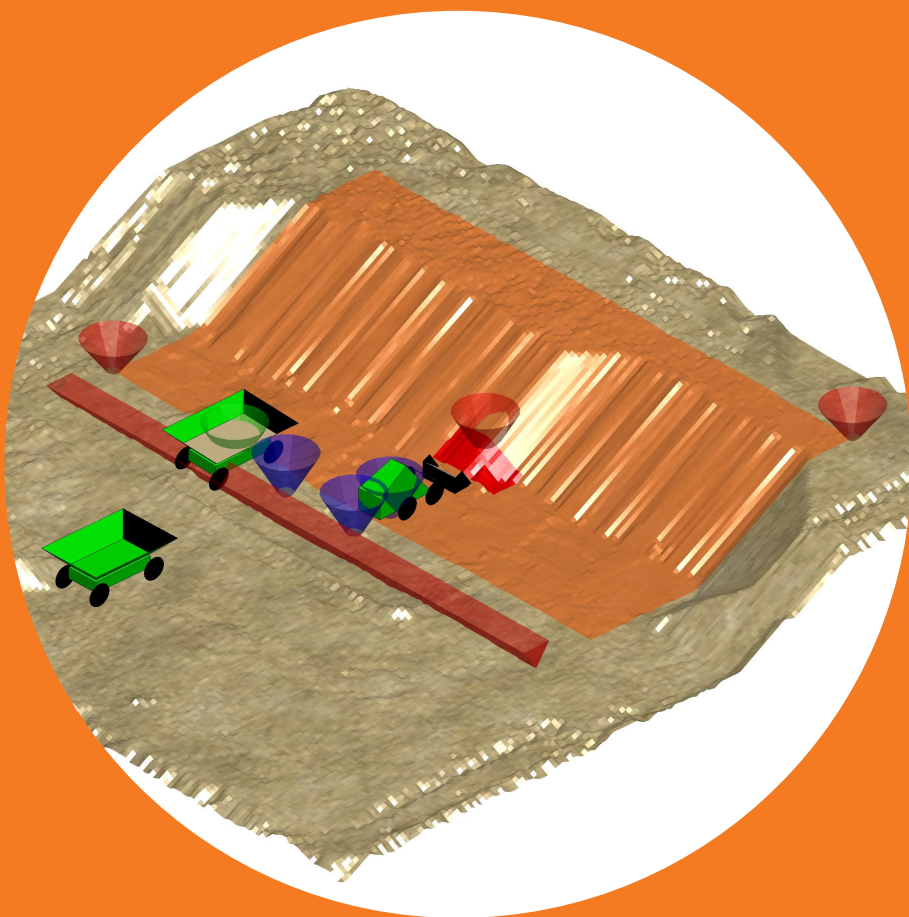


High-Level Job Planning for Automated Earthmoving

Eric P. Halbach



High-Level Job Planning for Automated Earthmoving

Eric P. Halbach

A doctoral dissertation completed for the degree of Doctor of Science (Technology) to be defended, with the permission of the Aalto University School of Electrical Engineering, at a public examination held in the lecture hall AS1 of the school on 17 June 2016 at 12 o'clock.

**Aalto University
School of Electrical Engineering
Department of Electrical Engineering and Automation
Generic Intelligent Machines (GIM)**

Supervising professor

Professor Ville Kyrki

Thesis advisor

Professor Emeritus Aarne Halme

Preliminary examiners

Professor Sanjiv Singh, Carnegie Mellon University, United States

Dr. Martin Magnusson, Örebro University, Sweden

Opponent

Associate Professor Reza Ghabcheloo, Tampere University of Technology, Finland

Aalto University publication series

DOCTORAL DISSERTATIONS 113/2016

© Eric P. Halbach

ISBN 978-952-60-6855-8 (printed)

ISBN 978-952-60-6856-5 (pdf)

ISSN-L 1799-4934

ISSN 1799-4934 (printed)

ISSN 1799-4942 (pdf)

<http://urn.fi/URN:ISBN:978-952-60-6856-5>

Unigrafia Oy

Helsinki 2016

Finland



Author

Eric P. Halbach

Name of the doctoral dissertation

High-Level Job Planning for Automated Earthmoving

Publisher School of Electrical Engineering

Unit Department of Electrical Engineering and Automation

Series Aalto University publication series DOCTORAL DISSERTATIONS 113/2016

Field of research Automation Technology

Manuscript submitted 18 December 2015

Date of the defence 17 June 2016

Permission to publish granted (date) 25 April 2016

Language English

☒ **Monograph**

☐ **Article dissertation**

☐ **Essay dissertation**

Abstract

High-level job planning strategies were developed which enable pile transfer and area clearing jobs to be performed autonomously by a robotic wheel loader. A job is first planned on a 3D surface model of a worksite by positioning graphical tools representing areas and approach directions for scooping, dumping and clearing material. The ground model can be from a recently-acquired surface scan, allowing the job to be configured ad-hoc without the prior need of a global map. Algorithms interpret the high-level plan and, based on an updated ground model, generate commands which ideally guide the job to completion with no further human input. Lower-level plans such as driving points are also represented graphically, allowing a remote supervisor to stay in-the-loop by monitoring the intentions of the machine and modifying the plans if necessary. Fully automated jobs were demonstrated in an earthmoving simulation environment developed using Matlab. The algorithms and search parameters for finding clearing paths and filling locations which worked in the simulator were also found to correctly generate commands using ground models obtained from manually-performed area clearing and filling tests using snow and gravel. As proofs-of-concept, a snow clearing test and two pile transfer tests with gravel demonstrated semi-automated work cycles with a robotic loader, whereby driving and joint actuation were computer-controlled, with transitions between separate actions commanded manually. The snow clearing test demonstrated updated paths being generated based on the changing state of the worksite. The planning tools and algorithms were also extended to jobs including dump trucks and multiple loaders, and applied to a large-scale simulated hillside excavation. Additional simulations evaluated the proposed alternative High Point (HP) method for generating scooping commands, which orients the loader towards the highest point in the pile or slope section from an adjacent stage point. This was compared with a Zero Contour (ZC) method which selects perpendicular scooping approaches along the bottom contour of the slope. Various excavation jobs with truck loading showed that assuming the same bucket filling efficiency, the HP method offers the advantage of a higher excavation rate due to its more limited driving pattern. For the larger plateau excavation jobs, the workspace was subdivided by scanning with the smaller rectangular Scoop Area (SA). It was found that compared with the ZC method, the HP method tends to achieve its maximum excavation rate with SAs which are narrower and longer. Factors which increased the amount of material to excavate per area, including a higher plateau and more surrounding slope collapse, were found to generally result in smaller SAs achieving higher excavation rates.

Keywords robotics, automation, earthmoving, excavation, wheel loader, job planning, supervisory control, 3D graphics, augmented reality

ISBN (printed) 978-952-60-6855-8

ISBN (pdf) 978-952-60-6856-5

ISSN-L 1799-4934

ISSN (printed) 1799-4934

ISSN (pdf) 1799-4942

Location of publisher Helsinki

Location of printing Helsinki

Year 2016

Pages 246

urn <http://urn.fi/URN:ISBN:978-952-60-6856-5>

Preface

The topic of this thesis is motivated by my interest in space exploration, and particularly Mars colonization, which I believe to be a worthy and necessary endeavour for humankind to undertake. This was also the motivation behind my Master's thesis, which was the development of a simulator for modeling robotic earthmoving tasks. I feel fortunate that I had the possibility to continue this work for my doctorate studies, focusing on the aspect of high-level job planning for automating robotic earthmoving.

The work herein was accomplished with the support of many individuals. Though they are too numerous to name individually, I greatly appreciate the camaraderie and friendship of all who I have spent time with since this work began in 2008, both in- and outside the academic sphere. Some individuals, however, I would like to thank directly.

First of all, I am grateful to Professor Aarne Halme, my thesis instructor and original supervisor, for the opportunity to stay at the Automation Technology Lab (ATL) after graduating from SpaceMaster, and for the guidance and feedback during my doctorate studies. I appreciated the working environment under his supervision, which offered both the freedom to develop ideas and support for implementing them.

I am also grateful to Professor Ville Kyrki, my supervisor, for facilitating a smooth transition to the new lab group, overseeing the conclusion of my research and his valuable advice and suggestions. Thanks also to Dr. Jussi Suomela who supervised my work, to the group leaders in the Generic Intelligent Machines (GIM) Centre of Excellence (CoE) for their guidance, and to all the GIM-ATL colleagues for the great working atmosphere.

I would like to thank my preliminary examiners, Professor Sanjiv Singh and Dr. Martin Magnusson, for their time and feedback which helped me to improve my thesis.

I would like to acknowledge the Academy of Finland for funding the

GIM CoE, which made my doctorate research possible, and also the Graduate School in Electronics, Telecommunications and Automation (GETA), which I was fortunate to be part of for 2 years. Special thanks to Marja Leppäharju for her work at GETA and support during the thesis submission process.

Thanks also to the GIM colleagues who made the Avant field tests possible, especially Dr. Mika Hyvönen and Dr. Reza Ghabcheloo from the Tampere University of Technology. I am also grateful to Professor Vadim A. Zhmud for supporting my research visit at Novosibirsk State Technical University, and to the colleagues there for their hospitality.

Outside of my studies, it has been a pleasure hanging out with friends in, among other places, Helsinki, Würzburg, Toronto, Vancouver and Lynn Valley, where it all started.

Finally, I would like to thank my parents, my sister and her family, and my relatives in Finland, Sweden, Germany and Switzerland for all the support over the years.

Helsinki, May 27, 2016,

Eric Halbach

Contents

Preface	i
Contents	iii
List of Figures	xi
List of Tables	xix
1. Introduction	1
1.1 Robotic Earthmoving	4
1.2 Problem Statement	6
1.3 Contribution of the Dissertation	7
1.4 Author's Contribution within the Research Group	8
1.5 Declaration of Previous Work	9
1.6 Thesis Outline	10
2. Related Work	11
2.1 Human-Robot Mars Operations	12
2.2 Site Modeling	14
2.3 Interactive 3D Graphical Planning	15
2.4 Construction Simulation	16
2.5 Autonomous Earthmoving	18
2.6 Multi-Robot Earthmoving	21
2.6.1 Behaviour-Based Resource Collection	23
2.7 Supervisory Control for Multi-Robot Systems	24
2.8 Augmented Reality	26
3. Job Planning Tools and Strategies	27
3.1 Simulation Environment	27
3.2 Robotic Machine Capabilities	33

3.3	Interactive 3D Graphical Planning Approach	35
3.4	Pile Transfer Job	36
3.4.1	Source Specification	37
3.4.2	Dump Pile Specification	39
3.4.3	Automatic Interpretation of Job Plan	41
3.4.4	High Point Scooping Method	44
3.4.5	Area Filling using Virtual Pile	49
3.4.6	Plan Modification and Scoop Area	54
3.4.7	Discussion	57
3.5	Area Clearing Job	58
3.5.1	Planning Features and Path Search Strategy	60
3.5.2	Effect of Path Search and Job Parameters	64
3.5.3	Additional Planning Features	68
3.5.4	Discussion	69
4.	Outdoor Tests Using Snow	71
4.1	Outdoor Test 1: Manually-Driven Area Clearing	72
4.1.1	Surface Model	74
4.1.2	Path Search	76
4.1.3	Discussion	78
4.2	Outdoor Test 2: Manually-Driven Area Filling	78
4.3	Augmented Reality Versions	82
4.3.1	Rendering and Interaction	82
4.3.2	Augmented Reality Area Clearing Tools	86
4.3.3	Augmented Reality Area Filling Tools	86
4.4	Outdoor Test 3: Manual Area Clearing	88
4.4.1	Path Search	90
4.4.2	Discussion	93
4.5	Outdoor Test 4: Robotic Area Clearing	93
4.5.1	Apparatus and Preparation	94
4.5.2	Experiment	97
4.5.3	Discussion	102
5.	Outdoor Tests Using Gravel	105
5.1	Outdoor Test 5: Manually-Driven Area Filling	105
5.1.1	Discussion	109
5.2	Outdoor Test 6: Manually-Driven Pile Loading	110
5.3	Outdoor Test 7: Robotic Pile Transfer	114
5.3.1	Analysis and Discussion	120

5.4 Outdoor Test 8: Robotic Pile Transfer	123
5.4.1 Analysis and Discussion	125
6. Multiple Vehicle Earthmoving	129
6.1 Pile Transfer with Dump Trucks	129
6.1.1 Modification of Plan	132
6.2 Mars Hillside Excavation	134
6.2.1 Robotic Excavation on Mars	135
6.2.2 Site Planning and Machinery Requirements	136
6.2.3 Automatic Plan Interpretation	138
6.2.4 Hauling Distance	140
6.2.5 Multi-Loader Excavation	142
6.2.6 Discussion	149
7. Comparison of Scooping Approach Method and Scoop Area Dimensions	151
7.1 Zero Contour Method	153
7.2 Simulation of Bucket Filling Effectiveness	157
7.3 Comparison of Scooping Approach Method for Single Pile . .	158
7.3.1 Discussion	162
7.4 Scooping Method Comparison for Plateau Section	163
7.4.1 Discussion	166
7.5 Optimal Scoop Area Dimensions and Workspace Division Strategy	167
7.5.1 Job 3 Results with Fixed SA Strategy	171
7.5.2 Job 3 with Variable SA Workspace Division	175
7.5.3 Discussion	178
7.6 Effect of Job Parameters on Optimal SA Dimensions	180
7.6.1 Effect of Plateau Height	182
7.6.2 Effect of Slope Angle	184
7.6.3 Effect of Surrounding Slopes	185
7.6.4 Volume per Combined Drive Time	187
7.6.5 Discussion	188
7.7 Multi-Loader Cases	190
7.7.1 Discussion	193
8. Conclusion and Future Work	195
References	201

Appendices	209
A. Avant 320 Forward Kinematics	211
B. Graphical Tool Mouse Interface	215
C. Virtual Pile Mouse Mapping	217

List of Acronyms

2D	Two-dimensional
3D	Three-dimensional
AR	Augmented Reality
CAD	Computer-Aided Design
CoE	Centre of Excellence
GIM	Generic Intelligent Machines
GMPP	GIM Machine Path Planner
GPS	Global Positioning System
GUI	Graphical User Interface
HIL	Hardware-In-the-Loop
HP	High Point
IHA	Department of Intelligent Hydraulics and Automation
IMU	Inertial Measurement Unit
ISRU	In-Situ Resource Utilization
MaCI	Machine Control Interface
MHP	Mars Homestead Project
NASA	National Aeronautics and Space Administration
NSTU	Novosibirsk State Technical University
ODE	Open Dynamics Engine
RANSAC	RANdom SAmple Consensus
RTK	Real Time Kinematic
SA	Scoop Area
SZ	Scoop Zone
TUT	Tampere University of Technology
ZC	Zero Contour

List of Symbols

Symbol	Units	Description
d	m	Cone diameter
\mathcal{F}_g	—	Global reference frame
\mathcal{F}_k	—	Camera reference frame
\mathcal{F}_r	—	Riegl laser scanner reference frame
f	m	Camera focal length
h	m	Pile height
kT_r	—	Transformation matrix from \mathcal{F}_r to \mathcal{F}_k
kx	m	X-coordinate in \mathcal{F}_k
ky	m	Y-coordinate in \mathcal{F}_k
kz	m	Z-coordinate in \mathcal{F}_k
l	m	Pile central section length
\overrightarrow{RK}	m	Camera position with respect to Riegl
r	m	Cone radius
rR_k	—	3D rotation matrix from \mathcal{F}_k to \mathcal{F}_r
rx	m	X-coordinate in \mathcal{F}_r
ry	m	Y-coordinate in \mathcal{F}_r
rz	m	Z-coordinate in \mathcal{F}_r
V	m ³	Pile volume
w	m	Pile central section width
θ	radians	Virtual Pile slope angle

List of Figures

1.1	Mars Hillside Settlement design by Georgi Petrov, 2004 [1] (used with permission).	3
1.2	Avant 320 skid-steered compact wheel loader.	6
2.1	Related research areas and interpretation of connections made in this thesis.	11
2.2	Virtual Avant 320 with simulated digging capability devel- oped using ODE.	17
2.3	Resource collection and delivery experiment with LEGO Mind- storms NXT robot [2].	24
3.1	Standard world model for single-machine simulations. . . .	28
3.2	Standard pile model for initializing workspace, composed of right-triangular prism with half-conical ends. Specifica- tions: central width w , height h , cone diameter d and volume V	29
3.3	Avant 320 model and 3D pose determination.	29
3.4	Checking for intersection of scoop cutting plane with ground during excavation sequence; red circular points are discretized cutting plane, with square blue points representing corre- sponding ground heights.	30
3.5	Pile loading without (<i>left</i>) and with (<i>right</i>) soil resettling be- haviour.	31
3.6	Simulated soil behaviour: recursive algorithm adjusts steep segments and checks neighbouring segments until stability is achieved.	32

3.7 Unloading of bucket: (<i>left</i>) one scoop-load deposited in pile shape on flat ground beneath bucket blade; (<i>right</i>) 2nd scoop-load added by finding pile-shaped surface which adds new volume and accounts for volume beneath.	33
3.8 Job-level control loop using 3D graphical interface.	35
3.9 Area Tool, used to demarcate source pile; arrows indicate clicking and dragging with mouse pointer.	37
3.10 Histogram of ground elevations in specified area (from Figure 3.9(d)); highest bar assumed to be at ground level. . . .	38
3.11 Dump pile height limits based on volume of one scoop-load (minimum) and geometry of scoop position and front wheel (maximum), with maximum reduced for driving error; dump offset is between scoop position and vehicle reference point.	39
3.12 Virtual Pile with slopes of angle θ (here 45°), used to specify goal state for earthmoving job.	40
3.13 Dump pile specification; arrows indicate clicking and dragging with mouse pointer.	41
3.14 Automated lower-level plan generation: A - Stage point for scooping; B - Scooping Destination; DA - Dump Approach side; C - Dump Stage point; D - Dumping Location; E and F - directions to next dumping locations. Buttons for planning and simulation: 1 - render Area Tool; 2 - render Virtual Pile; 3 - start; 4 - speed up simulation; 5 - pause.	42
3.15 Pile loading using the High Point (HP) method.	46
3.16 Range of scooping configurations resulting from adjustment of boom prismatic joint.	47
3.17 Scanning Virtual Pile for next dumping location, viewed from far end.	50
3.18 Filling Virtual Piles of various dimensions with same volume.	52
3.19 Modification of job plan to avoid obstacle.	54
3.20 Problem with High Point scooping strategy.	55
3.21 Scoop Area used to excavate smaller portions of total surface; Scoop Approach side used to specify search direction. .	56
3.22 End of pile transfer simulation with obstacle; last scoop-load being hauled.	57
3.23 Simulated snow ploughing with bucket: excess volume distributed mostly in front, some to the sides.	59

3.24 Control buttons and high-level plan specification for area clearing job.	61
3.25 Simulation of area clearing Job 1, with a search threshold of 0.3 times the layer thickness.	62
3.26 Computing average ground height along scanning line to find next clearing path.	62
3.27 Result of Job 2 and 3 with the same path search threshold of 0.3 times the layer thickness.	64
3.28 Snow clearing results for Jobs 1-3 and Job 1 with different parameters; zone description: a - more cleaning than heavy paths, b - alternating heavy and cleaning paths, c - more heavy than cleaning paths, d - only heavy paths.	66
3.29 Cleaning residual snow after Job 2 by switching clearing direction.	68
4.1 Experimental apparatus for Test 1: skid-steered Avant 320 compact wheel loader at left, camera, Riegl 3D laser scanner (fixed to cart).	73
4.2 Outdoor Test 1: manually-driven snow clearing.	74
4.3 Creation of surface model from point cloud.	75
4.4 Surface models from Test 1, with residual streaks visible.	76
4.5 Planning tools used with Test 1 surface models; next path found automatically by search algorithm.	77
4.6 Photos and ground models from outdoor area filling test: 9 scoop-loads dumped in a 3 x 3 square pattern, from left to right, and back to front.	80
4.7 Automatic job tracking of outdoor area filling test: next dumping location found by comparing ground heights with corresponding Virtual Pile heights.	81
4.8 Concept of earthmoving construction job planning on Mars using Augmented Reality graphical planning tools (Photo: NASA).	83
4.9 Photo (<i>left</i>) and ground model (<i>right</i>) of initial snow layer from Test 1, with 3 common points located.	83
4.10 Geometry of camera frame, image plane, laser scanner frame and three ground points.	84
4.11 AR tools for area clearing job.	87
4.12 AR Virtual Pile for specifying area filling job.	87

4.13 AR Virtual Pile for Test 2; next dumping location found automatically and marked with cone.	88
4.14 Outdoor Test 3: manual area clearing with snow shovel.	89
4.15 Ground models from Test 3 (manual area clearing with snow shovel); all paths found by algorithm, with alternating heavy and cleaning paths.	91
4.16 Robotic area clearing test apparatus.	94
4.17 Photo and map of Hervanta test area, showing points used to determine coordinate transformation from Riegl laser scanner frame to global frame.	96
4.18 Initial specification of robotic area clearing job and details of first clearing drive.	98
4.19 Clearing drives 2 and 3, for which changes were made to specified clearing area, path search height threshold and lateral positioning offset.	100
4.20 Strategies for maneuvering between clearing paths. Map is dilated version of Figure 4.17(b) used by Avant 635 for obstacle avoidance.	102
4.21 Clearing drives 4-8, all with same settings. The last 4 drives were generally too far to the right, indicating that the lateral offset may have been reduced too much after drive 2.	103
5.1 Outdoor area filling test with gravel; 19 filling actions over 7 locations (4 along back row, 3 in front).	106
5.2 Automatic tracking of pile filling job: next location found by comparing ground heights along line segment with corresponding Virtual Pile heights.	108
5.3 Manually-driven pile loading test with Avant 635 (trial 8 of 10).	112
5.4 Difference in front and rear wheel hydraulic pressure for other successful trials; characteristic spikes occur when wheels begin skidding.	113
5.5 Test 7 workspace, with loading region of pile and two unloading points indicated.	115
5.6 Ground model; pile transfer plan specification, automatic interpretation and user modification.	116
5.7 Five driving points defining job, with command button interface.	118

5.8	Test 7 command station and development of automated scooping logic.	119
5.9	Robotic pile transfer test: 3 full scooping and dumping cycles (Photos: Reza Ghabcheloo).	120
5.10	Pile placement compared with intended Dumping Location.	122
5.11	Verification that piles detectable with Virtual Pile tool.	122
5.12	Workspace for 2nd robotic pile transfer test, with loading and unloading locations labeled.	124
5.13	Initial surface with high-level plan. Arrows indicate that Base point (1) and Stage point (2) manually repositioned to create favourable driving path. Four new buttons for laser scanning added to control interface.	124
5.14	1st pile transfer cycle of Test 8.	125
5.15	Updated ground models in scooping area with Scooping Destination (cone at left) automatically located at highest point; other cone marks Stage point.	126
5.16	Photo and surface model of end piles, viewed from front.	126
5.17	Searching for Dumping Location with Virtual Pile, viewed from back.	127
5.18	Lower Virtual Pile with height of 0.20 m able to detect 2 piles.	128
6.1	Planning for pile transfer job with one loader and two dump trucks.	130
6.2	Simulation of pile transfer job with one loader and two dump trucks.	132
6.3	Modification of plan to avoid obstacles: user moves 2 cones, others reconfigured automatically.	133
6.4	Mars hillside excavation scenario; Area Tool used to specify slope section in vicinity of initial outpost.	137
6.5	Using Virtual Pile to plan worksite regolith storage with varying machinery capabilities.	138
6.6	Autonomous lower-level job planning based on high-level plan.	139
6.7	Large-scale Virtual Pile and use of conveyor belt spreader.	140
6.8	Two dump trucks insufficient for long hauling distance - loader stands idle while waiting for empty truck.	141
6.9	Autonomous Scoop Zone allocation pattern within Scoop Area for 2 loaders, and shift of Scoop Area when all zones completed.	143

6.10	Scoop Zone arrangement within Scoop Area for 3-7 loaders.	145
6.11	Long-term simulation with 8 loaders.	146
6.12	Division of workspace into Sub-Areas to separate teams.	148
6.13	Sub-Area sides adjusted automatically if neighbouring loader gets stuck, allowing other loaders to continue working while problem being solved.	149
7.1	Example of Zero Contour construction around pile on uneven ground, with 0.10 m (<i>left</i>) and 0.15 m (<i>right</i>) threshold above ground level.	153
7.2	Evaluation of convexity for possible scooping locations along Zero Contour; small separate contour at right assigned one possible location.	156
7.3	Checking approach traversability for Zero Contour scooping location.	156
7.4	Job 1: basic pile with 2 m-wide central section and 3.4 m base length; Area Tool 6.2 x 4.2 m.	159
7.5	Total driving (<i>left</i>) and turning (<i>right</i>) vs. minimum bucket fill ratio for Job 1a with HP and ZC method; 10 trials per data point.	161
7.6	Comparison of excavation rate with HP and ZC method for Job 1a; at left, plotted vs. minimum bucket fill ratio (10 trials per data point); at right, same data with each trial plotted vs. average fill ratio.	162
7.7	Excavation rate vs. minimum bucket fill ratio with HP and ZC method (10 trials per data point); at left, for Job 1a with 2x turning speed (60 °/s); at right, for Job 1b (same pile footprint, slopes of 30°).	163
7.8	Job 2: 6 x 1 m plateau section divided into two parts, 1 m high with 30° slope.	164
7.9	Excavation rate results for Job 2a (<i>left</i>), with 30° repose angle, and Job 2b (<i>right</i>), with 30-40° random repose angle; 10 trials per data point.	166
7.10	Job 3: excavation of plateau section 0.87 m high with a 30° slope.	168
7.11	Job 3a: excavation rate for different SA dimensions using HP method (<i>top</i>) and ZC method (<i>bottom</i>); 12 x 3 m workspace; 5 trials per data point.	171

7.12 Job 3b: excavation rate for different SA dimensions using HP method (<i>top</i>) and ZC method (<i>bottom</i>); 11 x 3 m workspace; 5 trials per data point.	173
7.13 Job 3c: excavation rate for different SA dimensions using HP method (<i>top</i>) and ZC method (<i>bottom</i>); 11 x 2.4 m workspace; 5 trials per data point.	175
7.14 Variable SA method attempts to partition 10 x 3.6 m workspace evenly based on desired “target” SA dimensions, here 3 x 1 m. Increments are limited to twice the grid resolution of 0.1 m.	176
7.15 Jobs 3a-c with Variable SA workspace division: excavation rate for varying SA target dimensions using HP and ZC method; 0.87 m-high plateau, 0.8 minimum scoop load ratio, 10 trials per data point.	177
7.16 Jobs for comparing same 11 x 2.4 m workspace with different plateau types, heights and slope angles.	181
7.17 Comparison of excavation rate with changes in slope height and plateau type; 11 x 2.4 m workspace with Variable SA division, 0.8 minimum scoop load ratio, 10 trials per data point.	183
7.18 Job 4c and 3e with Variable SA workspace division: 0.8 minimum scoop load ratio, 10 trials per data point.	185
7.19 Volume excavated per combined drive time for Jobs 3c-e; 11 x 2.4 m workspace with Variable SA division, 0.8 minimum scoop load ratio, 10 trials per data point.	188
7.20 Volume excavated per combined drive time for Jobs 4a-c; 11 x 2.4 m workspace with Variable SA division, 0.8 minimum scoop load ratio, 10 trials per data point.	189
7.21 Job 5 with eight loaders.	191
7.22 Job 5 excavation rates with 4 and 8 loaders: 0.8 minimum scoop load ratio, 10 trials per data point, Fixed SZ.	191
7.23 Job 5 volume per combined drive time with 4 and 8 loaders: 0.8 minimum scoop load ratio, 10 trials per data point.	193
8.1 Linear growth of Mars Hillside Settlement from Phase I to Phase II, for 48 inhabitants. Design by Georgi Petrov, 2004 [1] (used with permission).	200
A.1 Reference frames for Avant 320 forward kinematic model.	211

C.1 Interface for “Type 1” Virtual Pile: far cone stationary and near cone being dragged with mouse. Crescent of points marks permissible positions of moving corner cone given pile height and width restrictions and constant volume. Mouse positions outside crescent are mapped to crescent edge. Zone descriptions provided in Figure C.2. 218

C.2 Virtual Pile volume classification (*left*) based on limits defined by scoop width sw , maximum height $maxH$ and minimum height $minH$. Additional mouse position mappings (*right*) for Type 2-4 piles with zone descriptions (*bottom*). Far cone is stationary and mouse used to drag near cone. Length l is greater and width w lesser of pile’s top rectangular section, “inter” refers to intermediate height. 219

List of Tables

4.1	Summary of outdoor tests using snow.	72
4.2	Path search threshold analysis for Test 1.	78
4.3	Path search threshold analysis for Test 3.	92
5.1	Summary of outdoor tests using gravel.	105
6.1	Comparison of simulation results (1 loader, 2 dump trucks) with different loader scoop capacity; truck capacity 1 m ³ , all driving speeds 0.5 m/s.	142
6.2	Hillside excavation work rates for varying loader and fleet parameters.	147
7.1	Results for Job 1a with High Point (HP) and Zero Contour (ZC) method.	160
7.2	Results for Job 2a with High Point (HP) and Zero Contour (ZC) method.	165
7.3	Maximum Excavation Rates and associated SA dimensions for Job 3 with Fixed and Variable SA workspace division strategies.	178
7.4	Maximum Excavation Rate and Volume per Combined Drive Time for Jobs 3c-e and 4a-c.	187
7.5	Maximum Excavation Rate and Volume per Combined Drive Time for Job 5.	192

List of Algorithms

1	Automated adjustment of boom extension to keep bucket at ground level while scooping.	48
2	Algorithm for finding next clearing path.	63
3	Determining next scooping location from Zero Contour(s) in workspace	155

1. Introduction

The long-term survival of humanity will be enhanced by developing a true spacefaring civilization, distributed across several worlds. This will insure against natural and man-made planetary-scale threats, one example being asteroid strikes, a little-discussed yet constant natural threat facing Earth. By being distributed across several worlds, not only will humanity's survival be enhanced against such an event, but since an advanced spacefaring civilization would likely possess better capabilities of detecting and deflecting harmful asteroids, the chances of preventing catastrophic collisions would also be improved.

Another, more gradual, planetary-scale threat facing humanity is the slow deterioration of our life-sustaining biosphere. While some argue that expanding into space will reduce the incentive to care for our natural environment, a strong case can also be made that this expansion is the best way to preserve it. Concerns about pollution and environmental degradation are indications that the limits of the Earth are being reached for supporting our continued economic and industrial development. While degrowth may be one solution, expanding the economy into space could also ease the pressure on our biosphere. Aside from the development of physical resources, a far greater return may come in the form of new inventions, art and culture developed by settlers on other worlds.

Making the transition to a spacefaring civilization will require the construction of permanent settlements beyond Earth. The best candidate location for this to first take place is the planet Mars. The surface conditions on Mars are extreme, with cold temperatures (average -55°C), low pressures (average 0.6% of Earth sea level) and ionizing radiation from outer space, due to the lack of a substantial atmosphere and magnetic field. Despite this, Mars is the most Earth-like of the other planets in the Solar System, with an almost identical day-length, a similar axis tilt re-

sulting in seasons, and surface vistas analogous to Earth's arid regions. It will therefore be a challenging but welcoming place for human settlers. With minimum-energy Hohmann transfers from Earth only possible every 2.14 years, and the high cost of sending cargo, settlers will need to attain a high level of self sufficiency, using local resources for life support, construction and manufacturing.

Most human Mars exploration missions studied by the world's leading space agencies have included a return phase [3, 4, 5]. This increases the cost of the voyage, and also raises the possibility that no humans will actually stay there long-term, as happened after the Apollo Moon missions. More recently, however, concepts for one-way flights to Mars have been proposed, which would be highly advantageous due to the increased amount of cargo that could be brought - cargo that will be essential for constructing a permanent settlement [6].

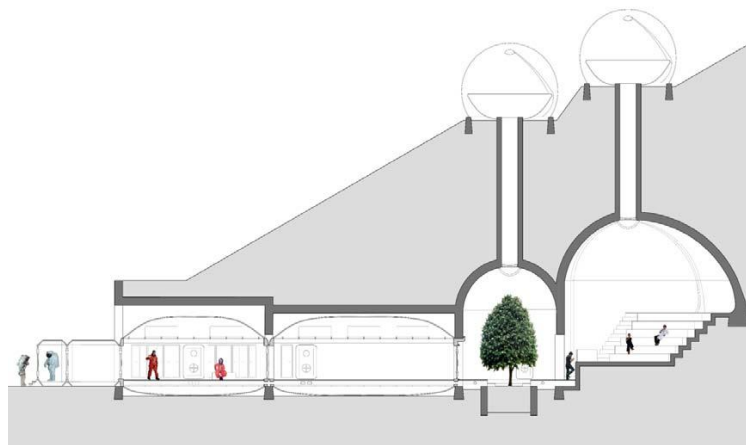
The prospect of one-way settlement missions to Mars comes as a shock to some, yet return flights to Earth would not be completely ruled out and could be a possibility later. Most settlers however would likely be going to stay. This scenario has similarities with the ocean voyages in the Age of Discovery starting in the 15th century, and even earlier with those of the Vikings and Polynesians. In some ways the settlement of Mars will be less extreme than these historical analogues, one example being the possibility of daily communication with a support base on Earth.

A detailed plan for establishing a permanent, growing settlement on Mars using mostly local resources is the Mars Homestead Project (MHP) by Mackenzie et al. of the Mars Foundation [7]. One of the main resources featured in this plan is the local regolith, or ground material, which can be used for resource extraction (water, chemicals), manufacturing building materials (bricks, glass, metals) and shielding from radiation (by burying inhabited spaces) [8].

The MHP architectural design by Petrov outlines the general construction plan, at a site carefully selected and surveyed beforehand [1]. Following the initial establishment of a 12-person outpost consisting of habitat modules brought from Earth, the settlers begin excavating into a nearby hillside. Masonry structures are then built in the dug-out area using locally-manufactured bricks and covered with a protective layer of regolith. The inhabited spaces are highly modular to protect against air leaks. Structures deep inside can be made to hold air pressure with a sufficient regolith overburden (>10 m), while masonry vaults near the exterior would



(a) Phase I of settlement for 24 inhabitants; habitat modules of initial outpost (front centre) connected via inflatable greenhouses.



(b) Section elevation; vaults closer to exterior (at left) encase inflatable modules, while those further inside can be pressurized due to regolith overburden. Sunlight collectors illuminate via lightwells.

Figure 1.1. Mars Hillside Settlement design by Georgi Petrov, 2004 [1] (used with permission).

contain inflatable modules, with thinner regolith layers (>1 m) providing at least some radiation shielding. A concept of the first phase of this settlement for 24 inhabitants is shown in Figure 1.1. Over time, the settlement would grow linearly along the slope to accommodate more inhabitants.

The slope section to be excavated in the first construction phase of the MHP is specified as being 45 m wide and 30 m deep (horizontally) [1]. Assuming a 30° slope, this corresponds to a regolith volume of 11691 m^3 which must be excavated, removed and stored. This figure would be even higher due to slope collapse surrounding the excavated section. Other

earthmoving tasks that may also be required on site include clearing and leveling areas (for landing pads, roads and infrastructure deployment), building blast-protection berms around landing pads, and harvesting regolith for In-Situ Resource Utilization (ISRU), which can be combined with the hillside excavation.

1.1 Robotic Earthmoving

The large-scale earthmoving work which would be useful for establishing a settlement on Mars, or also the Moon, would involve repetitive and lengthy surface operations in the hazardous outdoor environment, a task well-suited to robotic machines. The main reason for using robots for this type of work is to increase safety, given the radiation exposure and risk of depressurization that human drivers would face while operating machinery on the surface.

This leads to the question of how the machines would be controlled. If a human crew were located on site as in the MHP plan, direct teleoperation would be possible. In the case of a long telecommunication time delay such as between Earth and Mars (4-21 minutes one-way), only high-level control and monitoring would be possible, requiring machines with a high degree of autonomy. Even if a crew was located on site, the autonomous approach would be desirable to reduce operator workload and maximize the time available for other critical tasks, such as setting up and maintaining equipment (power generation, manufacturing, greenhouses), masonry construction, exploration and science.

For applications closer to home, similar motivations exist for the application of robotic earthmoving technology at terrestrial mining, construction and material storage sites. Separating humans from the worksite reduces exposure to threats such as landslides, collisions with other machines, harmful dust, and, in underground mines, toxic fumes. Another benefit is that costs can be saved if a human operator does not need to commute to the worksite, be it underground or in a remote desert, and can control and monitor the machines from the comfort of an office.

Some commercially-available mining systems are already able to fully automate certain tasks such as truck hauling, while RioTinto's Mine of the Future in Western Australia combines this capability with automated drilling and train transport, all monitored from an operations centre hundreds of kilometres away in Perth [9, 10, 11]. Other tasks such as exca-

vation and material loading usually still require the control of a skilled human operator, however this can be done remotely using teleoperation. Sandvik's Automine and Caterpillar's Command for Underground systems, for example, allow one ground-level operator to control several underground loaders by teleoperating loading (scooping) actions while the hauling and dumping segments of the work cycle are performed autonomously [12, 10].

Current research in robotic excavation and loading may lead to these tasks also becoming fully automated in the future. Assuming this does occur, and one or more robotic earthmoving machines are able to work completely independently for a long duration, it will still be necessary for a human to stay *in-the-loop*. The minimum input required would be the initial high-level plan which tells the machines what to do, while afterwards, for safety reasons and since unexpected problems may be inevitable, a human would likely need to remain present in a supervisory capacity.

This would be an effective combination of human and robotic capabilities, with robots doing the repetitive and potentially dangerous work while humans provide their cognition, by assessing the state of the work-site and conveying high-level plans to the machines. Work would be planned and monitored remotely, with humans providing occasional input when needed. Direct control of the machines would only occur in exceptional cases. This type of telematic control scheme has been described extensively by Sheridan as *supervisory control* [13]. The main feature of supervisory control is that lower-level control loops, such as for individual scooping actions and driving, are closed directly by the machines, while higher-level loops, such as deciding larger areas for digging, are closed by the human operator.

A similar idea of *task-level* control is described by Schubert and How with space applications in mind [14]. In task-level control, intelligent sequencing and large-scale planning is done by the human, while robots carry out simpler command sequences such as driving and joint manipulation, which are obtained after breaking down the higher-level task.

In this thesis some components of a supervisory control system for earthmoving jobs are developed which are focused on high-level planning, i.e. *where to dig* and *where to dump* material. The earthmoving actions which are studied include those which are possible with a compact wheel loader, which was chosen since one variant, a skid-steered Avant 320, was available to the author at Aalto University in Otaniemi, Espoo for hardware ex-

periments (see Figure 1.2) [15]. Two types of jobs are studied - pile transfer and area clearing - which include the three main actions a wheel loader is capable of: loading material from a pile, dumping material, and clearing along a path. Some experiments were also conducted with a centre-link-steered Avant 635, presented first in Section 4.5.



Figure 1.2. Avant 320 skid-steered compact wheel loader.

This work was part of the broader Future Worksite research effort at the Finnish Centre of Excellence (CoE) in Generic Intelligent Machines (GIM) Research, a collaboration between Aalto University and the Tampere University of Technology (TUT) from 2008 to 2013. The Future Worksite is a system architecture concept for coordinating a distributed multi-entity team of humans and robots working together at a construction site [16]. Various methods of controlling robotic machinery would be used in this architecture, including on-site gesture commands, and remote teleoperation and supervisory control.

1.2 Problem Statement

The general problem addressed by this thesis is the remote operation of one or more robotic wheel loaders without the use of direct teleoperation, either because it is not possible due to a long time delay, or undesirable to reduce operator workload. The problem is more specifically broken down in the following way:

- specifying a high-level job plan which represents enough work to last a time span of minutes to hours, which can be done visually on a newly-acquired ground model with no previous knowledge of the worksite
- automatically interpreting a high-level plan and generating lower-level plans based on the current state of the worksite, ideally keeping the machine(s) occupied until job completion without further human input
- making use of the cognition of the human, who remains in-the-loop in a supervisory role, by displaying the intentions of the machine and allowing the human to modify plans.

1.3 Contribution of the Dissertation

While studying the above problem, a number of contributions are made in this dissertation which are summarized in the following list:

- an algorithm and search parameters are proposed which automatically generate commands for a robotic machine to *fill* a specified dump pile with material
- an algorithm and search parameters are proposed which automatically generate commands for a robotic wheel loader to clear an area of snow using a bucket attachment
- the simple alternative “High Point” (HP) method, which scoops towards the highest point in the current workspace from a fixed point, is shown to offer the possible benefit of reduced driving compared with a *Zero Contour* (ZC) method, which selects a perpendicular approach along the bottom edge of the slope
- optimal rectangular workspace subdivision dimensions are found for even slope excavation and truck loading by a compact skid-steered wheel loader using the HP and ZC method, with the effect of slope height, angle and surrounding slope collapse also studied
- interactive 3D graphical tools are developed which allow a remote user

to visually specify high-level earthmoving job plans on a newly-acquired worksite model using a mouse interface, and which allow automatically-generated lower-level plans to be modified if necessary.

In addition to these contributions, some original ideas are presented including Augmented Reality (AR) versions of the job planning tools, and a proposed workspace division strategy for multi-loader hillside excavations which keeps the machines separated while attempting to excavate the slope evenly.

1.4 Author's Contribution within the Research Group

The work presented in this thesis was conducted from January 2008 to June 2014 at the Department of Automation and Systems Technology, School of Electrical Engineering, Aalto University (until 2010 the Helsinki University of Technology), and continued independently until December 2015. From October 2008 to February 2009 the author was absent for a traineeship at the European Commission in Brussels.

All graphical tools, planning algorithms and simulations presented in this thesis were developed by the author. The focus on pile transfer jobs by an Avant 320 wheel loader originated in the author's Master's thesis [17], while the idea of including snow clearing came from Dr. Jari Saarinen. Outdoor tests 1, 2, 3 and 5 in Chapters 4 and 5 were carried out independently by the author, while tests 4, 6, 7 and 8 were initiated and directed by the author and conducted in close cooperation with colleagues at the GIM CoE. In particular, Dr. Mika Hyvönen and Dr. Reza Ghabcheloo from TUT's Department of Intelligent Hydraulics and Automation (IHA) made these tests possible by arranging for the use of the robotic Avant 635. The following list mentions the names of those (aside from the author) who were directly involved in the outdoor tests:

- Test 4 (Section 4.5): Mika Hyvönen, Miika Ahopelto, Matthieu Myrsky, Antti Maula, Jussi Tervonen
- Test 6 (Section 5.2): Reza Ghabcheloo, Miika Ahopelto, Antti Kolu
- Tests 7 & 8 (Sections 5.3 & 5.4): Reza Ghabcheloo, Mika Hyvönen, Kimmo Rajapolvi, Antti Kolu.

Software developed by many colleagues was used for the outdoor tests, including GIMnet, the GIMnet Machine Control Interface (MaCI), the Avant 635 autonomous driving and scoop control, the GIMsim Hardware-In-the-Loop (HIL) Simulator, GIM Machine Path Planner (GMPP) for commanding the Avant 635 and MapSaver for building height maps from the Avant 635 tilting laser scanner [16, 18, 19, 20, 21].

Outdoor tests 7 & 8 were made possible through a close collaborative effort with Dr. Reza Ghabcheloo. This began with a joint analysis of the data from Test 6, from which the idea for the Avant 635 autonomous scooping behaviour emerged (see Section 5.2). Next, the Matlab pile transfer simulation user interface developed by the author (presented in Section 3.4.3) was modified and integrated with Matlab functions, scripts and Simulink models from IHA, enabling the direct receiving of data from the Avant 635 and sending of commands (see Section 5.3) [22]. Dr. Ghabcheloo developed the final autonomous scooping logic used, described in Section 5.3.

From October 2012 to January 2013 the author was abroad for a research visit at Novosibirsk State Technical University (NSTU), Russian Federation, under the supervision of Professor Vadim A. Zhmud, Chair of the Department of Automation at the Faculty of Automation and Computer Engineering. During this time, the author continued the thesis work remotely and also collaborated with colleagues at NSTU on writing two papers and in developing autonomous resource collection and unloading behaviours for a LEGO Mindstorms NXT robot (see Section 2.6.1) [2, 23].

1.5 Declaration of Previous Work

Some of the work presented in Chapters 3, 4 and 6 has been published in conference proceedings and a journal article, for which the author was the main contributor. These are summarized below:

- preliminary versions of the 3D graphical job planning tools, highlighting their potential for planetary base construction activities [24]
- presentation of the job planning tools in the context of the GIM Future Worksite (written in close collaboration with Tomi Ylikorpi) [25]
- Augmented Reality (AR) versions of the tools [26]

- full versions of the planning tools for pile transfer and area clearing jobs [27]
- worksite planning and simulations of a large-scale hillside excavation job involving dump trucks [23]
- further hillside excavation simulations and automated workspace management for more than one loader [28].

1.6 Thesis Outline

The next chapter begins by presenting the state-of-the-art in research fields related to this thesis, with some discussion how this compares with the approach taken here.

This main body of the thesis is divided into five chapters. First, Chapter 3 presents the simulation environment that was used, followed by the planning tools and algorithms that were developed for pile transfer and area clearing jobs by one robotic skid-steered wheel loader. These are demonstrated with full pile transfer and area clearing simulations.

Next, Chapter 4 presents four outdoor tests using snow as the ground material, followed by four more tests in Chapter 5 which made use of soil and gravel. These proofs-of-concept demonstrated that the available 3D laser rangefinders could be used with the planning tools for ground modeling, job tracking and command generation, and that these commands could be followed by a robotic Avant.

In Chapter 6 the planning tools and algorithms are extended to multi-machine scenarios, first by including two dump trucks in the pile transfer scenario. A large-scale hillside excavation job is then simulated, with a new machine added to increase the dump pile height. Workspace division strategies are also developed to accommodate more than one loader excavating along the hillside.

Chapter 7 then focuses on comparing the proposed High Point pile loading strategy with a Zero Contour strategy which was implemented, to investigate possible benefits due to reduced driving distance. Workspace division for slope excavation jobs is also studied, with simulations conducted to find the optimal subdivision dimensions for each scooping method. Conclusions and areas for future work are then discussed in Chapter 8.

2. Related Work

This thesis combines elements from several research areas. Figure 2.1 shows a categorization of these, and also an interpretation how they are connected by the work here. Current research in each of these areas is presented and discussed in the following sections.

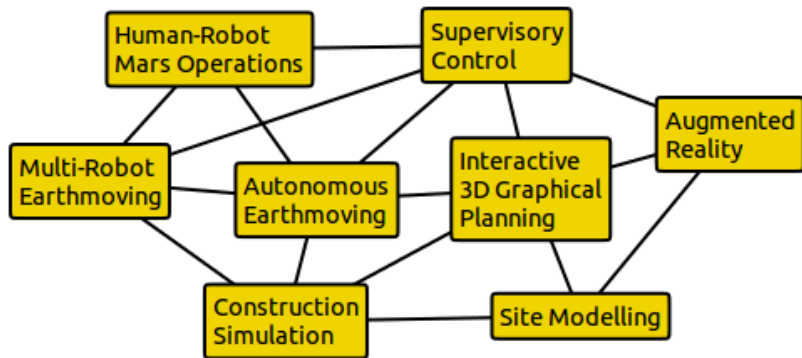


Figure 2.1. Related research areas and interpretation of connections made in this thesis.

Section 2.1 begins by discussing research related to remotely operating robots on Mars and other planetary bodies, which is the main motivation behind this thesis. Sections 2.2, 2.3 and 2.4 then discuss the areas of worksite modeling, 3D interactive planning and construction simulation, which are related to this thesis though not the areas of the main contribution. The main contribution is discussed in Section 2.5 in the area of automated earthmoving, with Section 2.6 discussing concepts for excavation with multiple robots. Sections 2.7 and 2.8 then discuss two further related areas, supervisory control and Augmented Reality.

2.1 Human-Robot Mars Operations

Important considerations for human-robotic operations on Mars are discussed by Mishkin et al. [29]. Two main motivations for using robots are identified: to minimize risk to human crews and maximize time available for critical tasks. Indeed, one man-hour on the surface of Mars may be the most valuable resource after humans arrive, so as many repetitive tasks as possible should be automated and/or managed from Earth.

The long communication time delay between Earth and Mars introduces difficulties, however. Even the nominal maximum delay of 21 minutes one-way can be stretched to hours in non-line-of-sight situations, if the necessary relay satellite network is not available in Mars orbit. Experience with remotely operating rovers on the surface of Mars has led to the use of event-based command sequencing, rather than time-based, the latter being possible with orbital probes operating in a more predictable environment. Event-based sequencing is necessary for surface operations due to the unpredictable interaction between the rover and ground while driving and acquiring samples, leading to uncertainty in task completion time. This will only increase with more interactive operations such as excavation and earthmoving [29].

With humans working on the surface of Mars, the traditional role of Mission Control on Earth will change to one of Mission Support, since updates and commands cannot be received instantly. Control will be shared with the ground crew, who may operate a local Mission Control to oversee robotic operations and assist fellow crew members. Strategic and tactical planning on the order of weeks to years will likely be done from Earth, while immediate actions on the order of seconds to minutes can only be handled locally. Short-term planning on the order of hours to days will be shared, which would include specifying goals and high-level commands to robots. A challenge will be to coordinate the interaction between these timescales and minimize conflicts when transferring control authority [29].

This scenario of a local mission controller operating on the Martian surface, supervising several human and robotic activities but without much time to devote to any particular one, is one motivation behind the 3D graphical planning approach taken in this thesis. This could also be analogous to a remote mining scenario on Earth. By being able to visualize high- to low-level plans, a human supervisor may gain a quicker under-

standing of the current situation pertaining to a particular robot or job after attention has been diverted for some time. If changes to the plan or more detailed instructions are needed, interacting with objects representing the plans by clicking and dragging with a mouse, for example, could provide an intuitive way to transfer the cognitive insight of the supervisor to a robot so as to minimize idle time. The 3D graphical planning tools and strategies developed in this thesis are presented as an original approach to this problem, though they are not the main contribution since a proper evaluation would likely require user tests.

A related scenario is a precursor mission where a robotic colony operates on Mars for years to prepare a site before humans arrive. This increases safety but also the degree of autonomy needed. Huntsberger et al. examine requirements for such missions, focusing on construction-related tasks including load transportation and handling, terrain conditioning and site preparation, and infrastructure servicing and repair [30]. A site preparation simulation was carried out in which multiple behaviour-based robot dozers cleared rocks from a 100 x 50 m area so that solar panels could be deployed. It was found that increasing the number of dozers reduced the total job time as expected, but as the number of agents increased they also began to interfere with one another [30].

L. Parker et al. studied the same multi-robot site preparation task, using a method of opportunistic adaptation to deal with uncertainties in the environment and changes in team configuration [31]. These changes are expected to happen as individual robot capabilities vary over time, such as wear and tear degrading the performance of older team members and newer models being introduced with higher performance. This was also a behaviour-based approach, with individual actions deduced by a function which attempts a globally optimal solution by minimizing energy use while maximizing the area cleared [31].

Although it is usually assumed in this thesis that a human supervisor is located on site for immediate input if problems occur, algorithms are developed which generate commands until job completion in the ideal case, making the job planning strategies applicable for long time-delay situations. For the multi-loader Mars hillside excavation simulation presented in Section 6.2.5, a strategy is also developed for autonomously dealing with a loader experiencing problems. In this case, the workspaces are rearranged such that the other loaders continue working around the stuck loader while the problem is being fixed, either locally or from Earth.

2.2 Site Modeling

The graphical job planning tools and algorithms presented in this thesis are based on the availability of a frequently-updated 3D surface model of the worksite. Some related research has focused on less-frequent scanning of entire construction sites from a few surrounding vantage points for monitoring large-scale changes. Cheok et al. used 3D laser rangefinders to make updated models of a construction site from two fixed locations after days which included earthwork, which were then registered and used to estimate volume changes for planning and management purposes [32].

As part of their Intelligent Excavation System (IES) project, Chae et al. developed a mobile 3D laser scanning platform and automatic scan registration system [33]. This was tested by combining 3 scans from around an 80 x 80 m site, with the aid of 5 spherical targets for the registration procedure. Each scan took approximately 30 min to collect, and the intention for the system was to make 3-4 site models per day.

In order to generate commands for automated excavation actions, more frequent surface updates are required at a local scale where changes are occurring. One solution is to collect these from a dedicated observation platform, which was the method used here in Outdoor Tests 1-5 with a 3D laser rangefinder mounted on a stationary cart or tripod (see Chapters 4 and 5). This can offer the advantage of high-quality scans with a favourable vantage point, however for real applications the platform would need to be mobile to follow the excavation work, thus increasing complexity by adding an agent to the system.

Another option is to use rangefinding sensors mounted on the earthmoving machines themselves, which was the method used in Outdoor Tests 7 and 8 presented in Chapter 5 [21]. Although adding simplicity by reducing the number of agents needed, this method has the problem that since work should not be interrupted by requiring a stop-and-scan approach, scans should be collected while the machine is driving, potentially causing distortion in the 3D point clouds collected due to rough terrain. A method for overcoming this problem was developed by Almqvist et al., which makes use of odometry data but does not require inertial measurement sensors or Global Positioning System (GPS) data [34].

2.3 Interactive 3D Graphical Planning

The graphical approach presented in this thesis has similarities with that taken by Kang et al. in their virtual reality system for road design [35]. This system allows for the visualization of cut and fill volumes associated with road construction, and can be used to compare different routes. Similarities also exist with commercial open-pit mine planning software which allows planners to visualize cut and fill volumes in 3D [36, 37, 38].

One benefit of this approach is that plans represented in 3D are easier for people with different expertise to interpret and understand. In construction applications, this can be important so that planners, workers and decision makers, all of whom may not be familiar with the same details, can see the plan presented visually and gain a common understanding. This case has been well made by Kamat and Martinez in their work on validating construction simulations using 3D visualization [39, 40].

One product making use of this approach is Xsite by Novatron, which makes it possible for an excavator operator to monitor their own progress in relation to a job plan without using any physical stakes to demarcate the workspace [41]. A 3D model of the worksite and the machine's location is rendered on a screen inside the cabin, also showing the elevation difference to the target surface. A wireless data link allows updated plans and site models to be sent between the machine and planning office.

The job planning tools presented in this thesis allow plans to be specified directly on a 3D site model using the mouse pointer. This is analogous to the idea of “drawing on the world” followed by Correa et al. in their multimodal tablet interface for interacting with an autonomous forklift, which allows the user to specify commands by sketching gestures onto live and synthesized camera images of the worksite [42]. Perhaps more similar is the “click to dig” concept presented by Dunbabin et al. in their 3D Graphical User Interface (GUI) for observing job progress and specifying a digging and dumping location for an automated dragline excavator [43]. Some advancements presented in this thesis include specifying *areas* for scooping and dumping, rather than just points; specifying these locations in 3D; and showing a predicted end state at a dumping site.

A reference must also be made to SimCity 2000, a computer game by Maxis [44]. At the start of the game, the player can initialize the land topography by clicking and dragging various tools directly on the terrain model which raise or lower the elevation.

2.4 Construction Simulation

The simulation of earthmoving and other construction tasks has become an important area of research, as it offers the possibility of testing plans before work begins in order to identify problems, select the proper machinery and lower costs and/or environmental impact. Systems such as VITASCOPE by Kamat and Martinez focus more on the visualization aspect, using 3D animation of construction processes to verify that no errors exist in a plan [40]. Another construction simulation and visualization tool is the Virtual Reality Modeling Language (VRML)-based system by Lipman and Reed, in which a variety of worksite objects can be represented including steel beams and backhoe excavators [45]. Such systems may allow for ground material to be removed and added, though not necessarily conserving the amount of material [45, 46].

Other systems focus more on optimizing machinery combinations and estimating costs, such as that developed by Shi and AbouRizk [47]. As an example, they compare excavation scenarios with varying backhoe capacity, two types of dump truck and different numbers of trucks, estimating productivity rates and finding the optimum. Production bottlenecks such as traffic queues can also be identified by simulation, as shown by Henriksen using Proof [48]. The Dynamic Interface Simulation for Construction Operations (DISCO) approach by Huang is more schematic, yet still has animation features with its dynamic block diagrams [49]. This has the advantage of not being computationally intensive (no 3D graphics) while still allowing the modeling and analysis of construction system dynamics.

A review of the simulation and optimization of earthmoving operation logistics is provided by Fu, who also develops a microscopic simulation model for some aspects of the work cycle [50]. This is in contrast to most systems which treat work cycle components macroscopically, for example by using deterministic or historical average values for truck loading time or hauling time. Fu's microscopic simulation uses more detailed models which include vehicle dynamics for the work cycle components investigated, which allows parameters such as fuel consumption for truck hauling to be optimized [50].

To develop automated excavation controllers or for training human operators on simulated machines, more detailed physics-based simulators are required for achieving realistic wheel-ground and tool-ground interaction with forces and ground deformation. Some examples of these are the sim-

ulators developed by Pla-Castells et al. for a wheel loader and Schmidt et al. for a bucket excavator [51, 52]. Advanced construction simulators are also available commercially, such as Vortex by CM Labs, which is marketed for operator training and product evaluation [53].

In early work by the author, a virtual Avant 320 wheel loader was developed in a dynamic simulation environment using the Open Dynamics Engine (ODE), shown in Figure 2.2 [54, 55]. A basic digging capability was developed, however the work was not continued after difficulties were encountered controlling the ease of the tool-ground interaction, thus it was not incorporated into the main body of this thesis.

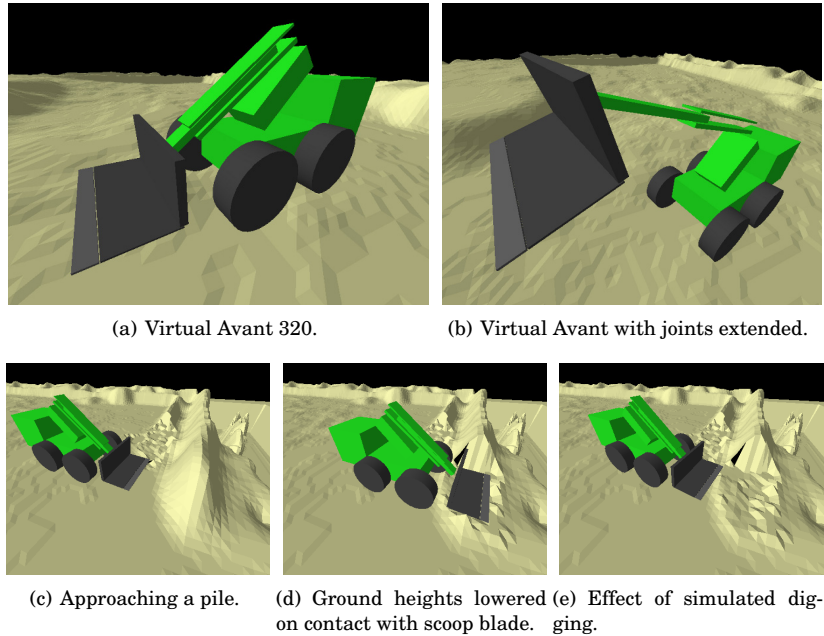


Figure 2.2. Virtual Avant 320 with simulated digging capability developed using ODE.

The earthmoving simulation environment used in this thesis, presented in Section 3.1, is purely kinematic and does not model forces. This simplification was made since the focus of the thesis is on high-level job planning aspects such as where to dig, and not developing controllers for scooping actions for example. The main features of the simulator are the ability to remove and add material while conserving total volume, and the maintenance of a maximum angle of repose at locations where ground heights have changed. It is similar to the pile model developed by Sarata, which was also used to develop algorithms for automatically determining where the next scooping location should be [56].

2.5 Autonomous Earthmoving

This section presents related work in autonomous excavation and earth-moving. Most results are related to wheel loader (or front-end loader) automation, which is the platform used in this thesis, but some results for backhoe excavators are also included as this is another commonly-used type of machine.

Given a high-level plan, the first step for an autonomous system is deciding where to dig in a larger area, after which the bucket loading must be controlled automatically. Singh and Cannon address this problem with their *multi-resolution* task planner, which includes a coarse planner for breaking down a specified workspace into sub-regions, and a refined planner for selecting scooping commands within a sub-region [57]. For the backhoe excavator, the dig is chosen after estimating the outcomes of digging actions at candidate locations using a forward controller [58]. For the wheel loader, the automated planner considers factors including side loading, to maintain lateral balance in the scoop; concavity, to assist penetration into the pile; and distance to the loading truck [57].

The Intelligent Excavation System developed by Seo et al. also includes a coarse planner, and automatically interprets a high-level site plan to generate commands for a backhoe excavator to follow [59]. Intelligent navigation strategies for this system presented by Kim et al. include automated path planning with obstacle avoidance, both for spatial coverage within work areas and transfer between areas [60].

The digging action itself is an especially challenging problem, since factors which are difficult to predict, including moisture, compaction and hidden rocks, can greatly affect the resistance of ground materials. To deal with such uncertainties, Lever and Wang developed a fuzzy controller using force and torque feedback which was intended for Lunar mining applications with a front-end loader [61]. This controller uses a fuzzy rule set to decide how to move the bucket based on current measured force and torque values. Another approach to digging is the iterative excavation process developed by Maeda et al. for a backhoe excavator [62]. With this low-level approach, the manipulator assumes a compliant behaviour, scraping away layers of soil until the target profile is reached. An admittance-based controller developed by Dobson et al. for a full-sized underground loader was found to load rock piles faster and with greater loads than an expert human operator, though using more work [63].

Work by Sarata et al. has focused on developing an automated pile scooping and truck loading system for an articulated-frame wheel loader, first tested with a small-scale laboratory platform [64, 65, 66]. Environment sensing is achieved via stereo cameras, with the pile modelled using an array of columns. An initial scooping point along the base of the pile is assumed to be known, with the scooping direction chosen to minimize the predicted lateral moment on the bucket, usually resulting in a perpendicular approach. A method for selecting the next scooping location is proposed whereby two locations are considered, less than one scoop width left and right of the previous point, with the location selected which minimizes the driving path length to the dump truck [65].

A further demonstration by Sarata, Koyachi and Sugawara involved a full-sized wheel loader [67, 68]. The approximate location of the pile was assumed to be known initially, which allows the loader to make a 3D model of the pile using stereo camera images. The bottom edge of the pile was then constructed using a height threshold of 15-20 cm above the ground. The scooping point is located at the intersection of the bottom edge with a line which passes through the gravitational centre of the pile and is 45° from the perpendicular to the side of the dump truck. As the loader nears the pile, onboard laser scanners are used for more precise relative positioning between the pile edge and bucket blade. Four full scooping and truck loading cycles were then conducted autonomously, all with the same scooping point [67].

Magnusson and Almqvist extend the work of Singh and Cannon for autonomous wheel loader bucket filling by increasing the complexity of the bucket model and considering how the 3D shape of the pile affects filling efficiency [69]. They propose using quadric surfaces to model a pile section, constructed directly from a 3D point cloud, which does not require the pre-processing needed for triangular mesh and elevation map models. They also present an algorithm for automatically classifying piles from ground scan data [69]. Previous work by Almqvist also focused on automated bucket filling for a wheel loader [70]. A key criterion for a successful filling action was to avoid any wheel slippage due to the increase in mechanical wear, and consequently the running costs.

Further work of Magnusson et al. provided an evaluation of autonomous bucket filling actions and coarse-to-fine loading strategies, which maintained the base contour shape of a large pile over many loading actions to ensure the availability of good scooping locations [71].

In the simulator developed by Schmidt et al. for testing an autonomous wheeled backhoe excavator, a behaviour-based approach is used for controlling the robotic machine, which is able to compensate for disturbance forces [52]. Their approach has similarities with this thesis in that surfaces representing the current ground and desired state are compared to generate digging location commands. In work closely related to this thesis, Bonchis et al. presented the development of an automated skid-steered compact wheel loader, tested in a simulation environment [72].

In this thesis some capabilities are added to those found so far in the literature related to robotic wheel loader automation. The first is the initial configuration of the job. In the examples mentioned above, the pile and/or initial scooping location is usually assumed to be known, but in case a new model of a changing worksite has just been obtained, a way is needed to specify where to work. Even if piles can be detected automatically, it may be necessary to specify which pile to load or to limit work to part of a large pile. A tool developed for this purpose is presented in Section 3.4.1. In case a pile has more than one side exposed, it may also be useful to specify the approach direction for scooping actions. Additional tools developed for this purpose are presented in Sections 3.4.3 and 3.4.6.

Another capability developed in this thesis is the ability to specify how to deposit material on the ground. In the related work pertaining to wheel loaders, the unloading portion of the work cycle usually takes place at a dump truck [57, 66, 68] or a stationary bin [72, 69]. Although these are common industrial scenarios, some jobs may require material to be stored directly on the ground in a new pile, therefore a tool developed for this purpose is presented in Section 3.4.2.

An additional contribution is an algorithm for finding paths to clear a rectangular area of snow, with the reasons for including this type of job explained in Section 3.5. In related work, Hess et al. and Saska et al. developed autonomous multi-vehicle formations for clearing airport runways and tarmacs of snow, based on area coverage whereby all areas are traversed once [73, 74]. The search algorithm developed here differs in that it checks the result of a clearing action and can detect spillage, requiring the machine to go back if the residual is over a certain threshold.

A simple alternative strategy for generating scooping approach vectors into a pile or slope is also proposed (Section 3.4.4). Comparisons are made in Chapter 7 with a Zero Contour method, showing a possible benefit in reduced driving distance.

Finally, an investigation is made into the optimal dimensions to use for subdividing a workspace into rectangular subregions, a strategy alluded to by Singh and Cannon [57]. In Chapter 7 simulations are conducted in order to find advantageous dimensions specific to a simulated skid-steered Avant wheel loader in a truck loading scenario.

2.6 Multi-Robot Earthmoving

Performing an excavation job using several robots offers the benefits of a potentially higher work rate, depending on sizing and numbers, and also the redundancy of a distributed system – if one robot fails, then the others can continue the job. This could be important for Mars applications as mechanical breakdown may occur frequently due to the harsh environment.

Coordinating several machines operating in the same workspace presents the dual challenge of avoiding collisions and ensuring that the efforts of individual machines do not interfere with each other, i.e. that progress is made on a global scale. One solution is to make use of a central planner which guides all agents and prevents conflicts. This was the approach favoured in the Nanorover Outpost Project at the National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL), which envisioned using a swarm of solar-powered nanorovers (14 cm long, 1.5 kg) for site preparation tasks on Mars [75].

An analysis of multi-agent coordination for general resource collection and delivery tasks was made by Zhmud et al., with special attention paid to the problem of deciding priority among agents if conflicts occur while driving, collecting and unloading material[2]. These situations are generally resolved using rules of the road such as the right-of-way, yet exceptions can be made if more progress towards the goal state is possible.

To avoid the need for machines to stop for others at intersections, Pecora et al. developed solvers which find conflict-free trajectory envelopes in both the temporal and spatial domain [76]. While not specifically for excavation work, this was demonstrated for a scenario of 7 Automated Ground Vehicles (AGVs) driving in the confined spaces of an underground mine, and also for cases of up to 10 vehicles driving in a 50 x 50 m area, with the constraints of a non-zero starting and minimum speed.

Some disadvantages of top-down, deliberative approaches are that the central planning may be computationally intensive, maintaining communication with all agents may be a challenge, and the system is highly

dependent on the functioning of the planner. Rather than controlling all actions centrally, some researchers have taken a bottom-up behaviour-based approach to this problem, looking to nature for inspiration.

C. Parker et al. developed a swarm construction algorithm called “Blind Bulldozing” for controlling robotic bulldozers [77]. This was inspired by the way in which certain ants clear an area on the ground for nest building. An ant begins by plowing debris from a given point, stopping when unable to continue due to the resistance force reaching a certain level. A new direction is then randomly picked for the next ploughing action. The final nest size is independent of the number of ants involved, and only depends on the initial debris density and pushing force of one ant.

This behaviour was duplicated using 1, 2 and 4 robots for clearing gravel spread on the floor. As the name of the algorithm suggests, the robots were essentially “blind,” with sensing capability limited to detecting the force on the ploughing blade and collisions with other robots. There was no direct communication between the robots; only stigmergy was used to coordinate the job, i.e. indirect communication mediated through the environment. The experiments resulted in roughly the same circular nest shape, with more robots expanding the nest at a faster rate [77]. A similar result was obtained in the simulations carried out by Huntsberger et al. mentioned in Section 2.1 [30]. C. Parker et al. observed that higher interference in the 4-robot case actually resulted in “accidental construction” which accelerated nest building by pushing debris with parts other than the plough, if robots were turning or reversing at the perimeter [77].

Thangavelautham et al. investigated autonomous multi-robot excavation for construction and resource utilization on the Moon [78]. The system used an Artificial Neural Tissue (ANT) architecture as a control system, inspired by the way social insects such as ants and termites construct superstructures without any high-level instruction. All that needed to be specified for this system was a single global fitness function, based on the desired end state of the ground map, and a set of allowable basis behaviours. An evolutionary selection process was then used to “breed” controllers for the specific task at hand. The control was not based on pre-programmed human expertise, but was allowed to evolve its own solution. This was because a human operator’s expertise would typically be limited to operating a single vehicle, but the global effect of many local interactions may be difficult to predict. In such systems, updates from a human operator may take the form of modifying weighting values or priorities of

different behaviours, however the results of these are likely non-intuitive and would likely require trial and error [78].

The multi-robot hillside excavation scenario presented here in Section 6.2 takes a central planning rather than behaviour-based approach, with machines operating in the same vicinity but restricted to separate workspaces to avoid interference.

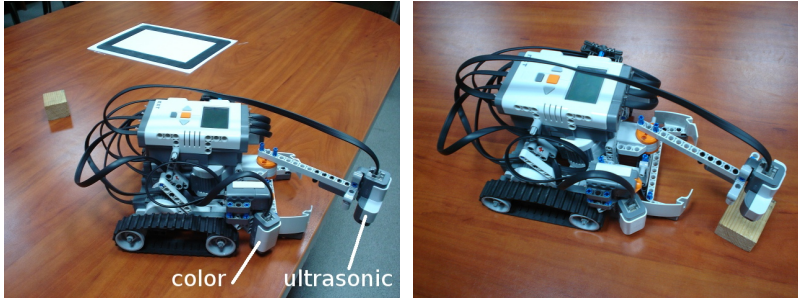
2.6.1 Behaviour-Based Resource Collection

The work by Zhmud et al. included contributions by the author during the research visit mentioned in Section 1.4 [2]. This involved the programming of autonomous resource collection and delivery behaviours using Python for a robot built using LEGO Mindstorms NXT (see Figure 2.3) [79, 80]. This robot had been programmed to drive in straight lines across a table and search for the edge via a downward-facing ultrasonic ranging sensor. The sensor was mounted on a front swing arm which panned back-and-forth through 180° with a frequency of about 1 Hz. If the edge of the table was detected, the robot would stop and turn, reversing a small amount if necessary, and continue driving in a new direction. This resulted in random area coverage of the table, mimicking the foraging behaviour of an ant.

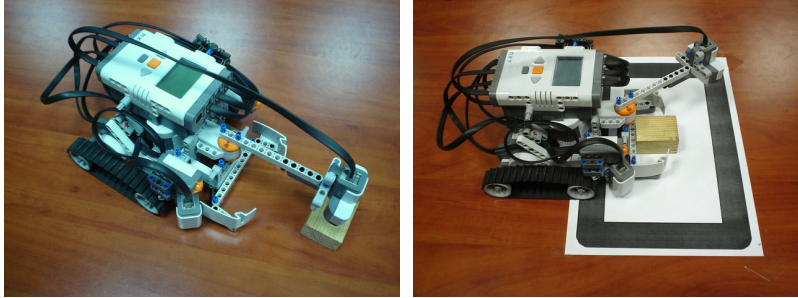
The resource collection task consisted of finding a block of wood, approximately 3 x 4 x 6 cm, and delivering it to an unloading zone represented by a white sheet of A4 paper with a 2 cm-thick black border. These are also shown in Figure 2.3(a). A behaviour was programmed which used the ranging information from the ultrasonic sensor to simultaneously search for the block as it was scanning for the table edge. If found (Figure 2.3(b)), the program would first assess the bearing angle of the swing arm when it detected the block. If sufficiently “in front,” it would drive forward to collect the block in the scoop. If not, the robot would rotate a small amount towards the block, then scan again and assess the bearing angle, repeating until it was aligned for collection (Figure 2.3(c)).

With the block collected, the robot would resume driving randomly, searching for the unloading zone. This was detected using a downward-facing colour sensor mounted on the side of the scoop (see Figure 2.3(a)). When the black border around the paper triggered the sensor (Figure 2.3(d)), the robot would stop and reverse, unloading the block.

Although the experiment was limited to a single-robot scenario, the independent behaviour-based control offers the possibility of scaling up to



(a) Experimental setup; edge of table detected with ultrasonic sensor on swing arm. arm bearing angle checked. (b) Block detected with ultrasonic sensor.



(c) Robot maneuvers until block in front, (d) Block unloaded when black edge of paper detected with color sensor.

Figure 2.3. Resource collection and delivery experiment with LEGO Mindstorms NXT robot [2].

several robots operating simultaneously, which would require some way to handle or avoid collisions. To make it more generally applicable to different environments, other ways of detecting resources and unloading zones would also need to be developed.

2.7 Supervisory Control for Multi-Robot Systems

This section presents research in supervisory control for a fleet of robots, addressing some of the operational considerations outlined in Section 2.1.

The development of a human telesupervision system for a fleet of autonomous robots operating on a planetary surface is described by Podnar et al., with an application scenario of wide-area mineral resource prospecting [81]. Human-robot interaction occurs in three ways: 1) humans make high-level plans and assignments for robots, 2) humans monitor progress of autonomous work, and 3) humans intervene when assistance is required. The user workstation would include several displays for monitoring and interaction, including high-level mission design with a GUI, and the overlaying of sensor data and predicted mineral concen-

trations over images [81]. The main features of this system match those presented in this thesis, with the main difference being the use case, i.e. mineral prospecting as opposed to earthmoving jobs.

Crandall et al. study the concept of “neglect tolerance” of a multi-robot system, and how this can be used to determine the maximum number of robots one human can supervise [82]. The basic factors are “neglect time,” which is the maximum duration that a robot can operate autonomously before its effectiveness drops below a certain threshold, and “interaction time,” which is the time needed by the human operator to bring a robot back to maximum performance. The theoretical maximum number of robots that one human can supervise is the ratio of average neglect time to average processing time, plus one.

Three interaction schemes were compared for a robot navigation task through a maze: direct teleoperation, point-to-point (instructions required at each intersection) and scripted (list of waypoints to follow autonomously). User tests confirmed that the scripted scheme had the highest neglect time, and furthermore maintained a constant effectiveness after long neglect. Teleoperation quickly reached a high effectiveness when in use, and quickly dropped to zero when neglected, while point-to-point shared similar effectiveness although with slower changes. A trade-off was identified between effectiveness and workload: the point-to-point scheme in some cases resulted in higher performance than the scripted scheme but required more attention, and vice-versa [82].

Mau and Dolan focus on scheduling algorithms for multi-robot supervision, for the scenario in which a computer would order robot interaction tasks for a human supervisor in order to minimize total robot down-time and maximize the number of robots a human can supervise. Whereas Crandall et al. determined the maximum fleet size by considering average interaction and neglect times over infinite time, Mau and Dolan show that in the case of finite job length and varying interaction times, the maximum fleet size can be increased [83].

This thesis is not directly related to the work of Crandall et al. and Mau and Dolan, since here it is assumed that the work cycle can be fully automated, without the need for direct human interaction after the job is specified. In the ideal case, the maximum number of robots would only be limited by the scale of the job, i.e. the division of the workspace into zones for each robot which are still big enough to work in effectively. In case frequent supervisor interaction was necessary however, an area for future

work could be to test if using interactive 3D graphical tools and/or Augmented Reality (AR) versions could reduce the interaction time vs. other methods, and increase the number of robots one human can supervise.

2.8 Augmented Reality

Augmented Reality (AR) is the superimposition of information on the naturally-sensed environment, which can be achieved for example by rendering graphical objects on a real-time video image or see-through display. The objects are typically registered with the point of view of the observer and fixed relative to the environment, making it appear that the information is connected to the real world.

The application of AR in the building and construction industry has become an active field of research, due to its potential for assisting on-site managers in visualizing plans and comparing them with the current situation. A software architecture presented by Woodward et al. makes this possible by allowing a mobile AR system at a construction site to access 4D Computer-Aided Design (CAD) and Building Information Model (BIM) information in real time [84]. Behzadan and Kamat developed an AR system that automatically prevents incorrect occlusions of real objects by virtual ones, in order to make the overall scene more compelling [85].

For controlling robots, Milgram et al. investigated the benefits of AR in remotely controlling a robotic manipulator arm [86]. Brujic-Okretic et al. present an AR application for remote vehicle maneuvering in case of temporary loss of vision or poor visibility, which constructs the nearby scene from laser ranging data and superimposes it on a video image [87]. Collett and MacDonald highlight the benefit of using AR to visualize a robot's understanding of the world, e.g. ranging sensor data, and comparing it with the "ground truth" provided by the camera image [88]. Discrepancies can be easily noticed, aiding the developer in coding and debugging.

So far, no results have been found specifically for using AR to guide robots for autonomous earthmoving. Steps towards this are made here in Section 4.3, by rendering and modifying graphical plans directly over a camera image of a worksite.

This concludes the review of work in areas related to this thesis. Some additional references are made to related work, however, throughout the document where they are deemed appropriate.

3. Job Planning Tools and Strategies

This chapter presents job planning tools and strategies that were developed to enable supervisory control for a robotic Avant 320-type compact skid-steered wheel loader with a front bucket attachment. Two types of jobs are included, pile transfer and area clearing, as these include the three main actions a wheel loader is capable of: loading (scooping) from a pile, dumping material and clearing along a path. The jobs are limited to a single machine working in a limited area, in order to begin with the simplest cases and to focus on scenarios that could be demonstrated using one of the available Avants in subsequent hardware experiments.

Section 3.1 begins by describing the simulation environment that was used to develop and test the planning tools and algorithms. This is followed by a discussion in Section 3.2 about what information from the workspace would be needed and what capabilities the robotic loader would have in order for the planning tools to be used in a real-world application. The 3D graphical approach which was taken is then described in Section 3.3, along with the kinds of scenarios in which it could be useful. Next, Section 3.4 presents the pile transfer job, which is the primary focus as this is the main type of work a wheel loader is meant for. Area clearing, on the other hand, may more typically be performed with a plough or bulldozing blade, yet for light material such as snow a bucket attachment can suffice, which is the scenario included here in Section 3.5.

3.1 Simulation Environment

The graphical tools were developed and tested in an earthmoving simulation environment developed previously by the author using Matlab [17, 22]. The simulator allows ground material to be removed at points of intersection with the bucket, and scoop-loads to be deposited on the ground

in a pile shape, while maintaining a maximum slope repose angle and preserving the total volume of material (thereby assuming constant density). It is purely kinematic and does not model forces, since its main purpose is not for developing control systems for the excavation action itself, but rather for generating commands specifying *where to dig* and *where to dump* based on the changing state of a worksite over time.

A worksite is represented by a digital elevation map which can be constructed from any 2D matrix of height values. The standard model used here is 20 x 19 m with a 0.1 m grid spacing, constructed from actual laser scans of a covered testing area at TUT (see Figure 3.1) [16]. This model is relatively flat but has an uneven surface typical of construction sites. For the larger-scale simulations in Sections 6.2 and 7.7, the grid spacing is increased to 0.2 m to lower computing requirements.

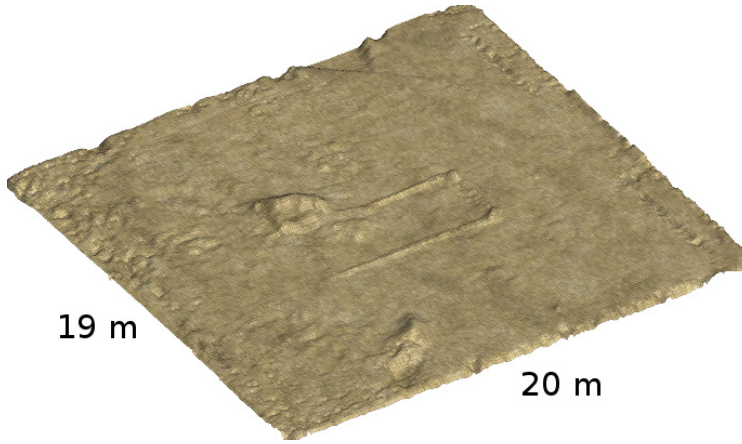


Figure 3.1. Standard world model for single-machine simulations.

Piles can be added to the ground model when initializing a simulation by specifying their location and orientation about the Z-axis. The standard pile model used in this chapter has slopes of 45° , with a right-triangular prism as the central section and half-conical ends, shown in Figure 3.2. The central width w and volume V are also specified, which sets the height h , and consequently the cone diameter d , according to Equation 3.1 for the sum of the cone and prism volume.

$$V = \pi h^3/3 + wh^2 \quad (3.1)$$

It was also made possible to specify piles with a different slope angle, and with a rectangular top section, with the width, length and height also being specified instead of the volume. Some of these different pile types

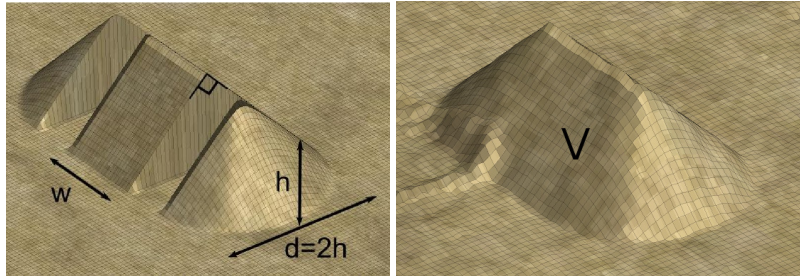


Figure 3.2. Standard pile model for initializing workspace, composed of right-triangular prism with half-conical ends. Specifications: central width w , height h , cone diameter d and volume V .

are used later in Chapter 7.

The Avant 320 model, shown in Figure 3.3(a), has a wheel spacing 0.80 m long by 0.79 m wide and scoop dimensions 0.89 m wide by 0.5 m long, with the scoop volume capacity set at 0.15 m³. These scoop dimensions were based on a different tool attachment than the one in Figure 1.2, which has a larger volume of approximately 0.20 m³. The wheels are rendered as 2D discs and represent the centres of the tires, therefore the total width at the outer edges of the tires would be greater than 0.79 m. Three joints are available for positioning the scoop: a rotary joint between the chassis and base of the boom, telescopic joint for extending the boom, and rotary joint between the end of the boom and the back of the bucket. The Avant 320 forward kinematic model which was used is available in Appendix A.

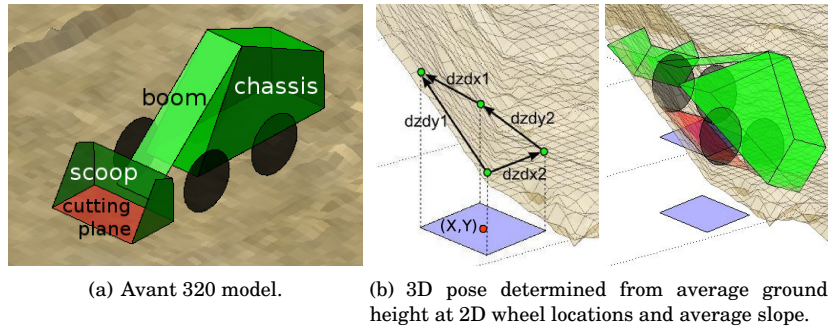


Figure 3.3. Avant 320 model and 3D pose determination.

The vehicle's position is defined in the horizontal plane by the (X,Y) coordinates of its reference point (centre of the 4 wheel-ground contact points) and heading about Z of the forward-facing direction. Driving occurs autonomously between points by first turning on the spot for alignment, then driving along a straight line. The driving speed is set at 0.5 m/s, and turning rate at 30 °/s, with one timestep in the simulator being 1/3 s. Changes to the joint angles for positioning the scoop occur in one timestep.

The vehicle's turning is made to reach the desired heading angle exactly, but driving stops when the position is less than or equal to 1.1 times the driving increment from the destination, a maximum of 0.18 m but average of 0.10 m. This can be thought of as a simulated driving error, though it is not random and driving paths would be repeated exactly given the same starting pose and destination. The current elevation and orientation of the vehicle is defined by the average of the 4 ground heights at the 2D wheel locations and the average plane between these heights, respectively (see Figure 3.3(b)).

Obstacles and traversability are generally not considered, as the focus here is on generating driving points to accomplish jobs within available work zones, with approaches being made from traversable ground. Later, in Section 7.1, a basic traversability check is implemented for a scooping approach method. If driving between zones, the machines are assumed to be able to check for hazards and collisions independently. As this chapter will show, however, the 3D graphical tools which are developed allow the remote human supervisor to visualize the job plan and verify that no major problems exist before work is allowed to proceed.

Tool-ground interaction works by checking for intersection between the ground surface and bottom plane of the scoop ("cutting plane" in Figure 3.3(a)). At each timestep, the forward kinematic model first calculates the position of the scoop based on the pose of the vehicle chassis and joint settings. Intersection is then assessed by discretizing the scoop plane into a grid of points and comparing each point height with its corresponding ground height above or below, illustrated in the excavation sequence in Figure 3.4. At any intersection point, the ground height is lowered and the column volume above added to the scoop load. The image shows a grid spacing in the scoop plane which is slightly less than that of the ground surface, but the spacing actually used is less than half to ensure coverage of all corresponding ground points.

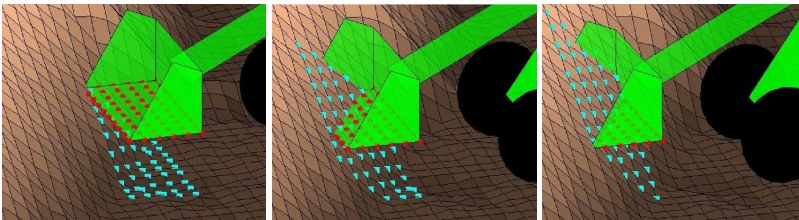


Figure 3.4. Checking for intersection of scoop cutting plane with ground during excavation sequence; red circular points are discretized cutting plane, with square blue points representing corresponding ground heights.

During a scooping action, material can be removed from the ground surface until the point at which adding the current timestep's newly penetrated material to the scoop load surpasses the scoop's volume capacity (the *full criterion*). The new material is then left in place and the scoop considered full, thus the amount of material removed by scooping actions constantly varies. With the scoop capacity of 0.15 m^3 , 0.1 m grid and 0.5 m/s driving speed, the average load extracted tends to be approximately 0.11 m^3 , or 73% full. An exception to this scoop-filling behaviour occurs if the current load is close to zero ($< 5\%$ full) and the newly-cut material surpasses the capacity in one timestep, which occasionally happens with a larger grid size of 0.2 m . In this situation, the scoop is filled exactly to capacity with the remaining material distributed over the penetrated area.

In Chapter 7, where scooping approaches are studied in more detail, a modification to the scoop filling is made whereby a certain minimum bucket fill can be specified when the full criterion is reached, with the remaining capacity filled randomly and the rest of the material left on the ground. This is so that a certain average filling effectiveness can be simulated for the scooping controller which is assumed to exist, but whose performance is unknown, with the randomness simulating the effects of tool-ground interaction. This is explained further in Section 7.2.

After a scooping action, an ideal ground behaviour maintains a maximum stable angle of repose at any locations where ground heights have been modified, simulating a homogeneous, granular material. Most granular materials have a repose angle in the range of $25\text{--}45^\circ$, with lower values for smoother spherical particles and higher values for rougher particles [89]. In this chapter 45° is used, though 30° is also used later in Chapter 7. Figure 3.5 illustrates the behaviour of the ground during pile loading without (at left) and with (at right) this resettling feature.

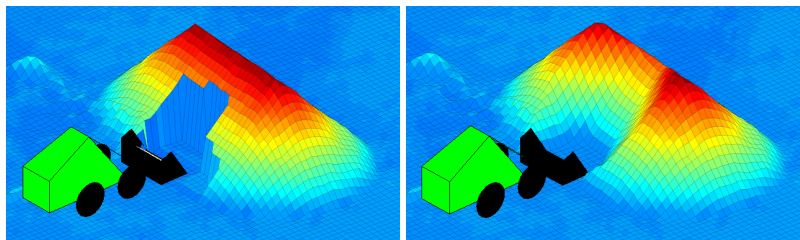


Figure 3.5. Pile loading without (*left*) and with (*right*) soil resettling behaviour.

The soil resettling behaviour is achieved by scanning in the X and Y di-

rections in the vicinity of a scooping action, and searching for segments steeper than the stable limit. To reduce the amount of computing required, a segment is only considered unstable if it is greater than 1.2 times the repose angle being used, making this the effective “static” repose angle. If a steep segment is found (as in Figure 3.6(a)), a recursive algorithm adjusts it to the nominal repose angle (Figure 3.6(b)), then checks neighbouring segments downhill (Figure 3.6(c)) and uphill until overall stability is achieved (Figure 3.6(d)). To reduce computation again, during this search phase a segment is only considered unstable if it is greater than 1.1 times the repose angle, making this the effective “dynamic” value. This algorithm is described in more detail in the previous work [17].

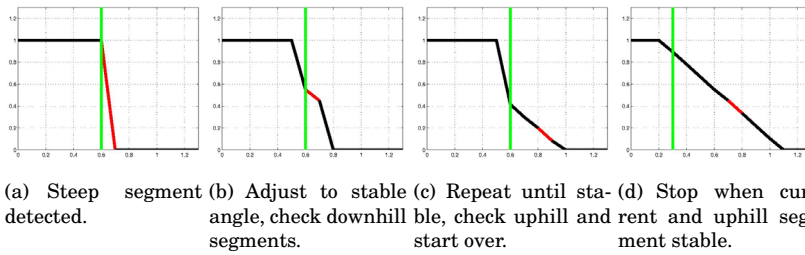


Figure 3.6. Simulated soil behaviour: recursive algorithm adjusts steep segments and checks neighbouring segments until stability is achieved.

The values for the effective static and dynamic repose angles (1.2 and 1.1 times the nominal) were chosen arbitrarily, however it is generally the case for granular materials that a static value exists which is higher than the dynamic [89]. Since this resettling behaviour only scans in the X and Y directions, it is limited in that the angle to the next diagonal cell could be higher than the stable limit.

Depositing material on the ground with the loader is simulated by adding the load volume to the surface below in the shape of the standard pile used earlier (see Figure 3.2). This pile has a central width equal to the width of the bucket, with the crest located beneath the blade, and a conical radius (and consequently, height) initially determined by the load volume. The new pile is first provisionally added to the surface model at the lowest elevation encountered in the new pile’s footprint. If the ground beneath happens to be flat, as at the left in Figure 3.7, then the provisional pile corresponds to the scoop load volume, and it becomes part of the surface model.

In case the ground below is not flat, an algorithm finds a pile-shaped surface with the proper dimensions which adds the new volume while ac-

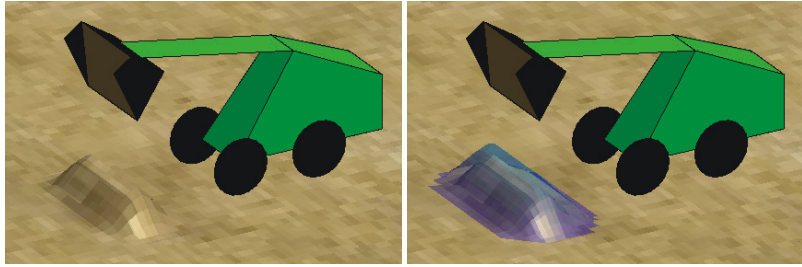


Figure 3.7. Unloading of bucket: (*left*) one scoop-load deposited in pile shape on flat ground beneath bucket blade; (*right*) 2nd scoop-load added by finding pile-shaped surface which adds new volume and accounts for volume beneath.

counting for the ground material beneath, as the right of Figure 3.7 shows when a second load is deposited over the first. This algorithm gradually increases the conical radius (and height) of the pile shape until the volume it encompasses equals the load volume plus the volume of the material already beneath, within a close enough margin. At the end a small corrective layer is added or subtracted to conserve the total volume, after which the slope resettling algorithm is applied to check for repose angle stability.

Checking how closely this simulator emulates real earthmoving work with sand or gravel remains an area for future work. This could be done by using scoop positioning data from hardware experiments, and also ground models from before and after scooping and dumping actions. The situations could then be recreated in the simulator, and differences between the real and virtual ground models before and after these actions compared to see if the ground behaviour was accurate. As it is, the ground behaviour in the simulator likely diverges somewhat from the real world, but it is assumed to be sufficient for testing the high-level job planning strategies developed in this thesis.

3.2 Robotic Machine Capabilities

One key assumption in the simulator is that a current surface model of the worksite is available at all times. In a real world situation, updated models of the whole site may be infrequent, although a local scan made by the machine after each scooping and dumping action is conceivable. This could be provided by a top-mounted 360° or perhaps a forward-facing laser scanner, or a combination of sensors. Another possibility could be to have a second “surveyor” robot would be on site to make updated models, however it would likely be preferable if sensors on the loader could make

sufficient updates during the work cycle.

Laser rangefinding sensors suffer from noise and occlusion, and high-quality variants can be expensive. Given current trends however, it is foreseeable that as robotic systems enter more industrial sectors, such environmental sensors will continue to increase in quality while decreasing in cost. A historical example is the Inertial Measurement Unit (IMU), which experienced a steep reduction in price over the last few decades. Whereas high-quality electro-mechanical IMUs were once prohibitively expensive, solid state and fibre-optic versions are now readily available to robotics researchers.

Another trend being seen today is the entry of sensors and platforms traditionally used for research purposes into the consumer market. These include ranging sensors for controlling video games and multirotor aerial drones which can be remotely controlled or even fly autonomously to follow and film their owner. One related construction application could be for these aerial platforms to provide surface models at worksites.

Another assumption in the simulator is that the machine is able to navigate and drive autonomously, which is already becoming commonplace with the help of high accuracy positioning data available via satellites. In scenarios where satellite positioning is not available such as underground mines or planetary surfaces, other navigation methods could be used such as landmark detection, dead reckoning, inertial sensing or, if available, a local positioning system deployed at the site.

In the heavy machine industry, positioning and environmental sensors may become more commonplace for mobile machines. These need not only be used by robotic machines but can also assist human drivers to follow plans more accurately. Future challenges will be Artificial Intelligence (AI) problems such as detecting and interacting with humans, dealing with even too much sensor information, and performing highly skilled tasks such as excavation, which may become commonly programmed via teaching by experienced operators. Even if robotic construction machines eventually match or exceed the capabilities of human operators, telling the machines what to do and monitoring their work will still be necessary, which is one contribution presented in this thesis.

3.3 Interactive 3D Graphical Planning Approach

The approach for implementing the supervisory control system was to begin with a 3D surface model of a worksite, which is assumed to be updated by ranging sensors and therefore represent the current state. The remote human supervisor plans a job by specifying areas and approach directions using interactive 3D graphical objects such as rectangular surfaces and prisms rendered over the site model, which are manipulated by clicking and dragging with a mouse. Since the plans can be made on a recently-acquired surface model, jobs can be configured ad-hoc without a prior global map or knowledge regarding the locations of piles.

An automated planner then interprets the high-level plan and, based on the surface model, generates low-level plans such as driving destinations which the machine(s) are assumed to be able to follow. These are also represented graphically, allowing the user to visualize and approve the plan before work begins. The machines follow these plans and interact with the worksite, with changes detected by the ranging sensors and the updated ground model used by the automated planner to generate the next commands. This process represents a *job-level* control loop, which is represented in Figure 3.8.

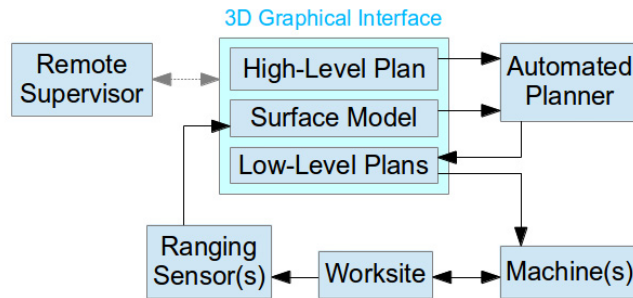


Figure 3.8. Job-level control loop using 3D graphical interface.

A key part of the loop is the 3D graphical interface which consists of the surface model with the high- and low-level planning objects. The remote supervisor can interact with the interface at various levels, ideally only specifying the high-level plan initially and then monitoring progress via updates to the surface model as the loop operates autonomously. Lower-level plans may also be modified by the user if necessary. The goal of the 3D graphical approach was to make the planning tools intuitive and easy to use, using a click-and-drag functionality familiar to desktop computer

users for the spatial management of icons and windows.

This interface can also be considered a type of *common presence* between the human operator and robotic machine. This is a concept proposed by Halme as a basis for developing a Human-Robot Interface (HRI), consisting of a virtual model of the environment which combines a geographical map with information related to various objects within it [90]. The common presence is used to transfer task-related knowledge between the robot and human, with the goal being to transfer the will of the user to the machine in a natural manner using symbolic information as much as possible.

The 3D graphical planning approach presented here could be used in several scenarios. One example is remote job specification and supervision, where limited information from the worksite is available and the human supervisors rely mostly on visual feedback from the site which helps to interpret 3D surface models on which the job is planned. If the site is well-characterized and robotic machines are following high-level plans already specified in a CAD model for example, displaying lower-level plans graphically could still be useful for monitoring purposes, so that the intentions of the robotic machines can be checked for potential problems.

Another use case could be where the human supervisor or foreman is standing beside a machine at a site, and wishes to specify a small- to medium-scale job. Even if it would be possible to plan the job from an indoor control station on a CAD model of the site, it may be more convenient to specify the job on the spot using a tablet computer, for example, by configuring the job visually on a surface model of the site. In this case an Augmented Reality interface may make the process easier, with concepts for this presented later in Section 4.3. The automatically generated plan would then be checked and work could begin.

3.4 Pile Transfer Job

This section presents the strategies that were developed to allow a remote human supervisor to specify a high-level plan for a pile transfer job, and to automatically interpret this plan to generate commands for a robotic wheel loader to follow. The main goals were to keep the method flexible, so it could be used for various worksite layouts, and to minimize the input required from the supervisor after the plan is made. Ideally the loader op-

erates autonomously for the duration of the job, yet if problems occur, the human remains in the loop and can provide further input, which remains above the level of direct teleoperation.

3.4.1 Source Specification

The first part of planning the job is to specify the location of the source material. The approach taken was to indicate an area similar to how humans may communicate a job to each other, i.e. by saying “take material from there,” or “remove that pile.” For this purpose the Area Tool was developed, a rectangular surface rendered over the worksite model that can be translated by clicking and dragging (Figure 3.9(a) to 3.9(b)), and resized by clicking and dragging one of the two corner cones (Figures 3.9(b) to 3.9(d)). In Figure 3.9(d) it is used to fully blanket a source pile, with dimensions 6.2 x 4.0 m. Appendix B contains more details about the click and drag mouse interface.

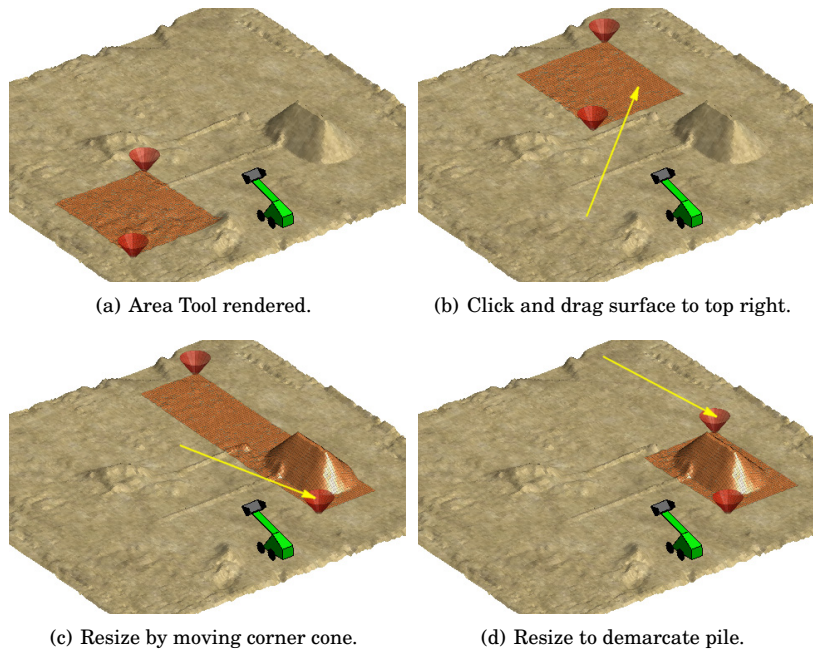


Figure 3.9. Area Tool, used to demarcate source pile; arrows indicate clicking and dragging with mouse pointer.

After the source pile is specified, two important values are obtained: the “ground level” surrounding the pile and the volume of the pile. Ground level is determined by creating a histogram of point elevations from the ground model grid section demarcated by the Area Tool, such as in Figure 3.10, with elevations grouped into intervals of 0.01 m or 0.02 m. With a

relatively flat ground, one interval will usually have more instances than the rest, which is assumed to be the ground level (0.25 m in this example). In order to find this level easily, it is important to include enough of the surrounding “flat ground” when specifying the pile (in Figure 3.9(d)).

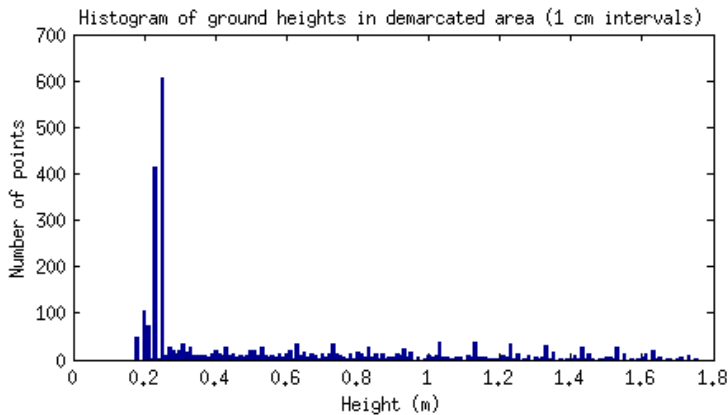


Figure 3.10. Histogram of ground elevations in specified area (from Figure 3.9(d)); highest bar assumed to be at ground level.

Only the bottom half of the histogram is evaluated, in case a pile with a flat top or plateau section were being specified, which might have a large number of points with the same elevation at the top of the slope. Once ground level is known, the volume of the pile is estimated by subtracting this value from each point in the discretized area, constructing a column at that point with the grid resolution base area (here 0.1×0.1 m), then summing the individual columns. Here the pile volume is found to be 7.65 m^3 .

One limitation of this histogram method is that it requires the ground to be flat. If this is not the case, and the pile rests on a sloped surface, an equation for the ground plane would be needed. One way to obtain this could be by selecting a portion of the ground plane using the Area Tool, then using coordinates from this area to estimate a plane equation using a least squares technique. This method is used later during the outdoor tests (see Section 4.1.1). A more sophisticated alternative could be a plane-fitting technique using RANdom SAMple Consensus (RANSAC) to find the plane surrounding the pile in the same surface section which blankets the pile [91]. Once a plane equation is obtained, the column heights needed to estimate the pile volume would be obtained by subtracting corresponding points on the plane from the surface heights.

3.4.2 Dump Pile Specification

The next stage in planning the pile transfer job is to specify the filling or dumping area. Only specifying the area where to dump the material was not considered sufficient for planning the job, since it does not allow the height of the pile to be visualized. Depending on the job, the operator may wish to control the height of the dump pile: to spread the material in a thin layer over a large area for example, or to make a high pile that occupies less space.

The pile heights that are possible to deposit with the wheel loader are assumed to range from a minimum, the height of a full scoop-load in the shape of the standard pile (0.346 m), to a maximum which the machine can build up at one location before it begins colliding with the pile (illustrated in Figure 3.11). In case a driving error exists, the maximum should be reduced so that the desired height can be achieved without driving into the slope, indicated by the dotted lines in the figure. Here the maximum height is set at 0.85 m. The “dump offset,” labelled in the figure, is the distance from the vehicle reference point to the point beneath the bucket blade in the dumping configuration. In a real system this would need to be determined experimentally, since another factor possibly affecting where the material falls to the ground is the angle of the bucket.

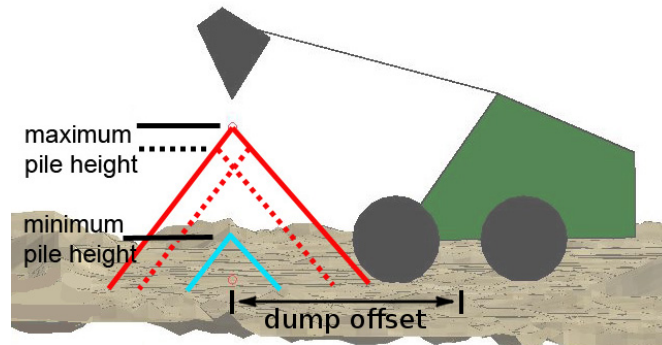


Figure 3.11. Dump pile height limits based on volume of one scoop-load (minimum) and geometry of scoop position and front wheel (maximum), with maximum reduced for driving error; dump offset is between scoop position and vehicle reference point.

As a way for the operator to visualize both the area *and* height of the planned dump pile, the Virtual Pile tool was developed. This consists of a surface with the shape of an ideal pile added onto the local ground surface (see Figure 3.12). The ideal pile has a central rectangular section with width w , length l and height h , surrounded by slopes with angle θ (here

45°), and a quarter-conical section at each corner with base radius r .¹ For the pile transfer job, the volume of the Virtual Pile is usually sized according to the volume of the source pile specified earlier, though other values can also be set.

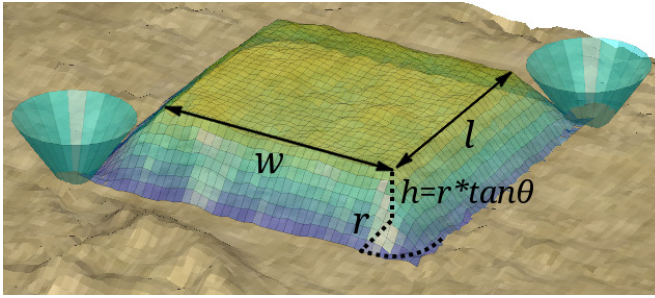


Figure 3.12. Virtual Pile with slopes of angle θ (here 45°), used to specify goal state for earthmoving job.

The Virtual Pile works much the same way as the Area Tool, and can be moved around the ground model by clicking and dragging, and resized by moving one of the two corner cones. By moving a corner cone, a new length and width is specified and the height is adjusted automatically, if necessary, to keep the total volume V constant. The new height is determined using the relation $h = r \tan \theta$ and solving the equation

$$V = \pi r^2 h / 3 + whr + lhr + wlh, \quad (3.2)$$

which is the sum of the individual Virtual Pile component volumes (cone comprised of four corner sections, triangular prisms at four sides and central rectangular prism). The height is limited by the minimum and maximum values described above, and an additional limitation is that either the width or length of the central section must be at least the width of the scoop, so that the loader could theoretically create the pile. All piles within these limits can then be specified by moving a corner cone. The functionality of the tool is illustrated in Figure 3.13, with additional details about the mouse interface presented in Appendix C.

A further consideration when determining if a certain pile shape is possible to deposit with the loader is the fact that since the scoop width is oriented perpendicular to the driving direction, material can only be deposited width-wise along increments of this distance. The width along which the material is distributed on the ground then depends on the cur-

¹The term *length* will generally be used throughout this thesis to describe the forward horizontal distance away from the observer, which could also be considered the *depth*.

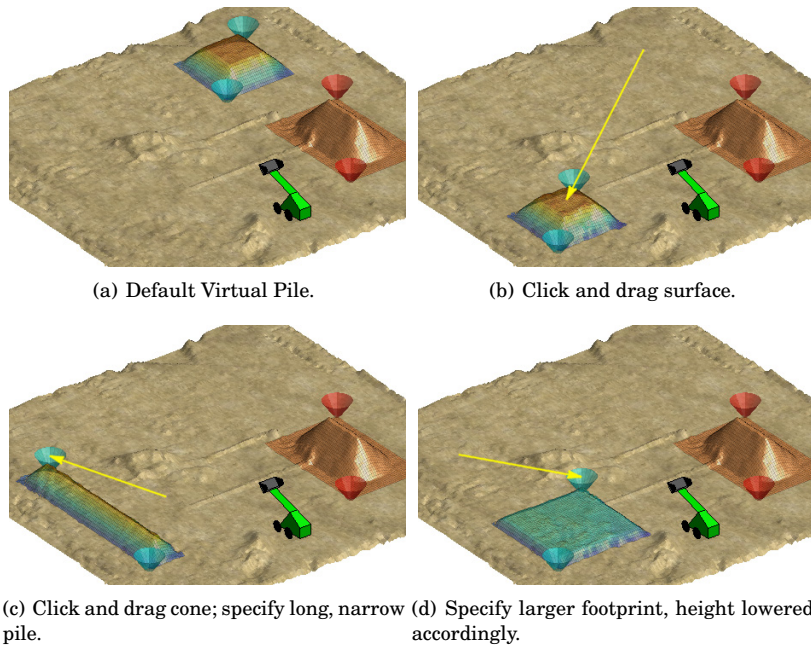


Figure 3.13. Dump pile specification; arrows indicate clicking and dragging with mouse pointer.

rent shape of the ground beneath. For high piles with side slopes wider than a scoop-width, the total width could be gradually increased by dumping on the slope. For low piles with slopes narrower than a scoop width, the total width can only be increased by increments on the order of the scoop width. Certain pile shapes could then be ruled out by taking this into consideration, however it would be difficult to know with certainty since the behaviour of the ground material when deposited may be unpredictable. As Section 3.4.5 will show, the Virtual Pile is not meant to be created perfectly by the loader, but rather to serve as a guide to be filled in approximately.

3.4.3 Automatic Interpretation of Job Plan

After the source material and dump pile are specified, the rest of the job is planned automatically and displayed to the user. Figure 3.14 shows how a basic high-level plan is interpreted here, with the source pile at right and Virtual Pile at left. In this figure the corner cones for both surfaces are omitted, while some of the planning and simulation control buttons are shown at the bottom left corner of the figure.

Such a plan could be interpreted in different ways, but to reduce unnecessary driving, each area should generally be approached from the side

closest to the other area, as long as no obstacles are present. One constraint regarding the dumping of material is the fact that deposited material should not hinder the machine, thus a specified area should generally be filled from the back to front.

Different ways of generating scooping commands are possible. Some strategies mentioned in Section 2.5 included aiming towards the centroid of the pile and adjusting the approach angle to be perpendicular to the bottom contour, or assessing all possible approaches along the bottom contour. Here a simple alternative strategy is introduced whereby all scooping actions begin and end at a point adjacent to the source area. This “Stage” point is the first lower-level plan component to determine. Four possible points are considered, off the midpoint of each side of the source area (marked with the red rings in the figure). The default distance of these points from the edge of the area is 2 m, which allows about 0.5 m of extra space between the scoop blade and the area edge when the loader is in scooping position. Of these four possible points, the one closest to the Virtual Pile centre is chosen and represented by the inverted cone A.

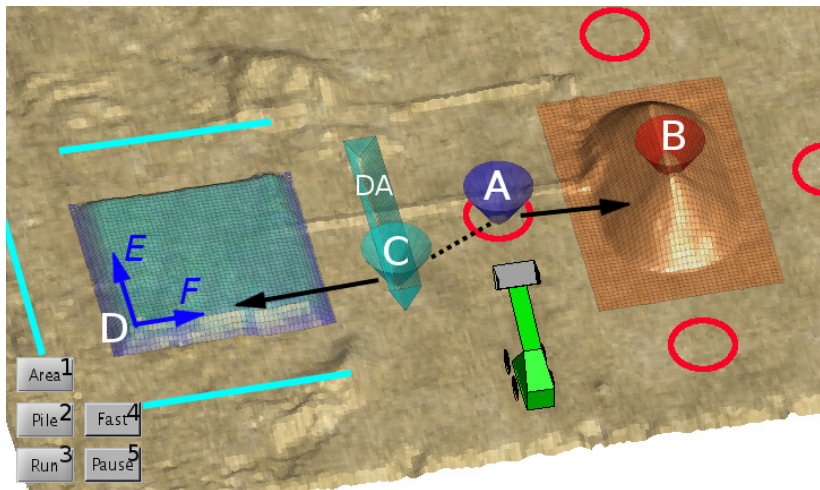


Figure 3.14. Automated lower-level plan generation: A - Stage point for scooping; B - Scooping Destination; DA - Dump Approach side; C - Dump Stage point; D - Dumping Location; E and F - directions to next dumping locations. Buttons for planning and simulation: 1 - render Area Tool; 2 - render Virtual Pile; 3 - start; 4 - speed up simulation; 5 - pause.

The scooping strategy employed was to begin all scooping actions at the Stage point, lower the bucket to a preset position and drive towards the Scooping Destination point located at the highest point in the pile (marked by cone B). A bucket-filling controller is assumed to exist which governs when the loader would stop driving and extract the bucket, after which the loader reverses back to the Stage point.

As explained in Section 3.1, with the default scoop filling behaviour in the simulator, the bucket is not filled to capacity each time but is extracted when the next penetrated volume increment would exceed its capacity. The machine is assumed to have a constantly updated ground model available, thus the Scooping Destination is updated after each scooping action in case the highest point in the pile has shifted. This strategy is discussed further in the next section.

The next step is to determine the Dump Approach side from which to approach the Virtual Pile for filling actions. All four sides around the surface are considered (marked with the blue lines in Figure 3.14), and the one closest to the source pile is chosen and represented with an inverted triangular prism (labelled “DA”). This is located 1.5 m from the edge of the surface, which is more than the minimum needed for depositing material at the closest possible location at the front. The filling strategy is to dump material from left to right (when facing the Virtual Pile) along vector E, and from back to front along vector F. The first Dumping Location is therefore at point D, and is approached from the Dump Stage point (marked with cone C) along a path perpendicular to the Dump Approach side. Points D and C are updated after each dumping action as the Virtual Pile is being filled in (described in Section 3.4.5).

All the points needed for the job are then defined, and after approving the plan the operator can command work to begin. During the simulation a finite state machine controls the driving between points, guiding the loader to repeat the Load-Haul-Dump cycle. This cycle is represented by point order A-B-A-C-D-C-A, with the machine driving forwards or backwards where appropriate. The job ends when either the source area is within a specified margin of ground level, or when the Virtual Pile is filled up and no dumping locations are available.

Since the volume of the Virtual Pile is based on the source volume, the standard goal is that either job completion criterion would occur around the same time. To ensure that the source area is levelled, a larger Virtual Pile volume can be specified to provide extra filling volume, whereas if a smaller dump pile is required, some source material may remain uncollected.

The job planning strategies presented here are kept simple but allow for plans to be generated which could function for various initial worksite layouts. One limitation is that only rectangular areas can be specified for scooping and dumping, with an arbitrary orthogonal orientation on the

worksite. The plan is also specific for a skid-steered machine which is able to turn on the spot at points A and C. The following sections will further describe the scooping and dumping portions of the job, and show the results from complete simulations.

3.4.4 High Point Scooping Method

As mentioned in the previous section, the automated scooping strategy used here is to begin each action at the Stage point and scoop towards the highest point in the pile, resulting in a fan-shaped coverage pattern over time as the highest point shifts due to slope collapse. Scooping commands are generated as long as any ground heights in the source area remain a certain threshold above ground level, guiding a job to completion.

A possible disadvantage of this strategy is that since it does not select approach vectors based on the shape of the bottom edge of the pile, or *Zero Contour*, it can result in non-perpendicular approaches into the pile. It also does not consider factors such as contour convexity along the approach path and asymmetrical side loading of the bucket. Other proposed scooping approach methods consider these factors as criteria for improving bucket-filling effectiveness and reducing mechanical wear (see Section 2.5) [57, 64, 69]. Although some experimental results have found that sideload and side moment measures had little relevance to the resulting bucket fill, it was also noted that symmetrical side loading is still important to avoid spillage when hauling [71].

The strategy used here, which will be called the High Point (HP) method, was not intended to be an improvement over others that have been proposed, but was initially developed as a simple way to generate commands for full pile transfer simulations in order to test the job planning tools. In case scoop filling using this method could work well enough in practice, however, which would need to be tested experimentally, it may offer some advantages.

The main possible advantage could be reduced total driving distance from a given starting point near the slope, due to the limited driving pattern of straight lines originating from the Stage point. Using a Zero Contour method, on the other hand, a loader might need to drive more around the edge of a pile to reach its chosen approach vectors. This will be investigated further in Chapter 7 by comparing simulation results using both the HP method and a Zero Contour method which was implemented.

Another possible advantage of the HP method is that the limited fan-

shaped driving envelope of the machine, which only enters a rectangular pile area from one side, may be desirable in confined spaces and/or if humans or other machines are working nearby. Furthermore, assuming that the Stage point is located on traversable ground, which would be verified during the planning phase, further traversability assessments during scooping actions are generally not necessary since the pile area is only entered by driving forward with the scoop lowered. Other scooping approach methods, however, may generate approach drives which involve maneuvering inside a pile area before it is fully cleared. This could increase the risk of non-traversable driving paths being chosen, thereby setting higher requirements for a robust traversability assessment capability.

Here the pile loading portion of the high-level plan made previously in Figure 3.14 is used to demonstrate the HP method. Figure 3.15(a) shows the end of the first loading action into the 7.65 m³ pile. Since the development of scoop control for optimal filling is not part of this thesis, the scooping action is simplified by keeping the cutting blade at a preset configuration, low and level, during the entire scooping action. In this simulator the scoop is able to penetrate the pile in this way until the full criterion is reached (see Section 3.1).

This technique was also used in Sections 5.2-5.4 during outdoor gravel loading tests, and was found to function for filling the scoop to a degree which was considered sufficient, though the load was not measured. In related work which has focused on developing scooping controllers, the scoop usually enters the pile horizontally, after which it begins rotating upwards while perhaps simultaneously following an upward ramp, before the machine stops to extract it vertically [67, 70, 63].

In the simulator, although the bucket is not generally filled to capacity when the full criterion is reached, the loader does not make another filling attempt and unloads at the Virtual Pile as described in Section 3.4.3 before returning for the next action. This could be because it does not have a direct load measuring capability, and/or because the average fill ratio achieved is considered acceptable for one action.

Returning to the loading simulation, Figure 3.15(b) shows how the machine initially focuses its efforts on the central section of the pile as the highest point shifts along the crest, forming the pile into a crescent shape. Highest point updates occur after the loader has extracted a load and reversed back to the Stage point, simulating the scanning of the pile with an onboard forward-facing laser scanner, for example. The Scooping Destina-

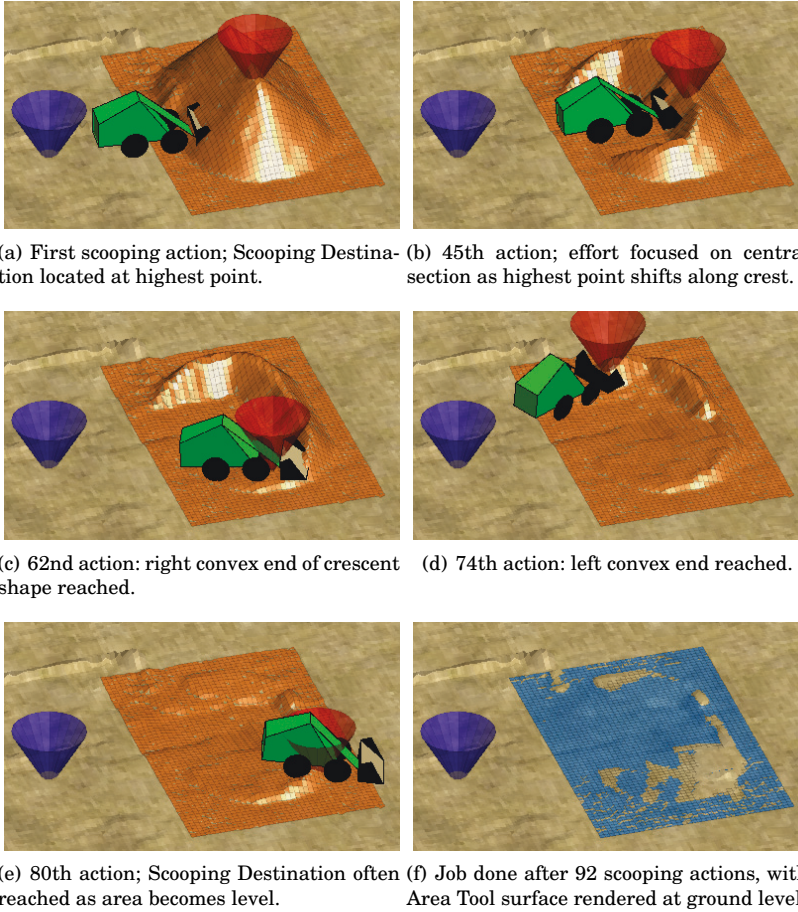


Figure 3.15. Pile loading using the High Point (HP) method.

tion cone's position is also updated to display the loader's next planned action to the remote user.

Eventually the convex ends of the crescent shape are collected, as Figures 3.15(c) and 3.15(d) show. As the source area becomes more level, with less material loaded by the bucket, the Scooping Destination is sometimes reached by the loader with the bucket filled under capacity, as in Figure 3.15(e). When this occurs, the machine begins the next scooping action after reversing to the Stage point, without unloading the material at the Virtual Pile until the bucket's full criterion is reached. Figure 3.15(f) shows the state of the workspace when the job is complete, after 92 actions. The criterion for completion was when all ground heights in the specified area were less than a threshold of 0.15 m above the initial ground level, found using the histogram technique from Section 3.4.1. In this last figure the source area surface is rendered as a flat plane at

ground level, to illustrate how some areas ended up above ground level and some below.

Of the 92 scooping actions made, 75 resulted in a full scoop-load of material being extracted, with an average load of 0.114 m^3 , while the other 17 resulted in partial scoop-loads. The total volume loaded was 8.59 m^3 , which is 12% higher than the estimated volume above ground level of 7.65 m^3 . Most of this extra volume was collected due to the scoop digging too deep in the source area, evident in Figure 3.15(f). Some of this extra volume is also collected outside the area when the loader reaches the Scooping Destination as in Figure 3.15(e), and the scoop ends up past the edge of the area.

The loader attempts to level the area at ground level by adjusting the extension of the prismatic boom joint, and hence the scoop height, used for a scooping action based on feedback from the previous scooping action. Figure 3.16 shows the range of possible scooping configurations, in which both rotary joints (at the base and end of the boom) are set to zero, orienting the bottom of the bucket horizontally. The boom extension ranges from 0 m to 0.24 m, which corresponds to placing the cutting plane 0.170 m to 0.010 m above ground level, respectively. The default boom starting extension is 0.10 m (bucket 0.103 m above the ground).

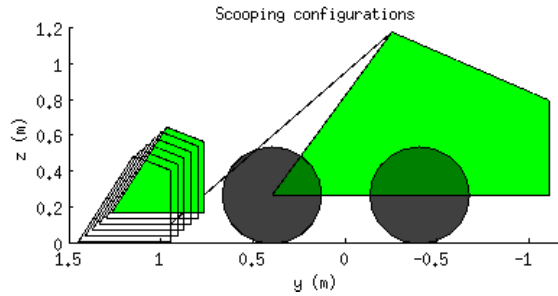


Figure 3.16. Range of scooping configurations resulting from adjustment of boom prismatic joint.

Algorithm 1 performs the scoop configuration adjustment by first determining *endHeight*, which in most cases is the height of the middle of the scoop blade's front edge at the end of a scooping action, when the full criterion has been reached. Assuming the loader knows its 3D chassis position and joint angles, this is found using the forward kinematics model. If the Scooping Destination has been reached with the bucket under capacity, *endHeight* is the resulting ground height at the target Scooping Destination, determined from the updated ground model after the loader has reversed from the pile. The reason for these two different *endHeight*

possibilities is that in the latter case, the bucket is not at the Scooping Destination but has moved past it, to ensure clearing all material there. In case the ground height at the target location remains too high, the bucket height can be adjusted properly for the next attempt.

Algorithm 1 Automated adjustment of boom extension to keep bucket at ground level while scooping.

```

if bucketFull                                % result of most scooping actions
    endHeight = bucketHeight;                % determine bucket height at end of action
else                                          % Scooping Destination reached by bucket
    endHeight = heightAtScoopDest;           % scan ground height at previous target
end
if endHeight > groundLevel + 0.01 % previous endHeight too high
    newExt = boomExtension + 0.01;           % extend boom 1 cm to dig deeper
    if newExt > 0.24                         % maximum extension reached
        newExt = 0.24;
    end
    boomExtension = newExt;
elseif endHeight < groundLevel - 0.01 % previous endHeight too low
    newExt = boomExtension - 0.01;           % shorten boom 1 cm to raise blade
    if newExt < 0                           % minimum extension reached
        newExt = 0;
    end
    boomExtension = newExt;
end                                           % if endHeight within 1 cm of ground level, no change

```

endHeight is then compared with ground level, and if it is above or below by some margin (here 0.01 m), the boom is extended or retracted a small amount (0.01 m), respectively, for the next scooping action. Even though the next Scooping Destination may be different than the previous one, this algorithm ensures that the bucket height will frequently be changing, always attempting to position the bucket so that the area in general ends up at the desired height.

The need to constantly adjust the scooping configuration during the job is mostly due to the nature of the simulator, which allows excavation to take place simply by the ground surface intersecting with any part of the scoop cutting plane. When driving with the scoop lowered on uneven ground, small pitch and roll motions by the chassis can cause the cutting plane to dip below the surface and consequently lower ground heights. A constant preset scoop position will therefore not work for all jobs, whereas in a real world setting, a preset configuration which lightly scrapes along the ground might generally work well.

Despite the use of Algorithm 1, simulations can still result in extra material being excavated below ground level, as in Figure 3.15(f). One way to reduce this could be to reduce the amount of driving with the scoop lowered. In this simulation, for example, the furthest distance from the Stage point to a Scooping Destination was 5.41 m, a rather long scooping

action due to the relatively large size of the workspace (6.2 x 4 m). Shortening the scooping drives can be achieved by breaking down a workspace into smaller sub-areas, which is a strategy used later in Sections 3.4.6 and 6.2.3, and in Chapter 7.

In this strategy only the boom extension is used to adjust the scooping configuration, whereas in real world applications a human driver may typically actuate all available joints during a scooping action. This was intentionally kept simple since the focus was not on developing scooping controllers, and the boom extension strategy worked reasonably well in the simulations. As an area for future work, however, it may be worth investigating if such a strategy, where only one parameter is adjusted, could work to regulate the elevation of the workspace.

One drawback of the HP method in general is the situation towards the end of the job, i.e. the distribution of the pile remnants in a crescent shape. This is likely undesirable since by spreading the material out it becomes more difficult to load. For larger jobs where there is usually a slope to push against, however, this would not be as much of a concern. The simulations in Chapter 7, which are used to evaluate the possible benefit of the HP method, are therefore based mostly on larger slope loading jobs.

A possible limitation of the HP method is that the bucket width may need to be at least as wide as the chassis (including the tires), so that the bucket could potentially clear all material in its path and prevent the wheels from driving up a slope adjacent to the path. This is discussed further in Section 3.4.7, with a solution that could work if the scoop was narrower than the full chassis width.

An additional capability which was developed, but not used in this example, was a *traded control* scheme (described by Sheridan [13]) which allows the remote user to reposition the Scooping Destination by clicking and dragging its cone, in case further control over the scooping strategy is desired. This causes the cone to remain static, with scooping actions repeated towards it either until the Scooping Destination is reached by the machine or until the user moves the cone or reverts to automatic updates. This capability is used in the outdoor pile transfer tests in Sections 5.3 and 5.4.

3.4.5 Area Filling using Virtual Pile

The area filling strategy, mentioned in Section 3.4.3, is to fill in the Virtual Pile from left to right (from the loader's perspective), starting along the

back row and working towards the front. Figure 3.18 shows the start of the filling portion of the job from the previous section, illustrating how the next dumping location is found using the Virtual Pile with a base width of 4.8 m and height of 0.59 m. The basic method is to scan along a line with the approximate width of the scoop, comparing the current ground heights (square black points) with the desired heights of the Virtual Pile (round blue points). If all of the desired heights are above their corresponding ground heights, the location is considered available.

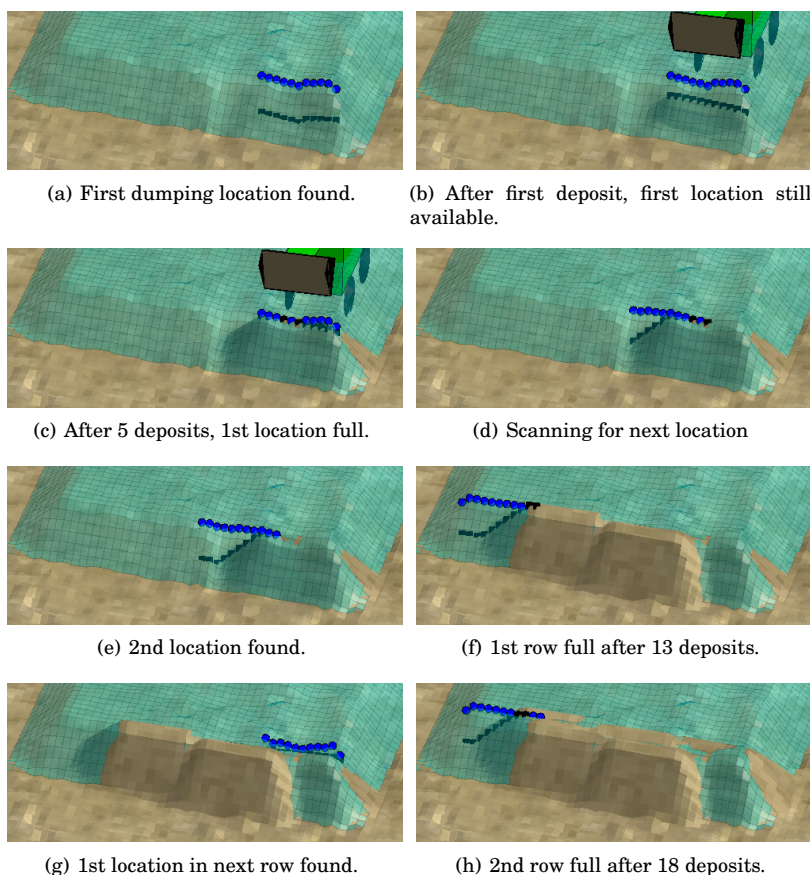


Figure 3.17. Scanning Virtual Pile for next dumping location, viewed from far end.

The main part of the scanning line used had a width of 0.8 m, rounded down from the scoop width of 0.89 m to ensure that the scoop blade could fully cover it. This portion begins at the back corner of the central rectangular section of the Virtual Pile. An extra point was then added to each end, making the line 1.0 m long, so that there would be some more lateral spacing between dumping locations.

Figure 3.17(a) shows the first scanning location, and since all desired

heights are currently above the actual heights, the location is free and the machine deposits the first scoop-load there. Afterwards, Figure 3.17(b) shows how the ground points have risen, but are still below the desired heights, making the location still available. Normally, scanning for the next location occurs after the loader has reversed from the Dumping Location (similar to the scooping procedure), simulating the use of a forward-facing rangefinding sensor. After the 5th deposit at the first location, some ground heights have risen above the desired heights (Figure 3.17(c)), thus the line begins to scan for the next location (Figure 3.17(d)), moving along increments of one grid point.

In Figure 3.17(e) the next location has been found for the 6th deposit, and the process is repeated until the row has filled up after 13 deposits (Figure 3.17(f)). The scanning line then moves 2 grid points (0.2 m) towards the front and searches from the initial side to find the next free location (Figure 3.17(g)). After 5 more deposits the 2nd row fills up (Figure 3.17(h)), requiring fewer loads than the first row because the material is being dumped on the existing slope. As the area is being filled, the Virtual Pile surface is kept constant by rendering it over the initial ground height matrix, meaning the plan can also be modified during the job relative to the original state.

Figure 3.18(a) shows the end state, with the Virtual Pile having a base length of 4.3 m and a volume of 9.18 m^3 , which was 1.2 times the estimated volume of the pile. This extra volume was added to ensure that enough space was available so that the job would end when the source area was levelled, without the Virtual Pile running out of dumping locations. This 20% margin also helps to account for extra material which may be excavated. Comparing the desired state represented by the Virtual Pile with the end state in Figure 3.18(a), it is evident that the filling was not perfect, with a somewhat uneven distribution and noticeable lack of material along the right edge (from the loader's perspective). This was apparently due to the Virtual Pile not being wide enough to allow an extra dumping location at the right. The overall shape of the dump pile, however, is reasonably close to that which was specified.

Figure 3.18(b) shows another result with a Virtual Pile 0.4 m wider (base dimensions 5.2 m wide by 4 m long), 1 cm lower (0.58 m) but with the same volume. Here the rows seem to fill more fully, but the amount of extra volume at the front appears about the same as before. Three more variations were simulated, all with the same Virtual Pile volume of

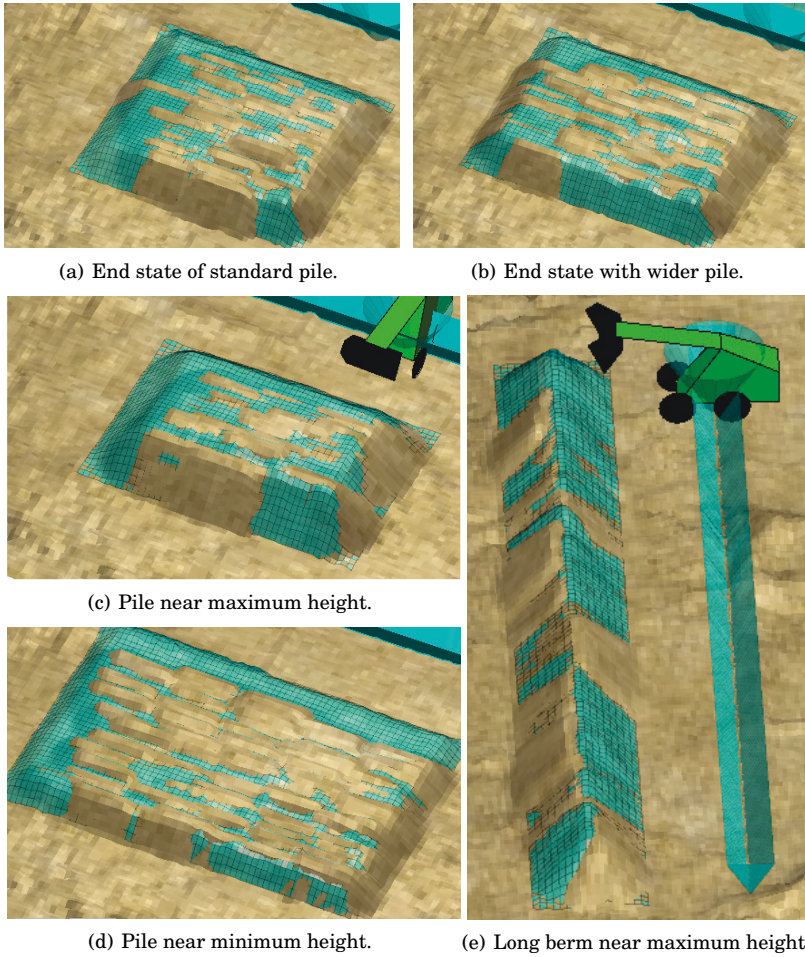


Figure 3.18. Filling Virtual Piles of various dimensions with same volume.

9.18 m³, to demonstrate how dumping locations can be found for different pile shapes. In Figure 3.18(c) the Virtual Pile was 5 m wide by 3.5 m long and 0.81 m high, close to the maximum height of 0.85 m, while the dimensions in Figure 3.18(d) were 6.6 m wide by 4.5 m long and 0.35 m high, slightly above the minimum height. The last example in Figure 3.18(e) is a 13.5 m-wide berm (1.7 m long and 0.80 m high). In relation to space applications, Chapter 1 mentioned how this type of structure could be useful for providing blast protection around planetary landing and take-off sites.

The scanning parameters which were used were chosen after some trial and error, but were then kept the same for all of these jobs. To investigate their effects, the first filling job in Figure 3.18(a) was repeated with some changes to the parameters. While each one was varied, the others were kept at their default values: 0.2 m step between scanning rows, 0.1 m

lateral step along rows, scanning line end extension of one grid point, and point threshold of 0 for a location being full.

First, the spacing between scanning rows was varied, with values of 0.1 m and 0.3 m yielding similar filling results, while 0.4 m caused the Virtual Pile to fill up before the source area was cleared of material. Next, the side step while scanning along rows was increased from 0.1 m, with values of 0.2-0.4 m yielding similar results but a step of 0.5 m causing the Virtual Pile to fill up early. The scanning line end extension was then increased from one grid point, with values of 2-5 (corresponding to total line widths of 1.2-2.8 m) giving a similar result, while a 6 point extension caused the Virtual Pile to fill up early. This parameter is limited by the height of the pile, i.e. the number of grid points that fit along the side slope.

Finally, the point threshold for determining if a location was full was increased from the default of zero. Thresholds of 1-2 yielded a similar filling result as the default, while higher thresholds of 3-8 caused more material to be loaded at the rear of the Virtual Pile, since the same locations were used for a longer time. Consequently, for this range of thresholds the material at the back rose significantly above the desired height while an increasing amount of space was left empty at the front of the Virtual Pile.

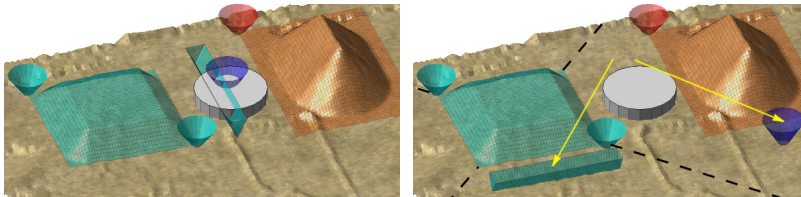
The wide berm in Figure 3.18(e) was also repeated with higher point thresholds, with values of 1-3 giving similar results but a value of 4 causing the job to fail. This was because of a problem related to the fact that the loader stops driving at an average distance of 0.10 m short of its destination. In this case the material was then dumped short of the desired location, causing a large pile to build up which the machine began colliding with before all of the desired ground heights were reached at the back of the Virtual Pile.

The dumping location search parameters were thus found to have some flexibility, with a range of possible values close to the defaults yielding similar results. In Sections 4.2 and 5.1 it will be shown how the search strategy used here is also able to find dumping locations for two manually-controlled area-filling experiments. An additional simulation in Section 6.2.3 shows how the Virtual Pile and dumping location search strategy can be extended to a much larger-scale job using a conveyor belt spreader.

3.4.6 Plan Modification and Scoop Area

Returning to the job planning phase, after a high-level plan is automatically interpreted, it must be approved by the remote user before work begins to ensure that no major problems exist. A quick glance at the objects which represent the driving destination points will reveal whether the plan makes sense. In case a problem is evident, the user may override the automatic interpreter and modify the plan. An example is illustrated in Figure 3.19.

In Figure 3.19(a) a standard pile transfer job plan has been made, with the Area Tool to the right specifying a pile with 11.64 m³ of material estimated over ground level, and Virtual Pile to the left specifying how to fill the material. The plan interpreter considers all four sides of these surfaces for staging the scooping and dumping actions and selects those which are closest to the other surface, representing them with the Stage point cone and Dump Approach side prism, respectively. These, however, conflict with an obstacle - a large cylindrical storage tank located beside the pile.



(a) Suggested high-level plan; Dump Approach side and Stage point conflict with obstacle. (b) User modifies plan by clicking and dragging objects.

Figure 3.19. Modification of job plan to avoid obstacle.

To modify the plan, the user clicks and drags the objects (represented by the arrows in Figure 3.19(b)) to new locations at the near side of their respective surface. The Stage cone follows the mouse position directly, whereas the Dump Approach prism snaps to one of the four possible sides based on the position of the mouse, which is mapped to one of the 4 sectors defined by the Virtual Pile corners (represented by the dashed lines). With the approaches to the scooping and dumping areas now free of the obstacle, work may proceed.

If work begins with the plan specified in Figure 3.19(b), a problem soon becomes evident with the High Point scooping strategy. During the pile loading in Figure 3.15, the Scooping Destination shifted frequently in the

lateral direction, resulting in a relatively balanced series of scooping actions. This is because the highest point was usually at the top of the slope section being excavated, and slope collapse at the base would soon affect the crest of the pile.

When approaching a pile from the narrow end, as in Figure 3.20, the highest point tends to stay in front of the loader for the main part of the job, resulting in a long, narrow channel through the middle of the pile. It may be possible to dig this way if the bucket is wider than the vehicle, though this situation is likely undesirable due to the potential for the bucket to get stuck or the loader to drive along a slope and tip over.

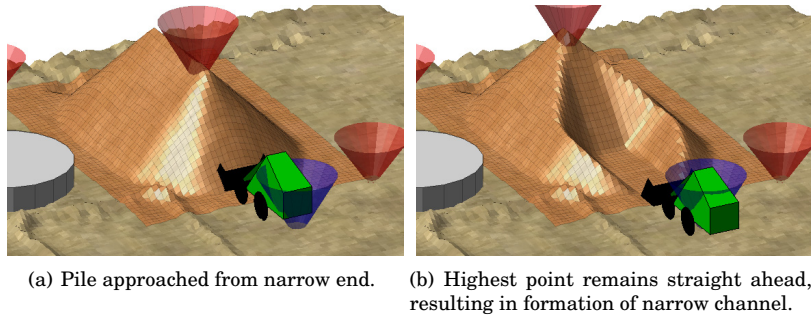
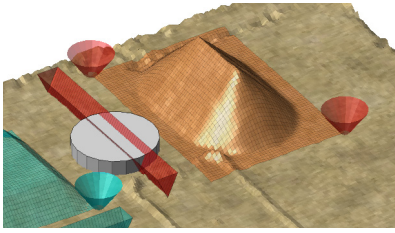


Figure 3.20. Problem with High Point scooping strategy.

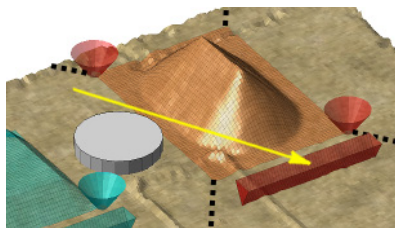
A solution to this problem was developed which is to first excavate a smaller rectangular portion of the total specified area, which will be called the Scoop Area, and which scans from front to back, stopping when ground heights at or above the 0.15 m threshold are found. Before the search can begin, the approach direction must be specified using the Scoop Approach side, an inverted triangular prism with the same functionality as the Dump Approach side used earlier. The automated planner renders this object by default on the side closest to the Virtual Pile (see Figure 3.21(a)). Since this again conflicts with the storage tank, the user clicks and drags to reposition it at the near side (Figure 3.21(b)).

With the Scooping Approach specified, the Scoop Area is then rendered over the front of the Area Tool surface, marked by the dotted line in Figure 3.21(c). Here the entire Area Tool has dimensions 4.6 m wide by 6.5 m long, with the Scoop Area having an equal width and arbitrary length of 2 m. Since the first Scoop Area location contains ground heights above the threshold, it is used to find the current Stage point and Scooping Destination, and work begins.

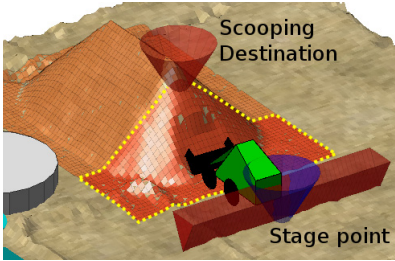
As before, the Scooping Destination is located at the highest point in the



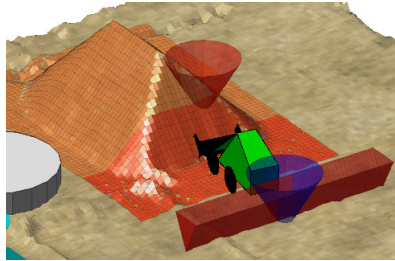
(a) Planner suggests Scoop Approach side.



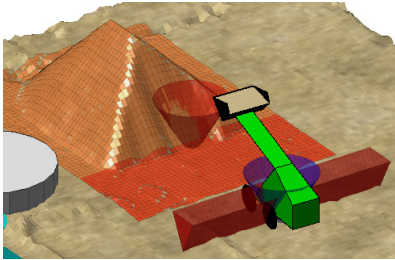
(b) User repositions by clicking and dragging; mouse position mapped to sector.



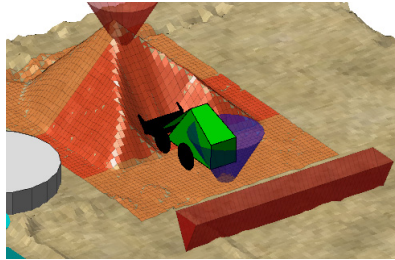
(c) Scoop Area begins at front; Scooping Destination at highest point.



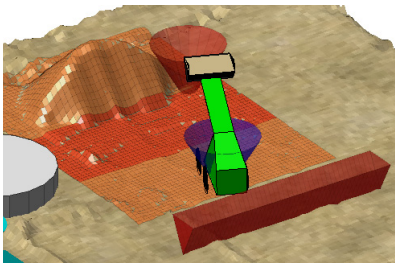
(d) Highest point shifts frequently, resulting in more balanced scooping.



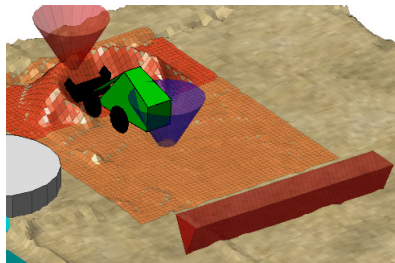
(e) 1st Scoop Area level.



(f) 2nd Scoop Area, one step back.



(g) 2nd Scoop Area level.



(h) 3rd Scoop Area.

Figure 3.21. Scoop Area used to excavate smaller portions of total surface; Scoop Approach side used to specify search direction.

area but shifts more frequently than in the problem case (Figure 3.20), resulting in more balanced scooping actions (see Figure 3.21(d)). When the Scoop Area is level (Figure 3.21(e)), it shifts one step back (Figure 3.21(f)), and the process continues (Figures 3.21(g)-3.21(h)) until the rear of the Area Tool surface is reached and the job is complete.

After extracting a load from the pile and reversing to the Stage point,

the loader continues reversing to the centre of the Scoop Approach prism before turning and driving to the Dump Stage point. This is to avoid colliding with the obstacle, and is a change that was made to the plan generation described in Section 3.4.3, where the loader drives directly between the Stage and Dump Stage points. The loader is still generally assumed to have a path planning and obstacle avoidance capability when driving between the two zones.

Figure 3.22 shows the end of the job with the last of 112 scoop-loads being hauled. 12.69 m^3 of material was removed, this time 9% higher than the pile volume estimated over ground level. The Virtual Pile, which had a 4.9 m square base and maximum height of 0.85 m, again had a volume 20% higher than the estimated source volume to ensure enough filling space was available, and as the figure shows, a short section at the front remains unfilled. The search parameters for the dumping location were the same as the default values used in Section 3.4.5.

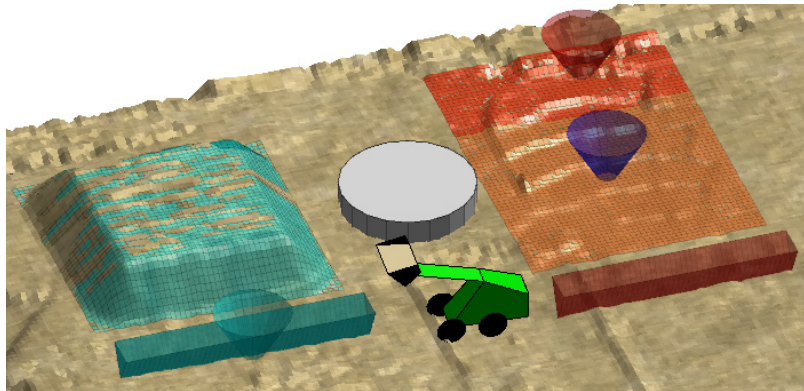


Figure 3.22. End of pile transfer simulation with obstacle; last scoop-load being hauled.

3.4.7 Discussion

This section presented interactive 3D graphical tools that were developed to allow a remote operator to specify a high-level plan for a pile transfer job, which is automatically interpreted to generate lower-level plans that a robotic wheel loader can follow to complete the job. The lower-level plans are also represented graphically and allow the remote user to monitor the intentions of the machine and modify the plans if necessary. Despite the tools being limited to specifying rectangular areas, some flexibility is provided by being able to change the approach side for scooping and unloading, as Figures 3.19 and 3.21 showed.

When the Scoop Area was introduced in Section 3.4.6, its length of 2 m was chosen arbitrarily, with its width equalling the full workspace width since this was relatively narrow. In Section 6.2.3 and Chapter 7, the Scoop Area is used for jobs with larger workspaces, requiring it to first scan along rows before moving forward. The Scoop Area dimensions which are used have important consequences for the job, since they define the amount of driving which must be done from the Stage point to clear all the associated material before scanning for the next location. This is investigated in Chapter 7 for a truck loading scenario, while the High Point scooping method presented in Section 3.4.4 is also compared with a Zero Contour method.

One factor which could have an effect on the Scoop Area dimensions to use for a particular loader is the width of the bucket in relation to the width of the chassis (including the tires). As discussed in Section 3.4.4, if the bucket is wider, then it would in theory clear a path for the wheels to drive in. If it was narrower however, then to avoid driving too far into the pile with the risk of driving into a slope, the Scoop Area should be sized such that all material can be cleared by the bucket without the wheels encountering a slope.

In the next two chapters (Sections 4.2 and 5.1), the Virtual Pile tool is tested with surface models from manually-driven filling jobs to verify that dumping locations can be correctly found using the search algorithm. The planning tools are also demonstrated in Sections 5.3 and 5.4 by configuring outdoor computer-controlled pile transfer jobs. First, however, the next section presents the use of the planning tools and a search algorithm developed for area clearing jobs.

3.5 Area Clearing Job

In the area clearing job, a layer of material in a rectangular area is cleared by the wheel loader using a preset scoop position. Here simulated snow clearing scenarios are presented, where the snow covers the ground in a uniform layer.

Although this is not the main type of job that is performed with a bucket attachment, it was included because it was thought to be easier to demonstrate with the available hardware in the near-term, before attempting pile transfer jobs with granular material. One reason snow clearing should be easier is that since it should not provide much resistance nor cause

damage to the machine, it should be possible to clear a path without much risk simply by driving with a preset scoop position. Thus for the time being, tool-ground interaction would not be a concern.

Another reason was that each clearing drive should result in a significant topographical change in the workspace, which should be easier to detect with a laser rangefinder than changes resulting from one pile loading or dumping action. The snow clearing job could therefore be used as an initial proof-of-concept of the job planning methods developed.

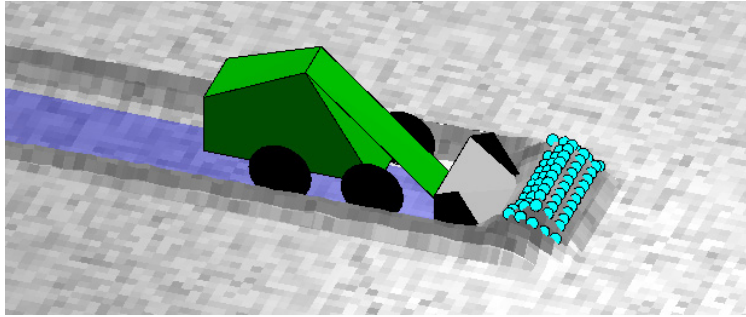


Figure 3.23. Simulated snow ploughing with bucket: excess volume distributed mostly in front, some to the sides.

The same Matlab-based simulator was used as with the pile transfer jobs, with some modifications made to model the snow layer and its behaviour. Here, the machine drives on a constant flat ground surface, while the bucket interacts with a 2nd surface above representing the snow. In the preset clearing configuration, the bucket blade is effectively at the driving surface level (less than 0.1 mm above), simulating a scraping action along the ground. When bucket-snow intersection is detected, volume is added to the bucket load as before and the snow layer heights reset. The surface grid resolution is again 0.1 m. Snow compression by the tires is not simulated, since most maneuvering occurs outside the area of interest to be cleared, while during a job the machine drives in paths cleared by the bucket.

If the bucket capacity is reached during a clearing path, the excess volume is distributed in a rectangular area in front of the scoop, represented by the blue circular points in Figure 3.23. This area is arbitrarily set at 0.4 m long and 1.2 times the scoop width (with the scoop being 0.89 m wide, this rounds to 1.1 m), to simulate the full scoop acting like a plough, pushing most of the snow in front and some to the sides. The effect of modifying the dimensions of this distribution rectangle is also analyzed, however. The resettling algorithm described in Section 3.1 is applied to

the snow layer after each driving increment to maintain a 45° angle of repose, thus even if all excess material is distributed in front within one scoop-width, some can still collapse to the sides.

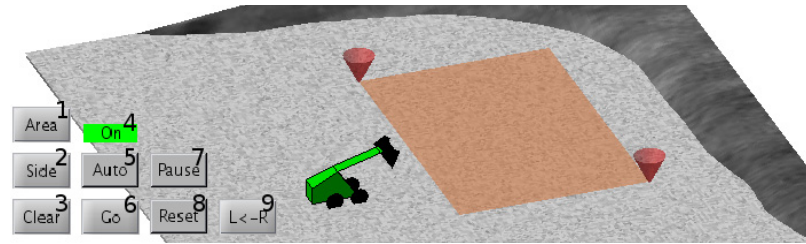
A sample job is first used to demonstrate the planning tools and path search strategy that were developed, with changes then made to the layer thickness, workspace length, path search parameters and snow distribution area for comparison.

3.5.1 Planning Features and Path Search Strategy

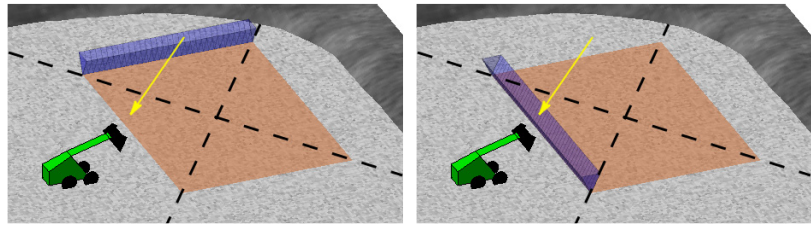
Figure 3.24(a) shows the initial state of the worksite for Job 1, with a 0.1 m-thick snow layer, and control buttons used to render the planning objects and manage the job. The first step is to specify the area for the machine to clear. This is done using the same Area Tool as before, also visible in the figure, here marking an area 9 m wide by 6.2 m long. Next, the clearing direction is specified using the Starting Side, an inverted triangular prism rendered along one side of the area (Figure 3.24(b)). This is basically the same as the Dump Approach and Scoop Approach tools used for the pile transfer simulations in Section 3.4.6. Any of the four directions can be selected by clicking and dragging the prism (represented by the arrow in Figures 3.24(b) and 3.24(c)). As the mouse is being moved, its position is mapped to one of the four sectors defined by the centre and corners of the rectangular area (marked by the dashed lines), and the prism snaps to the corresponding side. Once the direction is chosen, the rest of the job can proceed autonomously.

The clearing strategy is to make a series of straight, parallel drives from the Starting Side, beginning from the left or right edge and working across (here from right to left). This strategy was chosen in order to mimic how a human might ordinarily do such a job. To determine where the next clearing drive should take place, an algorithm scans along a line, computing the average height above ground level. Ground level, assumed to be unknown, is determined by first making one clearing drive at the right-most position (Figure 3.25(a)) and estimating the height in the resulting path using the histogram technique from Section 3.4.1. As with the pile transfer job, it is assumed that a ground model is available which is updated after any actions with the bucket.

One way the ground level might be known beforehand is if the machine had good 3D positioning information, and could then calculate the height at the bottom of its tires. Another way is if a previous ground model of the



(a) Clearing area specified with Area Tool; control buttons: 1-Render Area Tool, 2-Render Starting Side, 3-Begin clearing, 4-Auto mode indicator, 5-Auto mode on/off, 6-Resume/Next path, 7-Stop machine, 8-Delete plans, 9-Toggle left-right search direction



(b) Starting Side specifies direction; click (c) Starting Side repositioned by mapping mouse pointer to one of four sectors.

Figure 3.24. Control buttons and high-level plan specification for area clearing job.

worksite without any snow were available, or some exposed ground was visible. Here, however, it is assumed to be determined during the job.

During scanning, when the average height along the line reaches a certain threshold (here 0.03 m was used), the algorithm marks the location and continues. If subsequent scans all remain above the threshold and the distance from the first threshold location reaches the width of the scoop, the algorithm stops and decides to make a new clearing drive. This maximum-width case happens when the threshold is first reached at the edge of the main snow layer, as in Figure 3.25(b). Although the scoop width is 0.89 m, the path width is rounded to 0.8 m since its edges are determined symmetrically from the centre line, limiting width increments to twice the surface grid resolution.

When clearing through the fresh layer, some spillage of material may occur into a previously clean area, as to the right of the machine in Figure 3.25(c). If the spillage is substantial enough, it will be detected afterwards by the scanning algorithm. If the spillage swath is narrower than the scoop width, the next clearing drive will be centered on it to try and clear it most effectively, as in Figure 3.25(d). In this way, the algorithm alternates between “heavy” clearing drives of fresh material and “cleaning” drives of residual material. Since the next path is decided based on the current state which can be influenced by previous actions, the area clear-

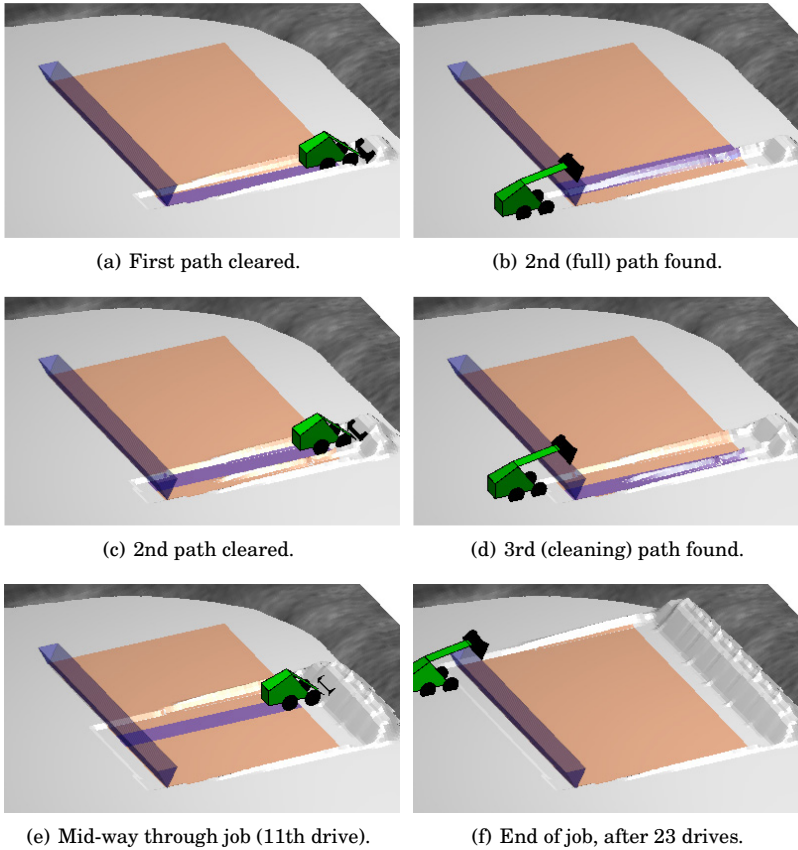


Figure 3.25. Simulation of area clearing Job 1, with a search threshold of 0.3 times the layer thickness.

ing job differs from classical “area coverage” problems where the only goal is to drive once over all parts of the area [73, 74]. Figure 3.26 illustrates the scanning line and corresponding ground heights, while a pseudocode version of the search algorithm is presented in Algorithm 2.



Figure 3.26. Computing average ground height along scanning line to find next clearing path.

The dark blue strip visible in Figure 3.25 represents the driving path chosen by the algorithm, and is rendered to show the current intentions

of the machine to the remote user. For the machine, the path consists of two driving points, one at each end of the path along the centre line. Since the vehicle reference point is at the centre of the wheelbase, the first point is offset back from the Starting Side (here by 1.5 m) so that when the scoop is lowered, its front edge is located at the front of the area. The second point is at the far end of the path, where the machine is located in Figure 3.25(a) for example. When this is reached, it raises and unloads the scoop, reverses to the first point (as in Figure 3.25(b)), and when the next path is found, it turns on the spot to drive to the first point of the new path, and the process repeats.

Algorithm 2 Algorithm for finding next clearing path.

```

threshold = 0.3*layerThickness;
thresholdDetected = false;
for position = rightEdge:leftEdge
    averageHeight = scan line at position;
    if averageHeight-groundLevel >= threshold
        if thresholdDetected == false
            thresholdDetected = true;
            firstPosition = position;
        elseif distance(firstPosition to position) >= scoopWidth
            return(path between firstPosition and position); // full path
        end
    elseif thresholdDetected == true // cleaning path found
        pathCentre = (firstPosition + position)/2; // find centre
        if distance(rightEdge to pathCentre) < scoopWidth/2
            return(right-most path); // too close to rightEdge
        elseif distance (pathCentre to leftEdge) <= scoopWidth/2
            return(left-most path); // too close to leftEdge
        else
            return(path at pathCentre); // path within main workspace
        end
    end
end
end

```

A threshold of 0.03 m was used here since a value of 0.3 times the layer thickness was found to work well for the two outdoor snow clearing tests in Sections 4.1 and 4.4. As Figure 3.25(f) shows, it also resulted in most of the layer being cleared for Job 1, after 23 drives. The remaining residual amount, mostly along the side edges, amounted to 0.049 m³ or 0.9% of the initial layer volume.

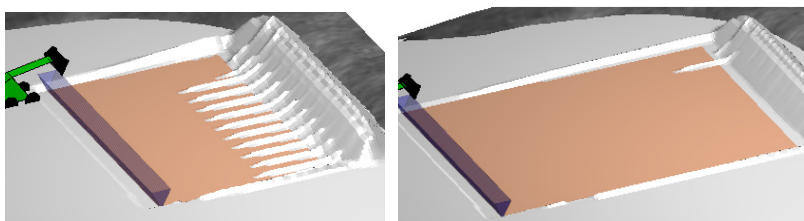
Using a threshold value which is a fraction of the layer thickness means that the original surface model needs to be saved, so that its thickness, taken as the average height of the layer, can be determined after ground level is known. In the strategy used here, both ground level and the layer height are updated during the job by evaluating them over the entire area cleared so far.

Using a fraction of the thickness as the threshold also means that the

definition of when an area is considered “clear” becomes relative. To check how this strategy works with different job parameters, two additional jobs are simulated in Section 3.5.2, first with a thicker snow layer and then a longer area. Results are then presented which investigate the effect of the search threshold by repeating Jobs 1-3 with a range of values. Additional effects which affect the outcome, including the dimensions of the excess snow distribution area and path width, are also examined by repeating Job 1 with changes to these parameters.

3.5.2 Effect of Path Search and Job Parameters

To determine how the path search threshold of 0.3 times the layer thickness would work with different jobs, Job 2 was first simulated which was the same clearing area as Job 1 but with the layer thickness doubled to 0.2 m. Job 3 was again a 0.1 m layer, with the area doubled in length to 9 x 12.4 m. Figure 3.27 shows the end states for these jobs, with the red area rendered just above the ground level plane to show the residual snow more clearly.



(a) Job 2: 0.2 m layer, same area as Job 1. (b) Job 3: 0.1 m layer, area twice as long.

Figure 3.27. Result of Job 2 and 3 with the same path search threshold of 0.3 times the layer thickness.

Job 2 (Figure 3.27(a)) required 26 drives and left a noticeable amount of material at the far end (0.674 m^3 , or 5.9% of the original layer volume). The pattern of path types followed during the job was also alternating heavy and cleaning paths. Job 3 (Figure 3.27(b)) required 44 drives and left a small residual amount of 0.212 m^3 at the end, or 1.9% of the original volume. This followed a pattern of one heavy path followed by 2 or 3 cleaning paths.

For Job 2, a higher clearance might have been achieved with a lower threshold, but this would probably have also had the disadvantage of increased driving. Conversely, less driving might have been needed for Jobs 1 and 3 if a lower clearance was acceptable. To further investigate the effect of the threshold and the tradeoff of clearance with driving amount,

Jobs 1-3 were all repeated with a range of threshold values, with the results shown in Figure 3.28. In these plots the thicker blue line shows the ratio of the volume layer that was cleared (left axis) and the thinner green line shows the corresponding number of drives to complete the job (right axis). Along the X-axis is the thickness ratio used as the search threshold, with the plots stopping at the maximum value which was able to find paths over the width of the workspace.

Three additional plots are also included for Job 1. First, changes were made to the distribution of excess snow, with the rectangular area (see Figure 3.23) first shortened to 1.1 x 0.2 m, then narrowed to 0.9 x 0.4 m. These types of changes might result from different snow properties, with light, dry snow possibly being lifted up easily and collapsing to the sides more than heavy, wet snow, which might tend to stick together and be pushed forward more. After this, the effects of a narrower search path of 0.4 m are studied.

Since there is no randomness in these simulations, a job is only run once for each search threshold. Randomness could be introduced by varying the excess snow distribution area at each timestep, and/or by introducing a random error to the loader's driving. Given the simplified nature of the simulations here, the results will serve only as approximations, however they will be compared with results from the outdoor snow clearing tests in Sections 4.1 and 4.4 in terms of which range of search thresholds yielded a particular clearance ratio.

As expected, a trend exists in the plots in Figure 3.28 whereby lower thresholds clear a higher portion of the volume, while using a higher number of drives. The plots have a general step shape, beginning with a plateau where high clearance is achieved when the thresholds are low enough to detect cleaning paths. The plots for Jobs 2 and 3 begin with a region where more than one cleaning path is required per heavy path (labelled "a" in the plots). Every plot then contains a range of threshold values over which the pattern consists of alternating heavy and cleaning paths (region "b").

After this, a sharp decline occurs when the increased threshold is unable to find most or any spillage, and the pattern changes to more heavy paths than cleaning paths (region "c" - only in the last plot), until only heavy paths are made without any cleaning (region "d"). The sharp decline usually happens around a threshold of 1.0. The reason that values above this are still able to find paths is due to the pushing of excess snow

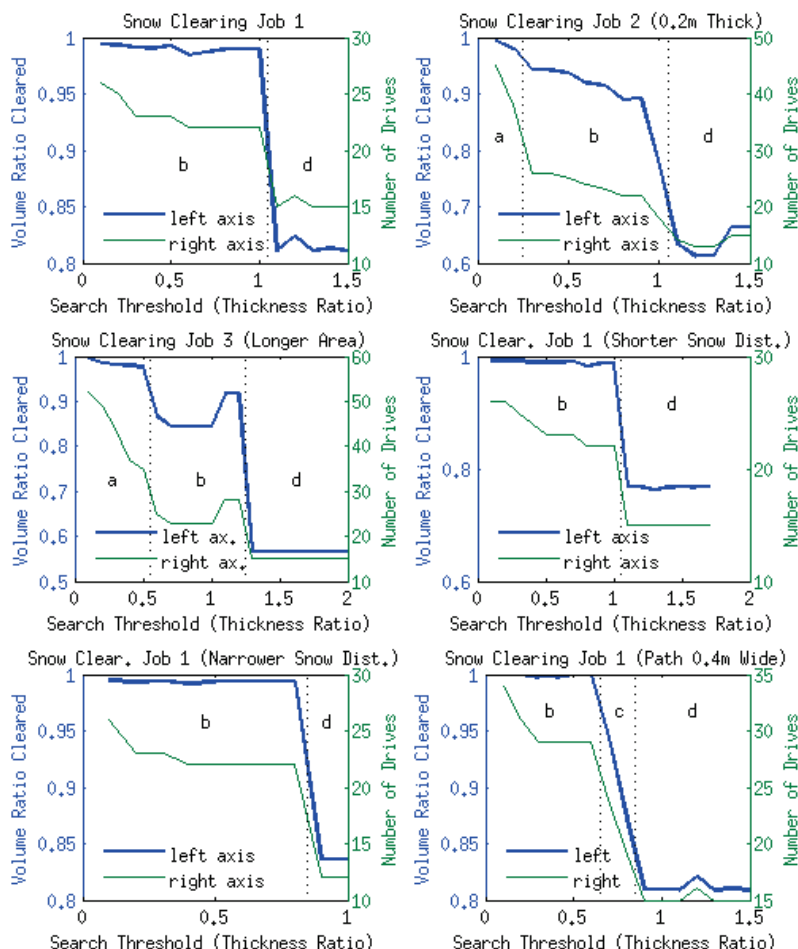


Figure 3.28. Snow clearing results for Jobs 1-3 and Job 1 with different parameters; zone description: a - more cleaning than heavy paths, b - alternating heavy and cleaning paths, d - only heavy paths.

to the sides by the scoop, which piles up higher than the original layer (as in Figure 3.25(a)).

Looking more closely at the results for Job 1 at the top left of Figure 3.28, all thresholds up to 1.0 result in a fairly constant clearance of at least 0.98. The thresholds 0.3-1.0 achieve this with the advantage of a few less drives (23-22) than the lower thresholds of 0.1 and 0.2 (26-25 drives). For Job 2 (top right of the figure), the clearance steadily declines from 0.99 to 0.89 over the thresholds up to 0.9, before the sharp decline begins. The best balance of clearance and driving seems to be with the thresholds of 0.3-0.5, which clear 0.94 of the volume and require 26-25 drives. Although the lower thresholds of 0.1 and 0.2 clear 0.99-0.98 of the volume, they require significantly more drives (45 and 38).

For Job 3, at the middle left of the figure, a plateau of high clearance of

at least 0.98 occurs in region “a” up to thresholds of 0.5, across which the number of drives decreases from 52 to 35. The clearance then decreases to a plateau of 0.84, after which it increases to 0.92 at thresholds of 1.1 and 1.2. The reason that these higher thresholds lead to an increase in clearance is that since they are higher than 1.0, they treat the start of the main layer as a cleaning path, only registering the crest of spillage which piles higher than the fresh layer. The path is centred on this crest and does not go all the way into the main layer, with the scoop positioned partly out in the already cleared region. This results in more of the path being collected in the scoop and reduced spillage, which is then cleaned to a higher degree than that created by the thresholds 0.6-1.0, which made heavy paths all the way into the fresh layer. The best balance here seems to be with a threshold of 0.5, clearing 0.98 with 35 drives.

At the middle right of Figure 3.28 is the plot for Job 1 with the shorter excess snow distribution of 1.1 x 0.2 m. With this distribution the same amount of material is still placed in front and to the sides, which is why the plot has a shape similar to the first version of Job 1, and also results in a clearance of at least 0.98 for thresholds up to 1.0. One difference with this distribution, however, is that since the material is piled higher, more collapses to the sides. This increase in spillage results in slightly more driving for most of the thresholds up to 1.0, with the thresholds of 0.5-1.0 offering the lower number of drives (23-22). Another difference from the first version of Job 1 is that in region “d” the clearance is below 0.8.

When the distribution of excess snow is narrowed to 0.9 x 0.4 m, the plot at the bottom left of the figure shows again a similar shape as the first version of Job 1. The main effect of this distribution is less spillage to the sides, which results in slightly less driving for thresholds of 0.2-0.5 (23-22 drives), and a high clearance of at least 0.99 up to a threshold of 0.8. The main difference here is that the transition to region “d” occurs after the threshold of 0.8, which is due to the decreased spillage not being detected by thresholds of 0.9 and 1.0. In this case the thresholds of 0.2-0.8 all offer about the same result, with 23-22 drives.

Finally, at the bottom right of Figure 3.28, the results with a narrower path of 0.4 m are shown. This has the effect that more of each path is collected in the bucket, causing less spillage and resulting in full clearance up to a threshold of 0.6. More driving is required, however, starting at 34 drives and then levelling out at 29 drives for thresholds of 0.3-0.6. The reduced spillage, as in the previous plot, results in an early transition

to region “d.” If the increase in driving is acceptable, changing the path width to this narrower value may be preferable to using the default of the scoop width.

3.5.3 Additional Planning Features

Returning to the result of Job 2 in Figure 3.27(a), the plot in Figure 3.28 showed how a lower threshold of 0.1 or 0.2 would result in less residual material at the end, but also more driving. Furthermore, these extra drives would mostly be across already cleared ground, since the residual amount lies along the far end. If this material is to be cleared, one way to reduce the extra driving could be to replan the job by specifying a new area only covering the remaining material. This could also be developed as an extension to the automated plan interpreter.

Another way to clear the remaining material, with a fewer number of paths, is for the user to change the clearing direction by moving the Starting Side as in Figure 3.29(a), assuming that enough maneuvering space is available at the new side. Repeating the strategy at the start of a new job, the machine clears the first path by default before evaluating ground level and the average thickness in the domain of this path. A new search then begins (Figure 3.29(b)), using the same threshold of 0.3 times the layer thickness, which is now less than the original 0.2 m layer.

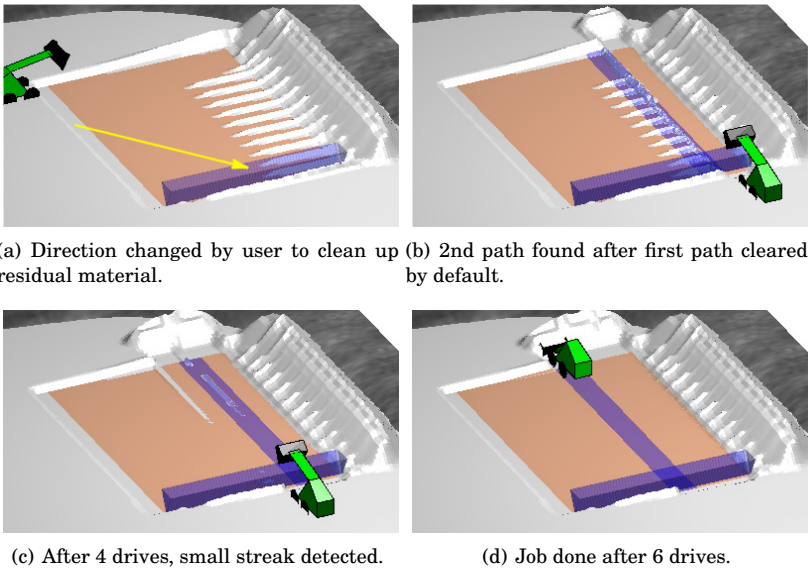


Figure 3.29. Cleaning residual snow after Job 2 by switching clearing direction.

As the job proceeds, the average layer thickness is re-evaluated, as be-

fore, over the domain of the newly-cleared area. Given the irregular topography of the residual snow, this average value changes as the area expands, ranging from 0.042 m to 0.025 m. One consequence of using a relative threshold is that since this residual snow was lower than the original layer, cleaning paths are found which would have been passed over the first time, such as those in Figure 3.29(c). A total of 6 drives is used in this 2nd phase of the job.

This example demonstrates one way the planning tools can be used to combine the high-level cognitive ability of a human to assess the state of work with the automated lower-level actions of the machine. An additional capability was developed which allows the user to specify a certain path for the machine by turning off Auto mode, then clicking and dragging the path itself to a new lateral location, however in this example it was not needed.

3.5.4 Discussion

In this section some of the same high-level planning tools used earlier for pile loading were used to plan area clearing jobs, by specifying the area and approach direction. As with the pile transfer planning, a simplification in the strategy presented here is that only rectangular areas having an arbitrary orientation on the worksite can be specified. A more advanced planner might generate a plan for any polygonal area, which could also contain obstacles. This type of area-coverage planner has been developed by Oksanen and Visala for agricultural applications, to generate driving swaths for machines operating on farm fields [92].

The plots in Figure 3.28 showed how the job and search parameters, and excess snow distribution, can all affect the volume clearance and number of drives needed to complete a job. Another factor not yet mentioned which could also affect the results is the driving speed. Depending on the properties of the ground material, driving faster might scatter it further to the sides which may be undesirable. If the material clumps together like wet snow, driving slower could help keep it together, enabling more to be pushed in front. On the other hand, in case the material provided significant resistance, the increased momentum of driving faster may be beneficial. This could not be studied directly since the simulator is non-dynamic, and the machine drove at the default speed of 0.5 m/s, however the changes to the distribution area for excess snow covered some of these possible effects.

For all of the plots in Figure 3.28, the threshold of 0.5 times the layer thickness generally offered the best balance as it used the least amount of driving to achieve a clearance of at least 0.98 for all cases except Job 2, for which the result was 0.94. The values of 0.3 and 0.4 achieved the same clearance and sometimes required more driving, while thresholds of 0.1 and 0.2 sometimes cleared more but usually required substantially more driving.

To ensure a high clearance, an alternative to using a relative threshold could be to use a fixed average height. Since all of the simulations other than Job 2 had the same snow layer 0.1 m thick, a fixed threshold of 0.05 m, for example, would have achieved the same result in those cases as the relative amount of 0.5 times the thickness. Furthermore, this fixed value would have resulted in a higher clearance for Job 2, corresponding to 0.25 times the thickness, although with more driving required.

One benefit of the relative threshold however, as the example in Figure 3.29 showed, is that if layers with a small average thickness are specified, paths can be found which may be noticeable relative to the thin layer but which might be insignificant relative to a thicker layer. Some fixed minimum threshold might be needed in any case to ensure that a job is not started if there is only a negligible layer covering the ground. The height resolution of the ground modeling sensors would also affect the minimum layer thickness which can be cleared.

To ensure even higher clearance of an area, another type of fixed threshold, rather than the average height along the scanning line, could be the number of points along the line, or in the area of the path, above a threshold height. This method, however, could result in much driving over already cleared ground if only small amounts of material are present. To reduce driving, a way of automatically shrinking the workspace as an area is cleared could be helpful, or also switching the clearing direction as in Section 3.5.3.

The cost of driving would be one important factor ultimately affecting the strategy used in practice, as well as how much residual material left in the area is considered acceptable. As Sections 4.1 and 4.4 show in the next chapter, the relative threshold of 0.3 times the layer thickness is able to discern quite well between paths of residual material which were considered significant enough to clear and those which were not, for two different outdoor snow clearing jobs. These experiments then lead to a full demonstration with a computer-controlled machine in Section 4.5.

4. Outdoor Tests Using Snow

The job planning tools and algorithms developed in Chapter 3 have thus far been validated in an ideal simulation environment, which assumes a perfect ground model, exact positioning and driving for the machine, and no friction or forces. To verify that the same tools and algorithms would work in the real world, where imperfect sensor information and positioning errors are the norm, a series of 8 outdoor tests were conducted for the same types of area clearing, area filling and pile loading jobs as in the simulations.

The strategy was to first test perception capabilities during manually-driven jobs, by building updated ground models from data obtained with a stationary 3D laser rangefinder. These models were used to confirm that progress could be tracked and commands correctly generated using the same planning tools and algorithms as in the simulations. Later, full proof-of-concept demonstrations were conducted in which commands generated on site were sent to a computer-controlled Avant 635. The last two tests made use of an onboard laser scanner for building ground models.

The focus was initially on area clearing and area filling jobs, since it was assumed they could be performed using preset joint positions and movements, and would therefore be easier to demonstrate in the near-term using a computer-controlled machine. In the end, scooping was also included, where the machine had to decide for itself when to stop the scooping action. Another reason for focusing first on area clearing was that each clearing path driven would result in a relatively large change in ground heights that should be easier to detect with the laser scanner than with single dumping or scooping actions.

The first tests used snow as the ground material, which was convenient to use in the winter as there was plenty available around the Avant storage facilities, making it easy to set up various jobs on the paved lots

nearby. After the tests were done, it could be pushed aside without extra cleaning required. The area clearing jobs were only intended for snow, which can be pushed by an Avant with bucket attachment. Proper bulldozing of soil would likely require a bulldozing blade, however this study was limited to a wheel loader + bucket combination. The later scooping and dumping tests made use of gravel, which was important to include since the planning tools are indeed intended for earthmoving jobs with granular ground material.

This chapter presents the first four tests which made use of snow. Basic details are summarized in Table 4.1, with the procedure, results and analysis for each one in the following sections. In addition, the photos and videos taken during these tests offered the opportunity to develop Augmented Reality versions of the job planning tools, which are presented in Section 4.3.

Table 4.1. Summary of outdoor tests using snow.

Test	Date	Location	Job	Machine/Tool	Control
1	27.03.2010	Otaniemi	Area clearing	Avant 320	Manual
2	27.03.2010	Otaniemi	Area filling	Avant 320	Manual
3	12.12.2010	Otaniemi	Area clearing	Snow shovel	Manual
4	30.03.2012	Hervanta	Area clearing	Avant 635	Computer

4.1 Outdoor Test 1: Manually-Driven Area Clearing

March 27, 2010, Aalto Campus, Otaniemi (3 °C)

In this experiment, the Avant 320 was driven manually to clear a layer of snow, with updated ground scans being made with a 3D laser rangefinder [93]. The purpose was to confirm if surface models built from the scans could be used with the planning tools from Section 3.5 and Algorithm 2 to track changes and correctly determine where the next driving path should be. This would then set the stage for future automated tests, by showing that commands could be generated for a computer-controlled machine to follow.

The experimental apparatus is shown in Figure 4.1. At left is the skid-steered Avant 320, in the middle the camera used to take photos and videos, and at right the Riegl LMS-Z210 laser scanner fixed to the cart, connected to a laptop computer running the manufacturer’s software in

Windows. The bucket dimensions of this Avant, with a roughly triangular cross-section, are 0.91 m wide, 0.69 m long and 0.64 m high, giving it a volume of approximately 0.20 m^3 . The location of the test was the Aalto Campus in Otaniemi, Espoo, on an asphalt parking lot covered with ice and slush, near the workshop where the Avant was stored.



Figure 4.1. Experimental apparatus for Test 1: skid-steered Avant 320 compact wheel loader at left, camera, Riegl 3D laser scanner (fixed to cart).

A rectangular layer approximately 5.8 m wide, 3.0 m long and 0.25 m thick (Figure 4.2(a)) was first set up using snow from the surrounding area, the snow being quite wet as the temperature was 3°C . After making an initial scan, the area was cleared with a series of 9 drives, with new scans being made after each clearing action. Figure 4.2 shows images from the experiment.

The first 7 drives were “heavy” paths into the fresh layer from right to left, similar to the simulation in Section 3.5. The first 6 resulted in negligible spillage to the sides, thus extra cleaning drives were deemed unnecessary. The low spillage was mostly due to the relatively short length of the paths, with a significant volume of each swath being collected in the bucket. After the 7th drive, 2 larger residual streaks remained, which were cleaned in the opposite order, from left to right. This clearing sequence was not planned beforehand, but simply seemed like the logical way to do the job at the time. The driving speed was kept slow and steady, around 0.6 m/s.

During the test, the laser rangefinder was kept stationary to avoid the

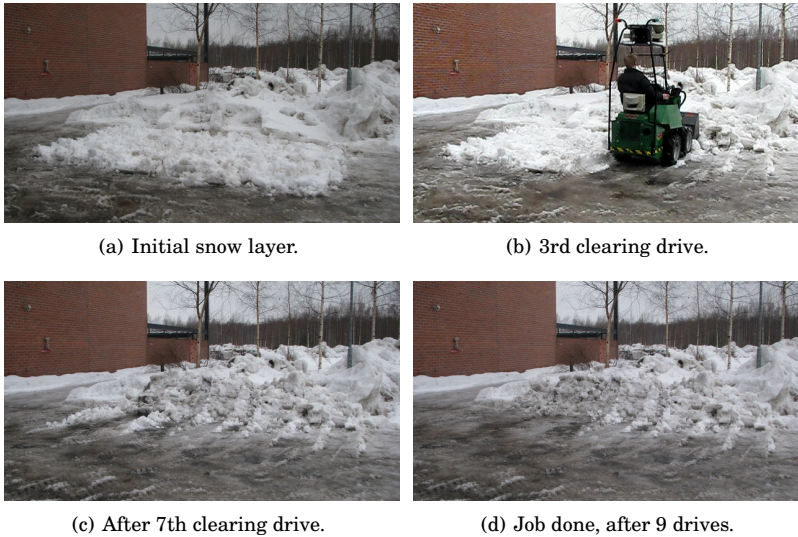


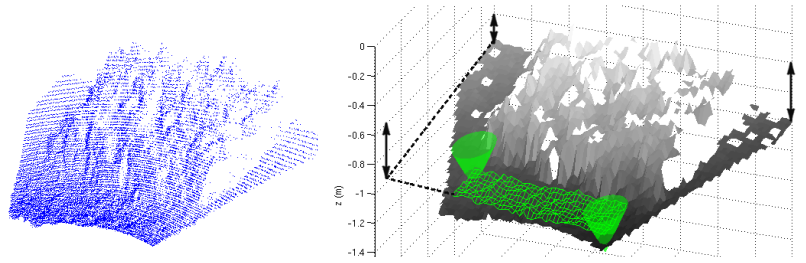
Figure 4.2. Outdoor Test 1: manually-driven snow clearing.

need for registering scans made from different locations. The Riegl LMS-Z210 scans in a vertical plane with an 80° field of view, while panning back and forth from 0° to 330° . Scanning time can be reduced by limiting these angular domains to the area of interest and by adjusting the resolution. For this experiment, a new scan took approximately 45 s, after limiting the domain to the workspace and using a medium resolution. Raw scans were obtained in the form of a 3D point cloud, with the list of coordinates saved in a text file on the laptop. These were processed afterwards to build surface models that could be tested with the planning tools. Next, the construction of these models from the point clouds will be described.

4.1.1 Surface Model

Figure 4.3(a) shows a raw scan of the initial snow layer (from Figure 4.2(a)), containing 23049 points. This is processed to create a heightmap with a 0.1 m grid resolution in the horizontal (X-Y) plane. Figure 4.3(b) shows the resulting surface, which contains holes mainly due to occlusion, and also has a trapezoidal shape due to the limited field of view used. The ground plane surrounding the snow layer is inclined in the laser scanner's coordinate system, indicated by the arrows in the figure. This is due to the cart/laser scanner standing on uneven ground or on a different slope with respect to the workspace.

For the surface in Figure 4.3(b), the point cloud was processed with an older procedure whereby if more than one point was located in the same



(a) Raw point cloud of initial layer (b) Inclined surface model with holes. Portion selected for ground plane estimation.

Figure 4.3. Creation of surface model from point cloud.

grid cell, only the last one iterated through was kept. All subsequent ground models in this section and Sections 4.2, 4.4 and 5.1, however, were constructed using a newer procedure whereby the value is calculated as the average height of all points falling within a grid cell. As an example, the highest number of cloud points in one grid cell in Figure 4.3(a) is 103, at the front centre of the layer, whereas many cells also contain only one point.

Rendering the graphical job planning tools requires a surface with no holes, since the geometrical objects are plotted relative to the local surface heights. Having a flat ground plane is also helpful, making the heights above ground level easier to compute. To finish constructing the model, the surface is rotated to make it parallel with the X-Y plane. This requires an equation for the ground level plane, which is estimated by first selecting a section of the flat ground with the Area Tool (shown at the front in Figure 4.3(b)). Around 100-200 evenly distributed points are then sampled from this area to construct the over-determined plane equation $z = ax + by + c$, whose coefficients are found by solving numerically using a least squares technique.

An alternative to using the Area Tool to select a ground plane section is to use the RANSAC plane fitting technique on the point cloud, with perhaps a verification step needed to make sure the right plane was found [91]. Once the ground plane is known, the rotation matrix from the ground surface to the X-Y plane can be constructed after finding the ground plane normal vector. Every coordinate in the surface model is then multiplied by this matrix to make the surface model level.

Holes in the scan data are filled using a nearest-neighbour averaging scheme, whereby for empty cells, the average of the 8 surrounding cells which have values is assigned. To ensure that all empty cells are filled, the

(X,Y) domain is iterated through twice, starting at opposite corners. Figure 4.4 shows the resulting ground models from this test corresponding to the photos in Figure 4.2.

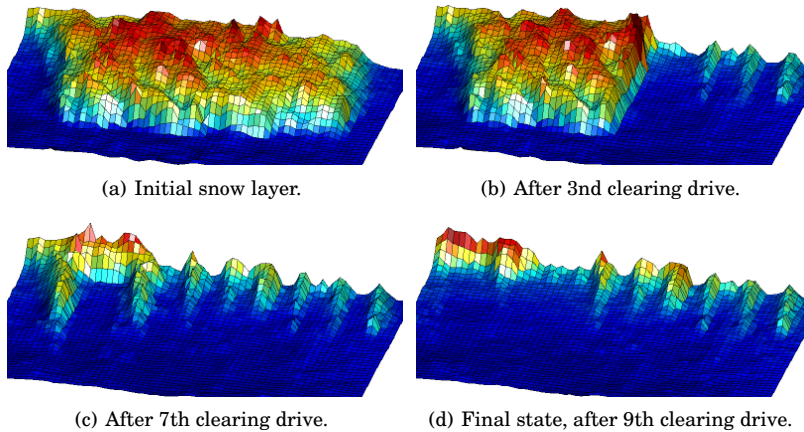


Figure 4.4. Surface models from Test 1, with residual streaks visible.

These surface models were rotated about the Z-axis a small amount so that the driving paths would align with the length axis, which is necessary since the tools and search algorithm are oriented along the X and Y directions. This alignment can be observed at the right edge of the layer in Figure 4.4(b), after the 3rd drive. After the 7 heavy drives all residual traces of snow are visible in Figure 4.4(c), of which only the two left-most were considered large enough to need clearing. Figure 4.4(d) shows the end state after 9 drives, with some of the piled up snow visible at the back left of the surface. The bottom of this slope is used in the next figure to define the back end of the clearing area.

4.1.2 Path Search

With ground models from the test available, the final stage was to see if the sequence of clearing paths could be automatically determined using the planning tools from Section 3.5 and Algorithm 2. Figure 4.5(a) shows how the Area Tool is first used to specify the initial snow layer. This job is somewhat different from the previous simulations since it is not a fresh layer of snow which covers the whole ground, but a limited layer surrounded by the exposed ground. An alternative way to plan the job could be to specify an area larger than the snow layer, determine ground level from this and then look for the first path using some fixed minimum threshold. Here, however, the job is treated the same way as before, by

specifying only the area which is to be cleared.

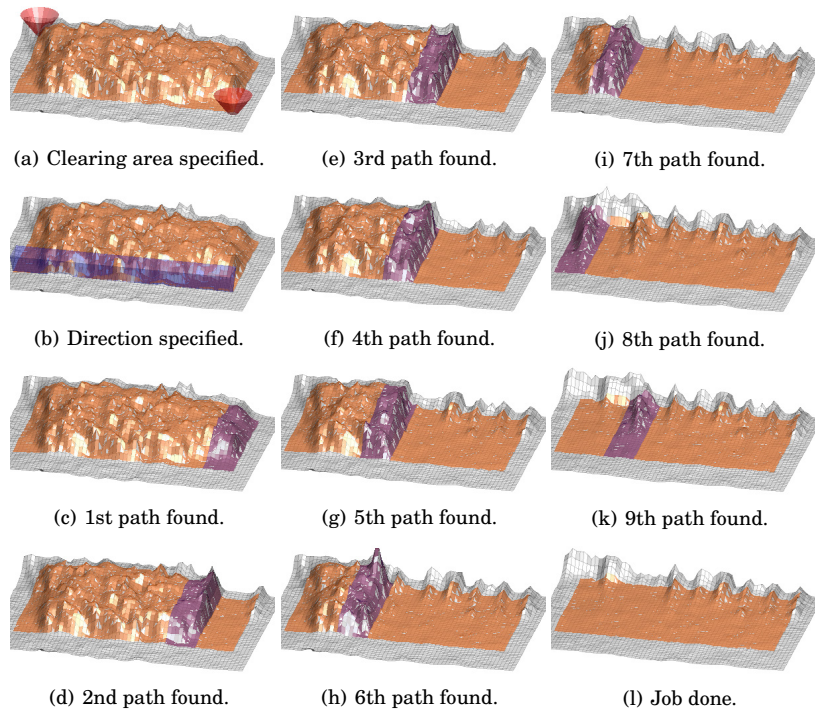


Figure 4.5. Planning tools used with Test 1 surface models; next path found automatically by search algorithm.

Next, the Starting Side is used to specify the clearing direction, shown in Figure 4.5(b). The first path is chosen by default (Figure 4.5(c)), after which ground level and the average layer thickness are first evaluated. The threshold from Section 3.5 of 0.3 times the layer thickness is then used again to search for subsequent paths, with ground level and the layer thickness updated as the cleared area increases. Over the course of the job, the average evaluated thickness ranged from 0.112 m to 0.202 m, which is less than stated thickness of 0.25 m (evaluated across the top of the layer) since some edges of the specified area reach ground level.

Figures 4.5(d) to 4.5(i) show how the next 6 paths are found and marked with the dark blue strip, all of these being heavy paths into the fresh layer. Once the left side is reached, the algorithm switches direction in a slight modification to mimic how the job was done, and the two larger streaks at the left are found (Figures 4.5(j) and 4.5(k)), thus completing the job. The volume of material above ground level before and after the job were estimated to be 3.301 m^3 and 0.195 m^3 , respectively, yielding a clearance ratio of 0.94.

This analysis was repeated with a range of threshold values to see which

ones could find all 9 paths in the same way, and where the upper and lower limits lie regarding the detection of paths. It was found that thresholds of 0.23-0.41 times the average layer thickness could find all 9 paths, while those below became stuck on small amounts of spillage considered insignificant, and values above began to skip over some of the paths. Table 4.2 summarizes the results.

Table 4.2. Path search threshold analysis for Test 1.

Threshold range (Layer thickness ratio)	Path search result
≤ 0.22	stuck on small residual, cannot find all heavy paths
0.23-0.41	all 9 paths found
0.42-0.61	7 heavy paths and first cleaning path found
0.62-1.29	7 heavy paths found
≥ 1.30	skips over some heavy paths

4.1.3 Discussion

This test confirmed that the available Riegl LMS-Z210 laser scanner could be used with the planning tools and Algorithm 2 to generate commands that an automated machine should be able to follow to clear an area of snow. It also provided an idea which threshold values could be used in practice to find paths containing enough snow to require a new drive. The thresholds of 0.23-0.41, which were able to find all paths and resulted in a clearance of 0.94, were values which also worked well and resulted in a clearance of at least 0.94 in the simulations of Section 3.5. This sets the stage for the computer-controlled test in Section 4.5, but first, the next sections will present two additional manually-performed tests.

4.2 Outdoor Test 2: Manually-Driven Area Filling

March 27, 2010, Aalto Campus, Otaniemi (3 °C)

This experiment was performed the same day as the previous one, and used the same apparatus as in Figure 4.1. The purpose this time was to verify if the Virtual Pile tool and dumping location algorithm from Section 3.4.5 could generate commands for an area filling job using scans from the Riegl laser rangefinder.

The job was first performed manually by driving the Avant 320 to fill

an area with 9 scoop-loads of snow in a 3 x 3 square pattern. With one scoop-load per location, this represented a “minimum height” pile. Since the snow was quite wet, each sub-pile had a rather irregular shape, sometimes mostly composed of one large block that had compacted together in the bucket. Nevertheless, all 9 deposits together formed a rectangular shape with a roughly uniform height.

The same filling pattern was followed as in the simulation from Section 3.4.5, i.e. first the back row was filled from left to right, then the next two rows in the same way. A laser scan was made of the initial workspace and after each filling action, with the scans processed afterwards to make surface models using the procedure from Section 4.1.1. Images and ground models from the job are shown in Figure 4.6, with some of the piled up snow bordering the rear of the area included in the surfaces.

The ground models could then be used with the Virtual Pile tool to see if the dumping locations could be automatically determined. The first step was to work backwards by finding a Virtual Pile which fit the shape of the final snow layer reasonably well. This is shown in Figure 4.6(h), with base dimensions 3.7 m wide by 2.2 m long, a height of 0.31 m and slopes of 45°. As in the simulation in Section 3.4.5, the Virtual Pile is actually rendered over a height matrix of the initial empty ground model, which allows it to be superimposed within the same 3D space as the final layer.

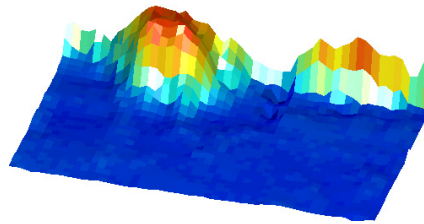
With the Virtual Pile set, Figure 4.7 shows how it was then used starting with the first ground model to search for available dumping locations. The same search algorithm as in Section 3.4.5 was used, with a line segment having the width of the scoop (rounded to 0.8 m) plus 0.1 m on each side, making a total of 1.0 m. The line starts at the back left of the top section, comparing the desired ground heights of the Virtual Pile (round blue points) with the corresponding actual ground heights (square black points). In Figure 4.7(a), since all the desired heights are above the actual ones, the first location is considered available.

After the first scoop-load is deposited, the scan for the next location passes over this location since at least one ground height has risen above its desired height (Figure 4.7(b)), with the line moving along increments of 0.1 m. When the back row has been filled up (Figure 4.7(f)), the scanning moves to the next row to find the next free locations, starting again from the left. Figures 4.7(g) to 4.7(r) show how these are correctly determined for the rest of the job.

To achieve this result, some adjustments were made to the Virtual Pile



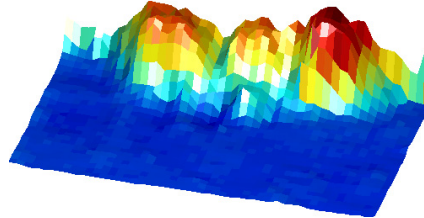
(a) 1st deposit.



(b) Ground model after 1st deposit.



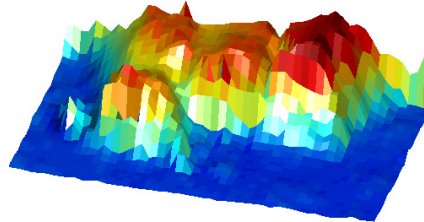
(c) 3rd deposit.



(d) Ground model after 3rd deposit.



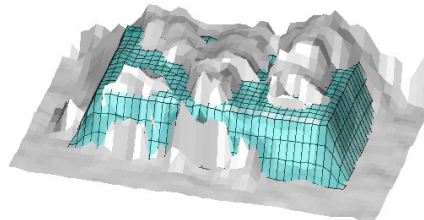
(e) 7th deposit.



(f) Ground model after 7th deposit.



(g) Final state - 9 deposits.



(h) Ground model of final state, with Virtual Pile fitted to snow layer.

Figure 4.6. Photos and ground models from outdoor area filling test: 9 scoop-loads dumped in a 3 x 3 square pattern, from left to right, and back to front.

dimensions until all 9 locations could be found successfully. The height of 0.31 m was equal to the average of the final layer over the domain of the top rectangular area of the Virtual Pile. Using the average height worked since the final layer was rather uneven, therefore the peak of each sub-pile protruded through the top of the Virtual Pile.

One change which was made to the search parameters from Section 3.4.5 was that the length step between search rows was increased from 0.2 m to



Figure 4.7. Automatic job tracking of outdoor area filling test: next dumping location found by comparing ground heights with corresponding Virtual Pile heights.

0.6 m, in order to match how the experiment was performed. The default value of 0.2 m would have also found free locations, though not all would have corresponded to those used during the experiment.

In the end, this experiment showed that scans from the Riegl could be used with the Virtual Pile and search algorithm to track job progress and generate commands that a robotic machine could follow to perform an area filling job autonomously. The assumption remains that the machine would be able to position itself and deposit the scoop-loads accurately enough. This was studied later in Tests 7 and 8 when pile transfer jobs were conducted with the computer-controlled Avant 635.

4.3 Augmented Reality Versions

The ground models and photos from Outdoor Tests 1 and 2 provided the opportunity to develop Augmented Reality (AR) versions of the graphical job planning tools, which are rendered directly over a camera image of the workspace. The goal here was to demonstrate a concept for increasing the situational awareness of a human operator who is remotely monitoring robotic earthmoving work, possibly by several machines, while perhaps also busy with other tasks. In this scenario the operator's attention would be divided, occasionally checking up on each machine but with minimal time to devote to it. By providing a current image of the workspace, the current intentions of the machine and the overall plan all in the same figure, an AR interface could help the operator to quickly assess the current situation and decide if any further attention is required before moving on.

For the Mars construction scenario, Figure 4.8 presents a concept of AR versions of the Area Tool and Virtual Pile rendered on the Martian surface. These could be used to specify areas to level or clear of rocks, and where to dump excess material, respectively. The long, narrow Virtual Pile specified here could be the type of shape required for a blast-protection berm around a landing pad for example.

This section will present the development of AR versions of the basic graphical tools used in this thesis using data from the previous outdoor area clearing and area filling tests. The AR versions are then also used to present the results of further tests in this chapter.

4.3.1 Rendering and Interaction

Since the graphical planning tools are rendered over a 3D surface model of the workspace, AR versions require that 3D information be mapped to the 2D plane of a camera image. Using the initial snow layer from Test 1

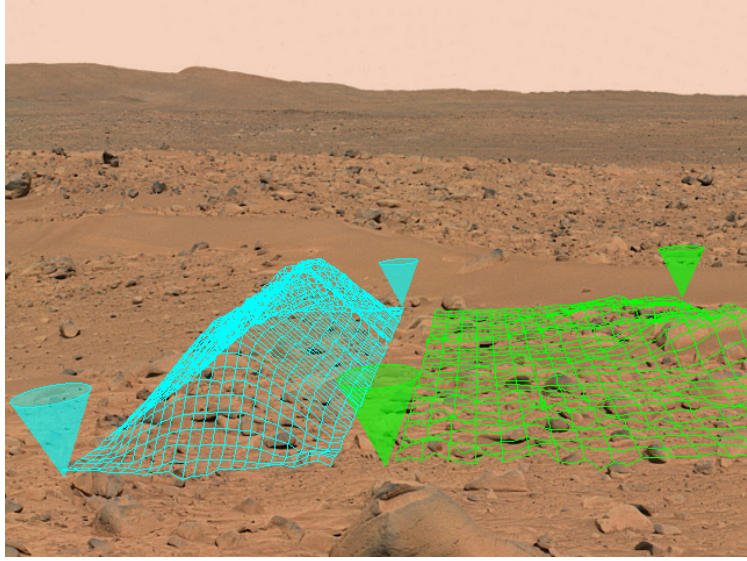


Figure 4.8. Concept of earthmoving construction job planning on Mars using Augmented Reality graphical planning tools (Photo: NASA).

as an example, this would mean transforming the points A, B and C in Figure 4.9 to their corresponding image locations at D, E and F properly.

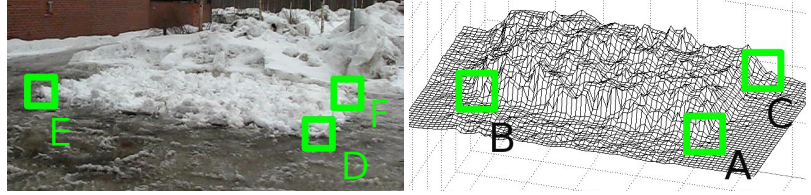


Figure 4.9. Photo (*left*) and ground model (*right*) of initial snow layer from Test 1, with 3 common points located.

This transformation is done using the pinhole camera model, illustrated in Figure 4.10. When the 3D coordinates of a point are known in the camera frame \mathcal{F}_k , its corresponding coordinates in the image plane are found using the pinhole camera equations:

$$\begin{aligned} u &= f * {}^k x / {}^k z \\ v &= f * {}^k y / {}^k z, \end{aligned} \quad (4.1)$$

where u, v are the image plane coordinates, f is the focal length of the camera and ${}^k x, {}^k y, {}^k z$ are the coordinates in \mathcal{F}_k .

All 3D information is obtained in the Riegl laser scanner frame \mathcal{F}_r , thus the transformation matrix ${}^k T_r$ is required to transform coordinates from the Riegl frame to the camera frame. This is complicated by the fact that the position \overrightarrow{RK} and orientation of the camera relative to the laser scan-

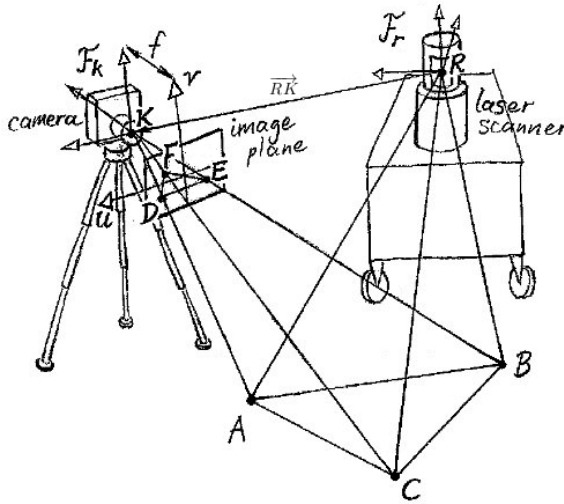


Figure 4.10. Geometry of camera frame, image plane, laser scanner frame and three ground points.

ner is unknown. It is possible to determine these, however, by calibrating the two frames with a geometrical triangulation procedure that requires three common points located in a surface model and corresponding workspace image. These points are represented by A, B, C and D, E, F in Figures 4.9 and 4.10, and were chosen since the corners of the snow layer were the easiest features to identify in the ground model. The calibration procedure consists of the following steps:

1. find distances \overline{DE} , \overline{EF} , \overline{FD} and \overline{KD} , \overline{KE} , \overline{KF} using known dimensions of image plane (i.e. physical dimensions of camera charge-coupled device (CCD)) and focal length f
2. find angles $\angle DKE$, $\angle EKF$, $\angle FKD$ using cosine law
3. calculate distances \overline{AB} , \overline{BC} , \overline{CA} from known coordinates in F_r
4. find distances \overline{KA} , \overline{KB} , \overline{KC} using cosine law and angles $\angle AKB$, $\angle BKC$, $\angle CKA$, which are congruent with angles in step 2. Requires solving a system of three equations numerically
5. triangulate camera position K in F_r using positions of A, B and C, sides \overline{KA} , \overline{KB} , \overline{KC} and 3D distance formula. Requires solving system of three equations numerically and generally has two solutions, however one can be eliminated by knowing rough camera position relative to scanner. This is represented by vector ${}^r(RK)$
6. find coordinates of D, E and F in F_r by constructing normalized vectors \overrightarrow{KA} , \overrightarrow{KB} , \overrightarrow{KC} and using known distances \overline{KD} , \overline{KE} , \overline{KF}

7. find orientation of \mathcal{F}_k , which is locked to image frame, using coordinates of D, E and F in \mathcal{F}_r by constructing image plane normal vector and u, v axis vectors
8. axis vectors of \mathcal{F}_k as columns form rotation matrix rR_k from \mathcal{F}_k to \mathcal{F}_r . The inverse of this yields kR_r , the rotation matrix from \mathcal{F}_r to \mathcal{F}_k
9. obtain camera position \overrightarrow{RK} in \mathcal{F}_k by

$${}^k(\overrightarrow{RK}) = {}^kR_r({}^r(\overrightarrow{RK})) \quad (4.2)$$

10. transformation matrix kT_r can then be constructed as

$${}^kT_r = \left[\begin{array}{ccc|c} {}^kR_r & -{}^k(\overrightarrow{RK}) \\ 0 & 0 & 0 & 1 \end{array} \right] \quad (4.3)$$

11. finally, coordinates ${}^kx, {}^ky, {}^kz$ in \mathcal{F}_k are obtained from ${}^rx, {}^ry, {}^rz$ in \mathcal{F}_r using

$$\begin{bmatrix} {}^kx \\ {}^ky \\ {}^kz \\ 1 \end{bmatrix} = {}^kT_r \begin{bmatrix} {}^rx \\ {}^ry \\ {}^rz \\ 1 \end{bmatrix}, \quad (4.4)$$

and 3D points can be properly rendered in the image using Equation 4.1.

Now that rendering is possible, the other key part of the interface is to transform input by the user into a position on the 3D surface model. The user selects and moves objects in the image plane by clicking and dragging with the mouse. To generate a surface coordinate that corresponds to an image coordinate, a ray is constructed from the camera origin K through the mouse position in the image plane, and the intersection of this ray with the plane at ground level is found using parametric line-plane intersection equations.

If either the camera or the laser scanner were to move during a job, they would need to be recalibrated by repeating the above procedure. To implement an AR system in practice, it would therefore make sense to mount the camera and laser rangefinder on the same platform. The fixed relative position between the frames might then be known from manufacturing plans, or if not, the calibration procedure would only need to be done once.

The above manual calibration method was used in order to take a first principles approach, though automatic methods have also been proposed. In work by Geiger et al., a camera and rangefinder could be automatically calibrated in a lab setting using a single image for each sensor [94]. Zhang and Pless presented methods for automatically calibrating various

sensors on a platform moving on a plane, while Levinson et al. developed a system geared for outdoor applications which automatically detects sudden miscalibrations and also corrects gradual errors such as sensor drift [95, 96]. In the latter, the sensors are assumed to be initially calibrated, and objects with edges also need to be visible.

4.3.2 Augmented Reality Area Clearing Tools

Here the AR versions of the tools for the area clearing job are presented using data from Test 1. As before, the job begins by specifying the area to clear using the Area Tool, shown in Figure 4.11(a). To move the cone, the user clicks and drags directly on the image. Here the rectangular surface is rendered as an empty mesh with a larger 0.2 m grid spacing, in order to occlude less of the background image and make it possible to see more of the snow layer. This demonstrates a further benefit of the AR approach: since the augmented surface shows the machine's current understanding of the world, the human operator can compare it directly with the camera image and verify that it is accurate.

After the area is chosen, the direction is specified using the Starting Side, shown in Figure 4.11(b). The AR version works the same way as before, by clicking the prism and dragging it to one of the four sides of the area. The job tracker then scans from right to left to find the next path, rendering it when found to show the remote user the intended actions of the machine. Figure 4.11(c) shows the 4th path, found after 3 drives.

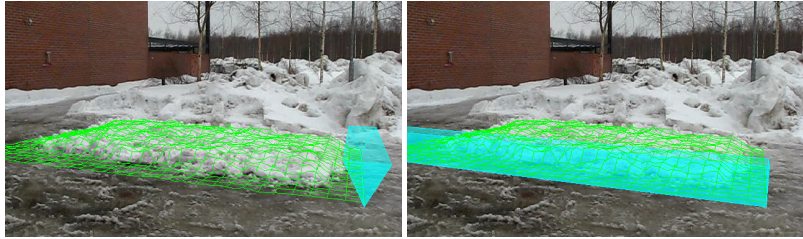
4.3.3 Augmented Reality Area Filling Tools

Data from Outdoor Test 2 was used to develop AR versions of the tools for area filling jobs. Figure 4.12 shows an example of the AR Virtual Pile, and how its height is automatically adjusted to keep the volume constant as the user moves the near corner cone to modify its footprint.

A Virtual Pile specified for Test 2 is shown in Figure 4.13(a), with the first dumping location found. Here the location is marked with a cone rather than the two lines of height-comparison points used earlier in Sections 3.4.5 and 4.2. This was due to the poor depth representation within the pile offered by the camera angle, which did not allow points to mark the location effectively. Figure 4.13(b) shows the job after five cycles, with the sixth dumping location found, and Figure 4.13(c) shows the final state, after all nine scoop-loads are deposited.



(a) AR Area Tool for specifying area to clear.



(b) AR Starting Side to specify clearing direction.



(c) After 3 clearing drives; 4th path found and displayed to user.

Figure 4.11. AR tools for area clearing job.**Figure 4.12.** AR Virtual Pile for specifying area filling job.

As the job proceeds, less and less of the Virtual Pile is shown; lines “under” the snow are not rendered, to avoid false occlusion of the material

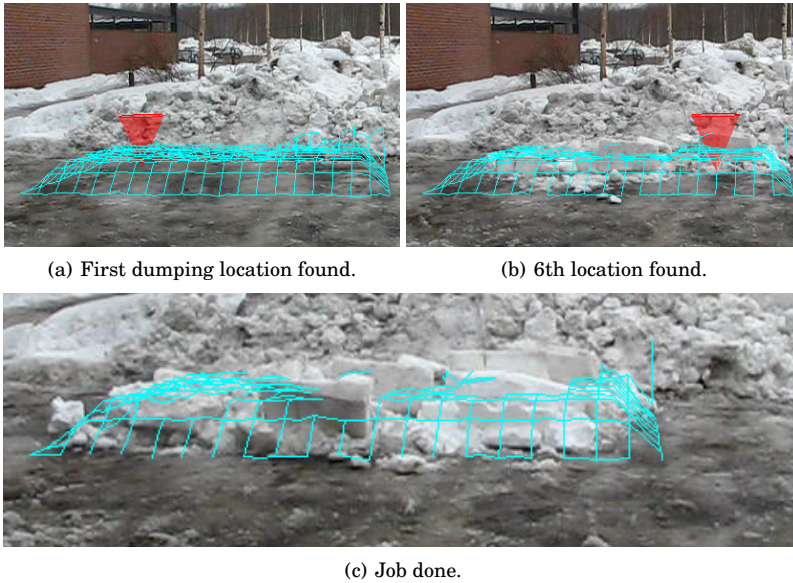


Figure 4.13. AR Virtual Pile for Test 2; next dumping location found automatically and marked with cone.

that has covered it. This was the main lesson learned when developing the AR versions of the tools, that it was important to render just enough for the tools to serve their purpose, but otherwise as little of the background image as possible should be occluded by extra graphics. The poor depth representation mentioned above also emphasizes the importance of camera position in providing a meaningful view of the workspace. In future systems this could be aided by the use of a mobile surveying platform.

4.4 Outdoor Test 3: Manual Area Clearing

December 12, 2010, Aalto Campus, Otaniemi (-14 °C)

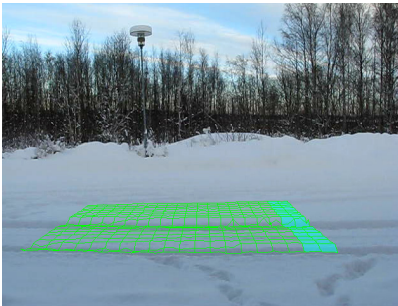
This experiment was another manually-performed area clearing job using snow as the ground material. The purpose was again to make updated ground models with the Riegl laser rangefinder and use these afterwards to test the automatic generation of commands. The original intention was to use the Avant 320 to repeat the procedure from Test 1, this time under different conditions since the ground was covered with a fresh layer of dry snow with a temperature of -14 °C. Due to a problem with the on-board electronics, however, the Avant was unavailable and it was decided to perform the job on foot using a snow shovel instead.

Using the snow shovel produced two unexpected benefits. First, it al-

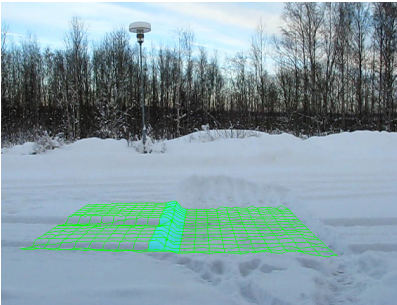
lowed the search algorithm to be tested when a different tool was used, with the only change in settings being the scoop width. And second, a snow behaviour similar to that exhibited in the simulations in Section 3.5 was recreated, i.e. while each fresh path of snow was being cleared, a significant amount would spill into the cleared section, requiring an extra cleaning path. This behaviour was likely due to the shovel having a small volume capacity and also the dry consistency of the snow, giving it a tendency to be deflected aside easily. This experiment therefore provided an opportunity to further test Algorithm 2 from Section 3.5, with its ability to find alternating heavy and cleaning paths.



(a) Apparatus.



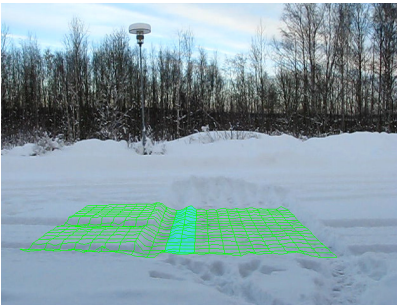
(b) Initial state; area to be cleared with first path.



(c) After 15 clearing actions - fresh path found.



(d) 16th clearing action.



(e) Subsequent cleaning path found.



(f) Job done, after 29 actions.

Figure 4.14. Outdoor Test 3: manual area clearing with snow shovel.

Figure 4.14 shows photos from the test with some results that were obtained afterwards, including surface models and paths chosen by the search algorithm. The workspace was the asphalt parking lot directly in front of the workshop where the Avant was stored, covered with a snow layer about 0.13 m thick. This location was advantageous since the Riegl scanner, laptop and camera could be set up inside the open garage door and kept warm via a forced air heater blowing down from above. This was important because the Riegl is only specified for operations to -10°C . Figure 4.14(a) shows the apparatus, with the 0.35 m-wide snow shovel that was used.

The workspace in its initial state is shown in Figure 4.14(b), with the area to be cleared marked by its surface model. The surface models and path search results shown here are older versions, though similar to the final versions in Figure 4.15, and are used here to illustrate the job. The area was 4.0 m wide by 3.8 m long, and was cleared as before using straight, parallel paths starting from the right side and working leftwards, with new scans being made after each clearing action. The shovel was pushed at a regular walking pace, approximately 1 m/s.

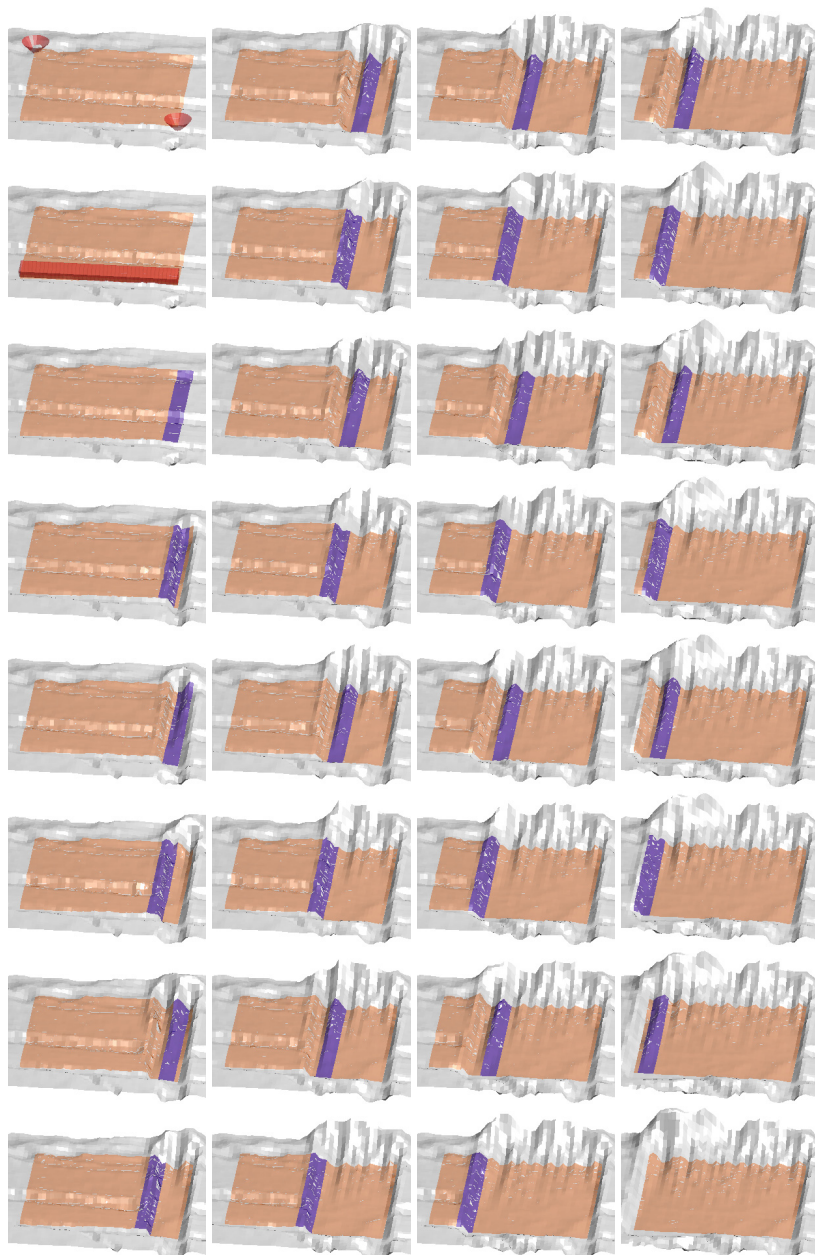
This scenario begins as in the simulation in Section 3.5, with the surrounding area also covered by the snow layer. Since “ground level” is unknown, the search algorithm would only estimate it after the first path is cleared, so the right-most path is chosen by default and also marked in the figure. In this figure a larger grid spacing of 0.2 m is used for clarity, and the path width is rounded from the scoop width of 0.35 m up to 0.40 m to display it more prominently.

When clearing the area during the test, no pre-planned criteria was followed for selecting the next path other than choosing the one which seemed necessary at the time. After the first path, the consistent pattern emerged whereby each fresh path caused enough spillage that one extra cleaning path was required. As an example, Figure 4.14(c) shows the fresh path found by the algorithm after 15 actions. After this is cleared (Figure 4.14(d)), the residual path is found (Figure 4.14(e)). In the end, a total of 29 actions were made to clear the area (Figure 4.14(f)).

4.4.1 Path Search

Figure 4.15 shows how the planning tools and search algorithm were applied to all surface models from the job, with a 0.1 m grid resolution. In the first 3 images at the top of column (a), the area and direction are first

specified, with the first path chosen by default. After this, a search threshold of 0.25 times the average layer thickness was used, which allowed the algorithm to correctly find all of the remaining 28 paths, alternating between heavy paths of fresh snow and cleaning paths of spillage.



(a) Area and direction specified, paths 1-6. (b) Paths 7-14. (c) Paths 15-22. (d) Paths 23-29, and end state.

Figure 4.15. Ground models from Test 3 (manual area clearing with snow shovel); all paths found by algorithm, with alternating heavy and cleaning paths.

In this figure the path width is again rounded up to 0.40 m to display it more prominently. This does not affect the path finding result since the algorithm scans along a 1-D line, however it does affect how far into the fresh layer the path is positioned. As before, the average layer thickness and ground level were evaluated after each clearing action over the area cleared so far, with the thickness values ranging from 0.10 m to 0.13 m. The volume of the initial layer above the final ground level was estimated to be 2.012 m³, while 0.145 m³ of material remained at the end, resulting in a clearance ratio of 0.93.

As with Test 1, this path search analysis was repeated with different threshold values to see where the upper and lower limits were for finding paths. It was found that thresholds of 0.23-0.28 were able to find all 28 paths (after the first default path), with values below 0.23 getting stuck on small residual streaks which were considered insignificant. Values of 0.29-0.38 found all except the 15th (cleaning) path, while values above 0.38 began to skip over more than one cleaning path. Only heavy paths were found with values of 0.60-1.12. Table 4.3 summarizes the results.

Table 4.3. Path search threshold analysis for Test 3.

Threshold range (Layer thickness ratio)	Path search result
≤0.22	stuck on small residual
0.23-0.28	all paths found
0.29-0.38	misses 15th (cleaning) path
0.39-0.59	misses more than one cleaning path
0.60-1.12	only heavy paths found
≥1.13	skips over some heavy paths

One parameter which affects these results is the rear edge of the clearing area. If this is placed further back into the snow pile, the average height along the scanning line is greater, resulting in higher threshold values finding the paths. Conversely, lower threshold values find the same paths if the rear edge is further forward into the cleared area. The rear edge is not clearly defined however, since the piled-up snow rises gradually without a clear bottom contour. When the job was performed, the area to be cleared was not marked beforehand, therefore it is difficult to know exactly what was considered the end of the area at the time. The location chosen in the end was that furthest to the rear where most of the valleys resulting after the clearing actions had a bottom ground height equal to

that along the front edge of the area, which was clearly defined.

4.4.2 Discussion

This experiment showed again that ground models obtained with the available Riegl laser scanner could be used with Algorithm 2 (Section 3.5) to generate commands which an autonomous agent should be able to follow to clear an area of snow, this time with a different tool and different snow conditions. Similar search threshold results were obtained as with Test 1, with 0.23 being the lowest value for both tests which was able to find paths successfully, and the range of 0.23-0.28 finding all paths for both jobs. This test resulted in an estimated clearance ratio of 0.93, which compared with the simulations in Section 3.5, was slightly lower than the clearances achieved with the same thresholds.

The threshold range of 0.29-0.38 found all paths for Test 1 and all but one path for Test 3. If this missed path can be overlooked, then a general result from both tests is that the range of 0.23-0.38 works well for finding paths with enough snow for clearing. The value of 0.3 used in the examples in Section 3.5 was chosen since it lies in the middle of this range. Another similar result from Tests 1 and 3 is the threshold transition at which only heavy paths were found: starting at 0.62 for Test 1 and 0.60 for Test 3.

Now that it has been shown that commands can be correctly generated for snow clearing jobs, the next step would be to send these commands to a robotic machine which is able to follow them autonomously. This is the goal of the next test which made use of the Avant 635 for snow clearing under computer control.

4.5 Outdoor Test 4: Robotic Area Clearing

March 30, 2012, TUT Campus, Hervanta (2 °C)

As a proof-of-concept of the area clearing planning tools and algorithms, and 3D graphical job planning concept in general, a snow clearing test was performed with the computer-controlled Avant 635 which was capable of driving autonomously outdoors. The experiment took place at the TUT campus in Hervanta, Tampere, and was a collaboration between members from both the Aalto and TUT research groups in the GIM CoE.

4.5.1 Apparatus and Preparation

Figure 4.16 shows the apparatus, set up on a paved lot in front of the IHA Mobile Laboratory. At right is the camera, next to the Riegl 3D laser rangefinder used in the previous tests, this time mounted on a tripod. Behind these is the rectangular layer of snow that was prepared (5.5 m wide by 4.2 m long, 0.19 m thick), which was wet and heavy as the temperature was 2 °C, yet the surface underneath still somewhat icy. Two laptops were used directly during the experiment: one for obtaining point clouds from the Riegl and running Matlab to construct surface models, plan the job and generate driving commands, and the other for transmitting commands to the Avant. A third was used by colleagues for monitoring the Avant.



Figure 4.16. Robotic area clearing test apparatus.

The Avant 635, visible at left, makes use of articulated frame steering, a form of Ackermann steering in which the wheels are fixed yet the chassis pivots around a centre joint. This has the advantage of avoiding skidding when turning, making it possible to use wheel odometry for dead-reckoning position estimation. This steering method also means that the Avant 635 requires more space to maneuver than a skid-steered machine, an important consideration when planning a job.

The autonomous outdoor driving capability of this Avant was developed at IHA, using a combination of GPS, inertial, centre-link angle and odom-

etry sensors, resulting in a driving accuracy of approximately 0.5 m [97, 19]. Driving commands are specified in the form of a desired 2D end pose of the vehicle frame in the TUT-IHA geographical coordinate system (here considered the global frame), and are followed with a driving speed of approximately 0.6 m/s. The vehicle frame is fixed to the front chassis and located at the centre of the front axle. Joint commands are sent in the form of a desired height, forward horizontal extension and pitch angle for the bucket frame at the rear of the bucket, relative to the vehicle frame. For this test the driving and joint commands were generated and sent using the GIM Machine Path Planner (GMPP) developed at IHA and the GIMnet Machine Control Interface (MaCI) running on the 2nd laptop [16].

Using 2 laptops for this experiment was necessary since the Riegl software was only available on the Windows laptop (which also happened to have Matlab installed), yet the GIMnet/MaCI software was developed for Linux. Driving commands generated on the first laptop were manually typed into the second laptop. Despite this manual step, the goal of the experiment was that driving coordinates would be generated automatically from the high-level graphical plan and search algorithms, and also that during the experiment all driving and scoop actuation would be done under computer control without any direct teleoperation.

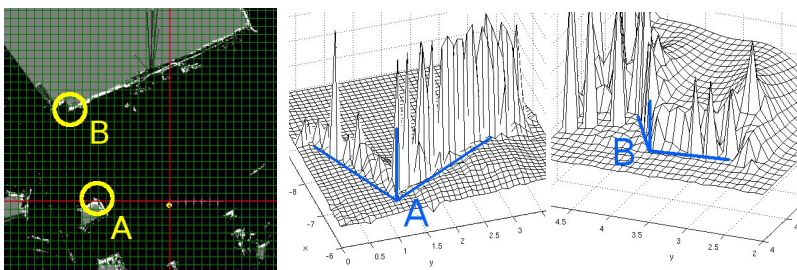
Since the Avant operates in the global coordinate system, yet the ground models, clearing paths and driving points for this test would all be generated in the local Riegl frame, the first step in this experiment was to obtain the transformation from the local to global frame. As the driving coordinates for the Avant are limited to the horizontal X and Y directions, the transformation could be reduced to a 2D problem as long as the laser scanner's Z-axis was aligned with the zenith direction. This was already approximately the case, with the Riegl standing vertical on the tripod. To account for any unwanted inclination, a surface model was first made of the asphalt in front of the Riegl which was assumed to be flat. The rotation procedure from Section 4.1.1 was then followed (as it was for the previous tests) to obtain the required 3D rotation matrix, which was applied to all scans taken thereafter to make the ground plane parallel to the Riegl's X-Y plane.

Two points known in both frames were then needed in order to determine the 2D coordinates of the Riegl and its rotation about Z in the global frame. Figure 4.17(a) shows how two building corners near the workspace were used for this purpose. These were chosen since they were visible to

the Riegl and it was possible to recognize them in two separate surface models made in their respective vicinities (Figure 4.17(c)). Their positions were also shown on a global map of the worksite which was available, although the coordinates could only be estimated with an error of approximately 0.5 m (Figure 4.17(b)). Once the corner coordinates were found, the Riegl was kept stationary for the rest of the experiment to avoid having to do so again. The two points A and B were then used to construct a 2D transformation from the Riegl to global frame by finding the rotation about Z and translation from one frame to the other.



(a) Workspace of robotic snow clearing test from far end, showing initial snow layer and building corner points used for coordinate transformation.



(b) Global map of Hervanta (c) Surface models of building corners to obtain points in test area. laser scanner frame.

Figure 4.17. Photo and map of Hervanta test area, showing points used to determine coordinate transformation from Riegl laser scanner frame to global frame.

With the transformation between the frames available, the planning tools and algorithms could then generate driving points that the Avant could follow. The Avant was then initialized, and a few short test drives

were commanded. The preset joint settings for clearing and dumping actions were also tested to make sure, for example, that the bucket blade was positioned just above ground level in the clearing configuration.

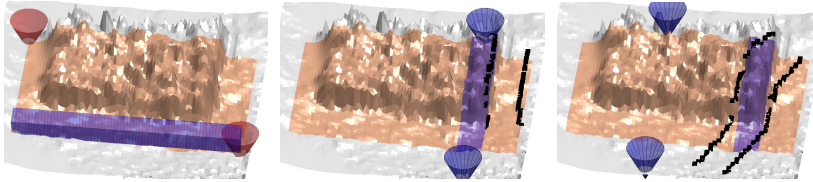
4.5.2 Experiment

When everything was ready and an initial surface model of the snow layer constructed, the actual specification of the job was done in a matter of seconds using the standard Area Tool and Starting Side (shown in Figure 4.18(a)). The specified area (7.2 m wide by 4.6 m long) stopped short of covering the entire layer at the back due to the offset between the Avant reference point at the front axle and the front edge of the bucket, about 1.15 m ahead of the axle in the clearing configuration (0.5 m to the bucket reference point plus the bucket length of 0.65 m). It was assumed that if the Avant drove to the back of the specified area, the portion of the snow layer beyond it would also be cleared.

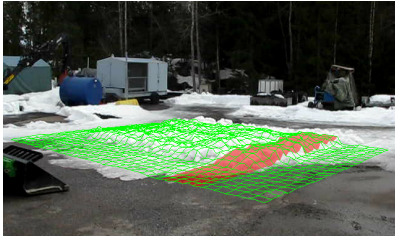
At the time of the experiment, some older versions of the planning software were being used. One of these was the older procedure for building surface models from raw point clouds mentioned in Section 4.1.1, whereby only the last point encountered per grid cell is used. Also, when specifying the snow layer with the Area Tool, some of the ground surrounding the layer was included because ground level was being determined using the histogram technique for piles (Section 3.4.1). In Sections 3.5 and 4.1, ground level was determined from the cleared area after first choosing the right-most path by default. A larger histogram interval of 5 cm was also being used to determine ground level, rather than 1 cm or 2 cm.

Finally, an older search algorithm was being used to find the clearing paths. This algorithm does not scan along a line as in Section 3.5, but first renders the whole path strip and compares changes in average height between adjacent locations as the strip scans laterally. The height thresholds for finding a path are therefore much lower than before since it is the change in height over the area of the strip, here 0.8 m wide. The actual bucket width was 1.04 m, however the path was made narrower since it was thought the extra width of the bucket would help account for positioning errors. This older search algorithm has the same basic capability of finding both narrow cleaning paths and fresh paths of material, yet if a narrow path is found it does not centre the drive at it and may approach it with one side of the scoop.

With a search threshold of 0.01 m, the first path to clear was found



(a) Initial surface model with (b) 1st path found, practice (c) Driving points offset to area and clearing direction drive (black lines) too far left, 1st clearing drive mostly right. covers intended path.



(d) Avant at 1st driving point of 1st clearing drive (only scoop visible).



(e) Driving autonomously, approaching layer.



(f) Driving autonomously, mid-way through layer.



(g) End of 1st drive.



(h) 1st step of automated bucket extraction.



(i) 2nd step of automated bucket extraction.

Figure 4.18. Initial specification of robotic area clearing job and details of first clearing drive.

(Figure 4.18(b)), with the associated driving points represented by the cones in the figure. The first driving point, represented by the near cone, is located 1.5 m in front of the edge of the area. This is due to the offset between the Avant reference point and scoop blade mentioned above, i.e. to allow enough space for lowering the scoop before driving into the layer.

When the coordinates for the first driving point were commanded, the Avant ended up about 1 m right of the layer. It was then commanded to

lower the bucket and drive to the 2nd point as a practice run, approximately following the black lines in Figure 4.18(b) and removing a small amount of snow at the edge. This error was assumed to be due to inaccuracies in the coordinate transformation process, i.e. the error in reading the building corners on the global map and the assumption that the initial ground plane modelled with the Riegl was flat.

To compensate for driving too far right, a lateral offset of 3 m to the left was introduced between the clearing path and the driving points. This was initially intended to be an over-compensation, which could be reduced after observing the outcome. This offset is evident in Figure 4.18(c), with the two cones now 3 m left of the path. When these coordinates were sent, however, the Avant cleared the path mostly on target, though it drove a somewhat angled path, again indicated by the black lines.

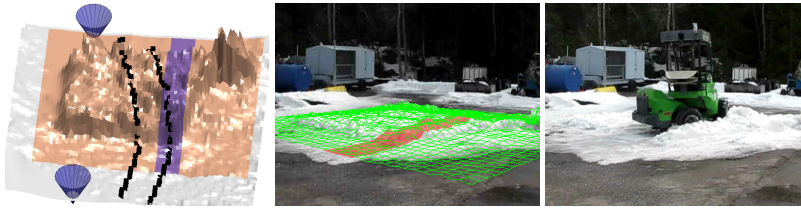
These lines are the estimated path of the front corners of the bucket, and were obtained by finding the coordinates of the right corner in 4 separate video frames spanning the driving action, then transforming these image frame coordinates into 3D surface coordinates using the ray tracing technique described at the end of Section 4.3.1. These 4 images are shown in Figures 4.18(d)- 4.18(g) to illustrate the first clearing action.

At the end of the drive, the Dump command was entered, in which an automated 2-step sequence was followed to extract the bucket from the collected snow. First the bucket was raised vertically by 0.3 m and angled forward by about 10° (see Figure 4.18(h)), then raised a further 0.6 m while rotated a further 50° (see Figure 4.18(i)). The Avant was then commanded to reverse to a standby point further back from the first driving point, where it waited while a new surface scan was taken to find the next clearing path.

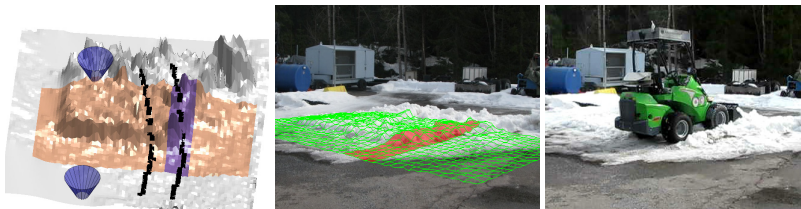
Although the first drive was mostly successful in the end, three settings were changed during the next two clearing drives in an attempt to improve the actions of the machine and generate commands as effectively as possible. These settings included the specified clearing area, the lateral offset between the path location and the driving points, and the average height threshold used by the search algorithm.

The surface model at the left of Figure 4.19(a) shows the result after the first clearing drive. The pile of accumulated material at the end of the drive was not quite pushed beyond the back of the layer, so for this reason it was decided to lengthen the specified area (from 4.6 m to 6.7 m) so as to fully drive through it. This lengthening also increased the area

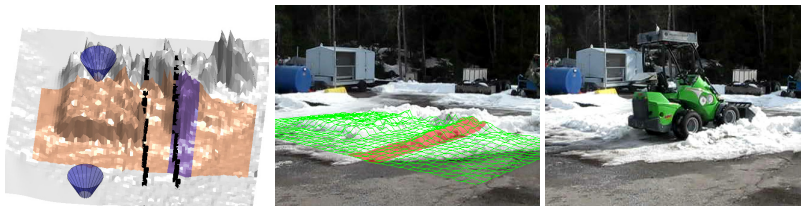
of the search path, which consequently caused a decrease in the average height difference between paths during the search. The height threshold of 0.01 m being used by the search algorithm was then too high to find a clearing path, so it was lowered to 0.007 m, allowing the path shown in the figure to be found. This threshold was kept the same for the rest of the job.



(a) 2nd clearing drive (1st attempt) with lengthened clearing area and new path search height threshold. Avant drove mostly left of path, then became stuck. Driven out manually, with gap in the layer filled in.



(b) 2nd clearing drive (2nd attempt) with shortened search area. Avant drove again too far left, leaving large streak to be cleaned.



(c) 3rd clearing drive (1st attempt). Large cleaning path found, but previous driving points sent accidentally. Good repeatability of previous driven path.



(d) 3rd clearing drive (2nd attempt), with reduced lateral offset. Cleaning path found again and cleared mostly on target. Current settings maintained for rest of job.

Figure 4.19. Clearing drives 2 and 3, for which changes were made to specified clearing area, path search height threshold and lateral positioning offset.

The black lines in the figure show how the Avant drove mostly left of the path. While in the middle of the snow layer, the Avant became stuck,

apparently from the resistance of the snow against the bucket and the icy surface underneath. The Avant was then shut down and driven out manually, and the gap in the layer filled in with extra snow.

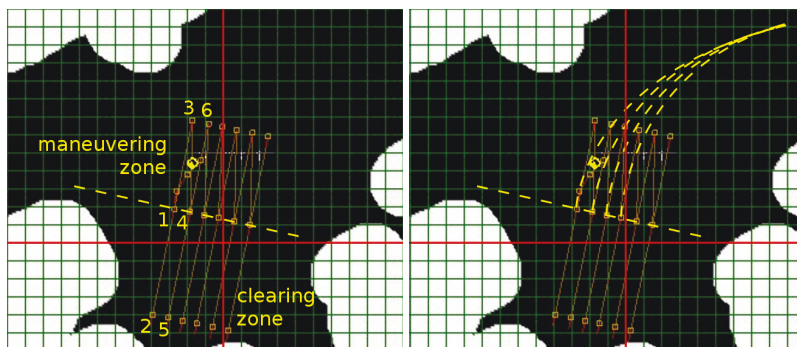
To avoid getting stuck again, the area was substantially shortened (from 6.7 m to 3.6 m) for the 2nd attempt of the 2nd drive (Figure 4.19(b)), so the Avant would not drive as deep into the layer. This area was kept constant for the rest of the job. The black lines on the surface model (at left) show how the Avant again drove mostly left of the path, as in the first attempt. The lateral offset was therefore reduced from 3 to 2 m, and kept the same for the rest of the job.

After the 2nd drive, a gap existed in the layer, leaving a significant “cleaning path” at the right edge which was found by the algorithm (Figure 4.19(c)). The previous driving points were accidentally sent to the Avant for the 3rd drive however, yet the black lines indicate good repeatability compared with the previous drive. For the 2nd attempt at the 3rd drive, the large cleaning path was again found (Figure 4.19(d)), and this time cleared relatively well. The settings used here were kept the same for the rest of the job.

Due to the Avant’s articulated frame steering and requirement for more maneuvering space than a skid-steered machine, a different strategy for driving to the next clearing path was required than in the simulation (Section 3.5). In that case, 2 points were enough to define each clearing path, with the machine turning on the spot to drive between the paths’ starting points and to align itself.

In the articulated frame case, the original plan for maneuvering was to have a 3rd point for each path, in line with the other two, further back from the start of the path. This is illustrated in Figure 4.20(a). Here, the snow layer would be in the area labelled “clearing zone,” with the first clearing path being from point 1 to 2. Afterwards, the machine would reverse to point 3 and wait for the next path to be found. The machine would then reposition itself laterally while driving to point 4, clear the next path to point 5 and reverse to 6, thus repeating the cycle. This was tested before the experiment in GIMsim, a Hardware-In-the-Loop (HIL) dynamic simulator for the Avant 635 [97, 19], and found to work with a maneuvering space of 5 m (e.g. between points 1 and 3).

During the experiment, this spacing was increased to 6 m as a conservative measure, however this was found to be insufficient and the Avant had trouble aligning itself properly when driving the repositioning seg-



(a) Initial plan for waypoint generation to maneuver between clearing paths. (b) Improvised strategy used during test for more effective maneuvering.

Figure 4.20. Strategies for maneuvering between clearing paths. Map is dilated version of Figure 4.17(b) used by Avant 635 for obstacle avoidance.

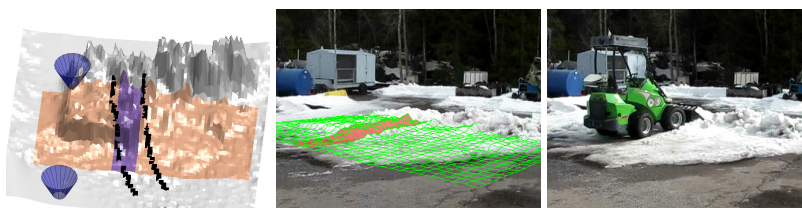
ment. On the suggestion of colleagues from IHA, the repositioning strategy was changed by locating the extra point further away and to the side, as in Figure 4.20(b). The longer driving distances along these arcs then allowed for effective repositioning between clearing paths.

After the 3rd drive, about half of the layer had been cleared. Figure 4.21 shows the remaining 5 drives needed to clear the rest of the layer. Drive 4 (Figure 4.21(a)) was largely on target, yet drive 5 (Figure 4.21(b)) drove too far right and missed the path. The same points were generated for drive 6 (Figure 4.21(c)) and the Avant drove a bit closer, removing some snow. Drives 7 and 8 (Figures 4.21(d) and 4.21(e)) were both also too far right but managed to clear enough to finish the job.

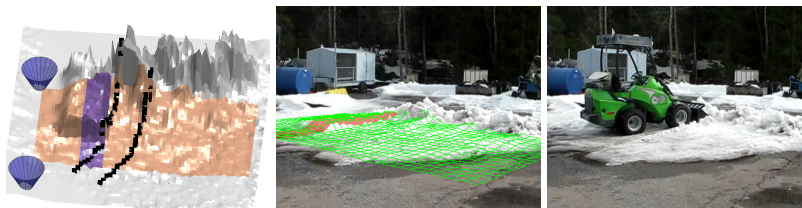
4.5.3 Discussion

During the last 6 clearing drives, the experiment was working as intended, with new commands being generated only from updated ground models and no changes to any settings. The Riegl and the search algorithm showed good performance, with changes detected even when only small amounts of snow were removed, such as after drive 6. Furthermore, in addition to finding “heavy” clearing paths into the main layer for most of the job, the narrower “cleaning” path resulting after drive 2 was also found. The overall concept was therefore validated by the experiment.

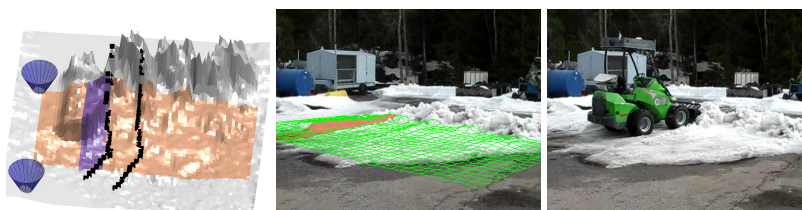
The fact that the last 4 drives were all too far to the right indicated that the change to the lateral offset made after drive 2, from 3 m to 2 m to the left, may have been too great. Ideally, more adjustments could have been made and then a new layer prepared, however after the first few



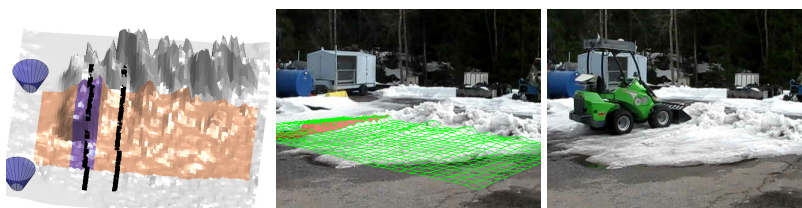
(a) 4th clearing drive, on target.



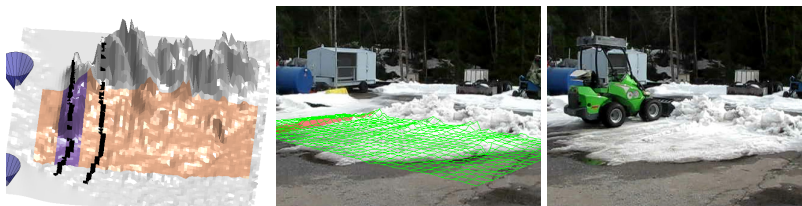
(b) 5th clearing drive too far right, missed path.



(c) 6th clearing drive mostly too far right but removed some snow.



(d) 7th clearing drive, slightly too far right.



(e) 8th clearing drive, mostly on target.

Figure 4.21. Clearing drives 4-8, all with same settings. The last 4 drives were generally too far to the right, indicating that the lateral offset may have been reduced too much after drive 2.

drives it was already late afternoon, the temperature was dropping and the experiment needed to be wrapped up.

The driving error for the Avant was confirmed to be about 0.5 m, given that with the final settings, some paths were cleared on target and some were missed by one scoop-width on the right side. Despite the combination

of errors from the coordinate transformation and driving, the job was still feasible since for most clearing drives at least some snow was removed. In total, 8 drives were needed to clear the area (not including the initial practice drive and the 1st attempts needed for drives 2 and 3). The layer was initially 5.5 m wide, and the bucket 1.04 m wide, therefore ideally it should have been possible to clear the area with 6 drives, depending on the amount of spillage.

The graphical job planning tools were easy to use during the job, and changes to the clearing area (e.g. after the 1st and during the 2nd drive) could quickly be made by clicking and dragging the corner cones of the Area Tool. The experimental procedure was slowed down however by several manual steps, such as reading coordinates from one laptop and typing them into the other, and operating the Riegl to obtain updated scans. Transitions between driving and bucket actuation segments were also commanded manually, though this was partly for safety reasons. For a system that operates hands-free after the initial plan is made, the scanning and transitions would all need to be automated. The main result demonstrated here however was that driving points were automatically generated based on changes in the workspace, and that all driving and bucket actuation was computer-controlled without direct teleoperation.

One lesson learned was the potential benefit of obtaining the coordinate system transformation faster, to allow more time for the actual experiment. The procedure used here was rather time consuming and also difficult to endure in the cold temperatures. One solution could be to mark the position on the ground where the tripod was standing so it could be used again, but would limit the workspace to the field of view. A mobile platform with self-localization capability could offer more flexibility by providing advantageous views as a worksite changes. Another approach would be to build ground models using laser scanners mounted on the Avant itself, thus benefiting from its localization information. This is in fact the method that was used in the final two outdoor tests in Sections 5.3 and 5.4. First, however, two more manually-driven experiments are presented in Sections 5.1 and 5.2.

5. Outdoor Tests Using Gravel

This chapter presents the next four outdoor tests which used gravel as the ground material. The first two were manually driven, to again analyze area filling and to investigate the possibility of developing an automated pile loading strategy. Finally, two computer-controlled pile transfer experiments are presented as a full proof-of-concept using the robotic Avant 635. Table 5.1 summarizes these tests, with more details and analysis in the following sections.

Table 5.1. Summary of outdoor tests using gravel.

Test	Date	Location	Job	Machine	Control
5	28.07.2012	Otaniemi	Area filling	Avant 320	Manual
6	05.07.2013	Hervanta	Pile loading	Avant 635	Manual
7	02.10.2013	Hervanta	Pile transfer	Avant 635	Computer
8	21.10.2013	Hervanta	Pile transfer	Avant 635	Computer

5.1 Outdoor Test 5: Manually-Driven Area Filling

July 28, 2012, Aalto Campus, Otaniemi, 25 °C

In this experiment an area was filled to test generating commands using the Virtual Pile tool and the search algorithm from Section 3.4.5, similar to Test 2 in Section 4.2. The main difference here was that more than one scoop-load was deposited at each dumping location, building up a higher pile, and also that gravel was used instead of snow.

Figure 5.1(a) shows the workspace, which was a storage lot in Otaniemi. The ground surface was compact mixed gravel and soil, and the material used for the test was taken from a pile of dry, uniform small-sized gravel. The apparatus was the same as in Tests 1-3, with the skid-steered Avant 320 and Riegl laser scanner being used. For this test a small generator



(a) 1st filling action (1st location).



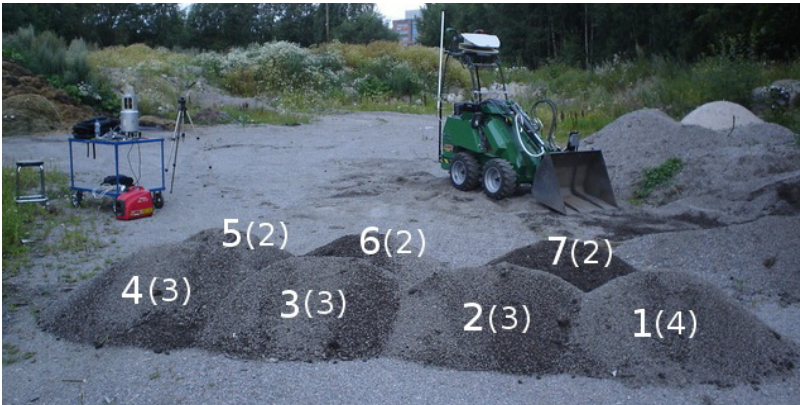
(b) 13th filling action (4th location).



(c) 17th filling action (6th location).



(d) 19th (final) filling action (7th location).



(e) Final pile from far end, with dumping locations numbered (number of deposits per location in parentheses).

Figure 5.1. Outdoor area filling test with gravel; 19 filling actions over 7 locations (4 along back row, 3 in front).

was used to supply electricity for the Riegl and laptop as no nearby power source was available.

The job followed the same approach as in Tests 1-3, i.e. that no detailed predetermined procedure was followed. The general plan was to fill gravel at a few locations along a line, with a few scoop-loads at each location resulting in a uniform overall height. The filling area had to stay within the field of view of the camera and laser scanner, which were kept stationary. A procedure was then followed which seemed logical - how a person might

typically perform such a job. First, one row of 4 piles was deposited from left to right, with the first pile consisting of 4 scoop-loads and the rest 3. The first and last filling actions in this row are shown in Figures 5.1(a) and 5.1(b).

After this, time was still available to continue, so another row of 3 piles was deposited in front, from right to left, each consisting of 2 scoop-loads (see Figures 5.1(c) and 5.1(d)). These were wider than the first 4 piles, and almost had the same total width as the first row. A total of 19 filling actions were thus made, with all 7 piles approximately 0.46 m high. Figure 5.1(e) shows the final state from the far end, with the dumping locations numbered and the number of deposits at each location in parentheses. The different numbers of scoop-loads at some locations was due to variation in the amount of gravel in each scoop-load, and also because some piles were built up on the slopes of previous ones and needed less gravel to reach the same height.

New scans were made after each deposit, which were then processed afterwards to build ground models that could be used with the Virtual Pile tool. The goal was to see if the dumping location algorithm could automatically determine when each location became full and suggest where the next location should be. Figure 5.2 shows the results.

As in Test 2, the first step was to find a Virtual Pile which approximately matched the final combination of piles. Here a slope of 30° was used, which was estimated as the repose angle of the gravel after calculating the slope of the final ground model along a few linear segments. The Virtual Pile which was then used had base dimensions 5.1 m wide by 2.2 m long and a height of 0.38 m, shown in Figure 5.2(x) with the final ground model.

The same scanning strategy was used as in Sections 3.4.5 and 4.2, with a linear segment approximately as wide as the bucket (0.8 m) which begins at the back left corner of the top section of the Virtual Pile, with a 0.1 m extension at each end. A dumping location is again considered full if only one ground point in the scanning line rises above the desired heights.

Figures 5.2(a)-5.2(c) show how the first location is available for the first 3 deposits. After the 3rd deposit, the present analysis already considers the location full (Figure 5.2(d)), however during the test a 4th load was deposited here (Figure 5.2(e)). After this the line begins scanning for the next location, moving along increments of 0.1 m.

Figures 5.2(f)-5.2(h) show how the 2nd location is available for loads 5-7, after which the search correctly determines the location to be full (Fig-

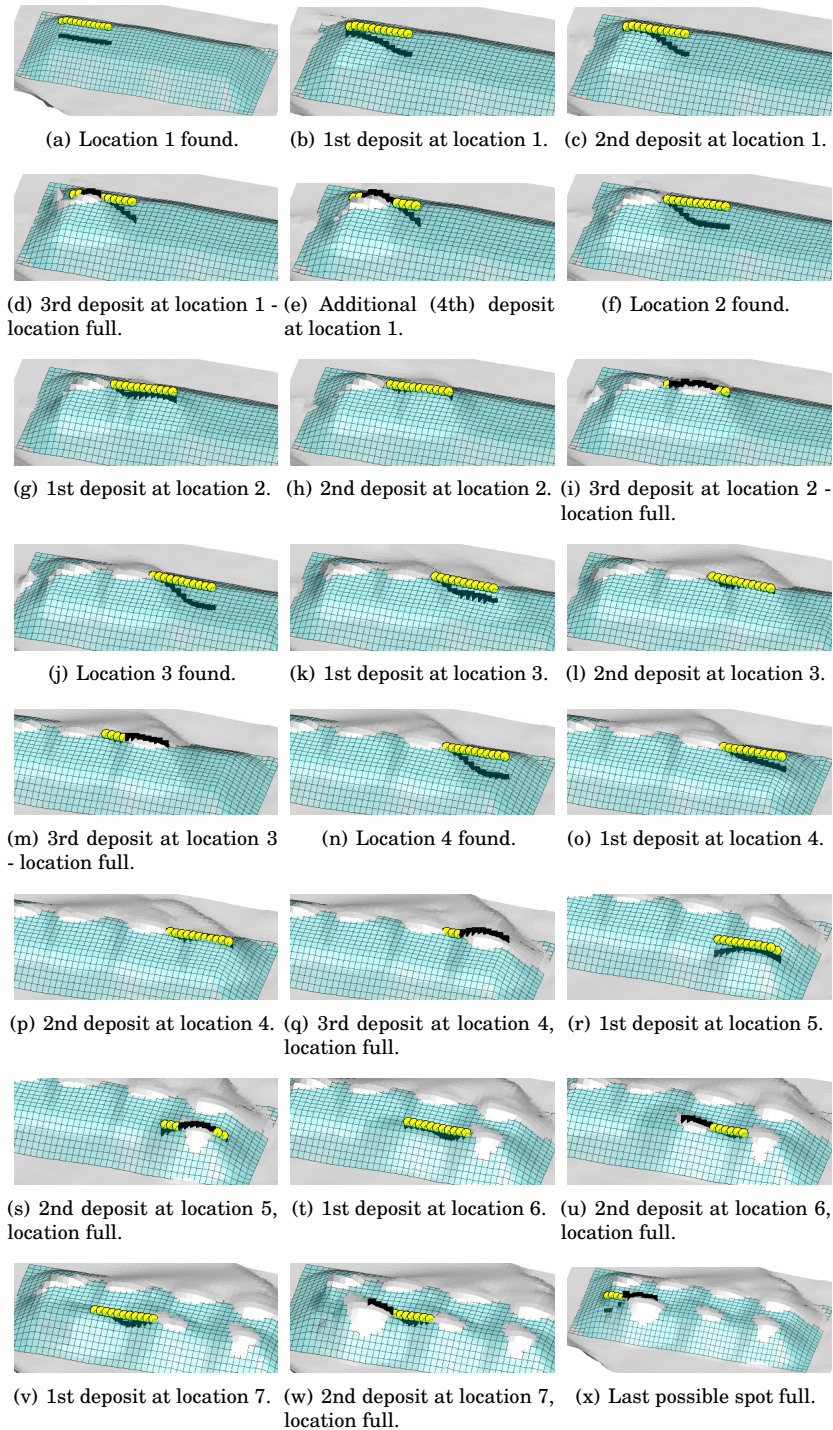


Figure 5.2. Automatic tracking of pile filling job: next location found by comparing ground heights along line segment with corresponding Virtual Pile heights.

ure 5.2(i). Figures 5.2(j) to 5.2(q) then show how this pattern is repeated for the next two locations, filling up the row.

With the first four locations filled up, a step of 0.7 m towards the front was used to find the first free location in the next row, this time searching from right to left. As with the first filling test in Section 4.2, the row step was one search parameter which was modified from the default of 0.2 m used in the simulations (Section 3.4.5), in order to match how the experiment was performed. A value of 0.2 m would have been able to find free locations for this experiment, however they would not have corresponded closely to those which were used. The search parameters within the row, however, are the same as before.

Figure 5.2(r) shows the 5th location (the first in the new row) after the first load has been deposited there, which is still available, after which it is filled by a 2nd load (Figure 5.2(s)). Figures 5.2(t)-5.2(w) then show the same pattern for the last two locations. Here it is evident that the last 2 suggested dumping locations do not match those actually used very closely in the lateral direction, due to the wider spacing between the sub-piles of the 2nd row. Nevertheless, the algorithm still correctly determined when each of the locations were full. Figure 5.2(x) then shows how the last possible location along the row is full at the end of the job.

In the final ground models a long slope can be observed extending from the back of the pile, which does not correspond to the shape of the actual pile shown in the photo. The reason for this is that the region behind the pile was occluded in the laser scan, and is filled in using the nearest-neighbour averaging scheme. Here it does not affect the results, since the scanning only occurs along the top of the pile, however to obtain a more accurate pile model, another method which attempts to more accurately predict the ground heights in occluded areas could be used. One example is the Gaussian Process (GP) method of Plagemann et al., which was shown to outperform an interpolated elevation grid for planning quadruped robot trajectories on uneven ground [98].

5.1.1 Discussion

The main result of this experiment, as for Tests 1-3, was the verification that scans obtained by the Riegl laser rangefinder could be used with the search algorithm to generate commands which would allow the job to be performed autonomously. Except for the case after the 3rd deposit at the first location, the algorithm was able to track the job correctly for all deposits, determining when the locations were still available and when they were full.

Building upon the results of Test 2 (the snow filling test in Section 4.2), this experiment showed that the concept works with granular materials, also when building up higher piles with several deposits at the same location. If a robotic machine had been following the commands generated here, the resulting pile would have no doubt looked different, since the suggested dumping locations did not correspond exactly to all the actual ones used. More important, however, was that the algorithm could decide when a location was full, thus it would be able to guide a machine to a new location when necessary.

As with Test 2, trial and error was used to find Virtual Pile dimensions and a row step which would maximize the number of deposits that could be tracked correctly, with 18 out of 19 the best result that was obtained. With more modifications to the search parameters, it was found to be possible to track all 19 deposits - this involved scanning with 2 parallel line segments and using a higher threshold of points to determine when a location was full. Here, however, the same single-line segment and threshold of 1 point was used as before to demonstrate a more general result.

When the job was performed, the 2nd row was begun far enough forwards to make a clear new row, which seemed logical at the time. If the goal had been to use up as little ground area as possible to store the material, the next row should have begun closer to the crest of the previous, to try and create a plateau of uniform height. In this case then the default row step of 0.2 m may have worked.

The next section presents the final manually-driven test which investigates the more difficult task of pile loading, before the full robotic pile-transfer demonstrations in the last two sections.

5.2 Outdoor Test 6: Manually-Driven Pile Loading

July 5, 2013, TUT Campus, Hervanta (25 °C)

The focus on snow clearing and area filling jobs thus far was due to the fact that they could be carried out using a combination of driving to specified points and joint actuation to preset positions. This makes them easier to demonstrate with a robotic machine; automated filling, however had so far not been tested.

The inclusion of pile loading tasks was until now avoided since developing a bucket filling controller was not the main focus of the research, and it was assumed to be too difficult to automate as it often involves complex

combinations of drive throttle and bucket joint movements. Human operators intuitively utilize these combinations in an attempt to fill a bucket to the maximum and/or to wrest a load free if the pile material is compacted or non-homogeneous.

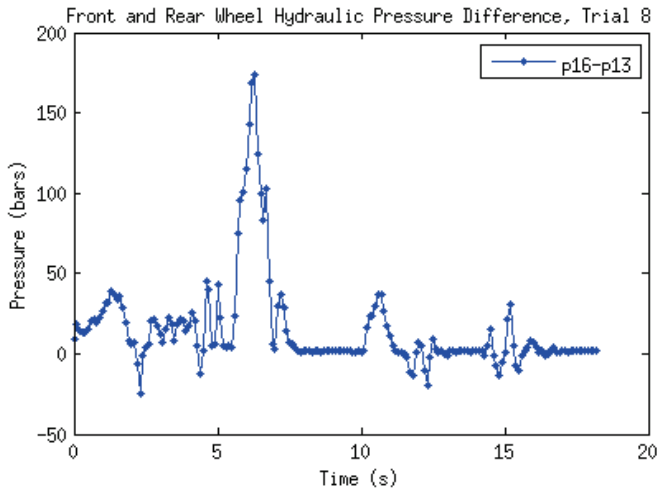
It was surmised, however, that for homogeneous materials such as gravel or dry soil, a simple loading strategy of driving straight into a pile with a preset bucket position may work to fill the bucket reasonably well. The main challenge, then, would be knowing when to stop driving. For the Avant 635, this might be determined by monitoring the available onboard sensors to detect either when the vehicle has stopped moving (due to the resistance encountered when pushing the bucket into the pile), or when the resistance has increased past a certain threshold. Since all bucket joints and also the driving is hydraulically-powered, the resistance encountered by the machine should be observable as increased pressure readings via the numerous hydraulic sensors. Other sensor information available includes inertial measurement, GPS positioning, wheel odometry and joint angles.

To investigate this strategy, manually-driven loading tests were conducted at the TUT campus in Hervanta with the Avant 635 while recording all available sensor signals. A pile of dry, clay-soil mixture was first built up, after which ten trials were conducted whereby the machine was driven at low speed into the pile, with the scoop low and level with the ground. Three of the trials encountered problems due to the pile being too small, with the machine pushing through the pile. The other seven, however, produced a consistent outcome, with the scoop penetrating into the pile, and the subsequent resistance against the scoop causing the vehicle to stop and one or more wheels to begin skidding. Shortly after any wheels began skidding, the driver released the engine throttle. An image from the 8th trial is shown in Figure 5.3(a) at the point where the front right wheel begins to skid.

After studying the resulting sensor signals, it was noticed that one value, the difference between the front and rear wheel hydraulic pressure, produced a characteristic spike every time the wheels began skidding upon contact with the pile. A plot of this signal corresponding to the 8th trial in Figure 5.3(a) is shown in Figure 5.3(c). While driving, the absolute value of the pressure stays approximately below 50 bars, with the spike reaching 174 bars. In the other 6 successful trials, a similar profile was observed, with the spike ranging from 106 to 268 bars (see Figure 5.4).



(a) Front right wheel skidding due to resistance from pile (Photo: Antti Kolu). (b) Load extracted with simple raising of bucket (Photo: Antti Kolu).



(c) Difference in front and rear wheel hydraulic pressure; spike of 174 bars due to wheels skidding on contact of bucket with pile.

Figure 5.3. Manually-driven pile loading test with Avant 635 (trial 8 of 10).

It was decided to try using this signal to program an automated scooping behaviour, whereby the machine would stop driving when the pressure rose above a threshold around 100 bars. This behaviour would then make it possible to test full pile-transfer work cycles without any direct teleoperation, and thus to further validate the job planning tools and algorithms developed.

A secondary purpose for these trials was to confirm that a simple raising of the bucket would be sufficient for extracting a scoop-load of material from the pile. This turned out to be the case - in all the trials the bucket was successfully extracted in one continuous motion by raising the boom mostly via the rotary base joint at the top of the boom (see Figure 5.3(b)), with perhaps some simultaneous extra pitch rotation of the bucket. Fine adjustments and maneuvering were not required.

For the case where a pile becomes too small to stop the machine, which

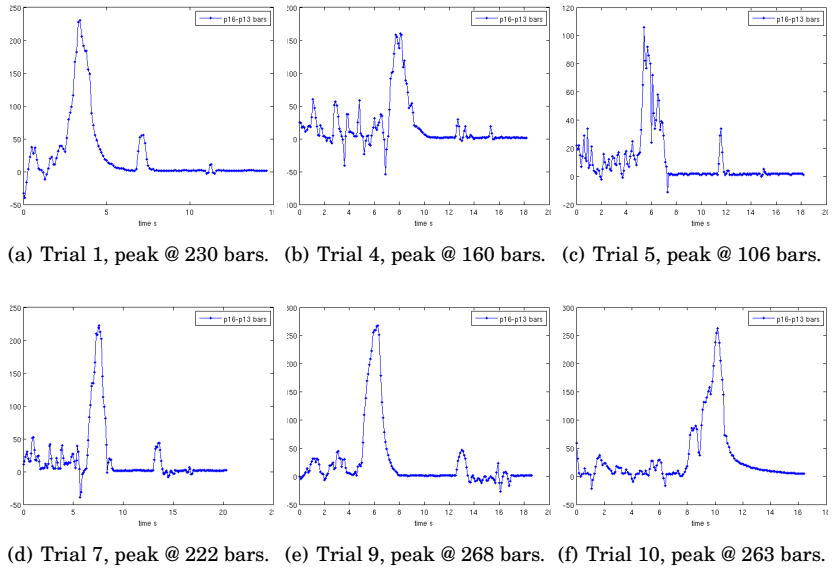


Figure 5.4. Difference in front and rear wheel hydraulic pressure for other successful trials; characteristic spikes occur when wheels begin skidding.

was encountered during these tests, the same strategy employed in the simulations (Section 3.4.4) could be used. In that strategy, if the bucket actually reaches the Scooping Destination point, which would be located inside or just beyond the pile, the machine stops and raises its bucket to collect the partial scoop-load. When the whole specified scooping area is sufficiently level, the job would be deemed complete.

One potential problem that was noticed during this test was that since the Avant 635 bucket is narrower than the chassis (including the tires), if scooping is repeated in the same direction, the loader will begin to drive over the streaks of material left over on either side before the bucket is filled (a version of the problem in Figure 3.20). As discussed earlier in Section 3.4.7, it would then be necessary to size the Scoop Area (Section 3.4.6) such that all material can be loaded before the wheels could encounter any part of the slope.

The results of this experiment were used to develop a scooping controller presented in the next section (and shown in Figure 5.8(c)), which was able to stop the machine during a scooping action as desired by monitoring hydraulic pressure values. A simple raising motion which was programmed then also succeeded in removing an amount of material which was considered sufficient. This strategy represents a simple alternative to others which have been developed, but would need to be tested in terms of bucket fill and mechanical wear [67, 70, 63].

5.3 Outdoor Test 7: Robotic Pile Transfer

October 2, 2013, TUT Campus, Hervanta (8 °C)

The purpose of this experiment was to achieve full pile transfer work cycles under computer control, whereby the Avant 635 would load material using a newly developed scooping behaviour and dump it in a designated area nearby, as in the simulations of Section 3.4. The main goals were to demonstrate the following:

- the graphical tools and algorithms could be used to plan a job and generate commands for pile transfer work cycles on an newly acquired ground model
- ground models from a new onboard laser scanner could be used to plan the job
- the new scooping behaviour could be used to load gravel
- the Avant 635 autonomous driving and joint control enable computer-controlled work cycles with no direct teleoperation.

The new laser scanner was a forward-mounted Sick LMS111, angled downwards, which scans in a plane and covers a portion of the ground in front of the machine using a sinusoidal pitch angle tilting motion [21]. When mapping starts, scans are integrated as the machine drives, building up a 3D ground model with 0.2 m grid resolution which is stored in a text file. The ground models are in the same reference frame as the Avant's autonomous navigation, meaning that no time-consuming coordinate transformation was required as with the Riegl scanner in Test 4. Another update was the Real Time Kinematic (RTK) GPS receiver being used by the Avant for positioning. With an accuracy of 0.1 m, this provided a significant upgrade over the previous version used in Test 4 which had an accuracy of 0.5 m. Figure 5.5 indicates the tilting laser scanner and RTK GPS antenna on the Avant.

The experiment was performed on an unpaved yard with a dirt/gravel surface, where a large pile of gravel was available for testing (see Figure 5.5). In order to have visual feedback of the machine while preparing and sending commands, a mobile workstation was set up on site (shown in

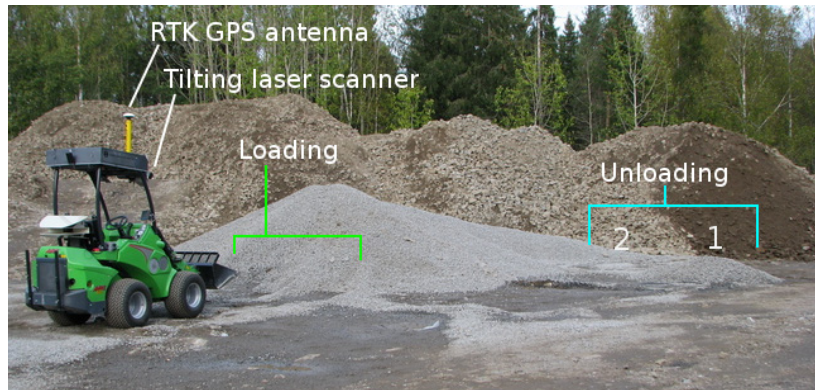


Figure 5.5. Test 7 workspace, with loading region of pile and two unloading points indicated.

Figure 5.8(a)), with power from the nearby covered test hall and portable antennae to establish a Wireless Local Area Network (WLAN) for communicating with the Avant.

The first step once all the equipment was set up was to build up a surface model of the workspace on which a job plan could be made. This was done by recording laser scanner data for a few minutes while manually driving the Avant 635 in the vicinity of the target pile. Figure 5.6(a) shows a raw scan, with the location of the pile indicated at left. Most of the pile is not modelled since the laser scanner is angled downwards and mostly sees the ground, however the pile can be identified by the gap in the scan and the lower portion of its slope. Holes in the raw scan were then filled with nearest neighbour average values to create a full surface model.

The surface could then be used to make a high-level plan as in Figure 5.6(b), with the Area Tool used to mark a 3 x 3 m area which included most of the visible front bottom slope of the pile. A Virtual Pile was specified which was relatively wide, short and low (base area 3 x 2 m, 0.5 m high), since it was hoped that if the first dumping location filled up after a few cycles, this might be detectable with an updated laser scan, allowing the search algorithm to automatically find the next location.

The high-level plan was automatically interpreted to generate the required driving points, shown in Figure 5.6(c). During testing, the plan needed to be generated multiple times in order to separately test driving, loading and unloading, and also in case the system needed to be reset, which was sometimes necessary if the Avant had to be restarted. In order to save time in setting up the plan, the directions from which to approach the scooping and dumping areas were preset to be on the front side where

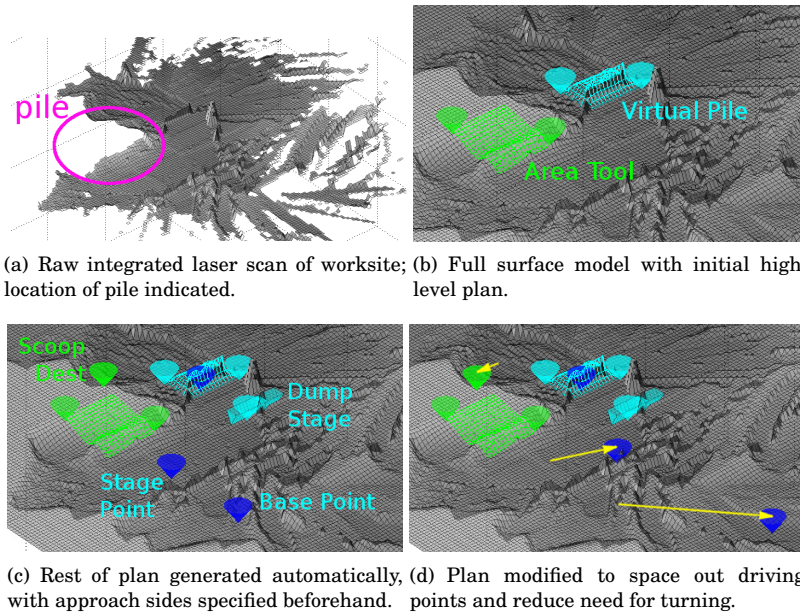


Figure 5.6. Ground model; pile transfer plan specification, automatic interpretation and user modification.

plenty of maneuvering space was available, whereas by default the interpreter would select the sides between the two areas, which were impractical in this case. Normally the plan would only be generated once, with modifications, if necessary, also being made only once. The Stage point and Dump Stage are both offset 3 m from their respective surfaces, to allow enough maneuvering space for the Avant. The Scooping Destination was moved 1 m further back from the highest point in the area, to ensure that it was far enough from the front of the pile to be used for several repeated scooping actions if necessary.

An extra driving coordinate called the Base point was added, 5 m in front of the Stage point, due to the steering requirements of the articulated-frame Avant 635. As it cannot turn on the spot, it needs to drive a V-shape path to change heading, which the Base point defines between the Stage and Dump Stage points. To make the maneuvering easier, the Stage and Base points were manually shifted by clicking and dragging (indicated by the arrows in Figure 5.6(d)), to narrow the angle defined by these 3 points. The Scooping Destination was also shifted a small amount, since it appeared too close to the right edge of the pile. These manual changes were made largely out of caution, to limit the driving corridor of the Avant during the experiment. In an open, less restricted workspace, the automatically-suggested driving points could have first

been attempted without any manual modification.

The five points defining the job are shown again in Figure 5.7. In order to command the Avant to drive between these points and to actuate its bucket, it was necessary to interface the Matlab job planning environment with the Avant 635 control software. The software used in Test 4 (GIMnet/MaCI and GMPP) was no longer easily compatible with the current software controlling the Avant, thus it was necessary to interface directly from the planning environment. This was simplified by the fact that Matlab and Simulink were being used at IHA to run the Avant, therefore it was relatively straightforward to add extra buttons to the planning window which called the appropriate Matlab functions, scripts and Simulink models. These were supplied by Dr. Reza Ghabcheloo from IHA, who assisted with the required modifications and overall integration work. In general, the situation was greatly simplified over that in Test 4 (Section 4.5), when two separate laptops were used: one for surface modeling and command generation, and the other for sending commands by typing in coordinates manually. Now everything could be done from the same laptop.

The new planning interface with a description of the extra command buttons is presented in Figure 5.7. Driving and scoop commands are sent in the same format as in Test 4 (Section 4.5), i.e. by specifying a desired end pose for the vehicle or bucket consisting of a 2D coordinate and heading or pitch angle, respectively. When a driving or joint button is clicked, a data packet with the desired pose is prepared and then sent via User Datagram Protocol (UDP). The basic functionality of the control interface is to click on a button to drive to one of the points or change the joint configuration, then confirm the command by pressing *Send*. Three joint configurations were programmed: *Scoop*, with the bucket low and level, *Raise*, which mostly lifts the bucket vertically to extract a load and move to a hauling position, and *Dump*, which pitches the bucket forward to unload any material inside.

Aside from sending commands, it was also possible to receive the current position of the Avant and display this on the planning surface by rendering an Avant model. When a driving point button was clicked, the planned path was also rendered before commanding the machine to drive it, which was a useful safety consideration.

This integration work was tested first in GIMsim, the Avant 635 HIL simulator featuring an interface with the virtual Avant which is the same

as with the real one [97, 19]. Testing then continued outdoors, where the Avant was commanded to drive between points and reposition the scoop. With the driving and joint commands verified to work properly, the next step was to test full pile transfer work cycles.

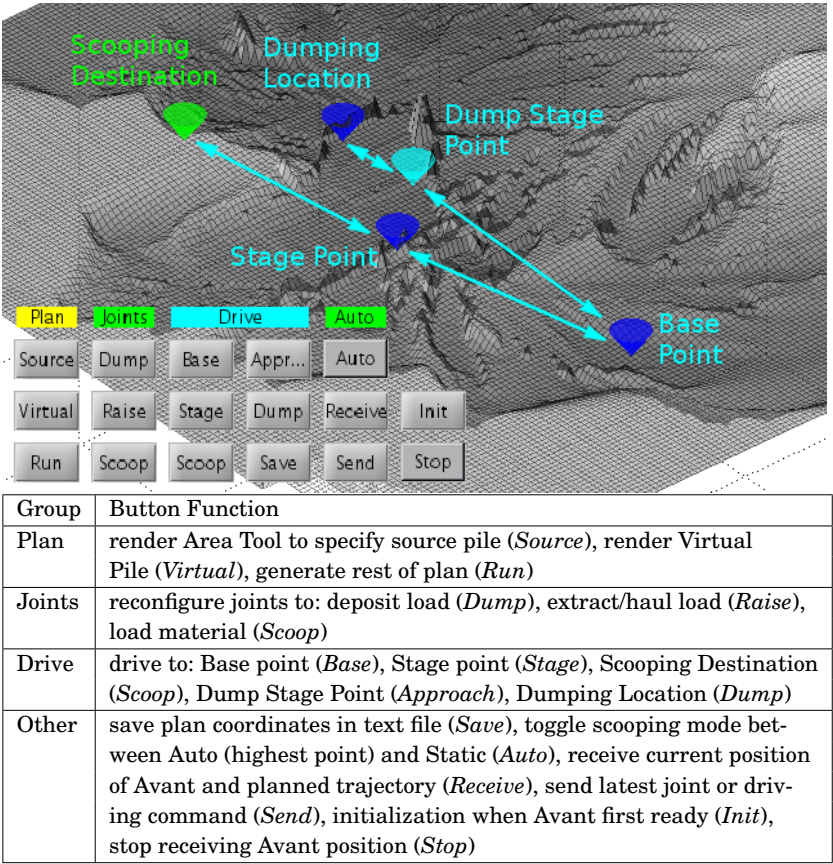


Figure 5.7. Five driving points defining job, with command button interface.

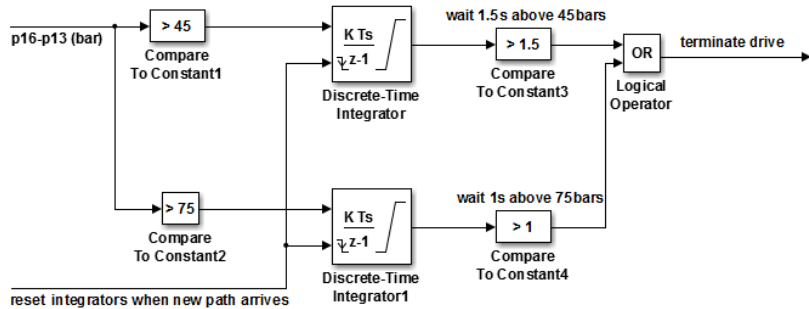
The final preparation step was to test the autonomous scooping behaviour. The initial idea, based on the results of the previous test in Section 5.2, was to lower the scoop to a preset position, drive towards the pile and stop when resistance against the bucket caused the wheels to begin skidding. This would be detected as the difference between the front and rear wheel hydraulic driving pressure passing a threshold around 100 bars. During testing however, it was found that the pressure values reached in this situation during computer-controlled driving were significantly lower than with a human at the wheel. When thresholds were set in the range of 100-150 bars, bucket resistance would stop the Avant while it continued driving with one or more wheels spinning.

The final strategy, devised by Dr. Reza Ghabcheloo from IHA, was to in-

tegrate the pressure values over time and stop the machine either when the pressure difference exceeded 45 bars for 1.5 s or 75 bars for 1 s. The Matlab Simulink block diagram of this logic is presented in Figure 5.8(c). The need to integrate for at least 1 s was so that brief spikes during regular driving would not terminate the drive, while the higher threshold for a shorter duration was meant to protect the machine and stop it sooner if very high pressures were encountered. This strategy was found to work well, with the Avant stopping consistently after the bucket had penetrated the pile (as in Figure 5.8(b)) and brief wheel skidding had occurred. During this trial phase, bucket extraction and unloading were also tested, with five scoop-loads of gravel being deposited at location 1 in Figure 5.5.



(a) Mobile planning and command station. (b) Testing of automated scooping behaviour.



(c) Logic for automated scooping developed by Dr. Reza Ghabcheloo from IHA.

Figure 5.8. Test 7 command station and development of automated scooping logic.

With all the points specified as required, and the scooping behaviour functioning, three full work cycles were then performed under computer control, whereby a scoop-load of gravel was collected and deposited (at location 2 in Figure 5.5). All transitions between driving actions and joint repositioning were manually commanded by clicking the appropriate buttons on the planning interface. This was partly for safety reasons, and also since the capability of autonomously deciding when an action was complete and when the next could begin had not yet been developed. The

main result here was that no direct teleoperation was used during driving and bucket actuation. Since only three cycles were performed, it was not necessary to update the points for scooping and dumping. This was the goal of the next and final experiment. A new ground model was made after the 3rd cycle however, which was studied afterwards.



(a) 1st scooping action.



(b) 1st dumping action.



(c) 2nd scooping action.



(d) At Base point, 2nd cycle.



(e) 2nd dumping action.



(f) 3rd scooping action.



(g) At Base point, 3rd cycle.



(h) 3rd dumping action.

Figure 5.9. Robotic pile transfer test: 3 full scooping and dumping cycles (Photos: Reza Ghabcheloo).

5.3.1 Analysis and Discussion

This experiment was generally a success, since the primary goal of achieving full work cycles with all driving and joint actuation under computer control was demonstrated. Since the same static driving points and surface model were used during the cycles, the full functionality of the planning tools and algorithms was not utilized. Had more time been available, updated scans and plans based on changes to the worksite would have

been attempted, however when everything was working, it was getting late in the day and there was only time for 3 cycles during the main test.

The automated scooping logic worked well, though its long-term effectiveness is unknown, i.e. if it can maintain a level ground surface using a preset bucket position and if it can fill the bucket enough using the range of approach vectors that the High Point method would select.

Precision in driving and joint actuation can be observed by comparing Figures 5.9(d) and 5.9(g). In these figures the Avant has driven to the Base point on two occasions, with the joints in “Hauling” mode. By observing the alignment of the machine parts with the larger wheel loader in the background, it is evident that it has repeated the action closely.

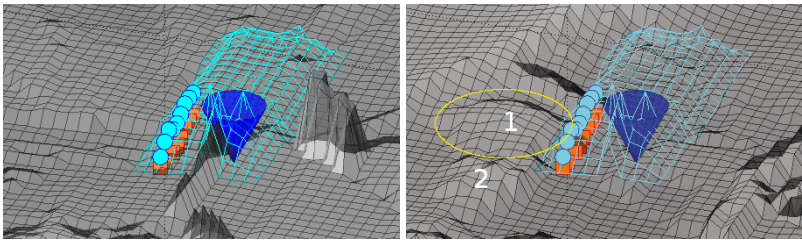
Driving precision was also observed in the placement of the bucket loads when unloading. Figure 5.10(b) shows the two resulting dump piles, with the lighter-coloured one on the right being the result of 5 dumping actions during the practice runs, and the darker one on the left being the result of 3 dumping actions during the main test. For each pile all deposits were placed on top of the previous loads, indicating good repeatability.

A surface scan made at the end of the job allowed the use of the Virtual Pile and Dumping Location commands to be checked. Figure 5.10(b) shows the initial ground model with the Virtual Pile used during the practice runs. The rows of points indicate the intended Dumping Location, with the inverted cone located approximately 1 m back to account for the predicted dump offset distance between the vehicle reference point and bucket position when unloading (as in Figure 3.11). Figure 5.10(c) shows the same plan with the surface model afterwards, with the ellipse indicating the resulting dump pile. It is evidently about 1 m in front of the intended location, though the lateral placement is correct. The same result is seen in Figures 5.10(d) and 5.10(e) from the main part of the test. This indicates that the driving accuracy was good, however the dump offset distance used was not long enough and should have been closer to 2 m to achieve unloading at the intended Dumping Location.

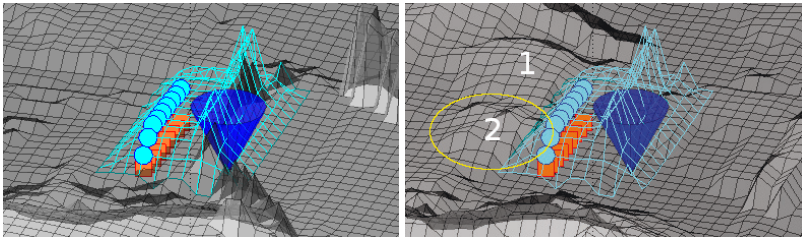
If the piles had been deposited inside the 0.5 m-high Virtual Pile which was used, it does not seem that they piles would have been high enough to fill the Dumping Locations. To see if a lower Virtual Pile could have been used during the test to obtain an automatically-updated Dumping Location, a lower, 0.2 m-high pile was rendered over the original ground height matrix, superimposed with the end surface model, shown in Figure 5.11. Since the tops of both piles are protruding through the Virtual Pile sur-



(a) End piles (Photo: Reza Ghabcheloo).



(b) Initial ground model; Virtual Pile and Dumping Location for practice trial. (c) Final ground model; 1st pile placed with good lateral alignment, off by approx. 1 m in forward direction.



(d) Initial ground model; Virtual Pile and Dumping Location for main trial. (e) Final ground model; 2nd pile placed with good lateral alignment, off by approx. 1 m in forward direction.

Figure 5.10. Pile placement compared with intended Dumping Location.

face, this confirms that updated locations could have been obtained using ground scans from the onboard laser scanner.

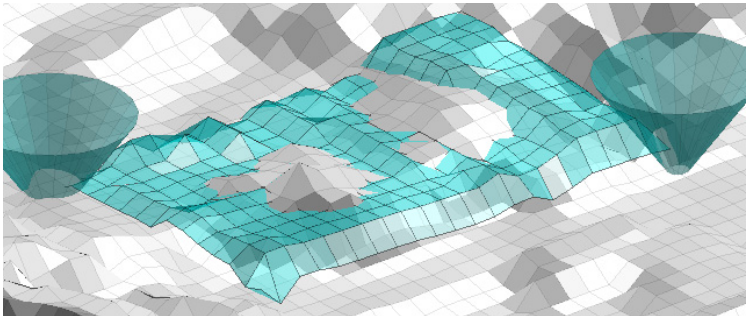


Figure 5.11. Verification that piles detectable with Virtual Pile tool.

Unfortunately, this analysis of the Virtual Pile and Dumping Location

was not done before the next test, which may have improved the results of the final experiment by determining the correct dump offset distance and Virtual Pile height to use.

This experiment was similar to that demonstrated by Koyachi and Sarata, which included 4 full work cycles, with scooping at the same location and unloading at a dump truck [68]. The level of automation here may have been lower, since the transitions were commanded manually, but one additional result was that the planning phase was demonstrated. Although making changes to the plans represented again a lower level of automation, one use of the graphical planning concept was demonstrated, which allowed the user to visualize the automatically generated plan and make modifications which were deemed necessary. The benefit of having a human in the loop when specifying the job was also demonstrated. Given that the initial surface model was quite rough and did not clearly show a pile shape, it was still evident where the pile was located due to the distinctive bottom contour of the slope surrounding the large gap.

5.4 Outdoor Test 8: Robotic Pile Transfer

October 21, 2013, TUT Campus, Hervanta, 2 °C

This experiment was performed 3 weeks after the previous test, at the same location and using the same equipment. The purpose was to repeat the computer-controlled pile transfer work cycles, this time with laser scans being made after each scooping and dumping action. With updated ground models from these scans, the goal was to see if automatic updates could be made to the Dumping Location, and possibly the Scooping Destination, as work proceeded based on the changing state of the worksite. Figure 5.4 shows the workspace mid-way through the test, with the loading area to the right and first two unloading locations to the left. The large gravel pile is in a different state than during the previous experiment.

As before, the first step was to manually drive the Avant in the workspace for a few minutes while building up a new surface model with the onboard laser scanner. Figure 5.13 shows the initial surface model (with 0.2 m grid resolution) and high level plan, similar to that in the previous test (Figure 5.7). One difference is that the approach directions for the Area Tool and Virtual Pile are perpendicular to each other. The Stage (2) and Base (1) cones were manually repositioned (indicated by the arrows) as before to create the elongated V-path for the Avant to realign itself. The Scooping



Figure 5.12. Workspace for 2nd robotic pile transfer test, with loading and unloading locations labeled.

Destination (3), Dump Approach (4) and Dumping Location (5) remain at their automatically-suggested locations.

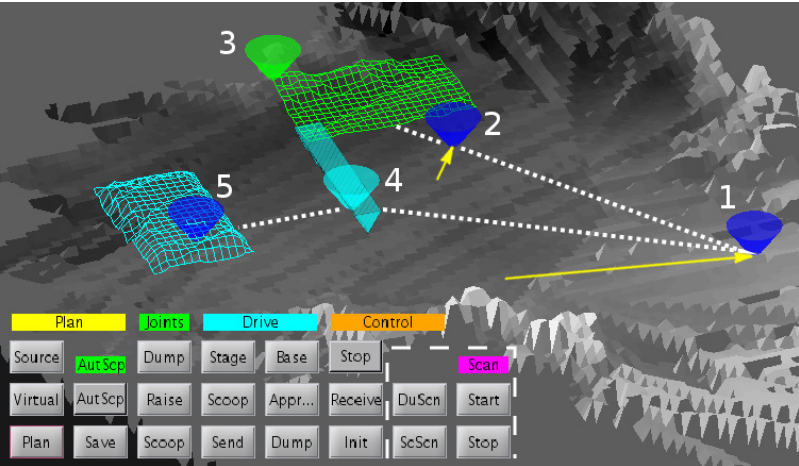


Figure 5.13. Initial surface with high-level plan. Arrows indicate that Base point (1) and Stage point (2) manually repositioned to create favourable driving path. Four new buttons for laser scanning added to control interface.

Four new buttons were added to the control interface, seen at the lower right in the figure. These are used to *Start* and *Stop* the laser scanner, and drive to extra Dump Scan (*DuScan*) and Scoop Scan (*ScScan*) points for scanning the dumping and scooping areas. These last two functions were not needed, however, since the Dump Approach and Stage points ended up being good locations for scanning these respective areas.

With the points ready, 4 full work cycles were performed. As before, all transitions between driving segments and joint configurations were man-

ually commanded, with the actions carried out under computer control without any direct teleoperation. Extra stoppages, joint movements and driving segments were required in order to make the updated scans after each action, with the bucket lowered so as not to occlude the scan. Figure 5.14 includes 4 screenshots from the first cycle, showing the scooping action, returning to the Base point, dumping action and stopping at the Dump Stage point afterwards to make a scan.



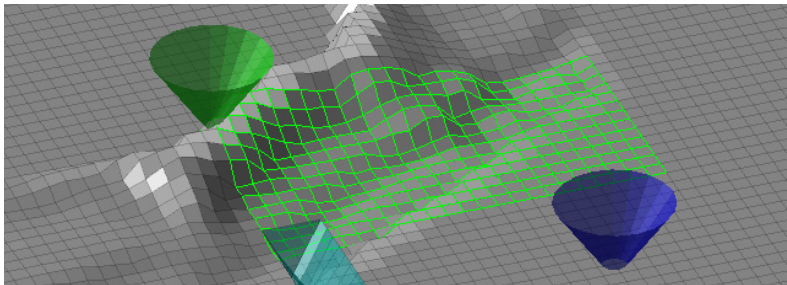
Figure 5.14. 1st pile transfer cycle of Test 8.

5.4.1 Analysis and Discussion

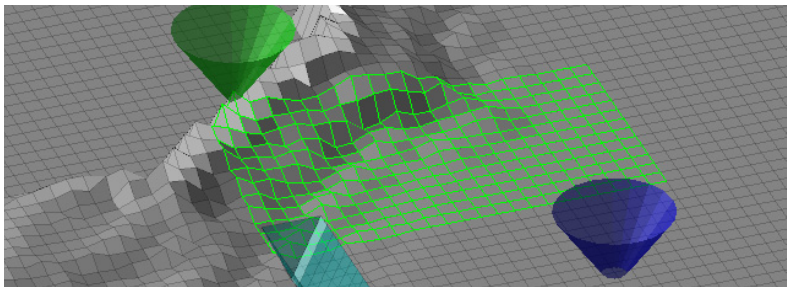
For the scooping part of the experiment, an attempt was made to place the Area Tool such that only part of its surface at the rear would be on the pile slope. This was so that the scooping actions would have a better chance of causing slope collapse which could be detected and result in an automatic update in the Scooping Destination. Given the rough shape of the surface model, however, it was difficult to see exactly where the slope started, so this placement could only be estimated. After the 2nd scooping action the Area Tool was shifted further away from the pile.

Figure 5.15 shows updated scans that were made of the scooping area after the 2nd (top) and 3rd (bottom) scooping actions. From the shape of

the slope, it is difficult to discern whether a scoop-load has been extracted. The Scooping Destination, marked by the cone at the left, does however shift slightly to the right for the 3rd scooping action. This could possibly be due to slope collapse, or also roughness in the ground model. Data from the other scooping actions are similarly inconclusive.



(a) After 2nd scooping action.



(b) After 3rd scooping action

Figure 5.15. Updated ground models in scooping area with Scooping Destination (cone at left) automatically located at highest point; other cone marks Stage point.

Figure 5.16 (at left) shows a close-up photo of the unloading area from the front, with the three dumping locations labelled. Two scoop-loads were first dumped at location 1, followed by the 3rd at location 2 and 4th at location 3. At right is a corresponding surface model, with the pile shapes discernible in the centre.

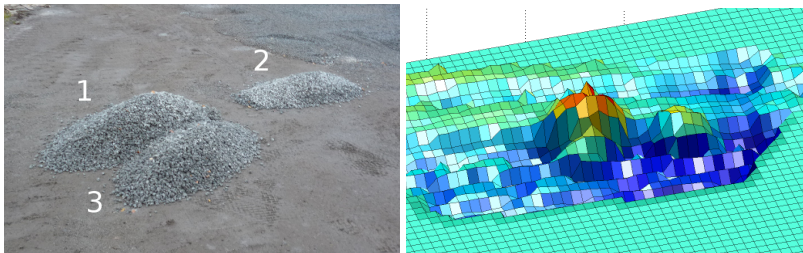


Figure 5.16. Photo and surface model of end piles, viewed from front.

To guide the filling actions, the intention was to specify a low Virtual Pile such that one or two scoop-loads of material would fit in one Dumping

Location, which would be detected and trigger an update. The minimum pile height being used at the time of the experiment was 0.33 m, based on the default scoop capacity of 0.15 m^3 , thus this height was used here. Figure 5.17(a) shows the Virtual Pile from the far end after the 2nd deposit. The rows of points mark the desired dumping location, with the dark blue inverted cone marking the corresponding driving location for the Avant, about 1 m back due to the estimated dump offset. The same dumping location was commanded for the 1st and 2nd deposits, with both ending up in the same location slightly too far back and to the right. Despite this apparent error in the plan, the repetition indicated good driving precision on the part of the Avant.

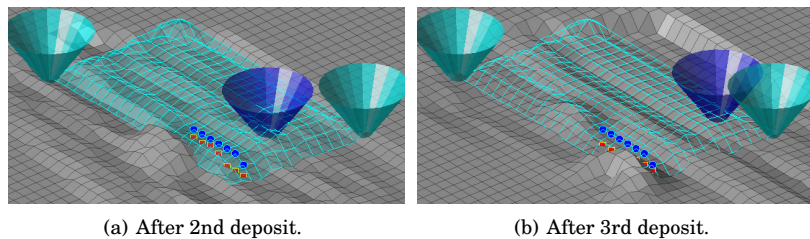


Figure 5.17. Searching for Dumping Location with Virtual Pile, viewed from back.

The fact that the first two deposits were too far back was likely related to the problem mentioned in Section 5.3.1, that the 1 m dump offset was too short for the Avant 635. Unfortunately, however, this was not known at the time. The deposits being too far to the right was unexpected, since good lateral positioning was demonstrated in the previous test.

To compensate for these errors, the dump offset was increased by 0.8 m, with a lateral offset added to align the desired location with the observed location. Figure 5.17(b) shows the plan and result for the 3rd deposit. The increased dump offset seemed to work here, with the shape of the deposit visible inside the Virtual Pile. The lateral offset did not work however, and an error to the right was repeated. The 4th deposit, however, was placed accurately. This followed the same plan as in Figure 5.17(b), and as Figure 5.16 shows, it was deposited on target at the foot of the larger pile from the 1st two deposits.

From Figure 5.17(a), it looks as though the higher pile from the 1st two deposits is just about as high as the Virtual Pile, yet perhaps would not have been detected as filling the 1st location if it had been on target. As in the previous test, it was not known that a lower Virtual Pile should have been specified to increase the chances of detection. In Figure 5.18,

for example, it is shown how a lower 0.2 m pile is able to detect the first larger pile and even the 2nd one consisting of only one scoop-load.

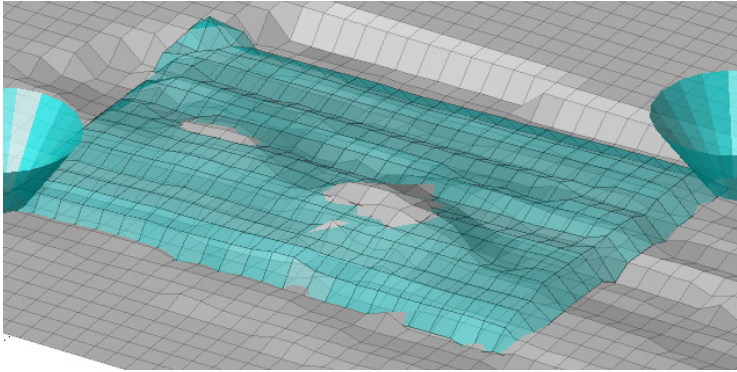


Figure 5.18. Lower Virtual Pile with height of 0.20 m able to detect 2 piles.

Overall, this experiment produced mixed results. The goal of obtaining automatic updates to the job plan based on updated ground models was not achieved for the dumping part of the job, while the results for scooping were inconclusive. After analyzing the data, it was confirmed that changes in ground height resulting from even one dumping action can be detected with the onboard laser scanner if the Virtual Pile is low enough. Another positive result was that the automated scooping logic developed during the previous test (Figure 5.8(c)) worked reliably for all scooping actions.

A challenge faced during this experiment was that a significant part of the day was spent doing preparation and setup, with a relatively short time available for the actual testing. Ideally, more than four cycles would have been included, since some learning and adjustment to the plan, procedure and dump offset was also done during these four cycles. More cycles would also result in greater cumulative changes with a higher chance of obtaining automatic updates.

Finally, one aspect of the experiment which validated the overall approach was that the worksite was in a difference state than during Test 7, with the large source pile in a different location and having a different shape. Since the job planning tools need no prior map and can be specified by the user on the current model of a changing site, the job plan was generated rather quickly once the setup was finished.

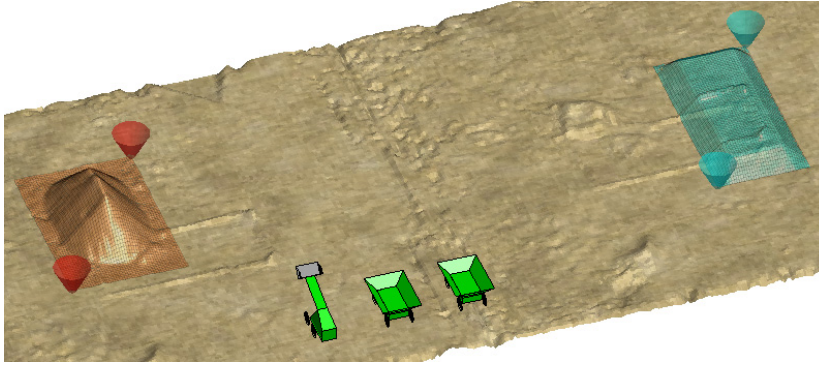
6. Multiple Vehicle Earthmoving

The earthmoving jobs presented so far have been limited to those performed by one wheel loader. This was done to start with the simplest cases and to focus on jobs which could be demonstrated in outdoor tests with one of the available Avant loaders. The job planning tools, however, are not necessarily specific to one wheel loader and can be followed by several machines as long as they are properly scheduled, and also by different types of machines as long as the correct commands can be generated. This chapter presents concepts showing how the tools can be extended in this way. First, the pile transfer job from Section 3.4 is extended to include dump trucks in Section 6.1. This is followed by a large-scale hillside excavation job in Section 6.2 motivated by the MHP construction plan from Chapter 1, which requires a conveyor belt spreader and can include more than one loader.

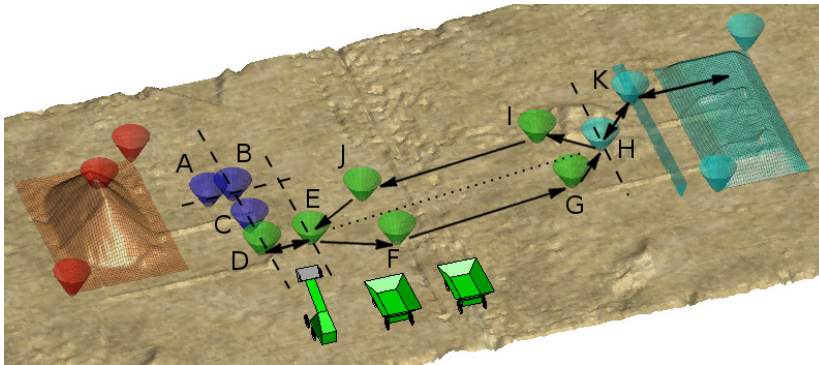
6.1 Pile Transfer with Dump Trucks

In this section the single-loader pile transfer simulation is extended by adding dump trucks, which can greatly improve the productivity of a wheel loader and is also a common industrial scenario. Determining the optimal combination of loaders and dump trucks for a job depends on several factors, including hauling distance, driving speeds, bucket volumes, load transfer time and unloading time. When comparing a single loader with a combination of one loader and one dump truck, a simple efficiency analysis shows that using the dump truck should generally improve productivity. This is because the hauling distance covered by the dump truck would have to be traversed N times by the loader to haul the same amount of material, where N is the number of scoop-loads needed to fill the dump truck.

With only one dump truck however, the loader stands idle during the hauling cycle. To further increase productivity by keeping the loader operating continuously, a minimum of two dump trucks is required, which is the scenario in this simulation. Three or more trucks would then be needed if the hauling distance is long enough such that one truck can be filled in less time than it takes for the other to complete a hauling cycle.



(a) High-level plan specifying source material and dump pile.



(b) Automatic interpretation of plan and generation of necessary driving points.

Figure 6.1. Planning for pile transfer job with one loader and two dump trucks.

Figure 6.1 shows the planning process that was developed for the same type of pile transfer job as in Section 3.4, this time with two dump trucks. Small-scale trucks were modelled which are also capable of turning on the spot, with a wheel spacing 1 m wide by 1.5 m long and a payload capacity of 1 m^3 . This was not based on any actual machine that was available, but rather designed for the Mars construction scenario presented in Section 6.2, where small-scale machines may be required due to the mass and volume constraints of planetary missions. The capacity of the loader bucket remains 0.15 m^3 .

First, the same high-level plan is made by the user (Figure 6.1(a)), with the Area Tool to specify the source material (to the left) and Virtual Pile

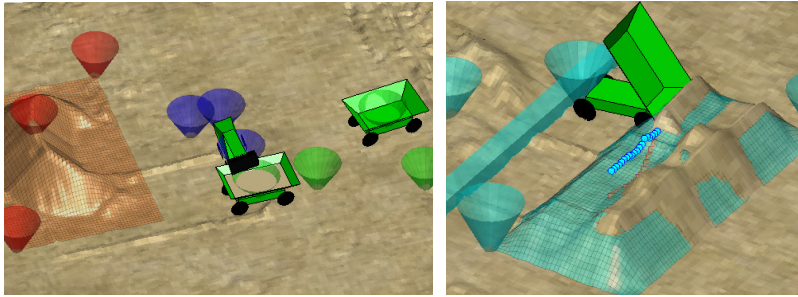
to specify the desired dump pile (to the right). This plan is then automatically interpreted to generate the driving points needed by the machines (Figure 6.1(b)). Point A is the Stage point used for beginning and ending scooping actions and point K is the Dump Approach point along the Dump Stage side used for filling the Virtual Pile - both of which are common with the single-loader planning from Section 3.4.3. All other points in between are new ones needed for the multi-machine case. These can be placed in two groups: points needed for load transfer between loader and dump truck (points B, C, D), and points representing the 2-lane road used by the dump trucks for hauling (points E, F, G, H, I, J).

Load-transfer points B, C, D are determined based on the load-transfer strategy: once the loader extracts a scoopful of material from the pile and returns to Stage point A, it reverses a small distance away from the pile to point B, turns 90° , and drives to point C to unload the material in the dump truck waiting at point D. Points B, C, D are therefore collinear, and perpendicular to line AB. The less the loader must turn, the faster it can return to scooping, thus positioning the dump truck closer to the pile may be more efficient, however this could also increase the risk of collisions while scooping, thus a 90° turn was considered an acceptable compromise.

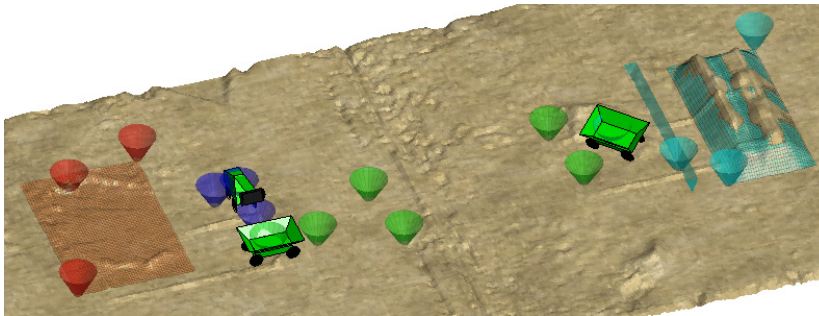
The load transfer event is illustrated in Figure 6.2(a), while the other truck waits at point J. Another benefit of this configuration is that for larger jobs, the dump truck could be loaded at the same time from the other side by another loader, a scenario presented in Section 6.2.5. Here, the dump truck is considered full after a scoop-load is transferred which surpasses the truck volume capacity, therefore the average truckload ends up being greater than 1 m^3 .

The hauling road is defined by points E and H. Point E is located a short distance from D such that line DE is perpendicular to line BCD. This is suggested so that the dump trucks approach and leave the transfer zone in the correct orientation, i.e. with the long side of the bucket facing the loader. If point E is moved by the user, however, to modify the road, the trucks can also turn on point D to orient themselves correctly. Point H is located off the mid-point of the Dump Stage side and stays stationary while point K occasionally shifts laterally with the Dumping Location as the Virtual Pile is being filled up. Points F and G are placed symmetrically alongside I and J to form a 2-way hauling road. A 2-lane road is not necessarily required with 2 trucks, as one truck may always return from a hauling cycle while the other is still being loaded. It was included,

however, to prepare for possible scenarios with 3 or more trucks when simultaneous driving in opposite directions may be necessary.



(a) Midway through job; truck being loaded, (b) Virtual Pile guides filling by dump truck; scanning line finds next location.



(c) End of job.

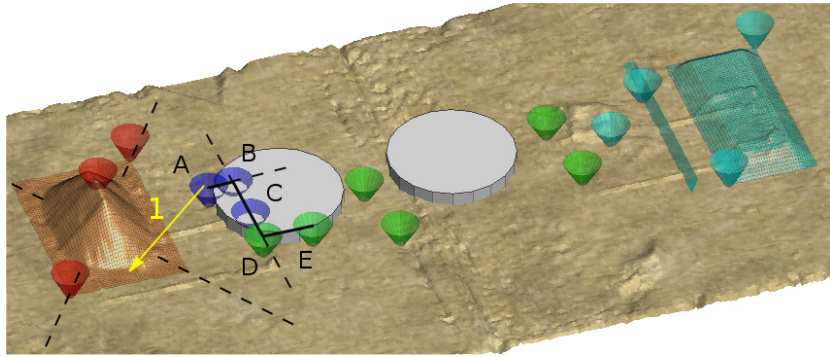
Figure 6.2. Simulation of pile transfer job with one loader and two dump trucks.

Figure 6.2(b) shows how the Virtual Pile guides filling actions by the dump truck in a similar way as for the loader in Section 3.4.5, with the line of scanning points indicating that the next dumping location has been found. Figure 6.2(c) shows the completed job at the end of the simulation.

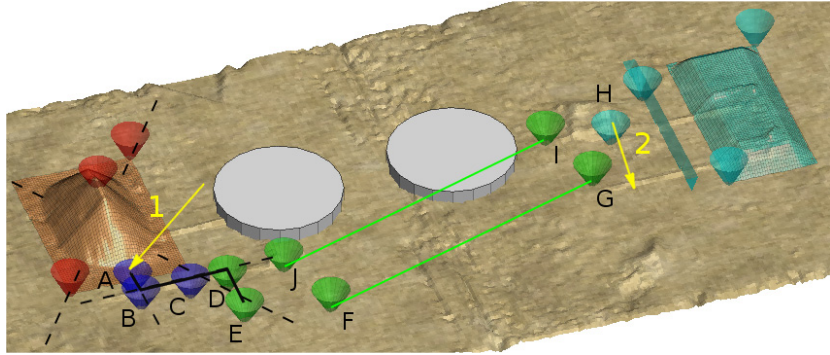
The overall plan presented here is only one possible strategy, and as before, is specific to machines capable of turning on the spot. The main goal is to show how a complete work plan for several machines can be generated automatically from a user-specified high-level plan and represented graphically. Furthermore, the multi-machine plan can be modified by the user if necessary, which is described next.

6.1.1 Modification of Plan

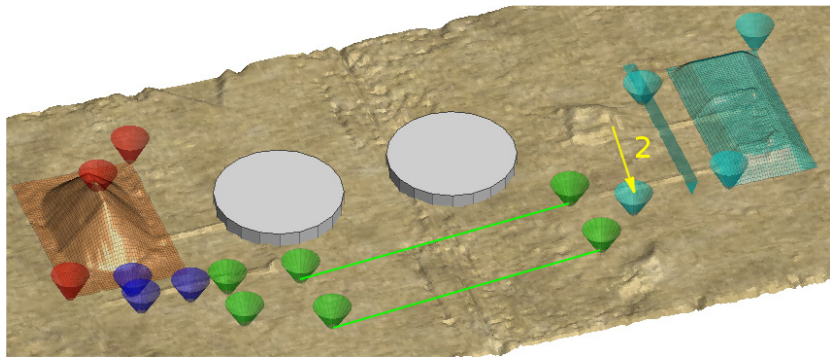
Figure 6.3 shows a scenario with the same high-level plan as previously, this time with 2 large obstacles in the way of the suggested lower-level plan. The user can modify this to avoid the obstacles in two steps. First, all the cones associated with the loading part of the job can be shifted by click-



(a) Obstacles in way of suggested plan.



(b) Move Stage cone (arrow 1) - associated cones reconfigured automatically.



(c) Move end of hauling road (arrow 2) - road cones reconfigured automatically.

Figure 6.3. Modification of plan to avoid obstacles: user moves 2 cones, others reconfigured automatically.

ing and dragging the Stage cone, indicated by arrow 1 in Figures 6.3(a) and 6.3(b), so as to approach the pile from the near side. The distances between cones A, B, C, D, E stay the same as cone A is being moved, however their orientation depends on which of the 4 Area Tool sectors cone A is in, marked by the diagonal dashed lines. These are the same as the sectors used when moving the Starting Line described in Section 3.5. As arrow 1 crosses from the right to the bottom sector, the cones are auto-

matically reconfigured such that vectors AB and DE point away from the pile, and vector BD points towards the Virtual Pile. Road cones F, G, I, J are repositioned along the new line EH.

Part of the road is then still being obstructed by one obstacle, so the user shifts the road by moving cone H along arrow 2 in Figures 6.3(b) and 6.3(c). The plan is now free of obstacles, and work can proceed. Point E could also be moved if necessary to modify the road, without affecting the other loading points A, B, C, D. A hierarchy therefore exists when moving groups of points, i.e. certain cones (in this case Stage cone A) cause all associated ones to be reconfigured, while others can be moved (road cone E) without affecting those above in the hierarchy.

6.2 Mars Hillside Excavation

In this section the planning tools and strategies developed in Sections 3.4 and 6.1 for pile transfer jobs are applied on a larger scale in a simulated robotic hillside excavation scenario, motivated by the Mars Homestead Project (MHP) construction plan outlined in Chapter 1 [1]. Although this is quite a specific target scenario, it is essentially an open-pit mining (or gravel pit) job, therefore the strategies developed here could find more general application in obtaining regolith on Mars or the Moon for In Situ Resource Utilization (ISRU), and also in automated mining on Earth.

The main scenario envisaged is one where human settlers are present on the Martian surface, but are mostly busy with other tasks, hence the desire for as much automation as possible for the excavation work. At least one human would be on hand monitoring all base activities including the construction work, and would be able to provide occasional updates while preferably avoiding constant monitoring. This is part of the motivation in developing the graphical tools: to allow a user to quickly understand the current plans of the machines after attention has been diverted for some time, and to provide updates if necessary which should keep the machines working effectively for as long as possible. Since the strategies developed do allow for long-term, fully autonomous operations in the ideal case, they could also be applied in scenarios with a long communication time delay, such as monitoring from Earth.

Section 6.2.1 begins with a discussion of considerations concerning the feasibility of robotic earthmoving work on Mars. In Section 6.2.2 high-level job plans are made beginning with the tools developed so far for the

combination of 1 loader + 2 dump trucks, which will lead to the identification of a new machinery requirement. Section 6.2.3 then shows how the commands are generated for loading, hauling and dumping, with the truck hauling aspect studied in more detail in Section 6.2.4. Section 6.2.5 concludes by presenting the addition of extra loaders to increase the excavation rate. Strategies are developed for automatically assigning workspaces among the loaders, allowing a remote supervisor to modify the workspace division, and automatically reconfiguring the workspaces in case one or more loaders encounter difficulties. Due to the larger scale of the simulations in this section, a surface model grid spacing of 0.2 m was used to reduce computing requirements.

6.2.1 Robotic Excavation on Mars

Heavy earthmoving operations on Mars present many challenges and unknowns. One limitation is the high cost of transportation from Earth, driving machinery to be as small and lightweight as possible. This conflicts, however, with the general requirement that excavation machines be robust and heavy enough to provide the reaction forces needed for interacting with the ground. Larger machines would also be desired to increase the bucket volume and work rate. Another interesting factor is the lower gravity on Mars (38% of Earth's gravity), which further increases the difficulty of applying force to the ground. One solution could be to use the regolith as ballast, filling large containers fixed to the machines to increase weight.

Power is another key consideration. Here it is assumed that machines are electric battery-powered with electric motors moving the joints. During operations, the machines would frequently need to recharge batteries from the base grid, which would ideally have a nuclear fission source for high power, but the possibility of solar power generation also exists.

Skonieczny et al. studied the problem of lightweight rover productivity for planetary excavation, comparing several parameters and making experiments with a prototype capable of carrying over 50% its weight in payload [99]. Payload ratio and driving speed were identified as the key parameters governing productivity, while others such as the number of wheels had less relevance.

Zacny and Bar-Cohen discuss several ways the extreme Martian environment presents challenges for earthmoving work [100]. These include the low temperatures and high temperature gradients which can cause

metals to become brittle and make lubrication difficult. Subsurface ice may also pose a problem as regolith-ice mixtures could be extremely hard. Friction from tool-ground interaction could cause ice to melt and refreeze later, which could make drilling difficult for example. Any earthmoving worksite would need to be carefully selected and surveyed beforehand, with detailed subsurface characterization [100].

An ideal hillside excavation site might be a debris apron consisting of loose, granular material which collapses uniformly during excavation, which is the soil behaviour exhibited in the simulation environment used here (Section 3.1). If slope fragmentation is required, conventional drill-and-blast techniques may be applicable [100].

One possible difficulty relating to slope excavation is that the lower gravity on Mars may affect the static and dynamic angles of repose of the regolith. An investigation by Kleinhans et al. of granular material behaviour on parabolic flights found that in lower gravity, the dynamic angle decreases while the static angle increases, which could strongly increase the effect of a landslide [89]. Contrary evidence was found by Atwood-Stone and McEwen however, who concluded from analysis of Mars orbital data that the angle of repose is independent of gravity [101].

With the unique environment and challenges on Mars, conventional machinery designs used on Earth may not be feasible or effective. Other possibilities include specialized dragline excavators designed for low gravity and ballistic transporters for moving regolith. The simulations here, however, assume conventional machinery: loaders modelled after the compact skid-steered Avant 320, dump trucks for hauling and also a conveyor belt spreader.

The nature of this work - repeated transfer of regolith in the abrasive dust and extreme cold of Mars - means that mechanical breakdown is likely inevitable. Fault tolerance would require having extra machines and parts. Parts commonality would also be an asset, perhaps with the loaders and dump trucks sharing the same chassis. Another balance will have to be found between the size of each machine and total number of machines. Extra machines arriving with periodic cargo flights could allow the production rate to ramp up over time.

6.2.2 Site Planning and Machinery Requirements

A concept model of the initial Mars hillside excavation site is shown in Figure 6.4, with an outpost of habitat modules and small fleet of ma-

chines. The Area Tool is used to specify the slope segment to excavate, which according to the MHP design should be 45 m wide and 30 m deep (horizontally), with an assumed slope angle of 30° . The volume of this section is equal to 11691 m^3 , however the volume that would need to be excavated could be more than twice this amount, depending on the extent and properties of the overall slope, since additional regolith would collapse in from the sides and further uphill.

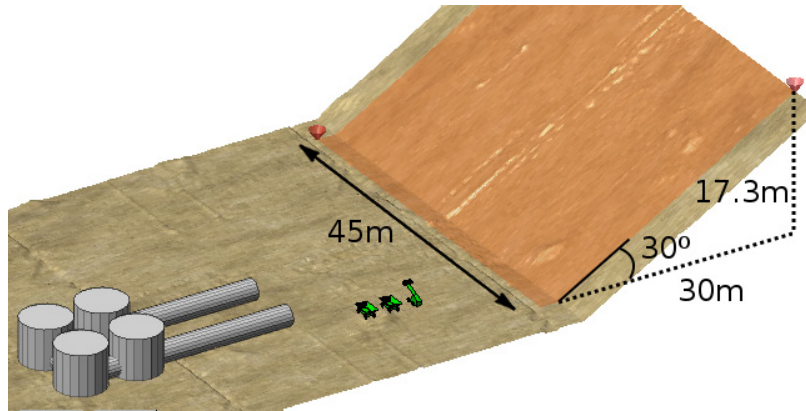


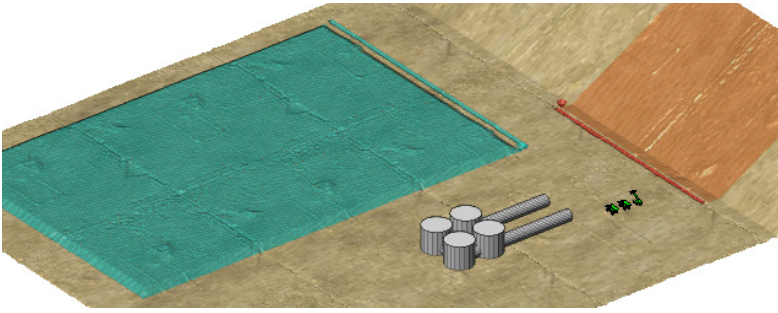
Figure 6.4. Mars hillside excavation scenario; Area Tool used to specify slope section in vicinity of initial outpost.

While some of the excavated regolith might be processed right away for ISRU and brick manufacturing, large volumes will likely need to be dumped in storage piles. Figure 6.5 shows how the Virtual Pile tool can be used to plan this aspect of the job, again with 30° slopes.

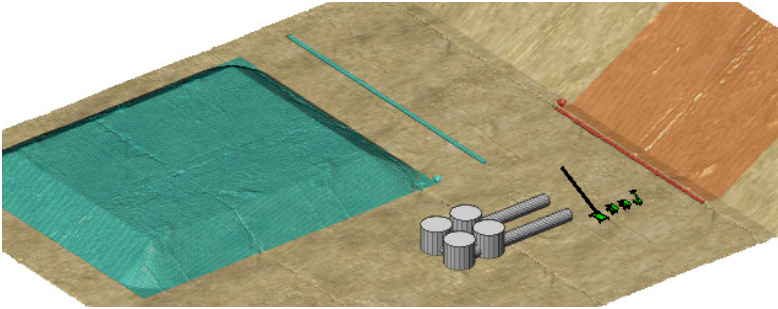
If only compact dump trucks of the type from Section 6.1 are available for depositing the regolith, the height of the storage pile may be limited to about 1 m, requiring a large area for storage as Figure 6.5(a) demonstrates. This pile has a volume only about half of the direct slope section, therefore it is evident that storing all the excavated material in a low pile would take up an exorbitant amount of space, possibly interfering with other site activities.

This leads to a machinery requirement for the ability to deposit regolith from higher above the ground, enabling more space-efficient piles as in Figure 6.5(b). Here over three times the material as the low pile is stored in a smaller area, with a pile height of 5.2 m. The machine chosen to meet this requirement was a conveyor belt spreader capable of accepting a truck-load of material and depositing it via a 12 m-long boom, which would be angled slightly less than the repose angle of the material.

Another alternative for storing the material could be to have dump trucks



(a) Low storage pile (1 m) possible with dump trucks stores 5143 m^3 in $60 \times 90 \text{ m}$ area.



(b) High storage pile (5.2 m) possible with conveyor belt spreader allows storing more volume (16207 m^3) in a smaller area ($60 \times 70 \text{ m}$).

Figure 6.5. Using Virtual Pile to plan worksite regolith storage with varying machinery capabilities.

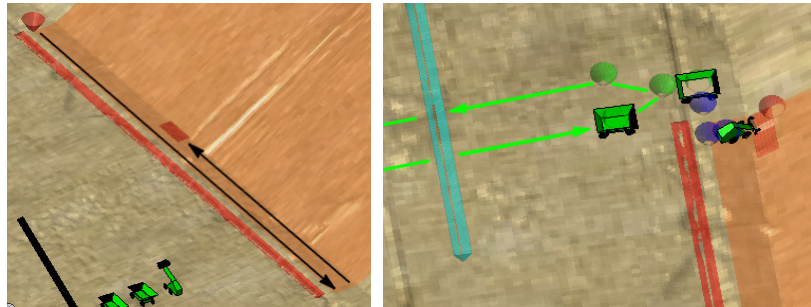
capable of driving over the regolith they deposit. They could then build up a ramp that increases in height with the volume dumped. For covering structures with protective regolith layers later in the construction plan, however, some capability would still be required of depositing regolith from high above the ground.

In Figure 6.5 the Scoop Approach and Dump Approach prisms are also shown to indicate from which sides the Area Tool and Virtual Pile are to be approached by the machines. The Dump Approach side for the higher pile has a larger offset distance from the Virtual Pile surface since it corresponds to the long reach of the conveyor belt spreader. These sides are automatically suggested to be closest to the other respective surface, but their position can also be pre-specified.

6.2.3 Automatic Plan Interpretation

After the high-level plan is specified, it is automatically interpreted to generate the points needed for loading, hauling and unloading, similar to the methods presented in Sections 3.4.3 and 6.1. The scooping location is first determined using the Scoop Area method from Section 3.4.6, illus-

trated in Figure 6.6(a) with a 2.6×1 m Scoop Area. The difference here from the example in Figure 3.21 is that the Scoop Area scans laterally from one side of the full workspace to the other before shifting one step back, illustrated by the arrows in the figure. The search stops when elevations are found a certain threshold above ground level, as in Figure 6.6(b), and the Stage point and load transfer points are positioned relative to this location. The loader works there until the area is level, after which the search for the next location would continue, with the loader and trucks then repositioning at the new location.



(a) Scoop Area scans rows front to back (b) Driving points automatically generated relative to Scoop Area and Dumping Location.

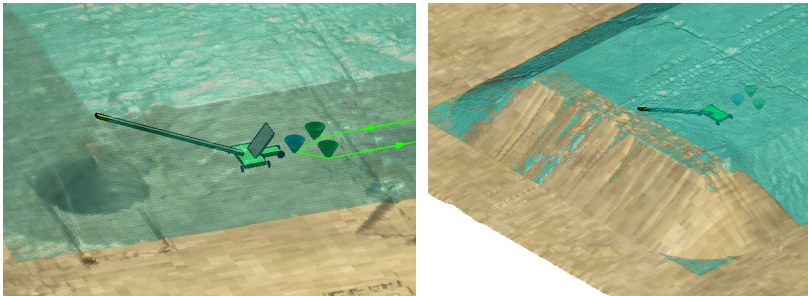
Figure 6.6. Autonomous lower-level job planning based on high-level plan.

Since the dump trucks also reposition themselves when one becomes full, one way to reduce some of the repositioning driving could be to continue working in the vicinity of a cleared Scoop Area until the currently loaded dump truck is full, so that partially-filled trucks do not need to drive to a new Scoop Area. In the strategy used here however, and also in the next chapter, the machines stop working and move on when the current Scoop Area is levelled, to give priority to excavating the slope evenly.

The dumping location for the conveyor belt spreader is determined as in Section 3.4.5, by scanning along rows of the Virtual Pile top central section with a line segment, starting at the back left corner, and comparing the desired heights with the actual corresponding ground heights. The line segment has a central width of 0.4 m, rounded down from the conveyor belt width of 0.5 m, with one grid point added at each end, making it 0.8 m wide due to the 0.2 m grid spacing. The step size along rows is again one grid point, and the step to the next row is 2 points or 0.4 m. Thus the same principle is followed as previously though with different values.

Figure 6.7(a) shows a load transfer operation at the spreader, while a

pile is building up at the first location. Here the hauling road is slightly different than in Figure 6.1(b), and is arranged relative to the moving scooping and dumping locations. This assumes that the whole workspace in between is traversable, however if this were not the case then obstacle avoidance or a multi-segment road with fixed sections would be required.



(a) Dump truck unloading at spreader; filling (b) Long duration simulation; 2315 truck loads deposited totalling 2487 m^3 .

Figure 6.7. Large-scale Virtual Pile and use of conveyor belt spreader.

Figure 6.7(b) shows a long-duration filling operation, with 2315 truck loads deposited totalling 2487 m^3 . This Virtual Pile is narrower than that in Figure 6.5(b), with a base width of 46 m, though it has the same height of 5.2 m. It can be seen that the shape of the dump pile follows that of the Virtual Pile quite closely. Long-duration results of the excavation portion of the job are presented in Section 6.2.5, but first certain aspects of the truck hauling will be discussed.

6.2.4 Hauling Distance

Here a brief analysis is made which highlights one benefit of the simulator for site planning purposes, and is related to some of the construction simulation research presented in Section 2.4. When making the initial high-level plan, a basic trade-off exists in the placement of the Virtual Pile for regolith storage. It should be sufficiently far away so as not to obstruct settlement construction and base activities, yet also be close enough to the building site and processing equipment so that the hauling distance does not unnecessarily decrease the work rate or require extra trucks. In the case of two trucks, Figure 6.8 shows a situation where the hauling distance is too great to keep the loader operating continuously, since the empty truck is still on its way back after the other one has already been filled.

The threshold distance at which a third truck would be required can be

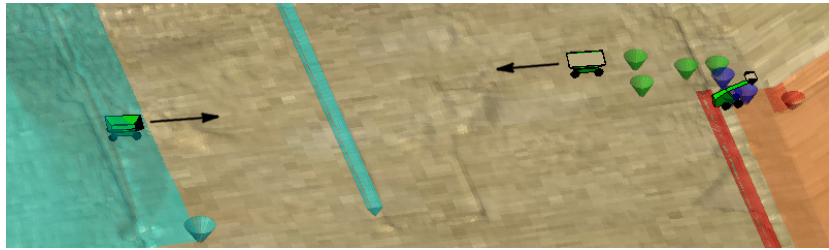


Figure 6.8. Two dump trucks insufficient for long hauling distance - loader stands idle while waiting for empty truck.

estimated by comparing the time needed to load a truck with the time needed for a complete hauling cycle. Loading time can be estimated, however some uncertainty will always exist, one reason being variability in the driving distance when acquiring a given scoop-load, and perhaps also in the scoop filling which affects the number of loads needed to fill a truck. One benefit of simulation is therefore being able to estimate average times for various phases of the work cycle, which can then assist with planning.

Assuming two dump trucks are available, one useful value to find out is the maximum hauling distance which still keeps the loader operating continuously. Hauling distance here means the distance between the end cones along the centre of the road (cones E and H in Figure 6.1(b)).

Besides hauling distance and the impact it has on site planning, another important parameter is the excavation rate, with a higher rate desired for speeding up construction. One way to achieve this is simply to have all the machines drive faster. With a given driving speed, limited perhaps by energy consumption, another way to increase the rate would be to increase the volume capacity of the loader bucket, which also has higher energy requirements. This then affects the hauling distance - assuming the dump truck capacity remains the same, a larger loader bucket will fill the dump trucks more quickly and shorten the hauling distance which allows continuous loader operation.

Two simulations were carried out which compare the effect of loader bucket capacity on the excavation rate and hauling distance, with the results shown in Table 6.1. Both simulations involved 1 loader and 2 dump trucks, with the scoop capacity first set at 0.15 m^3 and then doubled to 0.3 m^3 . Truck capacity was kept constant at 1.0 m^3 . As the scoop does not generally fill completely when loading, the average scoop-loads are shown, together with the average number of loads needed to fill a truck.

“Maximum Haul Distance” refers to the approximate maximum hauling distance which still allowed for continuous loading operations, and was

Table 6.1. Comparison of simulation results (1 loader, 2 dump trucks) with different loader scoop capacity; truck capacity 1 m³, all driving speeds 0.5 m/s.

Sim.	Scoop Capacity (m ³)	Average Scoop Load (m ³)	Average Loads per Truck	Work Rate (m ³ /h)	Maximum Haul Distance (m)
1	0.15	0.100	10.5	11.11	~55
2	0.30	0.238	4.69	24.78	~20

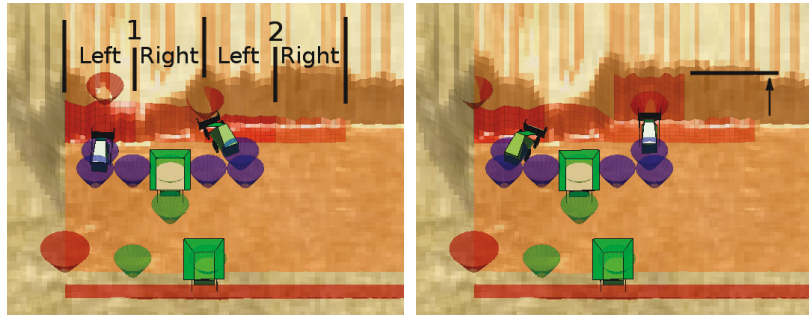
found by gradually moving the Virtual Pile further away during a simulation until the loader had to occasionally wait for an empty truck to return. In these simulations the hauling road terminated at the Dump Approach prism as in Figure 6.1(b), thus an additional hauling distance existed from the end of the road to the spreader. The Virtual Pile used had a base length of approximately 30 m.

The results in Table 6.1 show that in Simulation 2 the average scoop-load was more than twice that in Simulation 1, therefore as expected, the excavation rate was also over twice as high while the maximum hauling distance was less than half.

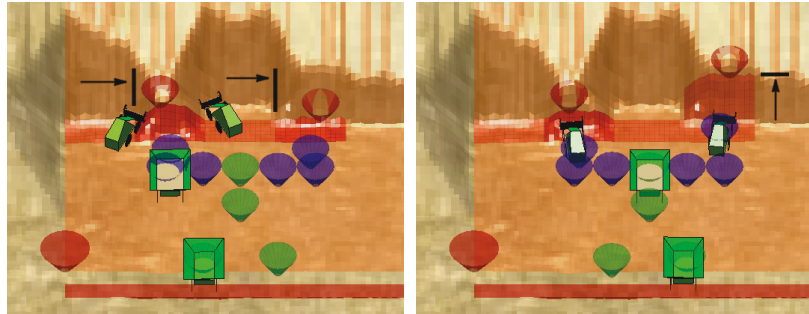
6.2.5 Multi-Loader Excavation

To increase the excavation rate and/or redundancy by adding more loaders, a strategy for automatically assigning workspaces was developed based on the Scoop Area from the single loader case. Two main requirements are assumed: the workspaces of the loaders should be separated laterally along the hillside face to avoid collisions, and the hillside should be excavated evenly along the face. A further assumption is that each dump truck filling location should be shared by 2 loaders by filling from both sides, to decrease the total number of dump trucks required. The number of trucks required would also depend on other factors such as volume capacity and hauling distance, as Section 6.2.4 explained. Figure 6.9 shows how this workspace strategy was implemented for a 2-loader case. In this section, although the hillside has an initial slope angle of 30°, a repose angle of 45° was used during the simulations to allow for more rapid progress into the hillside. One possible explanation for this, however, could be changing material properties over time.

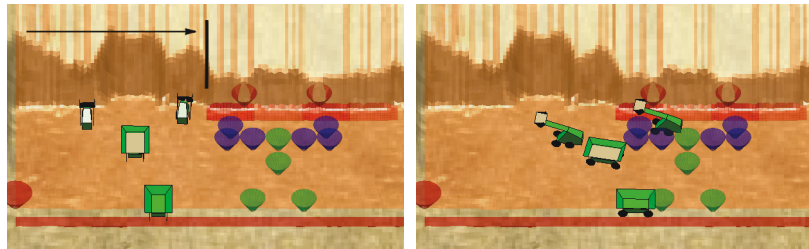
The method that was developed to automatically allocate workspaces was to have a Scoop Area which is divided into two times the number of



(a) Scoop Area divided into 4 zones; both loaders working on respective left zones. (b) Loader 2 finished, zone moves forward temporarily.



(c) Loader 1 finished, zones switch to right with driving points shifted. (d) Loader 2 finished again, zone moves forward temporarily.



(e) Loader 1 finished; Scoop Area finished, shifts right with driving points. (f) Machines drive to new points.

Figure 6.9. Autonomous Scoop Zone allocation pattern within Scoop Area for 2 loaders, and shift of Scoop Area when all zones completed.

loaders. Each of these divisions is called a Scoop Zone (SZ) and is the size of the workspace needed by one loader at any given time, for example 3 m wide by 2 m long. For 2 loaders, the Scoop Area therefore consists of 4 SZs, with each loader having a left and right zone.

When the Scoop Area finds a location to excavate with the same scanning pattern as the single-loader case, the loaders first work at their respective left SZs (see Figure 6.9(a)). The spacing between the loaders allows the dump truck to approach fairly close to the hill face, reducing the time needed to load the truck. In Figure 6.9(b), the 2nd loader has

finished its zone, so it is extended temporarily forward by one length increment. When the 1st loader has also finished, both loaders shift to their right zones, with the truck loading points also updated (Figure 6.9(c)). Soon after, the 2nd loader again finishes its zone, which is temporarily extended forward as before (Figure 6.9(d)). When the 1st loader is finished, the entire Scoop Area is also finished, so it scans for the next location (Figure 6.9(e)), and the machines drive to their new points (Figure 6.9(f)). In this way, the loaders stay separated and excavate relatively evenly along the hillside face.

Having no separation between current SZs could offer the benefit of allowing more loaders to work simultaneously in the same total space, however this would also make it possible for the intended scooping trajectories of the loaders to intersect. In this case a priority system would need to be in place requiring one to wait while the other proceeds. The action of one loader might then change the topography in the location where another loader was also intending to scoop, possibly causing it to reassess the workspace and decide on a different scooping location. Furthermore, no separation between the zones may require that the dump trucks wait further back since there would be less space available between the loaders. For simplicity, it was thus decided to separate the workspaces so that each loader can work within a certain zone until it is complete without the possibility of conflicts with other loaders.

The reason for dividing the Scoop Area into zones of equal width is so that ideally along a uniform hill face, each zone would take approximately the same time to level, and the loaders would not spend much time working on temporary extensions (as in Figures 6.9(b) and 6.9(d)). The volume of material to excavate in a zone depends on more than the zone width however, since extra material collapses in from the surrounding slopes. Nevertheless, after conducting numerous simulations, this method was observed to result in relatively even excavation along the slope without narrow sections being dug too far inwards.

For more than 2 loaders, different approaches are possible because since each loader team uses its own dump trucks, it is not linked to the other teams and could therefore work anywhere else along the hillside. The initial strategy decided upon, however, was to combine all loaders into the same greater team, which would ensure the workspace separation among the loaders by extending the pattern developed for 2 loaders. Figure 6.10 shows the workspace planning for 3-7 loaders. For an even number, two

loaders always share a truck for loading, while for odd numbers the right-most loader has its own truck.

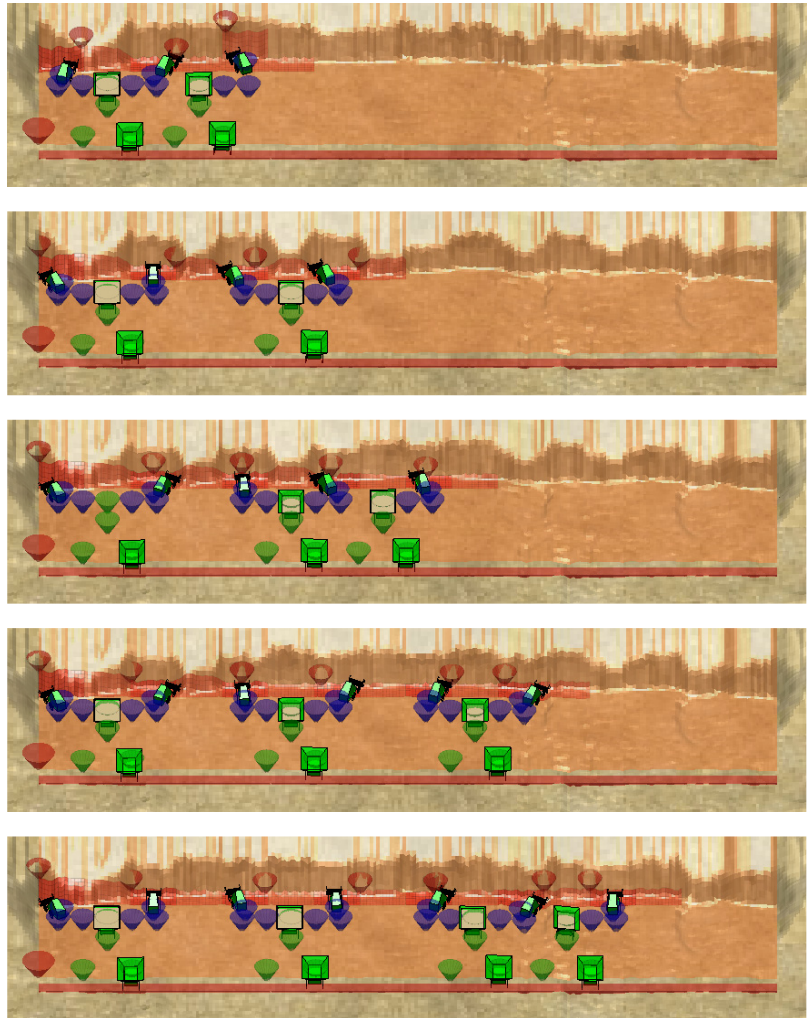


Figure 6.10. Scoop Zone arrangement within Scoop Area for 3-7 loaders.

Since the focus here is on workspace allocation for the excavation portion of the job, one aspect which was not developed for the case of 3 or more loaders was the hauling routes needed by dump trucks from more than one team. As the truck paths would need to merge at some point to use the same hauling road and/or to unload at the spreader, proximity detection and right-of-way protocols would need to be developed to ensure smooth traffic flow. Another way to coordinate the dump trucks could be to find conflict-free trajectory envelopes using the methods developed by Pecora et al. mentioned in Section 2.6 [76]. In these simulations, when a full truck reaches the equivalent of point F in Figure 6.1(b), the load is

deleted and the truck continues to point J.

The maximum number of loaders that can work along a given hillside section depends on the SZ dimensions. A narrower SZ allows more loaders to fit along the section, potentially increasing the excavation rate. The minimum zone width would ultimately be limited by the width of the machines and perhaps space needed for turning. If too narrow, the possible approach vectors with which to scoop into the hillside would also be limited. In these simulations, with the 45 m-wide hillside section specified in the MHP construction plan, the maximum number of loaders was usually set at 8. This was based on SZs 2.4-2.8 m wide, which allowed sufficient space for the machines.

Figure 6.11 shows an example of a long duration simulation with 8 loaders which began with the fresh hillside. Here SZ dimensions 2.4 x 0.8 m are used, which are shown later in Section 7.7 to result in the highest work rate compared with a few other dimensions for a similar simulation. After 19.21 hours of simulated time, 2330 m³ of material was excavated, with all workspace updates having been made using the automated allocation pattern outlined in Figure 6.9. As the figure shows, the hillside has been excavated quite evenly as the bottom contour is relatively straight.

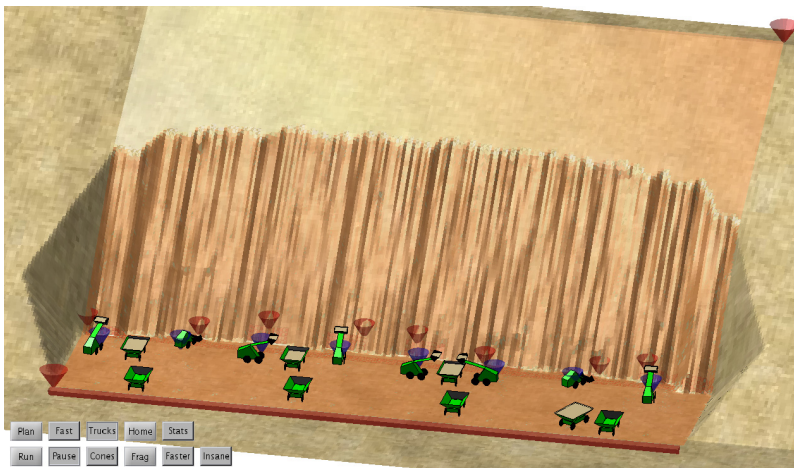


Figure 6.11. Long-term simulation with 8 loaders.

For this particular simulation, the flat ground at the base of the original slope was a perfectly level plane, which was to help minimize the unwanted excavation of material below ground level that can occur in this simulator. As the pile loading job in Section 3.4.4 showed, this can sometimes happen if the bucket blade dips below ground level during scooping actions due to driving on uneven ground. For large jobs such as this, the

gradual excavation of a downward ramp can sometimes result, which is avoided if starting with a flat surface.

The multi-loader hillside simulation was repeated several times in order to compare excavation rates while varying certain parameters, such as the number of loaders, loader bucket volumes and driving speed. The results are summarized in Table 6.2, with the last column showing an estimated job completion time for each simulation. This is based on a total excavated volume of 24162 m^3 , which is the direct $45 \times 30 \text{ m}$ section volume of 11691 m^3 plus an estimated additional volume from surrounding slope collapse. A relatively low effective work rate of 3 hr/day is assumed, which leaves time for interruptions due to battery loading, breakdowns, maintenance and possibly slope fragmentation. In the best case, spare, fully-charged loaders would be available to take over when other loaders run low on power or need to be fixed.

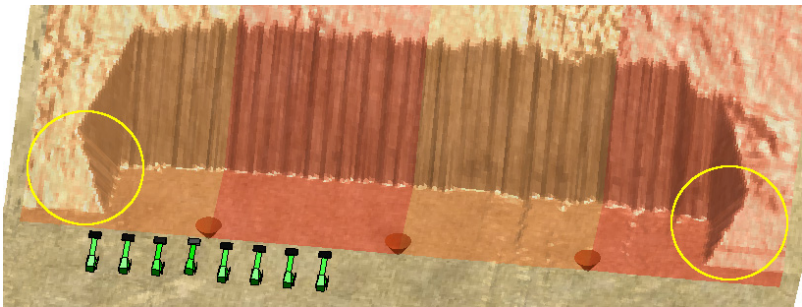
Table 6.2. Hillside excavation work rates for varying loader and fleet parameters.

Sim.	Loaders	Driving Speed (m/s)	Scoop Cap. (m^3)	Avg. Load (m^3)	Excav. Rate (m^3/h)	Job Completion Time (days), @ 3 hr effective work/day
1	2	0.5	0.15	0.105	31.67	254
2	4	0.5	0.15	0.104	62.34	129
3	8	0.5	0.15	0.105	121.29	66
4	8	0.25	0.15	0.123	94.41	85
5	8	0.5	0.20	0.148	172.92	47

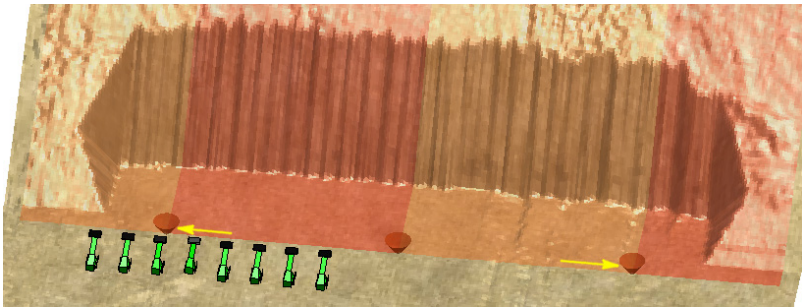
As expected for Simulations 2 and 3, increasing the number of loaders results in an approximately linear increase in the excavation rate compared with Simulation 1. For the 4th case with 8 loaders and the driving speed reduced by half, the rate is not reduced by half since the average bucket load increases significantly. This is due to the default scoop filling behaviour in the simulator described in Section 3.1, which is affected by the driving speed. With the same driving speed and the scoop capacity increased to 0.20 m^3 , the rate predictably increases.

While observing these simulations, it was noticed that it may be useful to have the capability of dividing the loaders into separate teams. One example would be if, after some progress is made along the 45 m -wide section as in Figure 6.12(a), it was desired to create a more gradual excavation contour by excavating the convex side corners bordering the cut-out (indicated by the circles). This could be accomplished by widening the total workspace, however with the previous strategy of all the loaders on

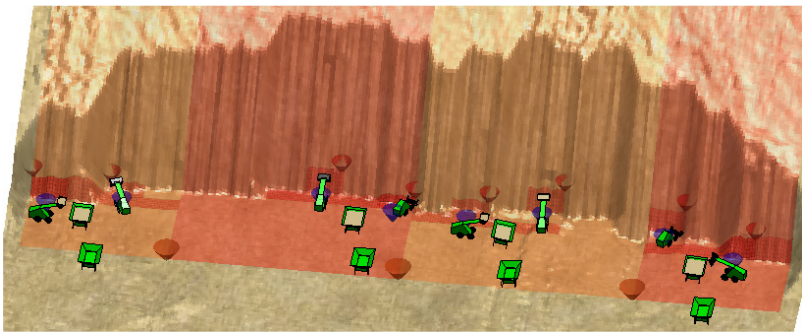
the same team, only one corner at a time could be reduced while the remaining loaders work nearby.



(a) Excavation area widened to include convex section edges, area divided evenly into Sub-Areas for each loader team.



(b) User clicks and drags divider cones to make Sub-Areas at sides narrower (indicated by arrows), focusing excavation effort on section edges.



(c) Convex edges reduced over time, creating more gradual cutout.

Figure 6.12. Division of workspace into Sub-Areas to separate teams.

Another strategy was developed whereby the initial Area Tool is divided into Sub-Areas for each loader team, with the division width by default proportional to the number of loaders per team. The lighter and darker sections in Figure 6.12(a) represent the Sub-Areas, with the cones along the front marking their boundaries. The user can adjust the Sub-Area widths by clicking and dragging the dividing cones (indicated by the arrows in Figure 6.12(b)), for example to focus excavation work by the two

outside teams on the convex sections. The Sub-Area has a minimum width equal to the loader team's Scoop Area.

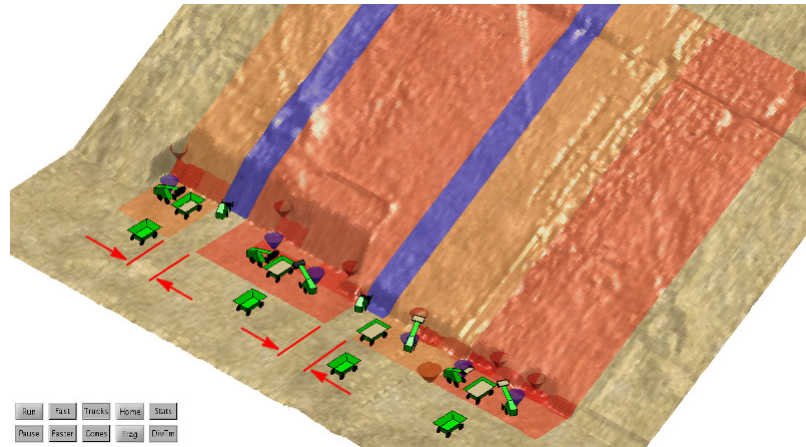


Figure 6.13. Sub-Area sides adjusted automatically if neighbouring loader gets stuck, allowing other loaders to continue working while problem being solved.

The Sub-Area strategy would also be useful if a loader got stuck or broke down, blocking a section of the workspace. For this situation, an algorithm was developed for automatically reconfiguring the Sub-Areas such that the other loaders can continue working along as much of the hillside as possible while the problem is being solved.

Figure 6.13 shows a case in which two loaders have become stuck, with narrow zones marked by the blue surfaces indicating the problem and demarcating sections of the workspace where other loaders cannot work. The Sub-Areas neighbouring these zones have been automatically adjusted, with some made narrower and some extended (indicated by the arrows), so that the other loaders can continue working in the remaining area.

6.2.6 Discussion

This section presented an autonomous multi-robot hillside excavation and material storage job, motivated by the MHP construction scenario. Since the job is of a much larger scale than the pile transfer jobs in previous sections, the Area Tool and Virtual Pile which are used to specify the high-level plan would likely be placed according to a detailed site plan, and not by clicking and dragging by one user as before. They still however serve to represent the plans visually and, together with the tools for lower-level plans, assist in monitoring overall progress once work begins. The Virtual Pile and dumping location search strategy were shown in Section 6.2.3

to also function on a much larger scale than previously in Section 3.4.5, using the same general search parameters.

The multi-loader workspace division strategies presented in Section 6.2.5 represent a top-down approach to multi-robot coordination, whereby each machine always has its own separate zone in order to avoid interference with the other machines. Automatic updates were made with a strategy which attempts to excavate the slope evenly, by assigning temporary new Scoop Zones in the forward direction until all loaders have cleared their zone assigned laterally in the rectangular Scoop Area. Here arbitrary Scoop Zone dimensions were used in the simulations, however in the next chapter these are varied to investigate their effect on the work rate.

As an area for future work, the workspace division strategy could be extended from the purely rectangular implementation to one which distributes the Scoop Zones along a general curved contour. Other planning approaches could also be applied such as the conflict-free trajectory envelope method of Pecora et al. [76]. This could offer an advantage by allowing the machines to work closer together without explicit zones and reduce overall driving distance, for example. Comparisons could also be made with bottom-up planning approaches such as the Artificial Neural Tissue (ANT) controller mentioned in Section 2.6 [78], or the market-based multi-robot coordination method of Dias et al. [102]. These may result in the machines organizing themselves into configurations which are non-intuitive yet beneficial for certain metrics such as maximizing the excavation rate or minimizing driving.

7. Comparison of Scooping Approach Method and Scoop Area Dimensions

In this chapter, two high-level planning considerations for automated loading jobs by skid-steered wheel loaders are investigated using simulations. The first is the method for selecting the next scooping approach within a given workspace, with comparisons made between the High Point (HP) strategy presented in Section 3.4.4 and a Zero Contour (ZC) method which was implemented. The latter is based on the work of Singh and Cannon, and attempts to increase bucket-filling efficiency by evaluating all possible scooping approaches along the bottom edge of the slope before selecting one based on a set of heuristics [57]. The former, as described earlier, represents a simple alternative resulting in more limited driving paths, although with the possible disadvantage of selecting non-perpendicular approaches into the pile.

The HP and ZC methods are compared mainly in terms of excavation rate for various jobs to determine if either one offers an advantage. All of the jobs are based on the truck loading scenario presented in Section 6.1, which means all of the loader's scooping drives start and end at a point beside the truck close to the slope being loaded. The main hypothesis is that if the same scoop filling effectiveness can be achieved, the HP method results in less driving by the loader and hence a higher work rate due to its more limited driving paths.

The second high-level planning consideration which is investigated is the dimensions (width and length) of the Scoop Area (SA), or smaller current workspace used by the loader within a larger workspace, both of these being rectangular. This division strategy was illustrated earlier in Figures 3.21 and 6.6(a), and was necessary for the HP method to avoid the problem of persistent scooping towards a point far in the pile (as in Figure 3.20). A ZC method would not have this problem since the search criteria result in scooping locations which are distributed along the bot-

tom contour.

If a loader using a ZC method was free to scoop anywhere along the workspace, which would be the case if it was transporting the load to a stationary point, the coarse loading strategy presented by Magnusson et al. could be used to maintain an even profile across the workspace [71]. If a truck were to be loaded which is independently located, then the coarse loading strategy presented by Singh and Cannon could be used, which selects a region bounded by two contour points near the truck [57].

For a large slope excavation job with truck loading, where the truck positions can be specified in order to excavate the slope evenly, assuming this is a goal, then a ZC method would also need to subdivide the workspace using the SA strategy. This is because the loader would at some point need to stop working on a given width-wise section of the slope to avoid digging too far inwards, and then move laterally to the next location.

For the larger simulations in this chapter, the length and width of the SA is varied for different jobs, using both the HP and ZC method, to determine how the dimensions affect the work rate, and if optimal values can be found. This may be interdependent on the method being used (HP or ZC), with each one achieving its maximum rate with different SA dimensions. Certain parameters of the job, such as slope height, angle and surrounding slope collapse may affect the optimal SA dimensions, which is also investigated.

A tradeoff is expected whereby smaller SA dimensions have the benefit of reduced driving within the workspace, increasing the work rate, however since they would be cleared sooner, more repositioning and therefore driving would also be required, lowering the work rate. Larger SAs, conversely, should have the benefit of less frequent repositioning, however more driving within the current workspace. In order to account for the disadvantage of repositioning all the machines when shifting the SA, a second type of rate is also used in Section 7.6 which is the volume excavated divided by the combined machine driving time. For all the simulations in this chapter which involve SA subdivision, the truck loading aspect of the job is kept constant, with the same volume capacities and strategy of repositioning all machines when the current SA is levelled.

To compare the performance of the HP and ZC methods, a small-scale pile-loading job similar to that in Section 3.4.4 is first used. This is contained in one workspace and therefore does not require subdivision with the SA. A small slope excavation job is then simulated which is subdi-

vided width-wise into two parts. These smaller jobs are used to get an initial idea of the different performance of the HP and ZC method. They are then followed by larger jobs requiring multiple subdivisions width- and length-wise, which are used to test the effect of varying SA dimensions.

Before the first comparisons are made, the next section presents the ZC method that was implemented. The simulations in this chapter again use a surface grid spacing of 0.1 m, except for the multi-loader case in Section 7.7 where 0.2 m is used.

7.1 Zero Contour Method

The Zero Contour (ZC) method used here is based on the work of Singh and Cannon discussed in Section 2.5, and uses the convexity of the contour as the main criterion for selecting the best scooping location [57]. This heuristic is based on the idea that with higher convexity, loading effectiveness is increased by applying more force along the cutting blade during the scooping action rather than along the walls of the bucket, which would be the case when penetrating into the pile along a straight or concave contour. Driving distance to the load transfer point is also considered here when choosing a scooping location.

One difference with the method developed is that convexity is determined in 2 dimensions by adding up positive and negative linear ZC components, instead of in 3 dimensions as a volume ratio, and another is that side loading is not considered. This version is not intended to be an improvement over previous ones, hence the simplifications, but is meant to consider a large number of possible actions and generate similar driving patterns, against which the HP method can be compared.

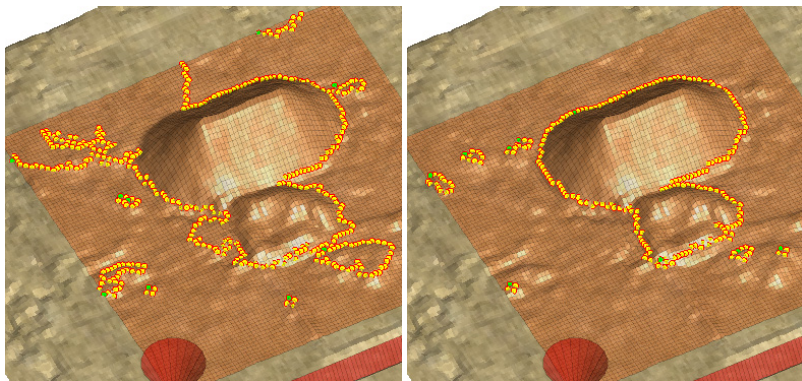


Figure 7.1. Example of Zero Contour construction around pile on uneven ground, with 0.10 m (*left*) and 0.15 m (*right*) threshold above ground level.

The first step in the ZC method is to generate the contour(s). Given a ground model, workspace and approach side, an algorithm scans the workspace domain from left to right, along each line from front to back, searching for a contour point. This is where the ground height increases to a value equal to or greater than a certain threshold above ground level. When found, the contour is traced either to its end, if it terminates at the domain border, or to the point farthest from the front, if it is closed within the domain. The contour is then built starting from this end point and the points saved, after which the original scanning skips to the next line. If a point is reached during scanning which already belongs to a contour found previously, the search also skips to the next line, therefore contours which are occluded by others in front may initially be disregarded.

A height threshold of 0.15 m above ground level is used for constructing the ZC. In the example in Figure 7.1, this results in a classification closely matching what one would consider to be “the pile” (at right), with a few smaller extra groupings. Lowering the threshold to 0.10 m, on the other hand, results in too much of the “ground area” being surrounded by the contours (at left). A minimum contour length can be specified in case single point spikes or small clusters should be disregarded due to sensor noise, for example, but here all points are included since the HP method also considers all points at this threshold above ground level.

Algorithm 3 details how the ZC method selects a scooping location given one or more contours. As with the method of Singh and Cannon, the scooping strategy is for the bucket to enter the pile such that the two front corners reach the ZC at the same time, resulting in a perpendicular approach towards the contour [57]. Scooping locations therefore consist of pairs of contour points separated by a distance approximately equal to the scoop width, which also define an approach heading since the line segment connecting the points corresponds to the bucket blade orientation. Searching for eligible locations is thus analogous to tracing a rectangle with the dimensions of the scoop cutting plane along the contour such that the two front corners touch the contour, illustrated in Figure 7.2. To find all possible locations for each contour point, each following point is checked until the end, which makes it possible to detect locations containing long protrusions.

During the search, point pairs which define scooping approaches with a backwards-facing component are disregarded, so as to minimize the maneuvering required of the loader. Pairs are also disregarded if they contain

Algorithm 3 Determining next scooping location from Zero Contour(s) in workspace

```

for contour = 1:end // iterate through contours
| locationFound = 0;
| eligibleLocation = 0;
| for i = contour(1):contour(end) // iterate through every contour point
| | for j = contour(i+1):contour(end) // check every point after i
| | | lineDist = distance(point(i) to point(j));
| | | if lineDist ~ scoopWidth // possible scooping location found
| | | | if scooping heading has backwards-facing component
| | | | | continue // disregard, to minimize excess maneuvering
| | | | end
| | | | for k = i+1:j-1 // iterate through contour points within location
| | | | | keep track of points outside scoop path due to possible bulge
| | | | | // determine convexity:
| | | | | sum up +/- segments intersecting and perpendicular to i-j segment
| | | | | end
| | | | if contour points outside scoop path > 2 // if bulges too much
| | | | | continue // disregard, not contained by scoop path
| | | | end
| | | | eligibleLocation = 1; // eligible location found
| | | | if convexity higher than or within 10% of best so far
| | | | | // This narrows down eligible new locations, so only check
| | | | | // traversability if location possibly better
| | | | | if location traversable
| | | | | | if convexity over 10% higher than previous saved maximum
| | | | | | | choose current scooping location over previous best
| | | | | | | locationFound = 1;
| | | | | | | save new maximum convexity;
| | | | | | else // convexity within 10% of previous best
| | | | | | | distanceToLoad = distance from stage point to load approach;
| | | | | | | if distanceToLoad shorter than for previous best
| | | | | | | | choose current scooping location over previous best
| | | | | | | | locationFound = 1;
| | | | | | | end
| | | | | | end
| | | | | end
| | | | end
| | | end
| | end
| end
| end
| if eligibleLocation == 0 // i.e. from small contour narrower than scoop
| | make location at centroid of points with approach from front
| | if location traversable
| | | choose this scooping location
| | | locationFound = 1;
| | end
| end
| if locationFound == 1 // location found in current contour
| | compare location with best from other contours, pick closest to front
| end
end

```

a contour segment which bulges too much (> 2 points) beyond the scoop width, since all the material would not be contained in the scoop path.

For the remaining eligible scooping locations, convexity is assessed by summing up the individual linear components of the contour segments which are perpendicular to the cutting blade line. Convex components protruding into the projected bucket are assigned positive values, while concave points retreating away from the cutting line are negative (illus-

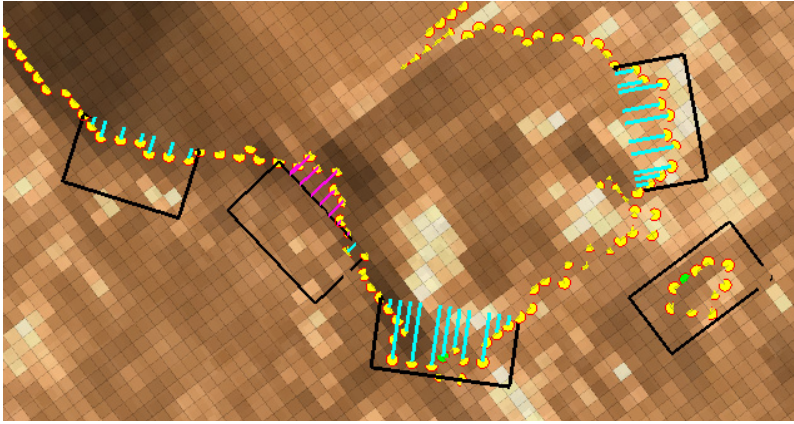


Figure 7.2. Evaluation of convexity for possible scooping locations along Zero Contour; small separate contour at right assigned one possible location.

trated by the light blue and magenta lines in Figure 7.2).

After the convexity of an eligible location is determined, it is only considered further if its value is greater than or within 10% of the best value saved so far in the contour. If this is the case, its traversability is then assessed by scanning ground heights in the region in front, to ensure that no slopes or piles would hinder the loader's approach (illustrated in Figure 7.3). If more than 4 points in the scanning region are greater than 0.15 m above ground level, the location is not considered further. A traversable location becomes the new candidate dig point if its convexity is over 10% greater than the previous best value saved (to increase scooping effectiveness). If not, it can still be selected if it is closer to the load transfer point than the current candidate (to reduce driving).

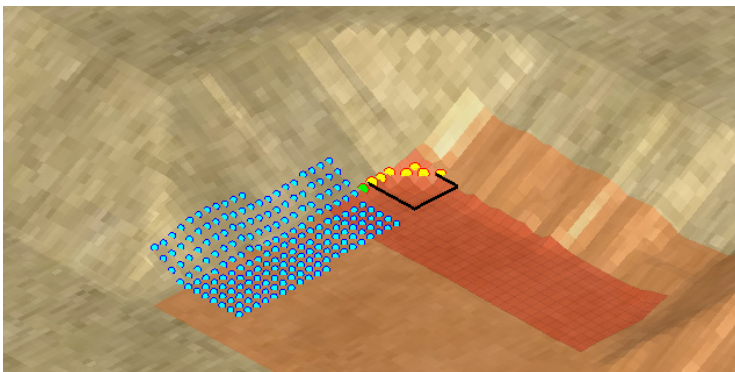


Figure 7.3. Checking approach traversability for Zero Contour scooping location.

For a scooping location, the Scooping Destination and Stage points needed by the loader are placed along the perpendicular centre line between the point pair. The Scooping Destination is located either between the point

pair, if the whole contour segment is convex, or at a distance towards the pile equal to the longest concave component. The Stage point is placed at a distance away from the point pair line equal to the longest convex component plus an offset of 1.8 m, which is the approximate distance from the vehicle reference point at the steering centre to the bucket blade in scooping configuration.

If no eligible location is found due to a contour being too small, with no point pair as wide as the scoop, a Scooping Destination is assigned at the centroid of the contour points with an approach from the front. An example of this case is shown at the right of Figure 7.2, and also in Figure 7.3.

In case a workspace contains more than one contour, the best scooping location from each one is assessed, and that closest to the front is selected in order to increase subsequent access to the rest of the workspace. If no acceptable scooping locations are found for any contour within the workspace, the HP method is used as a backup. As an example, this occurred 4.9% of the time for Job 3a which is presented later in Section 7.5.

7.2 Simulation of Bucket Filling Effectiveness

One feature that was added to the simulator for these tests was the ability to adjust the scoop filling effectiveness. The criterion for scoop filling described in Section 3.1, which was a simplified strategy and not a fully-developed bucket control sequence, resulted in an average fill ratio around 0.73. Automated scoop filling is a capability the loader is assumed to have, however it is unknown what its actual effectiveness would be. For this reason it was made possible to specify different scoop filling ratios.

Another limitation of the scoop filling behaviour used so far is that it results in a fully deterministic simulation; i.e. given the same ground model, plan and loader starting position, a simulation is repeated exactly the same each time. Repeating such a job in reality would certainly result in variations with each trial, particularly due to scoop-ground interaction. To simulate this, randomness was added to the scoop filling behaviour.

The specified scoop filling effectiveness and randomness were implemented as follows. As with the original behaviour, the scoop penetrates into the ground surface until the timestep at which the new volume increment would cause its capacity to be surpassed. With the new behaviour, the final load is then made to equal a minimum specified ratio of the vol-

ume capacity, with the rest of the capacity filled randomly and the remaining material distributed back on the ground below. This random amount is a uniform random distribution over the remaining bucket capacity, since it is assumed that any fill ratio from the minimum to 1.0 is equally likely to result from an action. Thus with a minimum load ratio of 0.8 for example, an average of 0.9 results over many actions.

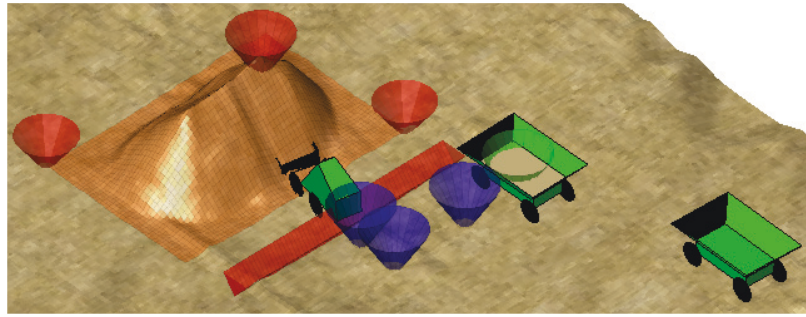
Jobs 1 and 2 in the next sections are simulated with minimum scoop fill ratios of 0.6, 0.7, 0.8 and 0.9, resulting in average fill ratios of 0.8, 0.85, 0.9 and 0.95. For the larger jobs when the SA dimensions are being varied, a minimum ratio of 0.8 is assumed in order to limit the number of variables. For most of the jobs, each case is repeated 10 times in order to obtain a distribution of results due to the randomness introduced. The first simulations of Job 3 in Section 7.5, however, are repeated 5 times due to the large number of cases and the time required to run the simulations.

7.3 Comparison of Scooping Approach Method for Single Pile

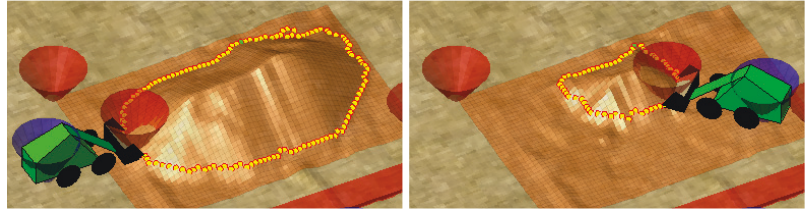
Job 1a, the first case used to compare the HP and ZC method, is similar to the pile loading portion of the jobs simulated earlier in Sections 3.4 and 6.1. This type of job was chosen first since it is a typical job performed with a wheel loader, and would give a first idea regarding the performance of each scooping approach method. The pile has a 2 m-wide central section and 45° slopes, shown in Figure 7.4(a) with the associated job planning tools and HP method being used. The Area Tool surface demarcating the pile has dimensions 6.2 m wide by 4.2 m long, with a total volume of 11.7 m³ initially detected above ground level. These pile dimensions were chosen arbitrarily, with the intention being to have a large enough pile which would not require subdivision with Scoop Areas (see Section 3.4.6).

Two dump trucks with capacities of 1 m³ are available for load transfer, as in Section 6.1, allowing the loader to operate continuously. The same load transfer strategy as before is also used, with cones B, C and D from Figure 6.1(b) rendered in Figure 7.4(a), and with the extra truck waiting at point J. Here the hauling and unloading portions of the work cycle are not included, and when a full truck reaches point F, its load is emptied and it drives to point J. The loader has the default volume capacity of 0.15 m³, and all machines drive at 0.5 m/s and turn at 30°/s.

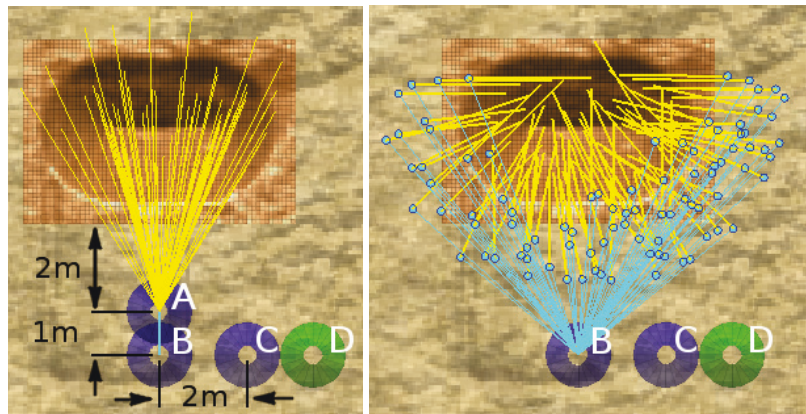
A second version, Job 1b, consists of a pile with the same footprint but with 30° slopes, to observe if a different repose angle would affect the re-



(a) Start of Job 1a (pile with 45° slope, 11.72 m^3 volume specified) with HP method.



(b) Job 1b (pile with 30° slope, 7.49 m^3 volume specified) with ZC method.



(c) Job 1a coverage pattern with HP method (*left*) and ZC method (*right*).

Figure 7.4. Job 1: basic pile with 2 m-wide central section and 3.4 m base length; Area Tool 6.2 x 4.2 m.

sults. This pile has a volume of 7.5 m^3 detected above ground level. In Figure 7.4(b), two images from this job with the ZC method show the different way the loader approaches the pile compared with the HP method. Towards the end of the job, at right, the remaining pile material is contained more than with the HP method, which as Figure 3.15 showed earlier, tends to leave remaining pile material in an crescent shape.

Figure 7.4(c) shows an example of the coverage pattern for Job 1a with each method (using a minimum scoop filling ratio of 0.8), by rendering all positioning and loading paths used for completing the job. Both methods share the same points for load transfer to the trucks (B, C and D), 3 m

Table 7.1. Results for Job 1a with High Point (HP) and Zero Contour (ZC) method.

Meth.	Min. Load Ratio	Average results after 10 simulations per Minimum Load Ratio							
		Load Ratio	Driv. (m)	Turn. (rad)	Num. Loads	Num. Actions	Vol. (m ³)	Time (min)	Rate (m ³ /h)
HP	0.60	0.801	1292	437	109.1	115.0	13.106	58.0	13.562
HP	0.70	0.850	1218	412	102.9	108.5	13.113	54.6	14.397
HP	0.80	0.896	1160	392	98.2	103.4	13.190	52.1	15.202
HP	0.90	0.951	1078	367	92.0	96.0	13.119	48.5	16.239
ZC	0.60	0.799	1469	599	102.0	104.6	12.219	68.9	10.645
ZC	0.70	0.845	1392	568	96.4	98.2	12.223	65.3	11.230
ZC	0.80	0.900	1301	530	90.4	92.0	12.206	61.0	11.999
ZC	0.90	0.951	1236	505	85.7	87.4	12.224	58.1	12.633

from the workspace (see also Figure 6.2(a)). The light blue lines mark the positioning paths from the Load Approach point B to the Stage point(s), which is stationary for the HP method (point A), and which has a new location for each scooping action with the ZC method, marked by the blue dots. The yellow lines mark the loading paths, driven with the scoop lowered, from the Stage point towards the Scooping Destination, terminating where the bucket blade stopped. It should be noted that after unloading at point C and reversing to point B, the HP method requires the loader to turn 90° every time to the Stage point, while with the ZC method the loader turns directly at point B by various amounts.

Table 7.1 summarizes the results for Job 1a. For each method, the job was simulated 10 times using each of the four minimum load ratios, with the table showing the average values recorded. The average scoop load ratio is, as expected, approximately half-way between the minimum and full, since the volume extracted for each load is randomly selected between these values as explained in Section 7.2. The Driving and Turning values confirm that the ZC method does require more total driving than the HP method for a given scooping effectiveness to complete the job (also reflected in the total time). These are plotted in Figure 7.5, with the error bars representing one standard deviation.

The lower amount of total driving for the HP method is despite the greater Number of Loads extracted. Table 7.1 lists the Number of Loads and Actions separately, since the former is for the number of times the scoop was filled, whereas the latter also includes actions which partially filled the scoop, after which further actions were made until the scoop was filled. These values show that the HP method required more actions per full load, most likely due to the situation at the end of the job when the

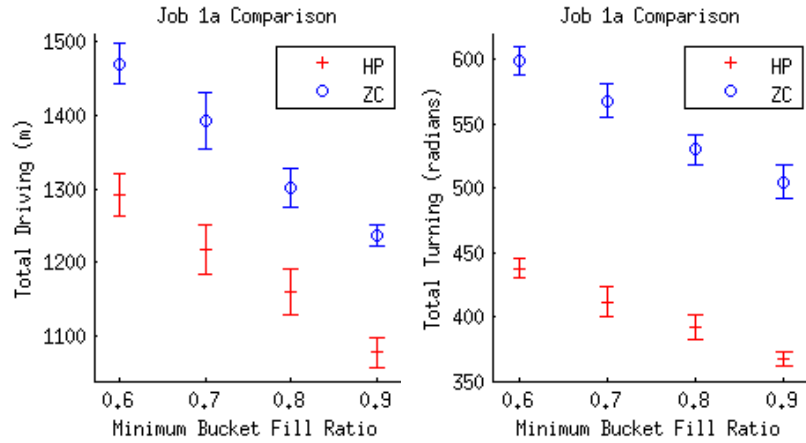


Figure 7.5. Total driving (*left*) and turning (*right*) vs. minimum bucket fill ratio for Job 1a with HP and ZC method; 10 trials per data point.

remaining material was spread out more than with the ZC method.

Related to the Number of Loads are the Volume values, which reveal that the HP method tended to excavate about 1 m^3 more than the ZC method to complete the job, with the ZC method excavating an amount closer to the 11.7 m^3 initially detected. This is despite both methods using the same general criterion for job completion (all points less than 0.15 m above ground level), and is likely a consequence of the HP method driving longer distances with the scoop lowered over already cleared areas of the workspace towards the end of the job. Since the simulator allows excavation to occur simply by the cutting plane intersecting the ground surface, extra material can be obtained when the scoop dips slightly below the ground while driving over cleared, uneven terrain. This effect can be reduced by dividing the area into smaller workspaces, which as the next sections show, also results in higher work rates.

Because of the difference in volume excavated to reach job completion with each method, the excavation rate can be used as a more general comparative value. This is plotted in Figure 7.6, vs. the minimum bucket fill ratio at left, and the same data with each simulation trial at right vs. the measured average bucket fill ratio. These also show an advantage for the HP method, with a rate more than $2 \text{ m}^3/\text{h}$ higher for a given scooping effectiveness. Even with the lowest minimum fill ratio of 0.6, the HP method has a higher rate than the ZC method with the highest ratio of 0.9 due to the total driving distance advantage.

One question when comparing these results is the separate contribution of driving and turning to the job completion time and work rate.

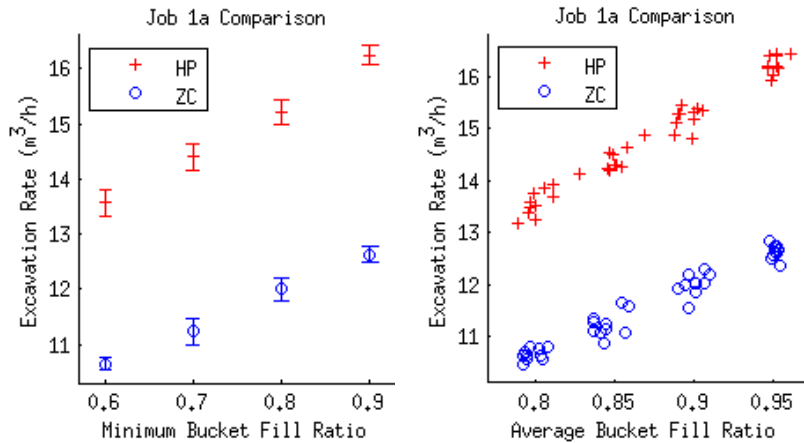


Figure 7.6. Comparison of excavation rate with HP and ZC method for Job 1a; at left, plotted vs. minimum bucket fill ratio (10 trials per data point); at right, same data with each trial plotted vs. average fill ratio.

Since a skid-steered loader is being used, all driving consists of turning on the spot at an arbitrary rate ($30^\circ/\text{s}$) and driving at an arbitrary speed (0.5 m/s) along straight segments. If one method tends to require more turning than the other, for example, then a higher turning speed should improve the performance of that method more than the other one. To investigate this, the same job was repeated with turning at $60^\circ/\text{s}$, twice the previous speed, with driving kept the same at 0.5 m/s . The results, plotted in Figure 7.7(a), show that the faster turning increased the rate of both methods by about $2\text{ m}^3/\text{h}$. For the rest of this chapter, the default driving and turning speeds of 0.5 m/s and $30^\circ/\text{s}$ are used.

The additional plot in Figure 7.7(b) is the excavation rate comparison for Job 1b, which is the pile with the same base contour but with 30° slopes. This shows a slightly lower rate than Job 1a, which intuitively makes sense since the pile has less volume distributed over the same area, so the loader clears the outer parts and must begin driving longer distances sooner than with the larger pile to reach new material.

7.3.1 Discussion

While the HP method was shown to offer an advantage over the ZC method for this particular job in this simulation environment, hardware tests would be needed to confirm the results. One reason is the situation towards the end of the job when the pile is getting low. Since the HP method tends to dig out the pile into a crescent shape, the resistance of the remaining material would decrease, likely making it difficult to collect with

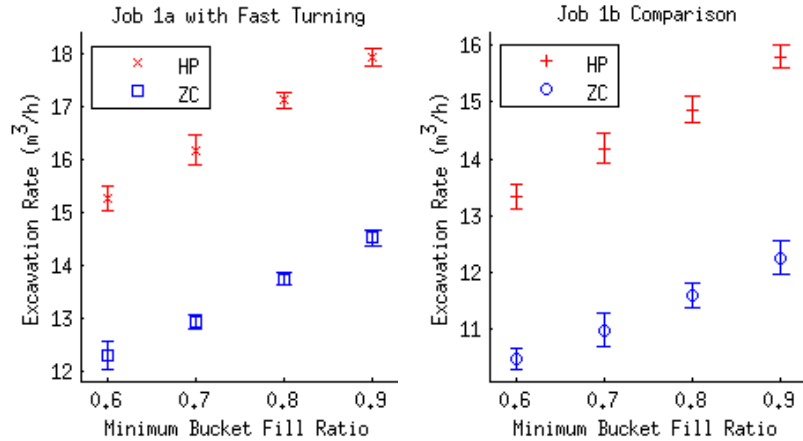


Figure 7.7. Excavation rate vs. minimum bucket fill ratio with HP and ZC method (10 trials per data point); at left, for Job 1a with 2x turning speed (60°/s); at right, for Job 1b (same pile footprint, slopes of 30°).

the bucket and resulting in some material being lost by being spread forwards. This effect, however, is not modelled in the simulator.

The ZC method, on the other hand, tends to keep the pile contained and would likely result in more effective loading and retention of material towards the end of the job. This would also depend on the type of material being loaded. Another factor that would need to be tested with the HP method is the effect of non-perpendicular loading approaches on scooping effectiveness and machine maintenance.

Although the HP method may be at a disadvantage due to the spreading of material at the end of a pile-loading job, it might be better-suited to jobs or phases of a job where there is sufficient material to push against, such as the hillside loading-type job presented in Section 6.2.3. This scenario will therefore be the focus for the rest of the chapter.

7.4 Scooping Method Comparison for Plateau Section

The next scenario, Job 2, consists of excavating into a 1 m high plateau with a 30° slope. The section excavated is 6 m wide by 1 m long, and is divided into two parts with a 3 x 1 m Scoop Area (SA) (see Figure 7.8(a)). Two versions of the job were simulated: Job 2a, with a constant 30° angle of repose, and Job 2b, with a random angle of repose from 30-40° selected after each scooping action. This variation was added to see if the effects of a more random soil behaviour could be noticed in the excavation rate.

For this job and those in the remainder of this chapter, the flat ground

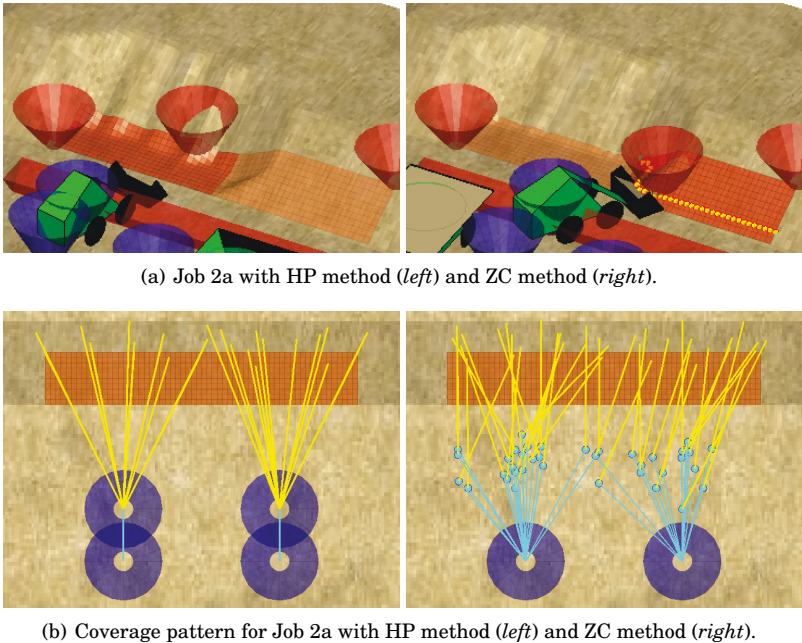


Figure 7.8. Job 2: 6 x 1 m plateau section divided into two parts, 1 m high with 30° slope.

from which the loader initially approaches the slope is a perfectly level plane. As with the large-scale hillside excavation in Figure 6.11, this helps to minimize the excavation of extra material below ground level which can occur in the simulator due to driving on uneven ground, so that the focus would be on the loading of the intended material.

The Area Tool surfaces used to mark the workspaces for this and the following jobs begin at the base of the slope, with ground level assumed to be known. This is so that when different SA dimensions are being tested later, the SAs in the first row contain a full amount of work without any flat ground at the front, which was required earlier to determine ground level.

Job 2 was made to consist of two parts in order to illustrate two types of slope sections that were identified for these kinds of jobs. “Type III” sections are those which have material collapsing in from all three surrounding sides, such as the left half of Job 2 (which is excavated first), while “Type II” sections, such as the right half of Job 2, only have material collapsing in from the front and one of the sides (here from the right). A “Type I” section would be a narrow protrusion with material only collapsing in from the front, as in the earlier simulation in Figure 3.21. Not all sections of the same type will have the same volume of material to excavate, as this also depends on the surrounding slopes, however these

Table 7.2. Results for Job 2a with High Point (HP) and Zero Contour (ZC) method.

Meth.	Min. Load Ratio	Average results after 10 simulations per Minimum Load Ratio							
		Load Ratio	Driv. (m)	Turn. (rad)	Num. Loads	Num. Actions	Vol. (m ³)	Time (min)	Rate (m ³ /h)
HP	0.60	0.802	433	191	47.1	47.1	5.662	21.4	15.883
HP	0.70	0.850	412	183	44.9	44.9	5.723	20.3	16.887
HP	0.80	0.902	382	170	41.8	41.8	5.654	18.9	17.920
HP	0.90	0.951	358	159	39.2	39.2	5.594	17.7	18.918
ZC	0.60	0.807	440	196	46.2	46.2	5.585	21.6	15.509
ZC	0.70	0.849	413	185	43.4	43.4	5.524	20.3	16.306
ZC	0.80	0.900	390	176	41.0	41.0	5.535	19.2	17.271
ZC	0.90	0.951	368	166	38.9	38.9	5.548	18.3	18.233

“Types” are used in an attempt to classify the jobs which are simulated.

Figure 7.8(b) shows examples of the Job 2a coverage pattern using each method, with a 0.8 minimum scoop fill ratio being used. While the radial pattern for the HP method resembles that for Job 1a quite closely, the ZC pattern is quite different than before since the sides of the workspace are not as exposed as with the pile. Approaches to contain the material from the side are therefore not required as much.

The results for Job 2a are summarized in Table 7.2, with these figures including the driving required to transfer to the second SA location after the first was completed. Comparing the driving, turning and job time, the HP method again has an advantage over the ZC method for a given filling ratio, though to a lesser degree than with Job 1a. One reason for these smaller differences is that the job itself is smaller, with less than half the volume of material to excavate. A similar trend is seen in the excavation rates, which are also plotted at the left in Figure 7.9. The main reason for the smaller advantage for the HP method compared with Job 1a is likely that the driving paths used by each method to complete the job are more similar, as Figure 7.8(b) showed.

Table 7.2 also shows that the Number of Loads and Actions are the same for each case, since the loader was pushing against the slope the whole time and no situation with small piles of leftover material and partial loads occurred. When comparing the volumes which were excavated to reach job completion, the HP values are again higher than the ZC values as in Job 1a, however this time only by about 0.1-0.2 m³. Since the initial approach ground was perfectly flat however, the unwanted excavation below ground level which occurred previously with the HP method should have been mostly avoided. The main reason the HP method still excavated

more material is likely due to its non-perpendicular approaches into the slope. In order to reach the material near the edges of the workspace, the HP scooping paths may cross beyond the workspace boundary more than with the ZC method, thereby excavating more material.

The excavation rates with both methods for Job 2a (ranging from 16-19 m³/h) are significantly higher than with Job 1a (11-16 m³/h). This is likely due to the smaller size of the workspace used by the loader (3 x 1 m vs. 6.2 x 4.2 m), meaning there is less driving to reach new material.

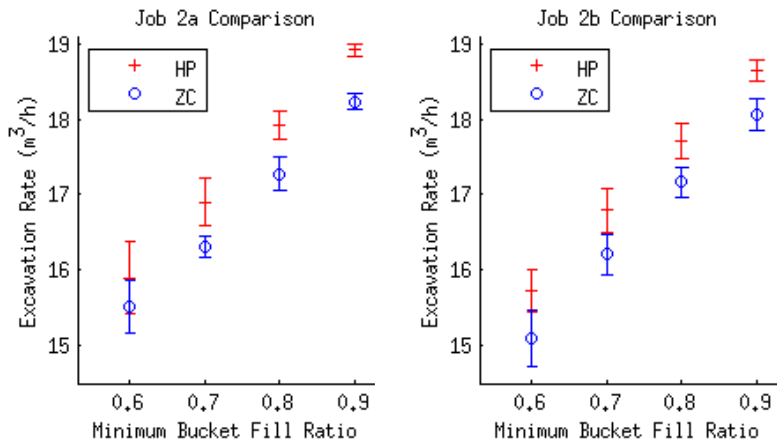


Figure 7.9. Excavation rate results for Job 2a (*left*), with 30° repose angle, and Job 2b (*right*), with 30-40° random repose angle; 10 trials per data point.

At right in Figure 7.9 are the excavation rates plotted for Job 2b, which was the same job with a random repose angle from 30-40° selected after each loading action. These rates are slightly lower (~0.1-0.4 m³/h) than in Job 2a, which could be explained by the higher average repose angle keeping material further from the front of the workspace, thus increasing the driving distance to reach new material.

In both plots a change in the error bars is noticeable, with the standard deviation growing smaller for the higher minimum filling ratios. This is because the filling ratio is randomly chosen from the minimum to 1.0, so the smaller the minimum, the greater the range of possible values. This effect was not observed in the Job 1 plots due to the larger number of scooping actions, which resulted in the average fill ratios converging.

7.4.1 Discussion

One main observation from Job 2 was that the differences between the HP and ZC excavation rates were much smaller than in Job 1, with the HP method maintaining an advantage for a given scoop filling ratio. These

smaller differences were likely due to the more similar coverage patterns used by both methods (see Figure 7.8(b)), with the ZC method not surrounding the workspace as when loading a pile and therefore driving less.

Another main observation was that the excavation rates with both methods were significantly higher than with Job 1. This seems due to the smaller current workspace (i.e. the 3 x 1 m SA) that was used, which reduced the driving distance to reach new material compared with the workspace of Job 1 (6.2 x 4.2 m), which was not subdivided. Since the current workspace size seems to have a significant effect on the excavation rate, the next section will repeat jobs with a range of SA widths and lengths in order to find those which result in the highest excavation rate for each scooping method.

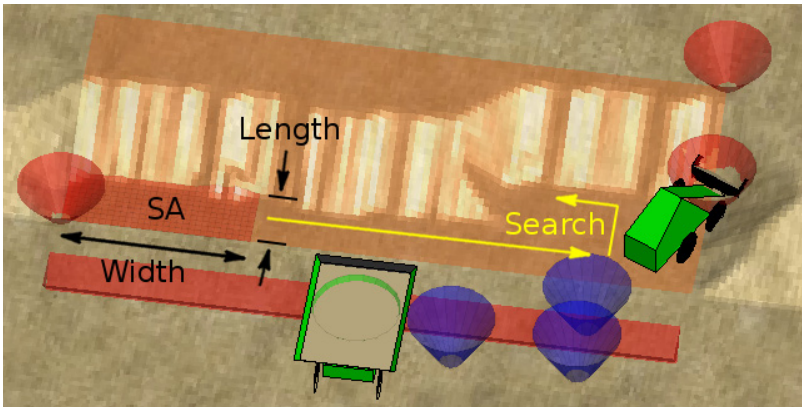
7.5 Optimal Scoop Area Dimensions and Workspace Division Strategy

In this section a job is repeated with changes to the width and length of the Scoop Area (SA), using both the High Point (HP) and Zero Contour (ZC) method, in an effort to find the dimensions which result in the maximum excavation rate and to continue comparing the two scooping methods. To find the optimal SA dimensions, a larger job which contains several SA locations should be simulated. This is because while a smaller SA at a single location may have a high excavation rate due to little driving needed to clear it, it will also be cleared sooner, requiring the machines (loader + 2 dump trucks) to reposition themselves more often at the next location. This extra driving would then lower the overall rate. To obtain a steady-state excavation rate for given SA dimensions, the job should therefore contain several SA locations so that the effect of the repositioning driving is included. For this reason Job 3 was simulated, which has a larger workspace than Job 2 and is subdivided both width- and length-wise using the SA strategy shown previously in Figure 6.6.

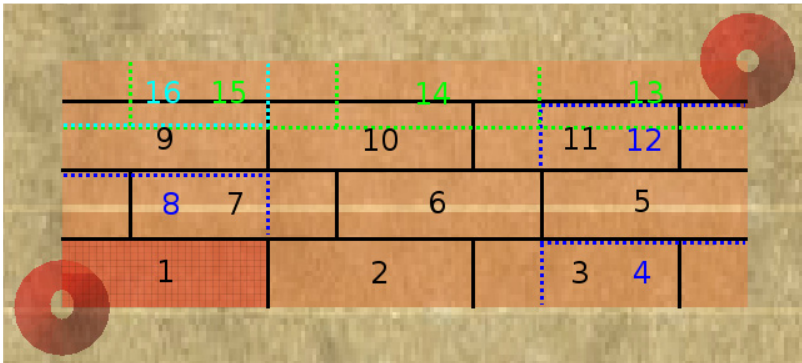
In addition, two different strategies for making the subdivision are compared, including with fixed SA dimensions which may not divide evenly into the total workspace, and with dimensions which are adjusted in an attempt to divide evenly. Since the SA dimensions are now the main variables for the remaining simulations, the scoop filling effectiveness is kept constant at a minimum fill ratio of 0.8 (average of 0.9). The results are therefore based on the assumption that both the HP and ZC method have

the same scoop filling performance.

Job 3 consists of excavating a rectangular area out of a plateau 0.87 m high with a 30° slope and constant repose angle, essentially a larger version of Job 2a but with a lower plateau height. Three versions of the job are simulated with changes to the overall width and length of the workspace, to test the effect on the subdivision strategies. Figure 7.10(a) shows an example of Job 3 and the scanning directions followed by the SA to find the next working location. In these single-loader jobs with workspace division, the truck is automatically positioned on the side of the loader which is towards the centre of the workspace. The trucks will thus always be stationed in level areas, making it possible for the machines to work in areas which are surrounded by slopes.



(a) Scoop Area (SA) scans workspace for next working location along raster pattern from front to back.



(b) Fixed SA area coverage pattern; can result in fractional workspaces.

Figure 7.10. Job 3: excavation of plateau section 0.87 m high with a 30° slope.

This job was set up as a constant-height plateau in an attempt to reduce the number of factors affecting the excavation rate, so the focus would be on the SA dimensions as the main variables. With such a plateau, each

row has the same amount of material collapsing in from uphill, whereas if the slope of a pile with a linear crest were being excavated, the amount of material collapsing in from uphill would be highest in the first row and decrease thereafter. Another factor affecting the amount of material collapsing in for a given row, however, is the geometry of the slope to the left and right. For Job 3, this would initially increase as the front slope is first excavated, then become constant when flanked by the plateau on both sides. Despite this initial transient phase, this job is used as a starting point to search for optimal SA dimensions.

Within each row excavated for Job 3, not every SA location will need the same amount of excavation work, since it will be either the Type III or Type II location mentioned in Section 7.4. The first SA in a new row would be Type III, with material collapsing in from three sides, while the remainder would be Type II, with material collapsing in from forwards and one side. The mix of different SA types should have an effect on the overall work rate, since Type III locations have more material to excavate for example. This is investigated in the next section with Job 4, which is the inverse of Job 3 and involves excavating a protruding plateau with Type II and Type I locations.

Another factor possibly resulting in different amounts of excavation work for the SA locations is that the SA dimensions being used may not divide evenly width- and length-wise into the overall workspace. Figure 7.10(b) shows how this is the case with a workspace 10 m wide by 3.6 m long and a 3 x 1 m SA, with the first “Fixed SA” subdivision strategy that was used. This strategy scans for the next work location (with ground heights 0.15 m or more above ground level) by shifting the SA along a row by one width-step, without any overlap, until the opposite side is reached, which may then overlap with the second-last location. Similarly, the next row is one length-step forwards.

In the figure, starting along the front (bottom) row from left to right, the black numbers and solid black lines indicate the first three SA locations at full sections of the slope, with location 4 indicated by the blue dotted line only used to cover the remaining slope portion at the right edge. The same sequence is repeated for the next two rows from location 5 to 12, after which locations 13 to 15 cover the partial back row. Location 16 then contains the smallest remaining slope portion in the corner.

This strategy is expected to be a disadvantage in cases such as this with uneven division, since it sometimes requires repositioning the machines

at a new location which only contains a partial amount of work. The larger the overall workspace, the less this would matter since most locations would be full, however it could have an effect on small- to medium-scale jobs. For a certain SA size, a partial division width-wise should not affect the amount of straight repositioning driving by the machines along a row, since with any division the full width will need to be traversed, however each repositioning at a new SA also requires turning which adds to the total driving. A partial division length-wise may have a greater effect since it adds a full extra row of driving. Nevertheless, the Fixed SA division method will first be used to ascertain its effect, with another strategy tested later which attempts to divide the total workspace evenly.

As noted in Section 6.2.3, one way to reduce repositioning driving by the machines could be to only reposition after one of the dump trucks is filled. This would require the loader to often temporarily extend its current working area, likely resulting in an uneven profile along the full workspace. Here it is assumed that the slope should be excavated evenly, therefore the machines reposition themselves when the current SA is cleared.

Since there is no overlap between the adjacent SA locations, slope collapse in the simulator can occasionally spill over into a previously-cleared location. If the spillage is above the height threshold, then it is found when rescanning for the next SA location, requiring the machines to drive back to that location for just one or a few scooping actions. This could be avoided by having an overlap, however it was not generally included in order to try to find the theoretical direct SA dimensions which are beneficial, and also because the spillage problem was found to occur relatively infrequently. As an example, for Job 3c presented later in this section, out of 120 simulations, with the workspace divided into 9-36 parts for a total of 2340 SA locations, spillage into a previous location occurred only 6 times. This occasional disadvantage can happen for both methods, therefore it should not affect the comparison of the HC and ZC method.

To study the effect of an overlap, a version of the SA with a 0.1 m extension forwards, at the trailing lateral side, and at the left and right edges of the full workspace was used for Job 3c. This virtually eliminated the spillage problem, which occurred twice out of 3803 SA locations, however it resulted in a reduction in the excavation rate due to the increased amount of driving needed to cover the larger SAs. If spillage occurred more frequently, which would depend on the ground material properties,

then including an overlap might be necessary.

7.5.1 Job 3 Results with Fixed SA Strategy

The first version of Job 3 simulated, Job 3a, is the excavation of a 12 m wide by 3 m long section of the plateau in Figure 7.10. The excavation rates resulting for this job with various SA dimensions are presented in Figure 7.11, using both the HP method (top) and ZC method (bottom), and the Fixed SA scanning strategy. The SA width is along the X-axis and excavation rate along the Y-axis, with separate lines plotted for each SA length. For each data point the simulation was repeated 5 times until job completion, i.e. until all ground heights in the workspace were less than the threshold height of 0.15 m above ground level.

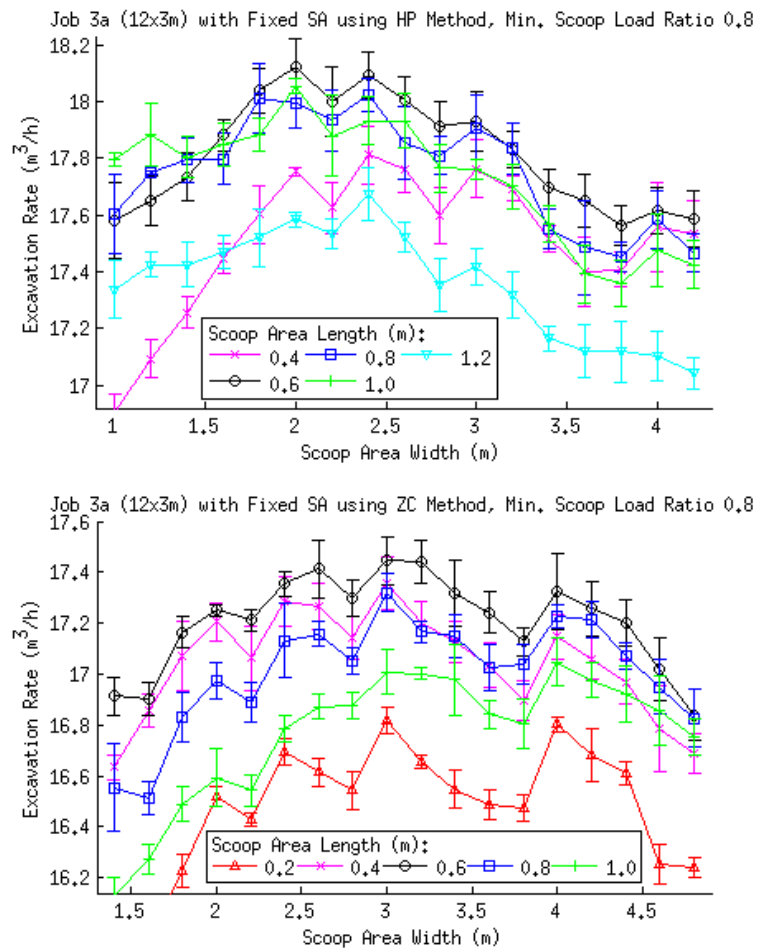


Figure 7.11. Job 3a: excavation rate for different SA dimensions using HP method (top) and ZC method (bottom); 12 x 3 m workspace; 5 trials per data point.

A range of SA dimensions was chosen for each method to show for which width and length values the excavation rate rises, reaches a maximum and then begins to decline, with increments in the values limited to 0.2 m. In general, the HP method achieves its highest rates with dimensions which are narrower and longer than with the ZC method, therefore the HP length values which were plotted range from 0.4-1.2 m while the ZC length values range from 0.2-1.0 m.

The highest rate with the HP method was 18.12 m³/h, achieved with an SA 2 m wide and 0.6 m long, with comparable rates obtained over the widths of 1.8-2.6 m. Over this width range slightly lower rates were achieved with lengths of 0.8 and 1.0 m. With the ZC method, the highest rate of 17.45 m³/h was achieved with a width of 3.0 m and also a length of 0.6 m, with relatively close values over the widths of 2.4-3.4 m and at 4.0 m. Over this range somewhat lower rates were achieved with lengths of 0.4 and 0.8 m.

Both plots of Figure 7.11 exhibit a general saw-tooth shape, with corresponding features occurring at the same width locations for both methods, such as the dips at 2.2, 2.8 and 3.8 m, and peaks at 2.0, 2.4, 3.0 and 4.0 m. This could have a relevant effect on the excavation rate, since some of these features occur in the regions of the maximum value. The reason for this shape seems to be a result of the division of the total workspace domain by the SA dimensions, since the peaks occur at SA widths which divide evenly into the total width of 12 m, while the dips occur at widths which divide with a remainder.

To investigate the possible effect of the width division further, a second version of the job, 3b, was simulated with the total workspace narrowed to 11 x 3 m. Figure 7.12 shows the results, with rates and a saw-tooth shape similar to before though with somewhat different features. A maximum rate of 18.08 m³/h was obtained with the HP method, this time with an SA length of 0.8 m and width of 2.2 m, though with similar values over the widths of 1.6-3.0 m and lengths of 0.6-1 m. With the ZC method, a maximum rate of 17.50 m³/h was obtained with SA dimensions 2.8 m wide and 0.6 m long, with relatively close rates over the widths of 2.2-4.0 m and lengths of 0.4-0.8 m.

By comparing some of the features of Figure 7.12 with Figure 7.11, it is confirmed that the width division is the reason for the saw-tooth shape. In the previous figure for Job 3a, for example, two prominent features in both plots are the dip and peak at 2.8-3 m and 3.8-4 m. As noted above, these

peaks (3.0 and 4.0 m) divide evenly into 12 m, while the dip locations at 2.8 and 3.8 m divide 4.3 and 3.2 times, respectively. In the current figure, these features have shifted to 2.6-2.8 m and 3.6-3.8 m, with the dip locations (2.6 and 3.6 m) dividing into 11 m 4.2 and 3.1 times, and the peaks (2.8 and 3.8 m) 3.9 and 2.9 times.

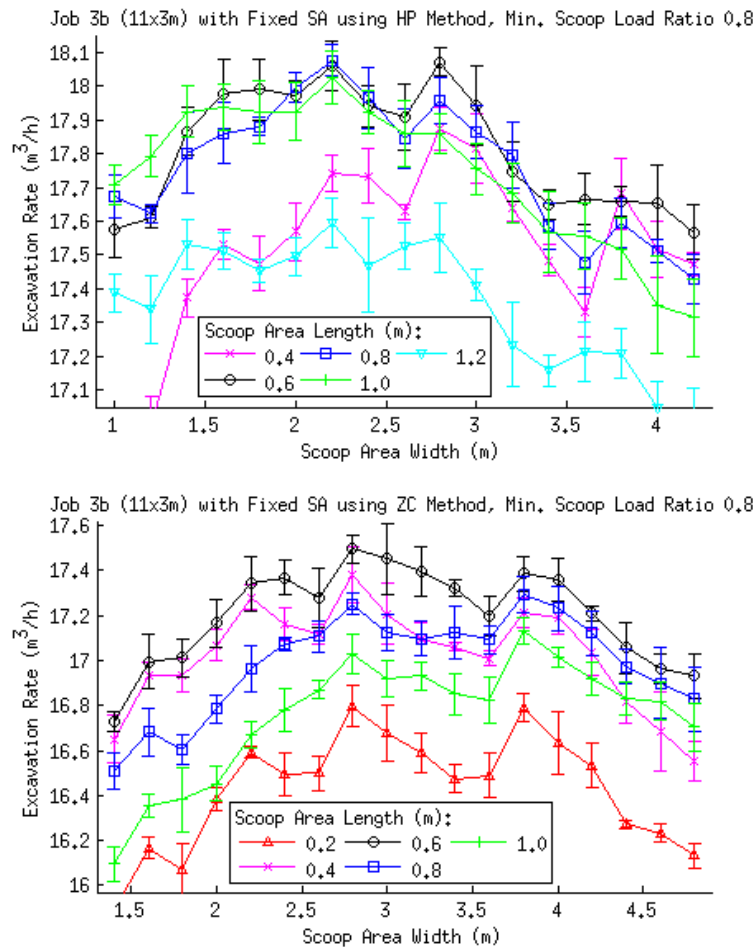


Figure 7.12. Job 3b: excavation rate for different SA dimensions using HP method (*top*) and ZC method (*bottom*); 11 x 3 m workspace; 5 trials per data point.

Thus with the Fixed SA division, the local peaks in excavation rate occur when the total width divides evenly by the SA width, or by just under a whole number, which corresponds to the last SA location along a row containing an almost-full amount of excavation work. Local decreases in rate, on the other hand, result when the width divides unevenly with a small remainder, corresponding to the last SA location along a row containing a small fraction of work, as in the example in Figure 7.10(b).

Another feature, visible in Figure 7.12(a) for the HP method, is that the

line for the 0.4 m length reaches its maximum rate at a wider value than the other lines. This makes sense since a shorter SA would need to be wider to contain enough work and avoid repositioning too often.

One additional version of this job, 3c, was simulated with the workspace shortened to 11 x 2.4 m, to check how a different length-wise division affects the rates. Figure 7.13 shows the results, with similar maximum rates as before. With the HP method, 18.18 m³/h was achieved with a width of 2.2 m and length of 0.8 m, with comparable values over the widths of 1.4-3.0 m and lengths 0.6 and 0.8 m, with the 0.8 m-length rates mostly higher. One difference this time is that the length of 1 m results in significantly lower rates than in Jobs 3a and 3b. This can be explained by the length-wise workspace division, since the previous total workspace length of 3 m divides evenly by 1 m, which is not the case with the new total length of 2.4 m. For the same reason, the length of 0.8 m, which now divides evenly, results in the higher rates than before.

With the ZC method, a maximum rate of 17.44 m³/h was achieved with a width of 2.8 m and lengths of 0.4 and 0.6 m. Similar values were obtained over the width range of 2.2-3.2 m and at 3.8 m, and lengths of 0.4-0.8 m, with 0.6 m usually highest. Effects of the length division are also evident here, with the lengths of 0.4 and 0.8 m resulting in higher rates than before since they now divide evenly into the total length (2.4 m), and the 1.0 m length resulting in somewhat lower rates since it no longer divides evenly.

Figures 7.11- 7.13 showed how the local minima and maxima in the excavation rate plots are a consequence of the Fixed SA workspace division strategy, i.e. the division of the total width and length of the workspace with the dimensions of the current SA being used. This intuitively makes sense since it would be undesirable for the machines to reposition themselves at a new SA if it only contains a partial amount of work. This could cause unwanted rate decreases in applications where the worksite dimensions are arbitrary but the machine always tries to use the same SA size which usually works well. To avoid this potential problem, a more general Variable SA strategy was developed, presented next in Section 7.5.2, which attempts to divide arbitrary workspace dimensions evenly. Jobs 3a-c were then repeated with this strategy for comparison.

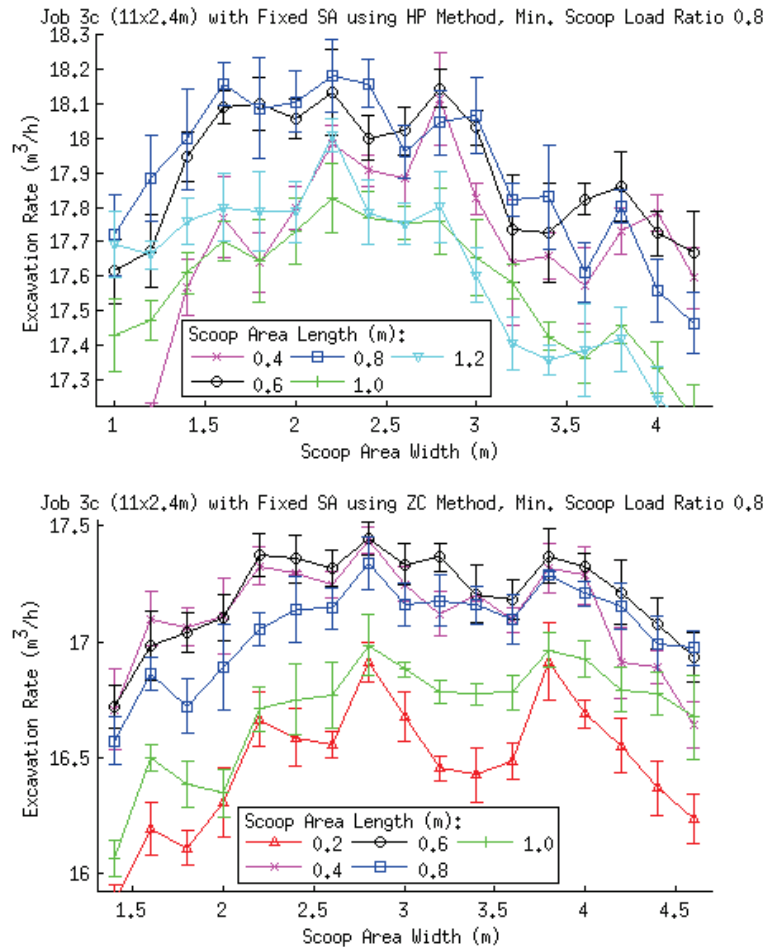


Figure 7.13. Job 3c: excavation rate for different SA dimensions using HP method (*top*) and ZC method (*bottom*); 11 x 2.4 m workspace; 5 trials per data point.

7.5.2 Job 3 with Variable SA Workspace Division

The Variable SA strategy is another way of dividing the workspace which attempts to avoid SA locations containing fractions of work and the associated effect on the excavation rate. It works by first dividing the total width and length of the workspace with a corresponding “target” SA width and length. The result for each dimension is then rounded, which is the number of partitions to be used. Remainders of 0.5 are rounded down to favour fewer partitions and less repositioning driving. Each workspace dimension is then divided by its partition number, resulting in the value for the first SA. A final rounding may be needed due to the world model grid resolution. For the next SA location, the remaining distance to the opposite workspace border is then divided by one less than the previous

partition number. Figure 7.14 shows an example, with the same 10 x 3.6 m workspace as in Figure 7.10(b) and target SA dimensions of 3 x 1 m.

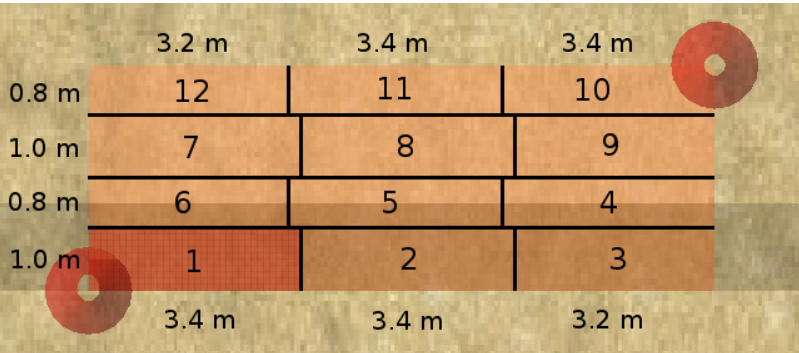


Figure 7.14. Variable SA method attempts to partition 10 x 3.6 m workspace evenly based on desired “target” SA dimensions, here 3 x 1 m. Increments are limited to twice the grid resolution of 0.1 m.

In this implementation, an SA location is defined at its centre and is rendered symmetrically, therefore width and length increments are limited to twice the grid resolution of 0.1 m, which is why the SA length in this example is not 0.9 m. Comparing with the Fixed SA strategy in Figure 7.10(b), the result is a division with 12 SA locations containing full amounts of work, rather than 16 locations, of which 9 were full and 7 fractional. Despite the advantage of less repositioning, a possible disadvantage is that the dimensions of the variable SA may stray from the optimal size. The larger the overall workspace and the more partitions, however, the closer the variable SA will generally be to its target size.

To compare the Variable with the Fixed SA division, Jobs 3a-3c were repeated with the new strategy, with the excavation rates plotted in Figure 7.15. For each data point the simulation was repeated 10 times. These plots have been reduced in scale compared with Figures 7.11-7.13 in order to save space, but still show the range of target widths and lengths over which the excavation rate increases, reaches a maximum, then decreases. The different range of SA lengths plotted for each method again reflects that the HP method tends to prefer longer SAs than the ZC method.

One difference in the Figure 7.15 plots, due to the new strategy, is that each width and length value which is plotted represents a range which is mapped by the Variable SA division algorithm to that value. The width ranges are represented by the dotted lines, thus in the top row for Job 3a for example, any target width between 2.1 and 2.7 m maps to 2.4 m. In the length direction, the target lengths from 0.2 to 1.0 m usually map to their own value, however with the Job 3c workspace length of 2.4 m, 1.0 m

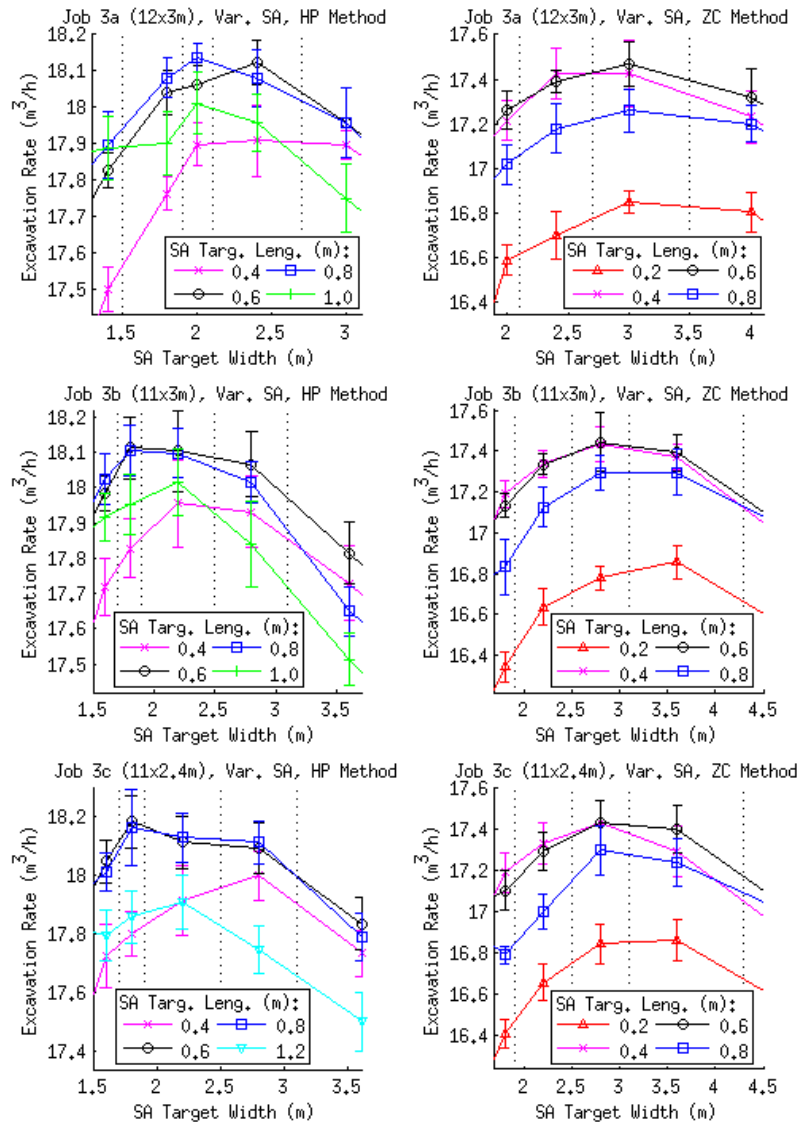


Figure 7.15. Jobs 3a-c with Variable SA workspace division: excavation rate for varying SA target dimensions using HP and ZC method; 0.87 m-high plateau, 0.8 minimum scoop load ratio, 10 trials per data point.

maps to 1.2 m, which is why this line appears in the bottom left plot.

Using the Variable SA strategy for Job 3, the maximum rates were similar to those achieved with the Fixed SA strategy, in the range of 18.11-18.18 m^3/h with the HP method and 17.43-17.47 m^3/h with the ZC method. The SA dimensions yielding these values, and values near the maxima, were also similar to before: widths of 1.8-2.8 m and lengths of 0.6-0.8 m for the HP method, and widths of 2.2-3.6 m and lengths of 0.4-0.6 m for the ZC method. Table 7.3 summarizes the maximum excavation

Table 7.3. Maximum Excavation Rates and associated SA dimensions for Job 3 with Fixed and Variable SA workspace division strategies.

Meth.	Job	Fixed SA			Variable SA		
		Max. Rate (m ³ /h)	Width (m)	Length (m)	Max. Rate (m ³ /h)	Target Wid. (m)	Target Leng. (m)
HP	3a	18.12	2.0	0.6	18.14	2.0	0.8
	3b	18.08	2.2	0.8	18.11	1.8	0.6, 0.8
	3c	18.18	2.2	0.8	18.18	1.8	0.6
ZC	3a	17.45	3.0	0.6	17.47	3.0	0.6
	3b	17.50	2.8	0.6	17.44	2.8	0.6
	3c	17.44	2.8	0.4, 0.6	17.43	2.8	0.4, 0.6

rates for all cases of Job 3 and the SA dimensions used to achieve them. In three instances, the same maximum rate was achieved with two different length values.

7.5.3 Discussion

One main result from Job 3 was that for each job version and division strategy, the highest maximum excavation rate was achieved using the HP method. The main assumption, again, is that both methods perform with the same scoop filling effectiveness (here a minimum 0.8 fill ratio), which would need to be tested with hardware experiments.

Another result was that each method achieved its maximum rate with different SA dimensions, with the HP method favouring narrower and longer dimensions (1.8-2.2 m wide, 0.6-0.8 m long) than the ZC method (2.8-3 m wide, 0.4-0.6 m long). The range of similar high rates for each method also followed this trend, with SA dimensions 1.6-3 m wide, 0.6-1 m long for the HP method, and 2.2-4 m wide, 0.4-0.8 m long for the ZC method.

The reason for this difference could be related to the heading angle at which the bucket tends to enter the slope. When entering close to the angle perpendicular to the bottom contour of the slope, the bucket should fill up relatively early since material will be encountered along the full width of the cutting blade. By filling the bucket early, such approaches would consequently reduce overall driving.

With the HP method, narrower SAs may result in more approaches which are closer to the perpendicular direction, since with wider SAs it would be expected that approaches towards the far ends would be more angled and result in uneven bucket filling. This would then result in more driving to fill the scoop and also more turning to reach the far ends. In

reality, it may be impossible to fill the bucket with extremely angled approaches due to high resistance encountered at one side of the bucket before it is filled. In pile loading experiments by Magnusson et al., it was found that with such angled approaches the automated bucket controller often failed to load any material [71]. If narrower SAs are advantageous for the HP method, then they would also need to be longer to contain enough material to avoid repositioning at new locations too frequently.

With the ZC method, all scooping approaches are generally close to the perpendicular angle due to the search process along the bottom contour. Wider and shorter SAs are then perhaps advantageous since scooping actions would fill up quickly with less forward driving into the slope, despite the extra turning and driving needed to reach the far ends.

Another factor affecting the optimal SA dimensions for each method could be the different situations after unloading at the truck. As Figure 7.4(c) shows, with the HP method the loader always turns 90° at point B to reach Stage point A. A narrower and longer SA then means that less additional turning is required at the Stage point. With the ZC method, the loader turns at point B to reach the current Stage point, and because the loader begins oriented towards point C, sometimes little turning is required if heading (in this example) to the right end of a wide SA. One area for future work could be to modify this configuration of load transfer points to find an optimum for different scooping methods.

When comparing the workspace division strategies in Table 7.3, the maximum rate was slightly higher with the Variable SA in three cases, equal in one case and slightly lower in two. A more important consequence of the Variable SA strategy, however, was that for a given method, it would have been possible to achieve the maximum rate for each job version using the same SA target dimensions. With the HP method, a target width of 1.9 m would have mapped to 2.0 m for Job 3a and 1.8 m for Jobs 3b and 3c, while a target length of 0.68 m would have mapped to 0.8 m for Jobs 3a and 3b, and 0.6 m for Job 3c. With the ZC method, any target width in the range of 2.7-3.0 m would have mapped to 3.0 m for Job 3a and 2.8 m for Jobs 3b and 3c, while a length of 0.6 m resulted in the maximum rate for each job version.

With the Fixed SA strategy, on the other hand, it was shown in Section 7.5.1 how a certain SA dimension which worked well for one job could result in a drop in the excavation rate for another job if it divided unfavourably into the total workspace. An example is the width of 2.0 m

for the HP method, which resulted in the maximum rate for Job 3a but then had decreased rates for Job 3b. Since the Variable SA method was shown to be more generally applicable, it is used for the next simulations in Section 7.6.

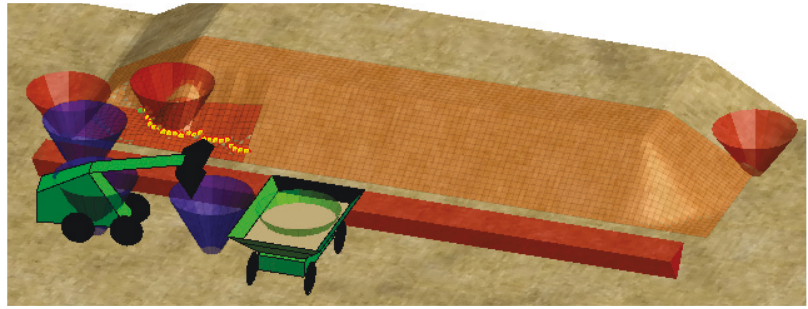
The results in this section are limited to this machine and also this particular plateau excavation job, with an arbitrary slope height and angle, yet they give a first idea about which SA size works well for each method. Job 4 in the next section will investigate the effects of various changes to the job parameters, including a different stand-alone plateau with less surrounding material collapse, a higher plateau and a steeper 45° slope.

7.6 Effect of Job Parameters on Optimal SA Dimensions

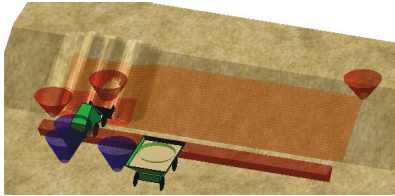
In the previous section, the optimal Scoop Area (SA) dimensions were found for a job with arbitrary specifications, including the flat, rectangular shape of the plateau section, the presence of surrounding slopes, the height and the slope angle. The optimal SA dimensions may be affected by some or all of these specifications, thus to obtain a more general result, a large number of random jobs should be simulated, consisting of piles and slopes of various shapes, uneven heights and different slope angles. This would be too computationally intensive for the present investigation, however in this section some variations to the job parameters are made to gain an insight into their effect on the optimal SA dimensions.

These job parameters include the plateau height, slope angle and presence of surrounding slope collapse. To compare the effect of surrounding slope collapse with Job 3, Job 4 is introduced, which is a stand-alone plateau whose width is fully contained in the workspace. It therefore consists of Type II and Type I SA locations, with material collapsing in from the front and either one side or neither side, whereas Job 3 was flanked by the surrounding plateau and consisted of Type II and Type III SA locations (“Types” introduced in Section 7.4). The first version, Job 4a, is otherwise the same as Job 3c, with a 0.87 m-high plateau, 30° slopes and 11 x 2.4 m workspace (see Figure 7.16(a)).

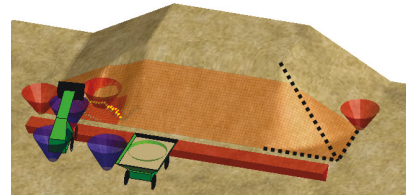
To investigate the effect of the plateau height, each type of plateau is doubled in height to 1.73 m in Jobs 3d and 4b (see Figures 7.16(b) and 7.16(c)), after which Jobs 3e and 4c will investigate the effect of a steeper 45° slope with the original plateau heights (Figures 7.16(d) and 7.16(e)). In all of these the workspace dimensions remain 11 x 2.4 m.



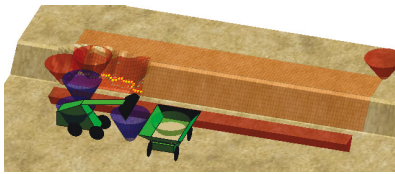
(a) Job 4a: 0.87 m-high stand-alone plateau, 30° slope.



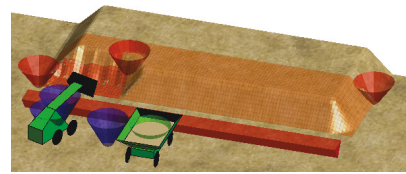
(b) Job 3d: 1.73 m-high plateau, 30° slope.



(c) Job 4b: 1.73 m-high stand-alone plateau, 30° slope (dotted line is alternate version).



(d) Job 3e: 0.87 m-high plateau, 45° slope.



(e) Job 4c: 0.87 m-high stand-alone plateau, 45° slope.

Figure 7.16. Jobs for comparing same 11 x 2.4 m workspace with different plateau types, heights and slope angles.

In this section a new measure, the *volume per combined drive time*, is also used to assess the performance of the SA dimensions. This measure divides the volume excavated during the job by the total amount of driving and turning time of all three machines. The reason for including this is to also consider that the highest possible excavation rate may not be the most important criterion for the job. As noted by Fu, fuel consumption and environmental impact have also become important considerations when optimizing earthmoving construction plans [50]. For space applications, the main limitation may be the amount of electrical energy available, while time may be plentiful.

The volume per combined drive time measure does not directly estimate energy consumption, but maximizing its value is similar to minimizing the total driving time by all machines, as a rough attempt to account for energy consumption. Instead of minimizing the driving directly, the volume is factored in since repeated simulations of the same job result in

small differences in the total volume excavated. It is expected that using this new measure, larger SA dimensions will become more beneficial since they contain more work and require less repositioning by the machines between SAs. Results using this new measure are presented in Section 7.6.4, but first comparisons will be made using the regular excavation rate.

7.6.1 Effect of Plateau Height

To observe the result of the higher plateaus, two comparisons can be made: from Job 3c to 3d, and Job 4a to 4b. Figure 7.17 shows the excavation rate results for Jobs 4a, 4b and 3d, all using the Variable SA strategy. The maximum rate for each case is summarized in Table 7.4 together with its corresponding SA dimensions. This table includes the previous excavation rate result for Job 3c to allow for easier comparison.

Comparing Job 4a with 4b in Figure 7.17, as the height of the plateau is doubled from 0.87 to 1.73 m, the maximum rate increases by about $0.4 \text{ m}^3/\text{h}$ for the HP method and $0.7 \text{ m}^3/\text{h}$ for the ZC method. These increases can be explained by the higher plateau, which results in more material to excavate per area and less driving to obtain a given volume.

With the HP method, the maximum rate for Job 4b is obtained with an SA of the same length as before (0.6 m), but a narrower width (1.8 m). A trend can be observed that the lines for shorter SAs rise higher in the plot. The 0.6 m line is now clearly above the 0.8 m line, whereas in Jobs 3a-c they usually overlapped, and the 0.4 m line has also risen closer to the top, above the 1.2 m line. Another trend is that the peak values for the 0.4-0.8 m lines shift to narrower widths.

With the ZC method the rates for the shorter SAs also increase, with the highest rate now with the shortest SA of 0.2 m, and the 0.4 m line clearly above the 0.6 m line. For each line the peak remains at the same width value, with the 0.2 m line's maximum rate at a very wide 5.6 m, corresponding to the workspace divided in half.

Comparing Job 3c (bottom of Figure 7.15) with 3d, the maximum rates achieved again increase as the plateau height is increased, by about $0.5 \text{ m}^3/\text{h}$ for each method. The same trend seen previously is also evident, that the lines for shorter SAs rise higher in the plot. With the HP method, the 0.4 m line rises near the top, above the 0.8 m line, while with the ZC method, the 0.4 m line rises above the 0.6 m line and the 0.2 m line above the 0.8 m line. The peak locations also shift to narrower values for some of the lines, including for those which achieve the maximum rates.

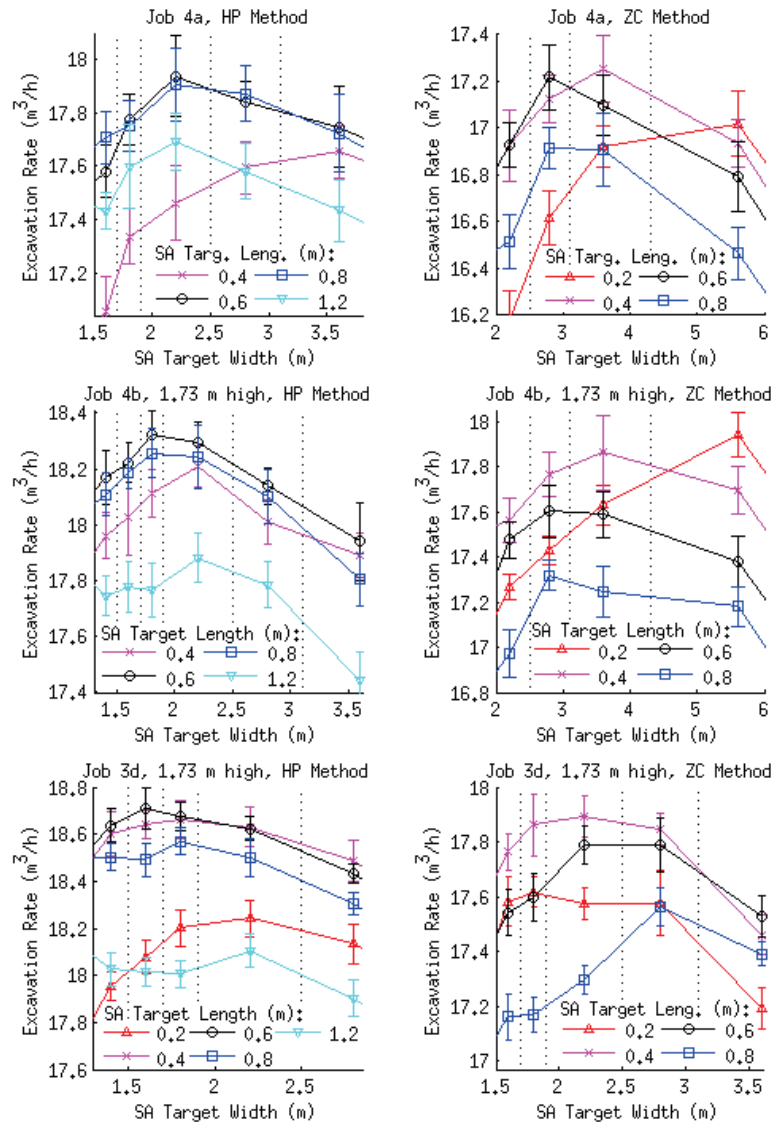


Figure 7.17. Comparison of excavation rate with changes in slope height and plateau type; 11 x 2.4 m workspace with Variable SA division, 0.8 minimum scoop load ratio, 10 trials per data point.

A general trend thus observed so far is that when the plateau height is increased, the maximum excavation rate is usually achieved with a narrower SA and in these cases, the same length. An exception was Job 4b with the ZC method, where the maximum was with a wider but shorter SA. As the rates for shorter SAs were seen to rise relative to longer values, then if the slope was made even higher, the highest rates might be reached with even shorter SAs. These trends towards smaller SAs make sense intuitively, since if a great deal of material collapses from uphill, then

enough material can be obtained in a small area and the extra driving needed to cover a larger area can be reduced.

One factor which affects the somewhat anomalous results of Job 4b with the ZC method is the shape of the bottom contour, which includes large rounded corners due to the conical corners of the plateau. If this curve is included in the SA, the ZC method could result in longer approach drives which attempt to contain the material from the side, as witnessed earlier with Job 1 in Figure 7.4. The 0.2 m-long SA may therefore be beneficial for Job 4b with the ZC method because it is too short to contain much of the curve, and would result in scooping approaches which are more from the front and not the sides. Since it is so short, it then needs to be wide (5.6 m) to contain enough material to avoid frequent repositioning. This factor could be a general reason why the ZC method prefers short SAs, since longer SAs could contain contours with more curvature and cause more approaches from the side which increase driving.

An alternate version of the Job 4b plateau, with the same height and which fits into the same space, could have a different corner geometry resulting in a rectangular footprint, indicated by the dotted lines in Figure 7.16(c). Although the full results are not included here, the excavation of this alternate version was also simulated, with a maximum rate of $17.72 \text{ m}^3/\text{h}$ achieved with a $2.8 \times 0.4 \text{ m}$ SA. The peak of the 0.2 m line also shifted to a narrower 3.6 m compared with Job 4a. These results therefore follow the trends observed otherwise, while it is clear that the shape of the plateau is another factor to consider rather than only the height.

7.6.2 Effect of Slope Angle

Next, the effect of a steeper slope is examined by comparing Job 4a with 4c and Job 3c with 3e. At the top of Figure 7.18 are the results for Job 4c, which was the same stand-alone plateau as Job 4a but with the slope increased from 30° to 45° . Looking first at the HP results and comparing with Job 4a in Figure 7.17, the same general plot shapes are observed, though with rates about $0.2 \text{ m}^3/\text{h}$ higher. This can be explained by the steeper slope, which helps the bucket to fill up sooner and therefore requires less driving for a given volume, also shown earlier when comparing Jobs 1a and 1b.

When comparing Job 4c with 4a for the ZC method, the maximum rate is slightly higher and is also reached with the 0.4 m long SA, though with a narrower width of 2.8 m. A change in the plots is noticeable which follows

the trends from the higher plateau, i.e. that the line for 0.4 m-long SA has risen clearly above the 0.6 m line, and that some peaks shift to narrower values (for the 0.2 and 0.4 m lines).

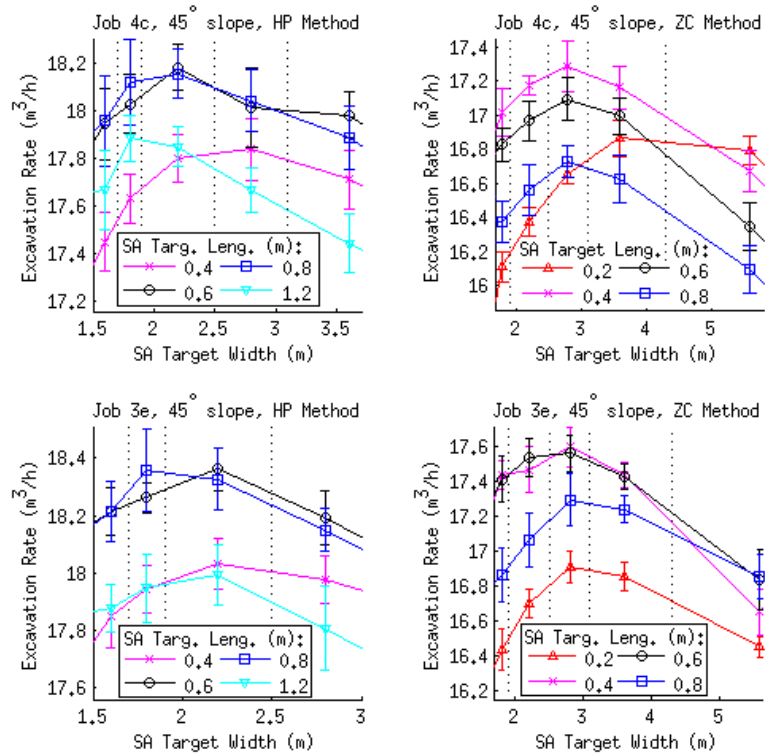


Figure 7.18. Job 4c and 3e with Variable SA workspace division: 0.8 minimum scoop load ratio, 10 trials per data point.

At the bottom of Figure 7.18 are the results for Job 3e, which was the same cut-away plateau section as Job 3c but with a 45° slope. The maximum rates for both methods are obtained with similar SA dimensions as Job 3c and are about 0.2 m³/h higher, due to the steeper slope.

In general, the steeper slope resulted in mostly the same SA dimensions yielding the highest rates. Some of the trends observed with the higher slope were also observed in the data, i.e. that when more material is present per area resulting in higher excavation rates, narrower and shorter SAs become more favourable.

7.6.3 Effect of Surrounding Slopes

To investigate the effect of the slopes surrounding the workspace, three comparisons can be made: from Job 3c to 4a (the original plateaus), from Job 3d to 4b (the higher plateaus), and from Job 3e to 4c (with 45° slopes).

Comparing the results for Job 4a at the top of Figure 7.17 with Job 3c from Figure 7.15, the maximum rates of Job 4a are about $0.2 \text{ m}^3/\text{h}$ lower, which can be explained by the lack of Type III SA locations and less material collapsing in from the surroundings than in Job 3c. Less material is therefore collected in the same area, resulting in lower rates overall.

These maximum rates were obtained with similar SA dimensions as with Job 3c, though with both methods shifting to a wider value: 1.8 to 2.2 m for the HP method and 2.8 to 3.6 m for the ZC method. With the ZC method, a value close to the maximum on the 0.6 m line remained at a width of 2.8 m. The trend of the maximum shifting to a wider value also noticeably occurs for the 0.4 m line with the HP method and the 0.2 m line with the ZC method. This could also be explained by the change in SA location types, since with less material collapsing in, a wider SA for a given length has the benefit of containing more material and requiring less repositioning. The longer SA lines, however, reach their maxima at the same width.

The change in the composition of SA types from Job 3c to 4a also has other effects which can be observed in the plots. In Job 4a with the HP method, the maximum rate for an SA length of 1.2 m rises higher than that for the length of 0.4 m, which was the opposite in Job 3c. With the ZC method, an interesting change is that relative to the other length values, 0.2 m rises significantly from Job 3c to 4a.

Comparing the plots for Job 4b with the plots for Job 3d in Figure 7.17, changes are evident when the plateau stays at the same height but changes to a straight section consisting of Type II and Type III SA locations. With the HP method, the maximum stays with an SA length of 0.6 m but shifts to an even narrower 1.6 m, while the line for 0.4 m has risen above the 0.8 m line. Meanwhile the 1.2 m line has sunk further relative to the others, and the 0.2 m line is included as it has resulted in rates comparable with the others. With the ZC method, the maximum rate is now achieved with a narrower-than-usual 2.2 m and length of 0.4 m, while the 0.2 m line is relatively further down than in Job 4b.

Comparing Job 3e with Job 4c in Figure 7.18, the two 45° versions, the maximum rates are obtained with similar SA dimensions but the rates are higher, due to more slope collapse from the surrounding plateau.

Some general trends thus far observed are that when more material is available per area from a higher slope, from more surrounding slope collapse and to some degree from a steeper slope, higher excavation rates are

Table 7.4. Maximum Excavation Rate and Volume per Combined Drive Time for Jobs 3c-e and 4a-c.

Job	HP Method			ZC Method		
	Max. Rate (m ³ /h)	Target Width (m)	Target Length (m)	Max. Rate (m ³ /h)	Target Wid. (m)	Target Leng. (m)
3c	18.18	1.8	0.6	17.43	2.8	0.4, 0.6
3d	18.71	1.6	0.6	17.89	2.2	0.4
3e	18.36	2.2/1.8	0.6/0.8	17.59	2.8	0.4
4a	17.94	2.2	0.6	17.25	3.6	0.4
4b	18.32	1.8	0.6	17.94	5.6	0.2
4c	18.18	2.2	0.6	17.29	2.8	0.4
	Max. Vol./ Comb. Drive Time (m ³ /h)	Target Width (m)	Target Length (m)	Max. Vol./ Comb. Drive Time (m ³ /h)	Target Width (m)	Target Length (m)
3c	15.52	2.8	0.8	15.00	3.6	0.6
3d	16.00	1.6/1.8/2.2	0.6/0.8/0.6	15.46	2.8	0.6
3e	15.53	2.2	0.8	15.00	3.6	0.6
4a	15.15	2.8	0.8	14.72	5.6	0.6
4b	15.52	2.8	0.8	15.41	5.6	0.4
4c	15.39	3.6	0.8	14.68	3.6	0.6

achieved with narrower and shorter SA dimensions. This is due to smaller SAs resulting in less driving within the SA, while the higher amount of material per area means that enough material is collected from each SA such that repositioning drives are not required too frequently. One exception to this was the wide (5.6 m), although short (0.2 m) SA that resulted in the ZC maximum rate in Job 4b, though this was affected by another factor, the shape of the bottom contour, which could be investigated further.

7.6.4 Volume per Combined Drive Time

Figures 7.19 and 7.20 show the results for Jobs 3c-e and Jobs 4a-c using the new *volume per combined drive time* measure described above. The maximum value for each case with its corresponding SA dimensions is also included in Table 7.4. When comparing these plots with the corresponding excavation rate plots for each job in Figures 7.15, 7.17 and 7.18, the peaks are observed to usually occur at wider values, while the lines for longer SAs also shift higher.

A general result from using the new measure is therefore, as expected, that the maximum values are obtained with larger SA dimensions than when plotting the regular excavation rate. This is due to the greater penalty for shifting to a new SA location, during which the driving of each machine

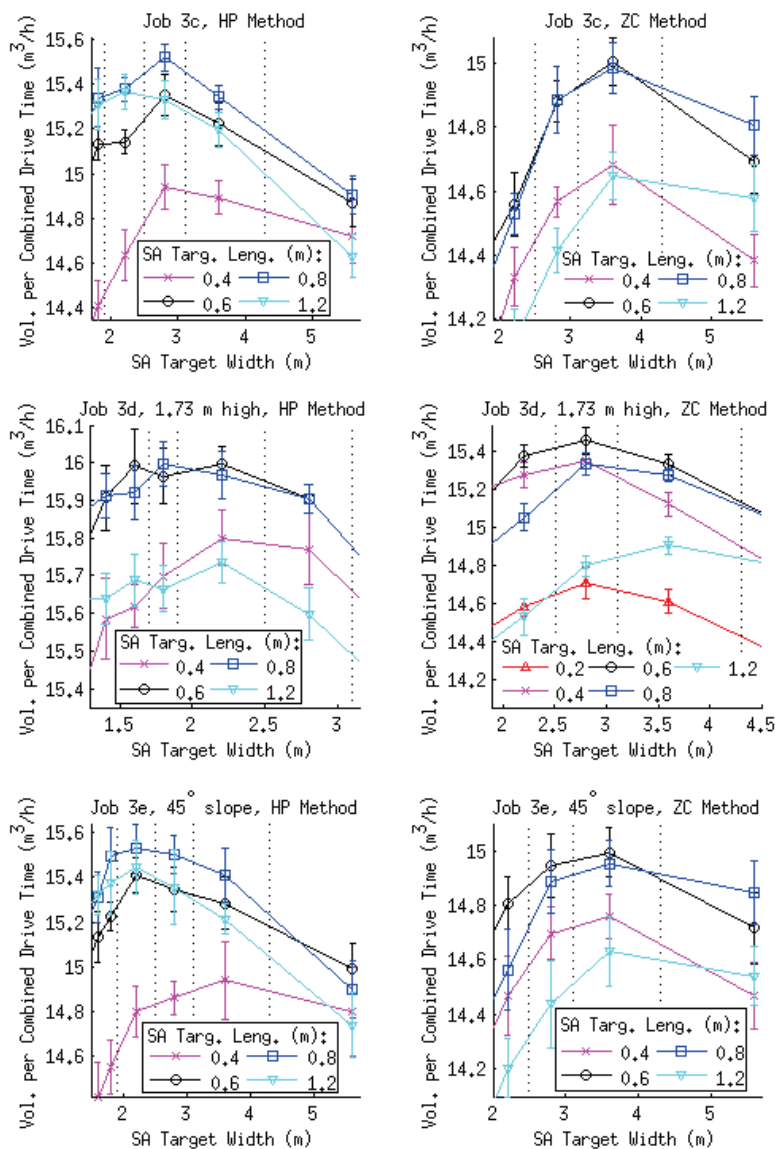


Figure 7.19. Volume excavated per combined drive time for Jobs 3c-e; 11 x 2.4 m workspace with Variable SA division, 0.8 minimum scoop load ratio, 10 trials per data point.

is added separately to the total. Larger SAs which contain more material and reduce the amount of repositioning are therefore beneficial.

7.6.5 Discussion

A general observation in this section was that whenever the amount of material to be excavated in the area increased, either from a higher plateau, steeper angle or more slope collapse, the maximum excavation rate and

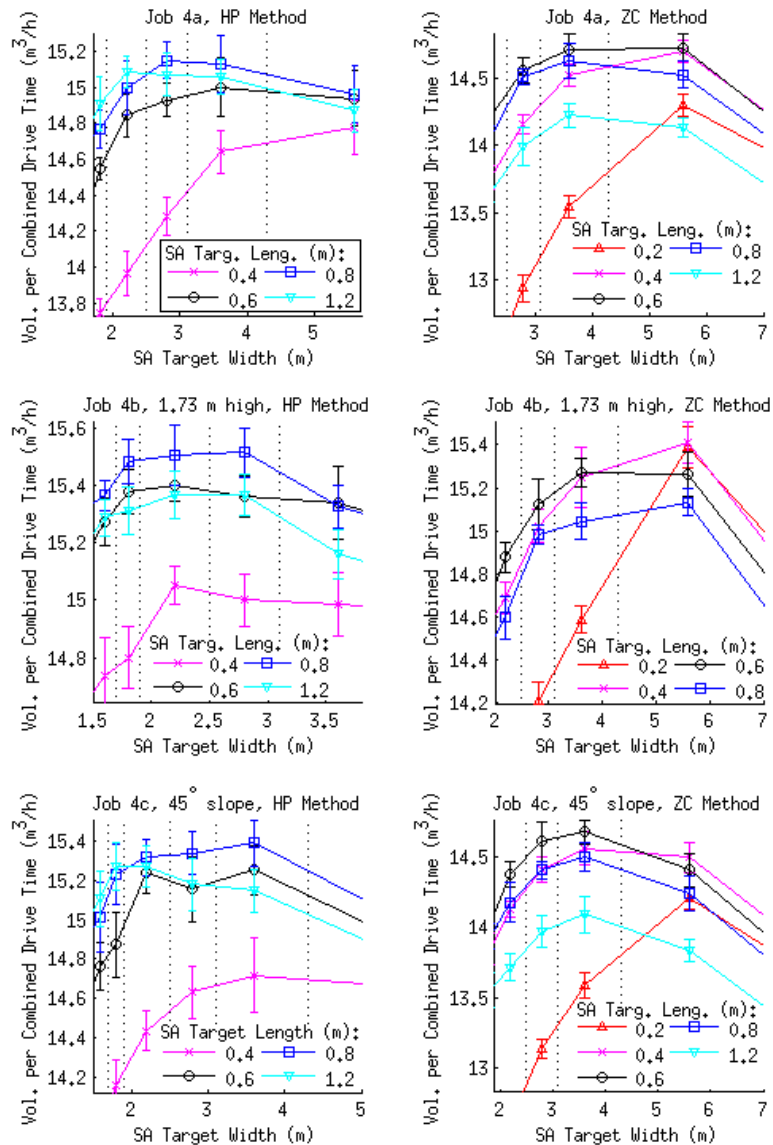


Figure 7.20. Volume excavated per combined drive time for Jobs 4a-c; 11 x 2.4 m workspace with Variable SA division, 0.8 minimum scoop load ratio, 10 trials per data point.

volume per combined drive time almost always increased (see Table 7.4). This was expected since with a higher material density, less driving should be required per volume collected. In one case the maximum volume per combined drive time stayed the same for both the HP and ZC method, from Job 3c to 3e. One seemingly anomalous result was Job 4b with the ZC method, which had a higher maximum excavation rate than Job 3d, despite having less surrounding slope collapse. This however was related to the shape of the bottom contour, since the maximum ZC rate of the al-

ternate version of the job mentioned at the end of Section 7.6.1 was lower than for Job 3d.

Continuing the comparison between the HP and ZC methods, higher excavation rates and volumes per combined drive time were achieved with the HP method in all cases, assuming the same scoop-filling effectiveness. When analyzing the SA target widths and lengths which yielded the maximum rates, the HP method again preferred narrower and longer SAs than the ZC method, as with Jobs 3a-c from the previous section.

For each method the maximum excavation rates for all six job versions were achieved with fairly consistent SA length values: 0.6 m with the HP method, and 0.4 m with the ZC method (except for Job 4b for which the length was 0.2 m). For the maximum volume per combined drive time, longer SA lengths were preferred by each method: 0.8 m for the HP method and 0.6 m for the ZC method (except again for Job 4b for which the length was 0.4 m). The changes to the slope parameters had a greater effect on the SA target widths which yielded the maximum rates. Usually, when the material density increased, the maximum rate was achieved with a narrower SA. One exception was Job 4b, for which a shorter (0.2 m) however much wider (5.6 m) SA yielded the maximum excavation rate.

As an area for future work, an automated planner could be developed which adjusts the current target SA dimensions being used in an attempt to maximize the excavation rate or minimize the total machine driving.

7.7 Multi-Loader Cases

The final simulation, Job 5, is a plateau excavation with 4 or 8 loaders (see Figure 7.21), to investigate if the same SA dimensions which worked well for a single loader would likewise function for multi-loader cases. The workspace division strategy presented in Section 6.2.5 is used to separate the Scoop Area (SA) into separate Scoop Zones (SZs) to prevent collisions between the loaders and to attempt to evenly excavate the slope face (see Figure 6.9). The SA and SZ dimensions are therefore fixed and no longer follow the Variable SA strategy from the previous two sections. Since the SA is now the larger area which contains the individual SZs for each loader, here it is the dimensions of the SZ which are investigated.

This plateau is 1.73 m high with a 30° slope and repose angle, and the specified workspace is 45 m wide and 5 m long. Here a 0.2 m surface grid size is used as in Section 6.2.5, due to the larger scale of the simulation.



Figure 7.21. Job 5 with eight loaders.

This width matches the MHP construction plan and hillside excavation jobs in Section 6.2.5, although this case does not include the high slope. For these simulations, the loaders do not work until job completion, but until a certain arbitrary excavated volume threshold is reached: 80 m³ for the 4-loader case and 100 m³ for the 8-loader case. This was an attempt to get an average rate while all the machines were working the whole time. If a job was simulated until completion, at the end some loaders would be standing idle until the last loader finished.

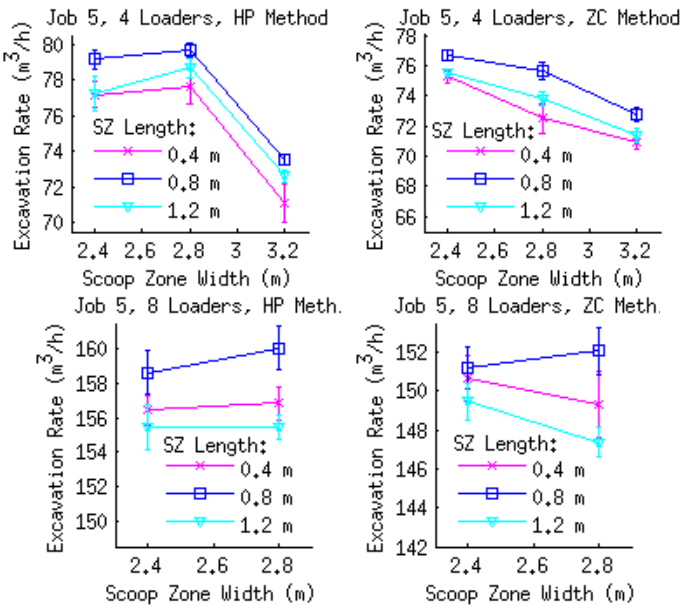


Figure 7.22. Job 5 excavation rates with 4 and 8 loaders: 0.8 minimum scoop load ratio, 10 trials per data point, Fixed SZ.

Figure 7.22 shows the excavation rate results for both the 4-loader and 8-loader cases, with the maximum rates summarized in Table 7.5 with their corresponding SZ dimensions. The minimum width is set at 2.4 m, which during the simulations appeared to be the narrowest SZ which still allowed enough maneuvering space for the machines. For the 4-loader plots, three width values are plotted until 3.2 m, since for both the HP

Table 7.5. Maximum Excavation Rate and Volume per Combined Drive Time for Job 5.

Num. Load.	HP Method			ZC Method		
	Max. Rate (m ³ /h)	Zone Wid. (m)	Zone Leng. (m)	Max. Rate (m ³ /h)	Zone Wid. (m)	Zone Leng. (m)
4	79.71	2.8	0.8	76.70	2.4	0.8
8	160.07	2.8	0.8	152.08	2.8	0.8
	Max. Vol./ Comb. Drive Time (m ³ /h)	Zone Width (m)	Zone Length (m)	Max. Vol./ Comb. Drive Time (m ³ /h)	Zone Width (m)	Zone Length (m)
4	17.84	2.4	0.8	17.24	2.4	0.8
8	17.94	2.8	0.8	17.11	2.8	0.8

and ZC method a downward trend exists from 2.8 to 3.2 m. For the 8-loader case the maximum SZ width is 2.8 m, which is the limit based on the 45 m workspace width, since the width of the full SA equals the SZ width multiplied by twice the number of loaders (44.8 m).

For all plots only three SZ lengths are included (0.4, 0.8 and 1.2 m), since they are sufficient for showing the rise, maximum and fall in the excavation rate for a given width. Dimension increments are limited to 0.4 m since the SZ is rendered from the centre, thus with a grid size of 0.2 m, the width or length must increase by 0.4 m to remain symmetrical.

As expected, the maximum rates for the 8-loader cases are approximately double those of their corresponding 4-loader cases. For both the 4-loader and 8-loader case, the maximum rate with the HP method is higher than with the ZC method, thus the advantage due to the more limited driving pattern of the HP method continues for these multi-loader scenarios. All cases achieve their maximum rates with SZ dimensions of 2.8 m wide by 0.8 m long except the 4-loader ZC case, for which the width is 2.4 m. This was somewhat unexpected, since in the previous simulations of this chapter for single-loader cases, the ZC method tended to prefer wider SAs than the HP method.

In Figure 7.23 the volume per combined drive time for the 4- and 8-loader simulations is plotted. The maximum values, also summarized in Table 7.5, were obtained with the same SZ dimensions as for the maximum excavation rates except for the 4-loader case with the HP method, when the width was 2.4 m. This was again unexpected, since previous results showed that the volume per combined drive time usually prefers a larger workspace to reposition less frequently.

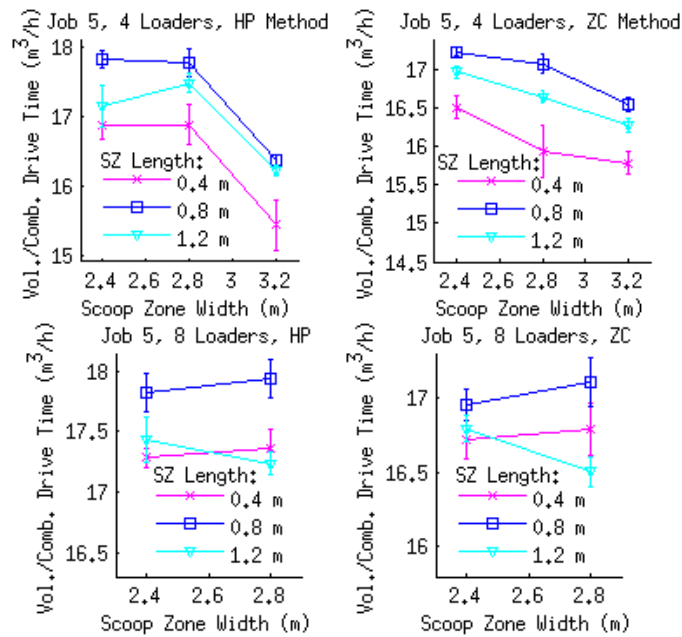


Figure 7.23. Job 5 volume per combined drive time with 4 and 8 loaders: 0.8 minimum scoop load ratio, 10 trials per data point.

7.7.1 Discussion

The multi-loader simulations in this section resulted in fairly consistent SZ dimensions which yielded the maximum rates, with lengths of 0.8 m in all cases and widths of 2.4 or 2.8 m. The trends noticed from the single loader simulations earlier in this chapter were not apparent here, since previously, the ZC method tended to favour wider and shorter SAs, and the maximum volume per combined drive time was usually reached with larger SAs.

One factor which may significantly affect the results is that with the workspace division strategy being used, the zone occupied by the dump trucks has the same width as the SZ being used by each loader. This means that a wider SZ increases the hauling distance for the loaders, which was not the case with single loader scenarios, where changes to the SA dimensions did not affect the load transfer configuration. An area for future work could be to modify the multi-loader workspace division strategy such that a constant load transfer distance is maintained by assigning the dump trucks a fixed narrower zone, and to see if this affects the slope profile shape over time.

The maximum excavation rate for the 8-loader case with the HP method can be compared with the result for Simulation 3 in Table 6.2, which was

the large hillside excavation with the same total workspace width and also SZ dimensions of 2.8 x 0.8 m (see Figure 6.11). The two main differences in the previous simulation were that it included a much larger slope, and that although the initial slope angle was 30°, the repose angle during excavation was 45°.

For Job 5 in this section, a rate of 160.07 m³/h was achieved with the 2.8 x 0.8 m SZ, while the previous simulation had a rate of 121.29 m³/h. This is an unexpected result, since the results so far in this chapter have indicated that higher excavation rates result from higher slopes (due to more material collapsing in per area) and steeper repose angles (due to the bucket filling sooner).

The reason for the lower rate from Section 6.2.5 could be the increased driving required due to temporary Scoop Zone extensions in the forward direction, which were illustrated in Figure 6.9. As the previous simulation was of a much longer duration, it included many of these extensions, which sometimes increment forwards more than once if a certain loader finishes its assigned zone much sooner than another. When a Scoop Zone is extended, the dump truck and load transfer locations remain unchanged, therefore the driving between scooping actions is significantly increased.

One reason SZs were frequently extended for the previous long-term simulation is that the left and right edges of the workspace were excavated slower than the centre portion, due to the material collapse from the sides. Thus while a single loader works at the right or left edge location, the others tend to finish their respective SZs, resulting in zone extensions. Since zone extensions result in a slower excavation rate due to the increased driving, the rate along the centre of the workspace would be slower than at the edge, which overall may help maintain the relatively even slope excavation witnessed in Figure 6.11. To maintain a high rate by avoiding zone extensions, dividing the loader pairs into separate teams as in Figure 6.12 should help, however this would likely result in an uneven, concave profile across the width of the workspace over time.

8. Conclusion and Future Work

This thesis presented high-level job planning capabilities which were developed to enable automated earthmoving performed by one or more robotic wheel loaders. Interactive 3D graphical planning tools were first developed in Chapter 3 which allow a remote user to specify a job visually on a surface model, including areas and approach directions for scooping, filling and clearing material.

Ways in which the planning tools could be developed further include making it possible to rotate the surfaces and to specify general polygonal shapes, in order to make them more flexible for various worksite layouts. Additional tools could also be developed to specify negative spaces for excavation below ground level. A more sophisticated Virtual Pile could have slope angles and dimensions which change automatically based on observed material behaviour and filling actions, to more accurately predict the end state of the job.

Algorithms were developed which automatically generate lower-level plans for a single loader scenario, ideally allowing the machine to work independently until job completion. These lower-level plans are also represented graphically, allowing a remote human operator to visualize and verify the current intentions of the machine. Pile transfer and area clearing simulations showed that the tools and algorithms are able to guide jobs to completion autonomously after the high-level plan is specified.

Among these algorithms was a method for finding the dumping location so as to approximately fill in the shape specified by the Virtual Pile, by scanning along rows with a line segment approximately as wide as the scoop. This was found to work for various pile shapes with a lateral scanning step increment of 0.1 m, row step of 0.2 m, and by deciding a location was full if only one ground point was above its desired height. For the larger-scale filling job in Section 6.2.3, the scanning increments were in-

creased to 0.2 m and 0.4 m, respectively, due to the larger grid size.

An algorithm for area clearing was also developed which finds the next clearing path by scanning along a line and comparing the average height with the average thickness of the layer cleared so far. A search threshold of 0.5 times the layer thickness was found to offer a good balance of a high volume clearance rate of at least 0.94 while avoiding excessive driving for the various jobs which were simulated.

Manually-controlled outdoor area clearing and filling tests in Chapters 4 and 5 verified that the planning tools and algorithms developed in the simulator could be used with surface models obtained with a 3D laser rangefinder to correctly track progress and autonomously generate commands. Data from these tests was also used to develop Augmented Reality versions of the planning tools, which may offer the possibility of increasing the situational awareness of a remote operator.

Ground models from the manual area clearing tests in Sections 4.1 and 4.4 were used to check the path search thresholds used in the simulations of Section 3.5. It was found that values of 0.23-0.38 times the layer thickness were able to find clearing paths which were considered to have enough material to clear, without getting stuck on negligible residual streaks. These values also resulted in a high clearance rate in the simulations, but sometimes resulted in a relatively high number of clearing drives.

In Sections 4.2 and 5.1, ground models from the manual area filling tests showed that the dumping location search algorithm from Section 3.4.5 was able to determine when locations were full using the same row search parameters as in the simulations. The algorithm also found free locations, though they did not always match those which were used, since during the experiments deposits were sometimes placed with a larger spacing. In particular, the length step between search rows had to be adjusted for each outdoor test to correctly track the job.

These tests led to three proof-of-concept demonstrations with a computer-controlled wheel loader, including an area clearing job and two pile transfer jobs. During these experiments full work cycles were achieved under computer control, with transitions between driving segments and scoop actions commanded manually. The tests validated the approach in general by allowing jobs to be specified using a recently obtained surface model, while plan modifications could easily be made on site by repositioning the graphical objects using the mouse. An automated scooping behaviour

based on monitoring driving pressure values was also developed.

The robotic snow clearing test demonstrated the full functionality of the area clearing tools, since updated driving commands were generated based on the changing state of the worksite. The robotic pile transfer tests showed that updates could have been made with the available data, but this was only confirmed in the subsequent analysis. The limited time available on site during the pile transfer tests only allowed for 3-4 full work cycles, which was mostly used to calibrate positioning and did not result in large-scale changes to the site topography.

If further tests were to be conducted, one goal would be to enable automated transitions, to allow for truly hands-free work cycles. Further pile transfer tests of longer duration, with more full work cycles, may make it possible to achieve automatic updates of the Scooping Destination and Dumping Location points. It could then also be verified if automated scooping with a preset bucket position can effectively level an area.

The pile transfer simulations were extended to a job with dump trucks and a multi-machine hillside excavation scenario in Chapter 6, showing how the same high-level planning tools could be applied on a larger scale. The benefits of construction simulation were demonstrated by the identification of machinery requirements when making a high-level site plan.

Automated workspace management strategies were also developed for the case of multiple loaders excavating along the hillside. The main goals were to keep the loaders separated to avoid collisions, and to excavate evenly along the slope face. These simulations allowed for the estimation of job completion time for various fleet and machinery parameters.

The simulation environment used for testing the job planning tools and algorithms was purely kinematic and did not model forces, thus its accuracy in modeling ground material behaviour could be tested with hardware experiments. A possible use for the simulator could be for predictive control in cases with remote supervision and long time delay. In this situation, observed ground behaviour could be used to update the simulator in an attempt to more accurately estimate the current state of work.

An area for future work could also be to evaluate the 3D graphical tools and/or Augmented Reality versions with user tests, to see if they offer an advantage for planning and supervising work by robotic machines, and modifying plans. A different approach to compare with could be entering coordinates to specify digging and dumping locations with only a 2D map available.

The 3D graphical planning interface developed here was limited to one remote user, yet the capability could be developed to allow multiple users to interact with and use it as a shared model to manage a job in a collaborative way. Some users may interact on an hourly or daily basis, while others in a management role may only provide higher-level input on a weekly or monthly basis. The same interface would be shared among robotic machines operating on site and be used to convey commands and receive updated sensor data, creating a *common presence* [90]. This concept can be linked to the Mars settlement construction scenario, with local supervisors on site and mission controllers on Earth using the same shared interface. The effect of a long communication time delay on issuing commands and updating information could also be an area for further study.

The simple alternative High Point (HP) method for generating scooping approach vectors was compared in Chapter 7 with a Zero Contour (ZC) method which was implemented. Simulations were conducted at different scales, beginning with a single pile and extending to larger hillside excavations requiring subdivision with the Scoop Area (SA), and also to multi-machine excavations with up to 8 loaders.

The HP method was shown to have a higher excavation rate for all jobs due to its more limited driving pattern, however in most cases this was with the assumption that both methods had the same average bucket filling effectiveness. This would need to be verified with hardware experiments, particularly due to occasional non-perpendicular approaches used by the HP method, which may make it difficult to fill the bucket.

The larger slope excavation jobs which required subdivision with the Scoop Area (SA) were also used to investigate two workspace division strategies. It was found that the Variable SA strategy was able to more generally result in higher excavation rates than the Fixed SA strategy, since the latter could sometimes result in an unfavourable fractional division which caused a decrease in the excavation rate.

From the larger simulations, it was found that the HP and ZC methods each achieved their maximum excavation rates with different SA dimensions, with the HP method favouring narrower and longer SAs than the ZC method. One reason for this could be that with the HP method, the loader always turns and drives forward to the same Stage point, thus a narrower SA helps to reduce further turning. With the ZC method, the loader turns by different amounts towards the current (moving) Stage point, thus it could reach the ends of wider SAs sooner.

Another reason for the different preferred SA dimensions could be that narrower SAs help the HP method achieve more perpendicular approaches (filling the bucket sooner), while since the ZC method by default achieves perpendicular approaches, shorter SAs help it to reduce forward driving. Short SAs containing contours with little curvature could also help reduce side approaches, and extra driving, for the ZC method. As an area for future work, the truck and driving point configuration beside the workspace could be adjusted for each method in an attempt to minimize extra driving and turning.

The optimal SA dimensions were also found to be affected by factors such as the height of the plateau being excavated and the presence of surrounding slope collapse, while the slope angle had less of an effect. In general, when more material was present per area, a trend was observed that higher excavation rates would be achieved with smaller SAs.

When the *volume per combined drive time* measure was used, larger SAs were generally preferred compared with the standard excavation rate due to the greater driving penalty of the machines repositioning at new SA locations. As an area for future work, since a given job may contain various slope shapes and heights, an automated algorithm could be developed which adjusts the SA dimensions being used, in an attempt to maximize the excavation rate or minimize energy consumption, for example.

The workspace subdivision dimensions were also investigated in the multi-loader simulations. The maximum rates were obtained with similar dimensions as with the single-loader cases, with little variation between the two methods and also between the two different rate measures. This comparison differed from the single loader case because increases in the Scoop Zone (SZ) width also increased the hauling distance to the trucks. To avoid this problem different division strategies could be used, which might also affect the shape of the excavation profile. Other approaches such as conflict-free trajectory envelopes [76] or bottom-up, behaviour-based strategies could also be implemented, which may result in beneficial non-intuitive machine formations.

All of the comparison results between the HP and ZC method, and SA dimensions which were found to be advantageous for the various cases, were specific to the machine being simulated. An important area for future work could therefore be to investigate how changes to the machine parameters, such as the bucket and chassis dimensions, or the use of centre-link steering, affect these results.

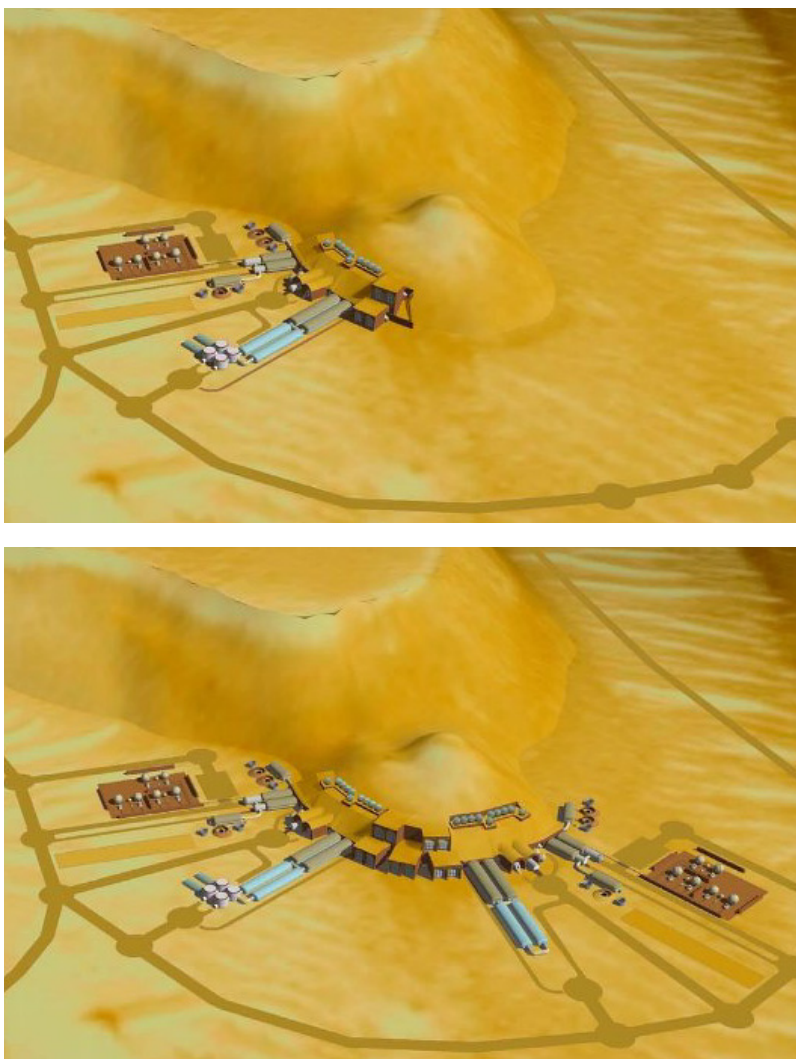


Figure 8.1. Linear growth of Mars Hillside Settlement from Phase I to Phase II, for 48 inhabitants. Design by Georgi Petrov, 2004 [1] (used with permission).

References

- [1] G. I. Petrov, A Permanent Settlement on Mars: The First Cut in the Land of a New Frontier, Master's thesis, Massachusetts Institute of Technology (2004).
- [2] V. A. Zhmud, E. Halbach, A. L. Pechnikov, V. G. Trubin, К вопросу об алгоритме управления снегоуборочными машинами (On the question of control algorithms for snow-clearing machines), Автоматика и программная инженерия (Automatics and Software Enginery) 2 (2) (2012) 65–71, <http://www.nips.ru>.
- [3] B. G. Drake, Human Exploration of Mars - NASA Design Reference Architecture 5.0, Tech. rep. (2009).
- [4] European Space Agency (ESA) CDF Study Report: Human Missions to Mars - Overall Architecture Assessment, Tech. rep. (2004).
- [5] E. Feichtinger, R. Charles, D. Urbina, P. Sundblad, C. Fuglesang, M. Zell, Mars-500 - A Testbed for Psychological Crew Support during Future Human Exploration Missions, in: IEEE Aerospace Conference, Big Sky, U.S.A., 2012.
- [6] R. Coppinger, Huge Mars Colony Eyed by SpaceX Founder Elon Musk (November 23, 2012).
URL <http://www.space.com/18596-mars-colony-spacex-elon-musk.html>
- [7] B. Mackenzie, B. Leahy, G. Petrov, G. Fisher, The Mars Homestead: a Mars Base Constructed from Local Materials, in: Space 2006, San Jose, U.S.A.
- [8] Y. Ishikawa, Utilization of Regolith for Manufacturing Construction Material on Mars, in: V. Badescu (Ed.), Mars Prospective Energy and Material Resources, Springer-Verlag, 2009, Ch. 19, pp. 543–550.
- [9] Komatsu Frontrunner (2016).
URL <http://www.komatsu.com>
- [10] Caterpillar Command for Underground (2016).
URL <http://mining.cat.com>
- [11] Rio Tinto Mine of the Future (2016).
URL <http://www.riotinto.com>
- [12] Sandvik AutoMine (2016).
URL <http://mining.sandvik.com>

- [13] T. B. Sheridan, *Telerobotics, automation and human supervisory control*, The MIT Press, 1992.
- [14] H. C. Schubert, J. P. How, *Space Construction: an Experimental Testbed to Develop Enabling Technologies*, in: SPIE 3206, *Telemanipulator and Telepresence Technologies IV*, Pittsburgh, U.S.A., 1997.
- [15] Avant (2016).
URL <http://www.avanttecono.com>
- [16] J. Saarinen, J. Suomela, A. Halme, *The Concept of Future Worksite - Towards Teamwork-centered Field Robotic Systems*, in: 18th International Federation of Automatic Control (IFAC) World Congress, Milan, Italy, 2011.
- [17] E. Halbach, *Development of a Simulator for Modeling Robotic Earth-Moving Tasks*, Master's thesis, Helsinki University of Technology (2007).
- [18] J. Saarinen, M. Hyvönen, J. Suomela, J. Vilenius, A. Halme, K. Huhtala, *Development of a Multi-Machine Remote Control Platform*, in: 10th Scandinavian International Conference on Fluid Power (SICFP), Tampere, Finland, 2007.
- [19] R. Ghabcheloo, M. Hyvönen, J. Uusisalo, O. Karhu, J. Järä, K. Huhtala, *Autonomous Motion Control of a Wheel Loader*, in: ASME Dynamic Systems and Control Conference, Hollywood, U.S.A., 2009.
- [20] R. Ghabcheloo, K. Huhtala, *Coordinated Control of Hydraulic Mobile Manipulators*, in: The 2nd IFAC Workshop on Multivehicle Systems, Espoo, Finland, 2012.
- [21] K. Rajapolvi, *Laserskannerin sinikommutoitu tasasähkömoottorikäyttö* (Sinusoidally Commutated Direct Current Motor Drive for Laser Scanner), Master's thesis, Tampere University of Technology (2012).
- [22] MathWorks MATLAB (2016).
URL <http://www.mathworks.com>
- [23] E. Halbach, V. A. Zhmud, A. Halme, *Simulation of Robotic Regolith Mining for Base Construction on Mars*, in: 12th Symposium on Advanced Space Technologies in Robotics and Automation (ASTRA), Noordwijk, The Netherlands, 2013.
- [24] E. Halbach, A. Halme, *Job-Level Control of Autonomous Earthmoving Machines for Base Construction on the Moon or Mars*, in: 61st International Astronautical Congress (IAC), Prague, Czech Republic, 2010.
- [25] E. Halbach, T. Ylikorpi, J. Suomela, J. Saarinen, A. Halme, *The Future Worksite Demonstrator: a hardware infrastructure for testing automated earthmoving for planetary applications*, in: 2nd Joint Planetary and Terrestrial Mining Sciences Symposium (PTMSS) and Space Resources Roundtable (SRR), Ottawa, Canada, 2011.
- [26] E. Halbach, A. Halme, *Augmented Reality Job Planning Interface for Robotic Construction on the Moon or Mars*, in: 11th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS), Turin, Italy, 2012.

- [27] E. Halbach, A. Halme, Job planning and supervisory control for automated earthmoving using 3D graphical tools, *Automation in Construction* 32 (2013) 145–160.
- [28] E. Halbach, A. Halme, Multi-Robot Hillside Excavation for Mars Settlement Construction, in: 65th International Astronautical Congress (IAC), Toronto, Canada, 2014.
- [29] A. Mishkin, Y. Lee, D. Korth, T. LeBlanc, Human-Robotic Missions to the Moon and Mars: Operations Design Implications, in: IEEE Aerospace Conference, Big Sky, U.S.A., 2007.
- [30] T. Huntsberger, G. Rodriguez, P. S. Schenker, Robotics Challenges for Robotic and Human Mars Exploration, in: Proceedings of ROBOTICS, Albuquerque, U.S.A., 2000, pp. 84–90.
- [31] L. E. Parker, D. Jung, T. Huntsberger, P. Pirjanian, Opportunistic Adaptation in Space-Based Robot Colonies: Application to Site Preparation, in: World Automation Congress, Maui, U.S.A., 2000.
- [32] G. S. Cheok, W. C. Stone, R. R. Lipman, C. Witzgall, J. Bernal, Laser Scanning for Construction Metrology, in: American Nuclear Society 9th International Topical Meeting on Robotics and Remote Systems, Seattle, U.S.A., 2001.
- [33] M. J. Chae, G. W. Lee, J. Y. Kim, J. W. Park, M. Y. Cho, A 3D surface modeling system for intelligent excavation system, *Automation in Construction* 20 (2011) 808–817.
- [34] H. Almqvist, M. Magnusson, T. Stoyanov, A. J. Lilienthal, Improving Point-Cloud Accuracy from a Moving Platform in Field Operations, in: IEEE International Conference on Robotics and Automation (ICRA), Karlsruhe, Germany, 2013.
- [35] L. S. Kang, H. S. Moon, N. Dawood, M. S. Kang, Development of methodology and virtual system for optimised simulation of road design data, *Automation in Construction* 19 (2010) 1000–1015.
- [36] GEOVIA MineSched (2016).
URL <http://www.geovia.com>
- [37] Ventyx MineScape (2016).
URL <http://www.ventyx.com>
- [38] Maptek Vulcan (2016).
URL <http://www.maptek.com>
- [39] V. R. Kamat, J. C. Martinez, Validating Complex Construction Simulation Models Using 3D Visualization, *Systems Analysis Modelling Simulation* 43:4 (2003) 455–467.
- [40] V. R. Kamat, J. C. Martinez, General-Purpose 3D Animation with VITASCOPE, in: Proceedings of the Winter Simulation Conference, Washington, D.C., U.S.A., 2004.
- [41] Novatron, Xsite (2016).
URL <http://www.novatron.fi>

- [42] A. Correa, M. R. Walter, L. Fletcher, J. Glass, S. Teller, R. Davis, Multi-modal Interaction with an Autonomous Forklift, in: 5th ACM/IEEE International Conference on Human-Robot Interaction, Osaka, Japan, 2010.
- [43] M. Dunbabin, P. Corke, G. Winstanley, J. Roberts, Off-World Robotic Excavation for Large-Scale Habitat Construction and Resource Extraction, in: T. Fong (Ed.), AAAI Spring Symposium, Palo Alto, U.S.A., 2006.
- [44] Maxis SimCity 2000.
URL <http://www.simcity.com>
- [45] R. Lipman, K. Reed, Using VRML in Construction Industry Applications, in: 5th Symposium on Virtual Reality Modeling Language (VRML), Monterey, U.S.A., 2000.
- [46] V. R. Kamat, J. C. Martinez, Large-Scale Dynamic Terrain in Three-Dimensional Construction Process Visualizations, *Journal of Computing in Civil Engineering* 19 (2) (2005) 160–171.
- [47] J. Shi, S. M. AbouRizk, A Resource-Based Simulation Approach with Application in Earthmoving/Strip Mining, in: Proceedings of the Winter Simulation Conference, Orlando, U.S.A., 1994.
- [48] J. O. Henriksen, Windows-Based Animation with Proof, in: Proceedings of the Winter Simulation Conference, Washington, D.C., U.S.A., 1998.
- [49] R.-Y. Huang, D. W. Halpin, Dynamic Interface Simulation for Construction Operations (DISCO), *Automation and Robotics in Construction X* (1993) 503–510.
- [50] J. Fu, Logistics of Earthmoving Operations: Simulation and Optimization, Licentiate thesis, Royal Institute of Technology (KTH) (2013).
- [51] M. Pla-Castells, I. García-Fernández, M. A. Gamón, R. J. Martínez-Durá, Interactive earthmoving simulation in real-time, in: Congreso Español de Informática Gráfica (CEIG) (Spanish Congress of Computer Graphics), San Sebastián, Spain, 2009.
- [52] D. Schmidt, M. Proetzsch, K. Berns, Simulation and Control of an Autonomous Bucket Excavator for Landscaping Tasks, in: IEEE International Conference on Robotics and Automation (ICRA), Anchorage, U.S.A., 2010.
- [53] CM Labs Vortex (2016).
URL <http://www.cm-labs.com>
- [54] E. Halbach, A. Hölttä, Virtual Skid-Steered Avant Wheel Loader with Simulated Digging Capabilities, in: Proceedings of GIMnet 2010, Espoo, Finland.
- [55] Open Dynamics Engine (2016).
URL <http://www.ode.org>
- [56] S. Sarata, Model-based Task Planning for Loading Operation in Mining, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Maui, U.S.A., 2001.

- [57] S. Singh, H. Cannon, Multi-Resolution Planning for Earthmoving, in: IEEE International Conference on Robotics and Automation (ICRA), Leuven, Belgium, 1998.
- [58] H. Cannon, S. Singh, Models for Automated Earthmoving, in: International Symposium on Experimental Robotics, Sydney, Australia, 1999.
- [59] J. Seo, S. Lee, J. Kim, S.-K. Kim, Task planner design for an automated excavation system, *Automation in Construction* 20 (2011) 954–966.
- [60] S.-K. Kim, J. Seo, J. S. Russell, Intelligent navigation strategies for an automated earthwork system, *Automation in Construction* 21 (2012) 132–147.
- [61] P. J. A. Lever, F.-Y. Wang, Intelligent Excavator Control System for Lunar Mining System, *Journal of Aerospace Engineering* 8 (1) (1995) 16–24.
- [62] G. J. Maeda, D. C. Rye, S. P. N. Singh, Iterative Autonomous Excavation, *Springer Tracts in Advanced Robotics*, 2014, pp. 369–382.
- [63] A. A. Dobson, J. A. Marshall, J. Larsson, Admittance Control for Robotic Loading: Underground Field Trials with an LHD, in: *Field and Service Robotics*, Toronto, Canada, 2015.
- [64] S. Sarata, Y. Weeramhaeng, T. Tsubouchi, Approach Path Generation to Scooping Position for Wheel Loader, in: IEEE International Conference on Robotics and Automation (ICRA), Barcelona, Spain, 2005.
- [65] S. Sarata, Y. Weeramhaeng, T. Tsubouchi, Planning of scooping position and approach path for loading operation by wheel loader, in: 22nd International Symposium on Automation and Robotics in Construction (ISARC), Ferrara, Italy, 2005.
- [66] S. Sarata, N. Koyachi, T. Tsubouchi, H. Osumi, M. Kurisu, K. Sugawara, Development of Autonomous System for Loading Operations by a Wheel Loader, in: 23rd International Symposium on Automation and Robotics in Construction (ISARC), Tokyo, Japan, 2006.
- [67] S. Sarata, N. Koyachi, K. Sugawara, Field Test of Autonomous Loading Operations by Wheel Loader, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Nice, France, 2008.
- [68] N. Koyachi, S. Sarata, Unmanned Loading Operation by Autonomous Wheel Loader, in: ICROS-SICE International Joint Conference, Fukuoka, Japan, 2009.
- [69] M. Magnusson, H. Almqvist, Consistent Pile-Shape Quantification for Autonomous Wheel Loaders, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), San Francisco, U.S.A., 2011.
- [70] H. Almqvist, Automatic bucket fill, Master's thesis, Linköping University (2009).
- [71] M. Magnusson, T. Kucner, A. J. Lilienthal, Quantitative Evaluation of Coarse-to-Fine Loading Strategies for Material Rehandling, in: IEEE International Conference on Automation Science and Engineering (CASE), Gothenburg, Sweden, 2015.

- [72] A. Bonchis, N. Hillier, J. Ryde, E. Duff, C. Pradalier, Experiments in Autonomous Earth Moving, in: 18th International Federation of Automatic Control (IFAC) World Congress, Milan, Italy, 2011.
- [73] M. Hess, M. Saska, K. Schilling, Autonomous multi-vehicle formations for cooperative airfield snow shoveling, in: 3rd European Conference on Mobile Robots, Freiburg, Germany, 2007.
- [74] M. Saska, M. Hess, K. Schilling, Efficient Airport Snow Shoveling by Applying Autonomous Multi-Vehicle Formations, in: IEEE International Conference on Robotics and Automation (ICRA), Pasadena, U.S.A., 2008.
- [75] K. Schreiner, NASA's JPL Nanorover Outposts Project Develops Colony of Solar-Powered Nanorovers, *IEEE Intelligent Systems* 16 (2).
- [76] F. Pecora, M. Cirillo, D. Dimitrov, On Mission-Dependent Coordination of Multiple Vehicles under Spatial and Temporal Constraints, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Vilamoura, Portugal, 2012.
- [77] C. A. C. Parker, H. Zhang, C. R. Kube, Blind Bulldozing: Multiple Robot Nest Construction, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Las Vegas, U.S.A., 2003.
- [78] J. Thangavelautham, A. Smith, N. Abu El Samid, A. Ho, D. Boucher, J. Richard, G. M. T. D'Eleuterio, Multirobot Lunar Excavation and ISRU Using Artificial-Neural-Tissue Controllers, in: Space Technology and Applications International Forum (STAIF), Albuquerque, U.S.A., 2008.
- [79] LEGO, Mindstorms NXT (2016).
URL <http://www.lego.com/mindstorms>
- [80] NXT-Python (2016).
URL <https://github.com/Eelviny/nxt-python>
- [81] G. Podnar, J. Dolan, A. Elfes, M. Bergerman, H. B. Brown, A. D. Guise-wite, Human Telesupervision of a Fleet of Autonomous Robots for Safe and Efficient Space Exploration, in: ACM/IEEE 1st Annual Conference on Human-Robot Interaction, Salt Lake City, U.S.A., 2006.
- [82] J. W. Crandall, M. A. Goodrich, D. R. Olsen Jr., C. W. Nielsen, Validating Human-Robot Interaction Schemes in Multitasking Environments, *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans* 35 (4) (2005) 438–449.
- [83] S. Mau, J. M. Dolan, Scheduling to Minimize Downtime in Human-Multirobot Supervisory Control, Tech. Rep. Paper 190, Carnegie Mellon Robotics Institute (2006).
- [84] C. Woodward, M. Hakkarainen, K. Rainio, Mobile Augmented Reality for Building and Construction, in: Mobile AR Summit @ the Mobile World Congress, Barcelona, Spain, 2010.
- [85] A. H. Behzadan, V. R. Kamat, Resolving Incorrect Occlusion in Augmented Reality Animations of Simulated Construction Operations, in: 15th International Workshop on Intelligent Computing in Engineering, Plymouth, UK, 2008.

- [86] P. Milgram, S. Zhai, D. Drascic, J. J. Grodski, Applications of Augmented Reality for Human-Robot Communication, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Tokyo, Japan, 1993.
- [87] V. Brujic-Okretic, J.-Y. Guillemaut, L. J. Hitchin, M. Michielen, G. A. Parker, Remote Vehicle Manoeuvring Using Augmented Reality, in: International Conference on Visual Information Engineering, Guildford, U.K., 2003.
- [88] T. H. J. Collett, B. A. MacDonald, Developer Oriented Visualisation of a Robot Program, in: 1st Annual Conference on Human-Robot Interaction, Salt Lake City, U.S.A., 2006.
- [89] M. G. Kleinhans, H. Markies, S. J. de Vet, A. C. in 't Veld, F. N. Postema, Static and dynamic angles of repose in loose granular materials under reduced gravity, *Journal of Geophysical Research* 116 (2011) 1–13.
- [90] A. Halme, Common Situation Awareness as Basis for Human-Robot Interface, in: M. O. Tokhi, G. S. Virk, M. A. Hossain (Eds.), *Climbing and Walking Robots*, Springer-Verlag, 2006, pp. 3–18.
- [91] M. A. Fischler, R. C. Bolles, Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography, *Communications of the Association for Computing Machinery (ACM)* 24 (6) (1981) 381–395.
- [92] T. Oksanen, A. Visala, Coverage path planning algorithms for agricultural field machines, *Journal of Field Robotics* 26 (2009) 651–668.
- [93] Riegl (2016).
URL <http://www.riegl.com>
- [94] A. Geiger, F. Moosmann, Ö. Car, B. Schuster, Automatic Camera and Range Sensor Calibration using a Single Shot, in: IEEE International Conference on Robotics and Automation (ICRA), St. Paul, U.S.A., 2012.
- [95] Q. Zhang, R. Pless, Constraints for Heterogeneous Sensor Auto-Calibration, in: *Computer Vision and Pattern Recognition Workshop (CVPRW)*, Washington, D.C., U.S.A., 2004.
- [96] J. Levinson, S. Thrun, Automatic Online Calibration of Cameras and Lasers, in: *Robotics: Science and Systems IX*, Berlin, Germany, 2013.
- [97] R. Ghabcheloo, M. Hyvönen, Modeling and motion control of an articulated-frame-steering hydraulic mobile machine, in: 17th IEEE Mediterranean Conference on Control & Automation, Thessaloniki, Greece, 2009.
- [98] C. Plagemann, S. Mischke, S. Prentice, K. Kersting, N. Roy, W. Burgard, Learning Predictive Terrain Models for Legged Robot Locomotion, in: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Nice, France, 2008.
- [99] K. Skonieczny, M. Delaney, D. S. Wettergreen, W. L. Whittaker, Productive Lightweight Robotic Excavation for the Moon and Mars, *Journal of Aerospace Engineering* 27 (4) (2014) 8p.

- [100] K. Zacny, Y. Bar-Cohen, Drilling and Excavation for Construction and In-Situ Resource Utilization, in: V. Badescu (Ed.), *Mars Prospective Energy and Material Resources*, Springer-Verlag, 2009, Ch. 15, pp. 431–459.
- [101] C. Atwood-Stone, A. S. McEwen, Measuring Dynamic Angle of Repose in Low Gravity Environments using Martian Sand Dunes, in: *44th Lunar and Planetary Science Conference*, The Woodlands, U.S.A., 2013.
- [102] M. B. Dias, D. Goldberg, A. Stentz, Market-Based Multirobot Coordination for Complex Space Applications, in: *7th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2003.

Appendices

A. Avant 320 Forward Kinematics

The Avant 320 compact wheel loader is modelled as a four-link kinematic chain, consisting of a chassis, two-part telescopic boom, and scoop. The chain involves six reference frames, the two redundant ones being for convenience. Figure A.1 shows the frames in relation to the vehicle, with the transformation matrices between them which make up the forward kinematic model detailed below. These matrices are used to transform a vector expressed in the outer frame into one that is expressed in the inner frame. A vector v expressed in Frame 2 for example, 2v , is left-multiplied by the 1-2 transformation matrix 1T_2 to express it in Frame 1 by

$${}^1v = {}^1T_2 {}^2v. \quad (\text{A.1})$$

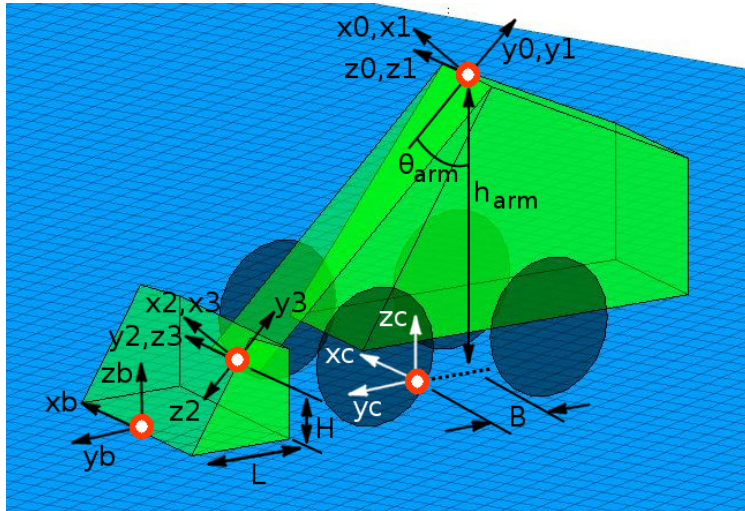


Figure A.1. Reference frames for Avant 320 forward kinematic model.

The vehicle reference frame "C" (for "Chassis") is located in the middle of the four wheels, in the plane defined by the bottom edges of the wheels. In this frame x_c points to the right while y_c points forward. The next frame,

"0", is located at the chassis-boom (or arm) hinge joint, and is related to Frame C by

$${}^C T_0 = \begin{pmatrix} 0 & 0 & 1 & 0 \\ \cos(\theta_{arm}) & -\sin(\theta_{arm}) & 0 & B \\ \sin(\theta_{arm}) & \cos(\theta_{arm}) & 0 & h_{arm} \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (A.2)$$

where $\theta_{arm} = 48.26^\circ$, the angle of the boom in home position (measured forward about z_0 from the downward direction), $B = -0.25$ m, the distance from Frame C to Frame 0 along y_c , and $h_{arm} = 1.18$ m, the height of Frame 0 above Frame C.

Frame 1 rotates with the boom about z_1 by joint angle θ_1 . Frames 0 and 1 are coincident in home position, when $\theta_1 = 0$, and are related by

$${}^0 T_1 = \begin{pmatrix} \cos(\theta_1) & -\sin(\theta_1) & 0 & 0 \\ \sin(\theta_1) & \cos(\theta_1) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (A.3)$$

Frame 2 is located at the boom-scoop hinge joint. x_2 and x_1 are parallel and separated by distance $d_2 + 1.367$ m, with d_2 being the prismatic boom extension. In home position $d_2 = 0$ m, the non-extended length of the boom, which places the scoop cutting plane 0.17 m above the ground when flat. In Figure A.1, the Avant is rendered with a boom extension greater than zero, placing the scoop closer to the ground. Frames 1 and 2 are related by

$${}^1 T_2 = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & -1 & -(d_2 + 1.367) \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (A.4)$$

Frame 3 rotates with the scoop about z_3 by joint angle θ_3 , which equals zero in home position when $\theta_1 = 0$ and the scoop blade is horizontal. Frames 2 and 3 are related by

$${}^2 T_3 = \begin{pmatrix} \cos(\theta_3) & -\sin(\theta_3) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -\sin(\theta_3) & -\cos(\theta_3) & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (A.5)$$

The last transformation is between Frame 3 and the scoop frame, called Frame B (for "Bucket"). This is located at the middle of the cutting edge and is oriented similar to Frame C when the scoop blade is horizontal.

Frames 3 and B are related by

$${}^3T_B = \begin{pmatrix} 0 & \sin(\phi) & -\cos(\phi) & D \cos(\psi) \\ 0 & \cos(\phi) & \sin(\phi) & -D \sin(\psi) \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (\text{A.6})$$

where $\phi = \theta_{arm} + \pi/2$, $D = \sqrt{L^2 + H^2}$, $L = 0.5$ m, the length of the cutting plane, $H = 0.1$ m, the distance between the scoop hinge and the cutting plane, and $\psi = \theta_{arm} + \arctan(H/L)$.

A vector expressed in the scoop frame can then be expressed in the chassis frame using the complete forward kinematic transformation

$${}^C T_B = {}^C T_0 {}^0 T_1 {}^1 T_2 {}^2 T_3 {}^3 T_B. \quad (\text{A.7})$$

In the scooping configuration, which had a variable bucket height (see Figure 3.16), $\theta_1 = 0$, $\theta_3 = 0$ and d_2 ranged from 0 m (home position) to 0.24 m. In the dumping configuration, $\theta_1 = 0.9$, $d_2 = 0.2$ m, and $\theta_3 = -1.9$.

In Figure A.1, H is actually equal to 0.2 m, the value used in the previous work by the author [17]. When experimenting with scoop height control strategies (see Section 3.4.4), this value was changed to 0.1 m, which ended up being used in this thesis. One consequence of this change is that at the lowest scooping configuration 0.01 m above the ground (see Figure 3.16), the scoop blade is 0.11 m further forward than it would be if H were equal to 0.2 m.

B. Graphical Tool Mouse Interface

The interactive translation and resizing of the graphical tools with the mouse, such as that presented in Sections 3.4.1 and 3.4.2, is made possible by using the Matlab Callback Functions *ButtonDownFcn* and *WindowButtonMotionFcn*. The former can be used to call a function when the left mouse button is clicked on a graphics object, while the latter can repeatedly call a function while the mouse is being moved within the figure window. Simplified examples of Matlab code are presented here to show how these Callback Functions were used to interact with 3D objects such as rectangular surfaces and cones.

In the following code segment, *object* is rendered using the arrays containing its X, Y and Z data. The function *StartDragSurf* is then assigned to be called whenever *object* is clicked, taking several arguments including the handles for the figure, axes, surface object and ground model. In the function *StartDragSurf*, the function *DragSurf* is assigned to be called repeatedly while the mouse button remains pressed and the mouse is being moved:

```
object=surf(figAxes,xData,yData,zData);
set(object,'ButtonDownFcn',{@StartDragSurf,figName,figAxes,object,ground});

function StartDragSurf(hObject,eventdata,figName,figAxes,object,ground)
set(figName,'WindowButtonMotionFcn',{@DragSurf,figAxes,object,ground});
```

Function *DragSurf* is then used to drag *object*, by modifying its geometrical data based on the position of the mouse. To achieve this, it is necessary to generate a ground model coordinate (x_i, y_i) corresponding to the mouse position in the screen. This is accomplished using the Matlab Axes Property *CurrentPoint*, which provides two points along a line perpendicular to the plane of the screen, passing through the location of the mouse pointer. These points, contained in the matrix *p*, are at the intersections of this line with the front and back of the axes volume. The coordinates

(x_i, y_i) indicated by this line are then found by approximating the ground surface with the plane $z = H$, the estimated ground level, and using a line-plane intersection equation. If the intersection strays beyond the domain of the ground model, x_i and/or y_i is set to its value at the nearest edge, thus all mouse positions are mapped to a valid coordinate. *object* is then moved or resized by first getting its current geometrical data with the *get* command, updating this data as needed using the (x_i, y_i) coordinates specified by the mouse, then assigning the new data with the *set* command:

```
function DragSurf(hObject,eventdata,figAxes,object,ground)
%% 1. Obtain mouse ground model coordinates:
Z=get(ground,'ZData'); H=ScanGround(Z); % ground level estimate
p=get(figAxes,'CurrentPoint'); % mouse - axes volume intersection points
% parameter "t" for mouse line intersection with plane "z=H":
t=-det([1 1 1 1;0 1 0 p(1,1);0 0 1 p(1,2);H H H p(1,3)])/...
det([1 1 1 0;0 1 0 p(2,1)-p(1,1);0 0 1 p(2,2)-p(1,2);H H H p(2,3)-p(1,3)]);
% coordinates of line-plane intersection:
xi=p(1,1)+(p(2,1)-p(1,1))*t; yi=p(1,2)+(p(2,2)-p(1,2))*t;
% limit range of drag point to domain of ground surface
%% 2. Process geometrical data of object:
curX=get(object,'XData'); curY=get(object,'YData'); curZ=get(object,'ZData');
% update curX, curY and curZ accordingly based on mouse coordinates (xi,yi)
set(object,'XData',newX,'YData',newY,'ZData',newZ);
```

Earlier, when the figure containing the ground model was first created, another Callback Function, *WindowButtonUpFcn*, would have been used to assign the calling of *StopDragFcn* whenever the left mouse button is unclicked. This function simply assigns nothing to *WindowButtonMotionFcn*, which causes *DragSurf* to cease its modification of *object*:

```
set(figName,'WindowButtonUpFcn',{@StopDragFcn,figName});

function StopDragFcn(hObject,eventdata,figName)
set(figName,'WindowButtonMotionFcn','');
```

Different versions of these functions were written for the various graphical tools and their components based on their functionality. For example, clicking on the rectangular surface component of the Area Tool would cause the surface to be translated together with the two corner cones, while maintaining the same dimensions. Clicking on a corner cone, however, would cause that cone to be translated while the other stayed stationary, and while the rectangular surface was resized to fit between the two cones.

C. Virtual Pile Mouse Mapping

When developing the Virtual Pile tool, the initial version of the interface only resized the base area of the pile if the mouse position corresponded to the corner coordinates of a permissible pile geometry, based on the height and scoop width restrictions and constant volume mentioned in Section 3.4.2. These permissible coordinates fall within a “crescent” of points, illustrated in Figure C.1. Here, the far cone is stationary while the near cone is being dragged to specify a new pile footprint. In Figure C.1(a) the mouse position at the near cone falls within the crescent, specifying a pile of intermediate height between the minimum and maximum. Normally this crescent is not visible, making it difficult for the user to know the domain of permissible positions. This initially resulted in an inconvenient user interface where the Virtual Pile tended to get stuck on a certain shape as the user was trying to change it by moving the mouse around, since the mouse position often strayed outside this crescent.

For a smoother functionality, it is necessary to map all mouse positions to a permissible pile shape. Figure C.1(b) shows the zones that were identified for this purpose. When the mouse falls within one of these zones, its position is mapped to the crescent edge. In Zone A1, the shape snaps to a square pile of maximum height. In Zone C the pile length is increased while the width is reduced and the height remains maximum. Length and width refer to the rectangular dimensions of the top section of the pile (see Figure 3.12), while in this appendix length will correspond to the greater of the two values.

As the mouse moves into Zone D (Figure C.1(c)), the length continues to increase while the width has reached zero and the height is decreased. In Zone F (Figure C.1(d)) the position is mapped to the near edge of the crescent, specifying a minimum-height pile. An additional zone “E” exists beyond the ends of the crescent (not shown here), which snaps to the

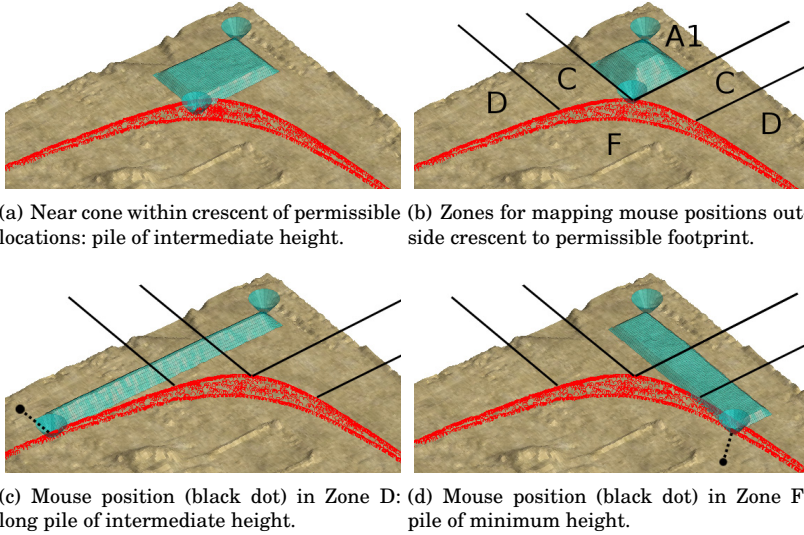


Figure C.1. Interface for “Type 1” Virtual Pile: far cone stationary and near cone being dragged with mouse. Crescent of points marks permissible positions of moving corner cone given pile height and width restrictions and constant volume. Mouse positions outside crescent are mapped to crescent edge. Zone descriptions provided in Figure C.2.

longest pile of minimum height. The result of this mapping strategy is a fluid interface that responds smoothly to all mouse positions as the Virtual Pile is being resized.

The mapping presented in Figure C.1 is limited to “Type 1” piles with $V \geq V_1$, the volume of the pile of maximum height with a central square section of scoop width sw , shown in Figure C.2(a). For piles with less volume, three additional mappings were identified as the volume is reduced past the threshold values in Figures C.2(b) and C.2(c) to the minimum of one scoop-load (Figure C.2(d)). With a 45° slope angle, vehicle specifications $sw = 0.89$ m, $maxH = 0.85$ m and $minH = 0.346$ m yield values $V_1 = 2.602$ m³, $V_2 = 1.286$ m³ and $V_3 = 0.531$ m³, with V_4 being equal to the 0.15 m³ scoop capacity.

The mappings for these additional pile types are illustrated in Figure C.2 on the right side. Here, the far cone is again stationary, with the near cone being dragged to specify a new length l and width w for the top section of the pile. l refers again to the greater and w the lesser rectangular dimension, indicated by the arrows in Figure C.2(e). When the mouse strays outside the red crescent of permissible locations into one of the zones, its position is again mapped to the crescent edge. A description of the values taken by l , w and height h in each mapping zone is provided in the table at the bottom of the figure.

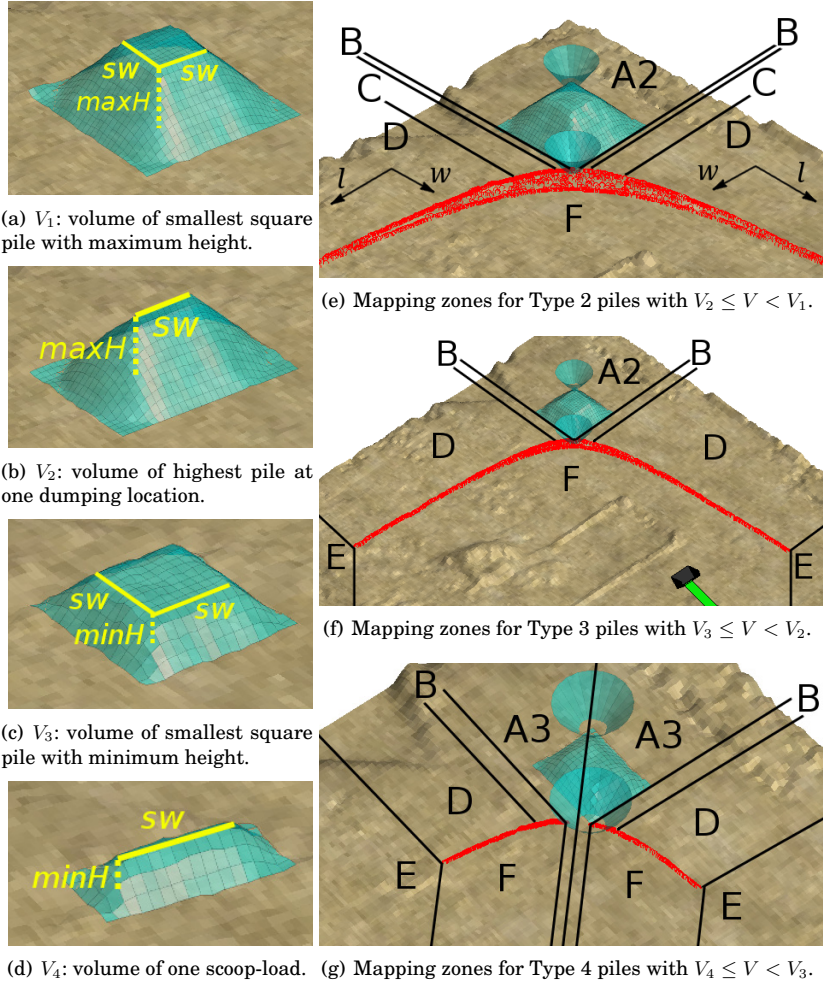


Figure C.2. Virtual Pile volume classification (*left*) based on limits defined by scoop width sw , maximum height $maxH$ and minimum height $minH$. Additional mouse position mappings (*right*) for Type 2-4 piles with zone descriptions (*bottom*). Far cone is stationary and mouse used to drag near cone. Length l is greater and width w lesser of pile's top rectangular section, "inter" refers to intermediate height.

For these mappings, the mouse position is taken to specify the distances $r + l$ and $r + w$ from the stationary cone, which correspond to the location of the near top corner of the pile's central rectangular section. If the mouse position is rather taken directly as the new corner cone position, corresponding to $r + l + r$ and $r + w + r$, two solutions for r will sometimes exist when solving Equation 3.2, and the mouse position cannot be used to uniquely determine the pile shape.

These additional pile mappings are not needed if the pile volumes encountered stay above the relatively low threshold of $V_1 = 2.602 \text{ m}^3$. If the maximum pile height is increased however, which could be the case with different machinery (such as in Section 6.2), threshold volumes V_1 and V_2 can increase substantially as they are proportional to the cube of the maximum height. Thus in some cases, the additional mouse position mappings could be useful when specifying a desired pile for depositing material.

Automated earthmoving offers the possible benefit of increasing both safety and efficiency, by separating human workers from potentially hazardous mining and construction sites, and allowing one remote user to supervise several robotic machines. In this dissertation, high-level planning for automated earthmoving is studied for jobs performed with a compact wheel loader, include scooping, dumping and clearing ground material. Algorithms are developed for determining where to dig, deposit or clear material based on the changing state of the worksite such that progress is made towards the goal state. The problem of excavating a slope face evenly while maximizing the excavation rate is studied in detail, with two methods for generating scooping approach vectors also compared.



ISBN 978-952-60-6855-8 (printed)
ISBN 978-952-60-6856-5 (pdf)
ISSN-L 1799-4934
ISSN 1799-4934 (printed)
ISSN 1799-4942 (pdf)

Aalto University
School of Electrical Engineering
Department of Electrical Engineering and Automation
www.aalto.fi

**BUSINESS +
ECONOMY**

**ART +
DESIGN +
ARCHITECTURE**

**SCIENCE +
TECHNOLOGY**

CROSSOVER

**DOCTORAL
DISSERTATIONS**