illustrations depict influence of time in four scenarios: (a) meandering channel producing by depletion, (b) straight channel produced by water displacement with one injector and one producer, (c) tortuous channel produced under the same conditions as for previous channel, and (d) straight and meandering channel produced together. | 63 |
| Figure 35: Illustrations of producing reservoirs using water flooding: (a) pattern water flood, and (b) early water breakthrough during pattern water flooding due to the presence of a high-permeability layer – figure redrawn from [120]. | 65 |
| Figure 36: Illustrations of juxtaposition styles – figure from [122]. | 65 |
| Figure 37: Effects of faults on reservoir connectivity – figure modified from [123]. | 66 |
| Figure 38: During the quality factor analysis, sand body cells of the models are examined in concentric cylindrical shells around a well insertion point, moving incrementally outwards – figure from [124]. | 67 |
| Figure 39: In this illustration, although the lengths of the straight lines connecting (c2, wp1) and (c3,wp1) are the same, the geometric average perms along these lines are different because of the shale zone. Thus, c2 and c3 provide different weights when calculating Qs - figure modified from [125]. | 68 |
| Figure 40: A schematic scenario of the fast-marching algorithm – figure redrawn from [128]. | 70 |
| Figure 41: Our workflow for performing static connectivity analysis. | 75 |

| | |
|--|----|
| Figure 42: The three main components of corner-point geometry: corners, pillars, and cells – figure from [9]. | 76 |
| Figure 43: Faults (left) and pinch-out cells (right) are typical structural irregularities incorporated in corner point grids – figure modified from [9]. | 76 |
| Figure 44: Illustrations of (a) regular cells with full contact interface and (b) corner-point cells with partial contact. | 77 |
| Figure 45: Pseudocode of Dijkstra’s algorithm. | 80 |
| Figure 46: Dijkstra’s algorithm has a propagation criteria that finds the shortest paths for all nodes that can be reached from the perforation node (in black). | 82 |
| Figure 47: Shortest path calculation along with minimum cost (delta time) parameter associated with nodes. | 83 |
| Figure 48: Illustrative example of shell factors providing a notion of time. | 85 |
| Figure 49: Illustration showing information loss as a side effect of an averaging procedure. On the left, original property values per grid cell; on the right, average property values per shell. | 87 |
| Figure 50: On the left, a theoretical balloon-shaped SAGD steam chamber; on the right, a more realistic, irregular chamber – figure from [138]. | 88 |
| Figure 51: In graph-based algorithms like Dijkstra, paths are restricted to graph edges in way that a vertex is updated from another vertex. | 90 |
| Figure 52: On the left, Euclidean distance computed with a 4-connectivity breadth-first search algorithms gives square level sets. On the right, the distance is computed with fast marching, giving almost circular level sets – figure from [142]. | 90 |
| Figure 53: On the left, Pollock tracing for Cartesian cell; on the right, a quadrilateral is mapped to a unit cell through an isoparametric transform – figures from [145]. | 91 |
| Figure 54: The fast marching simulates a front advancing from a source point and stores the computed distance values at the vertices (left). Yet, as the front arrives at an obtuse vertex v3, only the vertex v1 is already processed and has a valid distance value (right). Therefore, special handling is needed for this triangle. | 92 |
| Figure 55: Examples of single display systems utilized for analyzing large, multidimensional data in oil and gas domain. | 94 |
| Figure 56: Examples of different interactions for a multi-surface environment. On the top row, the “pouring interaction” transfers data from the mobile device to the tabletop; and the “camera interaction” captures tabletop data from camera images on the mobile device. On the bottom row, by placing a tablet device on top | |

of the tabletop, analytical and graphical data are automatically transferred to the tablet (a), (b), and (c), and by placing the tablet down vertically, the “slicing interaction” shows seismic data in the same orientation (d)... 98

Figure 57: Examples of techniques for assessing reservoir models in tabletop systems. From left to right: a method to probe individual cells and display property values; “splitting” and “peeling” gestures to view the inside of the reservoir while maintaining context; and a touch tap to select a well and to activate a cutaway view that removes cells occluding the chosen well path..... 99

Figure 58: On the left, SeisTablet integrates a shared monitor with tablets that support pen and touch interaction for seismic interpretation. On the middle and right, Slice Inspector combines a 2D multi-touch interface with a 3D large display for the analysis of seismic data..... 100

Figure 59: On the top row, the system allows remote experts to add waypoints and to highlight objects to guide novices situated in an immersive VE [166]. On the bottom row, CEE system give users the ability to annotate VEs and discuss them in real time (left), and Inside Reality [167] allows collaborative sessions in which avatars represent remote participants in the VE (right). 101

Figure 60: From left to right: geologist observing the topographic expression of valleys draining into Lake Tahoe; locations of faults, folds, and drainages precisely mapped directly onto the topographic surface; and results from the comparison between Stereo Analyst (left) and the VR mapping application (right). 103

Figure 61: Images to illustrate how high degrees of immersion may increase spatial understanding during well planning: by viewing regions of interest from many perspectives and at different scales to see their interdependency; and by performing direct measurements of locations, orientations, and intersections of often complex 3D structures. 104

Figure 62: Examples of systems using world-grounded interfaces. On the left and middle, the Snakey object is used for defining well path configurations. On the right, TanGeoMS permits users to hand-build models using material such as foam and modeling clay. Color maps are then projected over the models according to their format. 105

Figure 63: Examples of systems using tracked hand-held controllers. Flying Frustum uses a tracked pen for defining paths for an UAV to fly along in the real field (a), (b), and (c). CubicMouse is a novel input interface that may be used for editing the trajectory of an oil well (d). 107

Figure 64: Example of a system using a multimodal interface. The user may create seismic structures using voice commands (a) and (b), and may use gestures to manipulate the seismic volumes (c) and (d)..... 109

Figure 65: Images of off-shore facilities created and simulated using AVEVA AVP. 110

Figure 66: Examples of robots available in SimUEP-Robotics. 110

Figure 67: JackVR allows experts to train landing jack up oil rigs. The software uses graphical interface elements to inform them of the status of the simulation..... 111

Figure 68: Different learning platforms for drilling and rig operations provided by Drilling Systems. 112

Figure 69: Steps to create well trajectories using CMG Suite: (a) select a new well type (injector or producer) and a cross-sectional view; (b) select the first perforation cell on the top most layer of the reservoir model, (c) if needed, change the plane to select the other perforation cells; and (d) visualize the new well in the 3D window. 113

Figure 70: Isosurfaces (right) and flow lines (left) are commonly used to interpret flow simulation results.... 114

Figure 71: On the left, the model is colored based on the propagation time forecasted from the fast-marching method. On the right, the model is filtered using a time threshold. 115

Figure 72: Illustrations of the iterative process for placing vertical wells onto connectivity maps to optimize well placement in a minimum number of steps. 116

Figure 73: Two snapshots of the software showing the four different views that compose the environment (top) and a tooltip window displaying recovery measures per distance on a particular inter-well path (bottom). ... 117

Figure 74: User manipulates a tangible device for defining a well path. By rotating the tangible prop, the user can navigate through different reservoir layers. 118

Figure 75: User uses translation and rotation widgets for defining a plane and, then, creates a well path along the plane. 119

Figure 76: The opening angle on the cutaway is increased/decreased by rotating the tangible device. 120

Figure 77: The lollipop-up technique offsets cells to reflect the same shape as the actual well trajectory (left), and the information circles of multiple wells ‘hung’ from a horizontal pole (right). 120

Figure 78: Our application to support well placement studies running in different immersive mediums. 123

Figure 79: Our application explores three immersive environments: a CAVE with a tracked gamepad (left), a HMD with a tracked gamepad (center), and a HMD with a leap motion device attached to it (right). 124

Figure 80: User filters the reservoir model based on permeability values through the graphical menu. 127

Figure 81: Controller setup for performing reservoir manipulation. 128

Figure 82: Depictions of different subsurface and surface surroundings of our offshore oil reservoir. 129

Figure 83: Illustration of geological formations and facilities of an offshore drilling platform. 130

Figure 84: Shading levels indicate focus based on three parameters: focal point (pf), near (dnear) and far (dfar) distances. 131

Figure 85: Placing emphasis when the user is outside (left) and inside (right) the reservoir model. 132

Figure 86: Different combinations of color and line effects for stylized focus. 133

| | |
|---|-----|
| Figure 87: Illustrations of an octree, frequently used for partitioning volumetric objects such as a reservoir model. | 134 |
| Figure 88: Transfer function used for regulating focus..... | 135 |
| Figure 89: Transfer functions used in our focus lens. | 136 |
| Figure 90: User probing grid cells. | 137 |
| Figure 91: The technique for defining the transparency lens consists in scanning fragments of the reservoir model to evaluate whether they lie within a cone-shaped volume..... | 138 |
| Figure 92: Procedure to define a vertical well path: user hovers a cell (a) and presses the controller's button to confirm creation (b and c)..... | 140 |
| Figure 93: Process for defining a free-form well trajectory. | 141 |
| Figure 94: User probing grid cells for evaluating the connectivity graph. | 142 |
| Figure 95: By changing the time threshold through the slider widget, the connectivity graph is filtered according to it..... | 143 |
| Figure 96: User selects a specific perforation in order to assess only its associated connectivity graph..... | 144 |
| Figure 97: Mapping between the 3D ray-plane intersection point and the 2D virtual mouse click. | 144 |
| Figure 98: Each tube sub-mesh encodes its associated delta time as vertex data. | 145 |
| Figure 99: Example of an analytical connectivity panel..... | 146 |
| Figure 100: Same panel displaying data for different time thresholds. It can be observed variations in bar color and length in the second column. From left to right, $\delta = 59.6$, $\delta = 15.3$, $\delta = 0.6$ | 147 |
| Figure 101: The normalized bars ease the visual comparison between two or more panels. In the picture above, one can quick check that the total production volume on the well on the left is lower than the one on the right. | 148 |
| Figure 102: Examples of different panel arrangements. | 149 |
| Figure 103: On the left, the user is selecting the filtering operation on the graphical menu. On the right, the user is filtering the reservoir based on permeability values..... | 150 |
| Figure 104: On this gesture interface, hands must be fully opened to enable graphical controls. The left hand is used to select an operation – translation, rotation, or scale, while the right hand changes the operation value..... | 151 |

| | |
|---|-----|
| Figure 105: This interface extends the previous one by adding a sub-menu through which the user may choose the axis along which the currently selected operation is performed. | 153 |
| Figure 106: Sketches of the scaling operation. | 155 |
| Figure 107: Sketches of the rotation operation. The sizes of the axes disposed at the bottom left corners represent the magnitude of the hand motions in the associated directions. | 155 |
| Figure 108: Sketches of the translation operation. | 156 |
| Figure 109: Illustrations of our interactions for manipulating a reservoir model while using gestures: default interface when no operation is being performed (a); translation (b); rotation (c); and scaling (d). | 157 |
| Figure 110: Images of the focus lens when the user is situated inside (left) and outside (right) the reservoir model. | 159 |
| Figure 111: Illustrations of the alpha lens surrounding the user’s right index (left) and the information window associated with the probed cell. | 160 |
| Figure 112: The procedure for defining a vertical well trajectory: the user hovers a surface cell using the ray (a); after a time period, a graphical button is displayed (b); and the user hovers it with its left index in order to confirm the creation of the well (c and d). | 161 |
| Figure 113: Illustrations of participants being introduced to our application on the CAVE environment. | 162 |
| Figure 114: Participants testing our application using different immersive environments. | 165 |
| Figure 115: 2D and 3D views of the faulted reservoir model used during the evaluation sessions. | 166 |
| Figure 116: Minimum cuts may highlight potential reservoir segmentations (compartments) that lead to loss in production volume. | 177 |
| Figure 117: Conceptual prototype for capturing gestural movement over grid cells as a 3D sketch and translate it into a volumetric lens. | 178 |
| Figure 118: Theoretical design for selecting a series of contiguous cells along the well path to compose a perforation. | 179 |
| Figure 119: Illustrative set operations applied to connected regions on a reservoir model. | 180 |
| Figure 120: A grid-based snapping provides implicit vertical and horizontal lines that almost seem magnetic. When one drags an object on the grid, the object will align (“snap”) to the nearest intersection on the grid or to other objects. This is enormously useful for accurate placement of image elements. | 181 |
| Figure 121: Panel comparison using alpha blending would allow gradual transitions from one panel (a) to another one (c). | 182 |

Figure 122: Comparison between two panels using explicit encoding: the resulting panel (c) shows the complement between the first (a) and the second (b) panel..... 182

Figure 123: Immersive virtual environments have been considered promising mediums to address particular needs of oil and gas domain. In the figure above, users evaluate a petroleum reservoir field using a CAVE system. 184

1. INTRODUCTION




Figure 1: The same reservoir model being displayed in two immersive virtual environments: a CAVE and a regular desktop system.

Oil and gas reservoirs are subsurface pools of hydrocarbons encompassed by rock formations. The ultimate goal in the oil and gas industry is to explore these reservoirs and discover efficient and economic ways to produce oil and gas. However, as these petroleum products are trapped in reservoir rocks situated hundreds to thousands of meters underground, reservoir data acquisition is costly. Consequently, only limited information is available from a variety of different sources such as seismic geophones, well logs, pressure transducers, and core samples. These multivariate, multiscale data sources are combined using complex geostatistical methods and expert judgment to develop a realistic three-dimensional reservoir model representing both the structure and the stratigraphy of the subsurface volume.

A reservoir model typically consists of thousands to millions of cells, each of which is associated with various static and dynamic (that is, time-varying) properties such as porosity, permeability, pressure, and oil saturation. The structure of the three-dimensional reservoir model is often represented by corner point cells, which are irregular hexagonal geometries arranged along three dimensions (i, j, k). And since corner points are not required

to be regularly spaced nor spatially continuous, this gridding format has the ability of defining erosion surfaces that can act either as flow baffles or conduits, such as fractures, faults, and traps.

Besides the cell data, a reservoir model may also contain information about internal structures, such as well trajectories and fluid flow behavioral lines, since the model is the context in which well placement scenarios are defined and evaluated. In this sense, engineers formulate a placement scenario in a given reservoir model in order to specify a drilling recovery strategy that can be later settled in the real reservoir field. A placement scenario involves several variables such as the number of wells, the spacing between them, the locations and trajectories of the wells, the introduction or modification of injection fluids, and possibly many others. Afterward, engineers commonly run fluid flow simulation to calculate the expected outcome of the recovery strategy. The placement scenario is then evaluated based on the estimated economic production of petroleum products.

The placement of wells has a major impact in a development project. Optimal well placement improves the long-term and short-term performance of wells by maximizing hydrocarbon recovery, extending well life, and in cases of pressure support from injection, reducing water production. However, optimal well placement is challenging to determine due to the many variables involved and the static and dynamic uncertainties. These many parameters interact with each other, often in complex ways, resulting in a large solution space for fitting the well plan. As a result, engineers must define and assess a number of different placement scenarios in order to select optimal outcomes. This procedure is called well placement optimization. During optimization studies, engineers progress doing gradual changes in parameters, running fluid flow simulations, and evaluating the predicted reservoir performance.

For this reason, despite the aid of automated tools, the process of searching for optimal placement scenarios still remains heavily exploratory and rely on the analysis and interpretation of a series of specialists, who are the true driving force behind well optimization. Group work and analysis are also common for improving awareness of the data and reaching better decision making. Teams of engineers, geologists, geophysicists, and potentially

other specialists may elaborate on recovery strategies; review results from fluid flow simulations; check for inconsistencies or interdependencies on the data; showcase optimal strategies for project managers and stakeholders; etc.

Even though well optimization has a major impact on reservoir productivity, current commercial systems used by the oil and gas industry suffer from two limitations that may hinder the exploratory study of placement scenarios. One of them relates to the systems' hardware component. Although regular desktop displays along with keyboard and mouse are traditionally used, these physical interfaces provide little opportunity for serving specific demands from the oil and gas domain; this is the case of spatial abilities, collaborative interchange, and data integration. Engineers may then face difficulties when using these standard input and output devices for running well placement studies.

The second limitation, in turn, concerns the systems' software component. More specifically, it refers to the dynamic analytical methods commonly used for estimating reservoir response and production. Although they may lead to more accurate predictions, fluid flow simulations are computationally expensive and, as a result, engineers are forced to limit the number of different placement scenarios that they can test during optimization studies.

Given the relevance of well placement optimization for the oil and gas domain, this thesis addresses these two limitations in order to help advance research in well placement optimization in the oil and gas domain. To address the first limitation abovementioned, we investigate the potential benefits of more immersive virtual environments for facilitating tasks associated with well optimization. When it comes to the second limitation, we explore analytical methods that use static connectivity as a fast performance estimator of placement scenarios. The remainder of this chapter encompasses a brief overview of each of the two threads targeted by this thesis, followed by thesis goals, methodology, contribution, and lastly concluding with an overview of the remaining components of the thesis.

1.1. IMMERSIVE VIRTUAL ENVIRONMENTS

From the above mentioned, it can be observed that the oil and gas domain heavily rely on the analysis of three-dimensional, digital reservoir models. On this matter, coming from the beginning of the 90's, the incorporation of desktop systems represented a real breakthrough in reservoir geosciences and engineering [1] [2] [3]. In this sense, allowing real time visualization and manipulation of virtual reservoir models on conventional CRT monitors greatly improved understanding and eased analysis of reservoir data.

Currently, several commercial software are available to support the different activities performed throughout a field development project, and they have become integral to the oil and gas domain. Among the available commercial solutions, WIMP-based software packages for desktops have become the most pervasive tools in a manner that there is a multitude of solutions from different enterprises. CMG Suite [4] from Computer Modelling Group, Petrel E&P [5] from Schlumberger, and Nexus Suite [6] from Haliburton are among some of the most popular systems.

Even though numerous, these solutions provide standard visualization and interaction techniques for modelling, assessing, and monitoring reservoir data. Therefore, there has been a growing effort in some areas of reservoir engineering and geosciences to incorporate novel techniques to improve the ease, speed, and accuracy of these tasks, as can be seen in Figure 2. As an example, [7] suggested creating exploded view diagrams in reservoir grids, while [8] [9] proposed the use of a cutaway illustrative technique to improve the inspection and analysis of properties in reservoir models. These techniques emphasize important structures or parts of the model by selectively discarding occluding parts while keeping the contextual information in view.

As a second example, [10] argues that current computational tools applied to geology do not profit from the natural language of geologists; rather, they rely on WIMP paradigm that requires specialized training, and thus turn difficult for geologists to create conceptual 3D geo models from interpreting 2D geo maps. Therefore, the authors propose the modeling of 3D geological structures using sketch-based modeling (SBM) techniques.

These SBM metaphors try to mimic the natural manner that geologists communicate ideas through 2D drawings and translate it into 3D geo models.

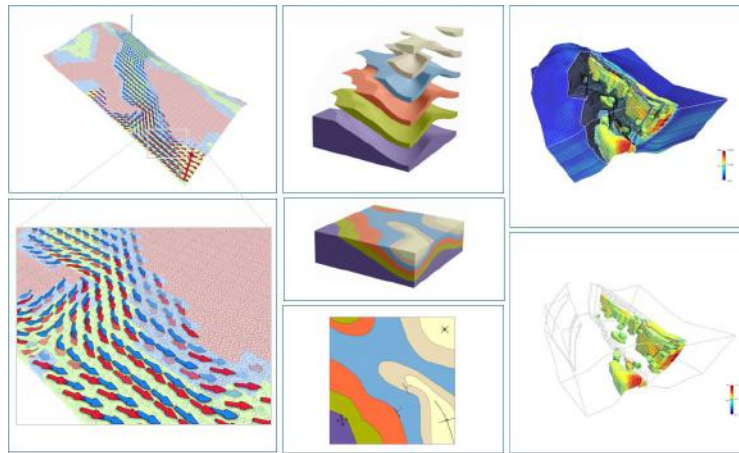


Figure 2: Examples of novel visualization techniques in oil and gas domain. From left to right: decal-maps for multidimensional data analysis, SBM for 3D geological modeling, and cutaway views for investigating internal reservoir cells.

As a last example, [11] introduced the concept of decal-maps for multivariate visualization. Decals are visual abstractions of attributes such as a pattern, a text, a glyph, or a symbol, that are transferred from a 2D-image to a 3D-surface upon contact. By combining the concepts of layers and decals, the authors propose a technique for the co-visualization of an increased number of attributes on arbitrary surfaces by layering different sets of decals, each of them representing one attribute. The authors demonstrate the applicability of the technique applied in the domain of oil and gas: to co-visualize several geological properties onto surfaces of reservoir models. This is relevant as engineers should, ideally, simultaneously evaluate geological properties for better prediction of oil recovery: they are strongly correlated to each other and their combined effects influence the overall recovery efficiency. Despite its relevance, co-visualization techniques such as this is still not supported by conventional commercial tools.

These traditional desktop systems have clearly demonstrated value by improving the efficiency of reservoir modeling and analysis; however, they still face a number of challenges for attending specific demands of oil

and gas industry. These include the co-visualization of multidimensional reservoir data to look for correlations and errors on the data; to facilitate collaborative interchange between domain experts; and the fundamentally difficult problem of handling complex spatial tasks within a two-dimensional interaction space provided by standard mouse and keyboard.

Particularly, desktop systems have been commonly adopted by industry for performing well planning; yet, tasks like the positioning of 3D well trajectories involve a series of manipulations often laid out in 2D maps or cross sections, making it difficult to understand the spatial relationships between the 3D reservoir model, the planned wellbores, and the estimated fluid flow model. This less-than-ideal mapping between interfaces and tasks, combined with the meticulous nature of formulating and assessing well designs, leads to a time-consuming and non-intuitive workflow. As a result, despite being an iterative process, well placement studies may actually prevent potentially interesting production scenarios to be explored due to the cumbersomeness of using current interfaces and their associated manipulations.

Therefore, given the inherent interdisciplinary, collaborative, and spatial nature of the activities performed in the oil and gas domain, there has been a growing attention towards systems with higher levels of immersion (see Figure 1). These systems may use a combination of different immersive interfaces to access their unique benefits for better performing activities contained in the oil and gas life cycle. Some benefits that may be particularly helpful for well planning include enhanced spatial understanding, increased information space, and more natural interactions. In this thesis, we explore immersive technologies' potential benefits for facilitating the execution of tasks associated with well placement optimization.

1.2. STATIC CONNECTIVITY ANALYSIS

As previously mentioned, to assess well placement scenarios, dynamic simulation techniques such as fluid flow simulations are performed in order to select the best scenario based on the predicted production performance.

There are, however, limitations to the number of scenarios that can be assessed using flow simulation. The time required by dynamic simulation can be very high, and increases significantly with the complexity of the scenario, the resolution and geometry of the reservoir grid, and the fidelity and duration of the simulation. It may thus require hours, days, or even months for the simulation to finish, results gathered, and the well performance evaluated.

Due to this, much research has been devoted to develop fast performance estimators as surrogates for flow simulation. These estimators do not aim to replace a full flow simulation; rather, their value lies in rapidly determining parameter sensitivities and screening reservoir models or production scenarios. In particular, some efforts focus on using a combination of static, geological information to quantify reservoir connectivity, a derived property that has already been proven to have a strong correlation with the efficiency of hydrocarbon recovery [12]. These static connectivity estimators are easy in concept, inexpensive in execution, and create an important intermediate level between the reservoir characterization and simulation studies for the assessment of reservoir productivity.

Hence, this thesis investigates analytical techniques that use static connectivity analysis as a proxy for dynamic flow simulation. And, although distinct tasks contained in a typical reservoir development workflow may be benefited from the addition of connectivity analysis – e.g., ranking geological realizations and looking for inconsistencies introduced by upscaling the fine-scale geological model into the coarse-scale simulation model, we focus on its use for assessment of optimum well placements.

1.3. GOALS

The main goal of this thesis is to investigate and engage in the design, implementation, and evaluation of new exploration paradigms for supporting well placement optimization. This broad goal can be broken down into two specific aspirations that complement each other. First, we aim to develop an analytical technique that forecasts

fluid movement within the reservoir and production profile for a certain production scenario configured by the reservoir engineer; yet, such technique should use solely static information from the reservoir to approximate the dynamic behavior of the fluid flow in order to be computationally efficient in terms of time. Secondly, by using the analytical method abovementioned, we aim to develop an application for the exploratory study of production scenarios; further to this, such forecasting tool should leverage unique benefits of more immersive interfaces for efficient visualization and interaction techniques that can facilitate tasks associated with well planning.

Hence, our goal focuses on conceiving alternatives to the current content creation and analysis paradigms used during well planning. After all, as previously mentioned, the number of placement scenarios investigated may be limited given the constraints imposed by the use of dynamic simulations and standard desktop, keyboard, and mouse interfaces. Thus, we expect to promote exploratory studies through the use of static connectivity analysis and more immersive interfaces, which may allow engineers “to test a number of different well placement scenarios; gather the most promising ones, and then run flow simulations with this best-small subset”.

It is worth noting that the two specific aspirations we aim to achieve require a multidisciplinary approach, building on theories and concepts from various disciplines, such as Human-Computer Interaction (HCI), Computer Graphics, Graph Theory, and Petroleum Engineering. Therefore, in this thesis, the reader will find elements from all these domains.

1.4. METHODOLOGY

In order to address our specific research goals, we had continuous meetings with a post-doctoral scholar in Petroleum Engineering who is part of our research group. These meetings were crucial to design both an algorithm and a system that are meaningful for reservoir engineers. Indeed, such meetings were essential to

understand which and how geological properties of the reservoir affect the flow of fluids during the production of oil and gas; this was a key factor for us to numerically model the expansive behavior of the fluid when elaborating our analytical technique. In addition, these meetings were important to acknowledge all the steps involved in well placement, starting from understanding engineers' decision-making on potential locations for placing wells to learning how they conduct the analysis of simulation results, as well as the difficulties engineers face when doing well placement studies using current commercial software.

Furthermore, we also surveyed existing analytical techniques that perform static connectivity analysis, always considering their technical strengths and limitations. Building on the knowledge acquired from the meetings with the petroleum engineer and the survey, we designed and implemented our own analytical technique for connectivity analysis.

When it comes to immersive environments, we investigated which and how immersive technologies have been addressing demands from the oil and gas industry. This raised our knowledge on novel immersive technologies that industry have been adopting, how they have been used, and how we could use them in our specific research. Further to this, we also surveyed academic and commercial systems that engineers may use for running well placement studies. This survey added to our knowledge on most commonly used interfaces for performing well placements as well as their technical limitations. Along with the knowledge acquired from the meetings with the post- doctoral scholar, all this study was useful to examine the potential of more immersive interfaces for well planning.

Following this study, we combined different interfaces in order to compose three immersive virtual environments: a CAVE with a tracked gamepad; an HMD with a tracked gamepad; and an HMD with a leap motion controller. For each environment, we elaborated on how to leverage the unique benefits of its technologies so we could define efficient visualization and interaction techniques for our application for well

optimization. Finally, we conducted a primary study with reservoir engineers in order to examine the relevance of our research for advance well placement optimization and for the broader area of oil and gas.

1.5. CONTRIBUTIONS

This thesis is multidisciplinary in its essence since it is a result of a combination of elements from Computer Science and the Oil and Gas domain. In this sense, it provides contributions to the literature concerning each of these broad areas. The main contributions of this thesis, however, can be considered to be the following:

- An analytical method that uses static connectivity analysis as a proxy for dynamic flow simulation to estimate production scenarios;
- An application to serve as an exploratory tool to run well placement studies, which makes use of our analytical method. Such software runs in three immersive environments, taking advantage of the unique benefits provided by the physical interfaces in each environment for creating techniques that eases the tasks associated with well planning;
- Reflections on engineers' needs, perceptions, and preferences on our analytical technique, as well as on our software and on the use of more immersive technologies for well placement optimization and for the broader area of oil and gas.

1.6. THESIS OVERVIEW

The purpose of this chapter was to explain why we decided to focus on more immersive virtual environments and static connectivity analysis, and where that stands within the context of well placement optimization. The following chapters are organized as follows:

- In *Chapter 2*, we will provide an overview on immersive virtual environments. We will first explain core concepts of virtual reality while exploring previous and current works in the area. Following this, we will

focus on immersion by giving its definition and explaining its relevance through the four potential benefits that it brings along: an increased sense of presence, a greater information space, more spatial understanding, and the possibility of more natural interactions.

- In *Chapter 3*, we will strictly focus on the oil and gas domain. We will first provide an overview of the field that discusses the oil and gas industry's goal and the three stages that comprise the oil and gas life cycle: exploration, development, and production. Following, we will target well placement optimization by explaining in detail the process of defining and assessing well placement scenarios, as well as discussing the benefits and drawbacks of the traditionally used reservoir fluid flow simulation. We will then introduce static connectivity analysis as a surrogate for flow simulation, and describe past and recent analytical methods that use reservoir connectivity for supporting different tasks contained in the oil and gas workflow. Finally, we will present our technique for static connectivity analysis as a possibility that may help reservoir engineers perform well placement studies.
- In *Chapter 4*, we will build on what that has been covered in Chapter 2 and Chapter 3. Hence, we will discuss how immersive virtual environments have been serving the oil and gas industry and illustrate the main advances through a number of examples from academia and industry. More specifically, we will approach works that have shown that higher levels of immersion are of great value for the oil and gas domain by exploring in detail each of the four benefits of immersion defined in Chapter 2 to challenges of the oil and gas industry explained in Chapter 3. In addition, we will emphasize how immersion has affected well planning and highlight the missing possibilities in that relationship. We will then describe works that tackle well planning, and elaborate on their limitations and on how these may potentially be mitigated by the addition of higher levels of immersion.

- In *Chapter 5*, we will present our application for supporting well planning, which uses our analytical connectivity method described in Chapter 3. We will do so by describing the hardware and software components employed by our application. For the former, we will present different input and output interfaces that were combined in order to compose three immersive virtual environments: a CAVE with a tracked gamepad; a HMD with a tracked gamepad; and a HMD with a Leap Motion controller. For the software component, we will present the design and implementation of visualization and interaction techniques for each of the different environments. These techniques aim to leverage the benefits provided by the different interfaces in order to facilitate the design and assessment of well placement scenarios.
- In *Chapter 6*, we will discuss the primary qualitative study we conducted with two main purposes: to evaluate the features of our application for well planning, and to investigate engineers' impressions and preferences in regards of the three immersive environments while using our application. We will do so by presenting the feedback we collected from reservoir engineers during evaluation sessions. Reflecting on such feedback, we will discuss the effectiveness of our application and the significance of more immersive interfaces for well placement studies and for the broader domain of oil and gas. Finally, we will present perspectives for future work drawing on the feedback we received during evaluation sessions.
- In *Chapter 7*, we will highlight the main contributions of this thesis by elaborating on what has been done in each of the previous chapters and provide final remarks.

2. IMMERSIVE VIRTUAL ENVIRONMENTS




Figure 3: Birdly [13] is a full-body simulator ride that uses multi-hinged flags, head tracking, and motion feedback to evoke the sense of flying. In addition, there is a fan in front of the user to simulate headwind which gives a sense of flying speed.

This thesis builds on the existing knowledge on immersive virtual environments to advance research in the oil and gas domain. Hence, in this chapter, we provide an overview of immersive virtual environments, discussing core concepts and outlining prior and current work in the area. This chapter is divided into two parts; the first of them concerns general terminology in virtual reality and the second part is centered on the concept of immersion. In the first part, we will discuss the different form of realities, ranging from the pure real environments to the pure virtual environments. We will then give a definition of virtual reality systems and a description of different forms of input and output interfaces that can be used as part of virtual reality systems.

In the second part, we will approach a concept that directly derives from the interface components of virtual reality systems and is named immersion. We will distinguish between the often-confounded concepts of presence and immersion: presence refers to a subjective, mental state of the user; immersion is an objective,

technology-related characteristic of the VR system. Following this, we will discuss that immersion should be understood as a continuum and that the combination of different interface technologies leads to different levels of immersion. The remainder of the second part will concern the potential benefits that may come along with higher levels of immersion – namely, an increased sense of presence, increased information space, enhanced spatial understanding, and the possibility of more natural interactions.

2.1. THE REALITIES

Reality may have many forms, ranging in a virtuality continuum from the pure real environment to the pure virtual environment [14]. The forms of reality that are somewhere between reality and virtuality are broadly defined as mixed reality, and can be further broke down into augmented reality and augmented virtuality. Figure 4 depicts different forms of reality along the virtuality continuum.

The pure real environment encompasses solely real objects, and some examples include direct viewing of a real world scene, or observing the same scene via a conventional video display system. The pure virtual environment consists of only virtual objects, and an example of such would be a conventional computer graphic simulation.



Figure 4: The virtuality continuum – adapted from [14].

The mixed reality environment is one in which real world and virtual world objects are presented together within a single display. This term covers a grey area in the center of the continuum that ranges according to the primary world being experienced. Mixed reality may be further classified either as augmented reality (AR) or augmented

virtuality (AV). Augmented reality refers to environments in which the real world is consistently supplemented by virtual objects, and an example is a video see-through head-mounted display with which computer-generated graphics are visually superimposed onto directly viewed real-world scenes. Augmented virtuality, in turn, is the result of capturing real-world content and bringing it into a virtual world. An example is a completely graphic display environment in which a video of a real world scene is added.

This thesis particularly addresses the right range of the virtuality continuum: virtual environments (VE), also generically referred to as virtual reality (VR) systems, such as Birdly, a virtual reality flying simulator that is illustrated in Figure 3. The following sections will discuss in more detail this form of reality.

2.2. THE VIRTUAL REALITY SYSTEM

A virtual reality experience refers to a human sensing and interacting with a virtual world – that is, a world in which real-world sensory stimuli is replaced by synthetic sensory stimuli. A virtual reality system, in turn, is the hardware and software system onto which virtual reality experiences are built upon. Figure 5 illustrates a user and a VR system segmented into its primary components: input, virtual world, and output.




Figure 5: A virtual reality system comprises an input interface, a virtual world, and an output interface – adapted from [15].

The input concerns the communication from the user to the VR system; it refers to the hardware interface that collects information from the user into the system, such as head and hands' motions, button presses, voice commands, etc.

The virtual world refers to the software application that defines the virtual objects and the rules and relationships governing those objects. Therefore, it includes both non-rendering and rendering aspects. In the first case, non-rendering routines include interaction techniques that convert user information collected from input devices into corresponding virtual actions. Other non-rendering processes include physics simulation models; data structures and algorithms, such as trees and graphs; among others. In the case of rendering aspects, rendering routines generate visual, aural, haptic, and other sensory signals that are sent to the output devices.

The output concerns the communication from the VR system to the user; it refers to the hardware interface that sends physical, sensory feedback from the system back to the user. These three components of a VR system form a cycle of input/output that continuously occurs for as long as the VR experience lasts. The subsequent sections address the two interface components of a VR system, to then discuss a concept that directly derives from them, named immersion.

2.3. THE INTERFACE TO THE VIRTUAL WORLD – INPUT

The user's interaction with the virtual world is an essential component a virtual reality system. In this sense, one or more physical input devices may be used to convey the user's actions in the virtual world to the VR system.

In this section, we explore a wide range of VR systems' input devices using a taxonomy similar to the one proposed in [15]. According to such taxonomy, as nowadays most VR systems' input devices are entirely based on the human's hands, they are broadly organized into two distinct classes: hand and non-hand input devices. In turn, each of these classes can be further arranged in two or more sub-classes, which are sets of input

devices that share similar characteristics. Each class of input devices is discussed in details in the next subsections.

2.3.1. HAND INPUT DEVICES

This class of input devices can be further classified as world-grounded, non-tracked hand-held controllers, tracked hand-held controllers, hand-worn, and bare-hand input devices. World-grounded input devices are designed to be constrained in the real world, and some examples are conventional keyboards and mice. Another example worth mentioning is a high-fidelity system designed by Disney Imagineering, in which people fly a magic carpet through a virtual world based on the animated film Aladdin [16]. Guests use an HMD, seat straddling a motorcycle-style seat, as shown in Figure 6, and grip a steering mechanism representing the front of the magic carpet. Turning left or right controlled yaw of the carpet, and tilting controlled the pitch of the carpet. Pulling the top of the wheel toward the driver pitched the carpet up while pushing it pitched the carpet down; pushing the entire mechanism forward, controlled velocity.




Figure 6: Examples of world-grounded devices developed in the early 1990's. (a) A person uses a real bulldozer steering wheel on the left that meets the virtual steering wheel on the right; (b) Disney Imagineering designed a steering mechanism to give guests the illusion of flying a magic carpet as in the film Aladdin; (c) surgeon practices a surgical procedure using an artificial cavity together with the correct instrument set coupled with force-feedback devices; (d) a researcher uses a force-feedback device to control the tip of a scanning tunneling microscope probe, exploring the surface of a sample of atoms.

Just as the case of the magic carpet illustrates, world-grounded devices may provide an intuitive physical interface as users may already have real-world experience with similar devices. They may also provide proprioceptive cues and be quite presence-inducing due to having physical signifiers, affordances, and feedback. Challenges in working with world-grounded devices, however, include the hardware accessibility to a wide audience and the difficulty to generalize the same device to work with a wide range of tasks.

Non-tracked hand-held controllers, in turn, are devices held in the hand that include buttons, analog sticks, triggers, etc. such as joysticks and gamepads commonly used along video games. An example is the MOGA Pro Controller [17], a Bluetooth gamepad that features dual analog sticks, bumpers, triggers, four face buttons, and a D-pad. Tracked hand-held controllers, also called wands, are devices tracked in 3D space and may present the same physical functionalities offered by non-tracked controllers. The Oculus Touch [18] is an example. It is a wireless controller with 6-DOF tracking along with an analog thumb-stick, two buttons, and an analog trigger. In addition to reporting input state, Oculus Touch controllers can provide haptic feedback through vibration.



Figure 7: Examples of hand-held controllers. Oculus Touch [18], Sony PS3 Move [19], THE VOID Nova IV Gun [20], and WorldViz Wand [21] are tracked controllers – (a), (b), (c), and (d), respectively. MOGA Pro Controller [17] is a non-tracked gamepad – (e).

Similarly to world-grounded devices, controllers may be an intuitive interface as many gamers already have experience with them. In addition, controllers provide proprioceptive cues as they can be physically felt; yet, wands may provide higher proprioceptive sensitivity as they are tracked and, thus, can be visually co-located with the real hands. In particular, tracked controllers are typically 6-DOF and thus are currently the most used devices in many 3D tasks due to their natural, direct mapping to the hand motion. A drawback of controllers however is that they do not provide the possibility of fully touching or feeling other objects in the virtual world.

Bare-hand input devices work via sensors aimed at the hands – mounted in the world or on the HMD. Figure 7 depicts the leap motion controller [22] that is a small USB peripheral device designed to be placed either on a flat surface near a computer or mounted onto an HMD. Using two monochromatic IR cameras and three infrared LEDs, these sensors tracks hands and fingers in a roughly hemispherical area, to a distance of about 1 meter.

Clearly, the main advantage of this input interface is that the user's hands are completely unburdened with any physical prop. Nonetheless, interaction with the bare hands has several usability challenges; these include not having a sense of touch, fatigue from holding the hands in front of the sensor, line-of-sight requirements, and consistent recognition of gestures across a wide range of users.

Hand-worn input devices include gloves, muscle-tension sensors, and rings. An example is Myo [23], an armband that reads the electrical activity from the muscles of the forearm and map them to gestures made with the user's hand. Another example is CyberGlove [24], a glove that uses resistive bend-sensing technology to transform hand and finger motions into digital joint-angle data. In addition, there are a number of tactile feedback options to be used in combination with CyberGlove, such as CyberTouch II [25], CyberGrasp [26], and CyberForce [27].

Similarly to bare hands, hand-worn devices have the benefit of having free hands to interact with physical or virtual objects. Yet, these devices may have advantages over bare hands, such as having haptic feedback

and not having line-of-sight, sensor field-of-view, or lightning requirements for not losing tracking. Nevertheless, a challenge that comes with hand-worn devices, just like bare-hands devices, is that hand and finger gesture recognition is still challenging due to the lack of consistent tracking accuracy.

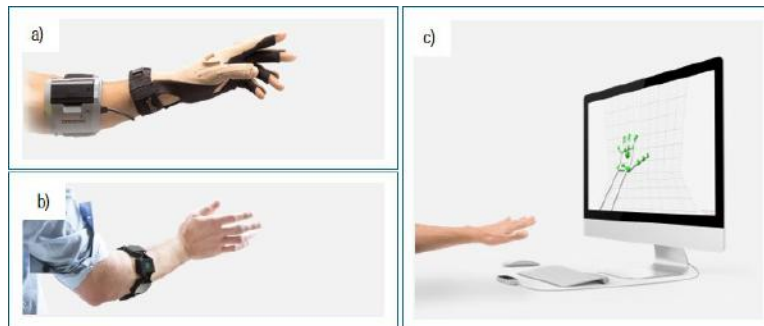


Figure 8: CyberGlove (a) and Myo (b) are hand-worn input devices. The leap motion (c) is a bare-hand input device.

2.3.2. NON-HAND INPUT DEVICES

This class of input devices can be further classified as head-tracking, eye-tracking, foot-tracking, full-body tracking, and voice input devices. Head-tracking input refers to tracking the position and/or orientation of the user's head. One example is used in the Oculus Rift [28], illustrated in Figure 9. This head-mounted display uses a gyroscope, an accelerometer, and a magnetometer to monitor head rotations – through a process known as sensor fusion, and apply the constellation positional tracking system to monitor its position in space. More specifically, the constellation tracking is an optical-based outside-in tracking technology where a series of infrared LEDs are embedded in the headset and, then, are monitored by a wireless sensor – a small microphone-shaped pole placed on flat surfaces.

In addition, it is worth noting that there are two common approaches to using head-tracking input. One of them is to passively track the user's head motions in order to provide context to the system. An example is to use the continuous tracking to allow the system to display the virtual world from a user-centric perspective. Another approach to using head-tracking input is to actively track a user's command made with the head in

order to trigger an action; an example is having virtual characters to respond whenever the user turns its head in their direction or approach them. Another example is doing head gestures to confirm a selection, such as nodding the head yes or no.

Eyes-tracking input, in turn, refers to tracking where the eyes are looking. FOVE [29] is an example of that type of input as it consists of a headset with two small infrared eye tracking systems with an accuracy of less than one degree. By using these high-precision tracking systems, FOVE aims to create more natural VR experiences by allowing virtual worlds to react to subtle eye movements and expressions. For instance, for aiming at, picking up, or throwing virtual objects, it would feel more natural if users only needed to look in a direction or at an object rather than physically motion their forehead towards the objects. This interaction is more natural, intuitive, and seamless as it closely mimics real-life interaction.



Figure 9: Oculus Rift [28] uses sensor fusion and Constellation tracking to provide head-tracking input – (c). FOVE [29] uses infrared sensors for eye tracking – (a, b, d). By using eye-tracking, the FOVE system allows two innovative features: eye contact monitoring and foveated rendering. Foveated rendering refers to loading high-resolution textures where users are actually looking, and optimizing out-of-focus areas with lower-resolution or fewer effects.

Similarly to head-tracking, eye-tracking input is usually used in two manners. One of them concerns passively tracking the eyes' gaze to provide context to the system – e.g., firing a gun where one is looking at; displaying a higher degree of detail in the direction of the tracked eyes' gaze. Another manner is to actively track eyes' gaze to trigger an action. An example is virtual character response whenever the user looks at or blinks to them.

Eye gaze selection is another example as it is a form of pointing where the user simply looks at an item of interest to select it.

Foot-tracking input refers to tracking one or both user's feet. Common approaches include stationary platforms such as bicycles, stair steppers, or treadmills, with the system causing the platform to move in accordance with the participant's steps. For instance, Cyberith Virtualizer [30] is a circular omnidirectional treadmill with sets of sensors in its flat base plate, pillars, and ring to monitor the movement of the feet. In addition, the system supports vertical movements such as jumping, crouching, sitting down, as well as 360-degree axial rotation.

It is also interesting to note that, although being mostly passively used as a means to determine the speed and/or direction a user wishes to travel, foot-tracking input also holds the potential to perform active, interactive actions. One example is Maximum VR [31], a game currently in development by AlienTrap. Recently, the company showed a game prototype in which the player strapped a pair of HTC/Valve controllers to his feet and has used them to add tracked-feet functionalities to a game – e.g., smashing building and hurling cars.

Contrary to the types of input that have been mentioned in this section so far, full-body-tracking input involves simultaneously tracking multiple body parts rather than just the head or the hands. One way of doing so is by using motion capture suits that are similar to what is used in film industry, such as PriorVR suit [32]. This system places a number of light inertial sensors in different parts of the body to provide motion tracking without the need for cameras, line-of-sight, or large equipment. Other less encumbering methods may use infrared (IR), RGB, or depth cameras such as Microsoft Kinect [33], a motion sensor consisting of a flat black box that sits on a small platform, placed on flat surfaces.



Figure 10: PriorVR suit [32] places inertial sensors onto the user's body, while Microsoft Kinect uses optical-based sensors placed on an add-on peripheral to provide full-body tracking input. – (a, b) and (c), respectively. Cyberith Virtualizer [30] is a treadmill that uses motion sensors in its base plate, pillars, and ring to provide foot tracking input – (d).

Finally, voice input may play an important role in creating a natural, non-encumbering communication channel with computer systems. This is especially true with VR systems, where the goal is to provide the most natural form of interface possible. Some recent hype over voice interfaces involved emerging technologies such as Microsoft HoloLens [34] and Google Glass [35]. Other extensive publicity over voice recognition systems were drawn upon intelligent personal assistants such as Siri [36], an application incorporated into Apple's iOS. Siri uses a natural language user interface for its voice recognition function to perform tasks exclusively on Apple devices, such as to make web searches, take pictures, record videos, send messages, make phone calls, etc. In a similar manner, Google Now [37] is a personal assistant that supports voice searches and commands for both iOS and Android devices.

Even though they provide an excellent opportunity for natural communication, technical issues faced by recognition systems may include the presence of either false negatives (that is, missed identifications) and false positives, like when the user is thinking aloud or when another person is talking. In addition, in regards of the

former matter, since speech input control is not as accurate as using devices like buttons, it is challenging to address tasks that require precisely timed inputs, including those that are to be correlated to physical movements. These tasks thus work best with physical control devices.

The issues mentioned above can be minimized through the use of microphones specifically designed for speech recognition that offers noise-canceling features for reducing ambient noise. In addition, in order to prevent unwanted operations, some methods of active selective listening may be considered to decide when the recognition system should attend to what the user is saying, such as “push to talk”, “name to talk”, and “look to talk” techniques [38]. However, even though these minor challenges may be mitigated by the application, an ultimate speech recognition system has yet to be developed. Such system would be capable of either processing a steady stream of speech from the user, and understanding context and use it to interpret the speech.

2.4. THE INTERFACE TO THE VIRTUAL WORLD – OUTPUT

The user’s perception of the virtual world is another fundamental component of a virtual reality system. As the human perceptual system consists of five senses providing information to the brain, the virtual reality system tricks one or more senses by displaying synthetic stimuli that replaces natural stimuli – note that the term “displays” broadly means a way of presenting information to any of the senses. In this section, we discuss the VR systems’ output display with a focus on the visual, aural, and haptic senses since these are the most commonly used channels to transmit synthetic stimuli in VR systems.

2.4.1. VISUAL DISPLAYS

Today’s visual displays are generally implemented in two ways: world-grounded and head-mounted displays. The following subsections will approach these two categories of visual displays by describing their unique features, strengths, and weaknesses and by providing examples.

2.4.1.1. WORLD-GROUNDED DISPLAYS

World-grounded displays render graphics onto surfaces that do not move with the head – that is, surfaces that are fixed in place. As can be observed in Figure 11, these displays can be further classified as monitor-based or projection-based displays. Monitor-based displays use standard monitors and are the simplest form of visual displays. Some classical models are known as fish tank VR – and the name comes from the similarity between observing through an aquarium glass and watching the virtual world. These displays differ from generic interactive computer graphics displayed on a monitor because a fish tank VR system tracks the user's head and the rendered scene changes in response to the tracked head movement.

Projection-based displays project the virtual view onto physical, not-necessarily-flat surfaces and most displays are rear-projected to prevent users casting shadows on the screen. Despite being currently more economical, in the future, rear-projection screens and projectors will probably be replaced by large monitor displays thin enough to be hang against a wall. Reasons for favoring panel displays include having higher resolution, demanding less maintenance, and requiring less room than typical rear-projected displays.




Figure 11: The EVL CAVE (a) and ImmersaDesk (d) are projection-based visual displays, while the ActiveWall (b) and ActiveDesk (c) are monitor-based displays designed by Virtalis [39].

A nice feature of world-grounded displays is that they may occupy a larger portion of the viewer's field of view than most of head-mounted displays. Other benefit refers to the fact that not much physical equipment is

required to be worn; rather, a pair of lightweight wireless glasses to create the stereoscopic views and a tracking device is usually the only requirement; that pair of glasses can be easily mounted right onto the display. On the downside, world-grounded displays often require more room space and maintenance. As a direct consequence, they tend to have higher cost when compared to head-mounted displays, thus leading to less frequent use.

Furthermore, contrary to most head-mounted displays, world-grounded displays do not isolate the viewer from the real world. There are advantages and disadvantages to this, many of which relate specifically to the nature of the application. For instance, for applications in which collaboration with coworkers is important, being able to see them is definitely an advantage.

2.4.1.2. HEAD-MOUNTED DISPLAYS

Head-mounted displays (HMD) are worn on or in some way attached to the user's head and move in conjunction with the head. Consequently, no matter which way the user turns their head, the displays move, remaining in a fixed position relative to the user's eyes. HMDs can be further classified into two types: non-see-through and see-through HMDs.

Non-see-through displays occlude all the real-world view in favor of the virtual view. These are currently the most common visual displays used in VR systems and, as can be observed in Figure 12 (top row), they are offered in a variety of sizes, shapes, resolutions, prices, etc. See-through displays combine real-world with computer-generated image sources and are usually used in AR and/or AV systems. As can be seen in Figure 12 (bottom row), these displays can be grouped into two categories: optical see-through (OST) and video see-through (VST).

With a VST HMD, the real-world view is captured with two miniature video cameras mounted on the head gear, and the computer-generated images are electronically combined with the video representation of the real world. With an OST HMDs, the real world is seen through semi-transparent mirrors placed in front of the user's

eyes. These mirrors are also used to reflect the computer-generated images into the user's eyes, thereby optically combining the real- and virtual-world views.



Figure 12: On the top row, HTC Vive [40], Samsung Gear VR [41], and Google Cardboard [42] are non-see-through HMDs. On the bottom row, Vuzix M100 [43] and NVisor ST50 [44] are OST HMDs, while Visette45 [45] is a VST HMD.

The most notable positive feature of head-mounted displays is that the field of regard covers the entire surrounding of the viewer. Thus, unlike world-grounded displays, head-mounted displays afford no gaps in imagery regardless of which direction the user looks. Furthermore, these displays tend to be more portable than most stationary displays, and require very little room space. On the negative side, some technical challenges faced by HMD systems include (i) difficulty in supporting multiple users, (ii) discomfort and fatigue due to their weight when users have to wear them continuously after a long time, (iii) adverse effects such as eyestrain and nausea, and (iv) the latency of the system. The latter, in particular, relates to a gap in time, referred to as “lag”, between the moment when the HMD positional data are measured and the moment when the virtual image for that position is fully rendered to the user. During this period of time, if the user moves its head, the discrepancy can break the illusion of the virtual reality and contribute to motion sickness.

2.4.2. AURAL DISPLAYS

Similarly to what occurs with visual displays, aural displays typically fall into one of two general categories: headphones and speakers. As can be observed in Figure 13, headphones are analogous to head-mounted visual displays as they move with the participant's head, are for one person only, and provide an isolated

environment. Headphones may be manufactured either as built-in accessories of HMDs, or as standalone devices to be integrated to different HMDs. Speakers, in turn, are similar to world-grounded visual displays as they are stationary display systems and work well as group presentation devices.



Figure 13: Types of aural displays: standalone and built-in headphones (left, middle) and loudspeakers (right).

A very appealing feature of aural displays in VR systems is the ability of making a sound seem to come from a particular location in space; this is called spatialization. By tracking the user's head, individual sounds can be processed by a computer and displayed to their ears in a way that seem to come from a particular spot. For instance, in Google VR SDK [46], to simulate spatial audio coming from different virtual objects, the listener is surrounded by a high number of virtual loudspeakers to reproduce sound waves coming from any direction in the listener's environment. The denser the quantity of virtual loudspeakers, the higher the accuracy of the synthesized sound waves.

On that matter, the greater proximity of the headphone display makes it much easier to control exactly what sound is presented to each ear than through speakers, and it is easier to determine the relative placement between the display and the ear. Yet, creating spatialized sounds using speaker technology may be challenging because speakers present a combination of direct and reflected sound waves, making it difficult to be controlled to ensure that each ear hears only the information intended for it.

Nonetheless, the stationary nature of speakers makes the sounds they produce world-referenced, which is generally used in VR systems and may be advantageous depending on the nature of the application. This is specially efficient when using surround sound systems such as Envelop [47], illustrated in Figure 13. In addition,

it is often more comfortable to listen to speakers for longer periods of time than headphones, which can be both uncomfortable and fatiguing due to their weight.

2.4.3. HAPTICS DISPLAYS

Haptics are quite powerful mechanisms to make users believe something is “real”. By coming into physical contact with an object, its existence is confirmed. Interestingly, the nature of the haptic interfaces dictate that it is a form of both input and output: as an output, it is physical stimuli displayed by the computer; as an input, it is the physical connection to the user transmitted to the computer. Haptics can be classified as passive (static physical objects) or active (controlled by the computer), tactile (through touch) or proprioceptive (through muscle spindles), and self-grounded (worn or held) or world-grounded (attached to the real world). Each one of these categories will be further discussed in the following subsections.

2.4.3.1. PASSIVE VS. ACTIVE

Passive haptics provide a sense of touch at a very low cost: one simply creates a real-world tangible object that matches the shape of a virtual object. The tangible model thus provides both a haptic and a visual representation of the object, though strictly speaking, the model is not particularly interactive and functions only as an output system. Nonetheless, passive haptics may increase presence and improve cognitive mapping of the environment. Perhaps the most compelling example to this day is the UNC Chapel Hill Pit’s study to assess presence [48].

As illustrated in Figure 14, for the UNC Chapell Pit’s study, a participant enters the left-most room, which is used for familiarization with the VR system such as learning how to move around or to select objects. The participant is then asked to move to the next room to pick an object that has been left on the plank that overlooks the pit. In some experiments that have been carried out, ‘static haptics’ were positioned such as small but real

ledges, which added to the effect of standing over a real pit. It has been shown that this passive haptics significantly increased the heart rate of participants compared with when the ledge was absent.



Figure 14: The 'pit room' using ledges as passive haptic feedback.

Active haptics, in turn, are controlled by a computer and are the most common form of haptics. They have the advantage of being dynamically controlled by the computer to simulate a wide range of forces of virtual objects. As an example, Gloveone [49] uses ten actuators distributed along the palm and fingertips; they vibrate independently at different frequencies and intensities in order to reproduce accurate touch sensations. Gloveone is an ongoing project that aims to enable users to feel and touch any virtual object: if a virtual apple is shown on the screen, the glove would enable users to feel its shape or weight, sense its physical features, and even smash it. The remainder of this section focuses on active haptics.

2.4.3.2. TACTILE VS. PROPRIOCEPTIVE

“Haptic” is loosely referred to as the sense of touch, but haptic perception actually relates to the combined sensations of proprioception and tactition. The term proprioception means stimulation from within the muscles, tendons, and joints of body. The term proprioceptive force, in turn, is often used as a synonym of force display in VR systems. Proprioceptive displays then generate synthetic forces to provide a sense of limb movement and muscular resistance.

Resistance and movements generated by the system can then be used by the user to interpret some aspect of the virtual world. Dexmo [50], for instance, is an exoskeleton glove with variable force feedback and motion capture abilities, as can be seen in Figure 15. The glove stimulates the sense of real touch by physically pulling back the user's fingers to fit the shape of the virtual objects, and it dynamically changes the force applied to the fingers to stimulate their stiffness. Therefore, the user not only feels the physical presence of a virtual object, but can also tell the difference between a virtual stone and a virtual rubber duck just by squeezing them.



Figure 15: Dexmo (a) provides force feedback, while Null Space (b) and KOR-FX (c) suits simulate tactile feedback.

The term taction, in turn, is the sense of touch that comes from sensors at the surface of the skin. Tactile displays thus focus on the skin's ability to interpret stimuli such as pressure, temperature, electricity, etc. For instance, NullSpace VR [51] is working on a vibrotactile display that stimulates tactile sensations using mechanical vibrations on the skin. It is a full upper-body haptic feedback suit and gloves with stand-alone motion tracking. The vest has 32 independently activated vibration pads placed on the chest, abdomen, shoulders, arms, and hands with 117 built-in haptic effects (see Figure 15).

KOR-FX [52], in turn, is an electrotactile display that induces tactile sensation via electrodes passing current through the skin. It is a gaming vest that transforms audio signals coming from games or other media into pinpointed, high-definition tactile responses that allows users to feel the on-screen action and environmental factors such as explosions or blade attacks. KOR-FX's acousto-haptic technology filters audio waves for creating haptic output using special transducers. This output, in turn, echoes through the user's chest cavity.

2.4.3.3. SELF-GROUNDED VS. WORLD-GROUNDED

Proprioceptive haptics can be either self-grounded or world-grounded. Self-grounded haptics are worn/held by and move with the user. The forces applied are also relative to the user. Gloves with exoskeletons are examples of self-grounded haptics. Hand-held controllers are also examples of self-grounded haptics. Such controllers might be simply a passive prop or might be rumble controllers that vibrate to provide haptic feedback to the user – e.g., to indicate that the user has put his hand through a virtual object.

World-grounded haptics, in turn, are physically attached to the real world. As an example, Widerun [53] is a made-for-VR cycling system, a bike trainer that aims to get users cycling through 3D environments (see Figure 16). The bike actuates forward and backward based on the terrain inclination; leans sideways based on user's turns; and has pedaling devices that become more resistive to the user's exertion to imply that they are going up a hill or are using a higher gear-ratio.



Figure 16: Widerun is compatible with any kind of bike. It gets inputs both from the game software and from the bike trainer to regulate back the electrical signals accordingly to the virtual experience the user is having, like steering degree, speed magnitude and ground resistance. Widerun is compatible with Oculus Rift and Samsung Gear VR.

2.5. THE INTERFACES TO THE VIRTUAL WORLD – IMMERSION

Given the existing wide range of powerful technologies, virtual reality systems have received extensive media coverage as it allows users to “be immersed” in high-fidelity, computer-generated worlds as if they were real

worlds, thus producing the sense of “being present” in users’ minds. Most media make use of the terms “being immersed” and “being present” as if they can be used interchangeably; these terms, however, refer to two different concepts: immersion and presence, respectively.

Presence is a subjective, psychological state that relates to the VR user. It is a feeling of being mentally involved in the VR experience. Immersion, in turn, is an objective, technology-related property of the VR system. It is then important to note that presence and immersion are logically separable, although empirically they are strongly related. Due to this recurrent misconception, the next subsections focus on clarifying these two concepts. Following this, we will discuss more closely immersion and its potential benefits to experience VR.

2.5.1. PRESENCE

As mentioned above, presence refers to the user’s feeling of “being there”; it occurs when the user “forgets” that his/her perceptions are mediated by other mediums. Beyond that, it is interesting to note that the concept of presence is not confined to the VR medium: reading a book, talking on the phone, and watching television can all engage some level of presence. For instance, suppose that someone close his/her eyes and starts listening to a quadraphonic sound system that is playing some orchestral music. The person may then feel “as if they were in the theatre where the orchestra is playing”. This feeling is a sign of presence [54].

Particularly in the VR medium, presence refers to a user’s subjective psychological response to a VR system. Presence is what happens when users feel they left the real world and are now “present” in the virtual world; and when the user is “present” their perceptual, vestibular, proprioceptive, and autonomic nervous systems are activated in a way that is similar to that of real life in similar situations. Even though cognitively they are aware that it is not a real life situation, they will potentially behave as if they were, and have similar thoughts [54]. Thus, given the same VR system, different people may feel different levels of presence; and different VR systems may engender the same level of presence in different people.

2.5.2. IMMERSION

Immersion stands for what the VR system's technologies deliver from an objective point of view; the more a VR system delivers hardware technologies that preserves fidelity in relation to their equivalent real-world sensory modalities, the more it is immersive [54]. Fidelity refers to the degree to which the synthetic stimuli produced by the VR system corresponds to real world sensory stimuli. For instance, in the real world, we observe the world in stereo with a field of view of approximately 180 degrees; we hear the world in stereo, which allow us to determine the direction from where the sound is coming; and we have thermoreceptors in our fingertips, through which we sense temperature changes whenever our hands get closer to a warm or cold surface or object. In turn, the closer a VR system's display is to matching such levels from the real world, the higher the fidelity will be for these corresponding components. A VR system's level of fidelity is thus dependent on its physical output interface.

Just as fidelity relies on the realism of stimuli produced by a VR system, it also depends on the realism of the interactions carried out by the user in a VR system when compared to the actions executed in an equivalent real-world scenario. For instance, in the real world, we physically turn in any direction to view more of our surroundings; while driving, we seat and use our feet to control the speed of the car with the pedals and our hands to steer the car with the wheel; while playing baseball, the batter is standing up and swings a bat in an attempt to hit the ball thrown by the pitcher. Hence, if we can also physically turn to view more of a VE, and perform similar actions while driving and playing baseball in a VE, then the level of fidelity for interaction control would be higher than if we only used mouse and keyboard input to virtually execute such activities. A VR system's level of fidelity is thus also a function of its physical input interface.

In addition, it is worth noting that the term 'immersion' does not stand for a single construct, an all or nothing as the terms immersive and non-immersive suggest; rather, as illustrated in Figure 17, immersion stands for a multidimensional continuum and, thus, it can be said that one system has a higher 'level of immersion' than

another [54]. In turn, a VR system's level of immersion depends on the combination of different hardware and software technologies, such as the system's field-of-view, field-of-regard, stereoscopic rendering, head-based rendering, display resolution and size, among others.




Figure 17: The continuum extended into a two-dimensional plane of virtuality and immersion.

2.5.3. BENEFITS OF IMMERSION

As was previously mentioned, although being conceptually distinct, presence and immersion are strongly related: higher degrees of immersion allows for more realistic virtual experiences, which may more easily induce user presence. On this matter, much of the excitement today for highly immersive VR systems is centered on obtaining high levels of presence. The most attractive benefit of immersion then seems to be related to persuading people to be psychologically in a place different than where they are physically located; such place may be a replica of the real world or may be an imaginary world that could never even exist. But how else can we benefit from VR technology? Is its actual use limited to providing great realistic experiences of a virtual world?

To illustrate this argument, consider the oil and gas industry, which has invested heavily in highly immersive systems in the last decades. These systems commonly provide a simulated representation of the real petroleum reservoir with the goal of performing visual and analytical analysis of complex three-dimensional structures. In

this scenario, presence is not even required since the sense of “being there” does not help experts to perform tasks more efficiently – and the reservoir rendering may even be purposefully unrealistic in order to help domain experts to better understand the reservoir’s critical features.

As can be observed from the example above, presence may not be the only benefit coming from high levels of immersion. As previously mentioned, a VR system’s level of immersion depends on the combination of different hardware and software technologies; hence, different forms of associating these immersive technologies may provide different benefits. In the following subsections, we particularly discuss four potential benefits of immersion that application may take advantage of: increased sense of presence, increased spatial reasoning, increased information space, and more natural interactions.

2.5.3.1. INCREASED SENSE OF PRESENCE

A potential benefit of immersion concerns the possibility of an increased sense of presence – that is, the sense of “being there”. It is useful for applications that highly depend on high-fidelity sensory stimuli for producing a realistic experience that effectively places the user *in* the simulated environment, as illustrated in Figure 18.

One application domain is phobia therapy – e.g., claustrophobia [55], aerophobia [56], and arachnophobia [57]. To virtually expose patients to their phobias may trigger the same fear structures in the person’s brain as if they were actually there in reality. This gives the therapist an opportunity to help shape the fear to a manageable and reasonable level. Moreover, phobia therapy using VR systems may be less expensive, time consuming, risky, and embarrassing.

Other application domains include military and medical training [58] [59] [60], and entertainment [61]. In military or medical training, the goal is to produce training transfer, that is, to cause the trainee to do the right thing when faced with situations in the real world. The goal is the proper mapping between the sensory stimuli

and the appropriate response. It makes sense, then, that training will only be as effective as the training system's sensory stimuli are realistic.



Figure 18: Some application domains rely on the realistic experience that VR systems may provide to the user, such as military and medical training (top) and phobia therapy (bottom).

2.5.3.2. INCREASED SPATIAL UNDERSTANDING

A second potential benefit of immersion is increased spatial understanding, which refers to the ability of discerning and judging features of objects in space such as shape, volume, position, and orientation, on an absolute or relative basis. Thus, spatial reasoning may relate to different types of tasks, such as searching, pattern recognition, property analysis, among others.

In this realm, although human eyes only senses 2D image projections of the environment, humans perceive a 3D, spatial environment. This occurs as the human brain reconstruct 3D images from the 2D images by exploiting depth cues such as stereopsis, motion parallax, perspective, and occlusion, and the perception of depth and shape are improved when multiple depth cues are present concurrently. In this sense, technologies used in VR systems may provide multiple, integrated depth cues – e.g., stereoscopy and head tracking permit users to apply their built-in capacity for understanding stereopsis and motion parallax [62]. Thus, it is not surprising that higher levels of immersion may lead to a greater spatial understanding.

There is a number of studies on how various systems characteristics affect spatial understanding in VR. A study in particular evaluated small-scale spatial judgments, which require careful visual inspection of objects that are small relative to the scale of the environment [63]. By using a CAVE system, the authors evaluated the effects of field-of-regard, stereoscopy, and head-tracked rendering on the performance – that is, time and number of errors – of the task of identification of collisions and gaps in complex underground cave systems. Overall, the results suggest that the addition of the higher fidelity system features support performance improvements in making precise spatial inspections of complex three-dimensional structures.

Another work also assessed the task of identification of object intersections; however, this work compared separate VR systems. It compared a standard desktop, a fish tank VR that used a desktop monitor, and a CAVE-type system with stereoscopy and head tracking. Again, the results suggested enhancement in performance due to the addition of higher fidelity system components [64]. Other studies evaluated spatial understanding by requiring participants to visually trace paths within three-dimensional graph structures [65] [66]. In general, the findings suggest performance improvements in path-tracing tasks due to the addition of either stereo and head tracking, with the best performance achieved with the combination of both features.

2.5.3.3. INCREASED INFORMATION SPACE

A third potential benefit of immersion is related to an increase in information space. Although high resolution computer monitor displays are becoming more common, the amount and variety of data that can be displayed and analyzed is limited, whereas an immersive environment may deliver increased field of view and field of regard that facilitates a larger information space so a larger amount of information and context may be considered at the same time. Additionally, such environments promote collaboration, allowing groups of people from different backgrounds to work together and to present the work results to stakeholders, managers, or general public.

Such benefits may be very valuable to aid in the exploration and analysis of complex, often large datasets. Particularly, they may be useful for scientific visualization that commonly assist researchers to understand numerical models describing constantly evolving natural phenomena. Such models, or datasets, can be the output of numerical simulations and often contain high-dimensional data in a three-dimensional volume.

As an example, authors in [67] mention that block caving operations generate large datasets from multiple work areas: 3D geologic data, 3D excavations and stress information, 1D time sequencing or mine sequencing, nD monitoring data, among others. These multitude of data need to be effectively integrated into a common mining process model to provide a platform for experts in various fields to extract value from the data collection and, so, to better design, monitor, and modify mine plans throughout a mine's life – from feasibility, to operation, to closure. Given this data challenge, the authors built the Block Cave Mining System Visualizer (BCMSV) software for running in the AVIE [68] environment, a 360-degree stereoscopic visualization system that includes motion tracking sensors and a multichannel audio system.

As another example, several case studies presented in [69] utilized the TESSIN VISLab, a projection-based four-screen VR system (see Figure 19). All case studies faced the need of analyzing and working with increasingly complex datasets generated by simulations of natural phenomena in environmental sciences. Most of them come from domains such as hydrology, geosciences and energy; yet, some projects also include climate prediction, urban planning, and biodiversity.




Figure 19: Climate data running in TESSIN VISLab – images from [70].

2.5.3.4. MORE NATURAL INTERACTIONS

As a fourth potential benefit, high levels of immersion afford more natural interfaces. The term “natural” in the concept of “natural user interface” is deceptive since it is often associated with the user interface and understood to mean a mimicry of the real world [71]. Interestingly, the term is not about the interface at all. Quite the opposite, “natural” refers to how the user acts and feels while using the interface to interact with the virtual world. Effective natural interfaces take full advantage of modern technology and find sweet spots in the interaction design in a way that technology behaves as a sort of appendage to an interaction that feels easy and fluid for users. Additionally, it is interesting to note that achieving the goal of creating an interface that feels natural to its user is not necessarily reached by limiting the design to a mimicry of real world, or by relying on familiar metaphors; rather, natural interfaces can set their own rules, ignore physical laws, and lay out novel interaction metaphors as desired.




Figure 20: Illustrations of spatial interactions for spatial construction – images from [72].

On this matter, although humans have managed to create virtual worlds with 3D structures that so closely mimics the world at present, interaction within these worlds is still difficult at the hands of interfaces built upon 2D interfaces; however, with 3D interfaces, where the dimensionality of the interaction space matches that of our natural space, offers something unique: a high level of expressiveness that better mirrors human capabilities. To illustrate, authors in [72] investigated interactions for performing 3D spatial construction using

four 3D interfaces: tongs to grab, stretch, and rotate objects; a ray gun to select and connected objects in space; a handle to drop objects; and the human hand to trace surfaces and to smoothly deform space (see Figure 20).

3. OIL AND GAS EXPLORATION, DEVELOPMENT, AND PRODUCTION

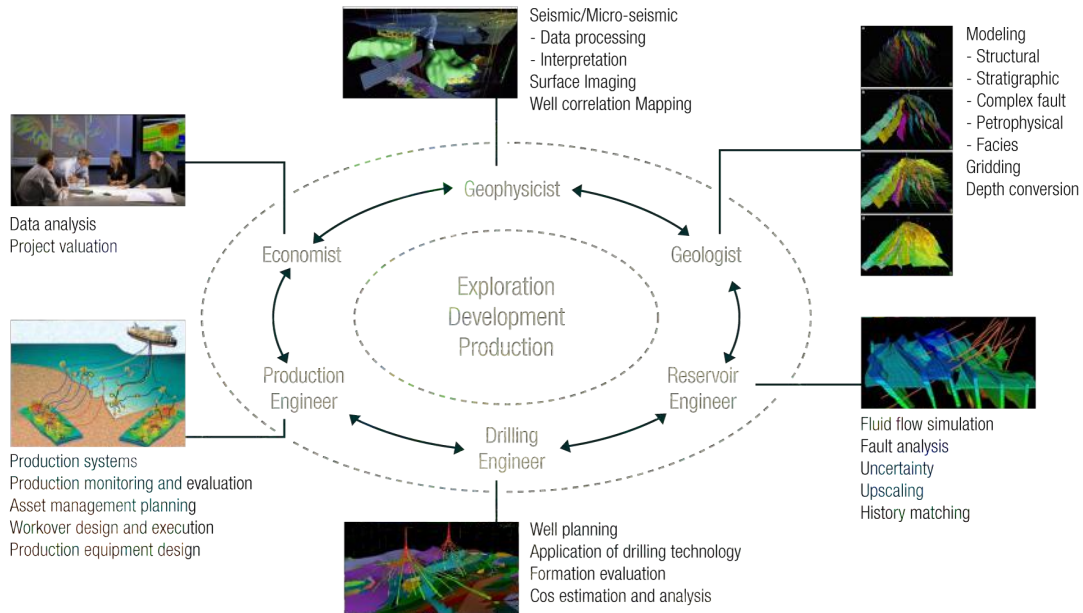


Figure 21: Oil and gas exploration, development, and production stages along with the various disciplines and activities involved.

As mentioned in Chapter 1, this thesis is multidisciplinary in its essence since it aggregates elements from both virtual reality and the oil and gas domain. In Chapter 2, we provided an overview of virtual environments and explained the relevance of immersion and its potential benefits; yet, in order to advance in our work, it is also necessary to get a general understanding of the oil and gas life cycle, and particularly of well placement optimization. This is essential not only to explain the context and relevance of our work, but also to acknowledge the challenges that are particular to this domain.

This chapter will thus strictly focus on oil and gas; we will first explain the three main stages of the oil and gas life cycle – namely the exploration, the development, and the production stages, in Section 3.1. Greater focus will be placed on the development stage as it encompasses the process of modelling the digital reservoir, which is the context in which well placement scenarios are formulated and assessed. Beyond that, emphasis

will be placed on this stage due to the fact that static connectivity analysis can be used as a proxy for dynamic simulation for different activities performed throughout this stage.

Once the three stages have been approached, we will explore well placement optimization in Section 3.2. We will explain how it has been done through dynamic flow simulation, which is the most traditional numerical simulation used to estimate reservoir performance; we will also explain its limitations. Given its deficiencies, we will explore numerous methods that use connectivity analysis as an alternative solution. We will then present our method of performing static connectivity analysis which we expect may help in well placement studies. Finally, we will provide a brief evaluation of local techniques for static connectivity analysis in order to better understand the strengths and weaknesses of this quite new approach.

3.1. OIL AND GAS E, D&P STAGES

Oil and gas are naturally occurring hydrocarbons found in underground geological formations. Oil and gas reservoirs, in turn, are subsurface pools of hydrocarbons encompassed by rock formations. In this realm, the ultimate goal in the oil and gas industry is to explore these reservoirs and discover optimal ways to produce oil and gas – that is, a way to maximize extraction in the most economical and environmental-friendly manner.

In the pursuit of this goal, the oil and gas lifecycle can be broadly organized into three main stages: exploration, development, and production (E, D&P). These are illustrated in Figure 21. Each one of these stages encompasses several sub-stages that comprise workflows with pipelined processes that require processing a large volume of data retrieved from a multitude of domains such as geophysics, geology, reservoir and production engineering, and economics [73].

3.1.1. EXPLORATION STAGE



Figure 22: Seismic studies and exploration wells are some of the efforts performed in the exploration stage.

Since oil and gas are trapped in reservoir rocks situated hundreds to thousands of meters underground – on shore/off shore, it is challenging to find field areas where there may be oil and gas reserves. Hence, the exploration stage aims to determine potential productive areas by assessing the likelihood of reserves in the subsurface of a certain field. For this purpose, as depicted in Figure 22, different assessment techniques can be performed to help gathering data about the referred field. For instance, geophysicists do seismic surveys and analyze the images of the field subsurface produced by seismic echography. Additionally, geologists may define the geological layering of the subsurface of the field, and then pinpoint potential reservoirs called “prospects”. As a final example, exploration wells can be drilled to either confirm or refute if an identified prospect does indeed contain petroleum or gas reserves.

3.1.2. DEVELOPMENT STAGE

Once the existence of reserves has been confirmed, three-dimensional, digital gridded reservoir models are elaborated and used to simulate the reservoir fluid flow behavior in different field development scenarios for production forecasting and planning. In this sense, as illustrated in Figure 23, the development stage addresses matters such as geological reservoir modelling, reservoir simulation modeling, static and dynamic uncertainty

analysis, and field development planning. Each one of these subject matter is further explained in the next subsections.

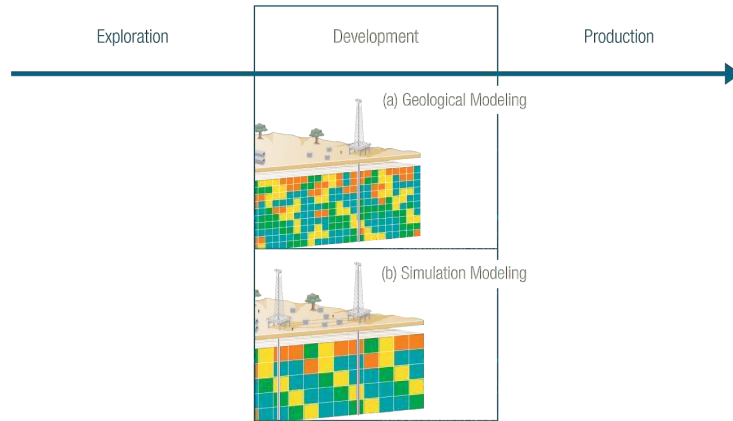


Figure 23: The development stage includes procedures such as geological and simulation modelling

3.1.2.1. GEOLOGICAL RESERVOIR MODELLING

By using the petrophysical, geological, and engineering data acquired during the exploration stage, geologists are able to elaborate a three-dimensional, digital model that represents the reservoir from a geological perspective – known as fine-scale, geological, or static reservoir model. For modelling the geological model, geologists have been using well-known commercial software used in the exploration and production sector of petroleum industry, such as Petrel [74], CMG Builder [75], and IRAP RMS [76]. These software typically follow a workflow of steps that are built upon each other. Figure 24 presents a workflow comprised of six steps that uses a surface-based geologic modelling approach. Steps 1, 2, and 3 define the reservoir structural model, whereas steps 4, 5, and 6 convert such model into the final fine-scale gridded model by adding stratigraphic and rock property descriptions of the reservoir. The steps are processed as mentioned below:

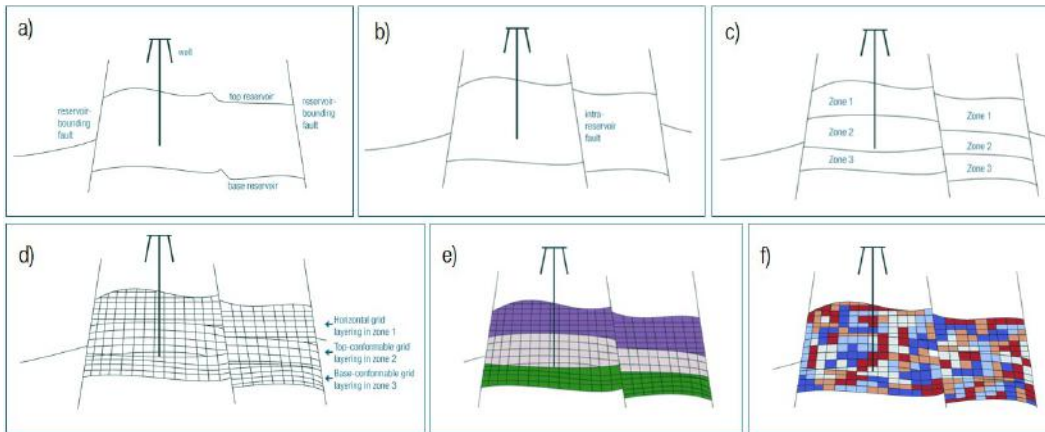


Figure 24: Sequential steps of a conventional workflow of geological reservoir modelling.

- Step 1 – Figure 24 (a): The reservoir volume is defined using surfaces representing the top and base of the reservoir, and surfaces representing key reservoir bounding faults. These surfaces are typically interpreted from seismic data.
- Step 2 – Figure 24 (b): The structural reservoir model is then defined by additional faults within the reservoir – called intra-reservoir faults, which are represented by additional surfaces across which the top and base surfaces may be offset. These surfaces are typically interpreted from seismic data.
- Step 3 – Figure 24 (c): Subsequently, the stratigraphic reservoir model is defined by subdividing the reservoir into geologically defined zones by one or more surfaces, which may be offset across the fault surfaces. These zones may be interpreted from seismic data or correlated between wells, and represent key stratigraphic surfaces, facies boundaries, or lithologic contrasts. Conventional reservoir models may contain tens to hundreds of these surfaces.
- Step 4 – Figure 24 (d): A gridding technique is applied on the stratigraphic model. In general, a grid is delineated within each reservoir zone, which subdivides the reservoir volume into a large number of smaller, discrete volumes called grid cells. The reservoir cells, in turn, typically are approximately in shape; although some gridding approaches can allow more complex geometries [77]. It is worth

mentioning that the aim of gridding a geological model is to turn the model of a continuous field into a discrete structure on which the fluid flow equations can be solved for flow simulations.

- Step 5 – Figure 24 (e): A unique facies code is then assigned in each grid cell, typically using geostatistical methods, such as sequential-indicator simulation (SISIM) [78], object-based modelling (OBM) [79], truncated Gaussian simulation (TGSIM) [80], or multiple point statistics (MPS) [81].
- Step 6 – Figure 24 (f): Finally, unique property values are assigned in each grid cell describing its geology such as rock type, permeability, and porosity. The property distribution typically uses simple interpolation (e.g., distance-weighted averaging) or one of a number of geostatistical methods such as kriging [78], sequential-gaussian simulation (SGS) [82], or multiple-point statistics (MPS), within each of the facies types previously modelled.

3.1.2.2. STATIC UNCERTAINTY ANALYSIS

As previously mentioned, as reservoirs lie under hundreds to thousands of feet on the subsurface, they cannot be seen directly or measured accurately. This makes the collection and modelling of data a challenging task and generates results wrapped in uncertainty. Indeed, uncertainty exists everywhere: in raw data measurements, raw data processing and interpretation, structural modelling, stratigraphic modelling, facies modelling, property modelling, among others. The uncertainty negatively affects the ability to fully understand the reservoir behavior and, thus, elaborate reliable production forecasts and drilling planning.

To reduce the uncertainty particularly during geologic modelling, a generally used approach concerns generating several alternative models – called geological realizations. In this sense, more than merely a mean of generating interpolated values of a property at unsampled locations, geostatistical techniques are also used to model the uncertainty of that property through stochastic simulations [83]. The stochastic simulation is the process of generating alternative, equally-probable images - called realizations - of the spatial distribution of a

property, all of them honoring the data available. A simple illustration of the process of producing geo-statistical realizations is shown in Figure 25.




Figure 25: Three geological realizations displaying the distribution of facies property. Even though there are significant differences between the realizations, parts of the reservoir where real well log data is available are honoured.

Once uncertainty has been modelled through the generation of many alternative models (possibly hundreds or thousands), how are these models processed for prediction purposes? For one, it is not feasible to run reservoir flow simulation on thousands of reservoir models to forecast production – e.g., to evaluate whether digging a new well is worth the cost. Furthermore, it would be extremely tedious to include hundreds of geological models into decision-making scenarios of a reservoir planning. Thus, regarding the last question, there is the need to be selective about the models that are processed for production forecasts.

For this purpose, it is a common approach ranking the stochastic geo realizations. More specifically, the term “applying a transfer function” has been widely used to refer to getting a single geo model and calculating a desired or “target” response from it. If several models have been created, the transfer function needs to be applied to each geo model, resulting in a set of alternative responses that model the uncertainty of the target response. In particular, regarding geo realizations, this target response typically is a volume-based measure such as the original oil or gas in place (OOIP or OGIP), the net porous volume (NPV), the gross rock volume (GRV), or the connected volumes using various connectivity criteria [84].

Afterwards, the models are sorted and a few representative models need to be selected based on the set of responses. Often, one would be interested in three models: a model with a low response – e.g., low estimated oil in place, a model with a “most likely” or typical response, and a model with a high response. More specifically, one often define a statistically representative selection by selecting the geo model whose response is such that 10% of the model responses are lower than the model response of this so-called P10 model – also known as lowest decile. In similar ways, one then selects the median model as the geo model such that 50% of models have smaller response and 50% have higher response, called P50 model. The P90 model is then the antipode of the P10 model.

Thus, considering a case where a transfer function calculated the cumulative oil production as the target response for hundreds of geological realizations, then one would ideally only need to select the P10, P50, and P90 realizations for further processing.

3.1.2.3. RESERVOIR SIMULATION MODELLING

At this point, there is at least one detailed gridded model that describes the static characteristics of the reservoir – that is, its geology. It is now necessary that the reservoir engineer creates at least one model that portrays the dynamic characteristics of the reservoir – that is, the behavior of its fluids. For this purpose, the engineer needs to first integrate into the geological model fluid-related properties such as relative permeability, capillary pressure, and oil and water saturation. These are called dynamic properties as they vary along time; yet, similarly to static properties, unique dynamic property values are stochastically assigned for each grid cell of the reservoir.

Following, the reservoir engineer needs to convert this fine-scale model into a coarse-scale one through a process called upscaling. This is required as subsequent procedures, called reservoir fluid flow simulations, are time and cost intensive when run in large grid models: common reservoir simulations can handle on the order of $10^5 - 10^6$ simulation cells; by contrast, geological models typically contain on the order of $10^7 - 10^8$ cells.

During upscaling, engineers typically apply an averaging method in which the static and dynamic characteristics of fine-scale cells are to be approximated by only one of a coarse-scale cell. Therefore, regions of fine, heterogeneous grid cells are each substituted with an equivalent coarse, homogeneous grid cell. A conceptual illustration of the upscaling process is shown in Figure 26.

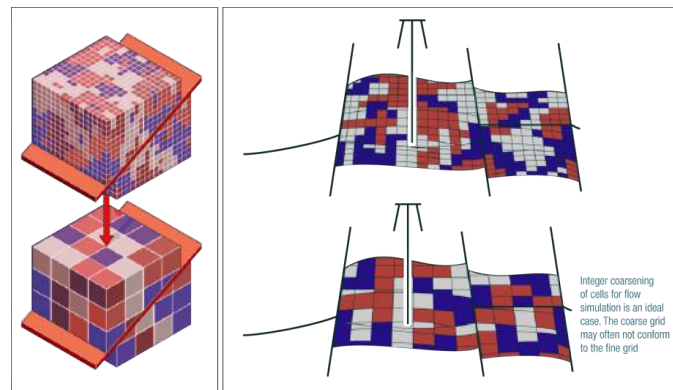


Figure 26: Upscaling is a process that converts a fine-scale model into a coarse-scale one.

At this point, there is at least one extended, upscaled model that represents the reservoir as a whole, with both its geological and fluid-related information – known as coarse-scale, simulation, or dynamic reservoir model. The reservoir simulation model is then used for dynamic uncertainty analysis and/or for field development planning.

3.1.2.4. DYNAMIC UNCERTAINTY ANALYSIS

In general, the uncertainty inherent to reservoir dynamic modelling may come from two different sources. For one, in a similar way to uncertainty generated during geological modelling, uncertainty also arises during simulation modelling as the dynamic property distribution is also driven by stochastic methods. In turn, ranking is once again a commonly used way to gauge this uncertainty; yet, two main factors differentiate the ranking of the geological from the simulation realizations. First, for the ranking of simulation realizations, engineers usually perform flow simulation in order to calculate target responses such as the oil or gas production rate at

breakthrough, or the partial pore volume injected at water breakthrough [85]. Second, as fluid flow simulation is typically used for calculating each target response, the number of dynamic realizations tend to be much smaller than the number of geo realizations.

Regarding the second source of uncertainty found during simulation modelling, as previously mentioned, to ensure that the CPU consumption of a flow simulation will be reasonable – i.e., within practical limits, an average-sized simulation model is built from the detailed static model. In turn, as displayed in Figure 26, because of the high property heterogeneity found in fine models, a common problem in conventional, averaging-based upscaling methods is the tendency to smear out the spatially continuous extremes.

Indeed, as the static models contain meticulous reservoir characteristics that are hard to capture properly with a reasonably coarser model, it goes without saying that some information loss in upscaling is inevitable. In this sense, there is a significant number of research proposing different upscaling methods focusing in minimizing the information loss. Some examples include renormalization techniques [86], pressure-solver methods [87], tensor methods [86], and pseudo function techniques [88].

3.1.2.5. FIELD DEVELOPMENT PLANNING

After the creation of one or more reservoir simulation models, engineers begin the establishment, adoption, and implementation of a field development plan. The formulation of a development plan is an extensive process that involves several decisions points. First, development studies are performed to define possible development scenarios that may produce the reservoir reserves in a technically, economically, and safety manner for both the operating companies and governments involved. In turn, a development scenario needs to define a number of critical information such as:

- the number of wells to be drilled;
- the well placement plan for each well to be drilled;

- the type and cost of installations, such as drilling platforms;
- the treatment systems needed to preserve the environment.

Afterward, the various development options are evaluated and a final decision is made to whether or not to approve one of the engineering studies. In case of acceptance, the sequence for approval is the implementation of the surface facilities and the drilling of development wells. From this point on, the development project enters the production stage.

3.1.3. PRODUCTION STAGE




Figure 27: History matching procedures and drilling of development wells are some of the efforts performed in the production stage.

At this stage, the hydrocarbon starts to be extracted by using drilling wells. In turn, every newly drilled development well provides production data under the current operating conditions, which are used to validate and review the digital reservoir model: geologists and reservoir engineers first analyze the differences regarding the previously simulated and the currently collected production data; then, they manually change one or more parameters at a time in order to improve the match – a process called history matching. As a result, the production performance of the digital reservoir model gets gradually closer to the one of the real reservoir, thus allowing more grounded, well-informed drilling decisions to be made during the revisions of the field development plan.

The time period over which hydrocarbons may be extracted varies between 15 to 30 years and may be extended up to 50 years or more for larger fields [73]. As illustrated in Figure 28, the lifetime of a reservoir is composed of different successive phases: a period of production increase; a stabilization phase, also called plateau; and the depletion period when hydrocarbon production declines progressively.




Figure 28: A common production profile of an oil field development.

Furthermore, petroleum production can be broken down into three phases: primary, secondary, and tertiary (or enhanced oil recovery). During the first phase of hydrocarbon production, the extraction occurs via the natural rise of hydrocarbons to the surface due to the high differential pressure between the reservoir pressure and the bottomhole pressure inside the wellbore; however, as the reservoir pressure declines due to production, so does the differential pressure. The primary recovery phase reaches its limit when the reservoir pressure is so low that the production rates are not economical. In oil reservoirs, typically only 10% of the hydrocarbons are produced during this phase.

Following this, in the second phase of hydrocarbon production, an external fluid (e.g., water or gas) is injected into the reservoir through injection wells that have fluid communication with production wells. The purpose of secondary recovery is to assist the hydrocarbon recovery by maintaining a satisfying reservoir pressure and by displacing hydrocarbons towards the wellbore. This phase reaches its limit when the injected fluid is produced

in considerable amount from the production wells and the production is no longer economical. Altogether, primary and secondary recovery in an oil reservoir produces about 15% to 40% of the original oil in place.

Traditionally, the third phase of production was carried out after the secondary process; however, today it can actually be initiated at any time during the productive life of a reservoir. Tertiary recovery (also called enhanced oil recovery and abbreviated as EOR) seeks to alter the hydrocarbon's properties to make it more conducive to extraction, rather than simply trying to force it out of the ground as did the previous two methods. The most common EOR techniques used are thermal methods, gas injection, and chemical flooding.

In particular, steam assisted gravity drainage (SAGD) is a thermal recovery method in which two horizontal wells are drilled about 5 meters apart, each of which with typically 1 km length. The extraction process consists of steam being continuously injected into the upper wellbore, and heavy oil or bitumen being produced from the lower wellbore. Thermal communication between the well pair commonly occurs within the first three months of steaming. After communication has been established, the steam chamber grows similar to a cone (see Figure 29). And, as the steam rises and expands, it heats up the oil thus reducing its viscosity. Gravity then forces the oil to drain into the lower well where it is produced.




Figure 29: Reservoir recovery techniques: primary (a), secondary with water injection (b), and tertiary with steam injection (c).

3.2. WELL PLACEMENT OPTIMIZATION

The definition of placement of wells, as above mentioned, is a fundamental aspect with major impact in a development project. The precise well placement improves the long-term and short-term performance of wells: maximizes hydrocarbon recovery, extends well life, and, in cases of pressure support from injection, postpone water production. However, optimal well placement is a challenging issue due to the various well plan variables involved and the typical static and dynamic uncertainties of initial fields. In particular, well planning is a function of several parameters such as:

- number of wells;
- location of each well;
- type of each well – injector or producer;
- trajectory of each well – the most common trajectories are depicted in Figure 30;
- perforations of each well;
- recovery technique to be used – primary, secondary, tertiary.

These many parameters, in turn, interact with each other to result in a wide solution space for fitting the well plan. Therefore, engineers must define and assess a number of different placement scenarios in terms of expected reservoir response and production in order to select optimal scenarios. This procedure is the so-called well placement optimization.

In this sense, the reservoir simulation models created earlier are used to describe the drilling plan under evaluation. Following this, the engineer runs fluid flow simulation in order to obtain the predicted reservoir performance for the specified scenario. During optimization studies, ideally, the engineer would progress doing gradual changes in parameters and running fluid flow simulations. Afterwards, the engineer would select the most adequate recovery scenario on the basis of the predicted reservoir performance.




Figure 30: Illustrations of the most common trajectories of wells: conventional wells, which are vertical or slightly deviated, and horizontal wells, which have inclination greater than 80-85 degrees. Horizontal wells can be further classified into multilateral, MRC, fishbone, and extended reach wells – figure redrawn from [73].

3.2.1. USING DYNAMIC RESERVOIR SIMULATION

A flow simulation is any mathematical method that uses physical equations to model the flow of fluid through a porous media. A reservoir flow simulation thus is a numerical module that uses one or more physical laws – e.g., conservation of mass, conservation of energy, and Darcy’s law – to predict the movement of fluids such as oil, gas, and water inside the reservoir in a certain production scenario configured by the reservoir engineer. Therefore, the output of the simulation is the predicted production data. As can be observed in Figure 31, these data are forecasted dynamic property variations over discrete time steps, and may change according to the recovery scenario specified by the engineer.

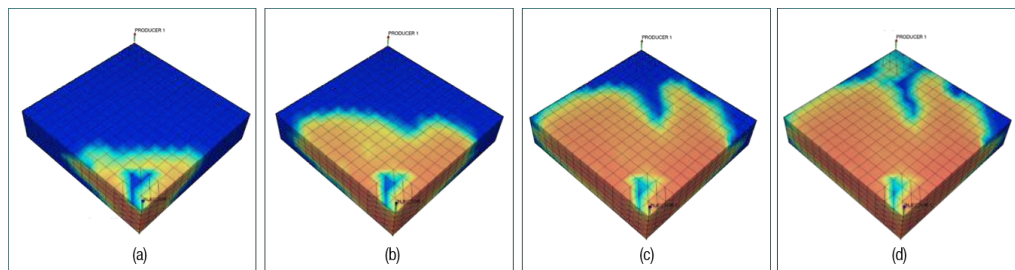


Figure 31: Fluid flow simulation results for water saturation: (a) 80 days, (b) 200 days, (c) 400 days, (d) 500 days - images from [89].

Although it leads to accurate production predictions, the primary disadvantage of flow simulation is that this process is computationally expensive. Depending on the size of the reservoir model, the complexity of the recovery scenario, and the desired result accuracy, it may require hours, days, or even months for the simulation to finish, results gathered, and the drilling plan's performance evaluated. Therefore, it is not feasible for engineers to run multiple simulations due to its computational time and, as a result, engineers are limited to run only a small number of flow simulations before choosing the "optimal" well placement scenario.

For this reason, a number of different solutions to soften this issue have been proposed throughout the years. One solution focused on using supercomputers for achieving high-performance computing [90] [91] [92] [93]. These supercomputers may use distinct architecture approaches – e.g., parallel computing, grid computing, and computer clustering, to form a high-performance basis to handle the large scale computing required during the simulation of fine-scale models with limited or no upscaling; yet, the trade-off for using this higher computational capability is that these are very large, expensive systems.

Another solution focused on the development of fast simulation methods, such as the streamline simulations [94], decline curve analysis [95] [96], capacitance-resistance models [97] [98] [99], and transfer function-based methods [100]. Although these methods provide reasonably accurate results, some of them can still be time-consuming demanding and may take several hours to days for running and obtaining results for one single analysis scenario. And some others, such as capacitance-resistance methods, rely on the available production data gathered during the production stage and, thus, cannot be used to test analysis scenarios neither for initial fields nor for fine-scale, geological models.

Another solution that has been receiving a growing interest investigates methods that use a combination of static reservoir properties to predict the approximate dynamic reservoir behavior. Despite the loss of accuracy due to the simplification of fluid equations, these heuristic techniques can provide information relatively fast and, therefore, can be efficient for sensitivity analysis and for screening reservoir models.

3.2.2. USING STATIC CONNECTIVITY ANALYSIS

In this sense, connectivity is a fundamental reservoir property that strongly correlates with the efficiency at which hydrocarbon is recovered [12] [101] [102]. In primary recovery, if a part of a reservoir is not connected to a producing well then the hydrocarbon present in that region cannot be produced. In secondary recovery using water injection, both producing and injecting wells need to connect to the same reservoir geobody in order to create better sweep zones. Connectivity is therefore a necessary condition for reservoir producibility and, due to its pertinence, there has been a significant amount of research devoted in developing methods to define connectivity as a function of static reservoir properties.

In turn, static connectivity measures have been proposed to be used as fast reservoir performance estimators in different activities in the oil and gas E, D&P, such as the assessment of well placements; the comparison between finer and coarser reservoir models to validate whether structural and stratigraphic characteristics have been preserved during upscaling; the ranking of different geological realizations; among other possibilities. In particular, for the assessment of optimum well placements, static connectivity analysis can be used to assess multiple production scenarios, find promising candidates, and fine-tune only the most promising scenarios by running dynamic simulations on them.

In addition, similarly to the definition of [103], static connectivity measures are here broadly classified either as global or local connectivity. Global connectivity metrics are defined entirely on the basis of static reservoir properties applied to the entire model, referred to as global static characterizations. Local connectivity metrics consider geological properties using more local characterizations in order to capture finer information on the connectivity across the reservoir model and, therefore, to provide a definition on reservoir connectivity that better resembles fluid flow simulations.

This section provides the findings of a survey of analytical functions to perform static connectivity analysis. It starts by describing global connectivity metrics and listing their main limitations. Then, it describes local connectivity measures that try to overcome some of these limitations.

3.2.2.1. GLOBAL CONNECTIVITY METRICS

Global connectivity metrics can be further classified into two distinct groups: those based on deterministic methods, and those calculated by using probabilistic approaches.

3.2.2.1.1. GEOBODY RATIO

Connectivity metrics that are commonly referenced in the literature are the so-called geobody and reservoir to well connectivity. For one, to determine geobody connectivity, one or more rock property values are first defined that differentiate the capacity for fluid to flow within the reservoir. In this sense, a (combination of) facies, permeability, or porosity cut-off(s) is specified so that differentiates rocks in which fluids could flow at some geologically reasonable rates from rocks that are essentially impermeable. For instance, a permeability threshold might be 0.1mD for a relatively light oil, or 10 mD for a heavier oil.

Once potential flow units are defined based on the user specified cut-offs, a mere propagation algorithm can be used to calculate the connected cells in the reservoir. For instance, a propagation algorithm may begin by scanning the grid until the first cell meeting some permeability cut-off criteria is found. It then propagates from this first seed to find all the grid cells that belong to the first group of connected cells. Looping through this process would then enable the algorithm to find the remaining groups.

After defining these groups of connected grid cells, called geobodies or sandbodies, they may be sorted by their volume in decreasing order. Lastly, the geobody connectivity may be defined as the ratio of the volume of the largest geobody to the sum of all the reservoir geobodies as illustrated in Figure 32.

Similarly, reservoir to well connectivity may be defined as the proportion of the volume of the reservoir geobody connected to one or more wells, as illustrated in Figure 32. During the primary depletion phase of producing reservoirs, reservoir to well connectivity refers to the volume fraction of the reservoir connected to producing wells. During the secondary phase on hydrocarbon recovery, this connectivity measure refers to the volume fraction of the reservoir connected to both injecting and producing wells. Thus, the reservoir to well connectivity may represent the drainage volume of oil or gas of the well.

A significant amount of research has been studying geobody and reservoir to well connectivity [104] [105] [106] [107] [108] [109] [110] [111]. Additionally, major differences between these works relate to the volume calculation – since the term volume might either be the gross rock volume, the net rock volume, the net porous volume, or the hydrocarbon volume; and the algorithms for computing and sorting the geobodies. volume of wells.

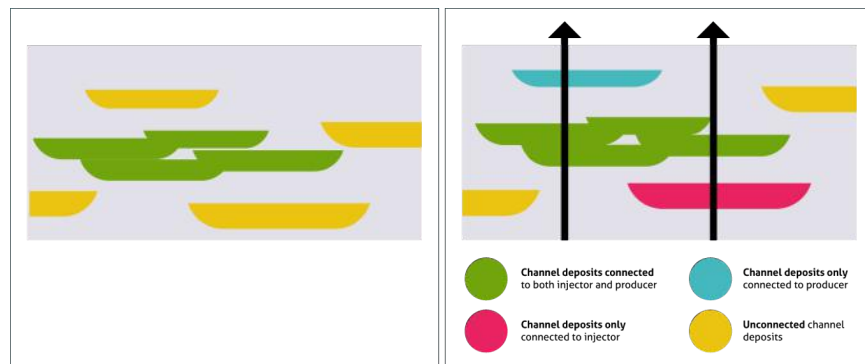


Figure 32: On the left, figurative geobodies representing individual channel deposits (yellow) and an amalgamated channel deposit (green). Geobody connectivity is defined as the proportion of the volume of the biggest geobody (green) to the sum of all geobodies (green and yellow). On the right: reservoir-to-well connectivity is calculated as the proportion of volume of geobodies connected to one or more wells to the sum of all geobodies - figures redrawn from [112].

3.2.2.1.2. PERCOLATED GEOBODY RATIO

Connectivity functions were also defined based on percolation theory. In statistical physics and mathematics, percolation theory studies connectivity and dynamics in random systems, and a typical question addressed by this probabilistic theory is: assuming that some liquid is poured on top of some porous material, is it probable

that the liquid will be able to flow from hole to hole and reach the bottom? Or, in other words, what is the probability that a flow path exists from top to bottom?

In turn, [113] first modelled the problem of connectivity of sandbodies as a percolation theory problem. To illustrate, consider a reservoir grid in which cells can randomly either be sand body with probability p or not with a probability $1 - p$. Now consider that p increases progressively. If $p = 0$, no cell is sand body. If p is relatively small, only a few cells are sand bodies and, thus, these cells are either isolated or in contact with only a few other sand body cells. If p further increases, the density of the sand body cells increases and, thus, small sand body clusters start to appear. If p continues to increase, the small clusters merge together and become larger clusters. However, there is one distinguishing p value when one giant cluster spans across the entire grid. This is called the percolation cluster and the p value is called the percolation threshold p_c . This phenomenon of the rise of the percolation cluster is illustrated in Figure 33.




Figure 33: Three examples of connected clusters colored in red for three different p_c values. For some specific value of p_c , the percolation cluster spans across the entire section – figure from [103].

Once determined the percolation threshold p_c , connectivity is then defined as the probability for a reservoir cell belonging to the single percolating cluster [103]. This may be estimated numerically by computing the ratio of the volume of the percolation cluster to the volume of all the sand bodies. In this sense, many numerical studies were conducted for describing and quantifying connectivity based on percolation theory [114] [115] [116] [117] [112] [103] [102] [118] [119]. For instance, [114] applied scaling laws from percolation theory to estimate the time for a fluid injected into an oilfield to breakthrough into a production well.

It is worth to mention, however, that in most percolation problems it is difficult to find the specific mathematical formulation for the percolation threshold p_c as it depends on a number of factors, such as the grid dimension and geometry, and whether the grid has face, edge, or vertex-connected cells [103]. In this sense, [115] defined four connectivity metrics based on percolation. By using simulations, the paper demonstrated the influence of grain and grid size on these connectivity characteristics, for both the 2D and the 3D grids. In addition, [112] showed the influence of both stratigraphic factors such as net: gross and non-stratigraphic factors, such as well orientation and density, on percolation-based connectivity.

It is also worth noting that this connectivity definition in percolation theory is very similar to the geobody connectivity commonly used during reservoir modelling – the ratio of the volume of the biggest geobody to the volume of all the geobodies. However, it is important to highlight a notable difference between the percolation theory and geobody connectivity. This difference refers to the geobody continuity in reservoir models. On one hand, in a percolation model, all cells are equally probable to be sand body or not, which corresponds to a maximum entropy state. On the other hand, geostatistical techniques attempt to constrain the simulated reservoir properties by reproducing typical reservoir length scales and by using different stratigraphic information. In this sense, sequential Gaussian simulations use variograms to create three-dimensional reservoir continuity, and multiple-point statistics simulations use training images to constrain the geobody structures and their probability of occurrence in specific locations of the model. These additional constraints, therefore, affect the probability that any given cell belongs to a geobody. Consequently, geostatistical models have lower entropy and the reservoir geobodies tend to be much more continuous than clusters in percolation theory [103].

3.2.2.2. LIMITATIONS OF GLOBAL CONNECTIVITY METRICS

As mentioned in [112], locating geobodies and computing the volume percentages are usually done early in life of a field as they may provide a good estimation on the hydrocarbon volume stored in a reservoir prior to

production – that is, the original oil or gas in place. However, once production starts, these volume percentages tend to be an overestimation of the effective recoverable volume, as it is usually only a fraction of the hydrocarbon volume percentage that can be extracted. In turn, this overestimation occurs because conventional measures do not consider issues such as time dependency on recovery, reservoir heterogeneity, and fluid transmissibility across structural reservoir complexities such as faults.

One of the limitations of traditional techniques concerns not defining connectivity as a function of time in order to model fluid behavior through reservoir rocks. Connectivity as a function of time addresses the fact that infinite time is not available to drain all the fluid in the entire channel deposit. As an example, consider a primary recovery scenario in which a channel deposit is connected to a production well as illustrated in Figure 34 (a). Even though the reservoir is completely connected according to the traditional reservoir-to-well connectivity measure, it is not clear whether this channel can be depleted when time is considered. In other words, considering a given time period Δt raises the question of if all the connected grid cells can be recovered within this time period.




Figure 34: Connectivity as a function of time refers to, at any given time, fluid may be actively flowing only in fractions of channel deposits. Illustrations depict influence of time in four scenarios: (a) meandering channel producing by depletion, (b) straight channel produced by water displacement with one injector and one producer, (c) tortuous channel produced under the same conditions as for previous channel, and (d) straight and meandering channel produced together.

As a second example, consider the three secondary recovery scenarios illustrated in Figure 34 (b, c, and d). Given a certain time t_1 , there is a single straight channel in which water displaces oil from the injecting (in blue)

to the producing (in green) well. For the second case, given a certain time t_2 , $t_2 > t_1$, there is a meandering channel producing under the same conditions. For the third case, the straight and the meandering channel are producing together. Further to this, at a given time t_1 , the injecting and producing wells are entirely connected on the straight channel, whereas the sinuous channel is largely less connected – and this will actually remain due to the early water breakthrough on the straight channel.

Therefore, by incorporating connectivity as a function of time, a more conservative definition of the influence of connectivity on recovery can be made. In primary recovery, by considering a time period away from the producing well may set a boundary between the connected and disconnected volume within the geobody. The connected volume may represent the region of the geobody that has already been produced. In secondary recovery using water flooding, a distance cut-off from the injecting well may set a boundary between the region already reached by the water front and the remaining unswept region.

Another limitation of traditional techniques relates to the strong influence of reservoir heterogeneity in the recovery efficiency. In this sense, reservoir heterogeneity refers to the fact that grid cells are heterogeneous formations whose stratigraphic properties differ from place to place. There is research to demonstrate that these variations strongly affect fluid behavior [101] [120].

As an example of the abovementioned, consider the two scenarios illustrated in Figure 35 in which oil is being produced using water flooding. In Figure 35a, the reservoir is somewhat homogeneous – also called isotropic, while in Figure 35b, the reservoir displays greater permeability variability and, so, is more heterogeneous – also called anisotropic. In addition, this reservoir displays a thief zone of high permeability that causes early breakthrough at the well and, as a result, lower sweep efficiency. Early water breakthrough may considerably influence the economics of a reservoir development plan as it increases costs associated with cycling water while reducing recovery efficiency.




Figure 35: Illustrations of producing reservoirs using water flooding: (a) pattern water flood, and (b) early water breakthrough during pattern water flooding due to the presence of a high-permeability layer – figure redrawn from [120].

A third limitation concerns not considering the effect of faults on connectivity. Faults are deformations in the reservoir structure and can be treated as hydraulically neutral planes across which rocks can be dissected or juxtaposed against other rocks. The effect of faulting on both migration and entrapment of fluid depends on the geological properties of the strata juxtaposed by the fault and on the structural attitude of the juxtaposed fault cells. The fault-plane sections illustrated in Figure 36 provide a view of the trapping and migration, and they illustrate three simple assumptions: a fault itself has no sealing properties; a fault is not an open conduit; and the trapping and migration relationships at a fault depend upon the fault juxtaposed stratigraphy, i.e., the thickness and spacing of permeable and impermeable units [121].




Figure 36: Illustrations of juxtaposition styles – figure from [122].

As an example, consider the two scenarios illustrated in Figure 37 in which oil is being produced in an unfaulted and a faulted model. According to some conventional connectivity measures, there is no distinction between the two scenarios concerning the volume of the geobody connected to the wells; however, on the faulted

reservoir model, fractions of the connected volumes are sedimentologically dissected by faulting and, thus, cannot be so easily sampled by the well.




Figure 37: Effects of faults on reservoir connectivity – figure modified from [123].

3.2.2.3. LOCAL CONNECTIVITY METRICS

Since predictions built on global characterizations may be limited, there have been attempts to propose measures that use local characterizations to obtain a more factual definition of reservoir connectivity. Such alternate approaches, named on this thesis as local connectivity measures, may be further classified into two groups: those that extend the previous global metrics by considering additional local parameters as graded factors on the calculation of the connectivity; and those that apply propagation algorithms for simulating the expansion of fluid through the reservoir.

3.2.2.3.1. WEIGHTED GEOBODY

A connectivity measure, proposed by (Gutteridge & Gawith, 1996) [124], correlates with the ultimate primary recovery, and is calculated according to the illustrations in Figure 38. For a certain three-dimensional reservoir model in which one or more producer wells are specified, the user must first define a cut-off of impedance in order to define geobodies. Following this, for each producing well w_p , the algorithm specifies incremental vertical shells of cells around the insertion point of w_p . The first shell then consists of the column of cells at the insertion

point of w_p , and has the thickness of one cell. The second shell consists of the column of cells forthwith surrounding the first shell, and so on. The formulation of shells continues according to the extent of the geobodies connected to the well w_p .

For each created shell $S_{i,p}$, the algorithm calculates a shell effective volume $SEV_{i,p}$ by considering the average permeability and net: gross of cells contained in the shell. Finally, the incremental shell volumes are summed over all shells until the extent of the geobodies connected to the well w_p is reached and a final quality factor Q_p is calculated for w_p according to (I). The

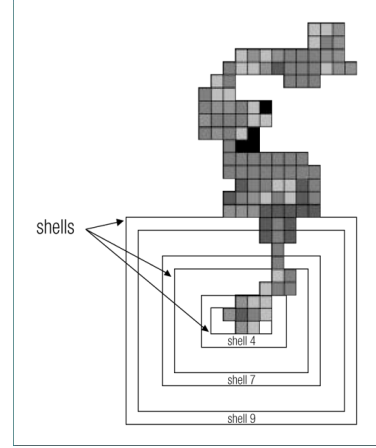


Figure 38: During the quality factor analysis, sand body cells of the models are examined in concentric cylindrical shells around a well insertion point, moving incrementally outwards – figure from [124].

authors mention that quality factor results may be used to rank suggested wells and to select the ones most suitable to be passed to the simulator and turned into actual development plans.

$$Q_p = \sum_{i=0}^n SEV_{i,p} / \sum_{i=0}^n HCPV_{i,p} \quad (I)$$

Another connectivity measure, named quality measure Q_s , is proposed by (Li, Deutsch & Si, 2012) [125] to rank three-dimensional geological realizations. To calculate Q_s for a certain realization in which one or more producer wells are specified, the user must first define thresholds for porosity, permeability, and water saturation in order to calculate the reservoir geobodies. Following this, the Q_s measure for a single geological realization and well design is computed according to (II).

$$Q_s = \sum_{w_p}^{n_w} \sum_{C=1}^{n_{wp}} V_c \Phi_c (1 - S_{w_c}) \left(\frac{d_{max}}{d_{c, wp}} \right)^{d_w} \left(\frac{k_{c, wp}}{k_{max}} \right)^{k_w} \quad (II)$$

As stated in the equation, Q_s is calculated as the cumulative sum of the traditional hydrocarbon volume connected to each production well w_p yet modified by two additional factors: the length of the straight path between each cell c and the nearest production well w_p – named d_{c,w_p} , and the geometric average permeability of the cells along this path – named k_{c,w_p} . According to the authors, these two new parameters may increase the sensitivity to reservoir heterogeneity and, thus, improve the correlation between the connectivity measure Q_s and fluid flow simulation results. Figure 39 shows a scenario to illustrate this assertion.




Figure 39: In this illustration, although the lengths of the straight lines connecting (c_2, w_{p1}) and (c_3, w_{p1}) are the same, the geometric average perms along these lines are different because of the shale zone. Thus, c_2 and c_3 provide different weights when calculating Q_s - figure modified from [125].

Another connectivity measure is proposed by (Fenik, Nouri & Deutsch, 2009) [126] to rank geo realizations and is particularly suited to steam-assisted gravity drainage (SAGD) production. As mentioned by the authors, to calculate the connectivity metric for a certain realization, the user first defines property cut-offs in order to specify geobodies, and a local window size which represents the potential SAGD drainage area. Afterward, considering the geobody that is either connected to both injector and producer wells and within the window, the connectivity metric is calculated as the traditional hydrocarbon volume modified by one factor: the average effective permeability weighted per reservoir layer. According to the authors, this additional factor enhances the correlation with simulation results by considering the reservoir heterogeneity.

3.2.2.3.2. TRAVERSED GEOBODY

Connectivity measures that use propagation algorithms to simulate the expansion of the drainable fluids inside the reservoir focus in either fast-marching methods or techniques that calculate paths of least resistance. On one hand, some connectivity measures are based on the fast-marching method (FMM), which was first

described by [127]. FMM is an upwind finite-difference approximation of the Eikonal equation that computes the evolution of monotonically advancing surfaces on a grid-based structure:

$$F(x) |\nabla \tau(x)| = 1, \quad x \in \mathbb{R}^n \quad (III)$$

With boundary condition $\tau = g(x)$ on $\Gamma \subset \mathbb{R}^n$, where is the τ arrival time, F is the speed function, and Γ is known as the source point – that is, a subset of \mathbb{R}^n in which $\tau = g(x)$. Essentially, this equation calculates the time τ needed for the interface to reach a certain location x as long as both the position x and the speed function F at that position are known.

The basic idea is then to sweep the front surface ahead in an upwind fashion in order to compute the time needed for the surface to reach a certain location. For this purpose, the fast-marching algorithm tags grid points as either *front*, *near band*, or *far away*, depending on whether they are behind the band, near its boundary, or outside the band, respectively. Thus, front points are those already reached by the front and, so, their propagation times are known; while narrow band points constitute the candidates to be subsequently reached by the front. In turn, the smallest narrow band point will be first reached by the front, while its neighbors will be added to the narrow band point set.

Figure 40 illustrates a simple scenario to explain how this method works. Consider a two-dimensional scenario in which the front propagates through a $N \times N$ grid which has cell dimension 1×1 and the speed value in each cell is the same as the cell index – that is, $F_x = x$, where x is the index of each cell as shown in Figure 40a. The front propagates from the grid point with index 13, points tagged as *front* are red colored, points tagged as *near band* are green colored, and remaining points constitute the *far away* set.

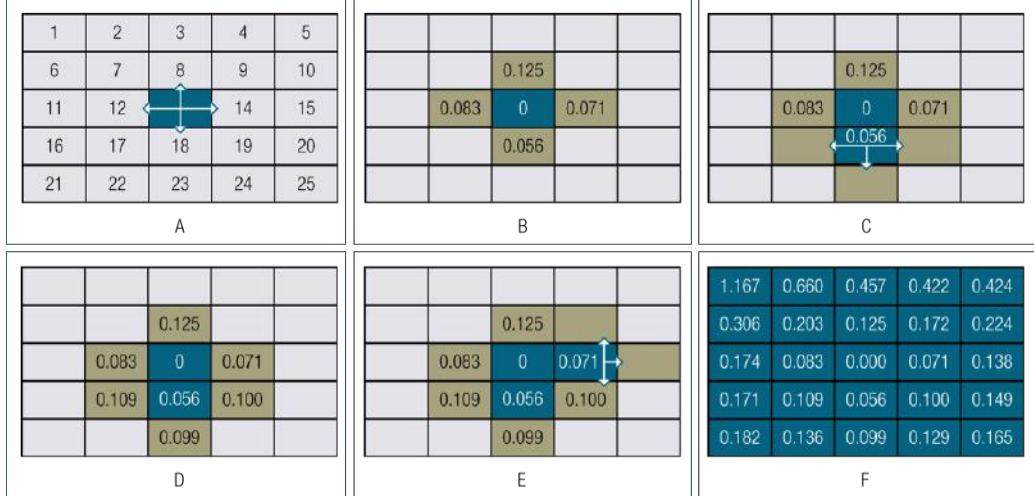


Figure 40: A schematic scenario of the fast-marching algorithm – figure redrawn from [128].

In this sense, a connectivity measure, proposed by (Sharifi et al., 2014) [128], uses the fast-marching method for ranking geological realizations in primary-depletion scenarios. For an anisotropic reservoir model, the velocity function is defined as:

$$F = \sqrt{\frac{k}{\mu c_t \Phi}} \quad (IV)$$

where k is the permeability in x -, y -, or z - direction, μ is the viscosity, and c_t is the total compressibility (summation of rock and fluid compressibility). Once the FMM calculates the arrival time τ_c at each cell connected to a production well, a threshold time value τ_{th} needs to be determined, and then the connectivity measure is calculated as the connected volume to the well considering only cells that have $\tau_c \leq \tau_{th}$.

To test the new methodology, the authors ranked 200 realizations. First, they ran FMM on one realization to choose a reference time threshold: the median time value of this first realization. By using this reference time, the total HCPV was calculated for all the realizations and used as rank measure on the basis of the new algorithm. Furthermore, results from reference flow simulations were used to evaluate the accuracy of the

ranking and production performance predictions. In this sense, the authors found strong correlation coefficients between the results from both methods, either on the ranking results ($R^2 = 0.97$) and on the cumulative oil production results after 10 years ($R^2 = 0.93$).

Other two connectivity measures were proposed by (Karlsen, Lie & Risebro, 2000) [129] and (Berre et al., 2002) [130] that use a fast-marching method for simulating immiscible, incompressible two-phase flow in oil reservoirs such as a water flooding scenario. However, these attempts were not very successful as the advancing-front methods revealed to have limited applicability due to two major problems. The first problem is related to the discretization procedure of the arrival time equation for discontinuous media. The second problem is that the discretization of the arrival time equation lacks the causality principle that the fast-marching method hinges on.

Later on, the work of (Berre et al., 2005) [131] continued the work of (Karlsen, Lie & Risebro, 2000) [129] on a marching method for simulating incompressible, immiscible two-phase flow of water and oil in a porous medium. The authors proposed an adaptive scheme to remedy the deficiencies of the previous work; however, even though the new marching method gives reasonably good results, the added computational cost of the adaptive algorithm is quite high. According to the authors, the efficiency of their method is comparable or even more computationally expensive than streamline simulation.

On the other hand, other connectivity measures were proposed using the concepts of least resistive path and of resistive index (RI), which was derived from the analogy between Darcy's law for linear, single-phase fluid flow, and Ohm's law for linear electric current.

Two connectivity measures were proposed by (Hird & Dubrule, 1998) [132] to estimate reservoir performance for primary and secondary recovery scenarios in two-dimensional cross sections. For primary recovery, the authors propose a connectivity measure to correlate with the drainable HCPV connected to the producing well. First, for a certain reservoir cross section in which a producer well w_p is specified, for each cell

$c_{i,k}$ the least resistive path is calculated between $c_{i,k}$ and w_p . In turn, the authors entitle the resistance metric as the resistive index:

$$I_{Rd}(i, k) = M \left[\sum_{\text{Injector}}^{(i,k)} \frac{\Delta L}{k_a A} \right] \quad (IV)$$

where $I_{Rd}(i,k)$ is the primary recovery efficiency at grid cell (i,k) , ΔL is the distance between the center of the adjacent grid cells, k_a is the average absolute directional permeability between the adjacent grid cells, and A is the cross-sectional area perpendicular to the direction of the movement – that is, for a vertical motion $\Delta L/A = \Delta L_x/\Delta L_z$, whereas for a horizontal motion, $\Delta L/A = \Delta L_z/\Delta L_x$. Then, the user indicates a cut-off value, $I_{Rd,cut}$, above which fluid cannot flow. Finally, the connectivity measure is calculated as the sum of the HCPV of all cells that have I_{Rd} less than $I_{Rd,cut}$.

For secondary recovery, the authors propose a connectivity measure to correlate with the flooded HCPV at water breakthrough. First, for a certain reservoir cross section in which an injecting w_i and producing w_p wells are defined, the RI is calculated between each cell $c_{i,k}$ and w_i :

$$I_{Re}(i, k) = M \left[\sum_{\text{Injector}}^{(i,k)} \frac{\Delta L}{k_a k_w(i) A} \right] \quad (V)$$

where $I_{Re}(i,k)$ is the secondary recovery efficiency at grid cell (i,k) and $k_w(i)$ is the estimated relative permeability to water for the i th column. Following this, the algorithm selects the producer cell with the lowest I_{Re} value. The authors argue that the producer grid cell that has the minimum I_{Re} represents the least resistive path connecting the injecting and producing wells and it is said to be the “time” when the water breakthrough should begin. The algorithm then uses the minimum producer I_{Re} value, $I_{Re,cut}$, as a threshold value to discard cells that were not

flooded at the referred breakthrough time. Finally, the connectivity measure is calculated as the sum of the HCPV of all cells that have I_{Re} less than $I_{Re,cut}$.

Furthermore, while the previous work applied the RI on 2D cross sections with vertical wells, (Ballin et al., 2002) [133] extended the work to use RI in more general cases: non-regular, 3D grids, with deviated or horizontal wells. By estimating recovery for two real fields producing gas by depletion, the authors compared this volumetric RI metric with two other static connectivity metrics and with streamline simulation results. Reference flow simulation results were used to validate the comparisons. According to the authors, the RI analysis required orders of magnitude less time than finite difference flow simulator, and reproduced the reference flow simulations results significantly better than the other techniques in both study cases. In this sense, the RI metric obtained a correlation coefficient of 0.95 for the gas recovery results, against 0.556, 0.741, and 0.864 for the other three metrics.

3.2.3. USING OUR STATIC CONNECTIVITY ANALYSIS

Our approach to perform static connectivity analysis uses foundations of graph theory due to three main reasons. One reason refers to the fact that graph theory affords a number of different methods that one can use for analyzing connectivity of reservoir models from different perspectives, such as shortest paths, connected components, minimum cut, minimum spanning tree, among others. And, although our current approach focuses on the shortest path method, these different graph algorithms may be investigated in the future for posing different aspects of reservoir connectivity.

For instance, the shortest path problem relates to finding the smallest length path from a source node to a target node. For the assessment of reservoirs, length may not necessarily be distance; it may represent other quantity such as cost, resistance, or time for fluid flow between two cells. Thus, minimum paths may be associated with breakthrough time and the volume accumulated along these paths may be related to sweep

efficiency, given that to the edges are assigned parameters relevant to the recovery processes such as permeability or porosity.

Furthermore, for a connected graph, the minimum cut problem refers to determining the minimal set of edges which, when removed, disconnect the graph into separate subgraphs. In a similar manner, the vertex cut problem is an analogous concept to minimum cut for vertices instead of edges; and the maximum cut problem relates to finding an edge cut (that is, a partition of edges) that is as large as possible rather than being as small as possible for disconnecting the graph. For porous media such as a petroleum reservoir, these cuts may represent potential bottlenecks whose loss may severe the continuity at these locations, thus prohibiting the flow of fluids and disconnecting potential drainable volumes.

As a second reason, since graph theory affords unique data structures, graph algorithms are fast and efficient for solving many problems modeled as graphs, such as the assessment of connectivity for different well placement scenarios. As a third reason, finally, it is worth noting that since a graph representation of a reservoir model is simply a matter of describing connections between neighboring cells and assigning property values to the connections, it is not limited in the types of models that are amenable to analysis. In this sense, graph theory may handle models of different dimensions; treat irregularities such as corner-point and unstructured gridding, faults, truncation, etc.; deal with both conventional and unconventional well trajectories such as deviated and horizontal wells, and operate them either individually or in pairs or groups.

Therefore, our approach adopts algorithmic graph theory to perform computationally efficient static connectivity analysis for primary and secondary recovery scenarios. The method follows the workflow illustrated in Figure 41. As can be observed, the workflow is composed of a number of sequential steps, and each one of them is described in details in the next subsections.




Figure 41: Our workflow for performing static connectivity analysis.

3.2.3.1. GRAPH CONVERSION

As illustrated in Figure 41, our method to perform static connectivity analysis first converts the reservoir grid into a graph by converting from a cell-based representation, in which each reservoir grid cell has one or more properties, to a node and edge-based representation. In this representation, each node has the same cell geological properties; however, the edges are associated with new, additional properties related to the connection between nodes. One such edge property is transmissibility. The term transmissibility as used in reservoir simulation refers to a property characterizing the easiness of flow in connections between two adjacent grid cells.

Our approach uses a transmissibility calculation algorithm that considers corner point grids as the discrete digital models of the structures of actual reservoirs. Corner-point models are composed of three main elements: corners, pillars, and irregular hexahedral cells. Cells are defined by eight corner points (vertices) and are the smallest volumetric unit in the grid. Corner points contain spatial position (x, y, z) and are identified by integer

coordinates (i, j, k) , where the i and j coordinates span for each layer, and the k coordinate runs along the pillars, as can be observed in Figure 42.

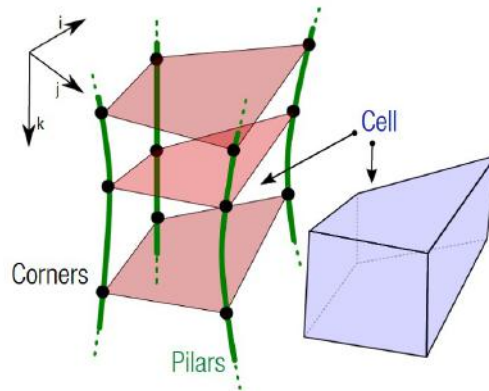


Figure 42: The three main components of corner-point geometry: corners, pillars, and cells – figure from [9].

Since corner points do not have to be regularly spaced nor spatially continuous, degenerated cells can be produced during the modeling phase, and some may also completely disappear, introducing connections between cells that were not initially neighbors. As illustrated in Figure 43 (right), if some of the corner points along the pillars collapse ($C1 = C5$, $C2 = C6$, $C3 = C7$, or $C4 = C8$), but not all four corners, then the cell is said to be partially pinched-out at one or more corners; and if all four corners are pinched, the cell is considered to be totally pinched-out and, therefore, it no longer participates in the simulation.




Figure 43: Faults (left) and pinch-out cells (right) are typical structural irregularities incorporated in corner point grids – figure modified from [9].

Additionally, these particularities of corner-point grids also allow for easily introducing discontinuities across faces, such as fractures and faults. As can be seen in Figure 43 (left), when there is a fault – i.e., a displacement within one or more rock layers as a result of earth movement, more than one corner will be sharing the same (i, j, k) while they will have different spatial positions.

In this sense, authors in [134] proposed an inter-cell transmissibility calculation scheme for corner point grids that is used by most reservoir simulators. To briefly present this transmissibility calculation, initially assume a simple case in which a cell A contacts another cell B through its front face, and that the two cells exchange fluid through the common area A_c , as illustrated in Figure 44a. The inter-cell transmissibility in the y direction is then calculated from the half connection concept: Darcy's law is used from the cell centroid to the contact interface and from the interface to the adjacent cell centroid.




Figure 44: Illustrations of (a) regular cells with full contact interface and (b) corner-point cells with partial contact.

For this simple case with regular cells with fully coincident interfaces, the half transmissibility factors in y direction can be calculated as:

$$T_A^y = K_A^y \frac{\Delta z \Delta x}{\Delta y_A / 2} = K_A^y \frac{A_C}{\Delta y_A / 2} \tag{VI}$$

$$T_B^y = K_B^y \frac{\Delta z \Delta x}{\Delta y_B / 2} = K_B^y \frac{A_C}{\Delta y_B / 2}$$

Then, the full transmissibility is equal to the harmonic average of the two half transmissibilities:

$$T_{A,B}^y = \frac{T_A^y \times T_B^y}{T_A^y + T_B^y} \quad (\text{VII})$$

Following this, assume a more complex case in which a corner point cell A contacts another cell B through its front face, and that the two cells partially share the contact interface A_c , as can be seen in Figure 44b. For determining the inter-cell transmissibility in y direction, the half transmissibility factor from cell A to the interface needs to be first defined. For this purpose, the vector AF that is parallel to the vector that goes through the centroid of the cell A and the centroid of the frontal face is calculated according to (VIII). The magnitude of vector AF is the length of the straight line that goes through the centroid of the surface built using points C2, C4, C6, and C8 and goes through the centroid of the surface created by points C1, C3, C5, and C7.

$$AF = [(C2 - C1) + (C4 - C3) + (C6 - C5) + (C8 - C7)] / 4 \quad (\text{VIII})$$

In addition, for non-orthogonal grids and/or grids with partial contact, the flow surface needs to be defined. In this sense, there is the need of not only a surface multiplier but also some type of projection of the surface. The effective area considered is then the perpendicular area to the flow between the two centroids of the cells. This area can be thought of as being the area of the face multiplied by the cosine of the angle between the face normal and the module of the center-to-center (AF) vector. The half transmissibility from the centroid of cell A to the shared interface can then be expressed as:

$$T_A^y = K_A^y \frac{A_c \cos \theta_A}{|AF| / 2} \quad (\text{IX})$$

Following this, the half transmissibility of the adjacent grid cell B is calculated and, then, the two values are combined using the harmonic average to calculate the inter-cell transmissibility between A and B, as was previously described.

3.2.3.2. GRAPH FILTERING: GEOLOGICAL PROPERTIES

After the graph conversion, the user may specify (a combination of) cut-off values for static properties in order to define geobodies in the graph. Although it is optional, this is a useful step since engineers usually already have some background knowledge about compositional characteristics of nonproductive portions of the reservoir. As an example, seal or cap rocks are relatively impermeable sedimentary rocks that impede the escape of fluids from the reservoir rock. Common seals include chinks, evaporites, and shales that typically have a permeability of 10^{-1} to 10^{-5} mD, with rare values as low as 10^{-8} mD. Therefore, engineers may define a range for permeability to filter out poor-quality rocks from the connectivity analysis.

It is worth mentioning that if the user does not specify one or more cut-off property values, the entire reservoir is considered to be the only existing reservoir geobody.

3.2.3.3. STATIC CONNECTIVITY

After defining geobodies in the reservoir model, the user may create a new well by specifying three parameters: type – either producer or injector, trajectory, and perforations. For the analysis of primary recovery scenarios, the user defines a producer well; for secondary recovery, the user specifies a pair of producer and injector wells. Dijkstra's shortest path algorithm [135] is then used to simulate the fluid behavior through the geobody connected to the newly wells.

In graph theory, this algorithm finds the shortest path between a given source node s and every other node in the graph. That is, for each other node u , the algorithm finds a path from s to u such that the sum of the costs of its constituent edges is minimized. Figure 45 depicts the algorithm's pseudocode. From this code, it

can be observed that the basic idea is to maintain S , a set of nodes for which the shortest path was already calculated, and Q , a queue of nodes for which a temporary, not necessarily the shortest path was found so far. At each step, the node u with the lowest cost from source is removed from the queue Q , added to S , and all edges from u are relaxed. Relaxing an edge (u,v) means testing whether the shortest path to v found so far can be improved by going through u . If so, the shortest path to v is updated to also consider this edge.

```

cost[s] ← 0 // cost to source node s is zero
for all v ∈ V - {s} do // set all other costs to infinity
    cost[v] ← ∞
S ← {} // S, the set of solved vertices is initially empty
Q ← V // Q, the queue contain all vertices
while Q ≠ {} do // while queue is not empty
    u ← minCost(Q) // select the vertex u with minimum cost. remove it from Q
    S ← S ∪ {u} // add u to S
    for all v ∈ neighbors[u] do
        If cost[v] > cost[u] + cost(u,v) then // if new shortest path found for v through u
            cost[v] ← cost[u] + cost(u,v) // set new value of shortest path
            previous[v] ← u // set u as predecessor node on the shortest path

```

Figure 45: Pseudocode of Dijkstra's algorithm.

Dijkstra's original algorithm does not use a minimum priority queue and has running time $O(N^2)$, where N is the number of nodes in the graph; yet, our approach uses a minimum heap-based priority queue for computing the shortest path since such data structure can lead to faster computing times than a basic queue.

In this sense, a priority queue is different from a "normal" queue because instead of being a "first-in-first-out" data structure, values come out in order of priority. In addition, although a priority queue can be implemented using a variety of data structures such as an array, a linked list, or a binary search tree, the heap

structure provides efficient operations for the calculation of the shortest path. More specifically, because the deletion, addition, and search for a minimum node value should be performed very efficiently to prevent the running time to be very long, the minimum heap structure makes these three operations very computationally efficient.

A minimum heap is, conceptually, a balanced binary tree with the order property that each node value is smaller than the value of its children. A minimum heap-based priority queue may then be implemented by using an array where each item in the array corresponds to one node in the heap:

- The root of the heap is always in array[1];
- If a node is in array[k], then its left child is in array[k*2], and its right child is in array[k*2 + 1];
- If a node is in array[k], then its parent is in array[k/2].

According to this definition, the root of the heap always contains the smallest node value. Additionally, the adding and deleting of node values to the heap are performed very efficiently through two processes of up-heap and down-heap. For the insert operation, the up-heap procedure starts by adding a value to the rightmost leaf in the tree, which corresponds to adding it to the end of the array; then it swaps values up the tree until the order property has been restored. In the worst case, the newly added node value follows a path all the way from a leaf to the root. This work is proportional to the height of the heap, which is $O(\log N)$ as the heap is a balanced binary tree. The remove operation is similar: in the worst case, the down-heap procedure follows a path down the tree from the root to a leaf. Again, the worst-case time is $O(\log N)$.

In our specific scenario, Dijkstra's shortest path algorithm is run on the graph for each well perforation along the well trajectory. For primary recovery scenarios, the method is employed along the producer perforations; for secondary recovery scenarios, the algorithm is run along the injector perforations. In addition, since the transmissibility represents a quantity of conductance for fluid flow between cells, the inverse of transmissibility is considered a quantity of resistance, or cost, and so is used as the edge cost. Thus, as illustrated in Figure

46, Dijkstra's method is applied in order to find the least resistive path between a perforation node and all other nodes that are reachable from the perforation. A node is not reachable either when no transmissible path exists connecting the node to the perforation or when it is outside the extent of the geobody connected to the well perforation.




Figure 46: Dijkstra's algorithm has a propagation criteria that finds the shortest paths for all nodes that can be reached from the perforation node (in black).

After running the shortest path method on a well perforation node, the output is:

- The "connectivity graph" that contains all nodes connected to the perforation;
- The total hydrocarbon pore volume connected to the perforation that represents the maximum recovery volume;
- For each connected node, the least resistive path connecting it to the perforation, and the cumulative cost along this path – that is, the minimum cost to go from the perforation node to this node;
- For secondary recovery scenarios, whether injecting and producing wells are connected;
- For secondary recovery scenarios in which there is inter-well connection, the minimum cost among all producer perforations that have been reached.

The minimum cost parameter associated with each connected node is referred to as "delta time". Not to be confused with absolute time, this measure represents the relative time to reach a node. For the scenario

illustrated in Figure 47, a perforation node w_p is connected to two nodes, n_1 and n_2 , with respective delta times δ_1 and δ_2 . If $\delta_1 > \delta_2$, it means that the path connecting w_p to n_1 is less transmissible than the one connecting w_p to n_2 and, thus, the oil in n_2 would be produced before n_1 .

In addition, for secondary recovery scenarios, the delta time parameter is considered to correlate with the wave front propagation. Therefore, the “wave front” $W(\delta)$ can be seen as the set of nodes that have been reached by the advancing water front at a given delta time δ . Beyond that, similar to [132], the minimum delta time among the set of perforated nodes at the producer is assumed to be a representative breakthrough time. Because the producer node that has the minimum delta time corresponds to the least resistive inter-well path, breakthrough should occur first at that specific location.




Figure 47: Shortest path calculation along with minimum cost (delta time) parameter associated with nodes.

3.2.3.4. GRAPH FILTERING: DELTA TIME

After running the shortest path in each well perforation, the user may then specify a cut-off value, δ_{th} , to filter out each connected node, n , that has $\delta_n > \delta_{th}$. Afterward, given the connectivity graph $G = (V, E)$ associated with the well perforation w_p , the connectivity measure for w_p is calculated as a weighted sum of the HCPV with the inter-cell transmissibility of the remaining connected cells, as according to (X). The overall connectivity measure for the well, or well pair, is the sum of the connectivity measures of all perforations in the well.

$$\sum_{(a,b)} [V_a \phi_a (1 - S_{w_a}) + V_b \phi_b (1 - S_{w_b})] T_{a,b} \alpha_{a,b} \quad , \forall (a,b) \in E \quad (X)$$

where

$$\alpha_{a,b} = \begin{cases} 1 & \text{if } \delta_a, \delta_b \leq \delta_{th} \\ 0 & \text{otherwise} \end{cases}$$

As can be observed, the connectivity measure estimates the performance of a certain well placement by calculating the total HCPV as a function of the spatial distribution of flow resistance to the well. For primary recovery, this measure correlates with the maximum hydrocarbon volume drilled by the producing well at a given δ_{th} . For secondary recovery, this measure refers to the fraction of connected cells located in regions behind the flood front – that is, the regions already contacted by water at δ_{th} .

3.2.4. EVALUATING LOCAL CONNECTIVITY METRICS

In this section, we present an evaluation of the local connectivity techniques introduced in Section 3.2.2.3, and our technique presented in Section 3.2.3. This evaluation is not comprehensive; there might have other relevant concerns to be evaluated when it comes to static techniques that are used as fast performance estimators. This is, however, an initial step to evaluate this relatively new approach of local connectivity techniques.

The evaluation of the methods accounts for different information: if the technique can be used during primary, secondary, or tertiary recovery stages; and if it addresses each of the three limitations introduced in Section 3.2.2.2. For the limitations, in particular, we adopt a three-level response to convey whether the technique fully addresses, partially addresses, or does not address them. Finally, the evaluation of the techniques is formulated per category: weighted and traversed geobody connectivity metrics.

3.2.4.1. WEIGHTED GEOBODIES

For the first category, the techniques presented in (Gutteridge & Gawith, 1996) [124] and (Li, Deutsch & Si, 2012) [125] address primary recovery scenarios while the technique in (Fenik, Nouri & Deutsch, 2009) [126] addresses tertiary recovery in SAGD reservoirs. Further to this, for the first limitation, we consider that the technique in (Gutteridge & Gawith, 1996) [124] partially addresses the issue of modelling connectivity as a function of time. The technique calculates a shell restriction factor per shell surrounding the wellbore. This factor uses an averaging of both permeability and net-to-gross ratio of all cells that are part of the shell. These parameters, in turn, are important for controlling production performance: permeability represents the ease with which fluid flow through rock, while net-to-gross ratio is a measure of proportion of clean sand within a reservoir interval. The higher the net-to-gross value, the better the quality of the sand.

Therefore, each shell restriction factor could provide a notion of relative time taken to pass through and reach the next shell. In Figure 48, permeable depositions on shells S2 and S3 surrounding the wellbore can be noticed; the higher fraction of shale content on shell S4, which contributed for decreasing the value of the associated factor, can also be observed. Thus, by evaluating the restriction factors, one could conclude that it would require more time for fluid to flow through S4 compared to the time needed for shells S2 and S3.




Figure 48: Illustrative example of shell factors providing a notion of time.

By following this approach, connectivity more closely portrays the dynamic reservoir behavior as it may, to some extent, address questions such as “given a certain time t_1 , what is the region connected to a producer well w_{p1} ?”; “given a certain time t_2 in a secondary recovery scenario using water flooding, does the water front already reached a producer well w_{p2} ?”; “given a certain reservoir area, what are the drainage areas for wells w_{p1} and w_{p2} ?”. However, we consider this a partial solution for the time-based issue since the technique does not provide local information: it computes (relative) time values for an entire shell rather than for each connected cell unit.

Additionally, the technique in (Li, Deutsch & Si, 2012) [125] computes an averaging permeability for each connected cell by following the straight line connecting the cell with its associated well. In a similar manner to the previous method, this measure could also represent the relative time required for fluid contained inside the cell to be drained from the well. Thus, we consider that this technique addresses connectivity as a function of time. Also, the technique in (Fenik, Nouri & Deutsch, 2009) [126] calculates an effective average permeability, K_{eff} , for each grid layer in j -direction. This measure is the sum between the arithmetic and the harmonic average of K along the line-of-sight of the producer on x - and z -directions, respectively. The K_{eff} measure, similarly to the two previous measures, may suggest a relative time for fluid movement – in this specific case, for fluid to flow between different reservoir layers. Yet, since the technique computes an effective average permeability for an entire grid layer rather than for each cell unit, we consider the method to partially address connectivity as a time-based function.

For the second limitation, we consider that all three techniques partially address the issue of accounting for heterogeneities in the reservoir. The technique in (Gutteridge & Gawith, 1996) [124] uses the shell restriction factor as a weight factor during the calculation of the shell effective volume. The authors argue that, by considering this local information, the method may determine a better capture of connectivity. On this matter, if permeability were the same at all places and in all directions – homogeneous and isotropic, then measuring

the average permeability per shell would reveal a realistic expansive behavior; yet, it is clear that information is lost when averages take place in anisotropic mediums such as reservoir models. Therefore, heterogeneous characteristics of petroleum reservoirs may not be properly portrayed. In Figure 49, a simple scenario shows how an averaging process may introduce losses of geological details that exist in the reservoir model.




Figure 49: Illustration showing information loss as a side effect of an averaging procedure. On the left, original property values per grid cell; on the right, average property values per shell.

The technique in (Li, Deutsch & Si, 2012) [125] calculates, for each cell connected to the well, the average permeability along the straight path connecting this cell to the well. Although this method considers the local information of the cells, it fails to assume a straight flow geometry between the cell and the well; rather, a more realistic flow behavior would account for rock properties such as permeability and porosity. Figure 34 illustrates cross sectional views of reservoirs with sinuous, permeable channels that contain hydrocarbons. These oil sand bodies are usually results from long-term erosional and depositional processes along sedimentary beds such as watercourses or rivers, and commonly have irregular depositional architectures – e.g., braided, meandering, and distributary channels.

For the technique in (Fenik, Nouri & Deutsch, 2009) [126], the authors argue that the K_{eff} parameter may lead to a better calibration of connectivity as permeability plays a crucial role in SAGD thermal recovery, especially around the wellbore. However, the method is deficient since it inspects heterogeneity only on the three axes of the producer well. As mentioned in [136], the entire SAGD steam chamber is heavily affected by

reservoir heterogeneity. On this matter, rather than the idealized balloon shape usually portrayed in theoretical steam chambers, actual steam chamber growth is much more irregular [137]. In many instances, inclined heterolithic strata (IHS), consisting of interbedded sand and silt deposited on point bars, can be dominated by mud and act as barriers to the vertical steam chamber growth at multiple horizons of a reservoir.

As can be seen in Figure 50, during the thermal recovery process, even thin mudstone beds may substantially slow and impede a steam chamber to grow within a reservoir section. These mud-dominated IHS may then considerably impact performance, increasing the steam-to-oil ratio (i.e., the amount of water needed to produce a unit of oil) and decreasing the amount of oil that can be produced [138]. Therefore, monitoring reservoir heterogeneity within the whole steam chamber is an important step to evaluate the reservoir development potential – by delineating the thickness and the extent of siltstone beds within IHS intervals, identifying potential thief zones, etc. – and designing an optimal reservoir development strategy.

Finally, for the third limitation, from the description of the three techniques, it can be observed that the connectivity calculations do not account for the impacts of faults on the flow of fluids on structurally complex reservoirs.




Figure 50: On the left, a theoretical balloon-shaped SAGD steam chamber; on the right, a more realistic, irregular chamber – figure from [138].

3.2.4.2. TRAVERSED GEOBODIES

For the second category, the technique in (Sharifi et al., 2014) [128] address primary recovery scenarios; methods in (Karlsen, Lie & Risebro, 2000) [129], (Berre et al., 2002) [130], and (Berre et al., 2005) [131] address secondary recovery scenarios; and the remaining techniques address both primary and secondary recovery scenarios. When it comes to the limitations outlined in Section 3.2.2.2, both fast-marching and graph-search based algorithms calculate connectivity in terms of a time: for each cell connected to either a producer or an injector-producer pair of wells, a connectivity measure that represents the easiness of flow between the cell and its associated well is calculated.

Yet, there are two remarks worth mentioning about how these two frameworks, fast-marching and graph-based methods, approximate the distance function for computing arrival times. One of them refers to the fact that graph algorithms such as Dijkstra are dependent on the metric of the graph that is used to discretize the reservoir, which can lead to metrication errors [139]. That is, it concerns to the fact that, for a given graph, the distance computed with graph algorithms may approximate poorly the correct geodesic distance.

As an example, assume we have a rectangular grid made of squares of unity size. Considered as a graph, the graph's vertices are the same as the grid's vertices, and the graph's edges are the edges of the grid's squares. We now want to measure the distance between points P_o and P_f illustrated in Figure 51. In the mathematical field of graph theory, the distance between two vertices in a graph is the sum of the lengths (or weights) of the edges appearing in the shortest path (also called graph geodesic) that connects the aforementioned points. In this sense, Dijkstra's algorithm yields the value 10 for the distance between P_o and P_f in this particular example, highlighted as a red line in Figure 51; in contrast, the correct Euclidian distance between P_o and P_f equals $5\sqrt{2}$, depicted as a green line in Figure 51.

Nonetheless, metrication errors can be improved by increasing the grid's resolution since both frameworks, FMM and Dijkstra, yields better approximations (to the geodesic distance) by using finer reservoir grids. Results

of the discrete graph-search algorithms can also be improved by taking a larger neighborhood (e.g., 6-connectivity neighboring) as a structuring element but they would still be not consistent; there will always be an error in some directions that will be invariant to the grid resolution [140].

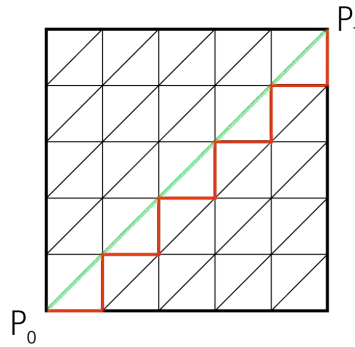


Figure 51: In graph-based algorithms like Dijkstra, paths are restricted to graph edges in way that a vertex is updated from another vertex.

On the other hand, the fast marching algorithm is also a Dijkstra's like graph search algorithm but is consistent in the continuous domain. The consistency condition of the fast marching guarantees that the solution of the discretized equations converges to the true one (the geodesic distance) as the grid is refined. The idea behind FMM algorithms is to compute the distance function as an approximation to the continuous partial differential equation, not the discrete network [141]. This yields better results in the case of typical structured grids used to model reservoirs (see Figure 52).




Figure 52: On the left, Euclidean distance computed with a 4-connectivity breadth-first search algorithms gives square level sets. On the right, the distance is computed with fast marching, giving almost circular level sets – figure from [142].

In regards of the second remark, for finite difference methods that compute arrival times and derived quantities in heterogeneous media, discretization errors may arise due to degenerated elements of the mesh geometry. On this matter, numerical methods commonly rely on the availability of meshes whose elements have the right shapes and sizes; yet, the accuracy of some techniques can be compromised by just a few bad elements that may exist in the mesh geometry [143].

Particularly in the oil and gas domain, it is not uncommon for numerical simulation methods to make assumptions regarding the reservoir mesh geometry. For instance, streamline-based simulation may use Pollock's tracing method [144] to determine the flow path within individual grid cells. Pollock's equation, however, assumes orthogonal grid blocks and, yet, very few real reservoir models use such strict Cartesian framework anymore. Indeed, with the widespread use of more complex grid geometries such as corner point gridding, more recent streamline-based techniques like in [145] handle this limitation by using isoparametric transformation. This method transforms corner-point geometry grids into unit cubes, applies Pollock's method, and then transforms the exit coordinate back to the physical space (see Figure 53).




Figure 53: On the left, Pollock tracing for Cartesian cell; on the right, a quadrilateral is mapped to a unit cell through an isoparametric transform – figures from [145].

In addition, authors in [145] argue that the isoparametric transformation can be applied in either quadrilateral corner-point grids as well as triangular unstructured grids. For the latter case, the authors propose a control-volume scheme: to decompose the triangle into sub-quadrilateral cells and, then, to compute the continuous

fluxes on the interior sub-quadrilateral faces. However, this scheme requires to keep track of the geometrical elements in 3D, which can be cumbersome and numerically expensive. Beyond that, algorithms for mesh generation may produce degenerated elements such as non-convex polygons, polygonal elements with more than four faces, or obtuse triangles for the case of triangle-based meshes.

These poorly generated elements may greatly affect the performance of the method. The authors in [143], for instance, demonstrated that the accuracy of finite element solutions on triangular meshes seriously degrades as angles are allowed to approach 180 degrees. Particularly, the original fast marching method for triangular meshes is suitable only for acute triangles [146]. If this requirement is not satisfied, a situation may occur when computing the distance of the obtuse vertex: the distance values of the two other vertices on the triangle are not yet available, as the propagated wave front could not “reach” the second supporting vertex (see Figure 54b).




Figure 54: The fast marching simulates a front advancing from a source point and stores the computed distance values at the vertices (left). Yet, as the front arrives at an obtuse vertex v_3 , only the vertex v_1 is already processed and has a valid distance value (right). Therefore, special handling is needed for this triangle.

Thus, when a mesh of simplicial elements (e.g., triangles or tetrahedral) is used to form a piecewise linear approximation of a function, the accuracy of the approximation depends on the shapes of the elements. It is thus essential to handle degraded geometrical elements to mitigate discretization errors. As a solution to the issue of non-acute triangles for fast marching methods, [141] proposed a scheme that splits obtuse angles by adding virtual vertices in such triangles. In the case of meshes for finite element’s simulation, there is literature

on meshing and mesh healing to improve the convergence of the numerical method. The authors in [143], for instance, propose error bounds and quality measures that estimate the influence of elements' geometry on accuracy, and that can guide numerical methods and mesh generation algorithms in creating and evaluating meshes.

For the second of the three limitations described in Section 3.2.2.2, all techniques address geological heterogeneities since they consider information from each cell unit for computing the expansive behavior of fluids within the reservoir. These information, in turn, strongly correlate with the fluid flow such as permeability, porosity, viscosity, among others. Finally, for the third limitation, none of the methods but the one proposed in this thesis explicitly state that they account for faulted reservoir models; however, it is worth to mention that there are cases where faults are modeled in the form of impermeable lithologies; thus, fault data can be stored in geological reservoir properties such as permeability and porosity. For these specific cases, all methods could handle these structural reservoir complexities.

4. IVEs IN OIL AND GAS E, D&P



Figure 55: Examples of single display systems utilized for analyzing large, multidimensional data in oil and gas domain.

As mentioned in the introductory chapter, standard desktop-based systems may face limitations in attending specific demands from the oil and gas domain. As a result, there has been a growing effort from the oil and gas industry towards more immersive mediums. In this sense, Chapter 2 examined virtual environments and potential benefits from immersion whereas Chapter 3 explained the oil and gas lifecycle with a focus on well placement optimization. Hence, this chapter builds on the two previous chapters to explore how more immersive technologies have been addressing particular needs from the oil and gas industry.

In the following sections, we will examine how the four benefits of immersion have been advancing the oil and gas industry in general. Each of these benefits will be approached separately as a topic; hence, Section 4.1 will discuss immersion's effects on increasing information space for reservoir data integration and collaboration. Sections 4.2 and 4.3, in turn, will provide an overview on how immersion has affected the potential to increase spatial understanding and to enhance more natural interactions, respectively; these benefits, as will be

explained, are of great relevance in the execution of intrinsically spatial activities throughout the oil and gas workflow. Section 4.4, finally, focus on the significance of an increased sense of presence for construction planning, simulation of industrial facilities, and training in the oil and gas industry. Academic and commercial systems employed in the area will be mentioned as examples throughout the sections. Although such examples will only be mentioned once, they may actually exemplify multiple benefits of immersion at once.

Following this, we will approach how immersive environments have been used for well optimization in Section 4.5. We will then describe systems that tackle well planning; some of these make use of analytical connectivity methods described in Chapter 3. We will then discuss some of the challenges engineers may face using these tools for formulating and assessing well designs. This discussion will provide a foundation for the relevance of our work on more immersive virtual environments for well optimization, which will be introduced in the following chapter.

4.1. FOR INCREASED INFORMATION SPACE

One important benefit of immersion for the oil and gas industry concerns increased information space, which allows users to concurrently view and interact with more multidisciplinary data, as well as to pursue and promote collaborative work with either co-located or remote teams.

4.1.1. FOR DATA INTEGRATION

It is widely acknowledged the large, multi-scale, and multi-faceted nature of the data handled in activities contained in the oil and gas E, D&P. Therefore, one challenge in reservoir data integration is: lots of data, in different scales, and from different disciplines, with a critical need to get the most clarity possible through integration.

For instance, drilling planning should ideally integrate all of the available geophysical and geological data in the same visual environment – e.g., seismic data, horizons, faults, attributes, existing well paths, log data, etc.

to allow more well-informed decisions of placement of wells; however, a single reservoir model may contain millions of grid cells, thousands of exploratory and development wells - for mature oilfields, thousands of fault surfaces to honour structurally complex reservoirs, large amount of production and geological data, etc. In addition, these different data may not be measured at the same scale and resolution – e.g., well logs collected from exploratory wells have higher resolution but are accurate for only a small area around the wells – around 10 feet, while seismic data covers much larger areas but the resolution is several hundred meters at best; yet, the ability to co-visualize datasets collected at these widely different scales, and see their interdependency, improve the return of scientific insight from reservoir data. Therefore, to address the challenge of handling the large amount of multifield reservoir data, two display system setups have been commonly used in both academia and industry.

The first of such setups uses a single large-scale display system to share information altogether. In this sense, large display systems hold the potential to provide a common three-dimensional space to visualize and evaluate these multidisciplinary data in a clear fashion, where any inconsistencies or errors in the data or in the geologic interpretation are glaringly obvious. For instance, in geosciences domain, [147] proposed an efficient setup for large display systems while discussing the relevance of combining 3D along with 2D geoscience data.

According to the authors of [147], during the analysis of geoscience data, it is essential to see the 3D geological structures in their original form; yet, much of the additional data that must be at the user's disposal during discussion are inherently better visualized in 2D, such as areal images, soil columns, log plots, among others. Thus, large virtual display systems are the ideal discussion platforms for investigating integrated 3D and 2D data, specially when coupled with stereoscopy and direct spatial interaction. Figure 55 (bottom right) shows this projection-based display system. On the rear and floor screens, a 3D model is rendered using a computer cluster, and on the side wings, additional 2D information is shown upon the user's request.

In the engineering domain, Environ [148] [149] is an application motivated by the demand in industry to use large industrial models, generated by CAD tools, for planning and simulating industrial projects. In this sense, although supported by desktop systems, Environ targets wide display systems in order to offer a realistic platform for CAD models to be used for virtual prototyping, collaborative design review, ergonomic studies, safety training, among other activities. Figure 55 (top left) shows this application running on a large-screen display.

In addition to academic research groups, oil and gas companies have also been relying on large display systems to address data integration challenges; an example is Barco's MegaCADWall, a wall-sized display installed at visualization centers owned by Kuwait and by OGX companies [150] [151]; and iScope [152], a room-sized display installed at Shell virtual reality center in the Netherlands.

The second of the two display system setups uses an integration of different displays so that they present complementary data. For instance, researchers in [153] and [154] presented a project whose major challenge was integrating microseep data with geophysical and geological datasets. For this purpose, the authors developed a set of interaction techniques deployed in a multi-surface environment – that is, a system where the interaction is divided over several different displays, which includes digital tabletops, wall displays, tablets, and mobile phones. Figure 56 depicts the system using the interaction techniques to display and seamlessly exchange information between tabletops and mobile devices.

Moreover, a system described in [155] integrates a tabletop along with mobile devices. The centerpiece of the system is a tabletop displaying a shared reservoir model, and accessible to all users; and the complementary pieces of the system are mobile devices that display an augmented reality layer that superimposes microseismic events over the tabletop, using a set of embedded markers for tracking. The rationale behind this architecture is that the microseismic dataset complements the understanding of the underlying model, but it may not be as relevant to the entire team as it is to the immediate expert. Following this rationale, every expert may have their

own complementary datasets and additional elements set up in their personal views, with direct spatial association to the reservoir models that aggregates the data.

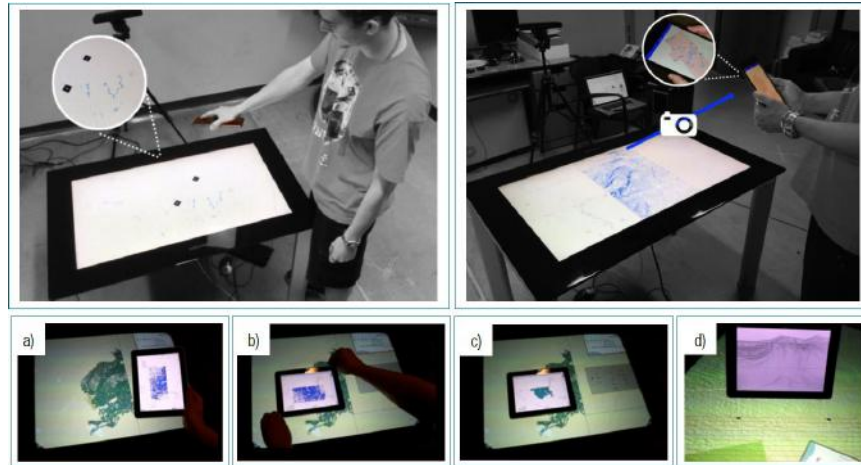


Figure 56: Examples of different interactions for a multi-surface environment. On the top row, the “pouring interaction” transfers data from the mobile device to the tabletop; and the “camera interaction” captures tabletop data from camera images on the mobile device. On the bottom row, by placing a tablet device on top of the tabletop, analytical and graphical data are automatically transferred to the tablet (a), (b), and (c), and by placing the tablet down vertically, the “slicing interaction” shows seismic data in the same orientation (d).

4.1.2. FOR COLLABORATION

Given the extent of their information space, desktop displays tend to direct individuals toward a “sequential” approach to get the job done – the geophysicist interprets the seismic data, the geologist integrates well and other geological data, they define a prospect and potential drilling locations, this is handed off to the drilling and reservoir engineers, etc. This sequential approach leads to inefficiencies in the process because a number of iterations or cycles in the process are necessary to arrive at a conclusion. Beyond that, it is widely acknowledged in the industry that collaborative interchange is not only common but essential throughout the oil and gas E, D&P life cycle [156].

Hence, once again, either multi-display or large display systems may address this issue by allowing an integrated team – the geophysicist, geologist, reservoir engineer, drilling engineer, etc., to view and work with

the same data at the same time. This affords a level of interaction and collaboration which simply does not occur with desktop displays.

For instance, for single display systems, tabletops have been employed due to their large screens and multi-user interaction, which together provide a democratic visualization platform for collaborative manipulation and analysis. In this realm, [157] proposed a set of techniques for supporting the visual exploration of reservoir simulation models in tabletops, as illustrated in Figure 57; [158] explored different interaction metaphors on a tabletop for exploring and annotating LiDAR geological outcrop data; and [159] developed two sketch-based modeling methods to permit experts to sketch reservoir volumes from microseismic point cloud data. The developed tools were integrated with a tabletop to create a collaborative work environment for the experts.

Further to this, authors in [160] [161] proposed a more context-aware tabletop system by integrating additional sensor data. By placing an in-depth camera above the tabletop, users around the table may be individually monitored and, thus, touch points on the tabletop may be associated with individual users. By taking advantage of the user identification for detected touch points, the authors modified a multitouch application for seismic data interpretation by associating a different menu control for each detected user. Therefore, menu-enabled operations such as selection, dragging, and annotation from different users became distinguishable.




Figure 57: Examples of techniques for assessing reservoir models in tabletop systems. From left to right: a method to probe individual cells and display property values; “splitting” and “peeling” gestures to view the inside of the reservoir while maintaining context; and a touch tap to select a well and to activate a cutaway view that removes cells occluding the chosen well path.

In terms of multi-display systems, Slice Inspector [162] interprets seismic data using a system with two main components: a multitouch tabletop, and a stereoscopic display integrated with a 6-DOF wand. Users can thus

create 2D slices of the 3D seismic data using the large display along with the wand, which are transmitted to the tabletop. Users can then create annotations on these seismic planes using the tabletop and results are automatically re-sent to the 3D environment. Similarly, SeisTablet [163] is an application for collaborative seismic discussion and interpretation using multitoch tablets and a shared display. Amongst different functionalities, the application supports direct viewport control of the shared monitor and viewport coupling between all the tablets, where a user can annotate on top of seismic planes and results are transmitted to the shared monitor to be presented to all other users. Figure 58 shows these two systems.



Figure 58: On the left, SeisTablet integrates a shared monitor with tablets that support pen and touch interaction for seismic interpretation. On the middle and right, Slice Inspector combines a 2D multi-touch interface with a 3D large display for the analysis of seismic data.

Furthermore, there have also been efforts to use immersive technologies to support remote collaboration. On this matter, it is interesting to mention that globalization has transformed oil and gas projects into a world-wide endeavor pursued by geographically distributed specialist teams. Offshore engineering projects are an example as they normally engage geographically distributed, highly-specialized engineering teams and require both improved productivity and reduced cost risks when reviewing plants and designs of deep-water oil and gas platforms.

Researchers in [164], for instance, developed a prototype for remote collaboration in close conjunction with an oil company. This allows users to place annotations directly on objects of interest in virtual environments. According to the authors, traditional annotations lack both flexibility and expressiveness to support multi-

threaded discussions; therefore, the proposed prototype provide distinct multimedia options to maximize expressiveness, and a tree-based structure to support traceability of discussions for decision making.

As another example, there is CEE [165], which is a collaborative environment developed to improve the execution of large-scale engineering projects, such as offshore engineering projects. According to the authors, throughout the conceptual design phase of an industrial plant, a number of simulations should be carried out and assessed to ensure the robustness and feasibility of the project. The environment permits collaborative sessions with support for creation and editing of facilities; execution of alternatives simulations by distinct users and have their results displayed; 3D annotations; 3D measurements; 3D manipulation of objects; videoconferencing; among others. The environment also supports virtual guided tours, where a session coordinator user creates virtual paths and annotations through critical spots in the scene, while the same camera movements along the paths are retransmitted to other session participants.




Figure 59: On the top row, the system allows remote experts to add waypoints and to highlight objects to guide novices situated in an immersive VE [166]. On the bottom row, CEE system give users the ability to annotate VEs and discuss them in real time (left), and Inside Reality [167] allows collaborative sessions in which avatars represent remote participants in the VE (right).

According to the authors of [165], guided tours are useful for design reviews, which is the process of inspecting the correctness and consistency of an engineering project, and making the necessary corrections to it – e.g.,

to assess the safeness of different emergency escape pathways in case of an emergency in the plant. In addition, at the end of each session, all users will have all the information exchanged during the session in their local systems. Figure 59 shows examples of systems for remote collaboration in oil and gas domain.

4.2. FOR INCREASED SPATIAL UNDERSTANDING

Another potential benefit of immersion for oil and gas domain is increased spatial understanding for performing activities that may involve search, pattern recognition, small-scale judgements, path tracing, among others. As an example, in the geosciences domain, geological mapping refers to the identification of structures such as faults, folds, layers of lithofacies, and geomorphic features either from data collected in the field, or from digital elevation data, or multi-spectral satellite and photographic imagery. During geological mapping, a geologist has to carry out detailed observations over large areas; view regions of interest from many perspectives and at different scales; measure the location and orientations of often complex undulatory 3D structures, defined solely by their intersection with the 3D surface topography; etc.

Given the challenging activities abovementioned, authors in [168] evaluated the impact of using more immersive technologies on geological mapping. They performed experiments where geologists created a map of the same geological area using two systems: Stereo Analyst, a desktop system that support stereoscopic glasses; and a CAVE-like system along with a 6-DOF wand. Although the informal study was not quantitative due to the small number of participants, the qualitative findings support the authors' hypothesis that the more immersive system enables scientists to make more accurate observations in less time, and to be more confident about their observations. Figure 60 depicts the results from the comparative evaluation.

In addition, examples of commercial software for field based geology and geophysics, as well as for geoscientists working with 3D digital outcrop models, include Virtual Reality Geological Studio [169] and GeoVisionary [170].




Figure 60: From left to right: geologist observing the topographic expression of valleys draining into Lake Tahoe; locations of faults, folds, and drainages precisely mapped directly onto the topographic surface; and results from the comparison between Stereo Analyst (left) and the VR mapping application (right).

In the engineering domain, in turn, making accurate small-scale judgments is important for well planning because it includes tasks such as evaluation of the geometry of objects, identification of object intersections, and determining whether different paths will cross each other or if open spaces exist between objects. Authors in [171] thus proposed an application for well planners to edit paths for wells in a mature field. According to the authors, mature oilfields can be very tortuous underground labyrinths of wells, creating a very complex spatial domain – see Figure 61.

Therefore, when planning a new well in a mature field, the planner must ensure that the new well does not collide with any existing wells – since a collision with another existing well could cause a blowout, which in turn can lead to fires resulting in the loss of the entire drilling rig. Interestingly, the authors examined the application performance in two immersive environments for well path editing. The authors compared task completion times and correctness between participants using an immersive CAVE-like environment and those using a desktop system. The more immersive environment provided stereo and head-tracked viewing, and used a tracked wand for navigation and direct pointing. The desktop version of the application used a stereoscopic computer monitor with a mouse, keyboard, and virtual widgets to support interaction. Although this study found speed and accuracy improvements in the more immersive environment when editing the path of a new well in a mature field, the differences in visualization or interaction techniques that may have caused the performance differences between the two conditions were not further explored.




Figure 61: Images to illustrate how high degrees of immersion may increase spatial understanding during well planning: by viewing regions of interest from many perspectives and at different scales to see their interdependency; and by performing direct measurements of locations, orientations, and intersections of often complex 3D structures.

4.3. FOR MORE NATURAL INTERACTIONS

The third benefit of immersion for oil and gas industry refers to the potential for more natural interactions. In this sense, the traditional desktop systems aim to solve a fundamentally difficult problem that is mapping from 2D interfaces to 3D actions, and an example includes defining a 3D well trajectory within an intrinsic 2D space embodied by the desktop monitor and the mouse. This less-than-ideal mapping between interfaces and tasks, plus the meticulous nature of defining well trajectories, may actually prevent potentially interesting placement options to be explored due to the cumbersomeness of their manipulation.

Hence, well planning analysis could benefit from an alternative, quicker way of manipulating these trajectories, allowing for more exploratory studies to potentially bring new solutions to light and affording better spatial mapping. By using associative 3D interfaces for 3D tasks, it is possible to present to people interaction metaphors that are more natural and intuitive when working with 3D datasets such as petroleum reservoir

models [172]. For this reason, there have been efforts for exploring the spatiality of (a combination of) different input interfaces.

Some research studies have focused on world-grounded devices for leveraging their unique advantage: physical familiarity. That is, interfaces that resemble the real structures by designing physical signifiers and affordances. To illustrate, in the engineering domain, Snakey [173] is a prototype that provides a tangible interface for well path editing. The interface consists of a graspable, semi-rigid physical device which mirrors the shape of an underground well path, and imposes a corresponding digital well path on a conventional display behind the tangible interface, as can be seen in Figure 62. ReservoirBench [174] is an interactive workbench to facilitate education of novice audience to teach them basic geological science and engineering tasks, such as arranging seismic planes accordingly to provided configurations and creating a vertical and/or horizontal well. The interface consists of two main components: a physical workbench that represents an underground petroleum reservoir, and a desktop monitor to display the user's physical interactions with the workbench.




Figure 62: Examples of systems using world-grounded interfaces. On the left and middle, the Snakey object is used for defining well path configurations. On the right, TanGeoMS permits users to hand-build models using material such as foam and modeling clay. Color maps are then projected over the models according to their format.

In the geoscience domain, in turn, TanGeoMS [175] [176] is a geospatial modeling system that couples a laser scanner, a projector, and a physical terrain model. The system projects real-world geographic data onto the

terrain model. The user can alter the model's topography by placing additional, physical objects on the surface of the model. The modified model is then captured by the overhead laser scanner, imported into a geospatial information system (GIS) for analysis and simulation of real-world processes, and the results are re-projected onto the surface of the model (see Figure 62). GeoTUI [177], in turn, is a system designed for geophysicists that provides props as tangible user interface on a tabletop vision-projection system for the selection of cutting planes on a geographical map of a subsoil model.

Furthermore, most systems use tracked hand-held controllers because of the number of potential capabilities, such as direct mapping with hand motions, tracking consistency, user mobility, additional physical buttons, and versatility for performing different tasks. To illustrate, Planwell [178] is a prototype to support remote collaborative drilling operations between an overseer (indoor user) and an explorer (outdoor user) using AR and 3D printing technologies. The overseer uses a tracked pen to draw paths; select areas and points of interests; and annotate data on top of a 3D printed model of a petroleum reservoir surface terrain, and send them to the explorer to provide guidance. The explorer, in turn, receives these instructions in its mobile device as visual annotations superimposed in the real terrain.

As another example, Flying Frustum [179] is a system that allows the control of semi-autonomous unmanned aerial vehicles (UAV) using pen interactions onto a physical terrain model. The user can define a path for the UAV by drawing a line upon the surface of the physical model. Then, video feed from the UAV's camera is superimposed on the terrain printout as a view frustum, which is situated according to the UAV's position and orientation on the real terrain (see Figure 63).

The Interactive Workbench [180], another interesting example, was designed to gather insights into the added value of virtual reality for the oil and gas industry. The system is composed of a large, horizontally mounted monitor that uses active stereo projection, along with a stereoscopic glass and a lightweight pen. By tracking the user's head and the pen held in the user's hand, the user's hands are overlapped by two virtual

hands, and a ray is attached to the virtual right index. The application supports a number of pen interactions: selection and manipulation of objects by using a method based on the HOMER [181] technique; clipping planes, for which the ray changes into a type of flyswatter so the user can define cutting planes on the volumetric data; creating, deleting, and dragging control points for defining and editing well trajectories; among other functionalities.




Figure 63: Examples of systems using tracked hand-held controllers. Flying Frustum uses a tracked pen for defining paths for an UAV to fly along in the real field (a), (b), and (c). CubicMouse is a novel input interface that may be used for editing the trajectory of an oil well (d).

Since the system abovementioned was initially a collaboration between the University of Houston's Virtual Environment Research Institute (VERI) and BHP Billiton's research branch, it was showcased and evaluated by many oil and gas professionals during different conferences, such as 1998 SEG Annual Meeting and 1999 AAPG. BHP also conducted an extensive in-house trial at its Houston office. Amongst the findings from the evaluation, authors mention that the clipping plane was considered to be a powerful tool. Geologists especially enjoyed the ability to place a clipped plane in any position and orientation, and immediately rotate the spatial scene to look at it from a different angle.

Other research studies that have explored wands applied to the oil and gas domain can be found in [182], in which authors developed a device named CubicMouse. It consists of a cube-shaped box with three main components: an embedded tracker to track the device's position and orientation; three perpendicular rods passing through its center, which can be pushed, pulled, and rotated; and nine application-programmable control buttons. CubicMouse was evaluated in different domains: medical field, car industry, and oil and gas industry. In the latter case, the device was used for performing activities such as cutting and slicing planes on volumetric reservoir data as well as for editing well trajectories (see Figure 63).

Additionally, since gestures are a natural supplement to human interaction, there have been some research investigating interactions using the user's bare hands. For instance, the work presented in [183] describes the concepts and implementation of a gesture recognition system; hand motions are captured from a camera streaming and sent as input to a method based on the machine learning technique Hidden Markov Models. This algorithm, in turn, tries to detect predefined gestures from the input data. In addition, the gesture detection system was integrated into an application for seismic data analysis since this work was supported by the VRGeo Consortium, whose members represent companies from the oil and gas industry as well as their software and hardware suppliers. According to the authors, an evaluation of the application using seven gestures showed accuracies up to 98.57% and the review from experts in the fields of virtual reality and geophysics was predominantly positive.

Further to this, authors in [184] adopted gestures and speech commands in order to evaluate the added value of a multimodal interface for the interpretation of seismic data sets compared to unimodal interfaces (see Figure 64). For this purpose, authors first created an application that supports both unimodal and multimodal interaction techniques for various interaction tasks, such as the creation, deletion, translation, and scaling of multiple slices and volume lenses; the navigation and zooming into the spatial scene using an orbit-navigation technique; among others. For performing each single task, unimodal interactions used either speech or gesture

commands individually, while multimodal interactions used a combination of both. Following, the unimodal and multimodal interaction techniques were assessed within a comparative user study. The authors aimed to evaluate whether the multimodal interface would reduce cognitive load. By measuring the time for fulfilling tasks and the number of errors, the findings of the evaluation showed that using the multimodal interaction users became highly significantly more efficient at complex interaction tasks, whereas the efficiency in less complex tasks did not result in a significant difference.

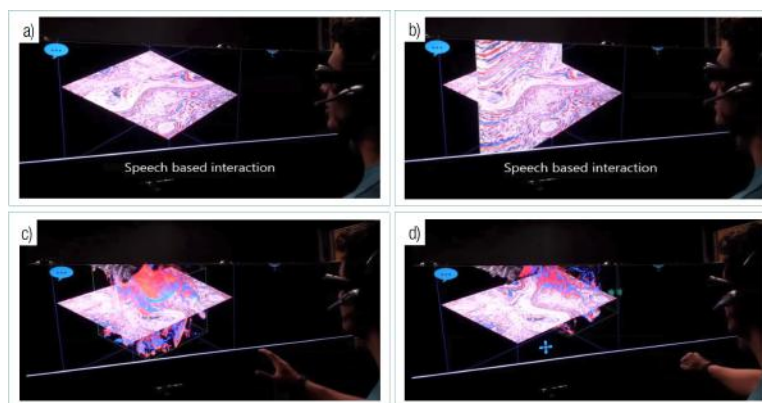


Figure 64: Example of a system using a multimodal interface. The user may create seismic structures using voice commands (a) and (b), and may use gestures to manipulate the seismic volumes (c) and (d).

4.4. FOR INCREASED SENSE OF PRESENCE

Another benefit of immersion is increased sense of presence, which is valuable for performing either construction planning and simulation of industrial facilities, or operations and maintenance training in the oil and gas industry.

Presence can be used for the planning and the simulation of industrial settings. Commonly, a number of different scenarios are designed during the conception of an industrial plant, and simulations are carried out and assessed to ensure the robustness and feasibility of each scenario for the project. Therefore, there are some software for building such realistic scenarios, as well as for running simulations as precisely as possible in order to provide the user a full understanding of the results.



Figure 65: Images of off-shore facilities created and simulated using AVEVA AVP.

An example of the type of software mentioned above is AVEVA Activity Visualization Platform (AVP) [185], which allows the creation of interactive multi-user virtual offshore rigs, allowing for multiple avatars to interact with the various pieces of equipment that would feature on real rigs. Downhole effects such as differential pressure, kick influx density, and bit balling are replicated through the system to further the sense of authenticity. The systems also support the rehearsal of safety and schedule critical activities, without putting either personnel or the plant itself at risk.




Figure 66: Examples of robots available in SimUEP-Robotics.

In addition, industrial robotics have become a trend in the oil and gas industry. On this matter, it is known that the oil and gas industry have seen increasing costs for finding and extracting hydrocarbons, especially in remote locations such as ultra-deep water reservoirs or in hostile environments [186]. These challenging exploration frontiers increase the production complexity and logistic costs. In such conditions, oil exploration is relying more and more on robotization to improve safety and efficiency and to reduce environmental impact. As a result, some software used for designing and simulating industrial plants are considering the use of robotic devices.

For instance, SimUEP-Robotics [187] is a project by Petrobras company [188] aimed to enable planning offshore platform robotizing, as can be seen in Figure 66.

Presence can be fostered for training simulators in order to produce training transfer; in other words, fostered to cause trainees to do the right thing when faced with new software or working environments, or with abnormal situations to enable them to recognize and recover anomalies and malfunctions, especially when they might lead to accidents. As an example, JackVR [189] is a simulator to train technicians to land jackup oil rigs. Jackup rigs are among the most common offshore drilling units for producing oil, but landing the rigs is mostly challenging because of the unpredictable sea and weather conditions, lack of clear vision, and the possible risk of damaging the ocean floor. JackVR then permits engineers to practice landing the oil rig within a safe and semi-realistic training environment. Furthermore, the system provides various superimposed spatial indicators to provide visual warnings on unexpected tasks conditions.




Figure 67: JackVR allows experts to train landing jack up oil rigs. The software uses graphical interface elements to inform them of the status of the simulation.

As another example, the research in [190] explores the use of virtual avatars and collaborative virtual environments to support team training. The agents act as the instructors while the humans act as trainees. Each trainee uses a head-mounted display, interacts with the virtual world via data gloves, and uses a headphone embedded on the HMD to receive broadcasted environmental noises based on the trainee's proximity to targets in the virtual world. The system allows any number of agents or humans to cohabit in the virtual world. The authors present as a simulated training environment the interior of a ship, complete with gas turbine engines,

a variety of consoles, and their surrounding pipes, platforms, and walls. In the most complicated scenario, five team members must work together to handle a loss of fuel oil pressure in one of the gas turbine engines. In turn, this task involves a number of subtasks, some of which are individual tasks while others involve subteams.

Other examples include commercial training platforms from EON Reality and TSC Simulation [191], Siemens Comos Walkinside ITS [192], and DrillsIM-6000 by Drilling Systems [193].



Figure 68: Different learning platforms for drilling and rig operations provided by Drilling Systems.

4.5. FOR WELL PLACEMENT OPTIMIZATION

The previous subsections described a number of virtual reality systems that target different activities throughout the oil and gas life cycle. This section, in turn, discusses systems that are useful particularly for well optimization studies. For each system, we describe how the it supports the design and assessment of different well placement scenarios by addressing three questions: “how does one create a well design?”, “what method does the system use to estimate the performance of the well design?”, and “how the estimation results are evaluated?”.

For instance, in the case of most desktop-based commercial software such as CMG and Petrel, there are typically two possible ways to define a well trajectory. One of them, illustrated in Figure 69, proceeds as follows. First, the user selects a cross-section plane among XY, XZ, and YZ options. Following, the user uses the mouse to select grid cells on the plane to compose the well trajectory. Once cells have been selected, the user may use a slider widget in order to modify the cross section being displayed and, then, continue to select more grid

cells. As a second way of defining well paths, the user may manually enter the three-dimensional coordinates of each of the cells that compose the well trajectory on a separate spreadsheet window.

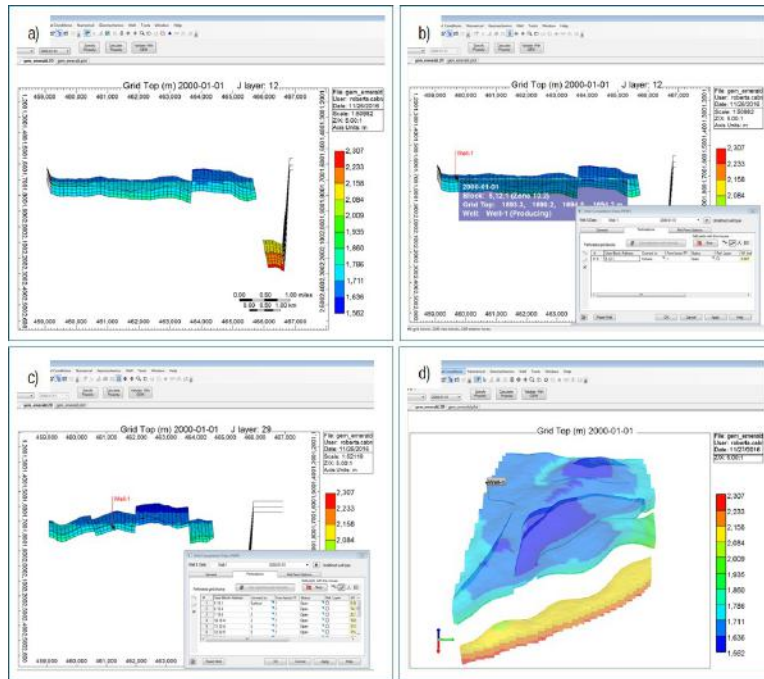


Figure 69: Steps to create well trajectories using CMG Suite: (a) select a new well type (injector or producer) and a cross-sectional view; (b) select the first perforation cell on the top most layer of the reservoir model, (c) if needed, change the plane to select the other perforation cells; and (d) visualize the new well in the 3D window.

After the procedure mentioned above is complete, the user then runs reservoir flow simulation to evaluate the performance of the well placement scenario. And, once the simulation results are calculated, the engineer typically executes two analysis tasks: a) the production profile, which is a production per time graph showing how the oil production is expected to behave through time; and b) the flow behavior, which in existing reservoir simulation software is often represented through visualizations such as isosurfaces or streamlines, as illustrated in Figure 70. Typically, a slider widget is used to move between different time periods to observe how the oil flow changes through time.

If there is a part of the reservoir that the flow does not reach, the user may try to understand why this happened by evaluating the structural features – e.g., “is there a fault or fracture in this region preventing oil to

flow?”, and the reservoir stratigraphy – e.g., “is the permeability low in this region that is prohibiting the oil to flow?”. In this sense, once again the user may need to switch back-and-forth between 3D and 2D views to try to map the problematic region found in 3D with the related 2D cross sections. Once founded, the cross sections may be evaluated using the cell probing feature: by clicking in one grid block, a small information window appears to show the location, IJK coordinates, and the property value of the referred block.




Figure 70: Isosurfaces (right) and flow lines (left) are commonly used to interpret flow simulation results.

As a second example of a system that is useful for well placement studies, there is CONNECT-DynaRank [194], which is a plug-in for Petrel for estimating ultimate primary recovery for ranking geological realizations. As this tool is a plug-in, it uses the features for defining a well trajectory innate to the host software. Therefore, wells may be created by using one of the two ways previously described. Once the user defines a well design for one single geological realization, the same design is considered for all other realizations.

Following the definition of wells, the tool uses the fast-marching method proposed by [128] for computing the distribution of the propagation time of the advancing fluid front for each geological realization. After running the method, each cell is associated with the time needed for the pressure interface to reach it. The user may then specify a time threshold, t_{th} , in order to calculate the drainable HCPV for each geo realization considering solely cells that have $t_c \leq t_{th}$. This measure represents the predicted production performance and is used for ranking the geo realizations.

As the result of the fast-marching algorithm, each geo-realization is colored according to the front propagation time scale, as displayed in Figure 71. As it is inherent to Petrel, each realization is individually displayed in either 3D or 2D cross-sectional windows. Additionally, the ranking results are summarized in a histogram that shows the distribution of the ranking metric and the P10, P50, and P90 models.

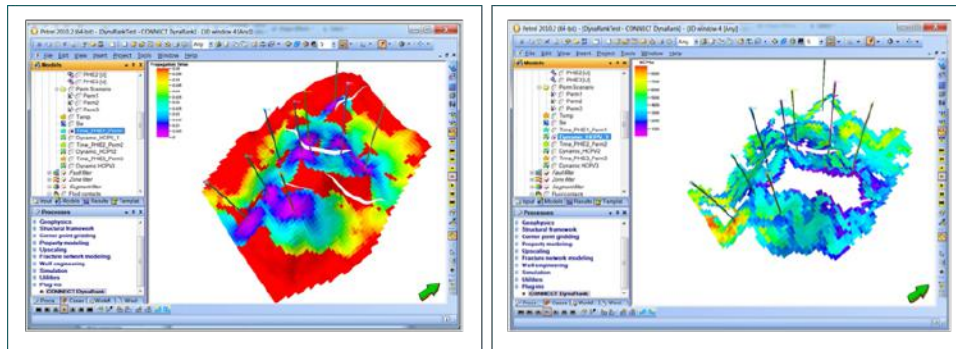


Figure 71: On the left, the model is colored based on the propagation time forecasted from the fast-marching method. On the right, the model is filtered using a time threshold.

As a third example, authors in [103] propose a software to support the placement of vertical wells for primary recovery of oil by using reservoir-to-well connectivity maps. As mentioned by the authors, such maps are 2D characterizations of connectivity, where each point on the map corresponds to one reservoir-to-well connectivity value for one vertical or one horizontal column of the reservoir grid. In this sense, as can be observed in Figure 72, the highest values on the map correlate with maximum connectivity and, therefore, correspond to optimal locations for placing a vertical well. Deviated and horizontal wells, however, are not supported.

In addition, after placing one vertical well by selecting one point on the 2D map, all of the geobodies connected to the newly well are discarded, and a residual connectivity map is calculated with the remaining geobodies to give the best location to place a second vertical well. According to the authors, this iterative process may then guide the placement of a minimum number of wells for covering the producible areas in the reservoir.

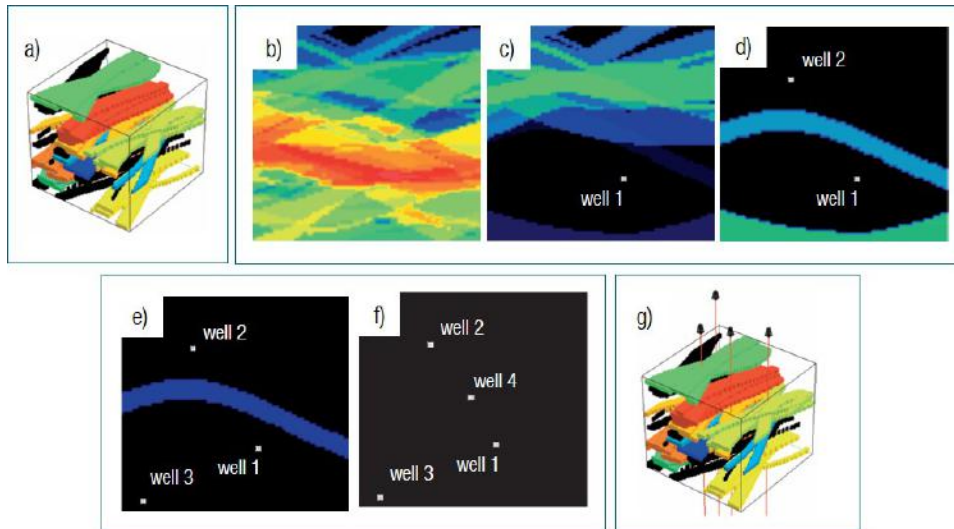


Figure 72: Illustrations of the iterative process for placing vertical wells onto connectivity maps to optimize well placement in a minimum number of steps.

As a fourth example, authors in [195] propose a desktop-based software to evaluate water flooding scenarios using static connectivity analysis. Figure 73 shows two snapshots of the software, which contains four different views. The first view allows one to see data distribution of reservoir properties and to filter the model based on one or more properties in order to determine geobodies. The second view, in turn, allows one to observe the reservoir model either as 2D maps per layer, a 2D average map, or a 3D model. It also allows one to define different well designs that may include injectors, producers, and cross-sectional inter-well paths. In this sense, to create a new well trajectory, one may select a specific grid cell on either the 2D or the 3D view, and then a vertical well is automatically defined for the entire grid column. Deviated and horizontal wells are not supported.

Following the definition of the well design, the software affords a combination of the fast-marching method proposed by [128] along with the shortest path algorithm proposed in this thesis for performing static connectivity analysis. Both methods calculate a time-based flow of the water front propagating from the injector wells to each cell along the inter-well path. The results of the analytical algorithms may be then observed in the

third and fourth views. The third view shows a list of all defined inter-well paths and the fourth view shows a list of paths selected from the third view.

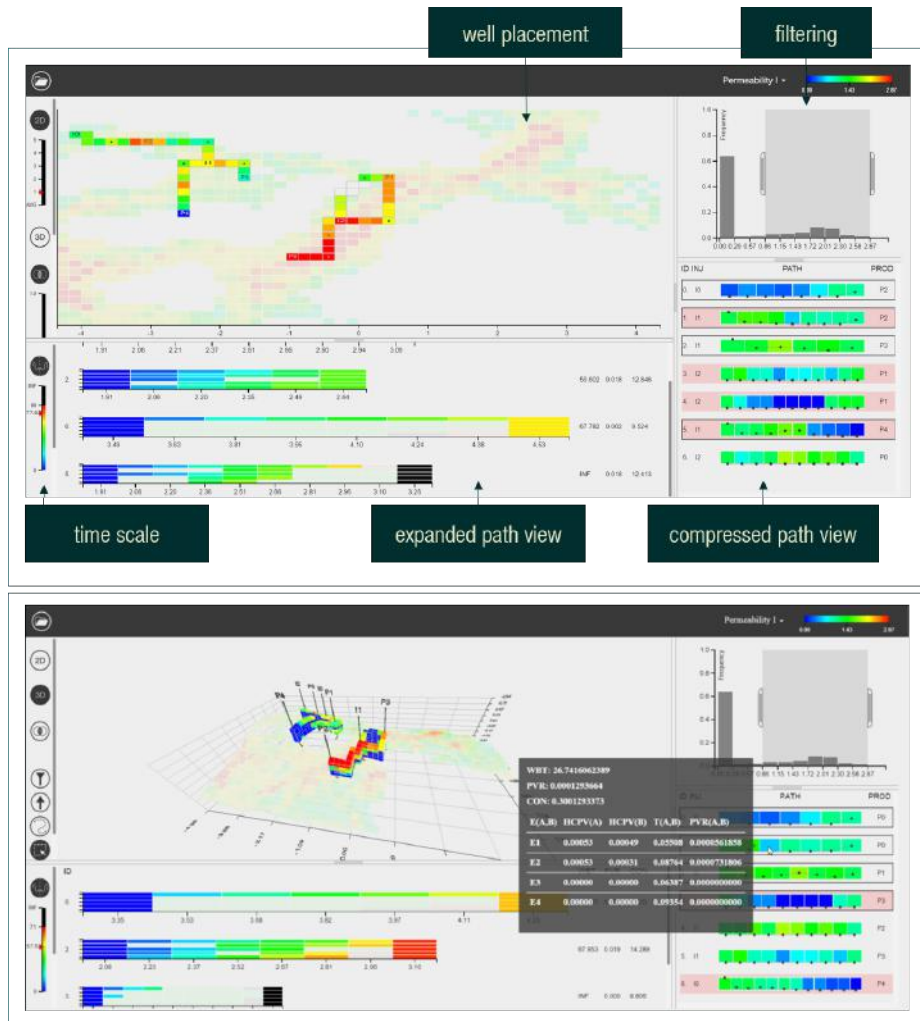


Figure 73: Two snapshots of the software showing the four different views that compose the environment (top) and a tooltip window displaying recovery measures per distance on a particular inter-well path (bottom).

In the third view, each path is portrayed as a row, where each row cell aggregates information of connected cells of one vertical column of the path. In this sense, in Figure 73 (bottom), a tooltip window provides recovery efficiency measures per grid column such as the water front reaching time, the connectivity value, the recoverable hydrocarbon pore volume, as well as further information for each pair of connected cells of the referred column like the inter-cell transmissibility. In addition, paths containing a red background are paths in

which there is no inter-well connection at a given time. The fourth view, as mentioned, shows a list of paths that were selected from the third view; each path is depicted as a cross section colored according to a water flooding time scale, and transparent path cells are cells where no water sweep occurred at a given time. For each path, this view also shows the water breakthrough time, as well as the overall recoverable hydrocarbon pore volume and connectivity value.

As a last example, [196] [197] explored three techniques to create well paths in tabletops. One technique, called the depth buffer based approach, uses 'guide lines' to navigate in space. In this sense, whenever the user selects a point on these guidelines, the 2D coordinates are converted to the corresponding 3D coordinates at that depth value. In addition, when the new 3D position is determined, the center of the guide lines is moved to this position in space – and, thus, the depth values of guide lines are updated. The user can then move over these guide lines to create free form sketches of well paths.

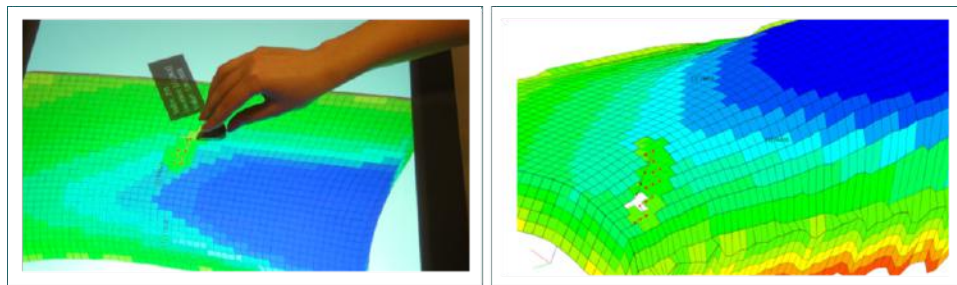


Figure 74: User manipulates a tangible device for defining a well path. By rotating the tangible prop, the user can navigate through different reservoir layers.

The second technique to create well paths in tabletops, called the dive-in approach, makes use of the layered nature of reservoir models to guide the well placement. Given that a reservoir model has a pre-defined number of layers and different layers correspond to different vertical depths, the authors proposed a method wherein the user can control the layer at which the well is being drilled by rotating a tangible device. More specifically, by default, the user starts at depth layer one and, by rotating the physical device by an angle greater than 30

degrees, the user can then move further down the depth layers. Figure 74 illustrates the placement of a well using this approach.

On the third technique, called the 2D planes approach, the user selects a plane (XY, YZ, or XZ) using a selection widget and can rotate it along its axis (X, Y, and Z, respectively) using a rotation widget. Following, the user can create a well path on the plane using its finger (see Figure 75).




Figure 75: User uses translation and rotation widgets for defining a plane and, then, creates a well path along the plane.

Furthermore, the authors of [196] also propose a simple measure, named uncertainty percentage, to support the assessment of different well paths. This measure is a function of the combination of two reservoir properties: capillary pressure and porosity. As mentioned by the authors, the rationale behind the formula is that wells are usually drilled in regions where the pressure is low and the rock is porous, in order to facilitate the recovery of hydrocarbons.

Beyond the definition of wells, the authors of [196] and [197] also propose two techniques for visualizing well trajectories. The first of them adopts foundations of focus and context rendering. In this sense, the user first selects a specific well through a touch tap on its approximate location. Following this, by placing a tagged tangible object on top of the tabletop, the rendering is triggered. This rendering involves a) a v-shaped cutaway that removes the cells located between the viewpoint and the chosen well, and b) a gradual increase in opacity on the remaining cells as the distance from the well increases. If the user rotates the tangible object, the cutaway may open or close (see Figure 76).




Figure 76: The opening angle on the cutaway is increased/decreased by rotating the tangible device.

The second technique, named lollipop-up, offsets the cells along the well trajectory to a certain height in order to turn them visible to the user. The offsets' height, in turn, reflect on the depth of the cells inside the reservoir model: cells closer to the surface have shorter offsets compared to the cells located in deeper locations. In addition, red lines are shown to connect the original perforated cells and offset cells. Finally, the authors propose a technique for evaluating well trajectories named information circles. By selecting a specific well trajectory, its related information circles appear, arranged in a linear fashion, equally spaced, and sorted by size. The size of an information circle, in turn, relates to the uncertainty percentage metric: the radius of the circle corresponds to the number of cells along the well trajectory with a particular value of this property.

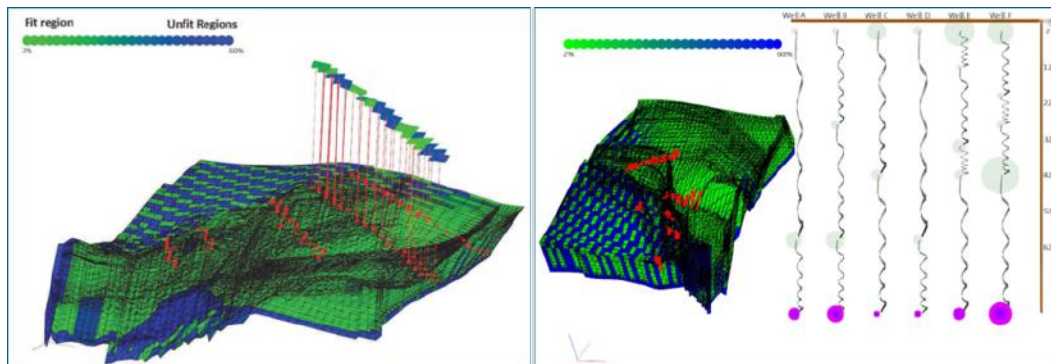


Figure 77: The lollipop-up technique offsets cells to reflect the same shape as the actual well trajectory (left), and the information circles of multiple wells 'hung' from a horizontal pole (right).

When it comes to the abovementioned applications, we can observe two major limitations which may negatively impact both the definition and the analysis of well placement scenarios. First, a strong limitation faced by all solutions is the use of 2D interfaces, such as a keyboard and mouse and a touch surface, for performing

intrinsically 3D tasks. To overcome this issue, the previous solutions commonly design the 3D well trajectory in 2D layouts, which may difficult the understanding on the spatial relationship between the 3D reservoir model and the 2D maps, as well as between the 3D and the 2D representations of the wellbore.

As an example, the works in [195] and [103] support the creation of well trajectories by selecting a grid cell either in the surface of the 3D model or in the 2D maps. This approach may be suitable for more simple configurations such as vertical wells, but it may be challenging to define for more complex, non-conventional types of well trajectories such as the ones specified in Section 3.2. As another example, when defining a well trajectory using Petrel software, the user cannot easily correlate the spatial positions that are being written on the tabular layout with the cells being displayed on the model window, and vice-versa.

In addition, the two well creation approaches proposed in [196] allow users to specify the trajectory by directly selecting grid cells onto the 3D reservoir model. If using the 2D-plane approach, users may easily specify trajectories that accommodate along a single plane such as vertical or horizontal wells; however, it may become time-consuming having to manipulate both the translation and rotation widgets often in order to modify the plane for fitting a deviated, multi-layered well trajectory. In a similar manner, it may become cumbersome having to successively rotate the tangible prop for accessing multiple reservoir layers when creating unconventional wells using the dive-in approach.

As a second limitation, as previously mentioned in Section 4.1, desktop-based solutions may suffer from limited data space. As an example, for assessing connectivity between well pairs, the work in [195] presents the inter-well paths in a list view; however, for even a small number of paths specified, there is insufficient space for displaying all the paths at the same time. Thus, as a result, the user needs to either scroll the list or rearrange the order of the paths on the list to be able to compare multiple paths.

Due to all the challenges mentioned in this section, we will explore in the next chapter how more immersive environments may contribute to increase engineers' efficiency for performing the design and assessment of

well trajectories. In this sense, given the potential advantages of immersion, some of them may be particularly helpful for well planning in the oil and gas domain. These include enhanced spatial understanding, increased information space, and more natural interactions.

5. IVEs FOR WELL PLACEMENT OPTIMIZATION



Figure 78: Our application to support well placement studies running in different immersive mediums.

This chapter presents our exploration on immersive virtual environments to help advance well placement optimization. It thus explains our application, which uses the analytical connectivity method we created and presented in Chapter 3, for performing well placement studies. Such application runs in different immersive environments in the search for efficient visualization and interaction techniques that may ease the tasks associated with well planning.

This chapter will proceed as follows. In Section 5.1, we will start describing the different technologies we use and how these were combined in order to compose three immersive environments. Once we have laid out the hardware components that can be used by our application, we will focus on the software component. Section 5.2 will then provide technical information regarding the programming language and supporting libraries used by our system, and Section 5.3 will describe the design and implementation of the application's features. These can be classified into two groups: general purpose features, such as system control, object manipulation, and

user navigation; and features specific for the creation and assessment of well configurations, such as the reservoir sectioning, cell probing, vertical and free-form well trajectory creation, etc. Section 5.3 will be divided into two subsections: the first of them will discuss the visualization and interaction techniques to use the features when operating a tracked gamepad as input interface; and the second will discuss the techniques when operating a leap motion controller.

5.1. HARDWARE SPECIFICATIONS

As illustrated in Figure 79, our application runs in three immersive environments. One of them consists of a CAVE with a tracked gamepad. The CAVE system consists of four screens powered by projectors with a resolution of 1280x1024. The system achieves stereoscopy using a Volfoni Edge IR 3D glass. The gamepad is a Logitech F710. Tracking of both the glass and the gamepad is accomplished using eight Vicon T160 infrared cameras. The second environment consists of an Oculus Rift DK2 with a tracked gamepad. The same model of controller (Logitech F710) and the same Vicon system are used for gamepad input. The third environment comprises a Leap Motion Controller along with the same Oculus Rift.




Figure 79: Our application explores three immersive environments: a CAVE with a tracked gamepad (left), a HMD with a tracked gamepad (center), and a HMD with a leap motion device attached to it (right).

5.2. SOFTWARE SPECIFICATIONS

The proposed tool is a Unity [198] application written in C# language. Some factors were considered when choosing this development platform. One of them was that Unity allows the developers to write multi-platform applications. A second factor considered was that Unity provides support for integration with MiddleVR [199], a generic immersive library that eases the portability of an application from one VR system to another. As its core, MiddleVR library provides an interface for different aspects of a VR application, such as visual and interaction devices, stereoscopy, and clustering. A third factor considered was that Unity allows the integration with the Hovercast UI toolkit [200], which provides a customizable menu interface for virtual reality applications.

5.3. INTERACTIONS USING GAMEPAD

5.3.1. GENERIC

Even though the application offers a number of domain-specific features, it also supports four generic features that form the basis for the other ones: system control, user navigation, and reservoir manipulation. The next subsections describe each of these features.

5.3.1.1. SYSTEM CONTROL

System control can be defined as the user interface in which a command is issued either to request the system to perform a particular function, to change the mode of interaction, or to change the system state. On this matter, according to the taxonomy for system control techniques for 3D interfaces introduced by [201], 1-DOF graphical menus are the 3D equivalent of the 2D menus that have proven to be a successful system control technique in traditional desktops. Due to their success and familiarity to users, this technique is the most used for system control in virtual reality applications [201], including our application.

It is worth mentioning that there are two main issues regarding the placement of graphical menus in 3D interfaces. First, the placement of the menu influences the user's ability to access the menu – as good

placement provides a spatial reference. Second, the amount of occlusion of the user's visual field. In this sense, the graphical menus used in our application consider both of these issues.

For the first issue, according to the classification provided in [201], we can consider graphical menus that are world-referenced, object-referenced, head-referenced, body-referenced, or device-referenced. World-referenced menus are placed at a fixed location in the virtual world, while object-referenced menus are attached to an object in the 3D scene. In addition, the authors in [201] mention that head- or body-referenced menus provide a strong spatial reference frame and so the user can easily find the menu. Furthermore, [202] explored body-referenced menus and found that the user's proprioceptive sense – that is, the sense of the relative locations of the parts of the body in space – can significantly enhance menu retrieval and usage. According to the authors, the body referenced menus may even enable eyes off usage, allowing users to perform system control tasks without having to look at the menu. Due to its pertinence, the body-reference approach is used in our application.

For the second issue, which concerns visual cluttering, the menu should not overlap too much with the workspace in which the user is working. This is a main concern specially when dealing with VR visualization platforms with small FOV such as an HMD. Thus, the graphical menus used by our application address this issue by using a menu hierarchy in which only two or three nested levels are presented to the user at a given time.

When using the gamepad, the application uses a 1-DOF hierarchical graphical menu. Whenever the user presses the Y button on the gamepad, this view-dependent menu appears positioned and oriented according to the user: it is always displayed in front of the wand held by the user and facing towards the user's head. The user can then select a menu item by hovering it with a ray attached to the controller and pressing the A button on the gamepad.

Furthermore, the graphical menu uses a three-level menu hierarchy. By default, the menu's root label is displayed on the top of the menu interface. Whenever the user selects a 3rd -level menu item that contains sub-items, the selected item then becomes the temporary root on the menu hierarchy and the next-level-down sub-items replace the previously shown top-level menu items (see Figure 80). To navigate back in the menu hierarchy, the user may select the back navigate button near the menu's root label. By using this approach, users may navigate through the menu hierarchy without either entirely losing context where they are in the hierarchy, and cluttering up the available visualization space.

When it comes to the menu settings, the user may use the menu in order to set system configurations such as to switch between interaction modes – either navigation or object manipulation; to set reservoir visibility as visible, semi-transparent, or invisible; to toggle on/off the transparency focal lens; among others. The user may also use the menu to perform certain functions such as filtering grid cells based on property values; modifying the property being displayed on the reservoir; setting the type of well trajectory to be created or deleted; loading different reservoir models or pre-defined scenes; etc.




Figure 80: User filters the reservoir model based on permeability values through the graphical menu.

5.3.1.2. RESERVOIR MANIPULATION

Users can either manipulate the reservoir or navigate through the scene; yet, they cannot perform both tasks at the same time and, therefore, they may switch between these two interaction modes by selecting the desired one on the graphical menu.

For reservoir manipulation, the user may use the gamepad to rotate, translate, or scale the reservoir as illustrated in Figure 81. By using player controls commonly used in avatar-based games as a basis, our interaction metaphors use the left stick allows the user to translate the reservoir on the x- and z- axis; while the left and right bumpers translate on the y- axis. The left and right triggers allow the user to scale the reservoir model up or down. The reservoir is equally stretched in all directions. Finally, the right stick allows the user to rotate the reservoir around the x- and y- axis.




Figure 81: Controller setup for performing reservoir manipulation.

5.3.1.3. USER NAVIGATION

Besides manipulating the reservoir model, the user can also navigate through the entire scene as can be observed in Figure 82. For user navigation, the same controls for translating and rotating the reservoir are employed for positioning and orienting the user's viewpoint within the virtual world; however, the viewpoint translation is performed relative to the direction the user is facing in order to be consistent with the behavior of a first-person view.




Figure 82: Depictions of different subsurface and surface surroundings of our offshore oil reservoir.

This is a fundamental feature since, rather than simply displaying a stand-alone reservoir, our system aims to deliver a holistic, “big-picture” model that covers the entire development field, ranging from subterranean reservoirs to surface facilities. This comprehensive model is considered to be an integrated asset model (IAM) and its value lies in allowing for a comprehensive assessment that accounts for the complexities of interaction between sub-surface and surface domains.

An example of such interdependency is well mapping within facilities. For drilling new wells in a production field, two essential questions need to be answered: “when should a new well be added?” and “which facility will the new well be routed to?”. The current stand-alone modeling approach is able to answer only the first question, but does not have the means to delineate to which facility the new well will be connected [203].

To further illustrate the point abovementioned, consider any typical hydrocarbon field such as the one depicted in Figure 83. The hydrocarbon pathway constitutes a continuous course of actions: the formation fluid flows from the reservoir to the wellbore, up to the well, through the surface gathering network, and then to the processing surface facilities [204]. Despite the interdependency between these processes, it is common practice to create stand-alone asset models for each discipline-related domain. It is also a common practice to

maintain separate models for each reservoir; however, there are many cases where the reservoirs may be producing through joint facilities and, therefore, share common surface constraints. On this matter, although our application is currently displaying only the reservoir model as a realistic asset, it is our intention to incorporate the remaining field data in future.




Figure 83: Illustration of geological formations and facilities of an offshore drilling platform.

5.3.2. DOMAIN-SPECIFIC

Our application offers four field-specific features that engineers may use for performing well placement studies: reservoir sectioning, cell probing, well trajectory creation and analysis. The next subsections describe each of these features.

5.3.2.1. RESERVOIR SECTIONING

Conventional commercial software typically use techniques that display a reservoir model solely as a surface. That is, these traditional visualization techniques render the external faces of the reservoir with their respective property values while removing its internal cells. However, since a reservoir model is a volumetric entity, it is important to evaluate its internal grid cells as well as structures arranged within the reservoir such as well

trajectories and lines of fluid flow. Therefore, visualization techniques capable of depicting and highlighting internal structures of the reservoir model may be extremely valuable, especially during well placement studies.

As discussed in Section 1.1, there have already been some efforts focused on addressing this issue, such as cross sections, cutaway views, exploded diagrams, among others. Our application, in turn, offers a technique through which the user manipulates a volumetric focus lens that partitions the reservoir in regions of focus and defocus in order to allow the examination of internal parts of the reservoir.

5.3.2.1.1. DESIGN

[205] first introduced the concept of three-dimensional focal point p_f from which the focus falls off radially in a three-dimensional scene. In a similar manner, we use the notion of focal point to define a focus volumetric lens for emphasizing and deemphasizing parts of the reservoir model. As illustrated in Figure 84, the lens is defined based on three parameters: the focal point, the front focal distance from the focal point (here called distance near), and the focal length of the lens (here called distance far).




Figure 84: Shading levels indicate focus based on three parameters: focal point (p_f), near (d_{near}) and far (d_{far}) distances.

In turn, the d_{near} and d_{far} parameters segment the reservoir in regions of focus and defocus. More specifically, given a point p on the reservoir and a focal point p_f , if the Euclidean distance between p and p_f is either below d_{near} or above d_{far} , then p is labeled as defocus; otherwise, if the distance between p and p_f is within $[d_{near}, d_{far}]$ range, p is labeled as focus.

Whenever a focal point is specified and a focus lens is set on the reservoir model, emphasis effects are applied on the reservoir grid, two color-related (transparency and saturation) and one line-related (width). It is interesting to note, though, that structures internal to the reservoir such as well trajectories and lines of fluid flow behavior are not affected by the focal effects.

When using the gamepad, the interaction technique for defining the focal point depends on whether the user is spatially positioned outside or inside the reservoir model. When the user is outside the model, we use a ray casting technique in a way that the position in which the ray intersects the reservoir is considered to be the focal point. When the user is inside the model, the focal point is set to be the user's spatial position, as can be observed in Figure 85. Note that the focal point is dynamically updated whenever the intersection point or the user's position changes.




Figure 85: Placing emphasis when the user is outside (left) and inside (right) the reservoir model.

Additionally, there are two ways of setting the $[d_{near}, d_{far}]$ parameters for each rendering effect: use slider widgets on the graphical menu, or use the D-Pad buttons on the gamepad to set d_{near} and d_{far} for the currently selected parameter. By pressing the mode button on the gamepad, the user may switch the currently selected effect.

Furthermore, the user can enable/disable the manipulation of the focus lens through the graphical menu. The user can also toggle transparency and visibility of the entire reservoir by using either the menu or the start and back buttons on the gamepad, respectively. Setting the reservoir semi-transparent or invisible allows the user to view internal structures within the reservoir such as the well trajectories and their connectivity graphs (see Figure 86 (d)). Semi-transparent mode allows a contextual view of the reservoir surface to be maintained.




Figure 86: Different combinations of color and line effects for stylized focus.

5.3.2.1.2. IMPLEMENTATION

It is important to note that whenever the user is positioned outside the reservoir model, we use a ray casting technique in order to determine the focal point p_f as being the spatial coordinate where the ray pointer intersects the model. In this technique, we start by extracting the outer (polygonal) surface of the reservoir model to

improve the search for intersected cell by using fewer resources, such as execution time and memory space. Following this, we triangulate the polygonal grid cells and store them in an octree data structure, which provides a computationally efficient mechanism to search for possible intersection candidates among the large number of cells.

As illustrated in Figure 87, an octree is a tree data structure in which each inner node has up to eight children nodes, and are most often used to partition a three-dimensional space. For this purpose, an octree's root node describes a bounding box that encloses a volumetric model and its child nodes (also called octants) hierarchically subdivide the volume into eight disjoint sections. The octree is then constructed by recursively subdividing space into eight octants until the remaining number of objects in each octant is under a pre-defined threshold, or until a maximum tree depth is reached. In our specific case, the octree structure uses an octree subdivision where each octant carries an indication of whether it is empty or not, and each leaf octant carries a list of the cell triangles inside of it.

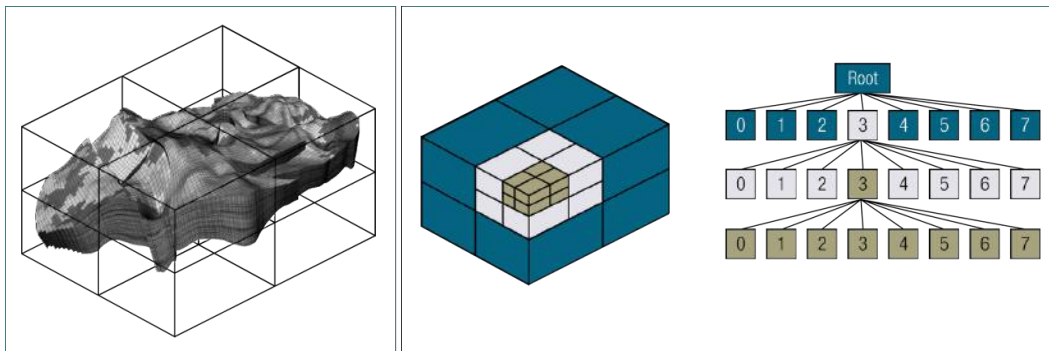


Figure 87: Illustrations of an octree, frequently used for partitioning volumetric objects such as a reservoir model.

To determine the cell being selected by the ray pointer, a ray-triangle intersection test [206] is performed. More specifically, while searching cell triangles in the octree, intersections are initially tested against the octants of the octree, and then linearly against the reduced set of triangles inside the final leaf octant. The test returns the intersection point (if any) and the cell which was intersected by the ray. Also, although the octree is constructed

during the pre-processing stage, it may be re-built considering that structural changes may happen to the reservoir model whenever property filtering is performed.

Once the focal point p_f is calculated, stylized emphasis is applied on the reservoir model using shading effects. As previously mentioned, we were inspired by the “stylized focus” technique proposed by [205]. According to this technique, for a particular point p , we calculate $f(p) = \|p - p_f\|$ where p_f is the focal point and $f(p)$ is called abstract focus and indicates how much emphasis to place at point p . Afterwards, a transfer function provides controls for adjusting the focus of $f(p)$ based on four parameters: e_{near} , e_{far} , d_{near} , and d_{far} . As can be observed in Figure 88, the parameters e_{near} and e_{far} indicate the degree of emphasis in the most and least emphasized regions, respectively. The parameters d_{near} and d_{far} control the fall-off regions.




Figure 88: Transfer function used for regulating focus.

This transfer function takes the form

$$T(p) = \text{clamp} ((f(p) - d_{near}) / d_{far} - d_{near})$$

The emphasis at a certain point p is then taken to be a linear combination of e_{near} and e_{far} , which represents the strongest and weakest emphasis, respectively.

$$E(p) = T(p) e_{near} + (1 - T(p)) e_{far}$$

For our specific case, three rendering effects are used for expressing emphasis on the reservoir cells: two color-related (transparency and saturation) and one line-related (width). These effects are generally used in

combination, though each effect has its own transfer function analogous to that of $E(p)$, with its own set of four parameters (see Figure 89).




Figure 89: Transfer functions used in our focus lens.

As previously mentioned, d_{near} and d_{far} parameters can be dynamically changed by the user; yet, e_{near} and e_{far} are pre-defined, empirically chosen values that we believe provide effective stylized emphasis rendering for reservoir models. More details on this technique can be found in the original paper [205].

5.3.2.2. CELL PROBING

The selection of a certain cell to view its information is a feature frequently used by reservoir engineers and geoscientists, also being common in commercial software. This feature is called cell probing and it allows a user to understand small-scale data variations along the grid cells. This is an essential tool especially when inspecting grid cells that may compose a well trajectory. Therefore, our application provides a similar feature, mapped to spatial interfaces.

5.3.2.2.1. DESIGN

To probe cells, the user presses and holds the A button on the controller. This action activates a small, white probing cube close to the controller that mirrors its position and orientation, as can be seen in Figure 90 (a). As illustrated in Figure 90 (b), the probing cube then acts as a three-dimensional cursor and whenever the cube collides with a cell, the grid cell is selected and rendered along with a small view-dependent window with

information about the cell such as the i, j, k coordinates of the cell, and properties strongly related to flow efficiency such as permeability in i, j, k directions; porosity; and transmissibility in i, j, k directions.

Additionally, a view-dependent, volumetric transparency lens appears around the cube in order to overcome the issue of surface occlusion when probing internal grid cells from outside of the reservoir surface. The lens allows the user to see into the reservoir and select internal cells. As can be observed in Figure 90, the lens removes grid cells surrounding the probing cube and only displays the reservoir's wireframe lines close to the cube and the cell being probed.

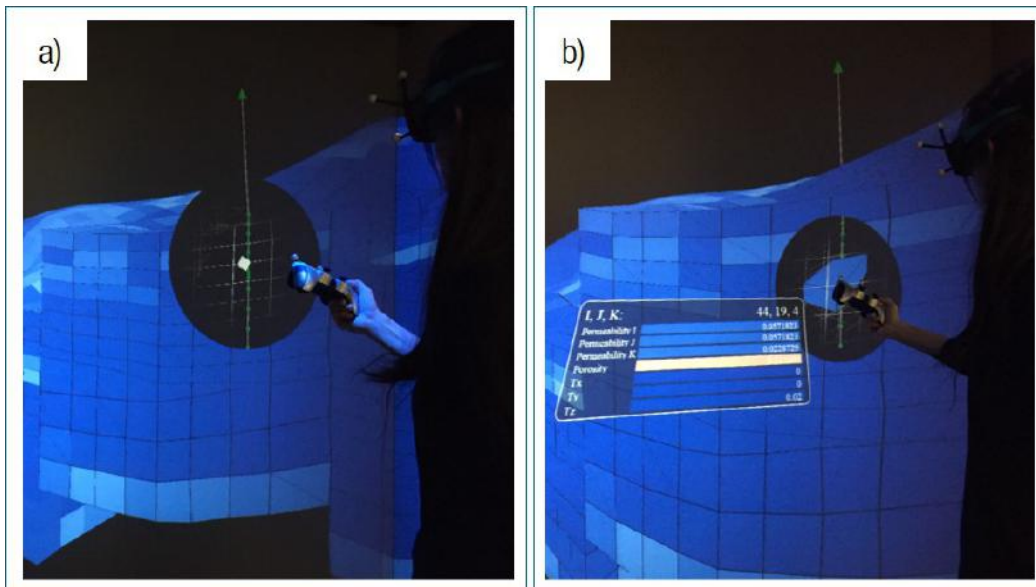


Figure 90: User probing grid cells.

5.3.2.2.2. IMPLEMENTATION

Whenever the white probing cube collides with the reservoir model, we use the same search technique introduced in Section 5.3.2.1.2 to find the intersected reservoir cell. However, in this case we build and conduct search operations in an octree that accounts for the entire reservoir model rather than solely its surface. This leads to a costlier searching procedure. Yet, considering that no structural changes affect this tree structure,

the octree construction occurs only once during pre-processing stage and, thus, allows for real-time search operations during cell probing.

To create the alpha lens' clipping effect, we use a shader programming technique. Figure 91 illustrates the scheme behind it. The lens has a right-circular cone shape, in which the yellow line segment represents the axis AC connecting the apex to the center of the circular base of the cone. C represents the center of the base, defined to be the spatial position of the probing cube; A refers to the apex of the cone, defined to be the coordinate of the user's head; and θ represents the half aperture of the cone.




Figure 91: The technique for defining the transparency lens consists in scanning fragments of the reservoir model to evaluate whether they lie within a cone-shaped volume.

The technique consists in culling all cell data lying within an unbounded version of this cone – that is, not limited by the plane of the base. For this purpose, in the shader, we convert the vertices of the reservoir model from their native 3D space to their final 2D position in screen space. Following, for each fragment F , we compute the angle δ between the vectors AF and AC . Then, we evaluate whether F is inside the cone by comparing the cosines of the angles θ and δ . If $\cos(\theta)$ is lower than $\cos(\delta)$, then F is located inside the cone.

Note that our approach discards not only fragments lying within the geometry, but also the ones located behind the round base of the cone. To remove only fragments positioned inside the cone, an additional step could be made to check whether the projection of the vector AF is shorter than the vector AC. Finally, we verify if the fragment F is a wireframe line by following the approach presented in [207]. If it is not an edge, the fragment is discarded; otherwise, the fragment is displayed with its transparency interpolated according to θ and δ angles:

$$\text{Opacity of F} = \text{lerp}(1, 0, \cos(\delta) / \theta)$$

This leads to a radial fall-off attenuation of the line intensity as more opaque lines are rendered near the center of the conical lens and the degree of transparency increases the closer the fragment is from the circular base boundary (see Figure 90).

5.3.2.3. WELL TRAJECTORY CREATION

A fundamental feature for performing well placement studies concerns the positioning of wells; however, as specified in Section 4.5, when performed in current commercial software, this procedure may be time-consuming and non-intuitive. As the design of well trajectories is often laid out in cross-sectional views, it is difficult for engineers to understand the spatial relationship between the three-dimensional reservoir model and the planned wellbore.

Due to its relevance and the limitations on standard software, our application provides the feature for defining well trajectories, mapped for spatial interfaces. In this sense, we take advantage of the inherent capacity of the tracked gamepad of interacting in all three dimensions to allow users to specify the well trajectory by directly selecting grid cells in the three-dimensional reservoir model.

5.3.2.3.1. DESIGN AND IMPLEMENTATION

Our application supports the creation of two types of wells: vertical and free-form wells. And the user may specify the type of well to be created on the graphical menu. For vertical wells, the user uses the ray attached to the gamepad to highlight cells on the surface of the reservoir. By pressing the A button on the gamepad, a vertical well is defined consisting of all cells along the vertical column below the currently highlighted grid cell. In addition, a well perforation is created for each of the cells along the well trajectory. Figure 92 illustrates the process of defining a vertical well trajectory.

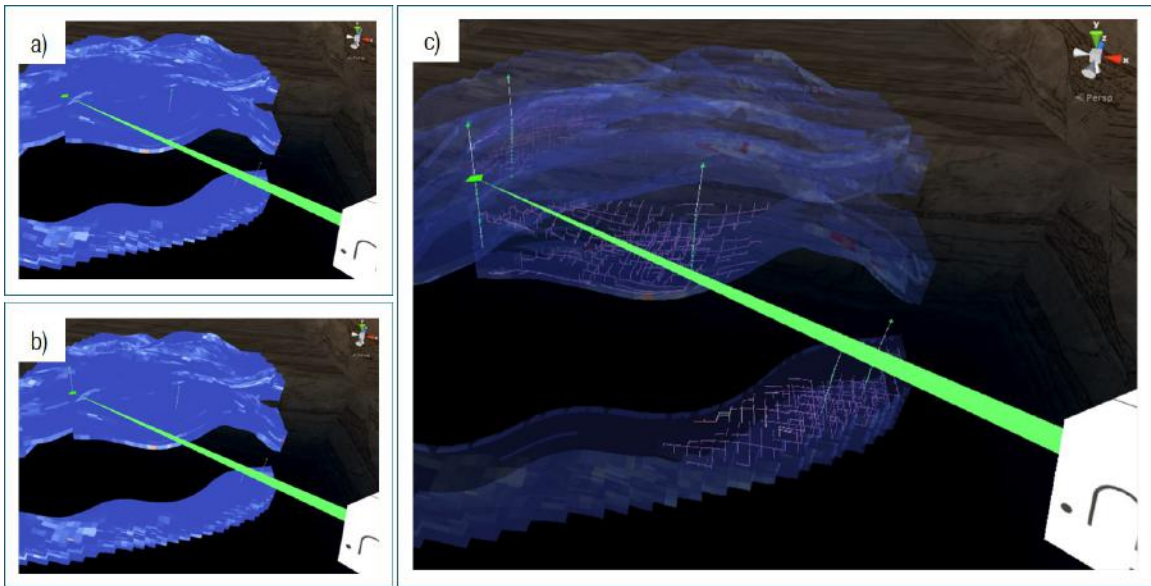


Figure 92: Procedure to define a vertical well path: user hovers a cell (a) and presses the controller's button to confirm creation (b and c).

When defining vertical wells, the ray casting technique described in Section 5.3.2.1.2 is once again used for finding the highlighted cell on the reservoir surface. Once the user presses the A button on the controller, we calculate an axis-aligned, vertical line segment that passes through the intersected surface point. Then, we compute the intersection points between this line segment and each of the remaining reservoir layers. In other words, we perform a layer-specific ray-triangle test to determine the intersected cell in each reservoir layer. These intersected cells then compose the well path.

To create free-form wells, the user must select each of the cells to compose the well trajectory. By holding down the B button on the gamepad, the probing technique is triggered and, whenever the button is released, the selected cell is defined as the first cell of the well path. By repeating this process, other cells are defined for the well. Re-selecting one of the cells signals the completion of the well path. Additionally, in a similar manner to the vertical well, a well perforation is created for each of the cells along the free-form well trajectory.

Figure 93 depicts the process of defining a free-form well trajectory.

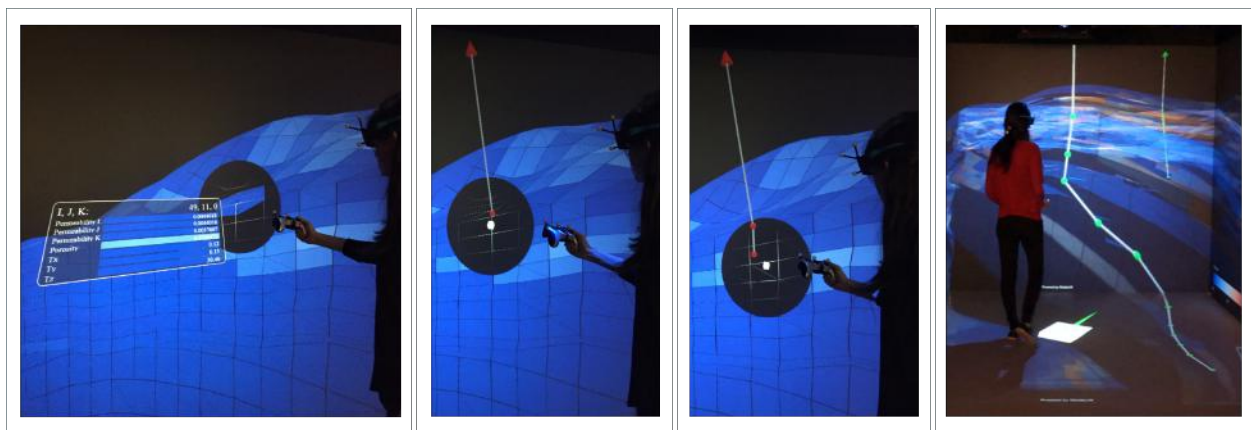


Figure 93: Process for defining a free-form well trajectory.

5.3.2.4. WELL TRAJECTORY ANALYSIS – GRAPH

To evaluate well trajectories, our application uses the method for performing static connectivity analysis described proposed in this thesis. The current version of our application offers connectivity assessment only in primary recovery scenarios. In this sense, as was previously mentioned, the algorithm computes a connectivity graph for each perforation along each well path. Each node included in a connectivity graph has an associated connectivity measure called “delta time”. The measure relates to the easiness of flow movement between grid cells, thus indicating the relative time that would be required for oil to flow from a cell to its associated well perforation.

5.3.2.4.1. DESIGN

When it comes to designing the connectivity graph, we decided to employ a diverging color scheme to the edges of the graph in order to encode “delta time” data. And the reason for this decision is the fact that this color map addresses a number of requirements important in scientific visualization, such as maximal perceptual resolution, minimal interference with the shading of three-dimensional surfaces, and matching between the perceptual interpolation and the underlying scalars of the color map [208]. As can be observed in Figure 94, edges with a deeper purple color have a lower delta time value and, thus, they have a stronger, more transmissible connection between the perforation and the cell. A brighter orange indicates a higher delta time value, and thus a relatively low connection to the cell.

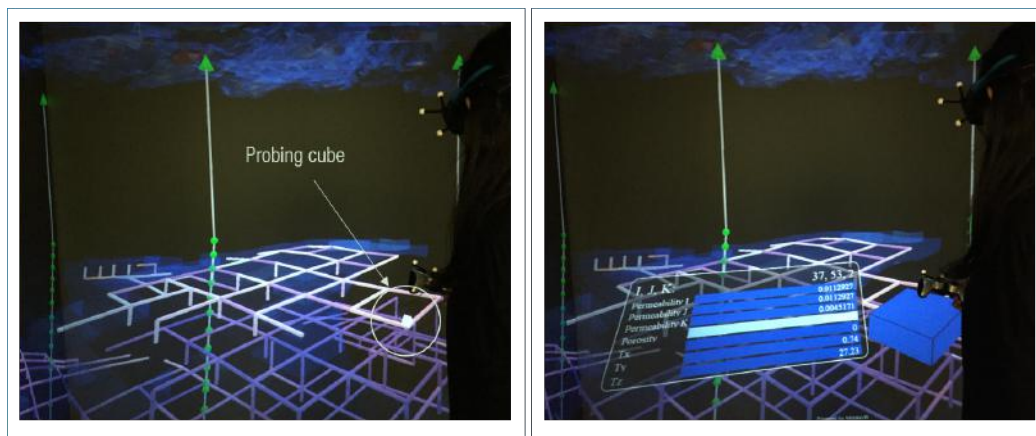


Figure 94: User probing grid cells for evaluating the connectivity graph.

The graph is represented as a tube that radiates spatially from a perforation cell and connects the centroids of each pair of adjacent cells. In this way, the graph is rendered within the reservoir model as this facilitates the evaluation of the stratigraphic and geological factors influencing the connectivity between cells. For instance, in a scenario where the connectivity changes abruptly, the user may use the focus lens or the transparency lens, or even set a semi-transparent reservoir in order to look for faults, fractures, or gaps on the cell geometries that may be acting as a barrier to the flow. Further to this, the user may also probe cells along the edges of the

graph where the connectivity suddenly changed in order to examine whether geological property values are negatively impacting the connectivity, as can be seen in Figure 94.

Also, as was previously mentioned, the static connectivity algorithm allows users to define a delta-time threshold value in order to filter out cells not reachable at that time limit. On this matter, the application uses a 1-DOF slider widget that ranges from zero to the maximum existing delta time. As illustrated in Figure 95, in order to define the threshold, the user places the ray pointer above the slider, holds the A button on the gamepad, and drags the ray along the slider. In turn, the spatial graph changes accordingly to the user-specified threshold.

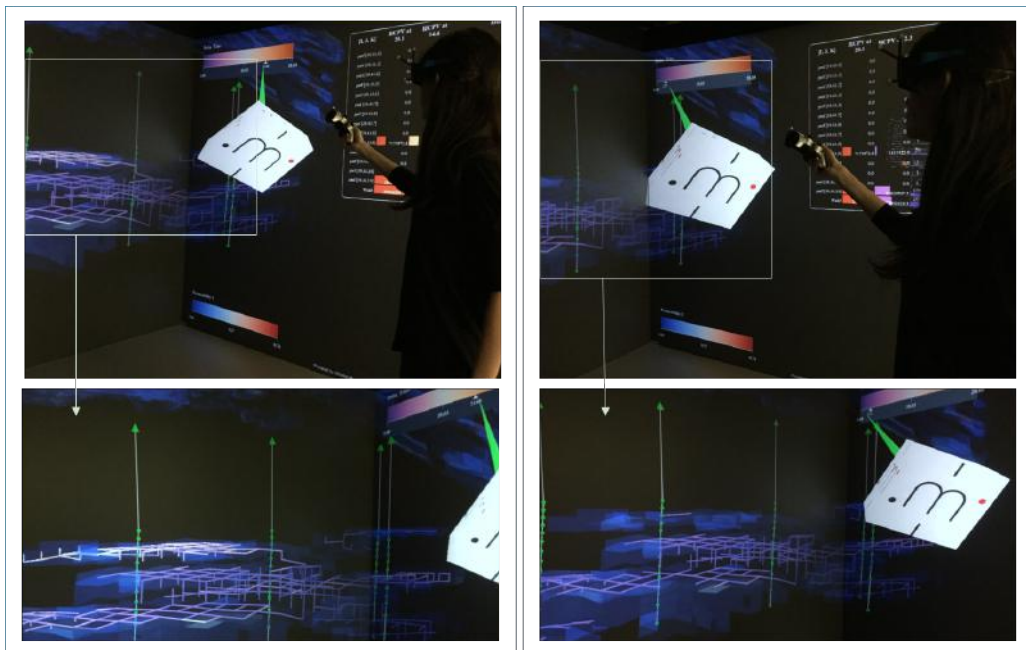


Figure 95: By changing the time threshold through the slider widget, the connectivity graph is filtered according to it.

Finally, as a reservoir model may contain multiple wells, and each well may have multiple graphs, the view may easily suffer from an excess of clutter. Thus, to help eliminate the clutter, the user may press and release the A button on the gamepad on a well head. This selection then results on only cells connected to the well to be

displayed. Similarly, by selecting a specific well perforation, only cells connected to that perforation will be displayed, as can be seen in Figure 96.




Figure 96: User selects a specific perforation in order to assess only its associated connectivity graph.

5.3.2.4.2. IMPLEMENTATION

The slider widget is a webview provided by MiddleVR: a plane in which an HTML page is applied as texture. In order to enable interaction control for defining a time threshold, whenever the user holds the A button on the gamepad while hovering the slider, the 3D position of the button press is mapped to a 2D virtual mouse click on the web page, as illustrated in Figure 97.




Figure 97: Mapping between the 3D ray-plane intersection point and the 2D virtual mouse click.

For this purpose, ray-triangle tests are first performed in order to determine the triangle intersected by the ray pointer. This provides information on the vertices and barycentric coordinates of the triangle hit. Following this, the texture (uv) coordinates of the intersection point is calculated by multiplying the vertex weights (barycentric coordinates) with the uv coordinates of the triangle hit. Finally, this point is scaled with the texture size in order to obtain the actual pixel within the webview.

Once the user determines a time limit using the slider widget, the connectivity graph is filtered according to this threshold. Hence, we elaborated a simple, light-weight technique for filtering the graph. First, while constructing the graph mesh, we logically separate it into disjoint sub-meshes, each of which connecting the centroids of two adjacent cells (see Figure 98). For each sub-mesh, in turn, we encode as its vertices' texture coordinates the time taken to reach the furthest cell between the two connected cells. By passing these extra data to the shader, we can perform a low-cost comparison between the currently user-specified time limit and the time taken to reach that particular vertex. If the reaching time is lower than the threshold, the vertex is discarded; otherwise, it is rendered. In this way, as can be observed, each sub-mesh is then either entirely rendered or culled based on the time threshold.




Figure 98: Each tube sub-mesh encodes its associated delta time as vertex data.

5.3.2.5. WELL TRAJECTORY ANALYSIS – PANEL

For each new well, along with the connectivity graphs, a tabular panel is created to display analytical connectivity data associated with the well.

5.3.2.5.1. DESIGN AND IMPLEMENTATION

Similarly to the slider widget, each analytical panel is a web view mapped as a plane texture. Each panel displays a distinct table. The table contains four columns, each row shows data associated with a well perforation, and the last row shows cumulative data for the entire well. The first column of the table shows the i, j, k coordinates of the perforations. The second column displays the total drainable hydrocarbon pore volume (HCPV) of all connected cells – that is, the cells associated with the nodes that compose the connectivity graph of a perforation. The third column shows the total HCPV of all connected cells that have delta time below or equal to the currently selected time threshold. The last column shows the distribution of time for all connected cells.

| [I, J, K] | HCPV at 59.6 | HCPV at 1.6 | Distribution Delta_Time |
|-----------------|--------------|-------------|-------------------------|
| perf [38,38,0] | 0.0 | 0.0 | 1 |
| perf [38,38,1] | 0.0 | 0.0 | 1 |
| perf [38,38,2] | 3753662.5 | 661980.2 | 208 14 4 |
| perf [38,38,3] | 0.0 | 0.0 | 1 |
| perf [38,38,4] | 0.0 | 0.0 | 1 |
| perf [38,38,5] | 0.0 | 0.0 | 1 |
| perf [38,38,6] | 0.0 | 0.0 | 1 |
| perf [38,38,7] | 0.0 | 0.0 | 1 |
| perf [38,38,8] | 0.0 | 0.0 | 1 |
| perf [38,38,9] | 3989388.6 | 3675362.7 | 185 4 |
| perf [38,38,10] | 3989388.6 | 3675362.7 | 185 4 |
| perf [38,38,11] | 3989388.6 | 3664843.1 | 185 4 |
| perf [38,38,17] | 0.0 | 0.0 | 1 |
| perf [38,38,18] | 0.0 | 0.0 | 1 |
| perf [38,38,19] | 5698212.1 | 4519931.2 | 141 8 1 |
| Well1 | 13441263.2 | 8857274.1 | 914 34 5 |

Figure 99: Example of an analytical connectivity panel.

To ease the individual evaluation of each panel, for each table column, the bar lengths are normalized across its rows. This allows a fast, graphical comparison of the performance of perforations that compose the well. As an example, for the first column, the bar length of a perforation p_i represents the fraction of pore volume

produced by p_r . Therefore, if p_r is capable of draining only 10% of the pore volume for the entire well, then its corresponding bar has 10% the width of the bar referred to the well. As another example, for the second column, if p_r produces 5% of the pore volume produced by the entire well at a certain time threshold, then its corresponding bar has only 5% the width of the bar relative to the entire well – and this width changes whenever the time threshold changes.

Further to this, the bars are colored according to their associated delta time. For instance, for the first column, the bars represent the drainable HCPV at maximum delta time and, thus, are always bright orange colored. For the second column, the bars correspond to the drainable HCPV at a certain time and, thus, are colored according to the currently selected time threshold – and, whenever the threshold changes, the color on the bars changes accordingly.



Figure 100: Same panel displaying data for different time thresholds. It can be observed variations in bar color and length in the second column. From left to right, $\delta = 59.6$, $\delta = 15.3$, $\delta = 0.6$.

To facilitate the comparison of different panels, the length of the bar charts is also normalized across all panels. For instance, consider that a well w_1 produces a total of V_1 of HCPV, which is the maximum drainable volume between all existing wells; consider also a well w_2 that produces only fifty percent of V_1 of HCPV. The graphical data displayed on the w_2 panel would then be displayed normalized according to such proportion. It is hence possible to compare graphical information from panels that are further away from the user or scaled down in a manner that the numerical data are not too readable.

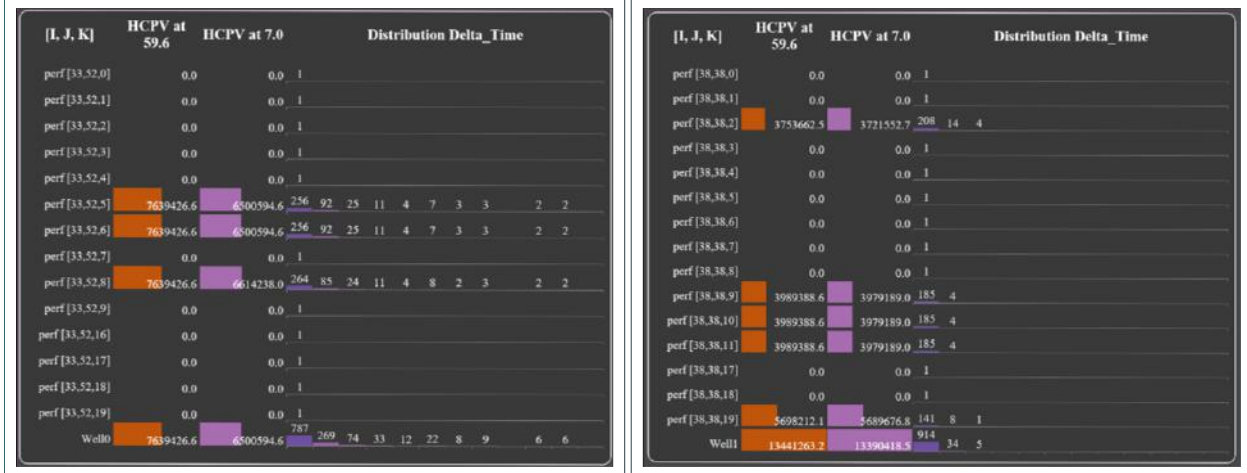


Figure 101: The normalized bars ease the visual comparison between two or more panels. In the picture above, one can quick check that the total production volume on the well on the left is lower than the one on the right.

Finally, the user may arrange each panel individually as can be seen in Figure 102. Panels may be expanded or compressed so that they may be brought to the forefront of attention or minimized to reduce information clutter. To scale a panel up or down, the user hovers over it with the ray and presses the up or down buttons on the D-Pad. Panels may also be positioned anywhere in space, stacked on top of each other, sorted, and organized according to the user's preference. To translate and to rotate a panel, the user places the ray on it and holds the A button on the gamepad in order to select the panel. Then, the panel follows the movement of the tracked controller.

In order to decrease the physical effort of arranging panels, though, the mapping between the movement of the controller and the selected panel is not 1:1. When moved using the gamepad, the panel translates by a factor of 5:1 and rotates by a factor of 2:1 relative to the physical translations and rotations using the controller. These factors were empirically determined, but we believe them to be effective in reducing physical workload while maintaining the required precision to efficiently perform the tasks.




Figure 102: Examples of different panel arrangements.

5.4. INTERACTIONS USING LEAP MOTION

When using the leap motion controller as an interaction interface, the user has access to only a subset of features: setting a system control menu; manipulating the reservoir model; navigating the user through the scene; sectioning the reservoir; probing the grid cells; and creating vertical well trajectories. This reduced feature set was implemented to collect initial feedback regarding gesture controls that could help designing the more complex features, such as the free-form well trajectory creation.

5.4.1. GENERIC

5.4.1.1. SYSTEM CONTROL

When controlling the application with the leap motion device using the Oculus, the application uses a 1-DOF hierarchical graphical menu attached to the user's left hand. This arc-shaped menu appears whenever the user rotates the left palm towards the face. The user can then select the menu items using the index finger of the opposite hand. To select an item, the user moves its fingertip to the menu item and hovers briefly. As illustrated in Figure 103, the menu uses a two-level menu hierarchy and the menu's root is displayed on the palm of the user's hand. And whenever a second-level item that contains sub-items is selected, it becomes a temporary root. Additionally, to navigate back in the menu hierarchy, the user may either close its left hand or select the back navigation button near the menu's root.




Figure 103: On the left, the user is selecting the filtering operation on the graphical menu. On the right, the user is filtering the reservoir based on permeability values.

5.4.1.2. RESERVOIR MANIPULATION

When using the leap motion, we investigated different sets of hand gestures for performing reservoir manipulation. These sets were sequentially prototyped, in a way that each set was built upon the lessons learned from informal usability evaluations of the previously designed set.

5.4.1.2.1. FIRST DESIGN

The first group of gestures was designed considering one major factor: ease of detection. In this sense, some hand motions and poses are tracked more consistently and, therefore, are easier to detect than others. This is important since motions that are easier to detect have lower rate of false positive and false negative detections, making the application experience more useable. In addition, the sooner the gesture is detected, the sooner the application can provide proper feedback. This leads to an application that feels more responsive and makes people feel more in control. Therefore, we focused on designing gestures that require users to keep their hands perpendicular to the leap motion Controller's FOV since this is known to be a very reliable VR tracking pose [209].

As can be observed in Figure 104a, for the first prototype, the user extends all fingers of the left hand in order to trigger a circular slider that allows the user to set the operation to be performed: rotation, translation,

or scale. By making a sphere-like rotation with the left hand around the wrist, the user controls the graphical widget in order to modify the currently selected geometric transformation.

Whenever the user extends all fingers of the right hand, this triggers a second widget. This circular slider shows in its center the currently selected operation and allows the user to set the direction and speed of the operation to be performed, as illustrated in Figure 104b. In terms of scaling, the reservoir is scaled up equally in all directions if the user rotates its right hand clockwise (CW), and is scale down likewise with a counterclockwise (CCW) rotation. For translation, CW and CCW rotations causes the reservoir to move to the right and to the left in the x-axis, respectively. For rotation, in a similar manner, CW and CCW rotations cause the reservoir to rotate positive and negative angles around the y-axis, respectively.

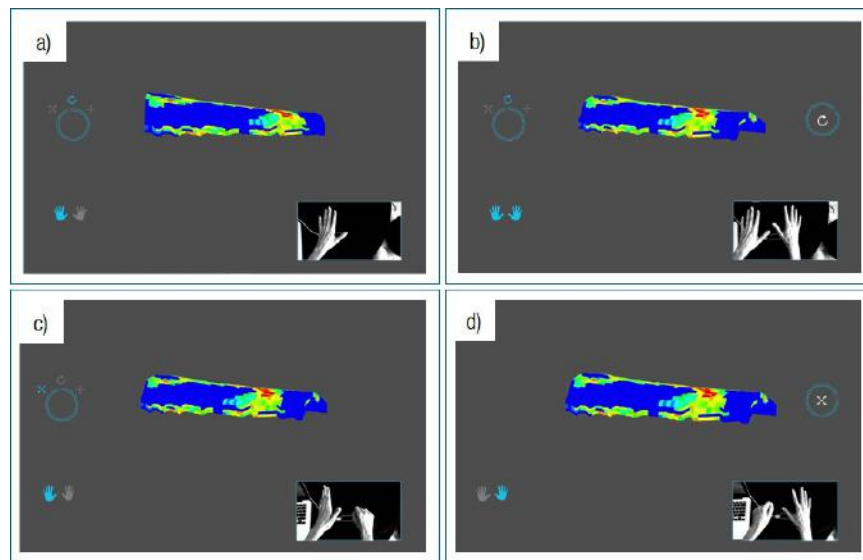


Figure 104: On this gesture interface, hands must be fully opened to enable graphical controls. The left hand is used to select an operation – translation, rotation, or scale, while the right hand changes the operation value.

In addition, two small hand images are displayed on the user interface to provide visual feedback regarding the hand tracking: whenever a “fully open hand” pose is tracked in either left or right hand, the associate hand image turns bright blue to provide positive detection indication. Besides making the application more responsive, we believe these interactive graphical elements also guide user learning to interact with our gesture-based

interface: by observing these visual elements, the user may empirically learn to keep the correct pose – that is, fingers splayed and perpendicular to the FOV of the device, and to keep their hands within the tracking zone.

By running informal usability tests, we found some limitations in our initial prototype. One major limitation regarded the fact that translation and rotation operations could only be performed in one, pre-defined axis. In our interaction design, even though we considered axes that are mainly used for manipulating the reservoir, we discovered them to be insufficient; therefore, operating a single axis becomes a strong limitation for engineers. As a second limitation, users had difficulty to interact using both hands at the same time, even though the use of both hands was not mandatory. On this matter, in situations where users had both hands detected, we observed some users modifying the selected operation by mistake by inadvertently rotating the left hand while focusing on performing the operation with the right hand, and vice-versa.

5.4.1.2.2. SECOND DESIGN

The second prototype was designed considering the lessons previously learned. To overcome the “operation along one axis only” limitation, we designed an arc-shaped sub-menu for the left menu where the user can select the axis over which to perform the operation. As seen in Figure 105, by hovering over one specific operation in the left menu, the sub-menu appears after some time period, and then the user may hover over a sub-item in order to select either x, y, or z axis. And, after a set time period of non-interaction by the user, the sub-menu automatically disappears. Furthermore, in order to prevent the user of setting operations unintentionally, we re-designed the left menu in a way in which it turned to being disabled whenever the right menu is being displayed.

However, when running tests to evaluate our reviewed design, we noticed that it was too cumbersome for users to access the left menu or sub-menu every time they needed to change the selection of either the operation or the axis. A second issue arose once we coupled the leap motion with the Oculus Rift. Due to the

fact that, in this new configuration, the leap motion Controller was now attached to the front of the Oculus, the tracking consistency was significantly affected for the same set of gestures previously used. Additionally, another limitation we found concerned the fact that the graphical widgets were occupying too much space on the already-limited FOV of the HMD.

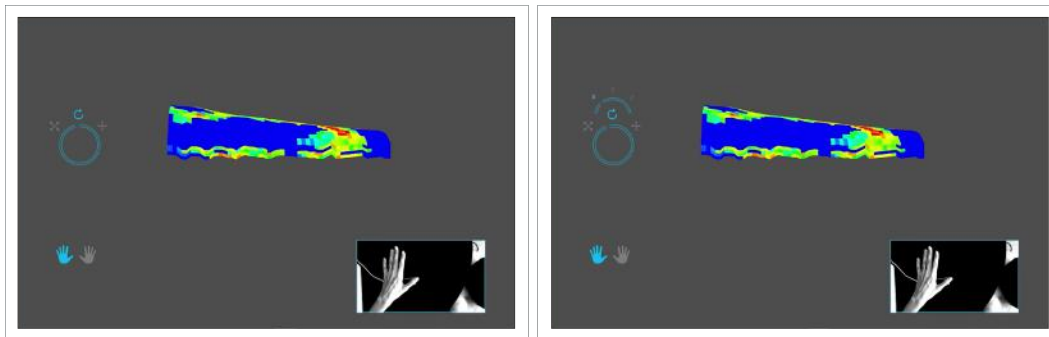


Figure 105: This interface extends the previous one by adding a sub-menu through which the user may choose the axis along which the currently selected operation is performed.

5.4.1.2.3. THIRD DESIGN

The third prototype was designed based on the lessons previously learned from the two designs we had investigated. For the new set of hand motions, we considered two additional factors that influence the usability of the application. A factor taken into account was ergonomics. In this sense, as the user navigation and object manipulation tasks are intended for extensive use, we considered whether the hand movements had the possibility to causing long-term harm or strain – also known as “gorilla-arm syndrome” in the scholarly literature [210]. Thus, we designed interactions to be non-repetitive in order to reduce repetitive strain injuries. Additionally, it was also important to see how tiring the interactions can be, and if they can be performed from comfortable positions. Therefore, we designed gestures considering that a comfortable position for people to use gesture-based applications is with their elbows resting comfortably on the table or arms of the chair [209].

Another factor considered was gesture transitions; there are two primary ways in which transitions can be an issue for the application experience: interaction overlap and ambiguous switching. For the former issue, this

is a situation in which two possible actions are too similar to each other. This can cause issues, both for people using the application and gesture detection algorithms. Thus, we spent the time designing gesture metaphors that are not overly similar.

In the case of ambiguous switching, where and how a gesture begins can have a lot of influence on how people tend to perform the gesture and on how in control people feel when using the gesture-based interface. On one hand, ambiguous switches can cause ergonomics issues where people tend to move in uncomfortable or overly exaggerated manners. On the other hand, if not properly designed, interactions may be initiated with gestures that are a part of casual movements – also known as the “Midas Touch” problem [211]. Thus, an important interaction design principle is that gesture-based interactions should be initiated with safe poses and/or specific gestures so that users can safely move through the device’s field of view without initiating unintentional interactions.

Therefore, for translating, rotating, or scaling the reservoir, we considered that an “initialization” (or “trigger”) speed magnitude from which an operation begins may be a good way to reduce the challenge of dealing with ambiguous switching. More specifically, after detecting a gesture pattern, the detection algorithm works as follows: first, the velocity of each of the n extended finger tips is measured as v_1, v_2, \dots, v_n . Then, the mean velocity is calculated as $v_m = (v_1 + v_2 + \dots + v_n) / n$. The algorithm then triggers the gesture-related action provided that v_m surpasses a certain speed threshold. And in case of an interruption of the gesture detection, the action stops after a set time period.

For scaling the reservoir model, as illustrated in Figure 106, the user must put both hands up with all fingers extended and with palms facing outwards. In order to trigger the scaling operation, the user must then move either forward or backward fast enough to surpass a pre-defined speed threshold. If the movement is fast enough and is forward (that is, positive in the z-axis), the reservoir is scaled up equally in x-, y-, and z- axis; if

the movement is fast enough and is backward (that is, negative in the z-axis), the reservoir is scaled down equally in all axes.




Figure 106: Sketches of the scaling operation.

For rotating the reservoir model, the user must put the right hand's fingers all extended and palm facing outwards. Next, the user must freely move the hand fast enough to exceed the pre-defined speed threshold in order to trigger the rotation operation.

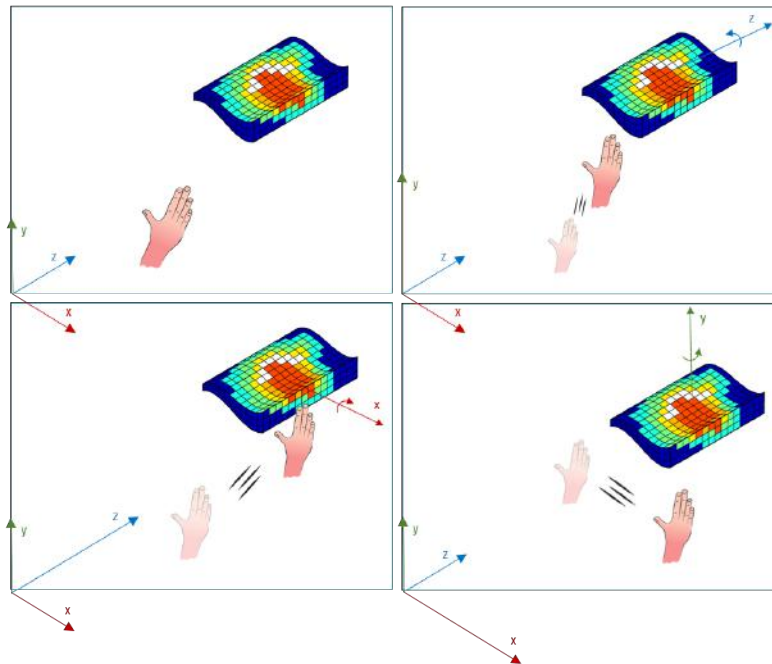


Figure 107: Sketches of the rotation operation. The sizes of the axes disposed at the bottom left corners represent the magnitude of the hand motions in the associated directions.

In case the operation is triggered, the hand motion is then filtered as follows: given the motion vector $v = (v_x, v_y, v_z)$, the filtered vector consists of the same vector with its two lowest values substituted to zero. The resulting filtered vector is then either $v_f = (v_x, 0, 0)$, $v_f = (0, v_y, 0)$, or $v_f = (0, 0, v_z)$. That is, the filtered motion has only one non-zero axis value, which means that the rotation is performed per axis.

Following that, the filtered hand motion's axis is mapped to another axis according to the mapping scheme depicted in Figure 107. This mapped axis is centered in the reservoir model's origin and is the one around which the reservoir is rotated. This mapping is applied to provide a more natural, consistent co-relation between the hand gesture and the reservoir behavior. For instance, as the moving the hand forward (+z) resembles a "push away" motion, the resulting mapped rotation tries to give the feeling that the reservoir is indeed being "pushed away". Finally, the angle degrees of rotation are defined as the magnitude of the mean velocity v_m .

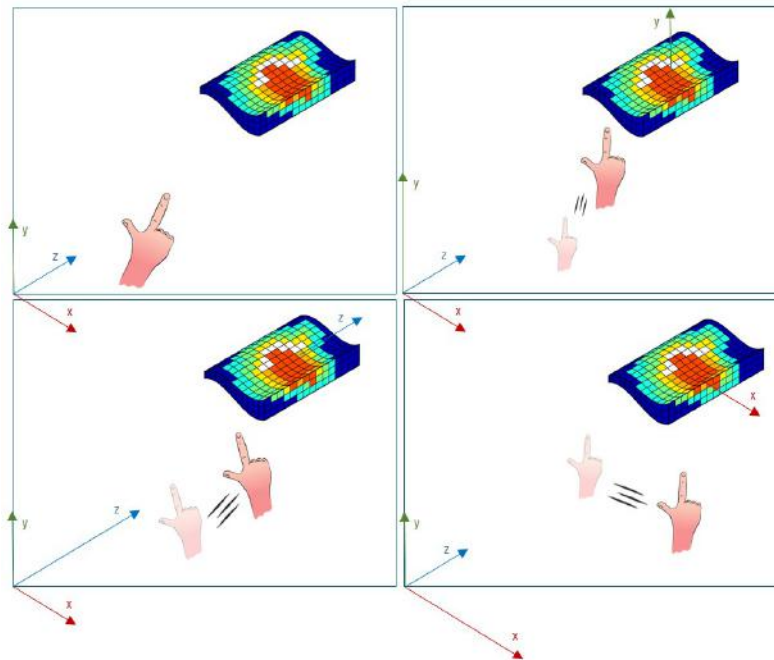


Figure 108: Sketches of the translation operation.

For the translation operation, the user must put the right hand's thumb and index fingers extended while the palm is facing outwards. Next, the user must freely move the hand fast enough to surpass the speed threshold

and, thus, trigger the operation. In case the operation is triggered, the same abovementioned procedure occurs: the hand motion is filtered in order to execute the translation in a particular axis; however, as illustrated in Figure 108, here an identity transformation is used to map between the filtered axis and the one along which the model is translated.

Figure 109 shows the rotation, translation, and scale operations applied to a reservoir model using our gesture-based interface in the application. Similarly to the two previous prototypes, graphical elements are used to allow the user to be aware of the currently performed operation. Thus, whenever an operation is detected and triggered, the associated image turns into bright orange. The user can also deactivate these images through the graphical menu.

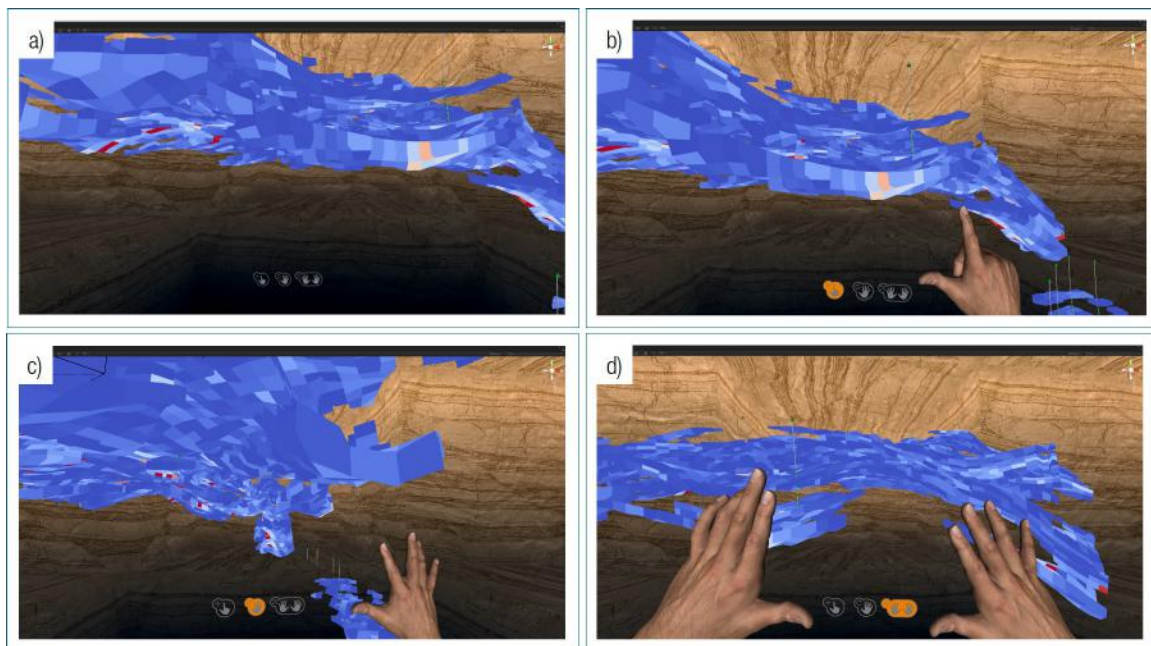


Figure 109: Illustrations of our interactions for manipulating a reservoir model while using gestures: default interface when no operation is being performed (a); translation (b); rotation (c); and scaling (d).

5.4.1.3. USER NAVIGATION

As mentioned in Section 5.3.1.3, when using the gamepad, the user can navigate through the scene using the same controls to translate and rotate the reservoir. In a similar manner, when operating the leap motion, the

user controls the viewpoint translation by using a procedure analogous to the one for translating the reservoir model: he/she first extends thumb and index fingers on the right hand and, then, moves the right hand above some speed threshold in order to initiate the movement. And the same occurs for controlling rotations of the viewpoint.

However, a major difference from the interaction style for reservoir manipulation is that, when navigating, the user does not have to maintain the hand gesture to continue the movement; rather, once the movement starts, it only stops when the user does the thumb-and-index-extended gesture once again. This design was chosen due to the matter that the user may want to look around the scene while moving through it. In this sense, the non-requirement for continuous gesture recognition is helpful because a) it avoids the challenge of making tracking fingers consistent while moving the head, and b) it prevents the virtual hands to block the user's view while navigating the earth. A second difference from the interaction interface for reservoir manipulation is that, similar to when using the gamepad, translation and rotation operations for the viewpoint are performed relative to the user's head orientation.

5.4.2. DOMAIN-SPECIFIC

5.4.2.1. RESERVOIR SECTIONING

As discussed in Section 5.3.2.1, in order to section the reservoir in regions of focus and defocus, the user must first specify a three-dimensional coordinate in the model to be considered the focal point. And, in a similar manner as occurs when using a gamepad, the focal point is set to be the user's spatial position if they are physically placed inside the model.

Additionally, if the user is spatially positioned outside the model, a ray casting technique is used. When using the gamepad, the ray pointer starts at the gamepad position and casts outwards – that is, in the direction of

the gamepad's forward direction. If using the leap motion controller, in turn, the ray starts at the user's head position and directs outwards, as can be seen in Figure 110.

It is worth mentioning that this design decision to emit the ray from the user's head was due to two main reasons. First, the head provides a more stable platform than the hand. On this matter, it is known that the task of selecting objects that have small visible areas or that are in highly cluttered environments is challenging when using direct pointing techniques because of they require high levels of precision. In addition, issues such as hand and tracker latency and jitter may also negatively affect the task performance [212]. Therefore, we designed our pointing technique to rely on the user's head rather than hands as it offers a number of benefits: it does not depend on the performance of the leap motion sensors; it does not suffer from the hand jitter factor that commonly impairs accuracy of fine hand motions; and it provides sufficient precision levels for selecting surface cells on the reservoir as focal points. As a second reason, by attaching the ray on the user's head, his/her hands remain free for performing additional actions such as manipulating the reservoir or changing the focus lens configurations.




Figure 110: Images of the focus lens when the user is situated inside (left) and outside (right) the reservoir model.

It should also be noted that for setting the $[d_{near}, d_{far}]$ parameters for the rendering effects of the focus lens, the user may use slider widgets on the graphical menu. Further to this, the user can enable/disable the focus lens, and set the transparency or invisibility of the reservoir through the menu.

5.4.2.2. CELL PROBING

As illustrated in Figure 111, the user must extend all fingers in the right hand with the palm facing outwards in order to trigger the probing function. The user can then freely move the fully-opened hand to probe different cells. Also, when using gestures, the right index finger acts as the three-dimensional cursor instead of a small white cube. Thus, the right index is surrounded by the transparency lens and whenever it collides with a cell, this grid cell is selected, causing it to be rendered along with the data window.




Figure 111: Illustrations of the alpha lens surrounding the user's right index (left) and the information window associated with the probed cell.

5.4.2.3. WELL TRAJECTORY CREATION

When using gestures, our current application allows the user to create only vertical wells. Similarly to what happens when operating the gamepad, the user uses the ray pointer to highlight cells on the surface of the reservoir; yet, here the ray is attached to the user's head rather than its hand, as illustrated in Figure 112 (a).

For our first prototype, we designed a time-based cell selection in a way that a certain grid cell was selected after being hovered for a pre-defined time period. Then, the same procedure occurs as when using the gamepad: a vertical well is created, made of all cells along the vertical column of the selected surface cell. However, while performing informal usability tests with this initial prototype, we observed that this time-based approach led to a number of accidental selections. More specifically, there were times when the user was simply observing the reservoir, looking for potential candidate cells, and unintentionally hovered and selected cells.

Therefore, for our second prototype, we tried to address this issue by restraining the validation of the selection to be determined by the user rather than by time. In this sense, whenever a given cell is hovered during a set time period, a graphical element appears in the center of the user's visual field, as can be seen in Figure 112 (b). This element acts as a confirmation button such that the user's left index must hover it in order to confirm the creation of a well trajectory along the currently highlighted cell, as shown in Figure 112 (c and d).




Figure 112: The procedure for defining a vertical well trajectory: the user hovers a surface cell using the ray (a); after a time period, a graphical button is displayed (b); and the user hovers it with its left index in order to confirm the creation of the well (c and d).

6. IVEs FOR WELL PLACEMENT OPTIMIZATION: AN EVALUATION




Figure 113: Illustrations of participants being introduced to our application on the CAVE environment.

This chapter investigates the relevance of our system, introduced in Chapter 5, to advance well placement studies from reservoir engineers' standpoint. It thus presents a qualitative primary study conducted with the purpose of evaluating both our application and the different immersive interfaces that can be used by our application. This chapter will proceed as follows: we will start by describing the structure of our study in Section 6.1. Following this, we will explain how we conducted the two-part evaluation sessions: one part for assessing our system's features, in Section 6.2; and another for investigating engineers' perceptions and preferences with regards to the different interfaces, in Section 6.3.

In addition, we will provide information about the participants' academic and professional background in Section 6.4, and concerning the dataset used on our study in Section 6.5. Section 6.6 will report our findings

grouped per theme: on the features and on the immersive interfaces. Section 6.7 will then provide our reflections on the findings in order to discuss the applicability of both the features and the different immersive technologies for well placement studies and for the broader domain of oil and gas; and, finally, Section 6.8 will discuss potential ideas for future work drawn upon the feedback received from our study.

6.1. EVALUATION SESSION

The evaluation sessions were divided into two parts: an evaluation of the features of the application and an investigation into the suitability and preferences of engineers with regards to different immersive environments when performing the inherently exploratory tasks of the application. Throughout each session, we performed a series of five standardized demos and semi-structured interviews with each participant. The detailed demo procedures, questionnaire, and discussion themes can be found in the Appendix A. The sessions were video and audio recorded for further qualitative analysis.

We conducted one pilot session with an expert reservoir engineer from our research group – who was not involved in the design of the interfaces, and five sessions with external reservoir engineering experts. Each of these six evaluation sessions lasted around 60 to 90 minutes.

Sessions started with a description of the goals of the evaluation session: the subjective judgement of the features of the prototype, as well as the reflection upon the different virtual reality mediums to be presented. Next, the participants were requested to answer a questionnaire to tell their fields of expertise, as well as their previous experience within the domain.

6.2. ON THE APPLICATION

The features of the application were first demonstrated using the CAVE environment. Since most participants were new to virtual reality technology, we were aware that the lack of familiarity could create an unfavorable or intimidating experience. To reduce biased results and allow feedback without having to account for lack of

training, each of the features were demonstrated by one of the two session facilitators. The facilitators used the tracked stereo glasses and thus had the optimal perspective and controlled the application using the gamepad. Participants used untracked stereo glasses and were requested to remain close to the facilitator to ensure their view perspective was not distorted.

A subset of the implemented features was then introduced and explained in the same sequence presented in this thesis: (1) reservoir manipulations, (2) cell probing, (3) vertical and free-form well trajectory creation, (4) well trajectory analysis using the connectivity graph, and (5) well trajectory analysis using the connectivity analytics panels. Figure 113 illustrates some of the application features being introduced to the participants.

At the end of each of the five demo sections, participants were asked to provide their opinion on the feature they had been exposed to in that section, and to reflect on its usefulness, potential problems, and suggestions for improvement. Once the five features were demonstrated, we also asked them for suggestions of other potentially useful features, and then finalized by welcoming additional comments if there were any.

6.3. ON THE IMMERSIVE INTERFACES

After the demonstrations and the interviews about the features of the application in the CAVE, a preliminary inquiry on other immersive displays and other interaction devices were performed, as can be observed in Figure 114. Participants were thus exposed to the same application running in an Oculus Rift DK2 along with the same tracked gamepad controller previously used in the CAVE. Then, participants were asked to reflect upon the differences they noted between their traditional desktop display, the CAVE, and the Oculus; what advantages or disadvantages they noticed about each environment; which environments they preferred and why, as well as general comments.



Figure 114: Participants testing our application using different immersive environments.

Following this, participants were introduced to the same application using an Oculus Rift DK2 along with a leap motion controller. They were prompted similar questions to gather feedback comparing their traditional keyboard and mouse, the tracked gamepad, and the leap motion controller.

6.4. PARTICIPANTS

The preliminary study relied on six individuals: four males and two females, ranging from 26 to 37 years old. All of them had B.Sc. degrees and one or more graduate degrees in Petroleum Engineering. Five of them were enrolled in a Ph.D. Program in Petroleum Engineering and one was a post-doctoral scholar. Five of them had between one and seven years of professional experience as reservoir engineers. One of the participants was currently employed as a reservoir engineer and two other participants were currently working in a reservoir simulation software company.

6.5. DATASET

The reservoir model used on the evaluation sessions is a hydrocarbon reservoir model constructed using a three-dimensional training image through multi-point facies statistics. The structure of the model consists of irregular hexagonal corner point cells with several normal and reverse faults. The model has 60*60*20 cells in i, j, and k directions, respectively. The model used in this evaluation contains two dynamic properties and ten static properties. MPS algorithm was used to populate the facies data including shale, silt, and sand. SNESIM algorithm has been used to populate other property data.

It is worth mentioning that the reservoir model represents a typical low permeable meandering channelized environment where the high permeable sand bodies are effectively isolated within the pervading very low-permeability reservoir facies. The porosity varies between 1% to around 0.35%. The permeability field is anisotropic and the horizontal permeability values are ranging from 0.001 mD for the background facies to around 750 mD for the channel deposits.

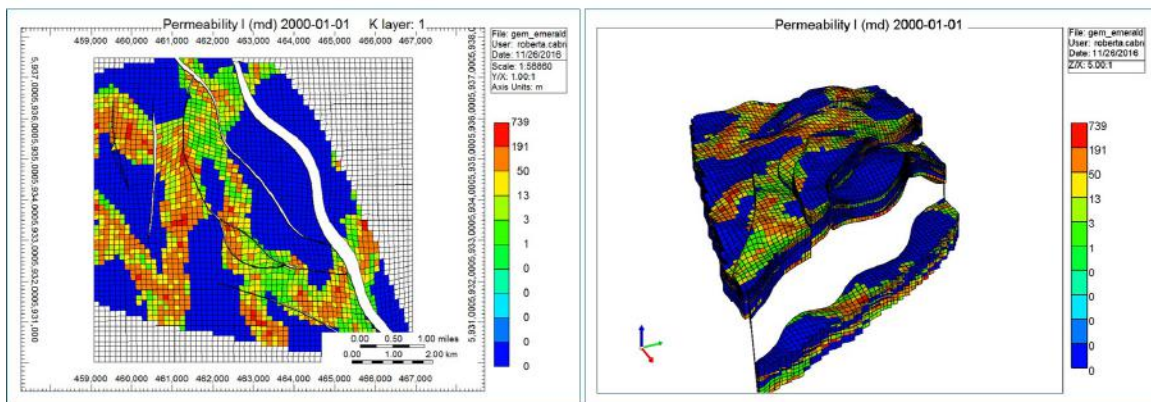


Figure 115: 2D and 3D views of the faulted reservoir model used during the evaluation sessions.

6.6. RESULTS

In this subsection, we present our findings from our qualitative preliminary study. For each of the user sessions, we turned the video and audio records into a written summary in order to highlight the points of discussion that

had been approached throughout the interviews. These key points were broadly organized into two groups: those concerning the features of the application and those related to the comparison between the different immersive interfaces. Then, for each group, the summary points were evaluated and grouped by similarity. And, as pertinent topics were arising, the video and audio records were once again revisited in order to collect meaningful quotes that could illustrate the discussion topics. The results of these efforts follow below in separate subsections for each of the defined groups.

6.6.1. ON THE APPLICATION

6.6.1.1. CELL PROBING

When asked about the usefulness of probing, participants claimed the probe was one of the most common functions used – e.g. “I probe all the time”. They appreciated the functionality of the transparency lens associated with the probe as it allowed them to efficiently see both the property values in the reservoir and the well configuration in 3D. When asked to compare with the existing, traditional modelling and simulation software, participants commented that, when using existing tools, it is required to either visualize the well path across multiple 2D cross sections or view a transparent overview of the 3D reservoir model. When prompted about improvements, two participants mentioned they would like to have the option to add or remove properties to be shown on the probing window.

6.6.1.2. WELL TRAJECTORY CREATION

When asked about the usefulness of the probing tool while selecting the well path, we received several mixed comments. One user commented that only seeing the selected cell is preferable. The main advantage of this, according to the user, is that the 3D geometry of the cell can be clearly observed to avoid the selection of shapes that could cause problems with the numerical simulation such as “pinch” cells. However, another participant commented that it would be preferable to have more contextual information surrounding the selected cell – e.g.

“I would like to see the surrounding grid blocks because I am not sure about the geometry that may surround the well and may impact the flow of oil” and “if I want to make this an exploratory tool, I want to see what I can do, where I can go, what are the options that I can try to create the well trajectory.” This user suggested that displaying the entire column or an adjustable number of cells neighboring the selected cell would be a potential improvement.

Also on that matter, instead of removing the cells inside the lens, one user claimed to prefer seeing them semi-transparent. According to the user, this would be useful for maintaining the reference of a targeted cell when shifting inside and outside the lens. The same user added that, even though he/she sees the wireframe of cells only at the center of the lens, it is still sometimes fuzzy to him/her to understand the geometry of the reservoir by looking at the wireframe inside the lens, specially in regions with a high number of cells. The user then suggested that “maybe, just to show the wireframe of the neighboring cells is sufficient”.

When prompted about the well path creation per se, participants claimed that it would be beneficial to them to be able to select specific ranges of cells to compose a single well perforation instead of considering one cell per perforation, as well as to remove perforations – e.g. “sometimes I already know which layers of the reservoir have oil, so I’ll perforate only those layers”, “the connectivity analysis shows me useless perforations, so I’d like to remove them”.

6.6.1.3. WELL TRAJECTORY ANALYSIS – GRAPH

When questioned about the usefulness of the connectivity graph, we received positive comments – “this is amazing. It’s like I can see where the oil comes from”. Some participants praised the benefit of being able to quickly assess the region that can be covered by a well. They claimed that this would help them find locations that cover larger areas. In a similar manner, they also commented that the connectivity graph would help with

finding locations in the reservoir not covered by existing wells. Users claimed these analyses would help them evaluate how many wells might be needed to cover parts of a development field.

Other participants praised the ability to see the connected cells associated with each perforation, as this would help them more accurately design a perforation scheme that was efficient and without perforations that are not needed or even detrimental to recovery.

When inquired about potential limitations and issues, some participants highlighted the need for higher contrast on the color scheme used on the graph encoding. When prompted for ideas for improvements, one participant in particular suggested it would be useful to highlight locations where the connectivity changed abruptly so barriers to flow could be easily identified. Another participant mentioned that it would be interesting if it also displayed a contour style visualization of the region connected to the well to provide a “general trend” or a “big picture of the behavior of flow” view as opposed to a “detailed” view using the connected graph.

Also, since wells are often associated with specific drainage areas, one participant requested a feature to filter the connectivity graph according to the distance from the well - “when wells are producing, each well is associated with a drainage area. (...) it’s important to see the volume connected to the well only in its share region.” Another participant mentioned that “if we have two wells not too far away from each other, it’s important to see the location and (delta) time when the regions connected to the wells start to interfere with each other”.

The comments mentioned above were based on the fact that if two wells are producing, each one has an amount of recoverable oil volume associated with it. However, if there is interference, wells may not be recovering the expected oil production since they might only produce from the region up to where interference occurs. The same participant also suggested that it would be useful to display the shortest path between a user selected cell and the associated well along the spatial graph. According to this participant, this would provide a visual indication of permeable channel deposits or saturation fronts.

6.6.1.4. WELL TRAJECTORY ANALYSIS – ANALYTICAL PANEL

The panels displaying analytical connectivity data were well received – e.g., “this is very good as it gives a snapshot of the performance of each perforation and the entire well”. All participants felt that it was very important to have both the visual and analytical data available so they could also access detailed information. Users appreciated the ability to move and organize the panels in three dimensional space, but some claimed that it would be interesting to have a way to help organize them, such as being able to set them to hover near their associated well, or automatically stack or sort based on user defined criteria.

Also, one participant indicated that it would be useful to be able to, for different well perforations, select ranges on the histogram of the delta time distribution in order to see the associated connected cells in the spatial connectivity graph. This action, in turn, would help defining the coverage areas for different wells and the inter-well spacing between them – e.g., “if two wells are producing and I see the connectivity graphs overlapping, I would need to decide how much to move one of them further away”.

6.6.1.5. FURTHER FEEDBACKS

When inquired about any additional comments or suggestions, most participants claimed that they wished they could be able to save well trajectory and perforation information in a file format that could be loaded in commercial reservoir simulation software. Requests for the tool to read in trajectories and perforation data from file formats commonly used by commercial reservoir simulation software were also common.

Furthermore, some participants claimed it was important to edit well trajectories and perforations so that different well configurations could be easily evaluated. Further to that, one participant suggested the creation of a feature to support the comparison between two or more configurations of the same well. Another participant indicated it would be useful to be able to save and recall well configurations, so that a saved well configuration could be “pasted” into arbitrary locations in the reservoir as often as needed.

One participant suggested to extend the current “well connectivity analysis” to “inter-well connectivity analysis”. The participant highlighted the fact that, for senior fields, engineers must commonly decide on the placement of injector wells based on the already existing producer wells. Thus, the evaluation of the connectivity between a pair of producer and injector wells would be even more helpful for his/her work.

6.6.2. ON THE IMMERSIVE INTERFACES

6.6.2.1. CAVE AND OCULUS

Following the interview regarding the application features, we captured each participant’s impressions and preferences regarding two immersive environments, the CAVE and the Oculus Rift, using the same tracked gamepad controller. Participants claimed the visualization of the reservoir was clearer when using the Oculus Rift. They reported that the perception of details was better than in the CAVE. Most participants (five out of six) appreciated increased detail, relating that to the feeling of actually being immersed in the environment – e.g., “it makes me feel more ‘into it’; as if things are closer”.

One of the participants claimed that what they saw in the Oculus looked more familiar – “things looked more aligned to what I know”. On the other hand, another participant argued that “immersion on the Oculus is so high that it may complicate the analysis of the reservoir as a whole”. The participant was alluding to the fact that the field of view is narrower in the Oculus compared to the CAVE system and this was a hindrance to visualizing the reservoir – “we are used to have a wider vision”.

When queried about the usefulness of both devices, most participants claimed they would prefer to work with the Oculus if they were working on their own; however, they would prefer to work with the CAVE if there were more individuals involved and the purpose concerned collaboration with co-workers, to showcase projects, or to make joint decisions. The participants did claim they tend to work individually, but acknowledged that a device for collaborative work (“which is common in industry”) is very important.

When it comes to difficulties and limitations, one of the participants claimed they “felt more in control with the Oculus when compared to the CAVE”. This participant pointed out that one of the shortcomings of the CAVE is that “it shows things according to a particular person (...) this may be the reason why I have felt some eye fatigue”. Another difficulty mentioned by two participants was regarding comfort while using the Oculus Rift. One of them mentioned that “technology may need to advance to be more practical” after noting the heaviness of the Oculus compared to the glasses used in the CAVE.

As a suggestion for improvement, one participant claimed that it would be ideal if more users could “control the environment” when using the CAVE or “share the environment” when using the Oculus. This participant suggested that equal autonomy for users in the virtual environment would improve both the quality of the experience and the potential for collaboration.

6.6.2.2. GAMEPAD AND LEAP MOTION

Finally, we captured impressions and feedback regarding experiences when using the Oculus Rift with two different interaction interfaces. Participants interacted with the application using a tracked gamepad first, and then hand gestures.

All participants claimed that gestures were more natural and intuitive for them – e.g., “your hands are part of you, so you know how to control it”; “(your hand) it’s you. (...) with the gamepad, you have the buttons as an additional layer”. Nonetheless, they also claimed there was a learning curve and some period of adaption would be necessary – which they did not think was problematic.

Although users agreed that gestures would feel more natural, they acknowledged that using a gamepad provided benefits that gestures do not allow. One of the participants noted that using gestures meant that hands were required to be “fixed in a particular spot” – that is, on the space covered by the gesture and motion sensor; whereas the gamepad allowed more flexibility. The same participant also noted that the use of gestures seemed

to lead to some glitches during the interaction process; something the user did not notice when using the gamepad.

Another participant claimed that the gamepad allowed users to see and reach farther than hands and arms physically allow when probing – “most of the time, I probe multiple cells in different regions, so I don’t want to be too close to a cell”. The user also claimed fatigue as another potential problem with using hands – “probing is a feature that I use all the time, so having my arms resting on the chair is more convenient”. The participant suggested that attaching a cursor to the index finger in order to probe cells further away could be an improvement. In spite of the drawbacks mentioned, this participant claimed he/she would prefer to use gestures instead of the gamepad.

6.7. DISCUSSION

6.7.1. ON THE APPLICATION

There were several suggestions and positive comments towards our application. A significant number of them were specifically directed to the use of connectivity analysis. This result is extremely positive as it shows that, even though static connectivity analysis was a new approach to them, they understood the purpose of it and its usefulness. Not only were participants pleased with the connectivity technique, but they also claimed that the three-dimensional perspective in both environments was superior to their traditional desktop environment. These results support our belief that higher immersion is indeed beneficial for our application.

Interestingly, several participants provided feedback that echoed some of the challenges VR has historically faced in the field of oil and gas mentioned in [213]. For instance, many participants requested the ability to save and load well trajectories, fault information, and other data from a software familiar to them. During the feedback sessions, it was clear that it was important for participants to easily transfer data between their existing workflows and our tool. In this sense, lack of compatibility with software packages had been identified in [213]

as a barrier to the use of VR in oil and gas. For this reason, we chose to use VTK as our data format since it is widely used. Also, we created tools to convert reservoir data formats to the VTK format we use, increasing the ease of importing data into our tool. However, we have not yet built support for saving or loading well trajectories. This could be added relatively easily and clearly it is desirable to do so.

6.7.2. ON THE IMMERSIVE INTERFACES

When it comes to immersive mediums, [213] noted that accessing a large system that would require users to leave their desks would be a barrier to widespread use of VR in oil and gas. However, there are currently high-quality virtual reality devices such as the Oculus Rift and the HTC Vive, and vendors are offering enterprise class VR workstations. It is thus now possible to affordably have a VR headset at a workstation for individual use or remote collaboration, and to move to a larger visualization platform such as a CAVE or CAVE variant only for co-located collaboration, training, and demonstration to groups. This is extremely relevant as many participants mentioned they would prefer to use HMDs such as the Oculus Rift for individual work whereas they would rather use the CAVE when performing collaborative work. The CAVE used in this primary study is relatively low resolution, but a high resolution system should address the concerns regarding clarity that were expressed.

Also on that matter, participants appreciated the ability to move between the head-mounted display and the large-scale immersive display, but were afraid that would require them to go through two types of training. As could be observed in Chapter 2, there are many different input and output interfaces, and to design a solution that provides user experiences on commodity HMDs and other immersive platforms is both a technical and a design challenge. We used MiddleVR to mitigate the technical challenges of cross-platform VR applications. However, there are several important design considerations when developing cross platform VR applications.

As we discussed in Chapter 2, there are various properties associated with input and output devices in VR systems, and many of these properties vary from one device to another. The strengths and weaknesses of (the

combination of) these properties among devices, in turn, may have a strong effect on the user interaction design as some interaction techniques may be more or less suitable across distinct devices, or may even only work with devices that have specific characteristics.

In the case of input devices, basic properties that should be considered are size and shape: larger devices that need to be worn or held by users are primarily controlled by larger muscle groups of shoulder, elbow, or wrist; whereas smaller devices use smaller, faster muscle groups as the fingers. Additionally, smaller devices may also decrease clutching – that is, the releasing and re-grasping of an object in order to complete a task due to not being able to complete it in a single motion [15]. Other properties that may also play a role on the interaction design include the number of degrees of freedom; the device tracking reliability; hands-free and/or eyes-free capability; haptics capability; user mobility; encumbrance; etc.

For output devices, in turn, a characteristic of visual displays that may strongly affect the interface design is FOV. The normal horizontal FOV for a human is approximately 200 degrees, with 120 degrees of binocular overlap. In a three-sided CAVE, the FOV is 170 degrees wide when the user faces forward; whereas head-mounted displays have narrower FOV that are no more than 100 degrees wide. In a narrow-FOV HMD, information density may easily lead to a cluttered display and, as the user might not be able to see all virtual objects at all times, they might have to move around their head a lot to compensate. Other properties to be considered include spatial and temporal resolution, room space, encumbrance, number of display channels, user mobility, latency tolerance, etc.

In this sense, even though our design accounted for interaction interfaces coherent in each environment individually, it is indeed essential to provide a coherent experience so that users are not confused when moving between platforms. Users should not require more than a trivial amount of training when moving between platforms, otherwise the effectiveness of moving between IVEs is compromised.

6.8. FUTURE WORK

In this section, some directions for future work will be discussed based on the feedback collected from the specialists during our evaluation sessions. The discussion points are organized according to the three elements investigated throughout the evaluation: the analytical method to carry out static connectivity analysis; the application features for performing well placement studies using the analytical method; and the immersive interfaces.

6.8.1. ON THE STATIC CONNECTIVITY ANALYSIS

There are a number of future perspectives for our proposed analytical connectivity analysis. One of them concerns extending our current approach to also handle secondary recovery scenarios. Another possibility is to expand our method to also consider connectivity analysis per distance. On this matter, some valuable questions to be investigated are “given a zone of coverage around the well-bore, what is the connectivity behavior within this area?”, “for a certain distance d from the well, what is the maximum delta time to cover the entire sector?”, “what is the minimum time to achieve the outer limits of the region?”, “what is the total drainable volume connected to the well within the belt?”, etc.

During our interviews with the subject matter experts, we discovered that the analysis of connectivity per distance may be useful due to the fact that characteristics related to the coverage areas of drilling wells may commonly drive the formulation of well plans. As an example, engineers may isolate or remove potential areas from the drilling operations such as complex mountainous areas. As another example, there may have situations in which two or more oil companies share a continuous producible region. Thus, when assessing well placement scenarios, a company may want to restrict the evaluation to its specific share.

Furthermore, other graph theoretical algorithms can be used as complementary methods to evaluate connectivity from different perspectives. As an example, in graph theory, a cut is a partition of the vertices of a

graph in two disjoint sets. An s-t cut is a cut that requires the source and the target nodes to be in different subsets. Additionally, any cut determines a cut-set, the set of edges that have one endpoint in each subset of the partition. And each cut-set determines a unique cut. A minimum cut, in turn, is a cut that is minimal in some sense such as minimum weight sum of the edges in the cut-set, or minimal number of edges in the cut-set.

For a porous media such as a reservoir model, one can determine the critical nodes and edges which, when removed, disconnect the graph. These pivotal locations represent potential bottlenecks whose loss prevents fluid to flow through the reservoir (see Figure 116). Thus, minimum cuts and derived metrics can be used to portray the vulnerability of reservoir connectivity. For instance, for a graph connected to a producer-injector pair that has multiple cut-sets, these locations can be used to estimate the likelihood of severing the inter-well connectivity and, thus, the potential loss of sweep volume.




Figure 116: Minimum cuts may highlight potential reservoir segmentations (compartments) that lead to loss in production volume.

Finally, another alternative for future investigation is to evaluate the performance of our current method by either running performance tests with larger reservoir models, or by comparing production estimations with flow simulation results.

6.8.2. ON THE APPLICATION

There are also many interesting future directions for our application for well placement optimization studies. Concerning the definition of free-form wells, we received requests to display more contextual information surrounding cells that constitute the candidates to compose the well trajectory. On this matter, a simple solution is to create a spherical lens around the probing cube to display neighboring cells, and the user may modify the radius of the sphere. This is somehow a straightforward solution as our current version of the application already supports the use of a volumetric focus lens with stylized variations in shading effects for reservoir sectioning, as explained in Section 5.3.2.1.

Another solution is to allow the user to manually specify the shape of the lens. And, by taking advantage of the grid-based structure of reservoir models, this shape could be defined as connected corner points of cells, selected by the user directly on the reservoir grid. The solution could also provide an interface for vertex- or face-level editing. Figure 117 depicts a conceptual scenario to illustrate this idea.




Figure 117: Conceptual prototype for capturing gestural movement over grid cells as a 3D sketch and translate it into a volumetric lens.

We also received a number of comments with reference to the selection of a sequence of cells to compose a well perforation; rather than our current “one cell per perforation” scheme, engineers suggested to also consider more realistic scenarios with multiple cells per perforation, as well as cells along the well trajectory that are part of no perforation. Since our current metaphor for defining free-form wells was well-received during our

evaluation sessions, a potential approach to define perforations is to allow users to select cell intervals directly onto the well trajectory (see Figure 118). Further to this, engineers also mentioned the need to edit and delete both wells and well perforations.




Figure 118: Theoretical design for selecting a series of contiguous cells along the well path to compose a perforation.

When it comes to the well analysis using the connectivity graph, some possibilities for future development include highlighting the shortest path, or the vertical cross section defined along the path, whenever a connected cell is probed. Another possibility is displaying data on the information window related to the connectivity of the probed cell, such as the distance between the cell to its associate perforation, the cumulative volume along the path, the connectivity measure associated with the cell, etc.

Moreover, in addition to displaying the graph for running a detailed connectivity analysis, visualizing the contours of connected regions may also be useful for performing rapid trend analysis. These include comparing the areas of connected regions or searching for overlapping regions. In this sense, set operations could be performed in order to extract interesting portions of the connected data. As an example, consider a case where an engineer wants to evaluate two potential configurations for a producer well w_p . By defining and running the connectivity algorithm on the two viable well trajectories, the engineer obtains the associated connected regions, C_1 and C_2 . Following, the engineer may wish to visually evaluate the production performance of the two

placement scenarios. By applying a set operation such as $C_1 - C_2$, the resulting set C_3 represents the region covered by the first but not by the second scenario.

As another example, consider a placement scenario illustrated in Figure 119 (a) with three different producer wells, w_{p1} , w_{p2} , and w_{p3} and their associated connected regions C_1 , C_2 , and C_3 , respectively. And, in this case, imagine that the engineer is searching for interference regions between neighboring drainage areas. As illustrated in Figure 119 (b, c, and d), operations such as $C_1 - (C_1 \cap C_2) - (C_1 \cap C_3)$ exhibit the region covered solely by the well w_{p1} , and resulting sets from operations such as $(C_1 \cap C_2) \cup (C_1 \cap C_3) \cup (C_2 \cap C_3)$ represent overlapping, problematic regions covered by the wells.

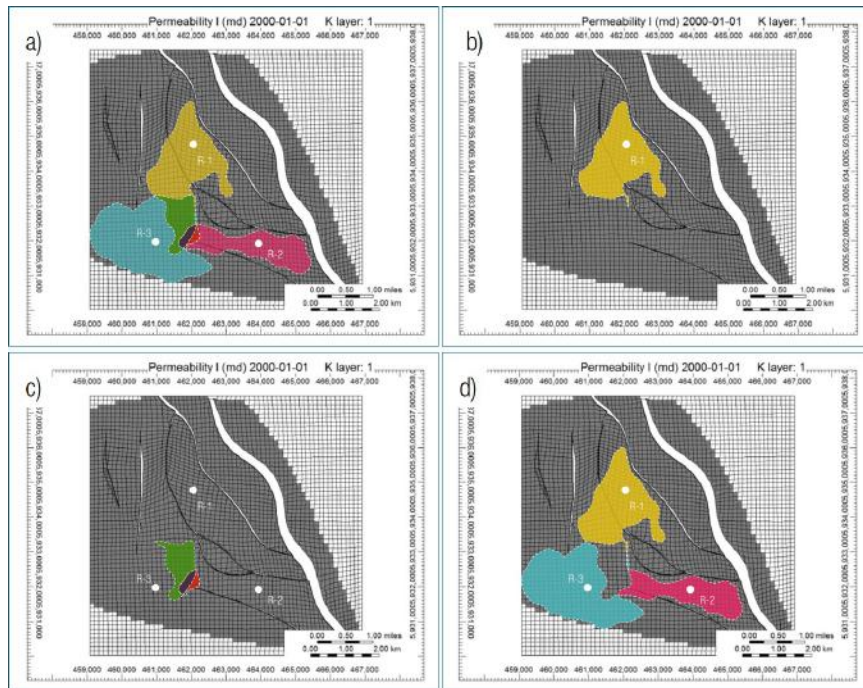


Figure 119: Illustrative set operations applied to connected regions on a reservoir model.

When it comes to the evaluation of wells using the analytical panels, an interesting extension would concern enhancing the interactivity between the panel and the 3D model. For instance, whenever the user selects a table row related to a particular well perforation w_p , the associated connected cells and w_p would be highlighted

in the model. In a similar manner, if the user selects a range in the histogram of delta times of cells connected to w_p , the cells within this interval would be highlighted in the 3D model.

Furthermore, since our tool aims to be inherently exploratory for well placement studies, multiple placement scenarios are to be created and compared. Therefore, future work can implement techniques to help users to accomplish comparison tasks more efficiently. Particularly for the visual comparison of the analytical connectivity panels, authors in [214] list three comparative visualization techniques: juxtaposition, superposition, and explicit encoding of relationships.

For juxtaposition, objects are placed separately in either time or space, and comparison relies on the viewer's memory to make the connections between objects. This is somewhat our current approach as users may freely arrange panels in space. Yet, our current approach of providing fully flexible arrangements may be advantageous to rapidly manipulate panels in order to group them or to disclose occluded parts of the scene; but, for comparison purposes, people commonly use borders or patterns to guide the objects arrangement. Therefore, our present approach could be extended to include the grid-based snapping in order to assist the alignment of panels to be compared, as illustrated in Figure 120.

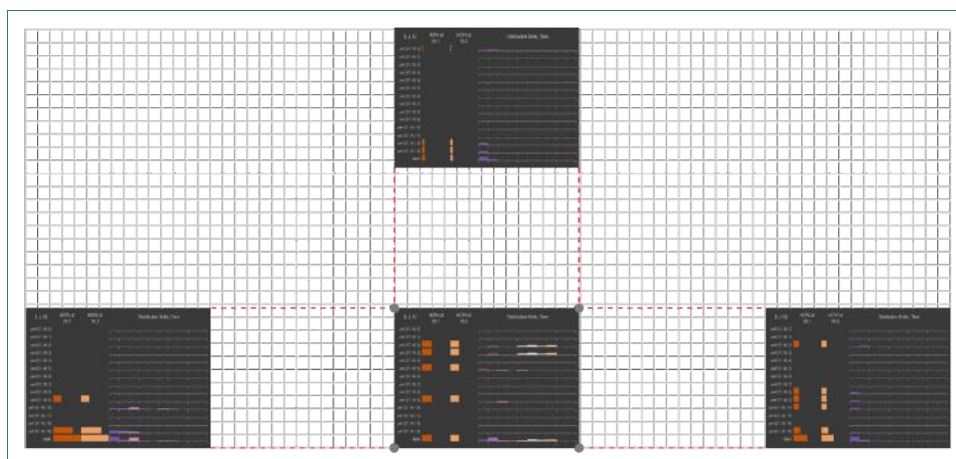


Figure 120: A grid-based snapping provides implicit vertical and horizontal lines that almost seem magnetic. When one drags an object on the grid, the object will align (“snap”) to the nearest intersection on the grid or to other objects. This is enormously useful for accurate placement of image elements.

For superposition designs, objects are presented at the same place and time. In our specific case, by selecting two or more panels, they could be automatically stacked on top of each other and occluded information would be revealed by alpha-blending the views. The fading effects could be either be automatically animated from fully visible to invisible, or the user could control the transparency levels, as can be seen in Figure 121. It is worth mentioning that, even though this overlaying design may offload the burden of side-by-side comparison from the viewer, it may have issues with clutter and scalability.

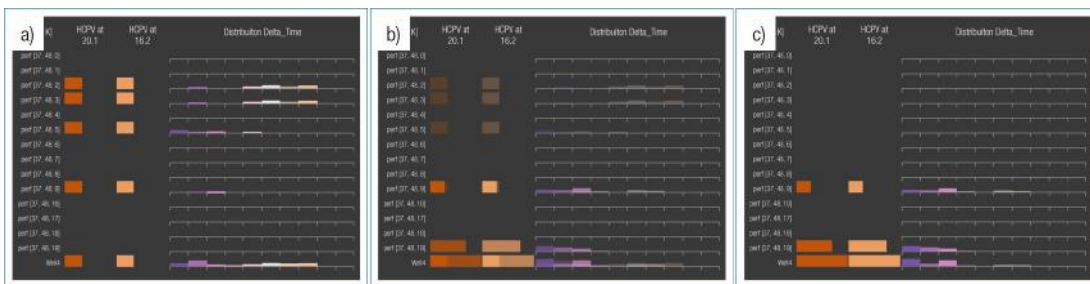


Figure 121: Panel comparison using alpha blending would allow gradual transitions from one panel (a) to another one (c).

For explicit encoding, the relationships between objects are computed and visual encoding of the relationships is provided. In our scenario, by selecting two or more panels, a relationship of our panels could express the relative differences of the graphical elements and metrics displayed along the table as illustrated in Figure 122. It is interesting to note that, although this explicit encoding design scales fairly well, it may require the relationships between objects to be known and may need to address issues of decontextualization [214].



Figure 122: Comparison between two panels using explicit encoding: the resulting panel (c) shows the complement between the first (a) and the second (b) panel.

6.8.3. ON THE IMMERSIVE INTERFACES

Particularly in terms of our investigation on the different immersive technologies, we suggest that some research directions for the short-term future should focus on designing the remaining set of features for the air-based gestural interface; and, for the long-term, running a more thorough user evaluation and, then, revising our interaction vocabulary for further improvements.

Other directions for the future can also be about transferring the application to other immersive mediums. Our research group recently purchased an HTC Vive and a Microsoft HoloLens. Therefore, it is the group's interest to port the application to these latest technologies.

Finally, other possibilities for future work could concern exploring the design challenges for cross-platform development in virtual reality. On this matter, an interesting leading research question is the following: "what should drive the cross-platform design: interaction consistency or user experience?". In the former case, it would be more important to maintain uniformity between interaction styles across different platforms. Therefore, the priority would lie on comparable instead of meaningful interaction designs. In the case of user experience, however, it would be more important to make the best use of the strengths of each interface to create an overall user experience greater than the sum of its parts. Therefore, designs that feel natural for the user should be primarily considered and, so, it would be absolutely not necessary for interaction styles to be identical across multiple environments. Nevertheless, in this case, further potential research questions would be the following: "how to create intuitive logical mappings between interaction styles?", "what factors should be considered when designing the mappings? ergonomics? learnability? memorability?", and "how to maximize or minimize these factors?".

7. CONCLUSION




Figure 123: Immersive virtual environments have been considered promising mediums to address particular needs of oil and gas domain. In the figure above, users evaluate a petroleum reservoir field using a CAVE system.

As discussed in the introductory chapter, the placement of wells is of great relevance for the production of petroleum products. Yet, current forms of optimizing well placement face difficulties due to factors associated with both hardware and software. The former concerns the use of standard desktop systems, which provides little opportunity to attend specific demands from the oil and gas industry, like handling inherently spatial tasks. Factors associated with software, in turn, refer to the use of reservoir flow simulation, which is a time-consuming numerical method used to estimate performance of production scenarios. These limitations force engineers to greatly reduce the number of placement scenarios that can be searched before selecting the optimal scenario; they thus prevent potentially good scenarios to be tested.

In this thesis, we put efforts towards defining forms of mitigating the two limitations abovementioned. We investigated and engaged in the design, implementation, and evaluation of an application that leverages the benefits associated with more immersive interfaces. We thus gained insight on the potential of more immersive

technologies to provide an improved user experience and possibly better decision making. Further to this, we also gained a deeper understanding on the usefulness of using fast performance estimators, such as static connectivity analysis, to be used as a surrogate for dynamic flow simulation. This methodology can be used for a number of applications, including ranking geological realizations, upscaling validation, and the optimization of well placement. By specifically targeting these two tasks, we aimed to take part in developing new exploration paradigms to support more exploratory well placement studies.

In this chapter, we will revisit the main contributions of this thesis as well as present our closing remarks. Specific possibilities for future work will not be outlined in this chapter. They were discussed in depth in Chapter 6, when we provided potential research questions and discussed how the feedback we gathered throughout our evaluation sessions can be transformed into guidance for future work. This discussion involved a number of perspectives concerning the analytical method we proposed in this thesis, our application's features to perform well placement studies, and the use of more immersive environments for well optimization and for the broader oil and gas domain. The main contributions of this thesis, in turn, are revisited below:

- **An analytical method to use static connectivity analysis as a surrogate for reservoir flow simulation in optimizing well designs;**

In Chapter 3, we presented our analytical method to rapidly estimate well placement scenarios. Our method adopts algorithmic graph theory for computationally efficient static connectivity analysis in either primary or secondary recovery scenarios; beyond that, it also accounts for the numerous existing graph algorithms that can be included in future for assessing connectivity from different perspectives. Our algorithm computes a time-based connectivity metric that is sensitive to both geological heterogeneities and fault structures of the reservoir model.

For the development of our method, we received guidance from a post-doctoral scholar in Petroleum Engineering; and we also investigated existing reservoir connectivity techniques, traditional dynamic reservoir simulation, and the drawbacks they present – all of these considerations were discussed in Chapter 3 prior to presenting our technique. In this sense, our method was proposed to be used as a light-weight proxy for costly dynamic simulation, thus allowing the exploration of multiple placement scenarios. And, after collecting the most promising scenarios, fluid flow simulations could then be used to assess this best-small subset.

- **An application, which uses our analytical method, to serve as an exploratory tool to run well placement studies;**

In Chapter 5, we presented our application to run exploratory studies of placement scenarios using our analytical connectivity method. Prior to the development of our application, we took into consideration which and how immersive technologies have been used by the oil and gas industry; and we also considered the systems that are currently available for engineers to perform well planning and what limitations they present – all of these factors were outlined and discussed in Chapter 4.

Furthermore, our application runs in three immersive virtual environments: a CAVE with a tracked gamepad; a HMD with a tracked gamepad, and a HMD with a leap motion controller. These immersive interfaces were investigated to be alternatives to the traditionally used desktop displays, mouse, and keyboard. By using the more immersive technologies, our proposed application leveraged their unique benefits to create techniques to view and interact with reservoir data in a way that facilitates the tasks associated with well planning.

- **Reflections on engineers' standpoint on our analytical technique, our software, and the use of more immersive technologies for well planning.**

In Chapter 6, we provided a qualitative primary user study reporting engineers' perceptions and preferences in terms of our analytical connectivity technique, our exploratory application, and the use of more immersive interfaces. Besides, we also elaborated on their impressions and suggestions in order to highlight the usefulness of either static connectivity analysis for fast flow diagnostics, as well as the applicability of more immersive interfaces for well optimization and for the broader domain of oil and gas. Building on the feedback collected and on our reflections, we also provided ideas for future work on each of the three themes investigated throughout the study: static connectivity analysis, features for planning well placement scenarios, and immersive technologies with high levels of immersion.

By considering engineers' point of view, we brought closer the two domains on which this thesis is built, namely Computer Science and Oil and Gas Engineering. Beyond that, this primary study was of great relevance as it offers practical insight on the challenges faced by well placement studies and helps guide multidisciplinary approaches as this one.

7.1. CLOSING REMARKS

Immersive virtual environments have been increasingly used in the last decade, assisting several domains through novel technologies that allows for experiences that would otherwise not be possible. The use of immersive environments with higher levels of immersion has allowed for an increased sense of presence, enhanced spatial understanding, larger contextual information space, and interactions with the virtual world that feel more natural to users.

This thesis has particularly focused on how immersive environments can push the boundaries of the oil and gas industry and help advance well placement optimization. More specifically, this thesis presented efforts towards two lines of research: analytical connectivity techniques that focus on computational speed for supporting interactive diagnostics of production scenarios specified by engineers; and beyond-the-desktop technologies that foster benefits of higher levels of immersion for facilitating tasks associated with well planning. By targeting these two fields of research, our work aimed to tackle more exploratory well placement studies.

Our research was designed and implemented building on a theoretical foundation which was obtained from surveys, the input and guidance of a post-doctoral petroleum engineer who is part of our research group, and, finally, feedback from other domain specialists during evaluation sessions. Knowledge obtained from these efforts provided awareness on key concerns in the design of immersive systems and fast performance estimators for well optimization and for the broader domain of oil and gas. Beyond that, it highlighted numerous possibilities for improvement within the rich, multi-faceted oil and gas field. Hence, we expect this thesis can prove to be useful to inspire and guide future endeavors towards new solutions for the domain.

APPENDIX A

This appendix contains all of the material used for conducting the qualitative primary study described in Chapter 6 of this thesis. It includes the document provided to the participants for introducing the study, the pre-study questionnaire for collecting background information about the participants, and the description of the five procedures that constituted the study.

Introduction: Presentation and Structure

Thank you for your participation in evaluating this prototype. We developed an application for the exploration of well trajectories of reservoir models using virtual reality, and we would like to hear your subjective insight on them.

There is no right and wrong, so express yourself freely. Tell us what you think about it. Your opinion is very valuable to us, and criticism is very welcome too. We are also seeking constructive opinions on how to improve and evolve our ideas. This study is neither task- or goal-oriented. You will not be asked to perform a specific task, but rather to see a set of tools, and to brainstorm with us about them. This evaluation will be divided as follows:

- 1) First, we would like you to fill out this simple questionnaire so we can know a bit more about your experience.
- 2) Then, we will present and explain the features of our application, and you can try testing them yourself if you want. Meanwhile, we will ask you to “think aloud” or, in other words, to constantly verbalize your thoughts. Your ideas will be recorded for further evaluation and consideration... again, there is no right or wrong, so please express yourself freely, whatever you are thinking. What might seem a trivial thought to you might be insightful for us. We will ask you a series of questions about your vision and opinion on the system to be presented.
- 3) Finally, we will conclude and welcome any additional comments, if there are any.

Questionnaire

- 1) Name:

- 2) Age:

- 3) Current professional/academic position:

- 4) Professional background/education and experience (include no. years):
 - a. Academic training (degrees, specializations):

 - b. Professional training (w/professional positions):

- 5) Res. engineering software tools used in the past, for how long and for which role.

- 6) How comfortable are you when using new interaction mediums/ new software?
(Scale 1-5; 1- Not at all comfortable, 5 - Very comfortable)

- 7) Previous experience/use of reservoir simulation models?

Procedure 1

⇒ BASIC MANIPULATIONS

- Procedure
 - Rotate
 - Translate => Go inside the model, walk inside it, then go outside
 - Scale

⇒ CELL PROBING

- Procedure
 - Probe cells with contrasting property values
 - Probe pinch-out cells

⇒ WELL TRAJECTORY: CREATION

- Procedure
 - Mention there are 2 types of wells that can be created
 - Create some vertical wells, put transparency on reservoir to show the wells, then remove transparency
 - Go inside the reservoir
 - Create one free-form well
 - Go outside the reservoir
 - Put transparency on reservoir, then remove transparency
 - Scale down the reservoir, create one free-form well from the outside of the reservoir
 - Put transparency on reservoir, then remove transparency
- Questions
 - What do you think about the two strategies for well path creation?
 - Easy to use / Cumbersome? To what degree?
 - Useful / Useless? To what degree?
 - What are the advantages of these strategies for well creation according to you?

- Do you see any potential problems with the way they currently work?
- Would you have any ideas on how they could be made better?

Procedure 2

⇒ WELL TRAJECTORY: ANALYSIS: GRAPH

- Load "Scene 2"

- Speech 1
 - Besides being interested in the creation of wells, we are also seeking to develop a simple method to evaluate the performance of wells.
 - Such method calculates a graph that shows the region connected to the well every time a new well is created.
 - The edge's color in the graph makes reference to "how fast" the existing oil in that area can flow to the well to which it is connected. You can observe there are different colors. The edges that are in purple indicate areas in which the oil would reach the well faster when compared to the edges in orange. The more purple, the faster; the more orange, the slower. In other words, given that there is oil in <location 1> and <location 2>, the oil in which region would be produced first? The oil present in <location 1> would be produced first than the oil in <location 2>.
 - This measure thus represents relative time, and we call this property "delta time" or simply "time".
 - We are able to analyze the behavior of the region connected to the well by altering the threshold of the delta time. In other words, we can choose a certain time limit if we want to see only areas that are below this particular value.

- Procedure 1
 - Move close to the well with the salient orange region
 - Change threshold value to see variation on the graph

- Speech 2
 - Now, let's observe this region in orange, meaning that if there is oil here, it will take longer for it to be produced by the well. We want to understand why it takes longer here.
 - 2.1 - So, we can start probing and analyzing these cells. Since these cells are neighbor cells in y direction, we look at Permeability J, Porosity, and Transmissibility Y. Note that, along the j direction,

- the cell is becoming less permeable, less porous, and with lower transmissibility capability. We know that these properties are strongly related to the easiness through which the oil flows in the reservoir. Meaning, we can conclude that the oil in this orange region takes longer to be produced because the values of these properties are gradually becoming smaller.
- 2.2 - Cell probing can also be useful to analyze why cells that are part of the boundaries of the region connected do not belong to the region. For instance, if we analyze these two cells <cell 1> and <cell 2>, we can see that the cell right below does not belong to the region connected to the well because the transmissibility in z direction of the cell above it is zero.
- Procedure 2
 - Stay close to the with the salient orange region
 - Put reservoir transparent
 - For 2.1, probe cells along (38, 47, 3) and 38, 43, 3)
 - For 2.2, probe cell (40, 49, 9) or similar to see that Tz is zero.
 - Questions 2
 - What do you think about this function (probe the cells on the region connected to the well)?
 - Easy to use / Cumbersome? To what degree?
 - Useful / Useless? To what degree?
 - Do you see any potential problems with the way they currently work?
 - Would you have any ideas on how they could be made better?
 - Speech 3
 - There are three ways of visualizing the region connected to the well.
 - Procedure 3
 - Pass the lens to see the graph over the opaque reservoir
 - Put the reservoir transparent to show transparent reservoir + graph, then pass the lens to see only the graph over the transparent reservoir
 - Put the reservoir invisible to see only the graph

- Questions 3
 - o What do you think about the three strategies to visualize the region connected to the well?
 - Easy to use / Cumbersome? To what degree?
 - Useful / Useless? To what degree?
 - o Which one do you prefer? Why?
 - o Do you see any potential problems with the way they currently work?
 - o Would you have any ideas on how they could be made better?

- Speech 4
 - o In case there are more than one well in a reservoir, we'll be able to visualize only the region connected to a particular well
 - o Besides visualizing the region connected to a particular well, we can visualize the region connected to a specific well perforation

- Procedure 4
 - o Put transparency on the reservoir
 - o Select one well head
 - o Select one well perforation of the previously selected well

- Questions 4
 - o What do you think about this function (well /well perforation selection)?
 - Intuitive / Not intuitive? To what degree?
 - Useful / Useless? To what degree?
 - o Do you see any potential problems with the way they currently work?
 - o Would you have any ideas on how they could be made better?

Procedure 3

⇒ WELL TRAJECTORY: ANALYSIS: PANEL

- Start separating one particular panel from one well

- Speech 1
 - Besides analyzing the graph of the region connected to the well, we can also analyze the performance of the well by looking at some quantitative information related to the region connected to the well.
 - For each newly created well, a new panel will be created to display this information.
 - Each panel is a table composed of four columns:
 - The first column lists the IJK coordinates of each of the well perforations that compose the well. Thus, each row on the table shows information of that particular perforation and the last row shows information of the entire well.
 - The second column shows the total hydrocarbon pore volume of the region connected to each perforation – and the last row shows the total hcpv of region connected to the entire well.
 - [Can you tell me which perforations are producible and which ones are not?]
 - The third column shows the hcpv of the region connected to each perforation at the specific time limit established.
 - Thus, currently the time threshold is < > and the region connected to this well has total hcpv of < >.
 - If we decrease the time threshold, what happens? We can notice that the region connected to the well is decreasing and so the total hcpv at this time.
 - The last column shows the time distribution of the region connected to each well perforation – and the last table row shows the time distribution of the region connected to the entire well.
 - [Can you tell me how many cells connected to the first perforation have low delta time?]
 - [Can you tell me which perforations have more producible regions? Why?]

- Questions 1
 - What do you think about this feature?

- Useful / Useless? To what degree?
 - Would you have any ideas on how this could be made better?

- Speech 2
 - Since you may have multiple wells on the reservoir, you may have multiple panels and you may want to compare the performance information of different wells.
 - Therefore, you can freely re-arrange each panel individually.
 - You can move the wells that have the best performance at one location, and the ones you still need to evaluate at another location.
 - You can organize the panels hierarchically based on some criteria – for example, the well performance.
 - You can also scale up a panel that you wish to further analyze, or scale up two or more panels to compare them.

- Procedure 2
 - Organize some panels on the right wall and organize some panels on the bottom wall
 - Organize the bottom panels hierarchically (just to give an idea)
 - Scale up the panel with highest performance, then scale up another panel

- Questions 2
 - What do you think about this feature (freely re-arrange the panels)?
 - Intuitive / Not intuitive? To what degree?
 - Useful / Useless? To what degree?
 - Do you see any potential problems with the way they currently work?
 - Would you have any ideas on how this could be made better?

Procedure 4

⇒ CAVE and OCULUS

- Procedure
 - Initially, experimenter controls the joystick
 - Load “Scene 2”

 - Follow these steps:
 - Move the user inside the reservoir
 - Ask the user to move his/her head
 - Probe cells of the connected tube
 - Move the user out of the reservoir
 - Ask the user to look up to see the GUIs
 - Scale up one of the GUIs
 - If the user wants, he/she can manipulate the joystick from now on

- Questions
 - What are the differences you can notice between the two environments?
 - Can you tell me advantages of using the CAVE environment? Can you tell me disadvantages?
 - Can you tell me advantages of using the Oculus? Can you tell me disadvantages?
 - Can you tell me scenarios where you believe the CAVE would be preferred to use?
 - Can you tell me scenarios where you believe the Oculus would be preferred to use?
 - Which environment do you prefer? Why?
 - [optional] Would you use the Oculus+Joystick in your daily work?

Procedure 5

⇒ OCULUS + LEAP MOTION

- Procedure
 - Initially, experimenter controls a joystick to move the user
 - Load “scene 2”

 - Follow these steps:
 - Ask the user to look at and move their hands
 - Right index to use transparency lens
 - Move the user close to one cell/graph
 - Right index to probe the graph
 - If the user wants, he/she can fully interact with the application from now on

 -

- Questions
 - What are the differences you can notice between this environment and the two previous environments?
 - Can you tell me advantages of using your hands instead of the joystick?
 - Can you tell me disadvantages of using your hands instead of the joystick?
 - Can you tell me scenarios where using your hands instead of the joystick would be preferred?
 - Which interaction do you prefer (using gestures or manipulating the joystick)? Why?

REFERENCES

- [1] A. S. Al-Naji and M. I. Gaddourah, "Exploration and Petroleum Engineering Computing in the 90's," Society of Petroleum Engineers, 1993.
- [2] A. Z. Austin, Y. K. Choo, M. E. Smith and D. G. MacDonald, "Application of 3D Visualization Software to Reservoir Simulation Post-processing," Society of Petroleum Engineers, 1992.
- [3] W. B. K and R. L. Brown, "An Interactive 3D Visualization Tool for Reservoir Simulation Data, Including Fluid Flow and Particle Tracking," Society of Petroleum Engineers, 1992.
- [4] "CMG Software Overview - 2016," Computer Modelling Group LTD., [Online]. Available: <http://www.cmgl.ca/software>. [Accessed 7 October 2016].
- [5] "Petrel E&P Software Platform," Schlumberger, [Online]. Available: <https://www.software.slb.com/products/petrel>. [Accessed 7 October 2016].
- [6] "Nexus Suite Reservoir Simulation," Halliburton, [Online]. Available: <https://www.landmark.solutions/Nexus-Reservoir-Simulation>. [Accessed 7 October 2016].
- [7] Z. M. Filho, E. V. Brazil and M. C. Sousa, "Exploded View Diagrams of 3D Grids," *Conference on Graphics, Patterns and Images, SIBGRAPI*, pp. 242-249, 2015.
- [8] F. M. de Carvalho, E. V. Brazil, R. G. Marroquim, M. C. Sousa and A. Oliveira, "Interactive Cutaways of Oil Reservoirs," *Graphical Models*, vol. 84, pp. 1-14, 2016.

- [9] Z. Martins, F. de Carvalho, E. V. Brazil, M. Sousa and R. Marroquim, "Cutaway Applied to Corner Point Models," in *Conference on Graphics, Patterns and Images, SIBGRAPI*, 2012.
- [10] R. Amorim, E. V. Brazil, F. Samavati and M. C. Sousa, "3D Geological Modeling Using Sketches and Annotations from Geologic Maps," in *4th Joint Symposium on Computational Aesthetics, Non-Photorealistic Animation and Rendering, and Sketch-Based Interfaces and Modeling*, 2014.
- [11] A. Rocha, U. Alim, J. Silva and M. Sousa, "Decal-maps: Real-Time Layering of Decals on Surfaces for Multivariate Visualization," *IEEE Transactions on Visualization and Computer Graphics*, 2016.
- [12] G. de Jager, J. F. M. Van Doren, J. D. Jansen and S. M. Luthi, "An Evaluation of Relevant Geological Parameters for Predicting the Flow Behaviour of Channelized Reservoirs," *Petroleum Geoscience*, vol. 15, no. 4, pp. 345-354, 2009.
- [13] "Birdly," SOMINIACS, [Online]. Available: <http://www.somniacs.co>. [Accessed 20 September 2016].
- [14] P. Milgram and F. Kishino, "A Taxonomy of Mixed Reality Visual Displays," *IEICE Transactions on Information Systems*, Vols. E77-D, no. 12, p. 1321–1329, 1994.
- [15] J. Jerold, *The VR Book: Human-Centered Design for Virtual Reality*, ACM Books, 2016.
- [16] D. A. F. S. T. S. i. V. Reality, "Pausch, Randy; Snoddy, Jon; Taylor, Robert; Watson, Scott; Haseltine, Eric," in *International Conference on Computer Graphics and Interactive Techniques*, New York, 96.
- [17] "MOGA Pro Controller," MOGA, [Online]. Available: <http://www.mogaanywhere.com/controllers/moga-pro-controller/>. [Accessed 21 September 2016].
- [18] "Touch by Oculus," Oculus, [Online]. Available: <https://www3.oculus.com/en-us/touch/>. [Accessed 21 September 2016].

- [19] "PlayStation Move," PlayStation, [Online]. Available: <https://www.playstation.com/en-us/explore/accessories/playstation-move/>. [Accessed 23 September 2016].
- [20] "Rapture Tech," THE VOID, [Online]. Available: <https://thevoid.com/Tech#tech>. [Accessed 23 September 2016].
- [21] "PPT Wand," WorldViz, [Online]. Available: <http://www.est-kl.com/es/products/motion-tracking/worldviz/ppt-wand.html>. [Accessed 23 September 2016].
- [22] "Leap Motion | 3D Motion and Gesture Control for Virtual Reality," Leap Motion, [Online]. Available: <https://www.leapmotion.com/product/vr>. [Accessed 21 September 2016].
- [23] "Myo Gesture Control Armband," Thalmic Labs, [Online]. Available: <https://www.myo.com/>. [Accessed 21 September 2016].
- [24] "CyberGlove III," CyberGlove Systems, [Online]. Available: <http://www.cyberglovesystems.com/cyberglove-iii/>. [Accessed 21 September 2016].
- [25] "CyberTouch II," CyberGlove Systems, [Online]. Available: <http://www.cyberglovesystems.com/cybertouch2/>. [Accessed 21 September 2016].
- [26] "CyberGrasp," CyberGlove Systems, [Online]. Available: <http://www.cyberglovesystems.com/cybergrasp>. [Accessed 21 September 2016].
- [27] "CyberForce," CyberGlove Systems, [Online]. Available: <http://www.cyberglovesystems.com/cyberforce>. [Accessed 21 September 2016].
- [28] "Oculus Rift," Oculus, [Online]. Available: <https://www3.oculus.com/en-us/rift/>. [Accessed 21 September 2016].

- [29] "FOVE | Eye Tracking VR Headset," FOVE, [Online]. Available: <http://www.getfove.com/>. [Accessed 21 September 2016].
- [30] "Cyberith Virtualizer," Cyberith , [Online]. Available: <http://cyberith.com/product/>. [Accessed 21 September 2016].
- [31] "MaximumVR with Tracked Feet - Youtube," AlienTrap, 19 April 2016. [Online]. Available: <https://www.youtube.com/watch?v=IdQDa-IMo7I>. [Accessed 21 September 2016].
- [32] "PriorVR: Suit up. Game on.," YEI Technology, [Online]. Available: www.priovr.com. [Accessed 21 September 2016].
- [33] "Kinect | Xbox 360," Microsoft, [Online]. Available: <http://www.xbox.com/en-US/xbox-360/accessories/kinect>. [Accessed 21 September 2016].
- [34] "Microsoft HoloLens | Official Site," Microsoft, [Online]. Available: <https://www.microsoft.com/microsoft-hololens/en-us>. [Accessed 17 November 2016].
- [35] "Google Glass," Google, [Online]. Available: <https://www.google.com/glass/start/>. [Accessed 17 November 2016].
- [36] "iOS - Siri - Apple," Apple, [Online]. Available: www.apple.com/ios/siri/. [Accessed 17 November 2016].
- [37] "Google Now Launcher - Android Apps on Google Play," Google, [Online]. Available: <https://play.google.com/store/apps/details?id=com.google.android.launcher>. [Accessed 17 November 2016].
- [38] W. R. Sherman and A. B. Craig, *Understanding Virtual Reality: Interface, Application, and Design*, Morgan Kaufmann, 2002.

- [39] "VIRTALIS: VR, Virtual Reality," Virtualis, [Online]. Available: <https://www.virtalis.com/>. [Accessed 15 November 2016].
- [40] "VIVE™ | Discover Virtual Reality Beyond Imagination," HTC Corporation, [Online]. Available: <https://www.vive.com>. [Accessed 15 November 2016].
- [41] "Samsung Gear VR - The Official Samsung Galaxy Site," Samsung, [Online]. Available: <http://www.samsung.com/global/galaxy/gear-vr/>. [Accessed 15 November 2016].
- [42] "Google Cardboard – Google VR," Google, [Online]. Available: <https://vr.google.com/cardboard/>. [Accessed 15 November 2016].
- [43] "M100 Smart Glasses," Vuzix, [Online]. Available: <https://www.vuzix.com/Products/m100-smart-glasses>. [Accessed 15 November 2016].
- [44] "NVisor ST50," VIRTALIS, [Online]. Available: <https://www.virtalis.com/blogs/products/nvis-2/>.
- [45] "Visette45," VIRTALIS, [Online]. Available: <https://www.virtalis.com/blogs/products/cybermind/>. [Accessed 15 November 2016].
- [46] "Spatial Audio | Google VR | Google Developers," Google, [Online]. Available: <https://developers.google.com/vr/concepts/spatial-audio>. [Accessed 15 November 2016].
- [47] "ENVELOP - 3D Sound," ENVELOP, [Online]. Available: <http://www.envelop.us/>. [Accessed 15 November 2016].
- [48] M. V. Sanchez-Vives and M. Slater, "From Presence to Consciousness through Virtual Reality," *Nature reviews Neuroscience*, vol. 6, no. 4, pp. 332-339, 2005.

- [49] "Gloveone: Feel Virtual Reality by NeuroDigital Technologies," NeuroDigital Technologies, [Online]. Available: <https://www.kickstarter.com/projects/gloveone/gloveone-feel-virtual-reality>. [Accessed 15 November 2016].
- [50] "Dexmo," Dexta Robotics, [Online]. Available: <http://www.dextarobotics.com/>. [Accessed 15 November 2016].
- [51] "NullSpace VR | Bringing sense to the virtual world," NullSpace VR, [Online]. Available: <http://nullspacevr.com/>. [Accessed 15 November 2016].
- [52] "KOR-FX 4DFX Haptic Gaming Vest," Immerz, [Online]. Available: <http://www.korfx.com/>. [Accessed 15 November 2016].
- [53] "Widerun The First Full Immersive Virtual Reality Fitness Experience," Widerun, [Online]. Available: <http://www.widerun.com/>. [Accessed 15 November 2016].
- [54] M. Slater, "A Note on Presence Terminology," *Presence Connect*, vol. 3, no. 3, 2003.
- [55] C. Botella, R. M. Baños, C. Perpiñá, H. Villa, M. Alcañiz and A. Rey, "Virtual Reality Treatment of Claustrophobia: A Case Report," *Behaviour Research and Therapy*, vol. 36, no. 2, pp. 239-246, 1998.
- [56] B. O. Rothbaum, L. Hodges, B. A. Watson, D. Kessler and D. Opdyke, "Virtual Reality Exposure Therapy in the Treatment of Fear of Flying: A Case Report," *Behaviour Research and Therapy*, vol. 34, no. 5-6, pp. 477-481, 1996.
- [57] A. S. Carlin, H. G. Hoffman and S. Weghorst, "Virtual reality and Tactile Augmentation in the Treatment of Spider Phobia: A Case Report," *Behaviour Research and Therapy*, vol. 35, no. 2, pp. 153-158, 1997.
- [58] R. M. Satava, "Virtual Reality and Telepresence for Military Medicine," *Computers in Biology and Medicine*, vol. 25, no. 2, pp. 229-236, 1995.

- [59] G. M. Reger, K. M. Holloway, C. Candy, B. O. Rothbaum, J. Difede, A. A. Rizzo and G. A. Gahm, "Effectiveness of Virtual Reality Exposure Therapy for Active Duty Soldiers in a Military Mental Health Clinic," *Journal of Traumatic Stress*, vol. 24, no. 1, pp. 93-96, 2011.
- [60] R. Pausch, T. Crea and M. Conway, "A Literature Survey for Virtual Environments: Military Flight Simulator Visual Systems and Simulator Sickness," *Presence: Teleoperators and Virtual Environments*, vol. 1, no. 3, pp. 344-363, 1992.
- [61] C. E. Hughes, C. B. Stapleton, D. E. Hughes and E. M. Smith, "Mixed Reality in Education, Entertainment, and Training," *IEEE Computer Graphics and Applications*, vol. 25, no. 6, pp. 24-30, 2005.
- [62] D. A. Bowman and R. P. McMahan, "Virtual Reality: How Much Immersion Is Enough?," *Computer*, vol. 40, no. 7, pp. 36-43, 2007.
- [63] E. D. Ragan, R. Kopper, P. Schuchardt and D. A. Bowman, "Studying the Effects of Stereo, Head Tracking, and Field of Regard on a Small-Scale Spatial Judgment Task," *IEEE Transactions on Visualization and Computer Graphics*, vol. 19, no. 5, pp. 886-896, 2013.
- [64] Prabhat, A. Forsberg, M. Katzourin, K. Wharton and M. Slater, "A Comparative Study of Desktop, Fishtank, and Cave Systems for the Exploration of Volume Rendered Confocal Data Sets," *IEEE Transactions on Visualization and Computer Graphics*, vol. 14, no. 3, pp. 551-563, 2008.
- [65] C. Ware and G. Franck, "Evaluating Stereo and Motion Cues for Visualizing Information Nets in Three Dimensions," *ACM Transactions on Graphics*, vol. 15, no. 2, pp. 121-140, 1996.
- [66] C. Ware, K. Arthur and K. S. Booth, "Fish Tank Virtual Reality," in *Conference on Human Factors in Computing Systems*, Amsterdam, The Netherlands, 1993.

- [67] J. Tibbett, F. T. Suorineni and B. Hebblewhite, "The Use of Virtual Reality Scientific Visualization for Investigation and Exploration of Block Cave Mining System Data," in *The SAIMM Virtual Reality and Spatial Information Applications in the Mining Industry Conference*, 2015.
- [68] "AVIE - Project Overview," iCinema - Center for Interactive Cinema Research, [Online]. Available: <http://www.icinema.unsw.edu.au/technologies/avie/>. [Accessed 14 November 2016].
- [69] L. Bilke, T. Fischer, C. Helbig, C. Krawczyk, T. Nagel, D. Naumov, S. Paulick, K. Rink, A. Sachse, S. Schelenz, M. Walther, N. Watanabe, B. Zehner, J. Ziesch and O. Kolditz, "TESSIN VISLab — Laboratory for Scientific Visualization," *Environmental Earth Sciences*, vol. 72, no. 10, pp. 3881-3899, 2014.
- [70] C. Helbig, H.-S. Bauer, K. Rink, V. Wulfmeyer, M. Frank and O. Kolditz, "Concept and Workflow for 3D Visualization of Atmospheric Data in a Virtual Reality Environment for Analytical Approaches," *Environmental Earth Sciences*, vol. 72, no. 10, pp. 3767-378, 2014.
- [71] D. Wigdor and D. Wixon, *Brave NUI World: Designing Natural User Interfaces for Touch and Gesture*, Morgan Kaufmann, 2011.
- [72] S. Schkolne, "3D Interfaces for Spatial Construction," 23 October 2003. [Online]. Available: <http://thesis.library.caltech.edu/2127/>. [Accessed 14 November 2016].
- [73] J. S. Gomes and F. B. Alves, *The Universe of Oil and Gas Industry: from Exploration to Refining*, Lisbon: Fundação Calouste Gulbenkian, 2007.
- [74] "Petrel E&P Software Platform," Schlumberger, [Online]. Available: <https://www.software.slb.com/products/petrel>. [Accessed 28 September 2016].
- [75] "BUILDER - Simulation Model Building Application," CMG - Computer Modelling Group LTD., [Online]. Available: <http://www.cmgl.ca/software/builder2016>. [Accessed 28 September 2016].

- [76] "Irap RMS," Roxar, [Online]. Available: <http://support.roxar.com/iraprms>. [Accessed 28 September 2016].
- [77] M. Mlacnik, L. J. Durlofsky and Z. E. Heinemann, "Sequentially Adapted Flow-based PEBI Grids for Reservoir Simulation," *One Petro*, vol. 11, no. 3, pp. 317-327, 2006.
- [78] A. G. Journé and C. J. Huijbregts, *Mining Geostatistics*, New York: Academic Press, 1978.
- [79] H. H. Haldorsen and E. Damsleth, "Stochastic Modelling," *Journal of Petroleum Technology*, 1990.
- [80] A. C. MacDonald and J. O. Aasen, "A Prototype Procedure for Stochastic Modeling of Facies Tract Distribution in Shoreface Reservoirs," in *CA 3: Stochastic Modeling and Geostatistics*, The American Association of Petroleum Geologists, 1994.
- [81] S. Strebelle, "Conditional Simulation of Complex Geological Structures Using Multiple-Point Statistics," *Mathematical Geology*, vol. 34, no. 1, pp. 1-21, 2002.
- [82] C. V. Deutsch and A. G. Journé, *GSLIB: Geostatistical Software Library and User's Guide*, 1998.
- [83] L. Mitas and H. Mitasova, "Spatial interpolation," in *Geographical Information Systems: Principles and Technical Issues*, pp. 481-492.
- [84] E. Gringarten, "Uncertainty Assessment in 3D Reservoir Modelling," Canadian Society of Petroleum Geologists, 2006.
- [85] C. V. Deutsch and S. Srinivasan, "Improved Reservoir Management through Ranking Stochastic Reservoir Models," in *SPE/DOE Symposium on Improved Oil Recovery*, Tulsa, 1996.
- [86] M. A. Christie, "Upscaling for Reservoir Simulation," *Journal of Petroleum Technology*, vol. 48, no. 11, pp. 1,004 - 1,010, 1996.

- [87] P. Audigane and M. J. Blunt, "Dual Mesh Method for Upscaling in Waterflood Simulation," *Transport in Porous Media*, vol. 55, no. 1, pp. 71-89, 2004.
- [88] M. A. Christie and P. J. Clifford, "Fast Procedure for Upscaling Compositional Simulation," *SPE Journal*, vol. 3, no. 03, pp. 272 - 278, 1998.
- [89] F. Marcondes, A. Varavei and K. Sepehrnoori, "An EOS-based Numerical Simulation of Thermal Recovery Process Using Unstructured Meshes," *Brazilian Journal of Chemical Engineering*, vol. 32, no. 1, pp. 247-258, 2015.
- [90] "GPUs Speed Reservoir Simulation with Cray CS-Storm Supercomputer," Inside HPC, [Online]. Available: <http://insidehpc.com/2015/10/gpus-speed-reservoir-simulation-with-cray-cs-storm-supercomputer/>. [Accessed 12 October 2016].
- [91] "Supercomputing in Reservoir Simulation," RFD - Rock Flow Dynamics, [Online]. Available: <http://www.rfdyn.com/technology/supercomputing/>. [Accessed 10 October 2016].
- [92] "Univa Corporation Provides Integrated Technology," Univa, [Online]. Available: http://www.univa.com/about/news/press_2014/06192014.php. [Accessed 12 October 2016].
- [93] "GPU Computing," Stone Ridge Technology, [Online]. Available: <http://stoneridgetechnology.com/gpu-computing/>. [Accessed 12 October 2016].
- [94] A. Datta-Gupta and M. J. King, *Streamline Simulation: Theory and Practice*, Society of Petroleum Engineers, 2007.
- [95] J. J. Arps, "Analysis of Decline Curves," *Transactions of the AIME*, vol. 160, no. 1, pp. 228-247, 1945.

- [96] R. O. Baker, T. R. Anderson and K. Sandhu, "Using Decline Curves to Forecast Waterflooded Reservoirs: Fundamentals and Field Cases," in *Canadian International Petroleum Conference*, Calgary, Alberta, 2003.
- [97] M. Delshad, A. Bastami and P. Pourafshary, "The Use of Capacitance-Resistive Model for Estimation of Fracture Distribution in the Hydrocarbon Reservoir," in *SPE Saudi Arabia Section Technical Symposium*, 2009.
- [98] A. Mamghaderi, A. Bastami and P. Pourafshary, "Optimization of Waterflooding Performance in a Layered Reservoir Using a Combination of Capacitance-Resistive Model and Genetic Algorithm Method," *Journal of Energy Resources Technology*, vol. 135, no. 1, 2012.
- [99] M. Sayarpour, E. Zuluaga, C. S. Kabir and L. W. Lake, "The Use of Capacitance-resistance Models for Rapid Estimation of Waterflood Performance and Optimization," *Journal of Petroleum Science and Engineering*, vol. 69, no. 3-4, p. 227-238, 2009.
- [100] M. Sayyafzadeh, P. Pourafshary, M. Haghghi and F. Rashidi, "Application of Transfer Functions to Model Water Injection in Hydrocarbon Reservoir," *Journal of Petroleum Science and Engineering*, vol. 78, no. 1, pp. 139-148, 2011.
- [101] D. K. Larue and F. Friedman, "The Controversy Concerning Stratigraphic Architecture of Channelized Reservoirs and Recovery by Waterflooding," *Petroleum Geoscience*, vol. 11, pp. 131-146, 2005.
- [102] J. M. Hovadik and D. K. Larue, "Stratigraphic and Structural Connectivity," *Geological Society of London*, vol. 347, pp. 219-242, 2010.
- [103] J. M. Hovadik and D. K. Larue, "Static Characterizations of Reservoirs: Refining the Concepts of Connectivity and Continuity," *Petroleum Geoscience*, vol. 13, pp. 195-211, 2007.

- [104] C. V. Deutsch, "Fortran Programs for Calculating Connectivity of Three-dimensional Numerical Models and Ranking Multiple Realizations," *Computers & Geosciences*, vol. 24, no. 1, p. 69–76, 1999.
- [105] B. Cappelaere, J. Touma and C. Peugeot, "A Recursive Algorithm for Connectivity Analysis in a Grid; Application to 2D Hydrodynamic Modeling in Heterogeneous Soils," *Computers & Geosciences*, vol. 26, no. 2, p. 121–135, 2000.
- [106] E. Pardo-Igúzquiza and P. A. Dowd, "CONNEC3D: A Computer Program for Connectivity Analysis of 3D Random Set Models," *Computers & Geosciences*, vol. 29, no. 6, p. 775–785, 2003.
- [107] J. Wu, "Reservoir Geobody Calculation," 2012. [Online]. Available: <https://www.google.com/patents/US20120150501>. [Accessed 12 October 2016].
- [108] T.-S. Lo and J. Chu, "Hydrocarbon Reservoir Connectivity Tool Using Cells and Pay Indicators," 1995. [Online]. Available: <https://www.google.ch/patents/US5757663>. [Accessed 12 October 2016].
- [109] F. M. Petit, P. Y. Biver, P. M. Calatayud, J.-L. Lesueur and F. G. Alabert, "Early Quantification of Hydrocarbon in Place Through Geostatistical Object Modelling and Connectivity Computations," in *SPE Annual Technical Conference and Exhibition*, New Orleans, Louisiana, 1994.
- [110] J. R. L. Allen, "Studies in Fluvial Sedimentation: An Exploratory Quantitative Model for the Architecture of Avulsion-controlled Alluvial Suites," *Sedimentary Geology*, vol. 21, no. 02, pp. 129-147, 1978.
- [111] J. R. L. Allen, "Studies in Fluvial Sedimentation: An Elementary Geometrical Model for the Connectedness of Avulsion-related Channel Sand Bodies," *Sedimentary Geology*, vol. 24, no. 3-4, pp. 253-267, 1979.
- [112] D. K. Larue and J. Hovadik, "Connectivity of Channelized Reservoirs: a Modelling Approach," *Petroleum Geoscience*, vol. 12, pp. 291-308, 2006.

- [113] P. R. King, "The Connectivity and Conductivity of Overlapping Sand Bodies," in *North Sea Oil and Gas Reservoirs—II*, Springer Netherlands, 1990, pp. 353-362.
- [114] P. R. King, S. V. Buldyrev, N. V. Dokholyan, S. Havlin, Y. Lee, G. Paul and H. E. Stanley, "Applications of Statistical Physics to the Oil Industry: Predicting Oil Recovery Using Percolation Theory," *Physica A: Statistical Mechanics and its Applications*, vol. 274, no. 1-2, pp. 60-66, 1999.
- [115] D. Allard and H. Group, "On the Connectivity of Two Random Set Models: The Truncated Gaussian and the Boolean," in *Geostatistics Tróia '92*, Springer Netherlands, 1993, pp. 467-478.
- [116] S. Sadeghnejada, M. Masihi, M. Pishvaie, A. Shojaei and P. R. King, "Utilization of Percolation Approach to Evaluate Reservoir Connectivity and Effective Permeability: A Case Study on North Pars Gas Field," *Scientia Iranica*, vol. 18, no. 6, p. 1391–1396, 2011.
- [117] M. Masihi, P. R. King and P. R. Nurafza, "Fast Estimation of Connectivity in Fractured Reservoirs Using Percolation Theory," *SPE Journal*, vol. 12, no. 02, pp. 167-178, 2007.
- [118] P. R. King, S. V. Buldyrev, N. V. Dokholyan, S. Havlin, Y. Lee, G. Paul, H. E. Stanley and N. Vandesteeg, "Predicting Oil Recovery Using Percolation Theory," *Petroleum Geoscience*, vol. 7, pp. 105-107, 2001.
- [119] O. Izgeca and C. S. Kabirb, "Quantifying Reservoir Connectivity, In-place Volumes, and Drainage-area Pressures During Primary Depletion," *Journal of Petroleum Science and Engineering*, vol. 81, p. 7–17, 2012.
- [120] J. Hovadik and D. Larue, "Predicting Waterflood Behavior by Simulating Earth Models with No or Limited Dynamic Data: From Model Ranking to Simulating a Billion-Cell Model," 2011. [Online]. Available: http://archives.datapages.com/data/specpubs/memoir96/m96_html/CHAPTER03/CHAPTER03.HTM. [Accessed 12 October 2016].

- [121] U. S. Allan, "Model for Hydrocarbon Migration and Entrapment Within Faulted Structures," 1989. [Online]. Available: <http://archives.datapages.com/data/bulletns/1988-89/data/pg/0073/0007/0800/0803.htm>. [Accessed 13 October 2016].
- [122] F. Neele, J. t. Veen, F. Wilschut and C. Hofstee, "Independent Assessment of High-capacity Offshore CO2 Storage Options," 2012. [Online]. Available: <https://www.globalccsinstitute.com/publications/independent-assessment-high-capacity-offshore-co2-storage-options>. [Accessed 13 October 2016].
- [123] W. R. Bailey, T. Manzocchi, J. J. Walsh, K. Keogh, D. Hodgetts, J. Rippon, P. A. R. Nell, S. Flint and J. A. Strand, "The Effect of Faults on the 3D Connectivity of Reservoir Bodies: A Case Study from the East Pennine Coalfield, UK," *Petroleum Geoscience*, vol. 8, pp. 263-277, 2002.
- [124] P. A. Gutteridge and D. E. Gawith, "Connected Volume Calibration For Well-Path Ranking," in *European 3-D Reservoir Modelling Conference*, Stavanger, Norway, 1996.
- [125] S. Li, C. V. Deutsch and J. Si, "Ranking Geostatistical Reservoir Models with Modified Connected Hydrocarbon Volume," in *9th International Geostatistics Congress*, Oslo, Norway, 2012.
- [126] D. R. Fenik, A. Nouri and C. V. Deutsch, "Ranking Realizations for SAGD Performance Predictions," Centre for Computational Geostatistics (CCG), 2009.
- [127] J. A. Sethian, "A Fast Marching Level Set Method for Monotonically Advancing Fronts," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 93, no. 4, pp. 1591-1595, 1996.
- [128] M. Sharifi, M. Kelkar, A. Bahar and T. Slettebo, "Dynamic Ranking of Multiple Realizations by Use of Fast-Marching Method," *SPE Journal*, vol. 19, no. 06, pp. 1069-1082, 2014.

- [129] K. H. Karlsen, K. Lie and N. Risebro, "A Fast Marching Method for Reservoir Simulation," *Computational Geosciences*, vol. 4, no. 2, pp. 185-206, 2000.
- [130] I. Berre, H. K. Dahle, K. H. Karlsen, K. A. Lie and J. R. Natvig, "Time-of-Flight + Fast Marching + Transport Collapse: An Alternative to Streamlines for Two-Phase Porous Media Flow with Capillary Forces," in *Proceedings of Computational Methods in Water Resources XIV*, 2002.
- [131] I. Berre, K. H. Karlsen, K.-A. Lie and J. R. Natvig, "Fast Computation of Arrival Times in Heterogeneous Media," *Computational Geosciences*, vol. 9, no. 4, p. 179–201, 2005.
- [132] K. B. Hird and O. Dubrule, "Quantification of Reservoir Connectivity for Reservoir Description Applications," *SPE Reservoir Evaluation & Engineering*, vol. 1, no. 01, pp. 12-17, 1998.
- [133] P. R. Ballin, R. Solano, K. B. Hird and R. F. Volz, "New Reservoir Dynamic Connectivity Measurement for Efficient Well Placement Strategy Analysis Under Depletion," in *SPE Annual Technical Conference and Exhibition*, San Antonio, Texas, 2002.
- [134] P. H. Sammon, "Calculation of Convective and Dispersive Flows for Complex Corner Point Grids," in *SPE Annual Technical Conference and Exhibition*, Dallas, Texas, 2000.
- [135] E. W. Dijkstra, "A Note on Two Problems in Connexion with Graphs," *Numerische Mathematik*, vol. 1, no. 1, pp. 269-271, 1959.
- [136] M. Fustic, B. Bennett, J. Adams, H. Huang, B. MacFarlane, D. Leckie and S. R. Larter, "Bitumen and Heavy Oil Geochemistry: A Tool for Distinguishing Barriers from Baffles in Oil Sands Reservoirs," *Bulletin of Canadian Petroleum Geology*, vol. 59, no. 4, pp. 295-316, 2012.
- [137] M. J. Peacock, "Athabasca Oil Sands: Reservoir Characterization and its Impact on Thermal and Mining Opportunities," in *Petroleum Geology Conference Series*, 2010.

- [138] S. G. R. Devriese and D. W. Oldenburg, "Feasibility of Electromagnetic Methods to Detect and Image Steam-assisted Gravity Drainage Steam Chambers," *Geophysics*, vol. 81, no. 4, pp. E227-E241, 2016.
- [139] L. D. Cohen and R. Kimmel, "Global Minimum for Active Contour Models: A Minimal Path Approach," *International Journal of Computer Vision*, vol. 24, no. 1, p. 57–78, 1997.
- [140] J. A. Sethian, *Level Set Methods and Fast Marching Methods - Evolving Interfaces in Computational Geometry, Fluid Mechanics, Computer Vision, and Materials Science*, Cambridge Monograph on Applied and Computational Mathematics, 1999.
- [141] R. Kimmel and J. A. Sethian, "Computing Geodesic Paths on Manifolds," *Applied Mathematics*, vol. 95, no. 15, p. 8431–8435, 1998.
- [142] C. Pêtrès, Y. Pailhas, Y. Petillot and D. Lane, "Underwater Path Planning Using Fast Marching Algorithms," in *Europe Oceans*, 2005.
- [143] J. R. Shewchuk, "What is a Good Linear Finite Element? Interpolation, Conditioning, Anisotropy, and Quality Measures," University of California at Berkeley, California, 2002.
- [144] D. W. Pollock, "Semianalytical Computation of Path Lines for Finite-Difference Models," *Ground Water*, vol. 26, no. 6, p. 743–750, 1988.
- [145] M. Prévost, M. G. Edwards and M. J. Blunt, "Streamline Tracing on Curvilinear Structured and Unstructured Grids," *SPE Journal*, vol. 7, no. 2, pp. 139-148, 2002.
- [146] M. Novotni and R. Klein, "Computing Geodesic Distances on Triangular Meshes," in *International Conference in Central Europe on Computer Graphics, Visualization and Computer Vision*, 2002.

- [147] B. Zehner, "Mixing Virtual Reality and 2D Visualization - Using Virtual Environments as Visual 3D Information Systems for Discussion of Data from Geo- and Environmental Sciences," in *Computer Graphics Theory and Applications, GRAPP*, 2010.
- [148] A. Raposo, L. Soares, G. Wagner, E. Corseuil, M. Gattass and I. Santos, "Environ: Integrating VR and CAD in Engineering Projects," *IEEE Computer Graphics and Applications*, vol. 29, no. 6, pp. 91-95, 2009.
- [149] I. H. F. dos Santos, A. Raposo and L. P. Soares, "EnViron: An Integrated VR Tool for Engineering Projects," in *Computer Supported Cooperative Work in Design, CSCWD*, 2008.
- [150] "KOC installs most advanced MegaCADWall in its industry," Barco, 2012. [Online]. Available: <http://az877327.vo.msecnd.net/~/.media/Downloads/Customer%20stories/2012/Case%20story%20pdf%20KOC%20Kuwait.pdf?v=1>. [Accessed 8 October 2016].
- [151] "Case story OGX Petróleo e Gás, Brazil," Barco, 2012. [Online]. Available: <http://www.barco.com/en/Data/References/2012/Q3/~/.media/DF05978B395048A694B1C5D58D6E5B53.pdf>. [Accessed 8 October 2016].
- [152] "Shell iScope," [Online]. Available: <http://gallery.tudelft.nl/index.php/start/album472/ta/spe-delft/2014/Visits/Shell-iScope>. [Accessed 8 October 2016].
- [153] T. Seyed, M. C. Sousa, F. Maurer and A. Tang, "SkyHunter: A Multi-Surface Environment for Supporting Oil and Gas Exploration," in *Interactive Tabletops and Surfaces, ITS*, 2013.
- [154] C. Burns, T. Seyed, K. Bradley, R. Duncan, A. Balasch, F. Maurer and M. Sousa, "Multi-Surface Visualization of Fused Hydrocarbon Microseep and Reservoir Data," *GeoConvention*, 2012.

- [155] P. Lapedes, N. Sultanum, E. Sharlin and M. C. Sousa, "Seamless Mixed Reality Tracking in Tabletop Reservoir Engineering Interaction," in *Advanced Visual Interfaces, AVI*, 2012.
- [156] J. Owens, "Effective Petrophysics models leading to Business Decisions," in *73rd EAGE Conference and Exhibition*, 2011 .
- [157] N. Sultanum, S. Somanath, E. Sharlin and M. C. Sousa, "'Point it, Split it, Peel it, View it': Techniques for Interactive Reservoir Visualization on Tabletops," in *Interactive Surfaces and Spaces (ISS)*, Kobe, Japan, 2011.
- [158] N. Sultanum, E. V. Brazil and M. C. Sousa, "Navigating and Annotating 3D Geological Outcrops through Multi-touch Interaction," in *Interactive Surfaces and Spaces (ISS)*, St. Andrews, Scotland, 2013 .
- [159] Y. Hajizadeh, R. Amorim, N. Boroumand, E. Brazil, D. Eaton and M. C. Sousa, "Interactive Sketch-based Estimation of Stimulated Volume in Unconventional Reservoirs Using Microseismic Data," in *13th European Conference on the Mathematics of Oil Recovery* , 2012.
- [160] M. Schlattmann, Y. Yu, N. Gruendl, M. Bogen, A. Kulik, D. d'Angelo, B. Froehlich and R. Klein, "User Awareness for Collaborative Multi-touch Interaction," in *8th Conference on Computer Graphics Theory and Applications (GRAPP)*, Barcelona, Spain, 2013.
- [161] D. d'Angelo, A. Kulik, M. Schlattmann and M. Bogen, "Enabling Context Aware Multi-User Interaction with the Microsoft Kinect," VRGeo.
- [162] "The VRGeo Slice Inspector," VRGeo, 12 July 2010. [Online]. Available: <https://www.youtube.com/watch?v=iMfRyGzntPo>. [Accessed 11 November 2016].
- [163] "SeisTablet," VRGeo, [Online]. Available: <http://www.vrgeo.org/index.php?id=598>. [Accessed 11 November 2016].

- [164] J. Guerreiro, D. Medeiros, D. Mendes, M. Sousa, J. Jorge, A. Raposo and I. Santos, "Beyond Post-It: Structured Multimedia Annotations for Collaborative VEs," in *The European Association for Computer Graphics (EUROGRAPHICS)*, Bremen, Germany, 2014.
- [165] I. H. F. dos Santos, L. P. Soares, F. Carvalho and A. Raposo, "A Collaborative VR Visualization Environment for Offshore Engineering Projects," in *International Conference on Computer Graphics and Interactive Techniques (GRAPH)*, Hong Kong, China, 2011 .
- [166] E. R. Silva and A. B. Raposo, "Colaboração em Ambientes Heterogêneos de Realidade Virtual para Aplicações de Treinamento," in *X Brazilian Symposium in Collaborative Systems*, Manaus, Amazonas, 2013 .
- [167] "Inside Reality Home Page," Schlumberger, [Online]. Available: <http://www.oilit.com/2journal/2article/0007/0005.htm>. [Accessed 11 November 2016].
- [168] O. Kreylos, G. Bawden, T. Bernardin, M. I. Billen, E. S. Cowgill and R. D. Gold, "Enabling Scientific Workflows in Virtual Reality," in *International Conference on Computer Graphics and Interactive Techniques (GRAPH)*, Hong Kong, China, 2006.
- [169] "Virtual Reality Geological Studio (VRGS)," VRGeoscience, [Online]. Available: <http://www.vrgeoscience.com/>. [Accessed 11 November 2016].
- [170] "GeoVisionary," British Geological Survey (BGS) and Virtualis Ltd, [Online]. Available: <http://geovisionary.com/>. [Accessed 11 November 2016].
- [171] K. Gruchalla, "ImmersiveWell-Path Editing: Investigating the Added Value of Immersion," in *IEEE Virtual Reality*, Chicago, 2004.

- [172] H. Ishii and B. Ullmer, "Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms," in *Conference on Human Factors in Computing Systems*, Atlanta, Georgia, 1997.
- [173] J. Harris, J. Young, N. Sultanum, P. Lapedes, E. Sharlin and M. C. Sousa, "Designing Snakey: A Tangible User Interface Supporting Well Path Planning," in *Human-Computer Interaction – INTERACT*, Lisbon, Portugal, 2011.
- [174] S. Somanath, A. Rocha, H. Hamdi, E. Sharlin and M. C. Sousa, "ReservoirBench: An Interactive Educational Reservoir Engineering Workbench.," in *Human-Computer Interaction - INTERACT*, Bamberg, Germany, 2015.
- [175] "Tangible Landscape (formerly TanGeoMS)," NCSU GeoForAll Lab, [Online]. Available: <https://geospatial.ncsu.edu/osgeorel/tangible-landscape.html>. [Accessed 13 November 2016].
- [176] L. G. Tateosian, H. Mitasova, B. Foglemann, B. Harmon, W. K and R. S. Harmon, "TanGeoMS: Tangible Geospatial Modeling System," *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, vol. 16, no. 6, pp. 1605-1612, 2010.
- [177] N. Couture, G. Rivière and P. Reuter, "GeoTUI: A Tangible User Interface for Geoscience," in *Tangible and Embedded Interaction*, Bonn, Germany, 2008.
- [178] A. S. Nittala, N. Li, S. Cartwright, K. Takashima, E. Sharlin and M. C. Sousa, "PLANWELL: Spatial User Interface for Collaborative Petroleum Well-Planning," in *International Conference on Computer Graphics and Interactive Techniques*, Kobe, Japan, 2015.
- [179] N. Li, S. Cartwright, A. S. Nittala, E. Sharlin and M. C. Sousa, "A Spatial Interface for Enhancing Human-UAV Awareness," in *International Conference on Human-Agent Interaction*, Daegu, Kyungpook, Republic of Korea, 2015.

- [180] C. Harding, B. Loftin and A. Anderson, "Visualization and Modeling of Geoscientific Data on the Interactive Workbench," *The Leading Edge*, vol. 19, no. 5, pp. 506-511, 2000.
- [181] D. A. Bowman and L. F. Hodges, "An Evaluation of Techniques for Grabbing and Manipulating Remote," in *Symposium on Interactive 3D Graphics and Games*, Providence, Rhode Island, USA, 1997.
- [182] B. Fröhlich, J. Plate, J. Wind, G. Wesche and M. Göbel, "Cubic-Mouse-Based Interaction in Virtual Environments," *IEEE Computer Graphics and Applications*, vol. 20, no. 4, pp. 12 - 15, 2000.
- [183] S. Seele, "Gesture Control in Virtual Environments," 2008. [Online]. Available: https://www.researchgate.net/publication/260401681_Gesture_Control_in_Virtual_Environments. [Accessed 13 November 2016].
- [184] J. Fiedler, S. Rilling, M. Bogen and J. Herder, "Multimodal Interaction Techniques in Scientific Data Visualization: An Analytical Survey," in *International Conference on Computer Graphics Theory and Applications*, Berlin, Germany, 2015.
- [185] "AVEVA Activity Visualisation Platform," AVEVA, [Online]. Available: http://www.aveva.com/aveva_avp.. [Accessed 15 November 2016].
- [186] A. Shukla and H. Karki, "Application of Robotics in Offshore Oil and Gas Industry- A Review Part II," *Robotics and Autonomous Systems*, vol. 75, no. PB, pp. 508-524, 2016.
- [187] I. H. F. Santos, G. M. Ribeiro, F. Coutinho, L. Hsu, A. Raposo, F. Carvalho, D. Medeiros, M. Galassi, R. Costa, P. Arroyo, G. M. Freitas, T. A. Antonio and F. Lizarralde, "A Robotics Framework for Planning the Offshore Robotizing Using Virtual Reality Techniques," in *Offshore Technology Conference (OTC)*, Rio de Janeiro, Brazil, 2013.
- [188] "Petrobras," [Online]. Available: <http://www.petrobras.com.br/pt/>. [Accessed 15 November 2016].

- [189] A. E. Mostafa, K. Takashima, M. C. Sousa and E. Sharlin, "JackVR: A Virtual Reality Training System for Landing Oil Rigs," in *International Symposium on Visual Computing*, Las Vegas, NV, USA, 2015.
- [190] J. Rickel and W. L. Johnson, "Extending Virtual Humans to Support Team Training in Virtual Reality," in *Exploring Artificial Intelligence in the New Millennium*, Morgan Kaufmann, 2003, pp. 217-238.
- [191] "3D Operator Training Simulator," EON Reality, [Online]. Available: <http://www.eonreality.com/press-releases/exxonmobil-awards-license-to-eon-reality-for-immersive-3d-operator-training-simulator-technology/>. [Accessed 15 November 2016].
- [192] "COMOS Walkinside ITS (Immersive Training Simulator)," Siemens, [Online]. Available: <http://w3.siemens.com/mcms/plant-engineering-software/en/comos-lifecycle/comos-walkinside/walkinside-its/pages/default.aspx>. [Accessed 15 November 2016].
- [193] "Drilling & Well Control - DrillsIM-6000," Drilling Systems, [Online]. Available: <http://www.drillingsystems.com/products/drilling-and-well-control/drilling-and-well-control/drillsim-6000>. [Accessed 15 November 2016].
- [194] Kelkar & Assoc, "CONNECT- DynaRank - Ocean | Plug-in Details - Schlumberger," Schlumberger, [Online]. Available: <https://www.ocean.slb.com/en/plug-ins/plugin/details?ProductId=PCDR-B1>. [Accessed 11 November 2016].
- [195] R. C. R. Mota, H. Hamdi, M. C. Sousa, E. Sharlin and Z. Chen, "A Visual Framework for Reservoir Connectivity Analysis," in *78th EAGE Conference and Exhibition*, 2016.
- [196] S. Somanath, "Exploring Tabletops as an Interaction Medium in the Context of Reservoir Engineering," September 2012. [Online]. Available:

- http://utouch.cpsc.ucalgary.ca/docs/ucalgary_2012_somanath_sowmya.pdf. [Accessed 11 November 2016].
- [197] S. Somanath, S. Carpendale, E. Sharlin and S. C, "Information Visualization Techniques for Exploring Oil Well Trajectories in Reservoir Models," in *Graphics Interface*, Montreal, Quebec, 2014.
- [198] "Unity - Game engine," [Online]. Available: <https://unity3d.com/>. [Accessed 5 December 2016].
- [199] "MiddleVR for Unity," MiddleVR, [Online]. Available: <http://www.middlevr.com/middlevr-for-unity/>. [Accessed 5 December 2016].
- [200] "Hover UI Kit," Aesthetic Interactive, [Online]. Available: <https://github.com/aestheticinteractive/Hover-UI-Kit>. [Accessed 5 December 2016].
- [201] D. Bowman, E. Kruijff, J. J. LaViola Jr and I. Poupyrev, 3D User Interfaces: Theory and Practice, Addison Wesley Longman Publishing Co., 2004.
- [202] M. R. Mine, "ISAAC: a meta-cad system for virtual environments," *Computer-Aided Design*, vol. 29, no. 8, pp. 547-553, 1997.
- [203] R. Torrens, A. Daoud, M. Amari, A. Sharifzadeh, R. Prakash, B. Al-Enzi and Q. Dashti, "Modeling from Reservoir to Export: A Compositional Approach for Integrated Asset Model of Different Gas Fields in North Kuwait Jurassic Carbonate Reservoirs," in *SPE Reservoir Simulation Symposium*, Houston, Texas, USA, 2015.
- [204] "Integrated Asset Modelling in Production Forecasting," Petro Wiki, [Online]. Available: http://petrowiki.org/Integrated_asset_modelling_in_production_forecasting. [Accessed 5 December 2016].

- [205] F. Cole, D. DeCarlo, A. Finkelstein, K. Kin, K. Morley and A. Santella, "Directing Gaze in 3D Models with Stylized Focus," in *Eurographics Symposium on Rendering (EGSR)*, 2006.
- [206] T. Möller and B. Trumbore, "Fast, Minimum Storage Ray-triangle Intersection," *Journal of Graphics Tools*, vol. 2, no. 1, pp. 21-28, 1997.
- [207] "Solid Wireframe | Nvidia," 2007. [Online]. Available: <http://developer.download.nvidia.com/SDK/10/direct3d/Source/SolidWireframe/Doc/SolidWireframe.pdf>. [Accessed 5 December 2016].
- [208] K. Moreland, "Diverging Color Maps for Scientific Visualization," in *International Symposium on Advances in Visual Computing*, Las Vegas, Nevada, 2009.
- [209] D. Freeman, R. Vennelakanti and S. Madhvanath, "Freehand Pose-based Gestural Interaction: Studies and Implications for Interface Design," in *Intelligent Human Computer Interaction (IHCI)*, 2012.
- [210] S. Boring, M. Jurmu and A. Butz, "Scroll, Tilt or Move It: Using Mobile Phones to Continuously," in *Computer-Human Interaction of Australia (OZCHI)*, Melbourne, Australia, 2009.
- [211] A. Colgan, "4 Design Problems for VR Tracking (And How to Solve Them)," Leap Motion, [Online]. Available: <http://blog.leapmotion.com/4-design-problems-vr-tracking-solve/>. [Accessed 5 December 2016].
- [212] F. Bacim, A. Silva, R. Balakrishnan, Y. Cao, C. L. North and N. F. Polys, "Increasing Selection Accuracy and Speed through Progressive Refinement," CiteSeerx, 2015.
- [213] G. L. Kinsland and C. W. Borst, "Visualization and Interpretation of Geologic data in 3D Virtual Reality," *Society of Exploration Geophysicists*, vol. 3, no. 3, pp. SX13-SX20, August 2015.

- [214] M. Gleicher, D. Albers, R. Walker, I. Jusufi, C. D. Hansen and J. C. Roberts, "Visual Comparison for Information Visualization," *Information Visualization - Special issue on State of the Field and New Research Directions*, vol. 10, no. 4, pp. 289-309, 2011.
- [215] A. E. Mostafa, S. Carpendale, E. V. Brazil, D. Eaton, E. Sharlin and M. C. Sousa, "FractVis: Visualizing Microseismic Events," in *9th International Symposium, ISVC*, Crete, Greece, 2013.
- [216] A. E. Mostafa, J. Cevolani, E. V. Brazil, E. Sharlin and M. C. Sousa, "PetroVis: Exploratory Visualization for Petrographic Characterization," in *Proceedings of IEEE VIS Extended Abstracts*, 2013.
- [217] J. T. Cevolani, A. E. Mostafa, E. A. V. Brazil, L. Costa de Oliveira, L. Goliatt da Fonseca and M. C. Sousa, "Computational Methodology to Study Heterogeneities in Petroleum," in *75th EAGE Conference & Exhibition*, 2013.
- [218] M. Sambridge, C. Beghein, F. J. Simons and R. Sneider, "How do We Understand and Visualize Uncertainty?," *The Leading Edge*, vol. XXV, no. 5, pp. 542-546, 2006.
- [219] H. R. Nagel, "Scientific Visualization versus Information Visualization," Citeseerx.
- [220] Christie Digital Systems, "Application Story - Lights On for Mine Rescue Training," 2010. [Online]. Available: https://www.christiedigital.com/CaseStudies/2013/Coal_Services_Case_Study.pdf. [Accessed 14 November 2016].