

Molten Regolith Electrolysis Reactor Modeling and Optimization of In-Situ Resource Utilization Systems

by

Samuel Steven Schreiner

B.S., University of Minnesota (2013)

Submitted to the Department of Aeronautical and Astronautical
Engineering

in partial fulfillment of the requirements for the degree of

Master of Science in Aerospace Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

© Massachusetts Institute of Technology 2015. All rights reserved.

Author
Department of Aeronautical and Astronautical Engineering
May 21, 2015

Certified by.....
Jeffrey A. Hoffman
Professor of the Practice of Aeronautics and Astronautics
Thesis Supervisor

Accepted by.....
Paulo C. Lozano
Associate Professor of Aeronautics and Astronautics
Chair, Graduate Programs Committee

Molten Regolith Electrolysis Reactor Modeling and Optimization of In-Situ Resource Utilization Systems

by

Samuel Steven Schreiner

Submitted to the Department of Aeronautical and Astronautical Engineering
on May 21, 2015, in partial fulfillment of the
requirements for the degree of
Master of Science in Aerospace Engineering

Abstract

In-Situ Resource Utilization (ISRU), the practice of leveraging space resources to support space exploration, has long been considered as a possible avenue for reducing the mass and cost of exploration architectures. In particular, producing oxygen from lunar regolith holds great promise for maintaining life support systems and enabling orbital refueling of chemical propulsion systems to reduce launch vehicle mass. Unfortunately, significant uncertainty as to the mass, power, and performance of such ISRU systems has prohibited a rigorous quantitative analysis.

To this end, parametric sizing models of several ISRU systems are developed to better understand their mass, power, and performance. Special focus is given to an oxygen production technique, called Molten Regolith Electrolysis (MRE), in which molten lunar regolith is directly electrolyzed to produce oxygen gas and metals, such as iron and silicon. The MRE reactor model has a foundation of regolith material property models validated by data from Apollo samples and regolith simulants. A coupled electrochemical and thermodynamic simulation is used to provide high-fidelity analysis of MRE reactor designs. A novel design methodology is developed that uses data from the simulation to parametrically generate mass, volume, power, and performance estimates for an MRE reactor that meets a set of performance criteria.

An integrated ISRU system model, including an MRE reactor, power system, excavator, liquid oxygen storage system, and other systems, is leveraged in a hybrid optimization scheme to study the optimal system design and performance characteristics. The optimized models predict that a 400 kg, 14 kW MRE-based ISRU system can produce 1,000 kg oxygen per year from lunar Highlands regolith. A 1593 kg, 56.5 kW system can produce 10,000 kg oxygen per year. It is found that the optimal design of an MRE-based ISRU system does not vary significantly with regolith type, demonstrating the technique's robustness to variations in regolith composition.

The mass and power of the optimized ISRU system exhibit an economy of scale, indicating that larger quantities of oxygen can be produced more efficiently. In fact, the production efficiency estimates of a lunar ISRU system provide initial evidence that lunar ISRU may prove beneficial in supporting a Mars Exploration campaign.

Thesis Supervisor: Jeffrey A. Hoffman

Title: Professor of the Practice of Aeronautics and Astronautics

Acknowledgments

As the end of this thesis-writing endeavor approaches, I would like to acknowledge those who have supported me along the way. First and foremost, I would like to thank my Lord and Savior Jesus Christ, the Creator of the Universe that I am so impassioned to explore. He has been the rock on which I have relied throughout graduate school here at MIT:

The Lord is my rock, my fortress, and my savior; my God is my rock, in whom I find protection.

- Psalm 18:2

Additionally, I would like to thank my wife Becca for her continual support and encouragement. She listened to my rants on the optimization of hypothetical reactors that produce oxygen on the moon with a patience that I do not fully understand, but for which I am eternally grateful. I also would like to thank my parents and family for instilling in me the value of hard work and supporting my academic endeavors, without their encouragement I would not be where I am today. I also thank my adviser, Jeff Hoffman, for giving me the freedom to select a research topic that I was passionate about and supporting my work wholeheartedly. My primary NASA collaborator, Jerry Sanders, also deserves recognition for supporting my research and introducing me to the field of ISRU. He opened up many doors for my summer site visits to NASA and introduced me to many people who enabled this research. I also would like to thank Laurent Sibille and Jesus Dominguez at Kennedy Space Center for taking the time to help mentor my work along the way. Without their expertise in Molten Regolith Electrolysis, this research would not have made it as far as it did. I thank Kris Lee for explaining the intricacies of the Hydrogen Reduction process to me. I would like to thank Rob Mueller, who made it possible for me to visit Kennedy Space Center and collaborate with the staff there. I would also like to thank Aislinn Sirk, Don Sadoway, and Bob Hyers for helping introduce me to the electrochemistry required to effectively model Molten Regolith Electrolysis reactors.

Ariane Chepko was instrumental in getting the ISRU systems analysis off the ground by donating her old file repository to me and answering my questions on her previous ISRU system modeling. I would also like to thank Cody Karcher for providing valuable feedback on the reactor design methodology section. I am grateful to Diane Linne for her feedback and expertise regarding ISRU system modeling. I would be remiss to not thank the administrative staff, including Jennie Leith and Liz Zotos, who cut through the red tape to let me focus on the research. Finally, this work was supported by a NASA Space Technology Research Fellowship (Grant #NNX13AL76H). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of NASA.

Contents

1	Introduction	21
1.1	Lunar ISRU Overview	21
1.1.1	Motivation: Why do ISRU?	21
1.1.2	The History of Lunar ISRU	23
1.1.3	Lunar ISRU Processes Overview	25
1.2	Molten Regolith Electrolysis	27
1.2.1	Process Overview	27
1.2.2	MRE Tradeoffs	29
1.2.3	Historical Development of MRE	31
1.2.4	The Need for MRE Reactor Modeling	33
1.2.5	Previous MRE Reactor Modeling	34
1.3	Integrated ISRU System Modeling	36
1.3.1	Motivation: Why model integrated ISRU hardware?	36
1.3.2	The history of ISRU system modeling	36
1.3.3	The necessity of integrated ISRU system modeling	38
1.4	Research Overview	39
1.4.1	MRE Reactor Parametric Model Development	39
1.4.2	Integrated ISRU System Model Development	40
2	MRE Reactor Model	43
2.1	MRE Model Overview	43
2.2	Regolith Material Properties	44
2.2.1	Composition	44

2.2.2	Density	45
2.2.3	Specific Heat	46
2.2.4	Thermal Conductivity	48
2.2.5	Electrical Conductivity	50
2.2.6	Optical Absorption Length	51
2.2.7	Current Efficiency	53
2.2.8	Gibbs Free Energy & Enthalpy of Formation	54
2.2.9	Latent Heat of Melting/Fusion	55
2.2.10	Miscellaneous Regolith Material Parameters	56
2.3	Regolith Throughput Requirements	56
2.4	Electrochemistry	59
2.4.1	Estimating Current	59
2.4.2	Estimating Voltage	60
2.4.3	Batch Profiles	61
2.4.4	Metal Production	62
2.5	Multiphysics Simulation	63
2.5.1	Simulation Motivation and Overview	63
2.5.2	Simulation Description	63
2.5.3	Constant Power Mode	68
2.6	Using the Simulation to Guide Reactor Design	68
2.6.1	Effects of Molten Mass and Operating Temperature: The Cut-off Line	69
2.6.2	Bounds on Reactor Diameter	73
2.6.3	Ensuring Operational Flexibility: The Design Margin	76
2.6.4	Designing a Feasible Range of Electrode Separations	77
2.6.5	Estimating Reactor Power and Operating Voltage	79
2.6.6	Estimating Reactor Mass	82
2.6.7	Reactor Model Flags	87
2.7	MRE Reactor Performance and Design Trends	90
2.7.1	Oxygen Extraction Efficiency and Current Efficiency	90

2.7.2	Mass and Power Estimates	92
2.7.3	Molten Metal Production	102
2.8	MRE Reactor Modeling Summary	103
3	Integrated ISRU System Optimization and Scaling	105
3.1	ISRU System Model Overview	105
3.2	System Model Description	106
3.2.1	Reactor	106
3.2.2	YSZ Separator	107
3.2.3	Excavator	109
3.2.4	Hopper and Feed System	110
3.2.5	Oxygen Liquefaction and Storage	112
3.2.6	Power System	112
3.3	Modeling the Integrated System	114
3.4	Optimization Technique	115
3.4.1	Hybrid Optimization Description	115
3.4.2	Bounding the Search Space	119
3.5	Optimized System Design	120
3.6	ISRU System Mass	125
3.7	ISRU System Power	127
3.8	Regolith Type Dependence	130
3.9	ISRU System Production Utility	135
3.10	ISRU System Visualization	137
3.11	ISRU System Modeling Summary	139
4	Conclusions	141
4.1	Drivers of ISRU System Mass and Power	141
4.2	Optimal MRE Reactor Design Characteristics	143
4.3	MRE Reactor Feedstock Sensitivity	144
4.4	ISRU System Mass and Power Scaling	144
4.5	The Utility of Lunar ISRU	145

4.6	MRE Reactor Future Work	146
4.6.1	Low Power Mode	149
4.7	ISRU System Future Work	150
A	Cutoff Line Design Justification	153
B	Derivation of Nonlinear Regression Equations	155
C	Regression Equation Coefficients	159

List of Figures

1-1	The silicate reduction reactor developed at Aerojet in 1963 (left) and 1965 (right), the earliest example of hardware development towards Lunar ISRU.	25
1-2	Schematic of a Molten Regolith Electrolysis Reactor that produces oxygen gas at the anode and molten metals at the cathode. The central molten core is insulated by solid “frozen” regolith around the reactor perimeter.	28
1-3	Three different anodes tested by Paramore [79], created by coating graphite geometries with iridium via electrodeposition.	32
1-4	(left) The multiphysics model of the MRE reactors tested at MIT, developed by Dominguez et al. [32] and (right) the multiphysics simulation of the joule-heated reactor design developed by Sibille and Dominguez [99].	35
2-1	A high-level schematic of the MRE reactor model, mapping user inputs to model outputs. The multiphysics regression model is used to generate realistic reactor design and performance estimates. Sections are referenced from the text.	44
2-2	The composition of lunar regolith by oxide type for three different regions: High-Titanium Mare (yellow), Low-Titanium Mare (cyan), and Highlands (red). Composition data from Apollo and Luna missions [111] and imagery data from Clementine UVVIS instrument [72].	45

2-3	The density of High-Ti Mare, Low-Ti Mare and Highlands molten lunar regolith as a function of temperature as calculated by the Stebbins density model [108], which includes both composition- and temperature-dependencies.	47
2-4	The specific heat model for lunar regolith. The data below 350K are from a model based on direct measurements of Apollo samples [51]. The data above 350K are based off of a composition-based model from Stebbins et al. [108].	48
2-5	The thermal conductivity data for solid lunar regolith simulant FJS-1 (left) and liquid silicates similar to lunar regolith (right) taken from Slag Atlas [38] and Snyder et al. [104].	49
2-6	Data of electrical conductivity for lunar regolith and similar materials [33]. A VTF fit [123, 113, 88] is overlaid for compositions similar to Highlands and Mare lunar regolith. The electrical conductivity exhibits a sharp increase around the melting temperature due to the increased ionic conductivity in molten regolith.	51
2-7	<i>(left)</i> Absorption length data from Sibille and Dominguez [99] fit with a step function (black solid line) overlaid with the Planck black body spectral energy density at two different temperatures (dashed lines). <i>(right)</i> The wavelength-dependent step function was integrated over the Planck curve to calculate the average absorption length as a function of temperature.	53
2-8	The Gibbs Free Energy (left) and Enthalpy of Formation (right) for each oxide specie in lunar regolith [21].	54
2-9	As electrolysis progresses, the composition and properties of the molten regolith vary. The liquidus temperature, calculated from Slag Atlas [38], and the Nernst decomposition potential for Low-Ti Mare lunar regolith initially increase as electrolysis progresses. As the operating temperature increases, the electrolysis can proceed farther to the right in the plot.	58

2-10	The liquidus temperature, calculated from Slag Atlas [38], for all three types of regolith throughout the electrolysis process. As the operating temperature increases, the electrolysis can proceed farther to the right in the plot.	58
2-11	The estimated voltage and current profiles over a single batch at an operating temperature of 1950K (left) and 2250K (right), along with the primary species being electrolyzed. The current is varied inversely with voltage to achieve constant power operation.	62
2-12	A side view of the cross-section of the multiphysics simulation of a cylindrical MRE reactor. The primary heat fluxes modeled are shown with different colored arrows, scaled by the same factor.	64
2-13	The predicted heat sinks due to the endothermic electrolysis reaction (left) and required power to heat fresh regolith feedstock (right) as a function of reactor operating temperature and current for High-Ti Mare.	66
2-14	The temperature (left) and voltage (right) profiles generated by the multiphysics simulation for a given reactor design. Current lines are shown in red and the phase boundary is shown in black.	67
2-15	The molten mass in an MRE reactor (left) and the operating temperature (right) depend on reactor diameter and electrode separation. Red X's indicate infeasible designs in which the side wall temperature gets too close to the melting temperature of regolith (1500K).	69
2-16	An illustration of how the current through the reactor affects the location of the cutoff lines between infeasible (red X's) and feasible reactor designs.	72
2-17	The molten mass (left) and operating temperature (right) for reactors on the "cutoff line" between infeasible and feasible designs. Each line represents a different wall thermal conductivity and a current of 500 A. Increasing the wall thermal conductivity decreases molten mass in the reactor and operating temperature.	73

2-18	The molten mass (left) and operating temperature (right) as a function of reactor diameter when designed on the cutoff line between infeasible and feasible designs. Each line represents a constant value for current and a wall thermal conductivity of 5 W/m-K. The minimum and maximum diameter bounds, resulting from the molten mass and operating temperature constraints respectively, are depicted.	74
2-19	The required electrode separation for a given reactor diameter. The electrode curve is affected by current (left) and wall thermal conductivity (right). A regression model, illustrated by the solid line, was fit to the multiphysics data, depicted by the data points connected by dashed lines ($R^2=0.975$, $RMSE=0.0032$ m).	78
2-20	The heat loss does not significantly depend on current (left), but is affected by wall thermal conductivity (right) when designing on the cutoff lines. A nonlinear regression model (solid lines) was fit to the data from a multiphysics simulation (dots with dashed lines) with $R^2=0.997$, $RMSE=0.51kW$	80
2-21	The oxygen extraction efficiency (left) and current efficiency (right) estimates for an MRE reactor.	91
2-22	The mass (top left), power (top right), specific mass (bottom left) and specific power (bottom right) of an MRE reactor over a range of oxygen production levels. Each line represents a different operating temperature and all designs have a margin of 1.5.	93
2-23	The mass and power of an MRE reactor at an operating temperature of 1850K, where each line shows a different design margin. Larger design margins clearly increase both mass and power, but enable higher production levels.	95
2-24	(left) Increasing the design margin opens up a larger range of feasible electrode separation values to enable operational flexibility. (right) Increasing the design margin requires less insulation and decreases the number of layers and mass of MLI on the reactor exterior.	96

2-25	The mass and power of an MRE reactor for different maximum wall temperatures. Increasing the maximum wall temperature definitively decreases reactor power (right-hand plot), but has a complex effect on reactor mass.	97
2-26	The height (left) and diameter (right) of an MRE reactor as a function of oxygen production level, for three different wall temperatures. . . .	98
2-27	The mass and power breakdown as a function of the wall temperature.	98
2-28	The mass and power of an MRE reactor for three different types of regolith with operating temperatures of 1850K (top), 2000K (middle), and 2300K (bottom).	100
2-29	The mass and power of an MRE reactor as a function of oxygen production and batch time. Longer batch times increase reactor mass and power.	101
2-30	The amount of metal produced by an MRE reactor operating on Mare (left) and Highlands (right) regolith. As operating temperature increases, more oxides can be reduced to produce more molten metal. . .	102
3-1	A diagram of the proposed YSZ separator used to purify oxygen from the exhaust gas from the Molten Regolith Electrolysis reactor.	107
3-2	A Computer-Aided Design (CAD) model of the hopper and auger used in the ISRU system model.	111
3-3	An N ² diagram of the ISRU system model within the optimization routine, showing how the subsystems are interconnected to generate a self-consistent estimate of system mass, which is then optimized. . . .	116
3-4	The mutation (red), crossover (blue) and elite (black) designs throughout multiple generations in the genetic algorithm optimization scheme.	117

3-5	A sample output from the genetic algorithm optimizer used on the ISRU system model, where the penalty value is the mass of the ISRU system (kg). The downwards trend in the blue data shows the effectiveness of the “natural selection” of better performing candidates from generation to generation.	118
3-6	(Top) The optimized ISRU system design variables over a range of production levels for Highlands regolith.	121
3-7	The mass breakdown of the ISRU system processing Highlands regolith, compared to other models from the literature.	125
3-8	The mass breakdown of a Highlands MRE reactor over a range of oxygen production levels. The refractory material contributes $\approx 50\%$ of the reactor mass.	127
3-9	The power predictions from the ISRU system model compared to four linear scaling laws from the literature.	128
3-10	The optimized mass and power of an MRE-based ISRU system for Highlands and High-Ti Mare regolith types.	131
3-11	The specific mass and power for each subsystem in the ISRU system for systems optimized for Highlands regolith (solid line) and Mare regolith (dashed line).	132
3-12	The optimal design characteristics for an MRE-based ISRU system for Highlands (solid lines) and Mare (dashed lines). Across most design variables, the two designs are reasonably similar.	134
3-13	The oxygen production level normalized by ISRU system mass (left) and ISRU system power (right) for both Highlands and High-Ti Mare regolith.	136
3-14	A CAD model of an ISRU system to produce 1,000 kg O ₂ per year., including the solar array power system (left), MRE reactor (front), regolith hopper and auger (front left), YSZ filter (front right) and oxygen storage system (right).	138

3-15	A CAD model of the ISRU system, including the solar array power system (left), MRE reactor (front), regolith hopper and auger (front left), YSZ separator (front right) and oxygen storage system (right). The 3x3 grid of solar arrays extends out of the image.	139
A-1	The heat loss (left) and heat loss divided by molten mass (right) for an MRE reactor. Infeasible designs, where molten material has touched the wall, are crossed out in red.	154
B-1	The radial temperature and phase indicator profiles from the multi-physics simulation of an MRE reactor with different wall thermal conductivities.	156
B-2	A linearized representation of the radial temperature profile.	157

List of Tables

2.1	The coefficients for the Stebbins density model (Equation (2.1)) applied to three types of lunar regolith.	46
2.2	The coefficients for the Stebbins heat capacity model (Equations (2.3) and (2.4)), in which oxide-specific coefficients are weighted by molar fraction. The summation represents a summation over all oxide species in lunar regolith.	48
2.3	The coefficients for the model of electrical conductivity of lunar regolith (Equation (2.6)), differentiated for Mare and Highlands regolith. . . .	51
2.4	The coefficients for the optical absorption length model (Equation (2.9)).	52
2.5	The modal mineralogical distributions for three types of lunar regolith [12], used to weight the latent heat of melting for each mineral [86] in order to calculate the latent heat of melting for each regolith type.	55
2.6	Experimental current density values from a range of molten regolith electrolysis experiments (with the exception of Kennedy [60]).	88
3.1	A review of various power system specific mass and area numbers for lunar surface systems in the literature.	113
3.2	The regression coefficients for the mass-specific and power-specific oxygen production performance of an MRE-based ISRU system.	137
C.1	The regression coefficients for Equation 2.20, which predicts the molten mass within the reactor based off of data from the multiphysics simulation presented in Section 2.6.	159

C.2	The regression coefficients for Equation 2.21, which predicts the operating temperature within the reactor based off of data from the multiphysics simulation presented in Section 2.6.	159
C.3	The regression coefficients for Equation 2.26, which predicts the required electrode separation distance to maintain thermal equilibrium in an MRE reactor based off of data from the multiphysics simulation presented in Section 2.6.	160
C.4	The regression coefficients for Equation 2.28, which predicts the expected heat loss from an MRE reactor based off of data from the multiphysics simulation presented in Section 2.6.	160
C.5	The regression coefficients for Equation 2.32, which predicts the expected operating voltage for an MRE reactor based off of data from the multiphysics simulation presented in Section 2.6.	160

Chapter 1

Introduction

1.1 Lunar ISRU Overview

1.1.1 Motivation: Why do ISRU?

One of the most significant barriers to space exploration is the burden of bringing all of the material resources from Earth required for a mission. The rocket equation [119] describes how a small increase in payload mass results in a dramatic increase in the total mass of the required launch system. This fundamental paradigm has limited space exploration in the decades since its birth. Today, typical launch costs are on the order of \$10,000/kg to low-earth orbit (LEO) [85].

A study by Eckart in 1996 [36] estimated the price to land hardware on the lunar surface to be \$75,000/kg to \$150,000/kg (2015 dollars), dramatically exceeding the cost of gold (\sim \$40,000/kg in 2015). A more recent study by the Colorado School of Mines (CSM) in 2005 [31] estimated lunar surface landing costs to be around \$110,000/kg (2015 dollars). To enable sustainable, affordable exploration of the solar system, the reliance on Earth’s resources must be reduced.

In-situ resource utilization (ISRU) is “*the collection, processing, storing and use of materials encountered in the course of human or robotic space exploration that replace materials that would otherwise be brought from Earth*” [89]. By producing resources outside of Earth’s gravity well, ISRU can provide

an avenue for reducing the launch mass from Earth. One form of ISRU is producing oxygen from lunar soil. Oxygen is a major component of launch vehicle, spacecraft, and lander masses – 80% of launch vehicle mass is fuel and oxygen [9], which translates to around 70% oxygen by weight. At the same time, oxygen is one of the most abundant lunar resources – lunar soil is around 44% oxygen by weight [9]. The production of this valuable resource outside of Earth’s gravity well can support lunar surface activities and enable orbital refueling to reduce mission mass and cost.

Studies have shown that oxygen can be produced via lunar ISRU at a lower cost than delivering it from Earth. In 1985, Michael Simon of the General Dynamics Space Systems Division conducted one of the first economic analyses of lunar oxygen production [102]. Assuming a 10-year amortization of capital costs along with operational costs, he determined that oxygen could be produced on the lunar surface and delivered to LEO at a cost of \$5,300/kg (2015 dollars), which is stated to be 1/3 the cost of delivering it using the Space Shuttle. A study by Eagle Engineering 3 years later [26] also included an amortization of hardware development costs and calculated a higher cost of \$8,095/kg oxygen delivered to LEO (2015 dollars). In 1993, Sherwood and Woodcock [98] included a spares analysis as well and calculated a cost of \$18,370/kg (2015 dollars) for producing oxygen on the lunar surface. The rising cost estimates over time can, to some degree, be attributed to the development of more detailed models that take more factors into account, but are also dependent on the assumptions made in each study. Sherwood’s model delved a level deeper by including hardware sizing models, a spare parts analysis, and hardware development costs of the ISRU system derived from a previous technical study by Woodcock et al. [127]. Sherwood also assumed a crew would be necessary on the lunar surface and included the crew and their support facilities in his cost estimates. The estimate by Sherwood and Woodcock [98] indicates that using lunar-derived oxygen in LEO would not be economically viable, at least in the near-term. Nevertheless, his cost estimate is still significantly lower than the price of \sim \$110,000/kg to launch oxygen to the lunar surface [31], suggesting that it is viable to use lunar-derived oxygen for lunar surface activities.

Lunar ISRU can significantly reduce the required launch mass and cost for certain missions. In 1993, the “LUNOX” (“LUNar OXYgen”) study sponsored by Johnson Space Center (JSC) [58] examined the impact of ISRU in the context of a lunar settlement. The study found that incorporating lunar oxygen production into a lunar settlement scheme reduced total program costs by 20% and launch vehicle costs by 50%. They concluded that *“emphasizing early production and utilization of lunar propellant has lower hardware development costs, lower cost uncertainties, and a reduction in human transportation costs of approximately fifty percent”*. A study by Duke in 2003 [34] found that lunar oxygen production with a propellant depot at the first Moon-Earth Lagrangian point could reduce the propellant delivered from Earth by 75% for lunar exploration missions similar to the Apollo program. A study out of the UK Space Agency in 2009 [120] found that between \$0.9 billion - \$3.8 billion could be saved annually, less the cost of oxygen production, if lunar oxygen was utilized to resupply four Altair lunar ascent flights each year.

1.1.2 The History of Lunar ISRU

The first recorded consideration of utilizing extraterrestrial resources was by Konstantin Tsiolkovsky in his science fiction works “On the Moon” and “Dreams about Earth and Sky” published in 1892 and 1895, respectively [118]. Arthur C. Clark wrote that *“The first lunar explorers will probably be mainly interested in the mineral resources of their new world, and upon these its future will largely depend”* [37].

The first appearances of utilizing lunar resources in the technical literature occurred as early as 1958 with K. Stehling’s work *“Moon refueling for interplanetary vehicles”* [110]. The earliest technological studies focused on extracting potential water from lunar soil and date to around 1962 with *“Water Extraction from Lunar Rock”* by R.W. Murray from the General Electric Missile and Space Division [76].

In February of 1963, Bruce B. Carr of the Callery Chemical Company examined extracting the potential water from lunar soil as well as chemically producing oxygen from lunar soil [19]. He proposed circulating hydrogen gas over lunar soil to produce water vapor by the reduction of iron oxides. Carr also pointed out that

direct electrolysis of the molten lunar soil could be used to extract larger amounts of oxygen. He created preliminary estimates of the mass and power of a particular oxygen extraction technique, called hydrogen reduction, and found that it would take 3-4 months to achieve mass payback (to produce the system's own mass in oxygen).

In March of 1963, a JPL working group published a set of recommendations for the utilization of lunar resources [57]. They identified the primary lunar resources as water, oxygen, hydrogen, raw soil, magnesium, iron, aluminum, nickel, and refractory materials. Their recommendations pointed to the need for a detailed systems analysis to assess the savings in cost, mass, and time associated with ISRU. The JPL group primarily investigated utilizing a solar furnace to extract water and, at significantly higher temperatures, directly dissociate the silicates in lunar soil.

In August of 1963, a group out of Aerojet General Corporation led by S.D. Rosenberg published their first quarterly report on a hardware project to facilitate the production of water from carbon monoxide and hydrogen - a critical step in reducing the silicates in lunar soil. **In December of 1965** they published experimental results from a reactor to reduce silicates with methane, demonstrating another critical step required to produce oxygen from lunar silicates. Figure 1-1 displays photographs of the hardware from their work in 1963 (left) and 1965 (right).

Significant work continued throughout the ensuing decades, with laboratory-scale development of several processes for extracting water and oxygen from lunar soil. **In 1988**, a study by Eagle Engineering [26] qualitatively ranked 13 oxygen extraction techniques based on technology readiness, number of processing steps, and process conditions. **In 1990**, a study by the Bechtel Engineering Group [3] conducted a more comprehensive evaluation of 16 oxygen extraction techniques based on feedstock material, oxygen production yield, the usability of byproducts, number of processing steps, operating temperature, required reagents, and estimates of the mass and power of the processing plant. **In 1992**, Taylor and Carrier III [115] evaluated 20 oxygen extraction techniques based on the same criteria as the Eagle Engineering study [26], with the addition of feedstock flexibility. These surveys all concluded that their rankings were preliminary because ISRU technology was at a low technology readiness




Figure 1-1: The silicate reduction reactor developed at Aerojet in 1963 (left) and 1965 (right), the earliest example of hardware development towards Lunar ISRU.

level (TRL) and the rankings would need to be updated as ISRU technology matured.

Research and development continued until the mid 2000's, at which point lunar ISRU received increased interest due to its role in the Vision for Space Exploration (VSE) [97]. This influx of funding supported phased ground development with a focus on maturing ISRU technology. Field tests were conducted in Mauna Kea, Hawaii from 2008 to 2012 that placed ISRU technology in a relevant operational environment [91, 92, 93]. These field tests demonstrated the *ROxygen* and *PILOT* that utilized the Carbothermal Reduction of Silicates process, as described in the following section. In parallel with the field tests, the Molten Regolith Electrolysis (MRE) process was developed at Massachusetts Institute of Technology (MIT) [103] in conjunction with Kennedy Space Center [100] and The Ohio State University [106], as well as at the Marshall Space Flight Center [30].

1.1.3 Lunar ISRU Processes Overview

Although over twenty different techniques to produce oxygen from lunar soil have been proposed in the literature [114], this work focuses on three processes in particular: Hydrogen Reduction of Ilmenite (HRI), Carbothermal Reduction of Silicates (CRS),

and Molten Regolith Electrolysis (MRE). These three techniques have undergone dramatic technology maturation in the past decade [92] and present perhaps the most likely processes to be implemented in the near future. A brief overview of each technique is given below.

Hydrogen Reduction of Ilmenite

The Hydrogen Reduction of Ilmenite (HRI) process heats regolith to around 900°C while exposed to hydrogen gas. The hydrogen reacts with the iron oxides in Ilmenite to form water, which can then be electrolyzed to form oxygen and recycle hydrogen [114]. This process can be expected to extract around 10% of the oxygen in lunar regolith in the equatorial Mare regions (4 kg of oxygen per 100 kg regolith) and 3% of the oxygen in lunar regolith in the highland regions (1.3 kg of oxygen per 100 kg regolith), though it presents some benefits in terms of a relatively low operating temperature which avoids the requirement of handling molten lunar regolith [92].

HRI reactors come in two primary flavors. The fluidized bed type uses a fast flow of hydrogen gas against the gravity vector to fluidize the lunar regolith to enhance reaction kinetics. The second reactor type uses a horizontal rotating bed to stir the regolith in the presence of hydrogen gas [3].

Carbothermal Reduction of Silicates

The Carbothermal Reduction of Silicates (CRS) process heats regolith past its melting point and exposes the molten regolith to methane gas. The methane reacts with silicates (and, with much slower kinetics, the ilmenite) in lunar regolith to produce carbon monoxide and hydrogen gas [114]. The carbon monoxide and hydrogen are then reacted over a nickel catalyst to produce water and reform methane. The water is then electrolyzed to produce oxygen and hydrogen gas, which is recycled to reform methane. This process can extract approximately 25-50% of the oxygen from lunar soil (10-20 kg of oxygen per 100 kg regolith), but at the added cost of process complexity, risk and a higher operating temperature [92].

CRS reactors typically involve a solar concentrator to direct high intensity solar

radiation towards a static bed of lunar regolith. Small pockets of regolith are melted using the concentrated radiation while methane gas flows overhead. After a given batch time, the solar concentrator is turned off, the bed solidifies, and the solidified, oxygen-depleted pockets are removed from the bed.

Molten Regolith Electrolysis

A third process, called Molten Regolith Electrolysis (MRE), also known as *Molten Oxide Electrolysis* or *Silicate Electrolysis*, has received significant technology maturation to date. This thesis places a special emphasis on MRE, therefore a more detailed description of the process is presented in the following section.

1.2 Molten Regolith Electrolysis

1.2.1 Process Overview

In the Molten Regolith Electrolysis (MRE) process, lunar regolith is fed into the reactor where it is heated to a molten state. Molten lunar regolith is conductive enough to sustain direct electrolysis [69], where two electrodes are immersed in the molten region and a voltage is applied. The applied voltage drives a current through the molten regolith, producing oxygen gas at the anode and molten metals and metalloids, such as iron, silicon, aluminum, and titanium, at the cathode [56].

Theoretically, this process can extract all of the oxygen from lunar regolith [103] (≈ 44 kg of oxygen per 100 kg regolith), but realistic operating conditions will likely limit oxygen extraction efficiency to lower values. Current efficiencies in excess of 95% can be expected, though iron-bearing regolith will decrease the expected efficiency [103].

Molten Regolith Containment: The Joule-Heated Cold Wall Solution

One design factor in an MRE reactor is the containment of molten regolith, which is extremely corrosive. The longest laboratory experiments lasted on the order of a

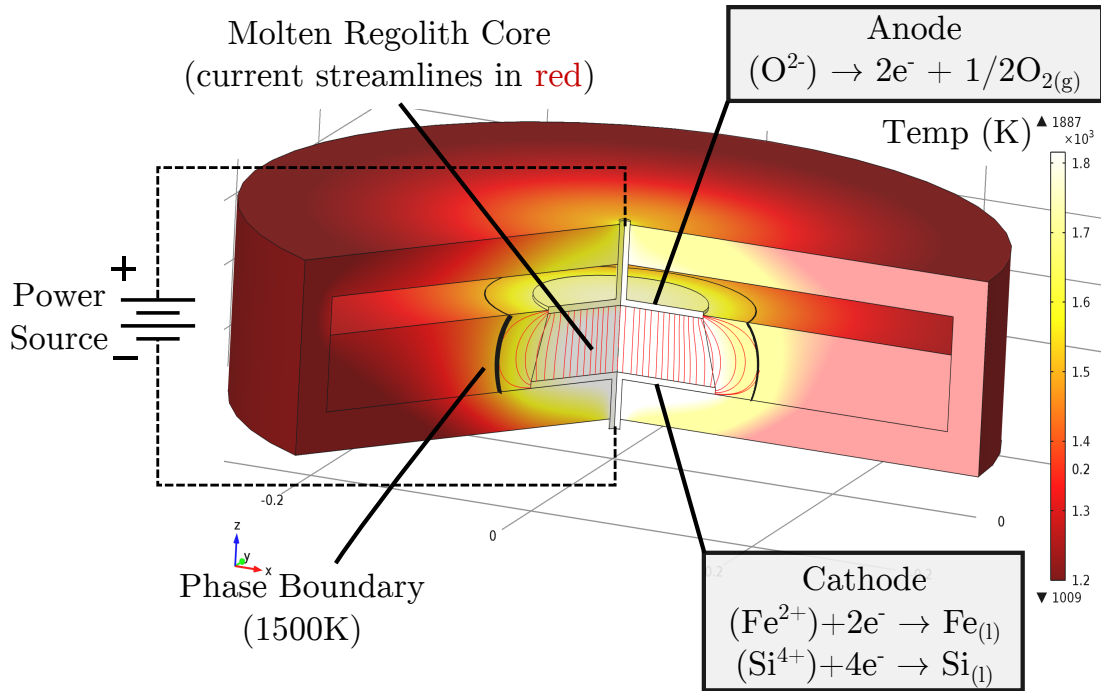


Figure 1-2: Schematic of a Molten Regolith Electrolysis Reactor that produces oxygen gas at the anode and molten metals at the cathode. The central molten core is insulated by solid “frozen” regolith around the reactor perimeter.

few hours [79] before the molten regolith eroded through the inner crucible, which is obviously not a solution to long-term oxygen production.

Paramore [79] notes that “*fortunately, this frustration [crucible failure] is purely an artifact of laboratory-scale experimentation.*” To solve the challenge of molten material containment, an MRE reactor can be designed similarly to industrial Hall-Heroult reactors [83] to support joule-heated, cold-wall (JHCW) operation. In JHCW operation, the the current traveling through the resistive melt generates heat via joule heating (I^2R) to maintain a central molten core, and the thermal gradients are designed such that the molten core is insulated by solid “frozen” regolith around the perimeter of the reactor, as shown in Figure 1-2. The solid regolith in contact with the reactor walls is not corrosive and enables long duration operation. The reactor shown in Figure 1-2 demonstrates the JHCW operation, with the central molten core surrounded by a phase boundary where the transition to the solid/glass phase occurs. The current streamlines are depicted in red and sample anodic and cathodic reactions are indicated.

1.2.2 MRE Tradeoffs

MRE Benefits

There is a strong impetus to explore the feasibility of an ISRU system with an MRE reactor, as there are many potential benefits to such a system. A study comparing various oxygen production methods (including water extraction from lunar polar craters) identified MRE as having the most favorable power consumption and the second most favorable mass throughput [73].

MRE does not require additional process materials, such as hydrogen, methane, or fluorine reagents. This means that an MRE reactor will not require additional systems to recycle the reagents, leading to a decrease in system mass and complexity. MRE can extract the majority of the oxygen from lunar regolith, which can dramatically decrease regolith throughput requirements. To provide some perspective, for every 100 kg of lunar regolith, Hydrogen Reduction can extract 1-5 kg of oxygen, Carbothermal Reduction can extract 10-20 kg of oxygen, and MRE can extract 16-40 kg of oxygen. The benefit of a higher oxygen extraction efficiency can have a profound impact in reducing total system mass and power requirements. Furthermore, the oxygen extraction ratio of MRE is relatively feedstock insensitive, meaning that one can expect an MRE reactor to extract similar amounts of oxygen at all locations on the lunar surface. Additionally, HRI and CRS reactors produce water from the primary chemical reaction, which must then be electrolyzed by a secondary system. MRE directly electrolyzes lunar regolith to produce oxygen and thus does not require the mass nor complexity of a secondary electrolysis system.

Perhaps the most attractive aspect of MRE is the fact that it produces a number of useful byproducts, including molten iron, silicon, aluminum, titanium, and glassy slag that can be readily cast after removal from the reactor. A number of studies determined that, with post-reactor processing, the byproducts of MRE could be used to produce infrastructure, spare parts and even solar arrays on the lunar surface [67, 66, 30, 9]. In fact, MRE is also being developed for environmentally-friendly metal production on Earth [43, 121]. Thus, developing MRE technology directly matures

technology that can be used on Earth to produce industrial metals with a zero-carbon footprint.

In light of the recent evidence in support of water in the polar lunar craters [27], there remain many potential benefits to using MRE on the lunar surface, perhaps even in parallel with a water extraction scheme. First, there is significant uncertainty as to the state and concentration of the water in lunar craters [27]. A resource prospecting mission is necessary to ascertain ground truth and is tentatively planned to launch in 2019 [5]. MRE may be concurrently developed using existing data from the Apollo lunar samples. Technical challenges associated with feedstock excavation from within permanently shadowed craters can also be avoided with the MRE process. Furthermore, previous ISRU studies have indicated that MRE may have an order of magnitude better efficiency, specific mass, and specific power than other production techniques, including a water extraction scheme [13].

MRE Design Challenges

Although there are certainly many benefits to using MRE on the lunar surface, a set of design challenges still remain. The problem of anode corrosion has been adequately addressed with a number of possible materials [2], but preventing the molten iron produced in the reactor from alloying with the cathode current collector (traditionally molybdenum) remains an open research problem. This problem may perhaps be overcome using a molten copper cathode concept proposed by Paramore [79].

For melts containing iron and titanium, the multivalent nature of these ions will likely create parasitic currents that reduce current efficiency. For certain oxides, such as sodium- and magnesium-oxide, the cations will form as a gas, rather than a liquid, at the cathode. These gases may bubble to the surface and recombine with the oxygen overhead to create a cyclic reaction that will also lead to current efficiency degradation. Previous experiments suggest that these effects may be mitigated with an anode oxygen collection tube [103], but further testing is required to determine how to best handle these species in lunar regolith.

Naturally, the high operating temperature (circa 1600°C) required by MRE poses

materials handling issues. Although many of these issues can be addressed by a joule-heated, cold-wall reactor design (see Section 1.2.1), hardware development is required to demonstrate this mode of operation on a lunar simulant. Finally, one of the most important but least tangible considerations is that MRE is at a lower technology readiness level (TRL) than Hydrogen Reduction and Carbothermal Reduction and thus requires relatively more technology development.

1.2.3 Historical Development of MRE

The process of electrolytically extracting metals from molten ores was first patented by Aiken [1] **in 1906** for the production of iron from raw ore. Carr [19] first suggested using a terrestrial electrolysis process to produce oxygen from lunar soil **in 1963**. The first experimental work on electrolytically reducing lunar soil simulants was conducted by Kesterke [61] **in 1970**. He extended industrial electrowinning practices to silica-bearing minerals and operated an electrolysis cell on terrestrial volcanic scoria, which is similar to lunar regolith. These experimental runs produced around 170g of oxygen as well as significant amounts of iron, titanium, aluminum, and magnesium using iridium as an anode and silica-carbide as a cathode. Over the next decade, electrolytic reduction of lunar regolith received considerable attention [56, 61, 69, 90].

In 1983, JPL published a report on space resources, including experimental efforts to electrolyze lunar simulants. They conducted experiments with current densities near 1.25 A/cm^2 for 1.4 hours and achieved current efficiencies in excess of 95%. This report mentioned the possibility of utilizing the thermal gradient in large-scale production reactors to effectively isolate the inner molten core from the containment crucible, which may be the first mention of the joule-heated, cold-wall reactor operation for lunar ISRU in the literature. In agreement with Kesterke [61], the report discussed producing not only iron, but also silicon and aluminum from lunar soil.

In 1992, Haskin et al. [45] studied the kinetics of the electrode reactions in an MRE reactor operating on feedstock with a composition similar to lunar regolith. They measured electrical conductivity of the lunar simulant to better understand power requirements of MRE reactors.

In 2006, Curreri et al. [30] at Marshall Space Flight Center carried out a 1 year hardware demonstration project for an MRE reactor. They tested a platinum 40% rhodium wire anode with three cathode materials: graphite, platinum 40% rhodium, and nickel-plated platinum rhodium. They demonstrated oxygen production from JSC-1 lunar simulant, though their apparatus required a fluxing agent to lower the melting temperature to 850°C. In the late 2000's, MRE research continued with significant efforts in hardware demonstrations and scale-up occurring at MIT [103, 121, 124, 2] in conjunction with Kennedy Space Center [99, 100, 101, 32] and The Ohio State University [107, 106].

The work at MIT addressed the problem of anode corrosion, which was one of the primary concerns of MRE reactors at that point in time [114]. After testing a group of platinum metals, iridium was found to be an acceptable anode material with a corrosion rate of less than 8 mm/year, which is well within the aluminum production industry standards [43, 124]. Reactor designs were scaled up from 0.1 A with 0.3 cm² electrodes to 10 A with 10 cm² electrodes [121, 103], making progress from laboratory-scale testing towards technology demonstration levels. Alternative anode materials were demonstrated, including iridium-plated graphite [79], 50-50 iridium/tungsten alloys [121], and iron-chromium alloys [2]. Novel anode geometries were investigated to increase current density limits [79], as shown in Figure 1-3. A counter-gravity molten metal removal system was developed at Ohio State University that leveraged the natural vacuum available on the lunar surface to remove molten metals from the reactor [106].



Figure 1-3: Three different anodes tested by Paramore [79], created by coating graphite geometries with iridium via electrodeposition.

1.2.4 The Need for MRE Reactor Modeling

Given the possible benefits of utilizing MRE on the lunar surface as described in Section 1.2.2, there is a large impetus to model such reactors. Altenberg [3] noted that two of the primary issues associated with MRE were **1)** the lack of specified optimal process, conditions, feed rate, and feedstock requirements and **2)** a poor understanding of the meaningful design parameters and oxygen extraction efficiency. These questions can be answered, in part, through extensive modeling of MRE reactors.

Indeed, models for HRI and CRS reactors have been extensively developed to better understand the design trades of such reactors. The HRI reactor model utilizes the shrinking-core physics formulation of fluid-particle chemical interactions to predict HRI reactor performance [48]. The chemical conversion rates, which include effects for particle size and reactor start-up, were validated with lunar simulant test data [47]. The CRS reactor model was also validated using lunar simulant test data [10]. The model can predict the conversion rate of a batch for a given temperature and batch time, which is a critical process parameter [11].

The CRS and HRI reactor models have been leveraged to better understand optimal reactor design. For both HRI and CRS reactor designs, the number of batches per day was evaluated to study its effect on reactor mass, process energy, and regolith throughput requirements [70]. The effects of combining multiple reactors in parallel with heat recuperation was studied using the HRI reactor model [71]. The tradeoff between shorter batches with faster kinetics and longer batches with more complete chemical conversions was studied to better understand optimal reactor batch time.

Although parametric models to predict the mass and power of reactors utilizing HRI and CRS have been developed, a similar model of suitable fidelity for an MRE reactor does not yet exist in the literature. This deficit has prevented quantitative comparisons between MRE and other processing techniques. In one of the most recent ISRU system studies, both HRI and CRS reactors were evaluated, but the MRE reactor model was not mature enough to be properly compared to the HRI and CRS models [23].

As discussed in Section 1.2.2, there are both potential benefits and drawbacks to utilizing MRE to produce oxygen from lunar regolith. To truly understand the tradeoffs associated with MRE reactors, these reactors must be quantitatively modeled. MRE modeling can answer important questions concerning the ideal operating conditions (batch time, operating temperature, etc.) and reactor geometry (diameter, electrode separation, etc.). Colson and Haskin [29] hypothesized that *“high melt resistivities coupled with the large distance between electrodes that would seem to be required to make the approach robust might make power requirements prohibitive [for an MRE reactor]”*. Teeple [116] surmised that *“the electrolysis techniques [including MRE], involve high temperatures, so one would expect high plant masses”*. Questions such as these, which attempt to understand the optimal design and performance of an MRE reactor, can be answered through parametric model analysis.

Furthermore, MRE modeling can guide hardware development. During the recent hardware development at MIT, the design of a JHCW reactor was avoided because *“at this stage [fabricating a JHCW reactor] was considered a superfluous enterprise, because the cell dimensions necessary to achieve sufficient joule heating would be extremely expensive to construct for an unproven process”*. Clearly, a better understanding of how to design JHCW reactors is needed in order to guide hardware development.

1.2.5 Previous MRE Reactor Modeling

There has been some amount of previous modeling work concerning MRE reactors. Dominguez et al. [32] created a multiphysics simulation of the laboratory MRE reactors at MIT using COMSOLTM, shown on the left in Figure 1-4. This simulation utilized material property data for lunar regolith including the density, electrical conductivity, and thermal conductivity as a function of temperature. The temperature and voltage profiles of the reactor were studied to better understand the thermoelectric topology inside an MRE reactor. The model included conduction, convection, and surface-to-surface radiation heat transfer, and the relative contribution of each heat transfer mode was studied.

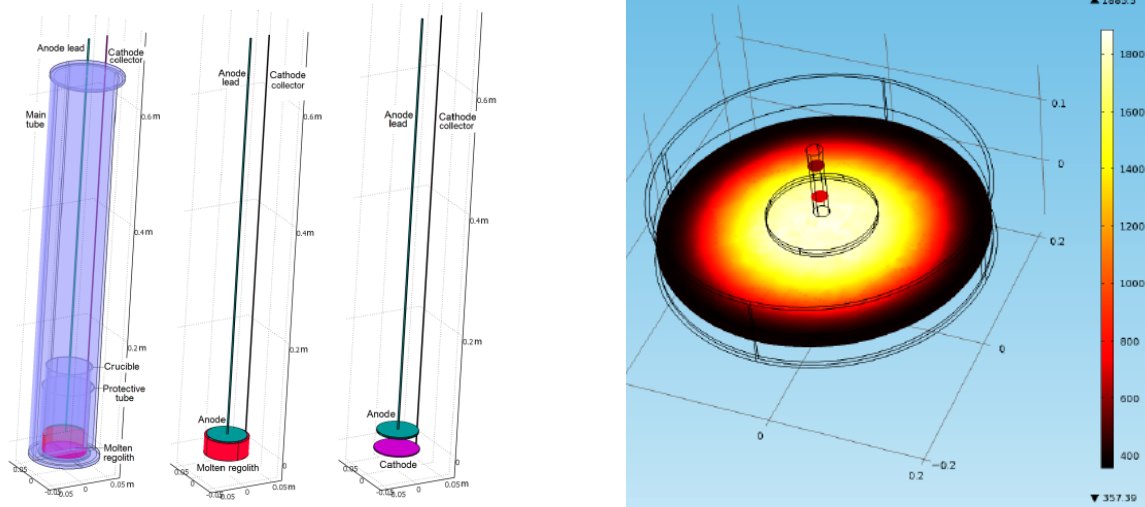


Figure 1-4: (left) The multiphysics model of the MRE reactors tested at MIT, developed by Dominguez et al. [32] and (right) the multiphysics simulation of the joule-heated reactor design developed by Sibille and Dominguez [99].

To assess the viability of a JHCW MRE reactor (see Section 1.2.1) with lunar regolith feedstock, Sibille and Dominguez [99] expanded the previous MRE reactor simulation. Their work demonstrated that it was indeed feasible to design a JHCW reactor for processing lunar regolith, shown on the right of Figure 1-4. Their simulation included an additional mode of heat transfer, called radiation in participating media. This mode of heat transfer accounts for the fact that molten lunar regolith is not completely opaque and therefore radiates, absorbs and retransmits energy when at high temperatures. This simulation effort also demonstrated that a “waffle” geometry anode can increase current capacity while also allowing for oxygen to readily escape from under the anode without disturbing the contact between the molten regolith and the anode surface.

Although previous simulation work demonstrated that a JHCW MRE reactor appears feasible [99], these simulations present a point design rather than a parametric model. That is, the simulation did not map oxygen production level to total reactor mass and power. Furthermore, the most recent simulations and experimental work have studied reactors that use tens of amps of current, while production-level reactors will need to use on the order of kilo-amps. The proper method for scaling up MRE reactor designs to these higher production levels has yet to be explored.

1.3 Integrated ISRU System Modeling

1.3.1 Motivation: Why model integrated ISRU hardware?

To truly understand the benefits and drawbacks of lunar ISRU, the holistic ISRU system design and performance must be quantitatively modeled. Simon’s 1985 parametric analysis of lunar oxygen production showed that, after the Earth-to-Moon transportation cost, the power required for ISRU had the biggest impact factor in the economic feasibility of lunar ISRU [102].

Sherwood and Woodcock [98] conducted an economic analysis of lunar oxygen production and noted that, “*the sensitivities [of their economic model] are modest, except for the mass of production hardware.*” They found that when the ISRU system mass was varied by a factor of two, the total cost of producing oxygen on the lunar surface varied from \$12,570 to \$29,857/kg (about a nominal value of \$18,370/kg) in 2015 dollars. **Thus, it is imperative to accurately model the mass, power and performance of the entire ISRU system to determine the utility of such systems.**

1.3.2 The history of ISRU system modeling

Lunar ISRU Hardware Modeling

In 1982, the first model of the life cycle of a lunar mission, called LUBSIM, was created [63]. LUBSIM utilized over 200 non-linear scaling equations with 600 constants and variables to simulate the mass flow, power, and manpower requirements of a lunar base.

In 1988, a technical report by Eagle Engineering Inc. [26] created detailed hardware designs for an HRI system (see Section 1.1.3). This design included an excavator, hauler, magnetic feedstock beneficiation, low-pressure and high-pressure hoppers, a hydrogen recycling system, water electrolysis cell, hydrogen makeup tank, oxygen liquefaction and storage, a thermal control system, and many other subsystems. Analytical component scaling models were integrated to predict the mass and power of

an HRI-based ISRU system. This report demonstrated the value of parametric sizing models by assessing the effect of feedstock ilmenite concentration, lunar base location, and other parameters on the mass, power, and performance of the ISRU system.

In 1990, Woodcock et al. [127] produced a detailed conceptual design of an entire lunar outpost dedicated to liquid oxygen production. Their design included many of the components included in the Eagle Engineering design [26] and expanded the scope to consider robotic maintenance, excavation, and base construction. The required launch manifest landed a total of 388 mT on the lunar surface to enable an oxygen production level of 100 mT/year within 3.75 years of the first landing.

In 1994, Hepp et al. [52] conducted a survey of the production of metals from lunar regolith. They compiled mass and power estimates for 11 chemical processing techniques, including HRI, CRS, and MRE.

In 1996, Eckart [36, 35] developed a steady-state model of lunar base systems that represented a simplification of the complex LUBSIM model developed by Koelle and Johenning [63]. This model included an in-situ oxygen production model to predict soil feed requirements, system mass, resupply mass, and power requirements, as well as a power system model. Eckart's model integrated the ISRU system with a full lunar base model to study the benefits and costs of manufacturing lunar products [37].

In 2007, Steffen et al. [109] began a new generation of ISRU system modeling that took a bottom-up approach. Their work developed analytical component sizing models and integrated the components together to study system mass and power. They modeled a set of HRI reactors that interfaced with a Knudsen flow hydrogen separator, a compressor, a solid oxide electrolyzer (to split the product water into oxygen and hydrogen), an oxygen liquefaction and storage system, and a variety of options for a fission power system. They quantitatively demonstrated that using two reactors operating out of phase in parallel would provide a more even power profile.

In 2008, Chepko et al. [24, 25] incorporated updated reactor models [10, 48] into a partially-integrated ISRU system model. Their model included a reactor, regolith storage hopper, and auger to insert regolith into the reactor, but did not include power nor excavation systems, two critical elements in designing optimal ISRU systems.

Their work demonstrated some of the first optimization of lunar ISRU systems, in which they leveraged a genetic-algorithm optimization scheme to locate the design variables and technology choices that led to the minimal system mass. This modeling demonstrated the power of optimizing analytical sizing models, in that it identified optimal process conditions and hardware configuration options.

In 2009, Linne et al. [70] utilized detailed reactor models for HRI and CRS to study the optimal number of batches per day. This work demonstrated the effects of feedstock parameters, such as particle size, on the design and operation of such reactors. Multiple reactors with heat recuperation were determined to be an effective avenue for reducing energy requirements by 20-40% [71].

Economic/Financial/Political Modeling

A number of studies have leveraged existing ISRU engineering models, either in the form of analytical scaling equations but more often as simple linear scaling laws, to assess the economic feasibility of lunar oxygen production [102, 98, 116, 14].

Some modeling approaches have even included political and legal aspects alongside engineering and economic models, such as the FERTILE MOON [13] and FULL MOON [7] ISRU studies. These two studies assessed the economic feasibility of producing hydrogen, oxygen, and water on the lunar surface, and considered five different ISRU methods, including HRI, CRS, and MRE.

1.3.3 The necessity of integrated ISRU system modeling

When modeling ISRU systems, it is essential to model the integrated system, including elements such as the reactor, power system, oxygen storage system, excavator, etc. A study by Linne et al. [70] examined the design and operation characteristics of an HRI reactor. They noted that operating at the optimal oxygen extraction energy (MJ/kg O₂) increased the regolith throughput from a nominal 200 kg per day to 333 kg per day, which will increase the mass and power of the excavation system. With an integrated ISRU system model, the tradeoff between

optimal reactor performance and optimal excavator design can be balanced to minimize total system mass. Linne et al. [70] also found that processing larger quantities of regolith in fewer batches each day reduced reactor power and regolith throughput requirements, but increased reactor mass. This tradeoff between mass and power appears often in reactor modeling and can be quantitatively studied using a model that includes a power system.

Hegde et al. [47] studied optimal HRI reactor designs and noted that the “*reactor must interface with the other sub-system processes such as upstream regolith extraction and beneficiation and downstream electrolysis and phase separation in a way that establishes the most favorable balance between efficiency, robustness, and equivalent system mass.*” Integrated ISRU system models can address the optimal design of these complex, coupled systems. One such tradeoff is the decision of the height-to-diameter (H/D) ratio of the hydrogen reduction reactor. Linne [71] noted that a lower H/D ratio was more efficient, but for a fluidized bed reactor this increases the required hydrogen flow rate to maintain fluidization. An integrated ISRU system model that incorporates the gas cleanup, condenser, compressors, and other subsystems could be used to quantify this tradeoff.

1.4 Research Overview

There are two primary objectives achieved via this research:

1.4.1 MRE Reactor Parametric Model Development

First, a parametric sizing model for an MRE reactor is developed.

This model has a foundation of lunar regolith material properties with both composition- and temperature-dependence, including thermal and electrical conductivity, density, etc. These data sets are integrated into a multiphysics simulation, created using COMSOL™, that simulates the electrical, chemical, and thermodynamic behavior of reactor designs. The multiphysics simulation is leveraged to create an extensive tradespace of reactor designs with different values for the diameter,

electrode separation, and wall thermal conductivity. A novel design methodology is implemented that determines the required reactor design (diameter, electrode separation, and wall thermal properties) that **1)** sustains the amount of molten mass and average current required to meet a given oxygen production level, **2)** maintains a given operating temperature within the molten core, and **3)** ensures that the reactor walls are insulated from the molten core by a layer of solid lunar regolith in the joule-heated, cold-wall concept (discussed in Section 1.2.2).

The sizing model presented in this work parametrically generates a reactor design and performance estimates for a given set of model inputs, including oxygen production level, operating temperature, and regolith feedstock type. This model can be used to **1)** guide MRE reactor design development and **2)** quantitatively compare it to other oxygen production techniques. This research objective seeks to address the following knowledge gaps:

1. What are the optimal MRE reactor design characteristics in terms of operating temperature, batch time, voltage, current, electrode separation, diameter, etc.?
2. What reactor geometry is required to maintain the joule-heated cold-wall effect?
3. How does the design of an MRE reactor scale with production level?
4. How is operational flexibility designed into an MRE reactor?

1.4.2 Integrated ISRU System Model Development

Second, an integrated ISRU system model is developed to study the optimal design variables that minimize the mass and power of an MRE-based ISRU system over a range of production levels.

Although the previous studies identified in the Section 1.3.2 evaluated the impact of ISRU systems on transportation, financial, and even political aspects of lunar exploration, these studies relied on either simple linear scaling laws or hardware scaling models that are several decades old for the ISRU system. With the new generation of ISRU component models available [11, 24, 48, 70], the time for an updated integrated ISRU system study is ripe.

The system model presented in this work expands upon previous work [23] to encapsulate a more complete system by including models of an MRE reactor, power system, excavation system, oxygen storage and liquefaction system, as well as a hopper and regolith feed system. By evaluating the integrated ISRU system, the holistic system performance may be studied and optimized, rather than just a subset of the entire system. A hybrid genetic algorithm/gradient-based optimization routine is developed and utilized to minimize the ISRU system mass over a range of oxygen production levels.

These optimized mass and power estimates can be leveraged in hierarchical models of a lunar base, transportation logistics, economic markets, and other aspects to better understand the impact and applicability of lunar ISRU. This objective attempts to address the following questions:

1. How do the mass and power of the holistic system grow with production level?
2. What are the design characteristics of the optimized system?
3. What is the optimal production level for a single reactor/how does the number of reactors scale with production level?
4. When does the ISRU system achieve mass pay-back?

The first objective, creating a parametric model for an MRE reactor, is described in Chapter 2. The second objective, optimizing the ISRU system model, is presented in Chapter 3. The conclusions are presented in Chapter 4.

Chapter 2

MRE Reactor Model

2.1 MRE Model Overview

This chapter presents the development of a parametric sizing model of MRE reactor. This model is used to generate an MRE reactor design, with mass, power, and performance estimates, for a given set of inputs, such as oxygen production level, operating temperature, batch time, etc.

The MRE reactor model is built on a foundation of composition- and temperature-dependent lunar regolith material property models presented in Section 2.2, which are validated using data from Apollo samples and regolith simulants. As described in Section 2.3, the reactor model calculates the regolith throughput required to meet the desired oxygen production level. The reactor model integrates electrochemical principles, presented in Section 2.4, with a multiphysics simulation, presented in Section 2.5, to model the performance of reactor designs. As described in Section 2.6, data generated by the multiphysics simulation is leveraged to create a novel reactor design methodology. This design methodology is built into a parametric sizing model to generate a reactor design that **1)** sustains the amount of molten mass and current required to meet a given oxygen production level, **2)** maintains a given operating temperature within the molten core, and **3)** ensures that the reactor walls are insulated from molten regolith by a layer of solid lunar regolith in the joule-heated, cold-wall (JHCW) concept described in Section 1.2.1.

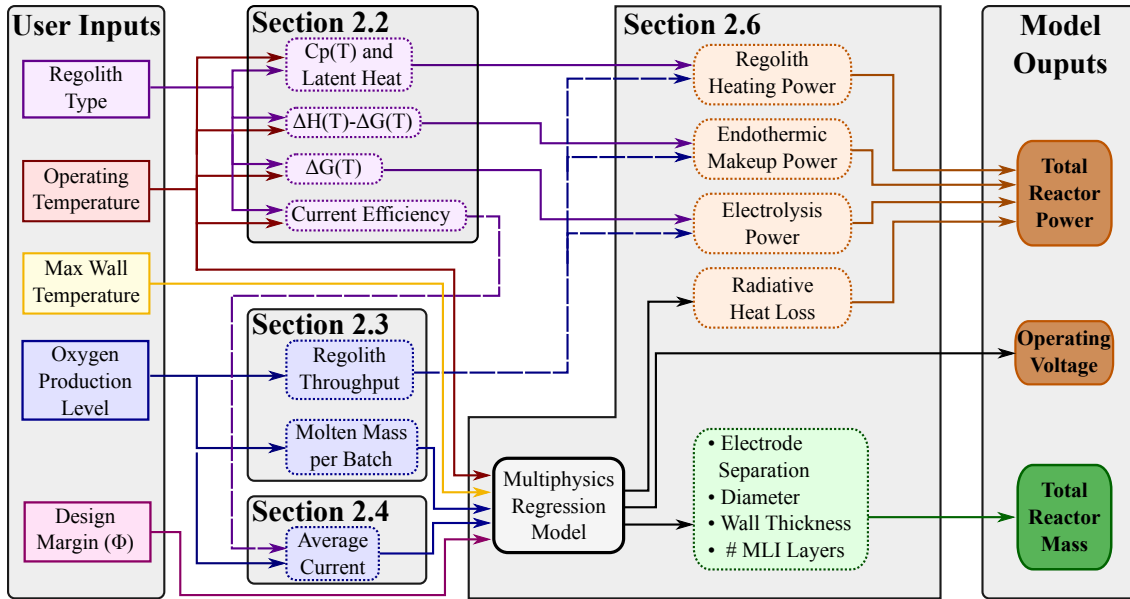


Figure 2-1: A high-level schematic of the MRE reactor model, mapping user inputs to model outputs. The multiphysics regression model is used to generate realistic reactor design and performance estimates. Sections are referenced from the text.

The reactor sizing model presented in this chapter parametrically generates reactor design and performance estimates that can be used to guide MRE reactor development and quantitatively compare it to other oxygen production techniques. A top-level schematic of the MRE reactor model data flow is shown in Figure 2-1, in which the user inputs are mapped to model outputs.

2.2 Regolith Material Properties

2.2.1 Composition

Lunar regolith is comprised of minerals such as plagioclase, pyroxene, olivine, and ilmenite, and the concentration of each depends heavily upon location on the lunar surface [49, 26]. Each of these minerals are composed of various oxide species, including FeO, SiO₂, Al₂O₃, etc. A modal analysis of the mineralogical distributions observed at the regolith at the Apollo and Luna landing sites grouped lunar regolith into three broad categories: Low-Titanium Mare, High-Titanium Mare, and Highlands [12, 111]. Figure 2-2 shows a map of the lunar surface with the lunar Highlands

(red), High-Ti Mare (yellow) and Low-Ti Mare (cyan) from the Clementine UVVIS instrument [72]. The regolith properties given below are often differentiated based upon these three categories.

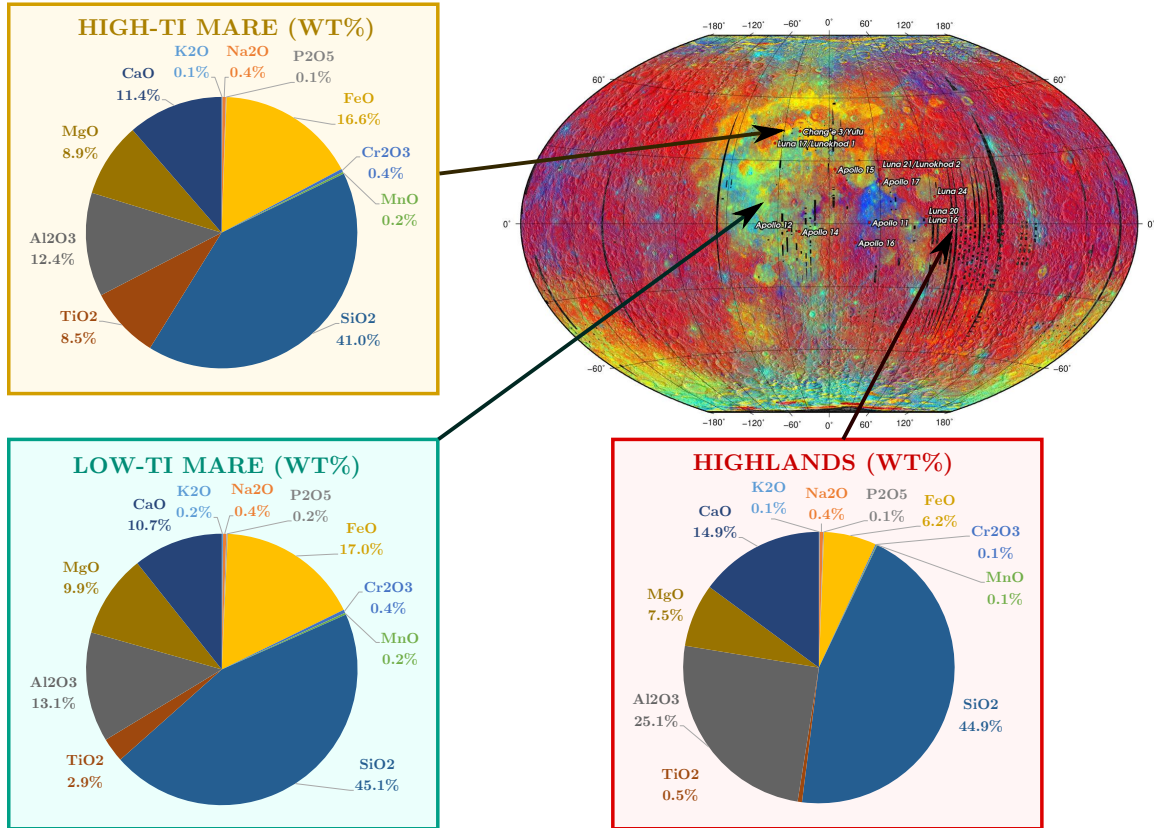


Figure 2-2: The composition of lunar regolith by oxide type for three different regions: High-Titanium Mare (yellow), Low-Titanium Mare (cyan), and Highlands (red). Composition data from Apollo and Luna missions [111] and imagery data from Clementine UVVIS instrument [72].

2.2.2 Density

The density of lunar regolith is used in the multiphysics simulation to map between the volume and mass of regolith in the reactor. The density of solid lunar regolith was taken to be 1500 kg/m³ [49]. For molten lunar regolith, the density varies with both temperature and composition. The Stebbins density model [108], shown in Equation (2.1), was used to model the density of molten lunar regolith because it could be easily adapted to the composition of lunar regolith and was valid in temperature

range of interest (1500K-2500K).

$$\rho = \frac{1000 r_1}{r_2 + r_3(T - 1873)}, \quad (2.1)$$

where ρ is the density (kg/m³), T is the temperature (K), and r_i are regression coefficients, shown in Table 2.1. The regression coefficients were calculated by weighting the oxide-specific coefficients from Stebbins et al. [108] by the oxide molar fraction for each type of regolith.

Table 2.1: The coefficients for the Stebbins density model (Equation (2.1)) applied to three types of lunar regolith.

	r_1	r_2	r_3
Highlands	63.45	24.11	0.001206
High-Ti Mare	63.33	22.48	0.001982
Low-Ti Mare	63.84	23.01	0.001612

The predictions from the density model are shown in Figure 2-3 for the range of temperatures relevant to an MRE reactor. The values predicted by the model compare well with the range of 2200 to 3200 kg/m³ given in the Lunar Sourcebook [49] and the range of 2490 to 2890 kg/m³ measured for molten JSC-1A lunar simulant [59].

2.2.3 Specific Heat

The specific heat of lunar regolith is used to estimate the energy required to heat lunar regolith to the operating temperature. For temperatures below 350K, a fourth-order polynomial fit of Apollo lunar sample data from Hemingway [51] was used:

$$C_p = A + BT + CT^2 + DT^3 + ET^4, \quad (2.2)$$

where C_p is the specific heat, T is the temperature in K, A , B , C , D , and E are regression coefficients from Hemingway et al. [51]. For temperatures above 350K, a model for silicate glasses developed by Stebbins [108] was employed to predict the the




Figure 2-3: The density of High-Ti Mare, Low-Ti Mare and Highlands molten lunar regolith as a function of temperature as calculated by the Stebbins density model [108], which includes both composition- and temperature-dependencies.

heat capacity of solid regolith:

$$C_p = \sum_i \chi_i a_i + T \sum_i \chi_i b_i + T^{-2} \sum_i \chi_i c_i, \quad (2.3)$$

where χ_i is the mole fraction of oxide specie i in lunar regolith, and a_i , b_i , and c_i are coefficients from Stebbins et al. [108]. Table 2.2 shows the resultant weighted coefficients, calculated by weighting the oxide-specific coefficients by molar fraction.

For molten regolith, the specific heat remains relatively constant with temperature but varies with composition, as predicted by the Stebbins specific heat model [108]:

$$C_p = \sum_i \chi_i \bar{C}_{p,i}, \quad (2.4)$$

where $\bar{C}_{p,i}$ is the coefficient of each oxide specie i in lunar regolith, formed from a linear regression of a database of 58 different compositions [108]. The resultant weighted coefficients are shown in the right-hand column in Table 2.2.

The combined specific heat models from Equations (2.2), (2.3), and (2.4) are shown

Table 2.2: The coefficients for the Stebbins heat capacity model (Equations (2.3) and (2.4)), in which oxide-specific coefficients are weighted by molar fraction. The summation represents a summation over all oxide species in lunar regolith.

	$\sum_i \chi_i a_i$	$\sum_i \chi_i b_i$	$\sum_i \chi_i c_i$	$\sum_i \chi_i \bar{C}_{p,i}$
High-Ti Mare	8.820E+02	3.083E-01	-2.278E+07	1531
Low-Ti Mare	9.093E+02	2.870E-01	-2.469E+07	1539
Highlands	9.530E+02	2.524E-01	-2.645E+07	1565

in Figure 2-4. The model from Stebbins [108] matches up well with the Apollo data fit at 350K. The data in Figure 2-4 is also in the same range at the model proposed by Colozza [28].

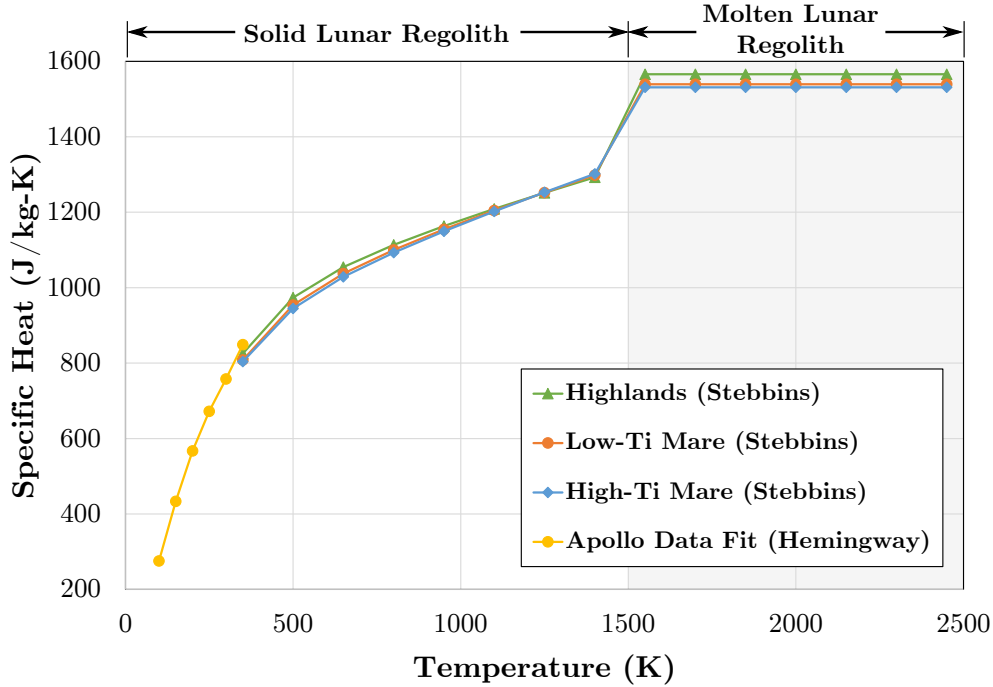


Figure 2-4: The specific heat model for lunar regolith. The data below 350K are from a model based on direct measurements of Apollo samples [51]. The data above 350K are based off of a composition-based model from Stebbins et al. [108].

2.2.4 Thermal Conductivity

The thermal conductivity of lunar regolith drives the thermal topology of an MRE reactor. For solid lunar regolith, thermal conductivity data was compiled for the FJS-1 and synthetic Apollo 11 lunar simulants [99]. Thermal conductivity data was

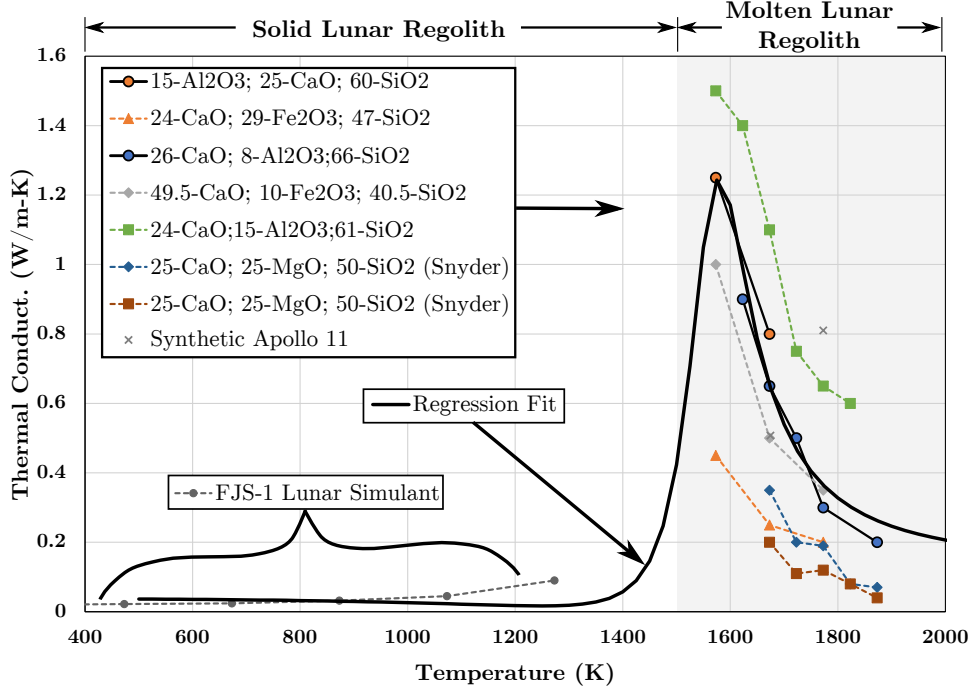


Figure 2-5: The thermal conductivity data for solid lunar regolith simulant FJS-1 (left) and liquid silicates similar to lunar regolith (right) taken from Slag Atlas [38] and Snyder et al. [104].

unavailable for molten lunar regolith, so similar minerals were selected from the Slag Atlas [38] and Snyder et al. [104], as shown on the right side of Figure 2-5.

The thermal conductivity rises around the melting temperature of 1500K and then decreases asymptotically at higher temperatures. Initially, separate fits were made for the solid and liquid data, but the discontinuity dramatically impeded convergence in the multiphysics solver. A continuous smooth fit (the solid black line in Figure 2-5), shown in Equation 2.5, was implemented over both the solid and molten data, which reduced the multiphysics convergence time by two orders of magnitude.

$$k = \frac{0.008865(T')^3 + 0.07197(T')^2 - 0.009113(T') + 0.005266}{(T')^2 - 1.053(T') + 0.2955}, \quad (2.5)$$

where T' is the temperature in K normalized by a mean 1289 and standard deviation 520.5. The fit coefficients were generated using the MatlabTM function `cftool` with bisquare weights for robustness. The thermal conductivity model in Equation (2.5) was used for all three types of lunar regolith in the MRE reactor model.

2.2.5 Electrical Conductivity

The electrical conductivity of lunar regolith affects both the operating voltage and reactor geometry required for an electrolysis cell [94], as it is inversely proportional to the resistive heating in the melt. For solid lunar regolith, electrical conductivity data from Olhoeft et al. [78] for Apollo samples 12002,85 and 65015,6 was used. For molten lunar regolith, data from Ducret et al. [33] was utilized with a composition of 25% CaO, 25% MgO, and 50% SiO₂ with varying amounts of FeO added. The mixture with 5 wt% FeO added was used to approximate both Mare types of regolith and data for the melt with 15 wt% FeO added was used for the Highlands, to match the compositions shown in Figure 2-2.

Intuitively, solid lunar regolith has a much lower electrical conductivity than the molten regolith. The dramatic increase in electrical conductivity around the melting temperature is due to the rise in ionic mobility in the molten material, which is a primary driver of electrical conductivity in molten oxides [94]. As the temperature increases, the molten regolith becomes less viscous and ionic mobility increases [29].

Figure 2-6 shows the electrical conductivity data [78, 33] with the model overlaid. Clearly the electrical conductivity increases with FeO concentration. In the context of lunar regolith, this means that one can expect a higher electrical conductivity in the Mare regions compared to the Highlands regolith. As with thermal conductivity, the electrical conductivity had to be fit with a continuous function for both solid and molten regolith to enable robust, rapid multiphysics solving. A Vogel-Tamman-Fulcher (VTF) fit [123, 113, 88] was implemented on the data, as it has been shown to match the temperature dependence of the electrical conductivity of ionic liquids quite well [122]. The VTF fit is given as follows:

$$\sigma_e = \frac{e_A}{T} \exp\left(\frac{-e_B}{T - T_0}\right), \quad (2.6)$$

where σ_e is the electrical conductivity in S/m, T is the temperature in K, and e_A and e_B are regression coefficients. These coefficients are presented in Table 2.3 and were generated using the MatlabTM function `cftool` with bisquare weights for robustness.

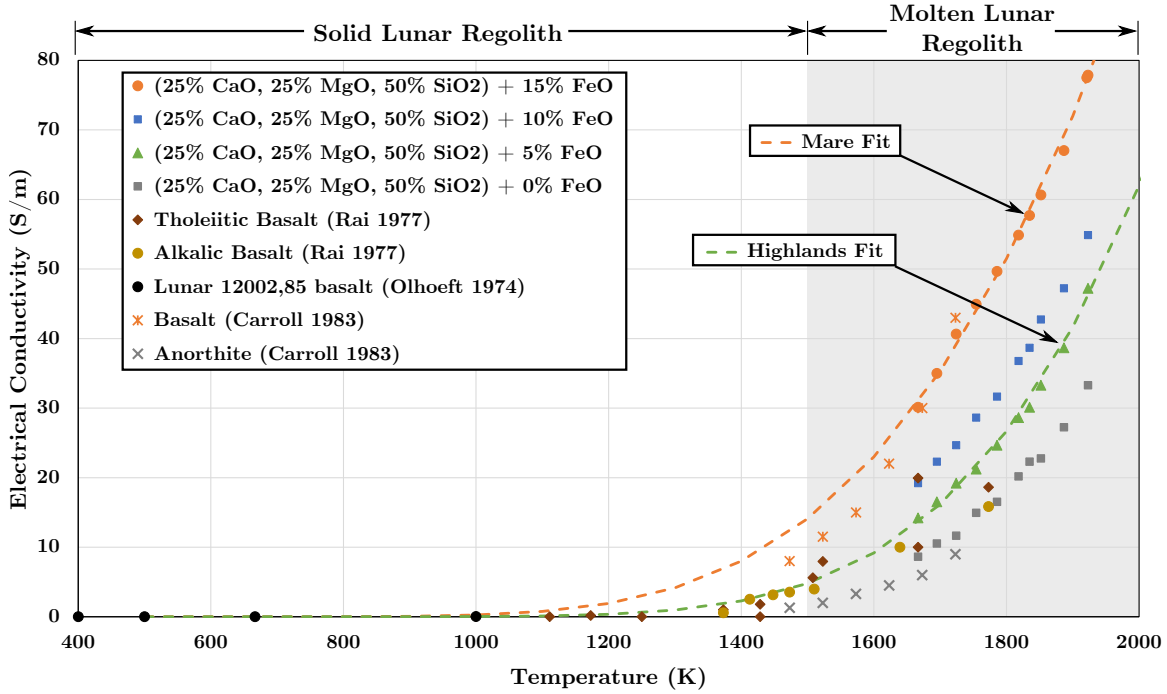


Figure 2-6: Data of electrical conductivity for lunar regolith and similar materials [33]. A VTF fit [123, 113, 88] is overlaid for compositions similar to Highlands and Mare lunar regolith. The electrical conductivity exhibits a sharp increase around the melting temperature due to the increased ionic conductivity in molten regolith.

Note that T_0 had to be set to zero to avoid asymptotes in the lower temperature regime. As observed in Figure 2-6, the model compares reasonably well with data for Tholeiitic and Alkali basalts from Rai [84], data on basalt and anorthite from Carroll [20], and data from Schiefelbein and Sadoway [94].

Table 2.3: The coefficients for the model of electrical conductivity of lunar regolith (Equation (2.6)), differentiated for Mare and Highlands regolith.

	Mare	Highlands
e_A	1.508E+8	6.341E+8
e_B	1.331E+4	1.708E+4
T_0	0	0

2.2.6 Optical Absorption Length

The optical absorption length of lunar regolith plays a large role in determining how radiation travels in the molten regolith core before being absorbed. Wavelength-

dependent data on molten glasses from Sibille and Dominguez [99] was first fit with a step function. This step function was then weighted by the energy density predicted by the Planck temperature-dependent black body radiation law and integrated over all wavelengths to determine the average absorption length at a given temperature:

$$\alpha(T) = \frac{\int_0^\infty P(\lambda, T)\alpha(\lambda)d\lambda}{\int_0^\infty P(\lambda, T)d\lambda}, \quad (2.7)$$

where $\alpha(\lambda)$ is the absorption coefficient as a function of wavelength, λ , and $P(\lambda)$ is the spectral energy density as a function of wavelength from Planck's Law [82]:

$$P(\lambda, T) = \frac{8\pi hc^2}{\lambda^5} \left(\frac{1}{\exp(\frac{hc}{k_B T \lambda}) - 1} \right), \quad (2.8)$$

where h is Planck's constant, c is the speed of light, and k_B is Boltzmann's constant.

The left plot in Figure 2-7 shows the step function for absorption length derived from Sibille and Dominguez [99] (black solid line) overlaid with the Planck black body spectral energy density curve at two different temperatures, 700K (blue) and 3000K (red). The right plot in Figure 2-7 shows the resultant absorption length as a function of temperature ($\alpha(T)$ from Equation (2.7)) with a regression fit using a summation of five Gaussian functions:

$$\alpha(T) \approx a_1 e^{-\left(\frac{T-b_1}{c_1}\right)^2} + \dots + a_5 e^{-\left(\frac{T-b_5}{c_5}\right)^2} \quad (2.9)$$

where a_i , b_i , and c_i are regression coefficients shown in Table 2.4, calculated using the MatlabTM function `cftool` with bisquare weights for robustness (Adjusted R² = 0.9997). These coefficients were used for all three types of lunar regolith in the MRE reactor model.

Table 2.4: The coefficients for the optical absorption length model (Equation (2.9)).

	Gauss 1	Gauss 2	Gauss 3	Gauss 4	Gauss 5
a_i	465	-0.650	163	0.325	1.45
b_i	-3670	200	528	175	150
c_i	5040	4.69	905	4.40	3.54

The trends in the right-hand plot in Figure 2-7 reveal that the absorption length decreases at higher temperatures. In the range of operation for an MRE reactor, radiation in the molten region will travel on the order of 7 mm before being reabsorbed. The model presented here is based on absorption length data from molten glasses and the uncertainty concerning the absorption length of actual lunar regolith warrants further examination.

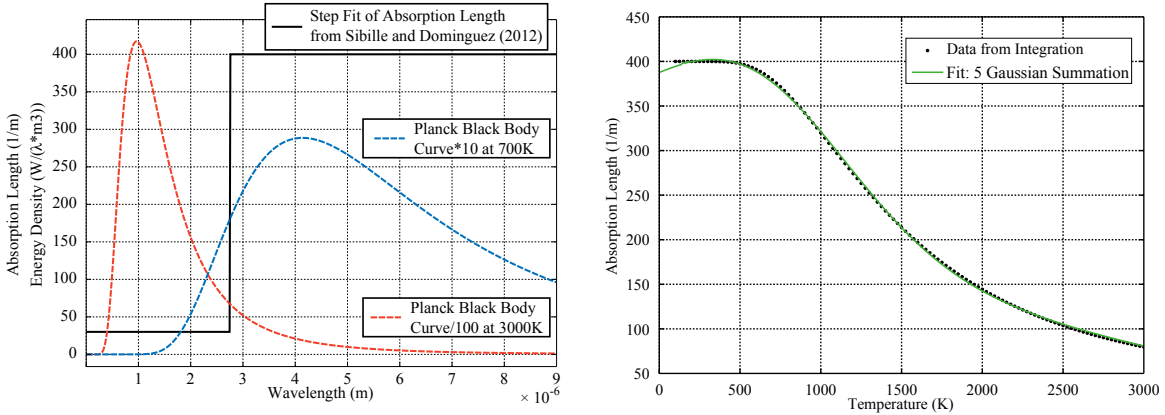


Figure 2-7: (*left*) Absorption length data from Sibille and Dominguez [99] fit with a step function (black solid line) overlaid with the Planck black body spectral energy density at two different temperatures (dashed lines). (*right*) The wavelength-dependent step function was integrated over the Planck curve to calculate the average absorption length as a function of temperature.

2.2.7 Current Efficiency

The current efficiency, defined as the ratio of the actual oxygen mass produced at the anode to the theoretical mass predicted by Faraday’s law [39], is critical in determining a realistic estimate of reactor current. A specific current efficiency was assigned to each oxide specie in lunar regolith using experimental data [103, 100]. The oxide species in lunar regolith are preferentially electrolyzed in the order of Gibbs Free Energy of Formation – thus at any given point in time one particular oxide specie will be the primary object of electrolysis with an associated current efficiency.

From experimental work conducted at MIT, the expected current efficiency while electrolyzing iron-bearing molten mixtures is 30-60% [103] due to the redox cycling

of multivalent iron ions [1]. As a conservative estimate, the lower bound of 30% was used for FeO. A current efficiency of 50% was used for TiO₂ due to a similar redox cycling phenomena [22]. Na₂O, P₂O₅ and MgO form gaseous Na, P, and Mg when electrolyzed [103], which will likely recombine with some of the product O₂ gas within or above the melt. This can produce a cyclic behavior, so an estimated current efficiency of 50% was assigned to Na₂O, P₂O₅ and MgO (though this affect can be mitigated with an anode collection tube [103]). For all other species in lunar regolith, a current efficiency of close to 100% can be expected [100]. As a conservative estimate, a value of 95% was used.

2.2.8 Gibbs Free Energy & Enthalpy of Formation

The Gibbs Free Energy (ΔG) of each specie describes the energy required to break the oxide chemical bonds to form gaseous oxygen and the product metal or non-metal. Similarly, the Enthalpy of Formation (ΔH) for a given oxide specie describes the energy required to electrolyze that specie in an isothermal manner. Data for these two values for each oxide specie in lunar regolith was obtained from the NIST-JANAF thermochemical tables [21], shown in Figure 2-8. Linear regression was used generate temperature-dependent models for $\Delta G(T)$ and $\Delta H(T)$, shown by the solid lines. Due to the fact that $\Delta G(T)$ and $\Delta H(T)$ often changed dramatically across the phase change, the data was separated into solid and liquid data and the relevant data set was used for the given operating temperature.

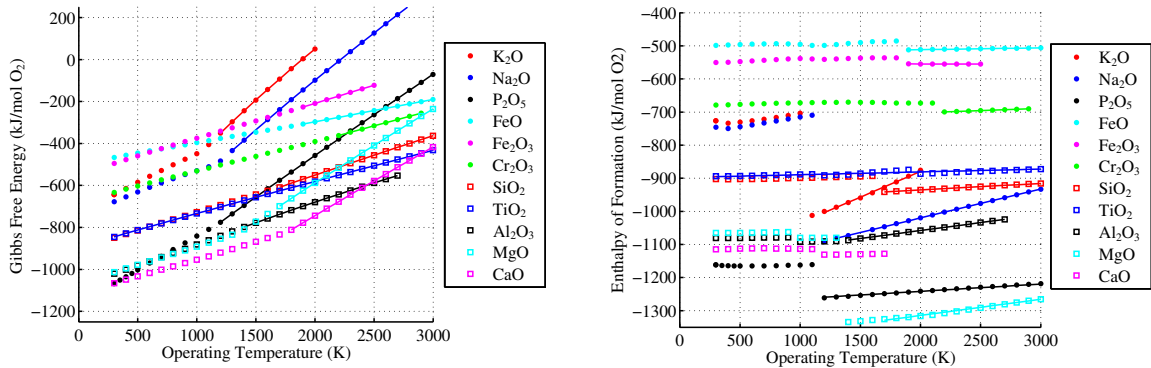


Figure 2-8: The Gibbs Free Energy (left) and Enthalpy of Formation (right) for each oxide specie in lunar regolith [21].

Although the source data for Gibbs Free Energy and Enthalpy of Formation is rooted in significant experimental work [21], the values will need to be adjusted based on the relative concentration of each oxide in the molten core. Future work can address this adjustment, though the effect is not anticipated to significantly change the model power predictions.

2.2.9 Latent Heat of Melting/Fusion

The latent heat of melting depends on regolith composition and therefore regolith type. The modal mineralogical distribution of plagioclase, olivine, pyroxene, and ilmenite from Basu and Riegsecker [12] is presented in Table 2.5 along with the latent heat of melting for each mineral from Richet and Bottinga [86]. To calculate the latent heat of melting for each regolith type, the latent heat of melting for each mineral was weighted by the molar fraction, the results of which are shown in bold in the bottom row of Table 2.5.

Table 2.5: The modal mineralogical distributions for three types of lunar regolith [12], used to weight the latent heat of melting for each mineral [86] in order to calculate the latent heat of melting for each regolith type.

	Molar Fraction (χ_i)			Latent Heat of Melting (kJ/mol)
	High-Ti Mare	Low-Ti Mare	Highlands	
Plagioclase	0.2830	0.2875	0.7320	134.3
Olivine	0.0465	0.0880	0.0230	83.2
Pyroxene	0.5945	0.5890	0.2440	65.1
Ilmenite	0.0760	0.0355	0.0010	21.7
Latent Heat of Melting (kJ/kg)	449.2	457.7	478.6	

As shown in Table 2.5, the latent heat of melting for lunar regolith is around 450-480 kJ/kg, depending on regolith type. These values compare well with the reported value of 506 kJ/kg for a Diopside (Di) - Forsterite (Fo) - Anorthite (An) system with Di:Fo:An = 49.0:7.5:43.5 wt% [64], which has a similar composition to lunar regolith. These predictions are also similar to the value derived from Colozza [28] of 400 kJ/kg (calculated by integrated the value of 1.429 kJ/kg-K over the melting range of 1373-1653K: $1.429\text{kJ/kg-K} * (1653\text{K} - 1373\text{K}) = 400\text{kJ/kg}$). It is worth

noting that the values reported here are significantly higher than the value reported for lunar simulant MLS-1 (161.2 kJ/kg) [87], although this value is for an average sample size of 33.2 mg and the uncertainty on this value is unknown.

2.2.10 Miscellaneous Regolith Material Parameters

For the multiphysics model described below, the surface emissivity of lunar regolith also played a role in determining how much heat was transferred from the regolith bed to the upper reactor wall. The surface emissivity of lunar regolith was taken to be 0.9 ([49]).

2.3 Regolith Throughput Requirements

A critical function in the MRE reactor model is calculating the amount of regolith that needs to be processed to meet the given oxygen production level. The regolith throughput requirement depends on the regolith feedstock type and the operating temperature. The amount of oxygen extracted per kilogram regolith, termed the “oxygen extraction efficiency” (η_{O_2}), was calculated as:

$$\eta_{O_2} \equiv \frac{m_{O_2}}{m_{regolith}} = \sum_i (w_i) \left(\frac{MW_{O_2}}{MW_{oxide,i}} \right) (r_{mol,i})(e_{frac,i}), \quad (2.10)$$

where m_{O_2} is the mass of oxygen produced, $m_{regolith}$ is the mass of regolith processed, w_i is the weight percent of oxide i in lunar regolith, $MW_{oxide,i}$ is the molecular weight of oxide specie i , $r_{mol,i}$ is the ratio of the number of moles of oxygen per mole oxide i ($\frac{mol_{O_2}}{mol_{oxide}}$), and $e_{frac,i}$ is the fraction of oxide i that is electrolyzed in each batch. Equation (2.10) sums the contribution of each oxide together to determine the total oxygen that can be extracted from lunar regolith, taking into account the fact that one might not electrolyze all (or any) of a particular oxide via the term $e_{frac,i}$. It is important to note that $e_{frac,i}$ was set to a maximum of 0.95 for all species to account for the fact that when the concentration of a particular specie becomes too low, the specie will no longer be preferentially electrolyzed in favor of more abundant species.

The weight percent (w_i) of each specie naturally depends upon the regolith type, as detailed in Section 2.2.1.

The fraction of each specie that is electrolyzed ($e_{frac,i}$) is dependent upon the operating temperature, because the melting temperature of the molten region generally increases as the composition changes throughout the electrolysis process. For a number of points throughout the batch electrolysis process, the liquidus temperature for the corresponding composition was derived using phase diagrams from Slag Atlas [38]. This analysis begins after most of the FeO has been electrolyzed, because the liquidus temperature through the entire FeO electrolysis ($<1700\text{K}$) is significantly less than the operating temperatures considered in this study ($>1850\text{K}$). As depicted in Figure 2-9, as electrolysis progresses the liquidus temperature first increases as the silicates (SiO_2) are reduced. After most of the silicates are electrolyzed, the reduction of MgO and TiO_2 leads to a general decrease in liquidus temperature. Finally, as Al_2O_3 is reduced the liquidus temperature first decreases and then increases to 2845K for a mixture primarily composed of CaO.

For a given operating temperature, electrolysis is allowed to progress until the liquidus temperature approaches the operating temperature. A 50K safety margin between the operating temperature and the liquidus temperature was used in this MRE reactor model. As can be seen in Figure 2-9, for a temperature of 1900K, one can electrolyze approximately 1/2 of the SiO_2 (0.19 g O_2 per kg regolith) before needing to insert fresh regolith into the reactor. But if the temperature is raised to 2100K, one can now electrolyze 5/6 of the SiO_2 (0.25 g O_2 per kg regolith). Temperatures above 2200K would open up the possibility of electrolyzing MgO, TiO_2 , and some amount of Al_2O_3 . Thus, higher operating temperatures increase the oxygen extraction efficiency by allowing the reactor to extract more oxygen per kilogram regolith.

Because the composition of lunar regolith depends on regolith type (as described in Section 2.2.1), the liquidus temperature curve depends on regolith type as well. The curve shown in Figure 2-9 represents the curve for Low-Ti Mare, and the curves for Low-Ti Mare, High-Ti Mare, and Highlands regolith are overlaid in Figure 2-10. Perhaps the most interesting feature seen in Figure 2-10 is that although the liquidus

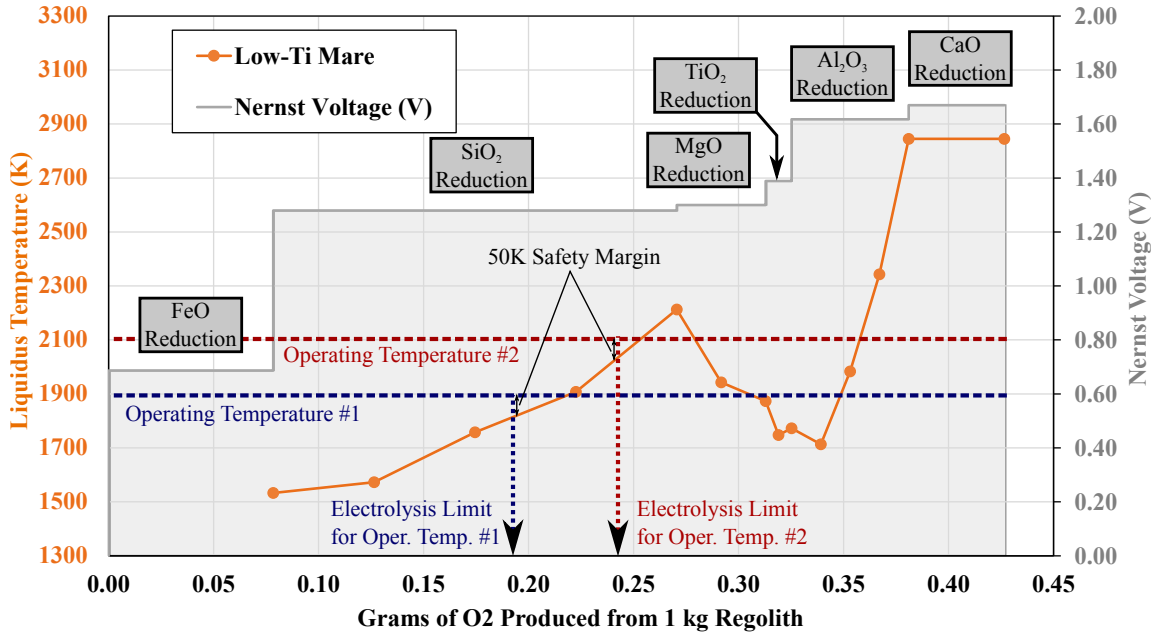


Figure 2-9: As electrolysis progresses, the composition and properties of the molten regolith vary. The liquidus temperature, calculated from Slag Atlas [38], and the Nernst decomposition potential for Low-Ti Mare lunar regolith initially increase as electrolysis progresses. As the operating temperature increases, the electrolysis can proceed farther to the right in the plot.

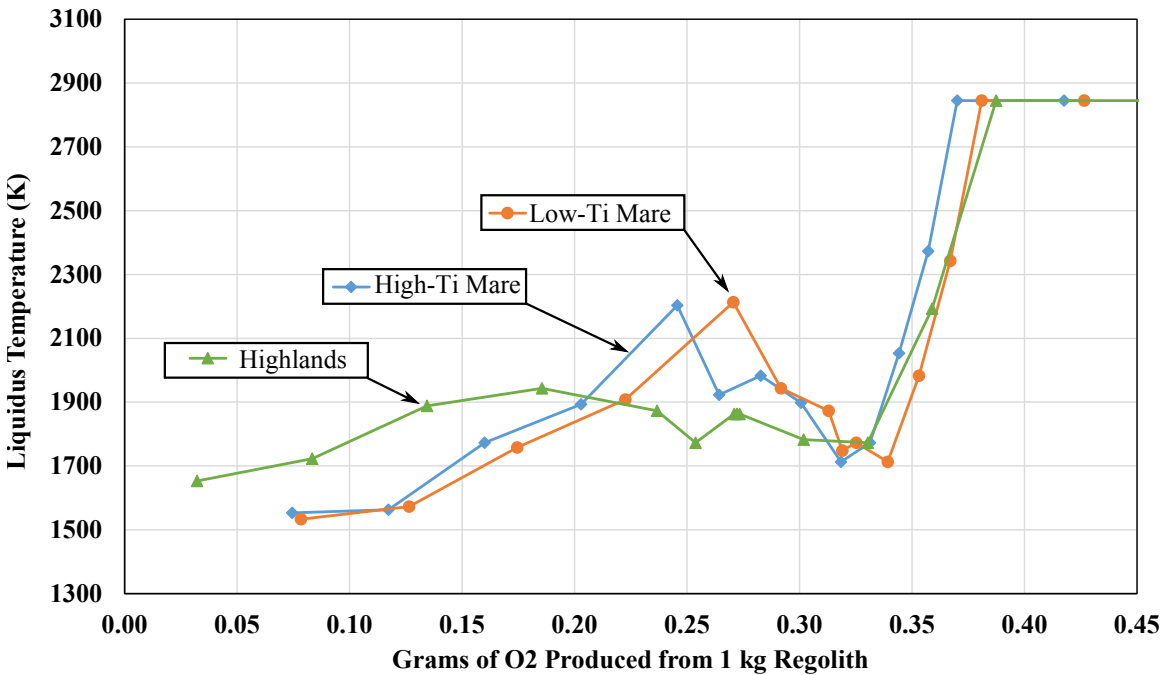


Figure 2-10: The liquidus temperature, calculated from Slag Atlas [38], for all three types of regolith throughout the electrolysis process. As the operating temperature increases, the electrolysis can proceed farther to the right in the plot.

temperature for Highlands regolith is initially higher than both Mare types, as the silicates are electrolyzed the Mare liquidus temperature rises significantly above that of the Highlands. This is due to the fact that Mare regolith has comparatively more MgO while Highlands regolith has comparatively more Al_2O_3 .

The oxygen extraction efficiency in Equation (2.10) was used to calculate the annual mass of regolith required to meet a given oxygen production level. The required amount of molten regolith in the central core for each batch ($m_{reg, batch}$) was then calculated as:

$$m_{reg, batch} = \frac{m_{regolith}}{n_{batches}} = \frac{m_{oxygen}/\eta_{O2}}{n_{batches}}, \quad (2.11)$$

where η_{O2} is the oxygen extraction efficiency from Equation (2.10) and $n_{batches}$ is the number of batches in a year, calculated as:

$$n_{batches} = \frac{t_{operation}}{t_{batch}}, \quad (2.12)$$

where $t_{operation}$ is the cumulative operating time for the reactor adjusted by duty cycle, and t_{batch} is the total batch time including any down time between batches.

2.4 Electrochemistry

2.4.1 Estimating Current

The MRE reactor model employs a number of equations to predict the electrochemical performance of the reactor. Faraday's law [39, 112], with an adjustment for the expected current efficiency, was used to calculate the expected number of coulombs required to produce the given mass of oxygen (m_{O_2}):

$$C = (n_{O_2}) \left(\frac{nF}{\bar{\eta}_I} \right) = \left(\frac{m_{O_2}}{MW_{O_2}} \right) \left(\frac{nF}{\bar{\eta}_I} \right), \quad (2.13)$$

where n_{O_2} is the moles of oxygen in the given oxygen mass, n is the number of electrons required per diatomic oxygen molecule (=4), F is Faraday's constant, $\bar{\eta}_I$ is the average current efficiency (see Section 2.2.7), and MW_{O_2} is the molecular weight

of diatomic oxygen. Using Equation (2.10) to substitute for $\left(\frac{m_{O_2}}{MW_{O_2}}\right)$ and including the oxide-specific current efficiencies in the summation yields:

$$C = (m_{regolith})(nF) \sum_i \frac{(w_i)(r_{mol,i})(e_{frac,i})}{(MW_{oxide,i})(\eta_{I,i})}, \quad (2.14)$$

where $m_{regolith}$ is the mass of regolith that needs to be processed to meet the given oxygen production level, w_i is the weight percent of oxide i from Figure 2-2, $r_{mol,i}$ is the ratio of the number of moles of oxygen per mole oxide i ($\frac{mol_{O_2}}{mol_{oxide}}$), $\eta_{I,i}$ is the current efficiency of oxide i , and $e_{frac,i}$ is the fraction of oxide i that is electrolyzed in each batch as discussed in Section 2.3.

Using the number of coulombs predicted by Equation (2.14), the average current (I_{avg}) could be calculated:

$$I_{avg} = \frac{C}{(t_{batch} - t_{feed})}, \quad (2.15)$$

where t_{batch} is the total batch time and t_{feed} is the downtime in each batch when fresh regolith is inserted. Although one can certainly envision operating an MRE reactor semi-continuously while inserting fresh regolith and withdrawing molten metals, the conservative assumption of a downtime of 5 min during each batch is made, during which electrolysis is not performed.

2.4.2 Estimating Voltage

The reactor was designed to be operated in a constant power mode, in which the power demand is spread evenly over time to minimize the required power system size. To maintain a constant power level, the current is varied (with an average value of I_{avg}) throughout the batch as the voltage changes due to different species being electrolyzed and the conductivity of the melt changing throughout the batch. Thus, for species like FeO, which have a lower ΔG and increase melt conductivity, the expected voltage will be lower than when electrolyzing other species, such as SiO₂. To maintain constant power, FeO will then be electrolyzed with a higher than average

current. The ramifications of the constant power mode will be further discussed in Section 2.5.3.

An estimate of the required voltage drop between the electrodes during the electrolysis of oxide specie i (ΔV_i) was calculated in a manner similar to that presented in Haskin et al. [45]:

$$\Delta V_i = \Delta V_{Nernst,i} + \Delta V_{joule-heated} + \Delta V_{endothermic,i}, \quad (2.16)$$

where $\Delta V_{Nernst,i}$ is the Nernst potential required to reduce oxide i [45], $\Delta V_{joule-heated}$ is the over-voltage required to offset the radiative heat loss to the environment and the extra heating required to heat fresh regolith to the operating temperature, and $\Delta V_{endothermic,i}$ is the over-voltage required to generate enough heat to offset the endothermic electrolysis reaction. Equation (2.16) is used to calculate the required voltage to sustain thermal equilibrium as well as perform electrolysis. It is assumed that the over-voltage need to maintain thermal equilibrium will be sufficient to drive the reaction at a suitably fast rate.

The estimated voltage, V_i , is supplied by setting a certain electrode separation (the $i \cdot R_{cell}$ term from Haskin et al. [45]). As described later in Section 2.6, a multiphysics simulation of an MRE reactor is utilized to determine exactly what electrode separation is required to maintain thermal equilibrium.

2.4.3 Batch Profiles

The MRE reactor model also produces “batch profiles”, which are plots of the estimated voltage and current throughout an entire batch to maintain a constant power level. As electrolysis progresses, species are preferentially electrolyzed in the order of increasingly negative $\Delta G_i(T)$. Batch profiles for an operating temperature of 1950K and 2250K are shown for High-Ti Mare regolith in Figure 2-11, including the Nernst voltage (dashed orange line), and the total estimated voltage, V_i (solid orange line). An operating temperature of 1950K (left) allows all of the FeO to be electrolyzed and approximately half of the SiO₂ before the melt solidification temperature rises

to within 50K of the operating temperature, as shown in Figure 2-9. An operating temperature of 2250K enables a more complete electrolysis of lunar regolith, including MgO, TiO₂ and some Al₂O₃. Note that at this higher temperature $\Delta G_i(T)$ for MgO actually rises above that of SiO₂, meaning these two oxide species will likely be electrolyzed concurrently – the incorporation of this phenomenon into the model remains a topic for future work, especially with regards to the liquidus temperature curves shown in Figures 2-9 and 2-10.

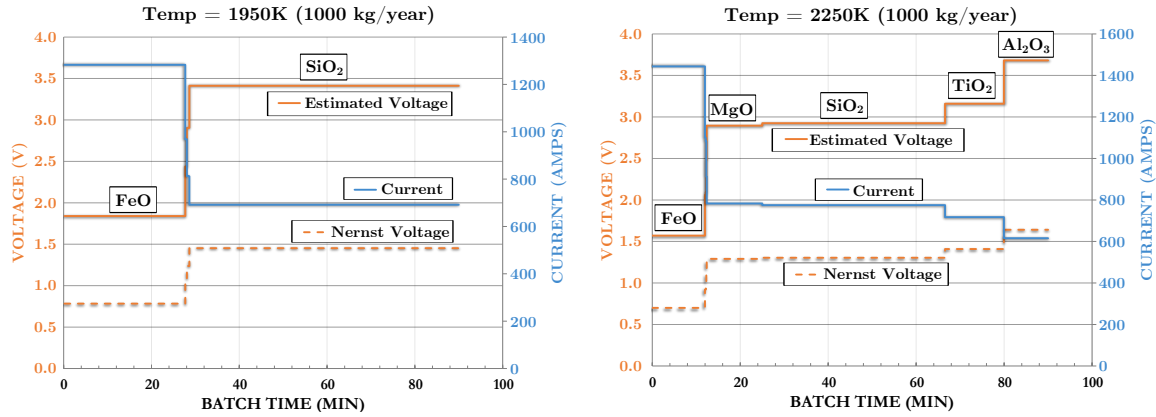


Figure 2-11: The estimated voltage and current profiles over a single batch at an operating temperature of 1950K (left) and 2250K (right), along with the primary species being electrolyzed. The current is varied inversely with voltage to achieve constant power operation.

2.4.4 Metal Production

Using the fraction electrolyzed ($e_{frac,i}$) of each oxide in lunar regolith, the amount of molten metal produced can also be calculated:

$$m_{metal,i} = (m_{regolith})(w_i) \left(\frac{MW_{metal,i}}{MW_{oxide,i}} \right) (e_{frac,i}), \quad (2.17)$$

where $m_{metal,i}$ is the mass of a particular metal produced, $m_{regolith}$ is the mass of regolith processed, w_i is the weight percent of oxide i in lunar regolith, $MW_{metal,i}$ is the molecular weight of the product metal i , and $MW_{oxide,i}$ is the molecular weight of oxide i reduced to produce the desired metal. To determine $m_{regolith}$, one can use

Equation (2.10) if oxygen is the primary product or Equation (2.17) if a particular metal is the primary product.

2.5 Multiphysics Simulation

2.5.1 Simulation Motivation and Overview

The previous sections have laid out the framework for the MRE reactor model, but the issue of molten regolith containment described in Section 1.2.1 must also be taken into account. An MRE reactor must be carefully designed to enable the joule-heated, cold-wall (JHCW) operation mode where the central molten regolith core is insulated by a frozen layer of solid regolith. Molten lunar regolith is extremely corrosive and cannot be contained for extended periods of time by traditional crucible materials [103], but solid regolith provides an excellent “crucible” for the molten core and can enable long-term operation [99].

Although JHCW operation is an elegant solution to the issue of molten regolith corrosion, this complex electro-thermodynamic requirement imposes significant constraints on the geometry and thermal topology of an MRE reactor. Although assuming a certain fraction of the regolith in the reactor is molten can provide a first-order estimate, it is not of high enough fidelity to be compared to the models of other oxygen extraction techniques. To provide higher fidelity designs, a multiphysics simulation was developed using COMSOL™.

2.5.2 Simulation Description

The reactor simulation includes chemical, electrodynamic and thermodynamic physics because they are tightly coupled. The electrical conductivity of the molten core strongly depends on temperature. However, this dependence is cyclic, as the temperature depends on the resistive heating within the molten region, which strongly depends on the electrical conductivity. Furthermore, although the Enthalpy of Formation of the reactor contains a weak temperature dependence, the Gibbs Free Energy

contains a strong dependence - thus, the amount of endothermic heating required and the Nernst voltage depend heavily on temperature. By modeling all three of these physical phenomenon concurrently, the simulation can capture all of these dependencies.

Figure 2-12 shows a cross-sectional side view of an MRE reactor design in the multiphysics simulation. To generate the 3D reactor design, this 2D profile is rotated around the axis of rotation (the dashed line to the left). An axis-symmetric simulation was chosen over a full 3D simulation as it cut the computation time down by several orders of magnitude. The simulation design includes an outer cylindrical shell and an anode (top) and cathode (bottom) comprised of a cylindrical shaft and plate.

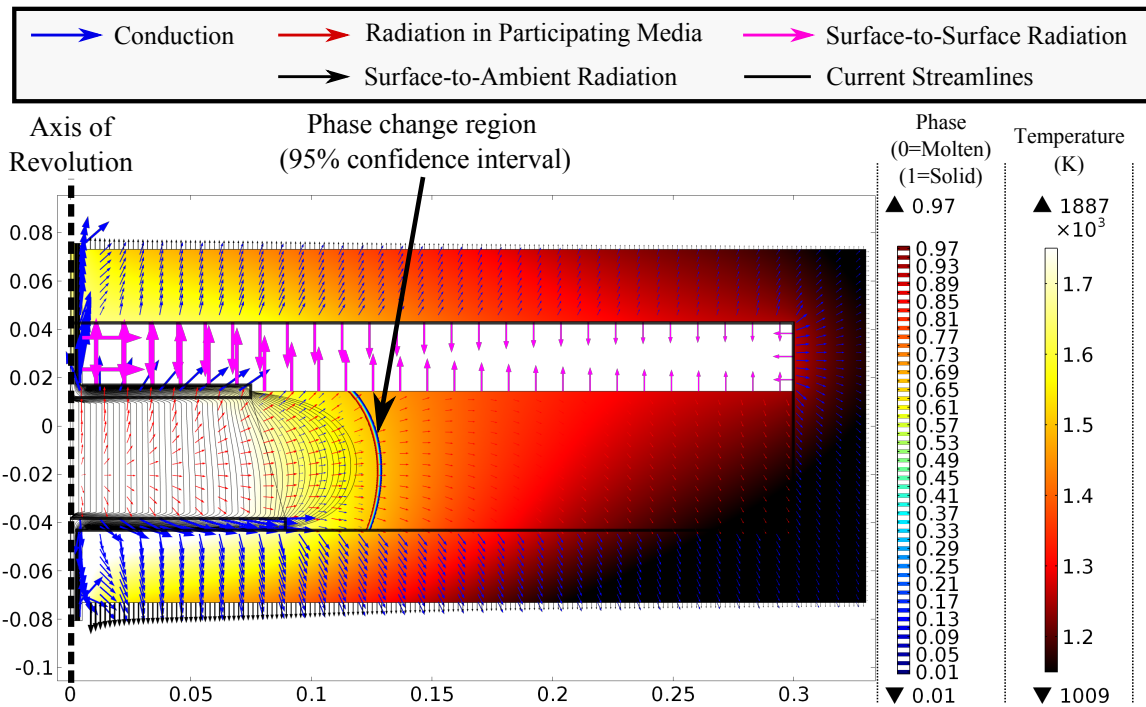


Figure 2-12: A side view of the cross-section of the multiphysics simulation of a cylindrical MRE reactor. The primary heat fluxes modeled are shown with different colored arrows, scaled by the same factor.

There are four primary modes of heat transfer in the simulation: conduction (the blue arrows in Figure 2-12), surface-to-surface radiation in the cavity above the regolith bed (magenta arrows), radiation to ambient on the exterior surface of the reactor (black arrows), and radiation in participating media (red arrows). Radiation

in participating media models the emission, absorption, and scattering of radiation within the high temperature translucent regolith [99]. The reactor was coated with Multi-Layer Insulation (MLI) [41], resulting in an effective surface emissivity of 0.05. The surface emissivity of the iridium anode was taken to be 0.3 [18]. All arrows in Figure 2-12 are logarithmically scaled by the same factor to allow for comparison between the various modes of heat transfer.

The endothermic nature of the electrolysis reaction, linked to the difference between the Gibbs Free Energy and Enthalpy of Formation ($\Delta H - \Delta G$), was modeled by introducing an additional heat sink in the central molten region. Although this value will change over time with the oxide that is preferentially being electrolyzed, the simulation is for a steady-state reactor and thus steady-state heat sink must be used. The average heatsink due to the endothermic electrolysis reaction ($\dot{Q}_{endothermic}$) was calculated as:

$$\dot{Q}_{endothermic} \equiv (\dot{n}_{O_2}) \text{ avg}[\Delta H - \Delta G] = (\dot{n}_{O_2}) \sum_i (\Delta H_i(T) - \Delta G_i(T)) \left(\frac{c_i(T)}{C} \right), \quad (2.18)$$

where \dot{n}_{O_2} is the average molar oxygen production rate (mols/sec), $c_i(T)$ is the number of coulombs used to electrolyze oxide specie i during a single batch, and C is the total number of coulombs used in a single batch from Equation (2.14). The left plot in Figure 2-13 shows the endothermic heat sink for High-Ti Mare as a function of operating temperature and average current. Equation (2.18) was not directly used in the multiphysics model, due to inaccessibility to several of the variables from within the multiphysics simulation. Rather, the data generated by Equation (2.18) was fit with a piecewise third-order polynomial regression model and integrated into the multiphysics simulation.

A second heat sink was integrated into the multiphysics simulation to capture the effect of periodically (or continuously) inserting fresh regolith in the reactor. The reactor model assumes that fresh solid regolith will start around 300K and will be partially preheated by exposing the fresh regolith to some of the heat loss from the reactor. This preheating step will bake off many of the volatiles observed in the




Figure 2-13: The predicted heat sinks due to the endothermic electrolysis reaction (left) and required power to heat fresh regolith feedstock (right) as a function of reactor operating temperature and current for High-Ti Mare.

Apollo samples, including H_2O , CO_2 , N_2 , and He [77], as well as some amount of the Na_2O and K_2O present in lunar regolith, which would somewhat simplify the electrolysis process. Although the regolith will be partially preheated, it will still require additional heating to reach the molten state and operating temperature. Thus, the preheated regolith will be inserted into the molten core and will then be heated to a molten state through heat transfer with the surrounding molten regolith. As a first-order estimate, it was assumed that half of the heating power would be provided by joule-heating in the reactor core:

$$\dot{Q}_{\text{regolith-heating}} = \frac{1}{2} \int_{300\text{K}}^{T_{op}} C_p(T) dT, \quad (2.19)$$

where T_{op} is the operating temperature and $C_p(T)$ is shown in Section 2.2.3. As discussed later in Section 2.6.5, 100% of the regolith heating power is included in the total power estimate, but for the thermal modeling of the reactor it is assumed that half of the heating power is done by a preheater prior to inserting fresh regolith into the reactor core. Thus, there is a factor of $\frac{1}{2}$ in front of the integral in Equation (2.19) but not in Equation (2.29).

The right plot in Figure 2-13 shows the heat sink related to the regolith heating power, which actually *decreases* at higher operating temperatures. Although the amount of heat per kilogram regolith increases with operating temperature, this effect

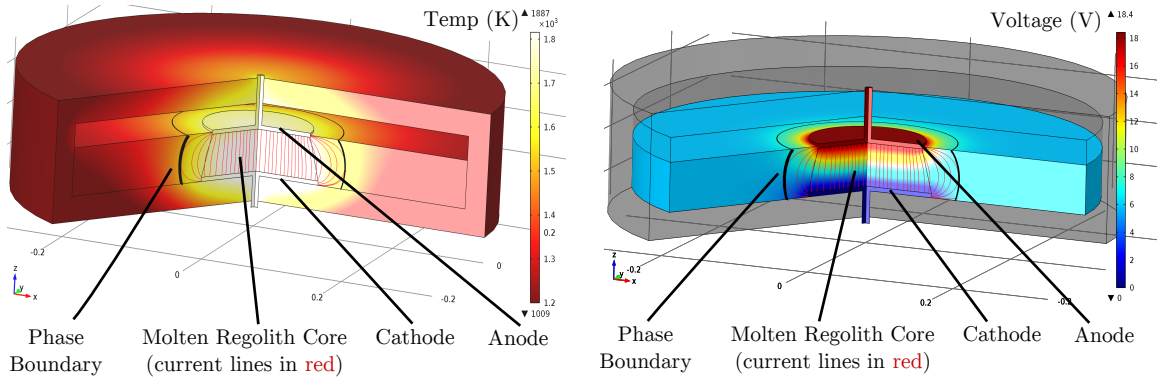


Figure 2-14: The temperature (left) and voltage (right) profiles generated by the multiphysics simulation for a given reactor design. Current lines are shown in red and the phase boundary is shown in black.

is apparently outweighed by the decreased regolith throughput requirement as a result of the increase in oxygen extraction efficiency at higher operating temperatures. This effect is described in more detail in Section 2.3.

The simulation generates both temperature and voltage profiles of the reactor as shown in Figure 2-14. The black border in the reactor indicates the phase boundary between the inner molten regolith and the outer solid regolith. Current streamlines through the central molten region are depicted in red and are primarily limited to the central region. Due to the fact that a continuous curve was used to model the electrical conductivity of lunar regolith (see Section 2.2.5), the electrical conductivity of solid regolith was modeled slightly higher than reality for the sake of dramatically decreasing model convergence time. This resulted in some small fictitious amount of current traveling beyond the phase boundary through solid regolith, but this “overflow” current was typically less than 0.1% of the total current passing through the reactor.

The multiphysics simulation was also used to measure several performance metrics for each reactor design. The heat loss was calculated by integrating the net radiative heat flux from the outer surface. The operating voltage was measured between the anode and cathode plates. The operating temperature was calculated as the average temperature within the molten region weighted by current density.

2.5.3 Constant Power Mode

As mentioned in Section 2.4.2, during conceptual operation of the reactor the voltage and current are varied inversely to achieve constant power consumption. For instance, the predicted operating voltage (see Equation (2.16)) is around 1.85 V for FeO and 3.45 V for SiO₂. The respective operating current for these two species is 1300 A and 700 A. Both of these operating conditions result in a power consumption of approximately 2.4 kW.

This constant power consumption is important for two reasons. First, it reduces the required power system size by minimizing peak power consumption. Second, it means that regardless of which oxide specie is currently being reduced, the reactor should behavior similarly from a thermal perspective. This is not strictly true, it is actually the over-potential multiplied by current which should remain constant throughout operation because the minimal Nernst voltage only goes towards electrolysis, not generating heat. Thus, operating in a constant total power mode will result in some small variations in the total heat generated within the reactor as different species are electrolyzed, but will reduce the peak power load and resultant power system size. Future work will have to address the impact of this operating mode in more detail.

2.6 Using the Simulation to Guide Reactor Design

There are five criteria that must be met by the MRE reactor design. The first two, molten mass within the reactor and current through the reactor, are derived from the oxygen production level using Equations (2.10) and (2.15), respectively. The third and fourth, operating temperature and maximum wall temperature, are set as model inputs. The maximum wall temperature, when set to less than the melting temperature of regolith, ensures that the reactor can support joule-heated, cold-wall operation (see Section 1.2.1). The fifth criteria is the design margin (described in Section 2.6.3), which dictates the amount of operational flexibility in a given design. To meet these five performance criteria, the reactor diameter, wall thermal conduc-

tivity, and electrode separation must be carefully designed. This section presents a novel design methodology developed to address this design challenge.

2.6.1 Effects of Molten Mass and Operating Temperature: The Cutoff Line

Over 100,000 reactor designs (consisting of unique combinations of reactor diameter, electrode separation, wall thermal conductivity, and current) were evaluated using the multiphysics simulation to generate a database of the associated performance characteristics (molten mass within the reactor, operating temperature, maximum wall temperature, and whether or not the joule-heated cold-wall condition was satisfied). Figure 2-15 shows a sample of the data for the mass of molten regolith in the reactor (left) and operating temperature (right) from the 100,000 designs evaluated using the multiphysics simulation. Both plots illustrate how the diameter and electrode separation affect the molten mass and operating temperature in a reactor for a fixed current of 2000 A and wall thermal conductivity of 5.0 W/m-K.

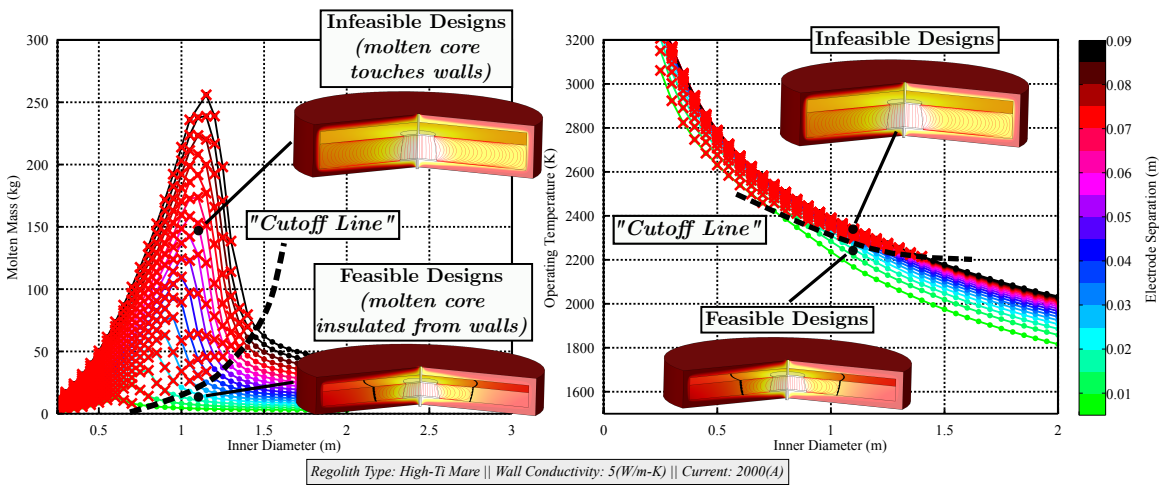


Figure 2-15: The molten mass in an MRE reactor (left) and the operating temperature (right) depend on reactor diameter and electrode separation. Red X's indicate infeasible designs in which the side wall temperature gets too close to the melting temperature of regolith (1500K).

The molten mass on the left of Figure 2-15 first rises with the diameter squared because the reactor interior is completely filled with molten regolith. Around a diam-

eter of 0.5 m to 1.0 m (for electrode separations between 1 cm and 9 cm, respectively), an inflection point occurs where the molten mass reaches a maximum value and then begins to decrease with diameter. In this region of the design space, solid regolith occurs around the outer wall of the reactor, decreasing the mass of molten regolith in the reactor. Beyond this inflection point, the molten mass decreases as reactor diameter increases. This is because there is now more surface area to radiate the resistive heat generated in the reactor core, so the outer surface can now be at a lower temperature while still radiating the required amount of heat. This temperature drop is inversely proportional to the fourth root of the outer surface area ($T_{out} \propto (D^2 + D)^{-1/4}$), by the Stefan-Boltzmann law. The drop in outer surface temperature creates a steeper temperature gradient in the reactor, which decreases the total molten mass in the reactor. Conversely, as diameter increases, the distance between the central molten core and the outer wall increases, which creates a less steep thermal gradient. This effect is apparently outweighed by the drop in outer wall temperature described above. Thus, as the diameter increases beyond the inflection point, the temperature profile in the reactor drops more quickly away from the inner core temperature, reducing the amount of molten mass in the reactor as shown in Figure 2-15.

Data points in Figure 2-15 with a red “X” overlaid represent infeasible designs, in which the layer of solid regolith insulating the molten core becomes too thin or nonexistent. Three different cutoff values were tested for identifying when this occurred: when the wall temperature was above 1200K, 1300K, or 1400K. The data in Figure 2-15 uses a cutoff value of 1400K to identify infeasible designs with too high of a wall temperature. A comparison of the three different cutoff values is provided in Section 2.7.2. The melting temperature of regolith is around 1500K [49], so all three cutoff values provide a factor of safety to ensure that no molten regolith touches the wall. Due to the fact that molten metal will pool at the bottom of the reactor and insulate the reactor wall from the molten regolith, temperatures in excess of 1500K were acceptable along the bottom wall.

The reactor design is chosen using the line that separates the infeasible designs (the red “X” region) from the feasible designs, termed the “cutoff line” shown in Figure 2-

15. The choice to design reactors on the cutoff lines was made after an initial analysis revealed that those designs had the lowest heat loss per kilogram molten regolith within the reactor. For more details on the justification of designing on the cutoff line, see Appendix A. By designing on the cutoff lines there is a one-to-one mapping between reactor diameter and electrode separation. As can be seen in the left-hand plot in Figure 2-15, as the reactor diameter increases the electrode separation must also increase (moving from the green lines towards the black ones) in order to stay on the cutoff line. Designing on the cutoff line provides a method for scaling up the diameter and electrode separation of a reactor while still maintaining the joule-heated, cold-wall constraint in which solid regolith insulates the reactor walls from the molten core. Of course, one could simply increase reactor diameter without increasing electrode separation, but this in fact results in less efficient reactors that lose more heat per kilogram molten mass in the reactor (see Appendix A).

It is critical to note that the cutoff lines illustrated in Figure 2-15 are for a fixed current of 2000 A and a wall conductivity of 5.0 W/m-K. The cutoff lines are also depend on the current and wall thermal conductivity. Figure 2-16 demonstrates how the reactor current affects the cutoff lines. The top three rows of plots depict cutoff lines for three different reactor currents with a constant value of wall thermal conductivity. The bottom plot in Figure 2-16 shows the cutoff lines from the top three rows, along with several other reactor current values, compiled onto a single plot, where the color of each line in the bottom plot now represents a different reactor current. Reactor current has a minimal affect on the molten mass within the reactor, as observed in the bottom left plot in Figure 2-16. As the reactor current increases, the mapping between reactor diameter and electrode separation changes (shown later in Figure 2.6.4). Thus, although the data in the bottom left plot in Figure 2-16 indicates that the current has a minimal effect on the mass of molten regolith in the reactor, this is achieved by reducing the electrode separation as current increases.

Conversely, operating temperature exhibits a direct increase with reactor current. This is undoubtedly due to the fact that as reactor current increases, the resistive heating within the molten core also increases to generate higher temperatures.

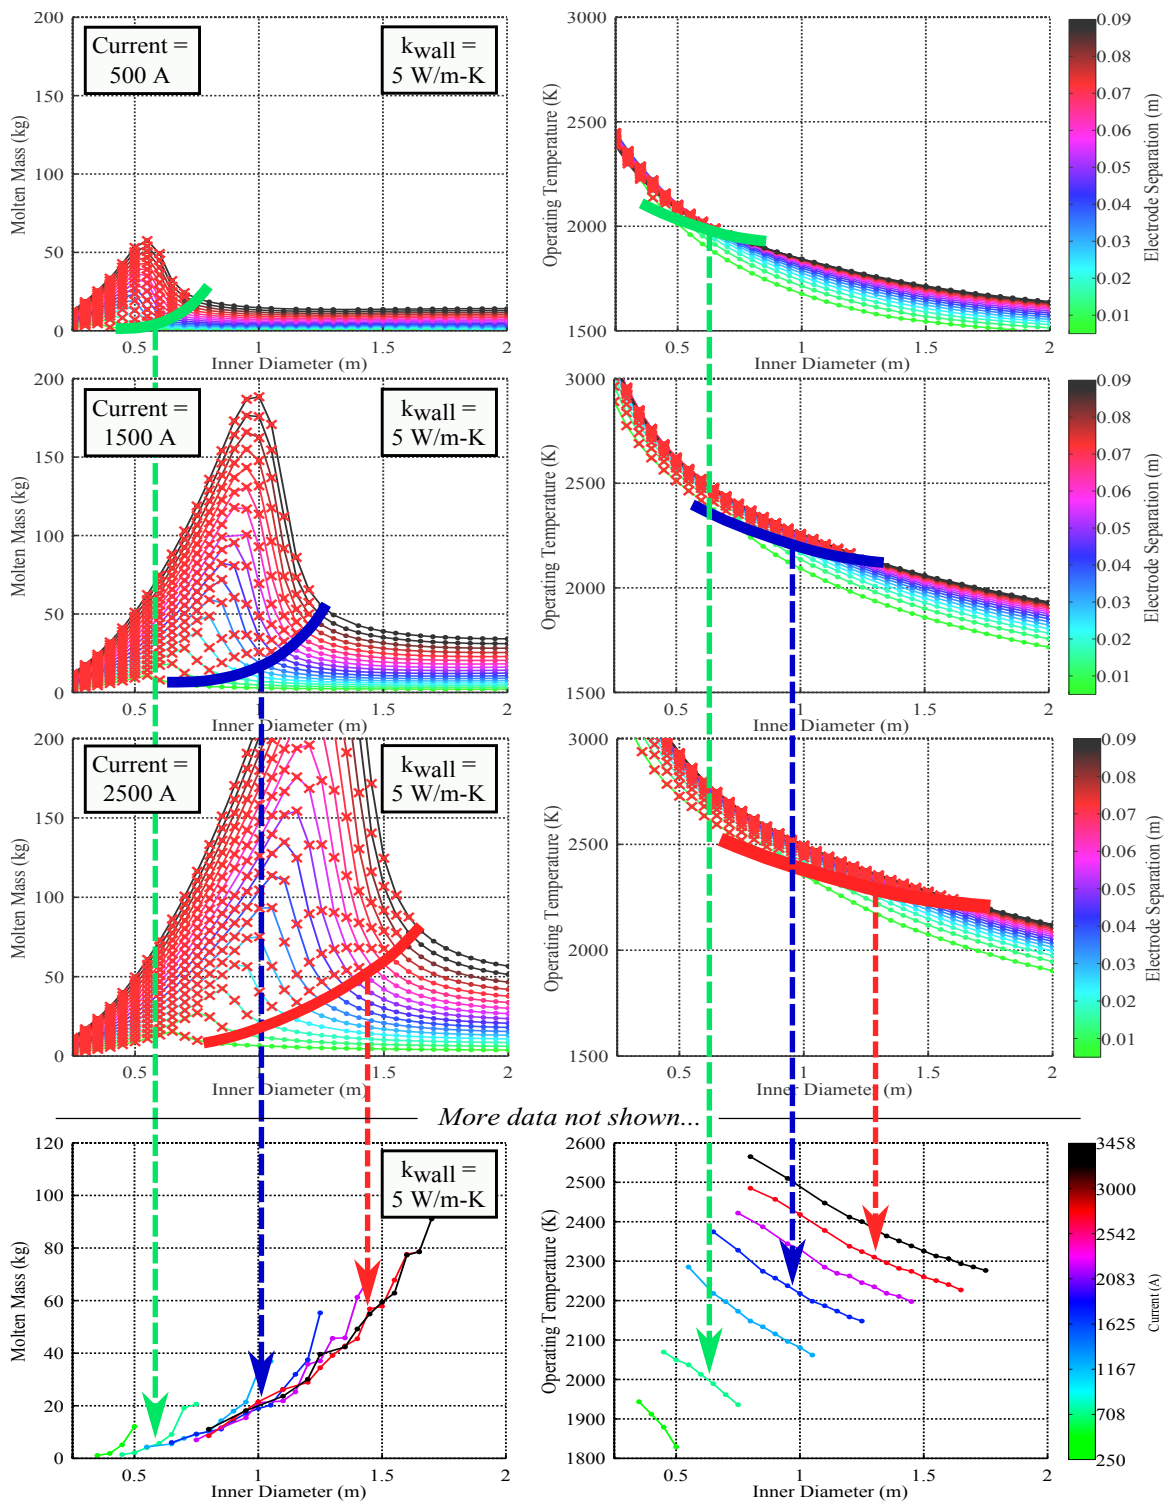


Figure 2-16: An illustration of how the current through the reactor affects the location of the cutoff lines between infeasible (red X's) and feasible reactor designs.

All of the data shown in Figures 2-15 and 2-16 are for reactor designs with a fixed wall thermal conductivity. Intuitively, the cutoff lines are also affected by the wall thermal conductivity, k_{wall} , as shown in Figure 2-17. One clear effect of reducing k_{wall} is an increase in the reactor diameter required to contain a given mass of molten regolith within the reactor. For instance, a molten mass of 4 kg can either be met by a reactor with ($D = 0.8\text{m}$ and $k_{wall} = 5.0 \text{ W/m-K}$) or ($D = 2.0\text{m}$ and $k_{wall} = 0.05 \text{ W/m-K}$). The operating temperature also displays a noticeable decrease with decreasing wall thermal conductivity.

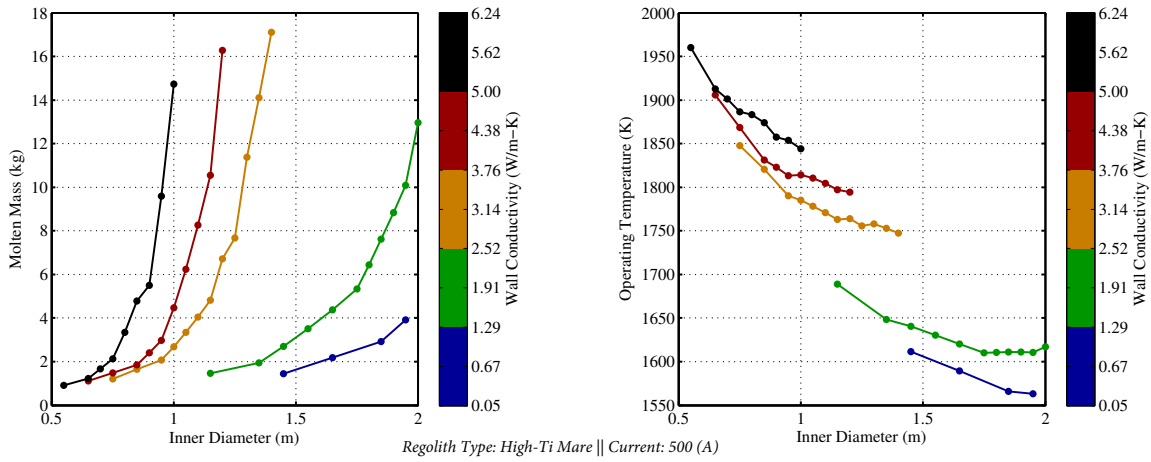


Figure 2-17: The molten mass (left) and operating temperature (right) for reactors on the “cutoff line” between infeasible and feasible designs. Each line represents a different wall thermal conductivity and a current of 500 A. Increasing the wall thermal conductivity decreases molten mass in the reactor and operating temperature.

2.6.2 Bounds on Reactor Diameter

The cutoff line plot at the bottom of Figure 2-16 proves useful for defining bounds on the reactor design and is shown in more detail in Figure 2-18. Succinctly, **1)** the molten mass requirement sets a minimum feasible reactor diameter, D_{min} , **2)** the operating temperature sets a maximum feasible reactor diameter, D_{max} , and **3)** D_{max} and D_{min} depend on the current in the reactor and the thermal conductivity of the reactor wall.

As can be seen in Figure 2-18, the cutoff line represents the minimal reactor di-

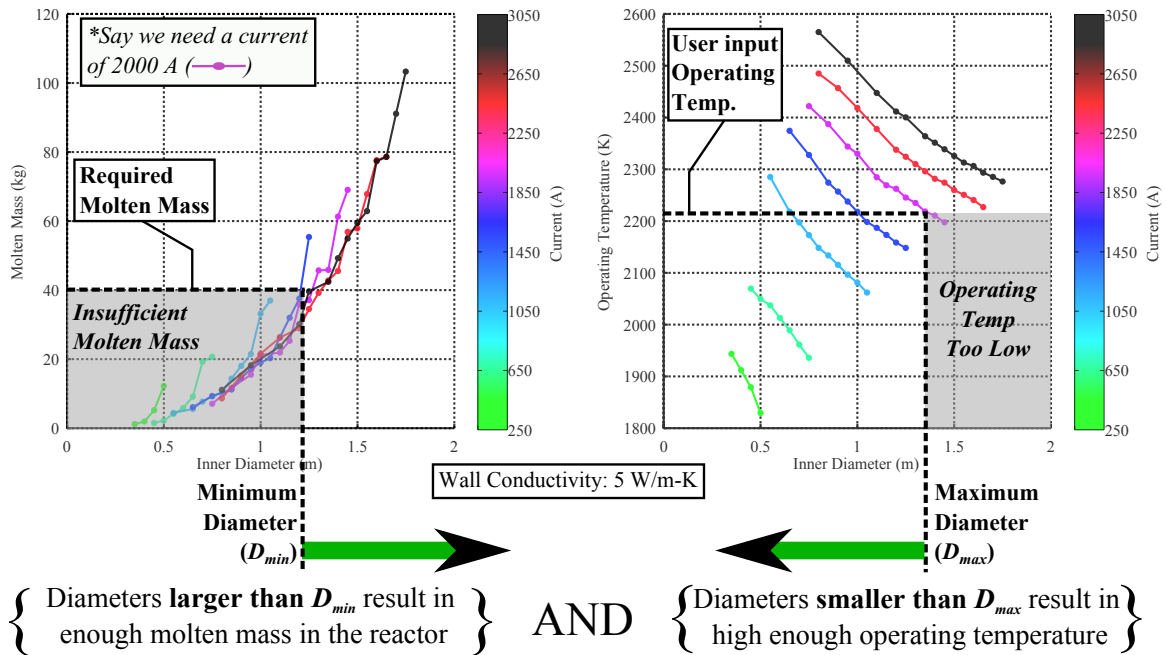


Figure 2-18: The molten mass (left) and operating temperature (right) as a function of reactor diameter when designed on the cutoff line between infeasible and feasible designs. Each line represents a constant value for current and a wall thermal conductivity of 5 W/m-K. The minimum and maximum diameter bounds, resulting from the molten mass and operating temperature constraints respectively, are depicted.

ameter (for a fixed wall thermal conductivity and current) that can meet the required amount of molten mass in the reactor while ensuring that no molten regolith touches the reactor wall. As shown in Figure 2-15, decreasing the reactor diameter (to the left of the cutoff line) crosses into the infeasible region, meaning that molten regolith will touch the reactor wall.

The cutoff line displays the opposite trend in the operating temperature data shown in the right-hand plot in Figure 2-18. Here the cutoff line represents the maximum diameter that satisfies the desired operating temperature (for a fixed wall thermal conductivity and current) while ensuring that no molten regolith touches the reactor wall. As shown in Figure 2-15, increasing the diameter to the right of the cutoff line crosses into the infeasible region, meaning that molten regolith will touch the reactor wall.

As an example, consider a scenario in which 40 kg of molten mass in the reactor and a current of 2000 A are required. An examination of the left-hand plot in Figure 2-18

reveals that this sets a minimum reactor diameter of about 1.2 m. More generally, the molten mass requirement imposes a certain minimum diameter on the reactor, and because there is a one-to-one mapping between diameter and electrode separation, the molten mass requirement also intuitively imposes a minimum electrode separation. Naturally, the diameter can be increased above this minimum threshold to design a reactor with more than sufficient molten mass inside.

From the data shown on the right of Figure 2-18, the operating temperature imposes a maximum feasible reactor diameter. In the MRE reactor model, the user inputs a desired operating temperature. As an example, consider a scenario where the user desires an operating temperature of 2200 K and the desired oxygen production level imposes a required current of 2000 A. An examination of the right-hand plot in Figure 2-18 reveals that these requirements impose a maximum feasible diameter of about 1.4 m, which again corresponds to a maximum electrode separation (discussed further in Section 2.6.4).

All of the cutoff line data (of which only a small portion is shown in Figures 2-15 through 2-18) were fit with multivariate, nonlinear regression equations using MatlabTM. That is, a set of regression equations were created that can predict the expected molten mass (MM) and operating temperature (T_{op}) for a given reactor diameter (D), current (I) and wall thermal conductivity (k_{wall}):

$$MM = f(D, I, k_{wall}) = I^{m_6} m_4 \exp\left(\frac{m_1 (D - m_5)}{I^{m_3} \left(m_2 + \frac{1}{k_{wall} + 0.15}\right)}\right), \quad \text{and} \quad (2.20)$$

$$T_{op} = f(D, I, k_{wall}) = t_1 + I^{t_5} t_2 \exp\left(-\frac{t_3 (D - t_7)}{I^{t_6} \left(1 + \frac{t_4}{k_{wall} + 0.15}\right)}\right), \quad (2.21)$$

where m_i and t_i are regression coefficients which depend upon the type of regolith being electrolyzed and the maximum allowable wall temperature. These coefficients are presented in Tables C.1 and C.2 in Appendix C.

Although the fits in Equations (2.20) and (2.21) are largely empirical, the dependence upon the thermal conductivity of the reactor wall does have some roots in

an analytical approximation presented in Appendix B. Note that, by designing on the cutoff lines, the electrode separation has a one-to-one mapping with the reactor diameter, so it is not a free parameter but rather $\Delta e = f(D, I, k_{wall})$. This is further discussed in Section 2.6.4.

Equations (2.20) and (2.21) can be algebraically manipulated to derive expressions for the minimum and maximum diameter bounds illustrated in Figure 2-18:

$$D_{min} = m_5 + \frac{I^{m_3} \left(\log\left(\frac{MM}{m_4}\right) - m_6 \log(I) \right) \left(m_2 + \frac{1}{k_{wall} + 0.15} \right)}{m_1}, \quad \text{and} \quad (2.22)$$

$$D_{max} = t_7 - \frac{I^{t_6} \log\left(\frac{T_{op} - t_1}{I^{t_5} t_2}\right) \left(\frac{t_4}{k_{wall} + 0.15} + 1 \right)}{t_3} \quad (2.23)$$

It is important to note that this step in the MRE model construction imposes the constraint that the regression equations for the molten mass and operating temperature must be invertible with respect to the diameter (D). If this condition is not met, a closed-form solution for D_{max} and D_{min} cannot be obtained.

2.6.3 Ensuring Operational Flexibility: The Design Margin

In Equations (2.22) and (2.23), the only free design parameter is the thermal conductivity of the wall; T_{op} is set by the user, $MM = m_{regolith}$ from Equation (2.10), and $I = I_{avg}$ from Equation (2.15). The wall conductivity must be chosen such that the diameter bounds are not reversed (i.e. $D_{min} \not\geq D_{max}$). To quantitatively control this requirement, a new variable was introduced, termed the “design margin” (Φ):

$$\Phi \equiv \frac{D_{max}(T_{op}, I, k_{wall})}{D_{min}(MM, I, k_{wall})}, \quad (2.24)$$

which represents the ratio of the maximum diameter bound to the minimum diameter bound. Imposing the constraint that $\Phi \geq 1$ ensures that the diameter bounds are properly ordered and $\Phi = 1$ corresponds to diameter bounds that are equal, in which only a single value for reactor diameter results in a feasible design.

By substituting Equations (2.22) and (2.23) into Equation (2.24), a closed-form expression for the required wall thermal conductivity can be derived (shown in Equation (2.25)) that **1)** satisfies the molten mass and operating temperature constraints **2)** sets the maximum and minimum diameter bounds within a certain range of one another as specified by the design margin, **3)** facilitates the joule-heated, cold-wall condition (no molten regolith touches the reactor wall), and **4)** ensures that the reactor can support the required current:

$$k_{wall} = -\frac{I^{t_6} m_1 t_4 \log\left(\frac{T_{op}-t_1}{I^{t_5} t_2}\right) + I^{m_3} \Phi t_3 \log\left(\frac{MM}{m_4}\right) - I^{m_3} \Phi m_6 t_3 \log(I)}{\beta},$$

$$\text{where } \beta \equiv I^{t_6} m_1 \log\left(\frac{T_{op}-t_1}{I^{t_5} t_2}\right) - m_1 t_3 t_7 + \Phi m_1 m_5 t_3 \dots \quad (2.25)$$

$$\dots + I^{m_3} \Phi m_2 t_3 \log\left(\frac{MM}{m_4}\right) - I^{m_3} \Phi m_2 m_6 t_3 \log(I) - 0.15,$$

where again, m_i and t_i are regression coefficients given in Tables C.1 and C.2 in Appendix C.

To parametrically design an MRE reactor, the first step is to pick a design margin ($\Phi \geq 1$) and then use Equation (2.25) to find the appropriate wall thermal conductivity. As mentioned in Section 2.6.6, although the thermal conductivity of the wall is considered a design variable in this methodology, in the MRE reactor model the desired k_{wall} in Equation (2.25) is translated into a wall thickness with real refractory and insulation materials, as well as a required number of MLI layers, to provide a thermal resistance equivalent to the 3 cm thick wall in the COMSOLTM model with the desired k_{wall} in Equation (2.25).

2.6.4 Designing a Feasible Range of Electrode Separations

By selecting designs on the cutoff lines illustrated in Figure 2-15, there is a one-to-one mapping between reactor diameter and electrode separation, which depends on current and the wall thermal conductivity. The following multivariate, nonlinear

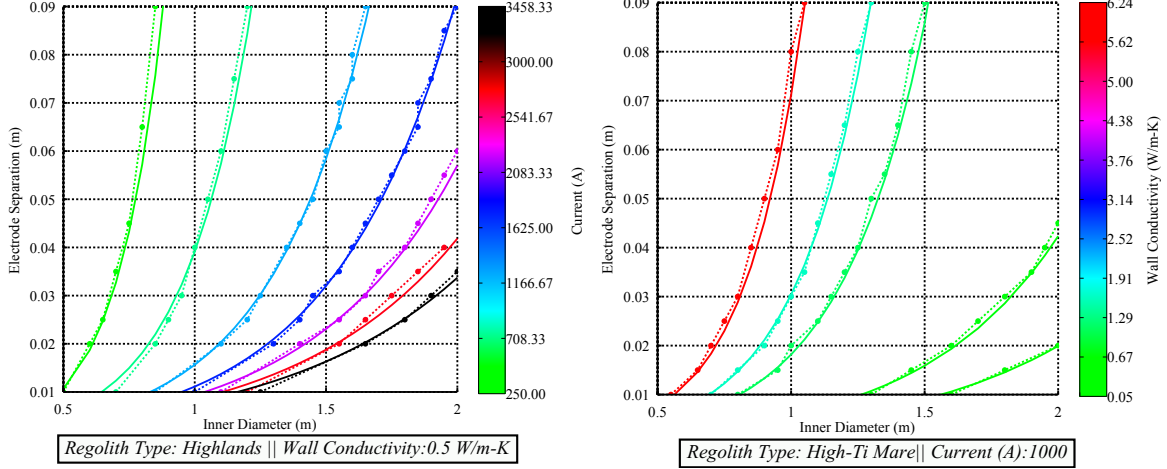


Figure 2-19: The required electrode separation for a given reactor diameter. The electrode curve is affected by current (left) and wall thermal conductivity (right). A regression model, illustrated by the solid line, was fit to the multiphysics data, depicted by the data points connected by dashed lines ($R^2=0.975$, $RMSE=0.0032$ m).

regression equation was fit to the data from the multiphysics simulation:

$$\Delta e = e_1 I^{e_8} (k_{wall})^{e_7} \exp\left(\frac{e_3 (D - e_2)}{I^{e_5} \left(\frac{e_4}{k_{wall}} + 1\right)^{e_6}}\right), \quad (2.26)$$

where e_i are regression coefficients presented in Table C.3 in Appendix C.

Equation (2.26) was fit to the data from the multiphysics simulation to develop a model that predicted the required electrode separation as a function of the other reactor design parameters. This regression equation is plotted in Figure 2-19 against a small sample of the data from the multiphysics simulation. As shown in Equation (2.26) and Figure 2-19, the electrode separation exhibits an exponential increase with reactor diameter. From a practical design standpoint, this means that to scale up an MRE reactor to a larger diameter, the required electrode separation should also increase according to Equation (2.26).

One intuitive finding from the left-hand plot in Figure 2-19 is the relationship between current and electrode separation to achieve joule-heated cold-wall operation. For a fixed reactor diameter, as current increases (moving downwards from the green line towards the black line) the required electrode separation decreases to ensure that the reactor does not generate too much heat and violate the cold-wall constraint.

Similarly, in the right-hand plot we see that as the reactor walls become less thermally conductive (moving from the red line to the green line) the required electrode separation also decreases to avoid generating too much heat in the reactor core.

As presented in Section 2.6.2 the bounds on the reactor diameter translate directly into bounds on the electrode separation as well. This is critical, as in the actual operation of an MRE reactor, one will desire to change the separation between the electrodes to provide control over the voltage and current running through the reactor.

2.6.5 Estimating Reactor Power and Operating Voltage

With the reactor design deterministically set using the scaling equations in the preceding sections, the total power must be calculated. As shown in Figure 2-1, the total power is estimated using the sum of four different terms:

$$P_{total} = \dot{Q}_{heat-loss} + \dot{Q}_{regolith-heatup} + P_{(\Delta G)} + \dot{Q}_{endothermic}. \quad (2.27)$$

Each of these four terms are detailed in the next four sections.

Radiative Heat Loss ($\dot{Q}_{heat-loss}$)

The radiative heat loss to the environment ($\dot{Q}_{heat-loss}$) was measured for each design in the multiphysics simulation. As shown on the left-hand plot in Figure 2-20, the current does not significantly affect the heat loss for a fixed diameter reactor, but this is for a reactor designed on the cutoff lines. Designing on the cutoff lines causes the electrode separation to decrease as current increases to maintain the joule-heated, cold-wall constraint, as shown in Figure 2-19. This decrease in electrode separation appears to work towards maintaining a constant heat loss. From Figure 2-20, we see that for a fixed diameter, the thermal topology required by the joule-heated, cold-wall condition imposes a certain heat loss on the reactor which does not significantly depend on the current.

Conversely, the heat loss does appear to depend significantly on the wall thermal conductivity, with a lower wall thermal conductivity intuitively resulting in a lower

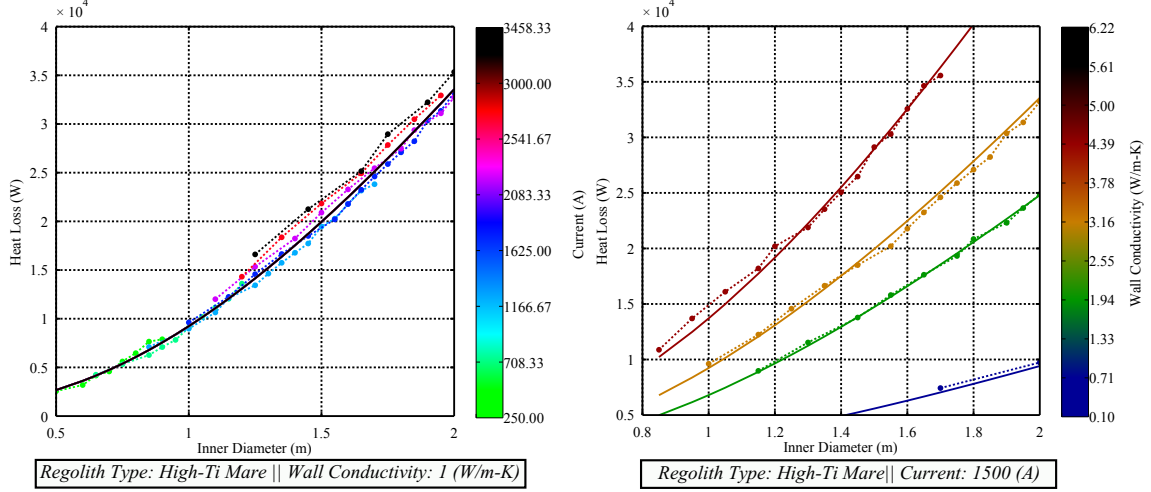


Figure 2-20: The heat loss does not significantly depend on current (left), but is affected by wall thermal conductivity (right) when designing on the cutoff lines. A nonlinear regression model (solid lines) was fit to the data from a multiphysics simulation (dots with dashed lines) with $R^2=0.997$, $RMSE=0.51kW$.

heat loss. The observed decrease in heat loss with decreasing wall thermal conductivity is a direct result of the decrease in electrode separation observed in Figure 2-19.

The heat loss data from the multiphysics simulation were fit with a multivariate, nonlinear regression model as shown in Figure 2-20 ($R^2=0.997$, $RMSE=0.51kW$). The dashed line and data points in Figure 2-20 show data generated by the multiphysics simulation and the solid lines show the predictions of the regression models:

$$\dot{Q}_{heat\ loss} = h_1 \left(\frac{\exp\left(\frac{h_2}{D+h_3}\right)}{h_4 + \frac{1}{k_{wall}-h_6}} \right) + h_5 k_{wall} \quad (2.28)$$

where h_i are regression coefficients given in Table C.4 in Appendix C. These coefficients depend upon the Regolith type and the maximum wall cutoff temperature.

Regolith Heating Power ($\dot{Q}_{regolith-heatup}$)

The power required to heat the regolith from ambient temperature ($\sim 300K$) to the operating temperature, is calculated as:

$$\dot{Q}_{regolith-heatup} = \dot{m}_{regolith} \left(\int_{300K}^{T_{op}} C_p(T) dT + L \right), \quad (2.29)$$

where $\dot{m}_{regolith}$ is the average regolith flow rate into the reactor, $C_p(T)$ (Section 2.2.3), and L is the latent heat of melting (Section 2.2.9). It is important to note that in contrast to Equation (2.19), which describes the regolith heating that must occur in the core of the reactor, there is no factor of $\frac{1}{2}$ in front of the integral in Equation (2.29). For the purposes of thermally modeling the reactor in the multiphysics simulation, our model assumes that a preheater will provide half of the heating power prior to inserting fresh regolith into the reactor and that the reactor itself must supply the rest of the heating requirement via joule-heating in the reactor core. For the purposes of total power accounting, the entirety of the regolith heating requirement must be included.

Endothermic Makeup Heat ($\dot{Q}_{endothermic}$)

The heat required to offset the endothermic electrolysis reaction is related to the difference between the Enthalpy of Formation and Gibbs Free Energy ($\Delta H - \Delta G$). Although the calculation for the endothermic heating power is presented in Equation 2.18, it is shown again here for completeness:

$$\dot{Q}_{endothermic} \equiv (\dot{n}_{O_2}) \text{ avg}[\Delta H - \Delta G] = (\dot{n}_{O_2}) \sum_i (\Delta H_i(T) - \Delta G_i(T)) \left(\frac{c_i(T)}{C} \right), \quad (2.30)$$

where \dot{n}_{O_2} is the average molar oxygen production rate (mols/sec), $c_{specie,i}(T)$ is the number of coulombs used to electrolyze oxide specie i during a single batch, and C is the total number of coulombs used in a single batch from Equation (2.14).

Chemical Electrolysis Power ($P_{\Delta G}$)

The electrical power required to perform the electrolysis chemical reaction is directly related to the Gibbs Free Energy (ΔG):

$$P_{\Delta G} = (\dot{n}_{O_2}) \text{ avg}[\Delta G], \quad (2.31)$$

where $avg[\Delta G]$ is the current-averaged Gibbs Free Energy value over all species electrolyzed and \dot{n}_{O_2} is the average oxygen molar production rate (mol/s).

Operating Voltage

The operating voltage was also evaluated for each design generated by the multi-physics simulation. Equation (2.32) shows the regression equation used to fit the voltage data ($R^2=0.987$, $RMSE=0.64V$).

$$V = \frac{1}{I} \left(\frac{v_1 \exp\left(\frac{v_2}{D+v_3}\right)}{v_4 + \frac{1}{k_{wall}-v_6}} + v_5 k_{wall} \right), \quad (2.32)$$

where v_i are regression coefficients presented in Table C.5 in Appendix C, which depend upon Regolith type and the maximum wall cutoff temperature. The form for this equation was taken directly from the regression equation for heat loss (Equation (2.28)), with the addition of the term $\frac{1}{I}$ to incorporate the expected relation of $\dot{Q}_{heat\ loss} \propto IV$ for joule heating.

2.6.6 Estimating Reactor Mass

Using the methodology described in Section 2.6.2, a reactor design was generated that satisfied the oxygen production level, the target operating temperature, and the joule-heated cold-wall constraint. This section describes how the mass of a given reactor design is estimated in the MRE reactor model.

Determining Wall Thickness and MLI Design

Although the desired reactor wall thermal conductivity is calculated in Equation (2.25), in the design of a real MRE reactor, the wall material and thickness are design parameters rather than wall thermal conductivity. In the MRE reactor model, the insulation and refractory materials are set by user input. A thickness ratio of $\lambda:1$ is used (insulation:refractory) to ensure that sufficient refractory material is used, while relying on the less dense insulation material to provide the bulk of the thermal insulation. For

this work $\lambda = 3$ was used on the reactor bottom and side walls and $\lambda = 6$ was used on the reactor top, as the contact with the molten regolith is assumed to be minimal on the roof of the reactor. With the thickness ratio, an effective thermal conductivity ($k_{effective}$) for the combined two layers can be calculated:

$$k_{effective} = \left(\left(\frac{\lambda}{(\lambda + 1)k_{insulation}} \right) + \left(\frac{1}{(\lambda + 1)k_{refractory}} \right) \right)^{-1}, \quad (2.33)$$

where $k_{refractory}$ and $k_{insulation}$ are the thermal conductivities of the refractory and insulation layers, respectively. Note that Equation (2.33) does not include the thermal resistance of the outer structural shell, as it is negligible compared to the insulation and refractory layers.

To translate the thermal design dictated by the COMSOL simulation into a real reactor design, the total thermal resistance of the COMSOL reactor ($R_{th-total}$), from the inner reactor wall to the ambient temperature of 298.15K, needed to be calculated:

$$R_{th-total} = R_{th-cylinder} + R_{th-radiation}, \quad (2.34)$$

where $R_{th-cylinder}$ is the thermal resistance from the inner reactor wall to the outer reactor wall and $R_{th-radiation}$ is the thermal resistance associated with radiating heat from the reactor outer surface through layers of MLI to a distant ambient temperature.

The thermal resistance through the wall of a three-dimensional cylinder is given by including the cylindrical side wall and the top and bottom walls in parallel [65]:

$$R_{th-cylinder} = \frac{1}{\pi k_{effective}} \left(\frac{(2(D_{in}/2 + \Delta x)^2)}{\Delta x} + \frac{2(H_{in} + 2\Delta x)}{\ln(1 + 2\Delta x/D_{in})} \right)^{-1}, \quad (2.35)$$

where Δx is the thickness of the insulation and refractory layers, D_{in} is the inner diameter, and H_{in} is the inner height of the cylinder.

The thermal resistance from the reactor outer wall through the layers of MLI to a distant ambient temperature is given by [65]:

$$R_{th-radiation} = (\epsilon^* \sigma (T_{out}^2 + T_{\infty}^2) (T_{out} + T_{\infty}) A)^{-1}, \quad (2.36)$$

where T_{out} is the temperature on the outer surface of the reactor (inside of the MLI layers), T_∞ is the ambient temperature, taken to be 298K, A is the outer surface area of the reactor, σ is the Stefan-Boltzmann constant, and ϵ^* is the effective emissivity of the layers of MLI, given by [42]:

$$\epsilon^* = \left(\frac{2N}{\epsilon_{MLI}} - N - 1 + \frac{1}{\epsilon_1} + \frac{1}{\epsilon_\infty} \right)^{-1}, \quad (2.37)$$

where N is the number of layers of MLI, ϵ_{MLI} was taken to be 0.03 for Mylar [42], ϵ_1 is the emissivity of the reactor outer surface (taken to be 1.0), ϵ_∞ is the emissivity of free space (taken to be 1.0). Although Gilmore and Donabedian [42] assert that the theoretical effective emissivity in Equation (2.37) is a dramatic underestimate, recent work by Hatakenaka et al. [46] demonstrated a new pin attachment method for MLI that allows for real hardware performance to approach the theoretical predictions of Equation (2.37). In fact, the James Webb telescope shield achieved an effective emissivity of 0.00038 [81].

The temperature on the outer surface of the reactor inside of the MLI layers (T_{out}) was calculated using the knowledge that all of the heat loss ($\dot{Q}_{heat-loss}$) calculated in Equation 2.6.5 must be radiated out of the reactor:

$$T_{out} = \left(\frac{\dot{Q}_{heat-loss} + A\sigma\epsilon_\infty(T_\infty)^4}{A\sigma\epsilon^*} \right)^{1/4} \quad (2.38)$$

Using Equations (2.33) through (2.38), the thermal resistance of the reactor in the multiphysics simulation could be calculated, using the required wall thermal conductivity from Equation (2.25) for wall thermal conductivity ($k_{effective}$), 3 cm for the wall thickness (Δx), and 0.05 for the outer wall emissivity (ϵ^*). With these values, a required thermal resistance could be calculated ($R_{th-required}$), which essentially describes the required thermal properties that a reactor needs to have in order to behave thermally similar to the reactor in the multiphysics simulation.

With the required thermal resistance as a design target, the MRE reactor model must then design a reactor with the same thermal resistance despite the fact that the

wall thermal conductivity ($k_{effective}$) is now based on real materials and will therefore be different from the value prescribed by Equation (2.25).

A previous version of the MRE reactor model [95] varied the insulation and refractory thickness (Δx) to achieve the required thermal resistance, while leaving the number of the MLI layers constant. Although this is a viable way to design MRE reactors, it incurs heavy mass penalties, as typical refractory materials have densities on the order of 3 g/cm³. Furthermore, the effect of adding more insulation to the reactor does not always increase the thermal resistance, as it increases the area for radiation to ambient [65].

To enable less massive, more robust reactor designs, a new methodology is proposed. This methodology starts with minimum wall thickness, $\Delta x = \Delta x_{min}$, set as a user input. The MIT experimental hardware used a two-layer crucible with a wall thickness of 1 cm each [43], so a minimum wall thickness of 2 cm was used for this study. For the minimal wall thickness, layers of MLI are added in an attempt to meet the insulation requirements set forth from the multiphysics model. If this fails to achieve the required thermal resistance (the effect of adding layers of MLI has diminishing returns), the wall thickness is increased to attempt to achieve the required thermal resistance. Increasing the wall thickness also has diminishing returns, and so the reactor model will throw a flag if it cannot produce a reactor design with the required insulation properties as discussed further in Section 2.6.7. The mass of the MLI was taken to be 0.02 kg/m² per layer [80] with a factor of two to be conservative and account for fasteners and other supporting material. This number is based off of an actual flight design at the Marshall Space Flight Center (MSFC) [80].

Although this thesis was not intended to focus on the material science factors in the design of an MRE reactor, some survey work was completed concerning appropriate refractory materials to gain a greater understanding of the approximate density and thermal conductivity one could expect out of such materials. The high purity alumina used in the MIT experiments was avoided due to its high density of 3.9 g/cm³. Representatives at Allied Mineral¹ were contacted in order to determine

¹<http://www.alliedmineral.com/>

if suitable alternatives with a lower density existed. A family of refractory materials with densities on the order of 3.0 g/cm^3 and thermal conductivities on the order of 1.5 W/m-K were identified. This model assumes the use of a similar material for the refractory layer. For the insulation material, Aeroguard HD^{TM 2} was selected for its excellent thermal properties at high temperatures and its high thermal stability. Furthermore, it has a low density on the order of 0.2 g/cm^3 , which is one reason it is commonly used in the aviation industry.

Structural, Insulation, and Refractory Mass

The structural thickness is set such that the hoop stress resulting from the operating pressure is less than the yield stress of the material with a safety factor (SF) of 3.5 ($\sigma_y = (SF) \left(\frac{Pr}{2t}\right)$). A constant ullage was used in all multiphysics simulations, which results in an inner reactor height that was approximately 1.5 times the maximum electrode separation. Due to the fact that the structure may encounter elevated temperatures, titanium was selected over aluminum. With the diameter, height, and thickness of the structural, insulation, and refractory layers determined, the mass of the cylindrical reactor shell can be calculated.

Anode and Cathode Mass

The mass of the anode and cathode are also taken into account in the MRE reactor model. The diameter of the cathode plate is fixed at 15% of the reactor diameter, after an initial analysis revealed that this ratio provided a good balance between larger molten masses and higher operating temperatures. The anode diameter is designed to be 19% less than the cathode diameter, using the same ratio as Sibille and Dominguez [99]. The cathode shaft length is long enough to pass through the bottom reactor wall and then connect to the power system. The anode shaft length is set to extend from the anode plate all the way out of the top of the reactor. The diameter of both the anode and cathode shafts is set to 2 cm plus $\frac{1}{25}$ th of the anode plate diameter. With these dimensions set, the mass of the anode and cathode are relatively easy to

²http://www.microthermgroup.com/landingpage/assets/TDS_AEROGUARD_V1-EN.pdf

calculate. Future work can address modeling the anode and cathode connections to the power system to provide even more accurate mass and power estimates.

Heater Mass

As mentioned in Sections 2.5.2 and 2.6.5, half of the regolith heating power occurs via joule-heating in the reactor core and half is provided by a preheater that bakes out volatiles and warms up the cold fresh regolith. The mass of this preheater was calculated using trends from a handbook on electric resistive heaters [117].

2.6.7 Reactor Model Flags

A number of flags in the reactor model serve to identify designs that may appear feasible, but due to concerns not directly modeled may not be achievable. The following sections detail each error flag that the MRE reactor model can throw. As discussed later in Section 3.4, these error flags are used in the optimization scheme to guide the reactor model away from infeasible designs to feasible design regions.

Electrode Separation Constraints

There are two constraints on the minimum and maximum electrode separation. These values can be calculated using Equation (2.26) when the design margin is > 1.0 , which will result in a range of acceptable reactor diameters and therefore a range of acceptable electrode separations. The average of the maximum and minimum electrode separation must be greater than or equal to 2.0cm, otherwise a flag is thrown. This value comes from two experiments performed at MIT that used electrode separations of 2.0 cm [62] and 2.22 cm [43]. If the electrodes are too close together, the metal dendrites that form [30, 62] may cause a short.

Furthermore, a range of 0.5 cm between the minimum and maximum electrode separation bounds was also imposed to provide a measure of operational flexibility. If the two bounds are closer than 0.5 cm, the reactor model throws a flag to indicate that the given design does not have enough operational flexibility.

Current Density Constraint

There is also a constraint on the current density. Because this value is not set explicitly within the design process, the reactor model throws a flag if the proposed reactor design has too high of a current density. Current density is simply taken on the anode surface by dividing current by the lower surface area of the anode plate. The appropriate limit for the current density of lunar regolith is somewhat uncertain. Some of the earliest work developing MRE achieved a current density of 1.25 A/cm² [20], but a large range of values have been achieved in the literature, as presented in Table 2.6.

Table 2.6: Experimental current density values from a range of molten regolith electrolysis experiments (with the exception of Kennedy [60]).

Source	Current Density (A/cm ²)	Anode Material	Notes/Comments
Kesterke [61]	15.5	Iridium	<i>At 17-19V</i>
Carroll [20]	1.25	Molybdenum	
Haskin et al. [45]	2.5	Platinum	<i>1560° C</i>
Gmitter [43]	0.641	Ir, Pt, Rh	<i>Increases w/ optical basicity</i>
Paramore [79]	7.35	Iridium	<i>Possibly 0.3 $\frac{A}{cm^2}$, y-axis unclear</i>
Sirk et al. [103]	0.8	Iridium	<i>At 0.2 A</i>
Wang et al. [124]	1.19	Iridium	<i>Increases w/ voltage</i>
Kim et al. [62]	11.7	Iridium	
Kennedy [60]	4 - 7	Graphite	<i>Electric slag furnace</i>
Allanore et al. [2]	2	FeCr _X	<i>At 3.8 V</i>
Allanore et al. [2]	5	FeCr _X	<i>At 6.0 V</i>

From Table 2.6, it is clear that the current density depends on the type of anode being used as well as the temperature, operating voltage and melt composition. Simulation work by Sibille and Dominguez [99] demonstrated a 33% increase in the current density (current divided by anode area facing the cathode) when using a “waffle”-shaped anode. Complementary experimental work by Curreri et al. [30] revealed a three-fold increase in current when using a coiled piece of metal for the anode rather than a simple plate. Voltages significantly above the Nernst potential are also expected to drive the current density higher, as demonstrated by Allanore et al. [2]. If we conservatively assume that more complex anode designs can result in a two-fold increase in current density over existing flat-plate experimental work, we can expect

current densities on the order of 20 A/cm^2 ($\approx 2 \times 11.7 \text{ A/cm}^2$ from Kim et al. [62]). To avoid violating this constraint, the reactor model will throw a flag if the current density surpasses 20 A/cm^2 .

Wall Thickness Constraint

Another constraint in the MRE reactor model concerns the thickness of the wall. As described in Section 2.6.6, the wall thickness is set such that the real reactor has the same thermal resistance as the one prescribed by the regression equations. If the wall thermal conductivity prescribed by Equation (2.25) is too high, the equivalent wall thickness in the real reactor may become too thin. An input was added to the model to prescribe the minimum allowable wall thickness. For this study a baseline of 2 cm was used, as the experiments at MIT utilized two 1 cm thick crucibles. If the required wall thickness drops below this value, the reactor model will throw a flag.

Wall Insulation Constraint

On the other end of the design space, the wall thermal conductivity prescribed by Equation (2.25) is too low, it may not be achievable by the materials chosen for the real reactor. Say that the Equation (2.25) prescribes a conductivity of $1\text{E-}5 \text{ W/m-K}$ and the effective thermal conductivity (from Equation (2.33)) of the real insulation and refractory materials is 0.1 W/m-K . Increasing the thickness of the real reactor wall beyond the 3 cm thickness from the COMSOL simulation will help raise the thermal resistance of the real reactor closer to the required value, but this increase is not monotonically increasing - after a certain point adding more wall thickness actually *decreases* the total thermal resistance of the cylinder. This is because adding more wall thickness increases the outer surface area of the reactor, which decreases the thermal resistance of the outer surface radiation given in Equation (2.36).

As detailed in Section 2.6.6, additional layers of MLI can also be added to attempt to achieve the required thermal resistance, but the effect of adding more layers also approaches a horizontal asymptote (see Equation (2.37)). Thus, there is a limit to the thermal resistance that can be achieved with real materials - if the thermal resistance

required is too high, the reactor model throws a flag.

Molten Mass Constraint

The required molten mass is calculated using Equation (2.11), but there are constraints on this value. Due to the fact that the regression equation for the molten mass in a reactor, given in Equation (2.20), is an exponential function without an intercept shift, it is actually possible for this equation to predict positive molten mass estimates when the diameter is zero (or even less than zero). Although alternative equations were explored, in the end the form shown in Equation (2.20) was found to fit the data best. To ensure that the reactor model does not predict negative reactor diameters, the molten mass is required to be above a minimum value such that $D_{min}(MM_{min}) > 0$ where $D_{min}(MM)$ is given by Equation (2.22). If the required molten mass is less than this limit, the reactor model throws a flag to indicate that the batch time must be increased to result in a suitable amount of molten mass within the reactor.

2.7 MRE Reactor Performance and Design Trends

2.7.1 Oxygen Extraction Efficiency and Current Efficiency

The oxygen extraction efficiency (η_{O_2}), calculated using Equation (2.10), describes the fraction of oxygen that can be extracted from lunar regolith by weight (i.e. the number of kilograms of oxygen that can be extracted per kilogram regolith). The left-hand plot in Figure 2-21 shows how η_{O_2} increases with operating temperature for three types of lunar regolith discussed in Section 2.2.1. The maximum value for this metric is 0.454 for Highlands regolith, 0.417 for High-Ti Mare regolith, and 0.423 for Low-Ti Mare regolith.

As discussed in Section 2.3, higher operating temperatures allow one to reduce more of the oxides in lunar regolith before the melting temperature of the molten regolith core approaches the operating temperature, causing the core to risk solidi-

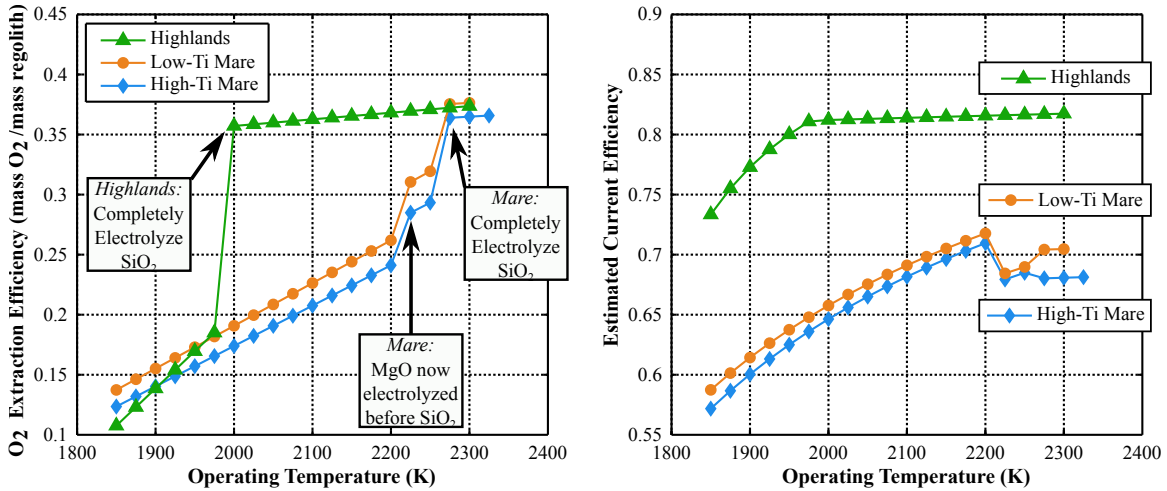


Figure 2-21: The oxygen extraction efficiency (left) and current efficiency (right) estimates for an MRE reactor.

fication and reactor shutdown. Thus by raising the operating temperature, one can extract more oxygen per kilogram regolith and increase η_{O_2} . Density considerations come into play as well, as one must ensure that the product metal does not become less dense than the leftover slag and rise to the top of the molten core, shorting out the reactor. Future work can further investigate this design consideration.

There are a few interesting differences between the Highlands and Mare regolith data in Figure 2-21. For Highlands regolith, η_{O_2} exhibits a sharp rise around 2000K, which corresponds to an operating temperature that is high enough to allow one to reduce all of the SiO₂ (see Figure 2-10: the SiO₂ peak is around 1950K + 50K safety margin = 2000K). For Mare regolith, a jump in η_{O_2} occurs around 2200K. This jump is not due to electrolyzing all of the SiO₂, but is rather due to a reversal in the electrolysis order of SiO₂ and MgO. As seen in Figure 2.2.8, around 2200K the Gibbs Free Energy for MgO rises above that of SiO₂, meaning that MgO will be preferentially electrolyzed over SiO₂ at temperature above 2200K. In reality, it is likely that both MgO and SiO₂ will be concurrently reduced at temperatures near 2200K, but for the MRE reactor model all oxide species are assumed to be discretely reduced in a specific order for simplicity.

The current efficiency (η_I) is shown in the right-hand plot in Figure 2-21 as a function of temperature for three different types of regolith. For all types of regolith,

the current efficiency starts at a lower value due to the relatively large amount of FeO being electrolyzed with a current efficiency of 30% (see Section 2.2.7). As operating temperature increases, more SiO₂ (and other species) are able to be electrolyzed with a higher current efficiency, increasing the average current efficiency over all species electrolyzed. For Highlands regolith, we observe a smooth increase from 74% at an operating temperature of 1850K to 82% above 2000K. For Mare regolith a similar increase is observed from 57% at 1850K to 72% at 2200K. As discussed in preceding paragraph, at temperature above 2200K MgO is reduced before SiO₂. Due to the fact that MgO is electrolyzed with an estimated efficiency of 50% (due to the possible recombination of gaseous Mg and O₂ as discussed in Section 2.2.7), this results in a drop in average current efficiency. It is important to note that Highlands regolith can be electrolyzed with a significantly higher current efficiency due to the fact that it has less FeO.

One may note that the current efficiency estimates presented here appear lower than many numbers cited in the literature, which are often around 90-100% [100, 121, 103]. These estimates are for melts without iron-, titanium-, or sodium-oxides, which is why the current efficiency is dramatically higher. Realistic current efficiencies over an entire batch, in which FeO, SiO₂, TiO₂, and other oxides may be reduced, will be lower.

2.7.2 Mass and Power Estimates

The MRE reactor model developed in this work was leveraged to generate estimates of reactor mass and power over a range of oxygen production levels from 1,000 kg O₂/year to 10,000 kg O₂/year. The effects of a number of design variables on reactor mass and power are explored in the following section.

The Effect of Operating Temperature

Figure 2-22 shows the trends in the reactor mass (top left), specific mass (top right), power (bottom left) and specific power (bottom right) over a range of oxygen pro-




Figure 2-22: The mass (top left), power (top right), specific mass (bottom left) and specific power (bottom right) of an MRE reactor over a range of oxygen production levels. Each line represents a different operating temperature and all designs have a margin of 1.5.

duction levels for High-Ti Mare regolith. From this figure, it is clear that, for a given oxygen production level, increasing the operating temperature decreases reactor mass. This is because higher operating temperatures increase the oxygen extraction efficiency (see Figure 2-21), which reduces the regolith processing requirement and reactor size. Prior to this work it was unclear whether higher operating temperatures would indeed result in smaller reactor sizes. Higher operating temperatures might have increased reactor size due to the need to accommodate a larger temperature drop radially before reaching the reactor wall (or perhaps a steeper thermal gradient resulting in more power loss). Furthermore, the density of molten regolith decreases with temperature (see Section 2.2.2), so one might expect higher temperatures to result in larger reactors to hold the same mass of regolith. There are numerous other

effects of increasing the operating temperature (as demonstrated by the temperature dependence of almost all regolith material properties in Section 2.2), but by integrating all of these effects into a single model, we are able to show that higher temperatures significantly reduce reactor mass.

Reactor power appears to slightly increase with operating temperature between 1850K-2150K, and then decrease significantly between 2150K and 2300K. This upper range corresponds to the temperature that allows one to electrolyze significantly more oxygen per kilogram regolith. The model shows that, for High-Ti Mare regolith, reactor power slightly increases with operating temperature until the temperature exceeds $\sim 2200\text{K}$, at which point dramatic power savings can occur. The reactor specific power monotonically decreases with production level, which indicates that from a power perspective, higher production levels are preferable.

Operating temperature also affects the oxygen production range achievable by a single reactor. As operating temperature increases, the minimum oxygen production level increases. We note that the trend does not hold between 2150K and 2300K, which corresponds to the point at which one is able to electrolyze significantly more oxygen per kilogram regolith as shown in Figure 2-10. It appears that crossing over this threshold enables one to design MRE reactors for smaller production levels due to the dramatic increase in oxygen extraction efficiency. The lower bound on oxygen production exists because below a certain oxygen production threshold, the reactor size becomes too small to maintain the joule-heated, cold-wall thermal design (discussed in Section 2.6) with suitable wall thickness. For instance, when the reactor model is used to design a reactor that produces 2000 kg O_2/year with an operating temperature of 2150K (yellow line), it requires a wall thickness of 1.5 mm, which is obviously not acceptable from a reactor robustness standpoint. Further mass savings may occur at temperatures above 2300K, but the validity of the reactor model at these higher temperatures is less certain due to material melting point considerations for the anode.

The Effect of Design Margin

The trends in operating temperature presented above are for reactor designs with a fixed design margin of 1.5. This means that the wall thermal conductivity is designed such that the maximum diameter set by the operating temperature is 1.5 times the minimum diameter set by the molten mass requirement (see Section 2.6.3). The design margin also affects reactor mass and power. As seen in Figure 2-23, increasing the design margin results in higher reactor mass and power, though it does increase the maximum production level for a single reactor.

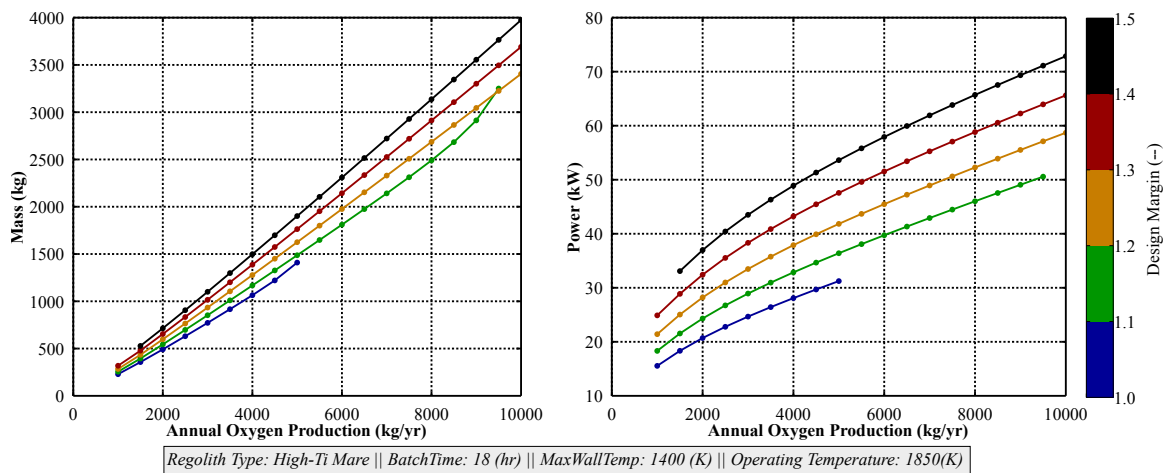


Figure 2-23: The mass and power of an MRE reactor at an operating temperature of 1850K, where each line shows a different design margin. Larger design margins clearly increase both mass and power, but enable higher production levels.

Perhaps the most important effect of design margin is on the acceptable range of electrode separation values. The design margin controls the acceptable bounds on reactor diameter and because there is a one-to-one mapping between reactor diameter and electrode separation, this also translates into bound on the electrode separation (see Section 2.6.4). The left-hand plot in Figure 2-24 demonstrates how increasing the design margin enables a range of electrode separations that **1)** generate enough heat to maintain the central molten core while **2)** not generating excessive heat to maintain the joule-heated cold-wall condition (no molten material touches the reactor wall) as described in Section 1.2.1. This is of the utmost importance, as having an acceptable range of electrode separation values will enable a control system to actuate

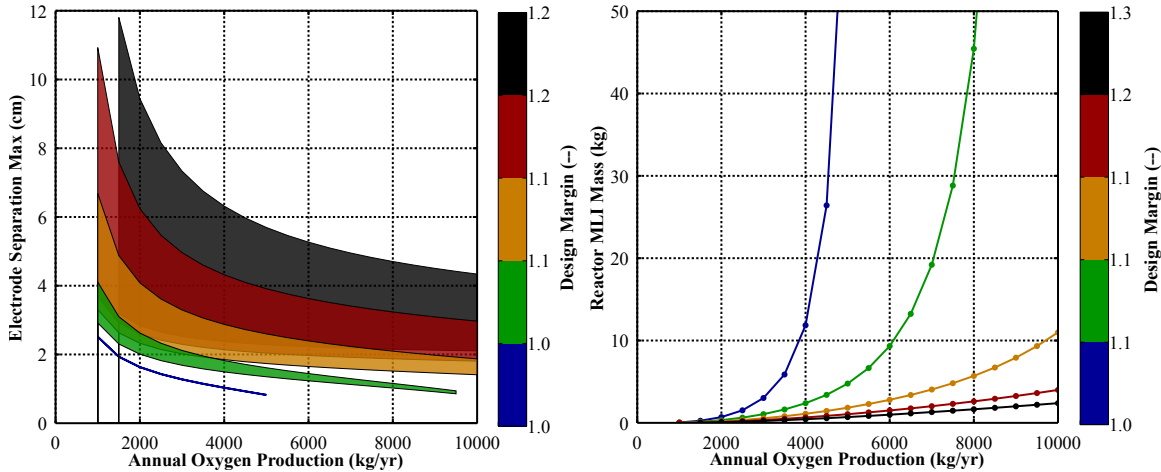


Figure 2-24: (left) Increasing the design margin opens up a larger range of feasible electrode separation values to enable operational flexibility. (right) Increasing the design margin requires less insulation and decreases the number of layers and mass of MLI on the reactor exterior.

the electrode separation to better control the reactor thermal topology during actual operation.

The right-hand plot in Figure 2-24 illustrates how increasing the design margin opens up larger production levels. For a design margin of 1.0, around a production level of 4500 kg O₂/year, the reactor wall must have more and more insulation in the form of more layers of MLI on the exterior of the reactor (see Section 2.6.6). At a certain point the reactor cannot be insulated well enough to satisfy the given design margin, so one must relax the design margin. This results in larger reactors with more conductive walls that are now feasible to design.

The Effect of Maximum Wall Temperature

Identifying infeasible designs plays a central role in determining the cutoff line location/equation. The data in Figure 2-16 crosses out infeasible designs that have a wall temperature above 1400K, to provide a 100K safety margin for the melting temperature of regolith at 1500K. This section evaluates three different maximum wall temperature cutoff values: 1200K, 1300K, and 1400K. For each value, the multivariate regression analysis described in Section 2.6 was carried out to generate reactor scaling equations. The form of the regression equations remained the same, but the

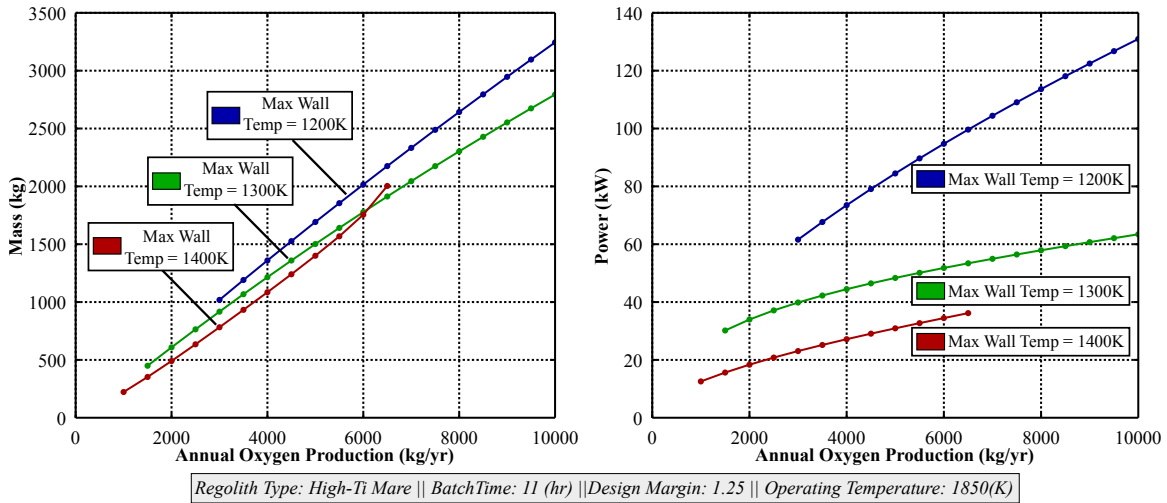


Figure 2-25: The mass and power of an MRE reactor for different maximum wall temperatures. Increasing the maximum wall temperature definitively decreases reactor power (right-hand plot), but has a complex effect on reactor mass.

regression coefficients were different, as shown in Appendix C. Figure 2-25 shows how the maximum wall cutoff temperature affects the mass and power of an MRE reactor.

As seen in Figure 2-25, a higher maximum wall cutoff temperature decreases the production range of a single reactor and generally decreases reactor mass and power. The change in production range is due to two factors. At low production levels, a higher wall temperature allows for smaller reactors because the reactor walls can get closer to central molten core. At higher production levels, higher wall temperatures require more insulation and after a certain production level the reactor can no longer be properly insulated (see Section 2.6.7).

To better understand why higher wall temperatures reduce mass and power, the reactor geometry was investigated. Figure 2-26 shows how lower wall temperatures push the reactor diameter slightly larger. This is simply due to the wall needing to be farther from the high temperature molten core. Due to the exponential relationship between reactor diameter and the required electrode separation (see Section 2.6.4), this slight increase in reactor diameter translates into a significantly larger electrode separation and therefore reactor height.

To further understand the trends in reactor mass and power observed in Figure 2-25, the breakdown of the mass and power was also examined, as shown in Figure 2-27.

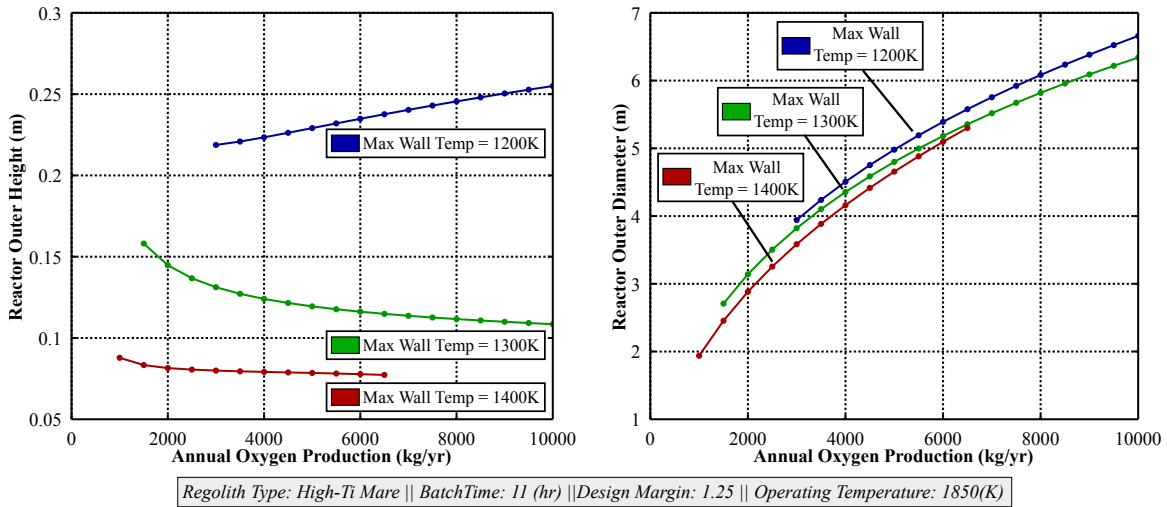


Figure 2-26: The height (left) and diameter (right) of an MRE reactor as a function of oxygen production level, for three different wall temperatures.

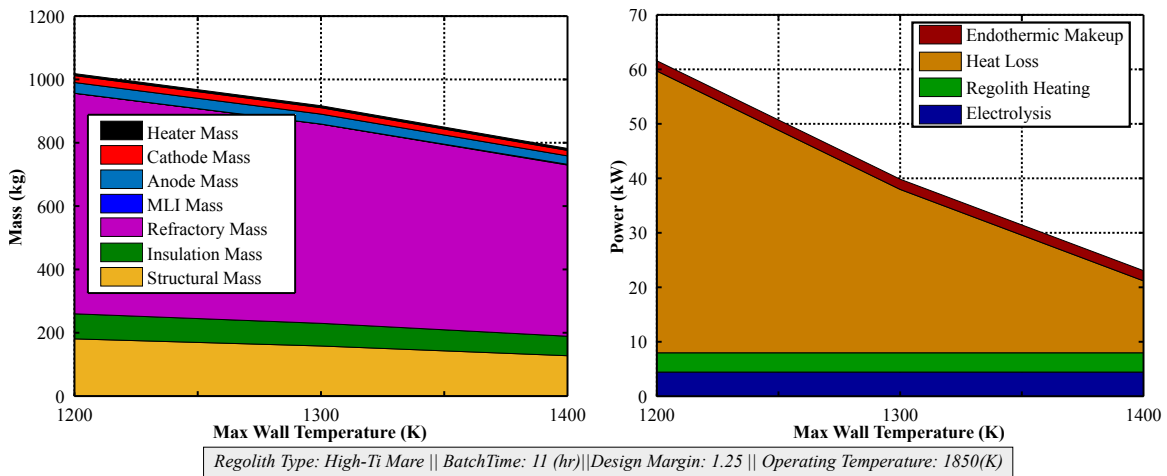


Figure 2-27: The mass and power breakdown as a function of the wall temperature.

Two clear trends in reactor mass appear: higher wall temperatures result in a significant reduction in the mass of the cylindrical shell (structural+insulation+refractory) due to the decrease in reactor height and diameter seen in Figure 2-26. Although almost imperceptible, there is a slight increase in the mass of the MLI around the reactor at higher wall temperatures, corresponding to the requirement for more insulation. The higher wall temperature also results in a dramatic decrease in the heat loss through the reactor wall, as shown on the right-hand plot in Figure 2-27.

The Effect of Regolith Type/Composition

The MRE reactor model has a strong dependence on regolith type. Of the regolith material properties listed in Section 2.2, composition, density, specific heat, electrical conductivity, current efficiency, latent heat of melting, endothermic power, and electrolysis power all depend upon regolith type. Furthermore, sets of multiphysics data were generated for Highlands and Mare regolith (the differences between High-Ti and Low-Ti Mare are not significant in the multiphysics simulation). From these data sets, two sets of regression coefficients were generated for the equations presented in Section 2.6, as shown in Appendix C.

Figure 2-28 shows how the regolith type affects the mass and power of an MRE reactor. The operating temperature plays a large role in determining how well a reactor can process a given type of regolith, so data is shown for operating temperatures of 1850K (top row), 2000K (middle row), and 2300K (bottom row). From these results, it appears that the liquidus temperature curves shown in Figure 2-10 play a large role in the regolith type dependence. That is, the type of regolith strongly affects how much oxygen can be extracted from lunar regolith before having to stop to avoid solidification of the inner molten core.

At 1850K (top row of Figure 2-28), one can extract less oxygen per kilogram Highlands regolith compared to Mare regolith so Highlands reactors must process more regolith, which increases the power required to meet a given oxygen production level. Furthermore, Highlands regolith has more SiO_2 , which must be electrolyzed at a higher voltage and subsequently drives the power requirement higher for Highlands regolith. A secondary effect appears to result in less massive reactors for Highlands regolith at higher production levels. This is due to a decrease in the required insulation for Highlands reactors compared to Mare reactors, as determined by the multiphysics simulation scaling equations.

The trend in liquidus temperature reverses around 2000K (middle row of Figure 2-28), at which point one can now electrolyze more oxygen from Highlands regolith compared to Mare. This results in a dramatic decrease in reactor mass (33%) and

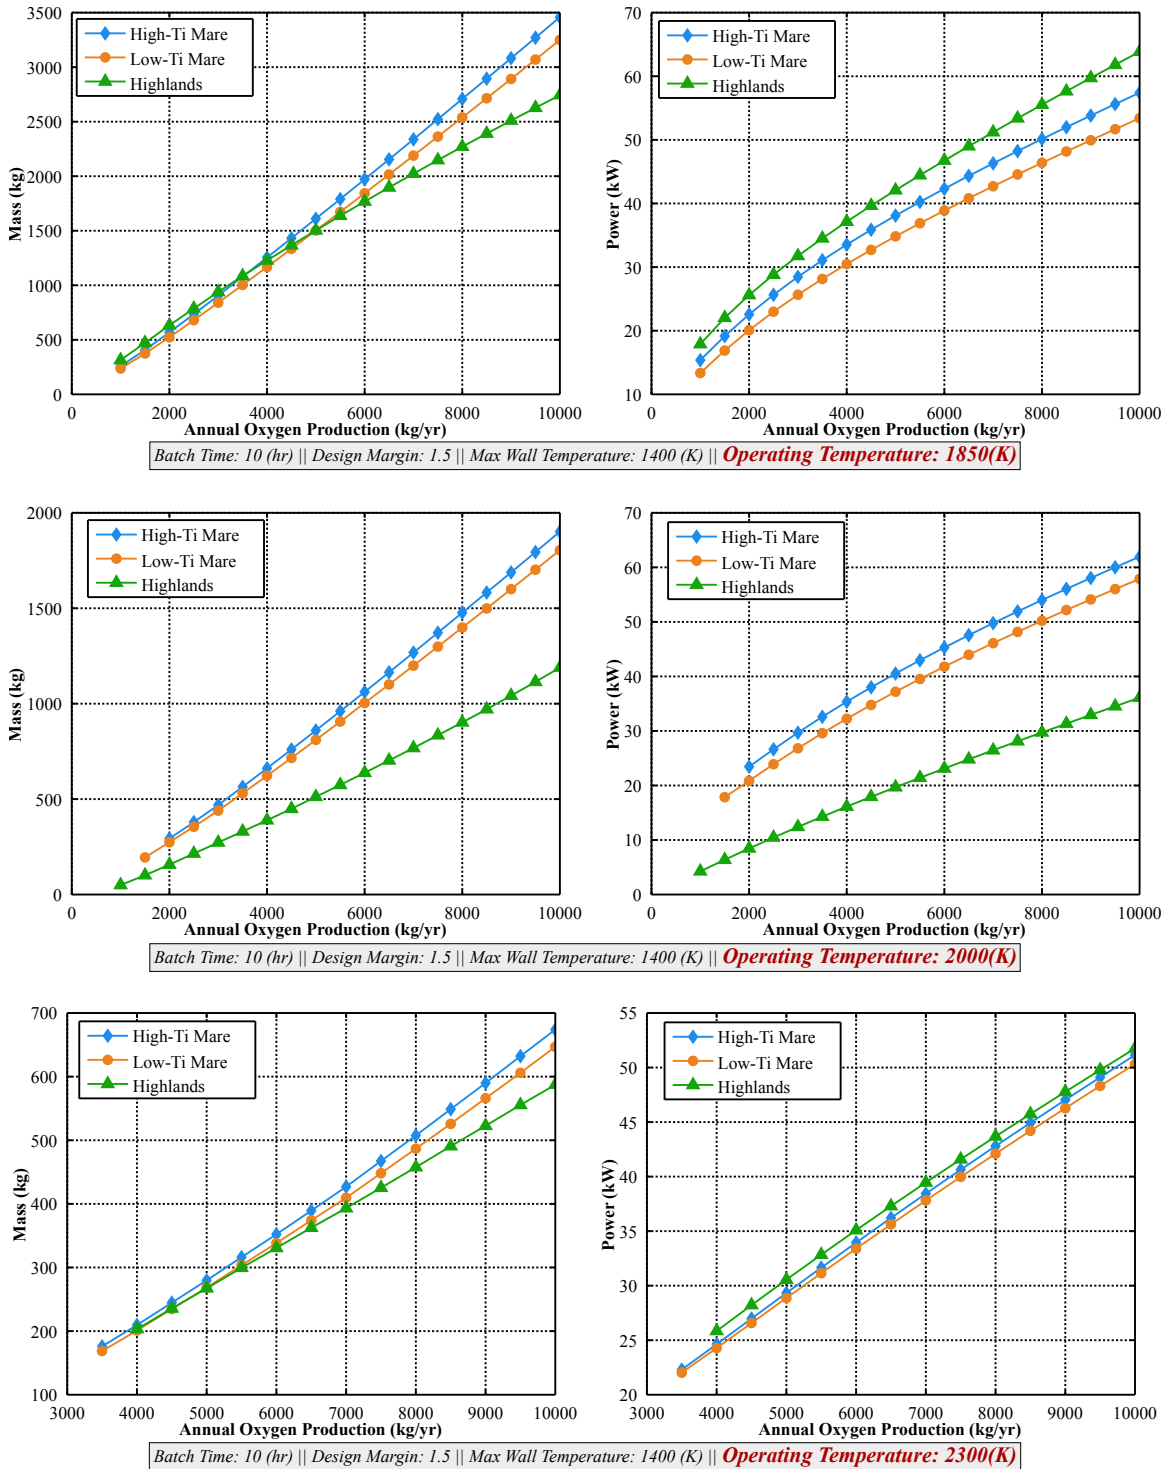


Figure 2-28: The mass and power of an MRE reactor for three different types of regolith with operating temperatures of 1850K (top), 2000K (middle), and 2300K (bottom).

power (60%) for Highlands regolith compared to Mare for operating temperatures circa 2000K. At operating temperatures circa 2300K, the mass and power of MRE reactors is again somewhat similar for both Highlands and Mare reactors. Highlands reactors are slightly less massive as higher production levels (due to less insulation requirements) but slightly more power intensive.

The Effect of Batch Time

In the MRE reactor model, the batch time plays an important role in determining reactor size and power. Through Equations (2.11) and (2.12), higher batch times require more molten regolith to be held in the reactor during operation. Figure 2-29 illustrates the impact of batch time on reactor mass and power. Naturally, longer batch times result in larger reactors that require more power. This is due to the fact that the reactor must now maintain a larger volume of molten material at the operating temperature for a longer period of time. Although shorter batch times reduce reactor mass and power, they also clearly limit the maximum production level of a single reactor, as observed in Figure 2-29.

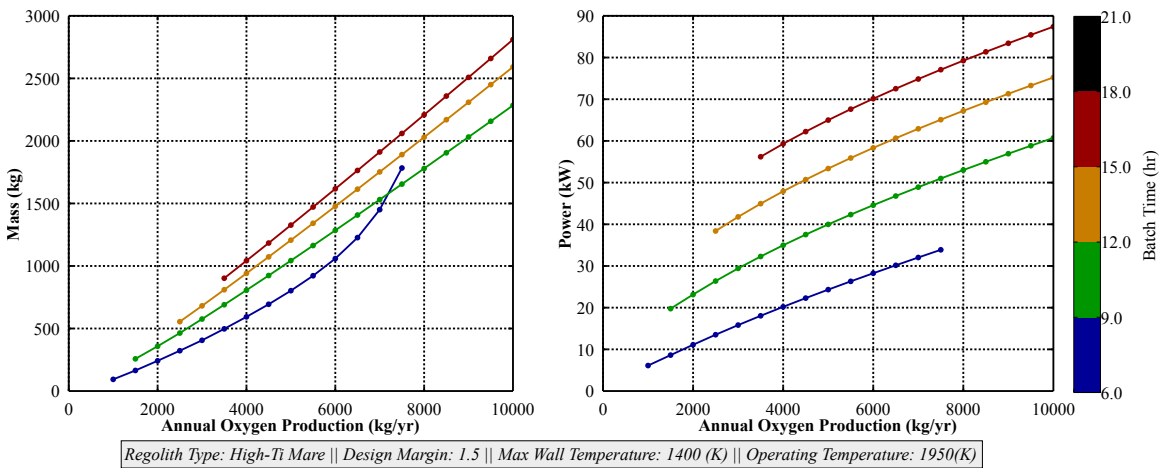


Figure 2-29: The mass and power of an MRE reactor as a function of oxygen production and batch time. Longer batch times increase reactor mass and power.

One attractive option for an MRE reactor is semi-continuous operation, in which regolith is continuously added while molten metals and slag are periodically withdrawn. This study of batch time is still relevant for continuous operation reactors,

as it describes the effect of the ratio of molten mass to reactor current. The results shown in Figure 2-29 demonstrate how lower ratios of molten mass to current (shorter batch times) are optimal, but they limit the production level of the reactor.

2.7.3 Molten Metal Production

As mentioned in Section 2.4.4, the MRE reactor model can also be used to predict the amount of molten metal produced by the reactor. Figure 2-30 illustrates the estimated metal produced from an MRE reactor as a function of operating temperature for Mare (left) and Highlands (right) regolith. As temperature increases, the metal production increases while the mass of leftover slag decreases by approximately 35% from 1850K to 2300K (for Mare) or 2000K (for Highlands). Reactors processing Highlands regolith can produce aluminum at lower operating temperatures because the liquidus temperature peak occurs around 2000K for Highlands regolith and around 2250K for Mare (see Figure 2-10). Reactors processing Mare regolith can produce more iron and titanium per kilogram regolith, while reactors processing Highlands regolith can produce more silicon and aluminum per kilogram regolith. This is directly due to the compositions presented in Figure 2-2.




Figure 2-30: The amount of metal produced by an MRE reactor operating on Mare (left) and Highlands (right) regolith. As operating temperature increases, more oxides can be reduced to produce more molten metal.

The data in Figure 2-30 is for an oxygen production level of 10,000 kg O₂/year. At higher operating temperatures (>2000K for Highlands and >2300K for Mare), the cumulative amount of metal produced is almost equal to the oxygen production level.

2.8 MRE Reactor Modeling Summary

This chapter presented the development of a parametric sizing model for a Molten Regolith Electrolysis reactor to produce oxygen and molten metals from lunar regolith. Section 2.2 gave a comprehensive overview of the regolith material property data that was integrated into the model. Section 2.3 describes how the regolith throughput requirements are derived from the oxygen production level. Section 2.4 describes the electrochemical equations behind the MRE reactor model. Section 2.5 provides the details on a multiphysics simulation of an MRE reactor, developed to increase the fidelity of the MRE reactor model. Section 2.6 describes how the data generated by the multiphysics simulation was leveraged to parametrically design MRE reactors.

Section 2.7 explored some of the results from the MRE reactor model. Perhaps the most important finding from the model is the tightly coupled nature of an MRE reactor design. Changing the diameter of the reactor has a dramatic impact on the operating temperature and molten mass within the reactor, and maintaining constant values for these quantities while varying reactor diameter requires the modification of wall thermal conductivity and electrode separation. Higher operating temperatures appear to always reduce reactor mass and sometimes reduce power (Figure 2-22), but have a strong dependence on the type of regolith being used (Figure 2-28). Lower design margins reduce reactor mass and power, but limit the maximum production level of a single reactor (Figure 2-23). Higher design margins are found to be an effective method for increasing reactor operation flexibility by enabling a range of feasible electrode separations (Figure 2-24). Batch time has a similar effect on reactor mass and power: longer batch times increase mass and power, but open up higher production levels.

These trends reveal the complex, coupled nature of MRE reactors. All of the

design variable effects listed here contain some cross-dependence between the various design parameters. For example, we expect the optimal design margin to strongly depend on the batch time and vice versa. Clearly, an optimization scheme must be employed to better understand the optimal MRE reactor design.

The MRE reactor model predicts that an MRE reactor is able to produce oxygen with a specific mass on the order of ≈ 0.1 (kg reactor)/(kg O₂/year) and ≈ 5 (W)/(kg O₂/year). These values provide initial evidence that an MRE reactor can be a viable option for producing oxygen from lunar regolith to resupply ECLSS consumables and provide oxidizer for chemical propulsion systems. Although these numbers differ from those presented in Schreiner et al. [95], they represent higher fidelity designs and are therefore considered more accurate. Despite these promising results, these mass and power predictions are for the MRE reactor alone. As described in Section 1.3.3, the integrated ISRU system must be studied to understand the true mass and power associated with producing oxygen on the lunar surface as well as the optimal ISRU system design. This integrated ISRU system is presented in the following chapter along with the holistic ISRU system mass and power estimates.

Chapter 3

Integrated ISRU System Optimization and Scaling

3.1 ISRU System Model Overview

As discussed in Section 1.3.1, modeling ISRU systems is critical to understanding the proper applicability of such systems. The MRE reactor model developed in Chapter 2 shed some light on the effect of various design variables, such as operating temperature, wall temperature, and batch time, but the optimal design of an MRE reactor cannot be completely understood without integrating it with other subsystem models. Naturally, a power system model is critical to determining the optimal balance between reactor mass and power, but other subsystem couplings can also prove important. Consider the operating temperature of an MRE reactor. The results shown in Section 2.7.2 indicate that higher operating temperatures are optimal, but operating temperature restricts the production range and therefore has a complex effect on the feasible values for batch time and design margin. Furthermore, higher operating temperatures result in the oxygen liquefaction and storage system using more power to cool the product oxygen to cryogenic temperatures. The proper balance of these design trades can only be understood when subsystem models are integrated together.

This chapter lays out an integrated ISRU system model that includes subsystem models of a reactor, power system, excavation system, regolith storage system, re-

regolith feed system, and oxygen storage and liquefaction system. A hybrid genetic algorithm/gradient-based optimization routine is developed and leveraged to minimize the mass of the holistic ISRU system over a range of oxygen production levels.

Section 3.2 provides an overview of the subsystem models. Section 3.3 presents the integrated system model and the N² diagram of the subsystem connections. Section 3.4 provides an overview of the optimization technique implemented on the ISRU system model. Section 3.5 explores the optimized system design over a range of oxygen production levels. Sections 3.6 and 3.7 present the estimated mass and power of the holistic ISRU system over a range of production levels. Section 3.8 explores the robustness of the MRE-based ISRU system to changes in regolith feedstock composition. Section 3.9 contrasts the oxygen production performance against system mass to assess the utility of an MRE-based ISRU system. Section 3.10 presents visualizations of the optimized ISRU system. Section 3.11 concludes with some key aspects of the system optimization analysis.

3.2 System Model Description

3.2.1 Reactor

Although a variety of reactor models can be integrated into the ISRU system model, this work utilizes the Molten Regolith Electrolysis (MRE) reactor model presented in Chapter 2. There are a number of key design variables in the reactor model that the optimization scheme can vary, as described in Section 3.4. These include the operating temperature, batch time, design margin, maximum wall temperature, as well as the number of reactors. Future work can address optimizing additional parameters, but these variables were chosen because they are the primary drivers of MRE reactor design.

3.2.2 YSZ Separator

A Yttria-Stabilized Zirconia separator is included in the system model to separate oxygen from the MRE reactor exhaust gas. Although the molten electrolysis process produces pure oxygen by electrolyzing oxides into oxygen gas and liquid metals, certain species (Na_2O , P_2O_5 , K_2O and MgO) will evaporate after electrolysis and will likely become entrained in the oxygen flow as contaminants. Additionally, trace gases such as H_2 , N_2 , CO_2 , and Helium will also be released as fresh regolith is heated up to a molten state [77]. Although a portion of these gases will be baked off by the preheater (see Section 2.6.6), it would be overly optimistic to assume that no volatile gasses are released once the regolith enters the high-temperature MRE reactor. To this end, a YSZ separator was included to ensure higher oxygen purity.

Yttria-Stabilized Zirconia (YSZ) is a ceramic material composed of zirconium dioxide (ZrO_2) stabilized by the addition of yttrium oxide (Y_2O_3). YSZ is commonly used as an electroceramic to measure oxygen content by monitoring the voltage across conductive platings on each side of the solid YSZ electrolyte. To act as a separator, an active voltage is applied across the electrodes while the gas flow encounters the cathode as shown in Figure 3-1. At the cathode, oxygen is ionized to O^{2-} and then transported through the YSZ electrolyte via the electric field between the plates.

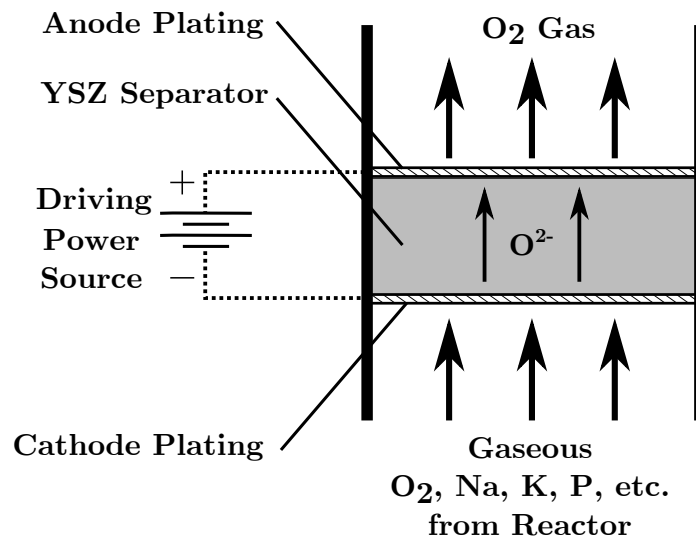


Figure 3-1: A diagram of the proposed YSZ separator used to purify oxygen from the exhaust gas from the Molten Regolith Electrolysis reactor.

The power demand of the separator is estimated by determining the required current and voltage across the YSZ ceramic. The current is directly proportional to the amount of oxygen that needs to be transported through the separator. The Nernst equation (shown in Equation (2.13)) with a current efficiency of one ($\bar{\eta} = 1$), as we assume no other species are transported through the separator, was used to calculate the expected current.

The electrical conductivity of YSZ drives the required voltage across the separator. Data on the temperature-dependent conductivity of YSZ [50] was fit with the equation:

$$\ln(\sigma_e(T)) = A_{YSZ} \exp(B_{YSZ} T) \quad (3.1)$$

where σ_e is the electrical conductivity in S/cm, T is the temperature in K, and the fit coefficients are $A_{YSZ} = -23.4(\pm 4.8)$ and $B_{YSZ} = -0.00259(\pm 0.0003)$. The temperature dependence in the YSZ conductivity couples the YSZ separator model and the MRE reactor model: a higher operating temperature in the reactor results in product oxygen with a higher temperature, which increases the temperature of the YSZ ceramic. This increase in temperature leads to an increase electrical conductivity, which decreases the voltage and power required for the separator. For simplicity, the temperature of the YSZ electrolyte was taken to be 75% of the reactor operating temperature. This was intended as a preliminary estimate to couple reactor temperature and YSZ temperature, while also accounting for some heat loss between the reactor and separator. Future work can generate a more accurate model of the expected temperature at the separator as a function of reactor temperature.

The electrical conductivity affects the resistance of the YSZ ceramic (R_{YSZ}):

$$R_{YSZ} = \frac{\Delta YSZ}{\sigma_e(T) S}, \quad (3.2)$$

where ΔYSZ is the thickness of the YSZ separator (conservatively assumed to be 0.25 cm), $\sigma_e(T)$ is the YSZ electrical conductivity calculated from Equation 3.1, and

S is the required cross-sectional area of the YSZ separator calculated as:

$$S = \frac{I_{YSZ}}{j}, \quad (3.3)$$

where I_{YSZ} is the required current through the YSZ separator and j , the limiting current density, was taken to be 0.4 A/cm² [125]. The power of the YSZ ceramic was then estimated using the current and resistance (I^2R). The dimensions of the separator with a 304 Stainless Steel case are used to calculate the mass of the YSZ separator.

It should be noted that the current YSZ separator model is a simplified version with the intention of determining the power needed for oxygen separation with only first-order estimates of mass and volume. It is believed that the power requirement of the YSZ separator will play a much more significant role than its mass in the ISRU system optimization. Solid oxide electrolyzers similar to this concept are commonly composed of multiple packed tubes or stacked wafers, which could reduce the mass and volume estimates, but not significantly change the power requirement, compared to this simplified YSZ model. A more realistic mass model will need to be created in a future iteration.

3.2.3 Excavator

The excavator system model, developed at the Glenn Research Center [40], predicts the mass of a mobile excavation platform sized to deliver the regolith throughput requirement to the reactor. A force module utilizes the Balovnev force equations to generate estimates of the force and torque involved in excavating lunar regolith. A hole depth of 25 cm with cut depths of 2.5 cm was used to size a front-end loader in this design iteration, though future work can optimize these values. The excavation force estimates are used to size the excavation actuators using commercial off-the-shelf (COTS) actuators and controllers from Danaher¹. The force module also determines the vehicle reaction and traction forces. A mass module conducts a structural analysis

¹<http://www.danahermotion.com>

to ensure that the excavator chassis can support the regolith weight and that the digging mechanism can support the expected excavation stresses. The locomotion motors are modeled after the Maxon motors used on the Mars Exploration Rover [68].

An excavator speed of 0.5 m/s and a plant distance of 100 m are used to properly size the mobility platform for the excavator. Information on the excavator operating duty cycle based on the power system charge/recharge cycle is also incorporated into the model. The excavator model utilizes all of this information to generate an excavator design that can meet the regolith delivery requirements from the reactor while withstanding the excavation forces and regolith load requirements.

3.2.4 Hopper and Feed System

The hopper model calculates the mass of the storage hopper required to hold regolith prior to inserting it into the reactor, shown in Figure 3-2. The main driver in the hopper model is the buffer capacity, or the amount of regolith the hopper had to hold in terms of days of reactor operation. A buffer capacity of 3 days was chosen to ensure that the hopper could hold enough regolith for continual reactor operation if the excavator needed repairs. Furthermore, a buffer capacity of 3 days effectively decouples the excavation system scheduling from reactor batch mode operation (i.e. although the reactor may operate on a 1 hour batch time, the excavator can deliver regolith with a lower frequency). An ullage of 20% volume was used with a height-to-diameter ratio of 1.25. The hopper walls were assumed to be made out of 3 mm aluminum sheets, as the hopper will be safely removed from the high-temperature environment of the reactor by the feed system. A single central hopper was designed to support any number of reactors.

Developed at the Glenn Research Center and modified for this work, the feed system model calculates the mass and power of the system required to insert fresh regolith from the hopper into the reactor, shown in Figure 3-2. An auger/screw was chosen for this design iteration, but other methods, such as a pneumatic feed system [75], may be modeled in the future. The feed system model sizes an auger that extends from the reactor through a cylindrical sleeve and into the bottom of

the hopper. Using estimates of the cohesion, internal and external friction angles, and soil-tool adhesion values for lunar regolith, the feed system model estimates the expected torque on the auger and the resultant power consumption. The number of feed systems is set equal to the number of reactors, as each reactor will likely require its own feed system. The sleeve and auger are made out of Titanium C110M, due to the interface with the high-temperature reactor. A fill ratio of 0.8 was used, meaning that 80% of the auger is assumed to be filled with regolith.



Figure 3-2: A Computer-Aided Design (CAD) model of the hopper and auger used in the ISRU system model.

One assumption built into the feed system model is that a 10 cm diameter auger rotating at 10 rpm would be adequate to insert a full batch of regolith in the feed time set as an input in the reactor model. A 5 mm thick sleeve with a 1.5 cm thick blade at a pitch of 9.6 cm were also fixed within the feed system model. That is, for smaller amounts of regolith per batch, the feed system was not parametrically sized down,

due to limitations in the model design. Future work can expand the feed system model to dynamically size the radius and rotational rate of the auger system to meet the required regolith rate of insertion. Furthermore, a motor to actuate the auger was not included in the feed system model; future work can address this deficiency.

3.2.5 Oxygen Liquefaction and Storage

The oxygen liquefaction and storage model utilizes oxygen production data from the reactor to size both the liquefaction and storage systems.

The storage system is comprised of a cylindrical tank with hemisphere end caps. A storage capacity of 6 months was chosen to allow for sufficient propellant production to support two refueling missions per year. The number of layers of MLI can be chosen to balance heat loss with system mass and a value of 50 layers was chosen for this study. Although a variety of materials can be chosen, for this study aluminum was used for the tank wall material. The storage tank is sized such that the yield stress is greater than the hoop stress with a factor of safety of 2. Two tanks, each with an ullage of 7% and a storage pressure of 306.8 kPa, were used in this study to provide a measure of redundancy. The tank size and number of layers of MLI directly impact the boil-off rate due to expected heat leakage into the tank.

The liquefaction system determines the mass and power of the system required to liquefy the oxygen coming from the reactor. A separate cooler was designed to re-liquefy oxygen that has boiled off in the storage system. The duct work required to link the liquefaction, storage, and reactor systems is not currently included in the system mass estimate.

3.2.6 Power System

The power subsystem is parametrically sized from the total power requirement summed over all of the other subsystems, as shown in Figure 3-3. For the results shown in this work, the system design was restricted to be solar arrays that provide power to the ISRU system for day-only operation. Future work can assess the impact of using

other power systems, such as solar arrays with fuel-cell energy storage to enable lunar night operation, a Stirling radioisotope generator, or a fission surface power system.

The most important factor in the power system is the specific mass (kg/kW). This parameter plays a large role in whether or not an ISRU system is in fact viable, but unfortunately values for this parameter vary widely depending upon year, source, and application. To this end, a literature review, shown in Table 3.1, was conducted to better understand the appropriate value to use for this study.

Table 3.1: A review of various power system specific mass and area numbers for lunar surface systems in the literature.

Source	Year	Specific Mass (kg/kW)	Specific Area (m ² /kW)	Comments
Simon [102]	1985	10.0	–	<i>Ballpark value, not cited</i>
Mendell [74]	1985	5.3	–	
Hepp Conservative [52]	1994	15.1	–	<i>Based on hardware</i>
Hepp Adv. Tech. [52]	1994	3.3	–	<i>Assumes advanced tech</i>
BVAD [4]	1995	20	–	
Hoffmann [54]	2000	5.1	9.7	<i>From hardware model</i>
Space Mission Eng. [126]	2011	14.0	5.0	<i>From satellite missions</i>
ATK ROSA [16, 17]	2013	7.3	2.8	<i>Based on hardware</i>
ATK MegaFlex [16, 17]	2013	7.2	1.8	<i>Based on hardware</i>

Some numbers used in the early years of lunar base modeling were ballpark estimates on the order of 5-10 kg/kW [74, 102]. The value listed in the 2015 Baseline Values and Assumptions Document [4] is actually based on work published in 1995 and is significantly higher than the numbers used in the decade prior. In 2000, Hoffman et al. [54] conducted a trade study of thin-film solar arrays and created a hardware sizing model that predicted a specific mass of solar arrays around 5 kg/kW. This model included the wiring, photovoltaic cells, substrate, boom/structure, and mechanical components. The Space Mission Engineering textbook [126] gives a value on the order of 14.0 kg/kW. This value is somewhat higher than the previous data, but it is based on existing satellite designs rather than lunar surface designs.

The most recent numbers (≈ 7 kg/kW) came from two conceptual design studies of high-powered solar electric propulsion systems that requires 30 kW [16] and 300 kW [17]. Both numbers are based off of actual hardware that is currently be-

ing fabricated by ATK for use on the Orion space capsule. Due to the high-fidelity, hardware-driven nature of the values for the ATK MegaFlexTM solar array, values of 7.2 kg/kW and 1.8 m²/kW, which include a 29% margin, were used for the power system model in this study.

A comparison of the specific mass value used in this study with values used in previous ISRU studies provides greater context. The model developed at the International Space University [13] used a value of 2.2 kg/kW, though this value seems somewhat low and it not cited. The modeling conducted by Chepko [23] used a value of 72 kg/kW, though this value seems somewhat high and is not cited. Compared to this work, the model presented in Schreiner et al. [95] used a more conservative value of 20 kg/kW from the Baseline Values and Assumptions Document [4, 44], but assumed a 70% duty cycle in the Shackleton crater rim area. The work presented here assumes a 50% duty cycle unless it is explicitly stated that the simulation is for a base near the polar regions.

It is worth noting that the specific area value used in this study (1.8 m²/kW) is significantly lower than previous estimates. This is due to the dramatic increase in the efficiency of thin-film solar array technology over the past decade, as evidenced by the monotonic decrease in array specific area observed in Table 3.1.

3.3 Modeling the Integrated System

The subsystem models described in the preceding section were integrated together into a holistic system model. By linking the subsystems (reactor, excavator, power, etc.) together into a self-consistent model, the entire mass and power of an ISRU system can be estimated. The self-consistency of the model allows the tradeoffs between subsystem designs to be studied. For instance, shortening the batch time of an MRE reactor is one avenue for reducing reactor mass, but this reduction in reactor mass comes at the cost of an increase in reactor power due to the increase in total down time between batches which reduces total operational time. The integrated model enables a more complete study of the optimal batch time, as one example, by

including the mass of both the power and reactor subsystems.

Another important design variable to optimize is the reactor operating temperature. The results from Section 2.7.2 revealed that higher operating temperatures minimize both reactor mass and power, but this also drives a higher power requirement for the oxygen liquefaction system. Furthermore, Figure 2-22 shows how reactor specific mass increases with production level while specific power decreases with production level. This means that higher production levels require relatively larger reactors but relatively less power per kilogram oxygen produced. To truly determine the optimal production level and operating temperature for a single MRE reactor, the integrated system must be studied.

Figure 3-3 depicts an N² diagram of the ISRU system. The primary subsystem couplings are shown, with some secondary connections left out for clarity. It is evident that the reactor is a strong driver of many other system designs. It contributes to the total power requirement and also determines the regolith processing requirement which directly affects the excavator, hopper and feed systems. The power requirement from each subsystem is summed together and used to size the power system. After the power system is sized, the mass of each subsystems, including the power system, is summed together to generate an estimate of the total ISRU system mass. Although many of the models have rudimentary volume estimates, at this point in time these estimates were not considered sound enough for publication so they are not presented.

3.4 Optimization Technique

3.4.1 Hybrid Optimization Description

A genetic algorithm (GA) optimization routine was used with the integrated system model to optimize the ISRU system design by varying subsystem design variables. A GA method was chosen, rather than purely gradient-based optimization techniques, for a number of reasons. First, the objective function of the ISRU system contains a large number of local minima and the general location of the global minimum is not

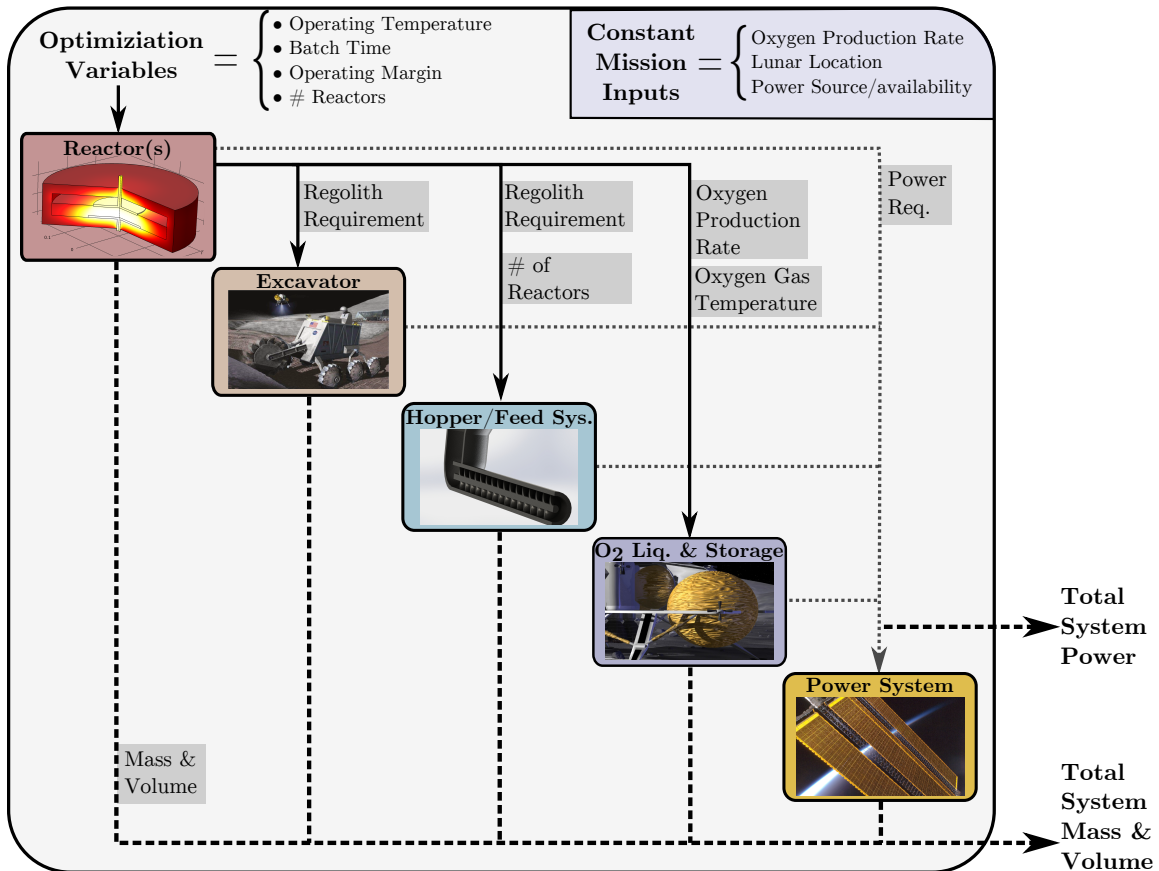


Figure 3-3: An N² diagram of the ISRU system model within the optimization routine, showing how the subsystems are interconnected to generate a self-consistent estimate of system mass, which is then optimized.

well known a priori. Second, the problem has an integer component: although some parameters were continuous, such as operating temperature, other parameters were discrete, such as number of reactors, wall temperature (discrete choices), number of excavators, and material selections.

A GA is a heuristic search method that attempts to mimic natural selection by generating a population of candidate designs in what is called a generation. The fitness (or goodness) of each generation is evaluated and three methods are used to create the subsequent generation. First, a set number of the best “Elite” designs are directly carried over. Next, a certain fraction of the remaining designs (the “Crossover Fraction”) are mixed together to create children for the next generation. Finally, the remaining designs (the “Mutation Fraction”) are mutated by randomly changing a




Figure 3-4: The mutation (red), crossover (blue) and elite (black) designs throughout multiple generations in the genetic algorithm optimization scheme.

design variable from some parent designs to introduce variety into the population. This process is illustrated in Figure 3-4, in which the genealogy of the GA is illustrated. Red lines link parents to mutated children, blue lines link parents to crossover children, and black lines show how elite designs are directly passed down. Note that the y-axis, “Individual”, is arbitrary and does not have a relationship to the fitness of a given individual.

The GA solver terminates when the fitness function does not significantly change over a set number of generations, called the “Stall Generation Limit”. For this study, the fitness function was simply the total ISRU system mass. A Stall Generation Limit on the order of 10-200 was chosen with a population sizes of 1000 to 10 designs in each generation (a range of optimization options were tested). An Elite Count in the range of 5-15 was also used with crossover fractions between 0.3-0.6.

The reduction in system mass with subsequent generations in the GA solver is shown in Figure 3-5. The “Mean penalty value” markers depict the mean system mass within the entire population of systems designs in a given generation. The




Figure 3-5: A sample output from the genetic algorithm optimizer used on the ISRU system model, where the penalty value is the mass of the ISRU system (kg). The downwards trend in the blue data shows the effectiveness of the “natural selection” of better performing candidates from generation to generation.

“Best penalty value” shows the lowest mass ISRU system in a given generation.

Although GA is a suitable technique for optimization over discrete variables, it is not particularly well suited to optimized a large number of continuous parameters. To enable a more efficient optimization, a gradient-based optimizer was implemented that used the final GA solution as a starting point with the integer variables fixed. The ISRU system model is nonlinear and contains no analytical gradient, so the solver used finite difference approximations for the gradient. In this manner, the GA optimizer was used to find what was likely the global minimum region while the gradient-based optimizer was used to hone in on the exact minimum.

As discussed in Section 2.6.7, many of the subsystem models contained error flags that identified infeasible reactor designs, vehicle slippage, and a number of other system model errors. A set of soft constraints were implemented by penalizing the mass of systems with error flags by adding 2000 kg and multiplying by a factor of 5 (the sample GA output in Figure 3-5 does not include this penalty). These soft constraints guided the optimization scheme away from system designs with error flags.

3.4.2 Bounding the Search Space

The optimization scheme described above was used to study the optimal ISRU system design. The batch time, operating temperature, design margin, wall temperature and number of reactors were used as free design variables in the optimization scheme. The batch time was constrained between 1 and 40 hours, after a preliminary evaluation revealed the optimal batch times to typically be in that range.

The operating temperature was constrained between 1850K and 2300K. The melting point of iron is $\approx 1810\text{K}$, so a lower bound of 1850K will avoid producing solid iron. It is worth noting that after the reactor has produced some silicon, it will create an alloy with the iron with a lower melting point - thus, we can imagine a reactor setup that operates below 1850K and still maintains a liquid metal product. Initial results from the optimization indicated that optimal designs always have operating temperatures significantly above the lower bound, so this bound was not relaxed for future runs. The upper bound of 2300K was used to avoid approaching the melting point of iridium ($\approx 2700\text{K}$), which could possibly be used as an anode material. The operating temperature is calculated as the current-averaged temperature within the melt, so an operating temperature of 2300K can correspond to higher peak temperatures within the reactor. Thus, a bound of 2300K was used to avoid material melting issues.

The design margin was constrained between 1.0 and 6.0. The lower bound of 1.0 is discussed in Section 2.6.3 and the upper bound of 6.0 was chosen after an initial optimization run revealed that the optimal design margin was always significantly below this value. The wall temperature could only take discrete values of 1200K, 1300K, and 1400K as discussed in Section 2.7.2, so an integer variable was used to set this parameter. Finally, the number of reactors was an integer variable bounded between 1 and 6, based on results from a previous study [96].

3.5 Optimized System Design

The optimization scheme described in Section 3.4 was used to generate optimal designs for a range of production levels between 1,000 kg O₂/year and 10,000 kg O₂/year. We note here that the multiphysics data presented in Chapter 2 extended to a maximum current of 3 kA, which corresponds to an oxygen production level of ≈ 2500 kg O₂/year. The confidence in the reactor scaling equations derived from the multiphysics data allows the parametric reactor model to extrapolate reactor designs beyond the range of the multiphysics data. Future work can generate more multiphysics simulation data at higher production levels to further anchor the design trends.

Figure 3-6 shows the the growth of the ISRU system mass and power over a range of oxygen production levels in the top two plots. The results shown in the following section are for Highlands regolith unless otherwise stated. The remaining graphs (with labels) depict the optimized system design tradespace, including the number of reactors (a), operating temperature (b), reactor diameter (c), molten mass per batch (d), average reactor current (e), operating voltage (f), current density (g), the MRE design margin (h), electrode separation range (i), and the batch time (j). The current, reactor diameter, and molten mass per batch are for a single reactor.

The system mass and power are discussed in the following sections, here we discuss the optimal design parameters shown in the various plots in Figure 3-6. The optimal number of reactors (*plot a* in Figure 3-6) exhibits a clear preference towards one large reactor as opposed to multiple smaller reactors. This value was not set a priori, but rather selected by the optimization scheme. This contrasts previous work [96] in which the number of reactors increased with production level. The main reason for this difference is the fact that the MRE reactor model now scales better at higher production levels. The reactor model used previously [95] exhibited asymptotic mass behavior at larger production levels, which led to a natural increase in number of reactors to meet higher production levels beyond the vertical asymptote. With the improved reactor scaling equations presented in Section 2.7, the optimal system consists of a single reactor over the production levels studied in this work.




Figure 3-6: (Top) The optimized ISRU system design variables over a range of production levels for Highlands regolith.

The optimal operating temperature (*plot b* in Figure 3-6) displays a monotonic increase with production level. The bounds on operating temperature, discussed in Section 3.4.2, are illustrated by the black dotted lines. The optimal operating temperature begins around 1900K at 1000 kg/yr and then rises towards the 2300K ceiling for higher production levels, though this rise is more gradual than previous results [95]. The increase towards higher temperatures is likely due to the fact that higher operating temperatures allows more oxygen to be extracted per kilogram regolith, which reduces regolith throughput requirements and reactor size as discussed in Section 2.7.2. Higher temperatures also reduce the power consumption of the YSZ separator. Prior to this work, it was unclear whether these benefits would be outweighed by the increased thermal gradient requirements, increased regolith heating requirement (per kilogram regolith), decreased regolith density, and potential increased power system size. The integrated system model showed that operating temperatures higher than the traditional 1850 K do indeed result in a lower total system mass.

The reactor diameter (*plot c*) naturally grows with oxygen production level. At a production level of 1,000 kg/year the optimal diameter is around 1.75 m and at 10,000 kg/year the diameter is as approximately 3.5 m. A power-law regression equation was fit to the reactor diameter using the `fit` function in MatlabTM:

$$D \approx 0.235[\pm 0.028] N^{0.293[\pm 0.014]}, \quad (3.4)$$

where D is the reactor diameter, N is the oxygen production level in kg/year, and 95% confidence bounds are given. One would expect the diameter to grow with the square root of production level: with a constant electrode separation the molten mass is proportional to diameter squared, so for a linear increase in molten mass the diameter should grow with the square root of production level. The observed power in the diameter scaling law is less than the expected value because because of the sub-linear growth in molten mass (*plot d*) as operating temperature increases. The molten mass per batch grows from ≈ 25 kg at 1,000 kg/year to ≈ 150 kg at 10,000 kg/year. Previous work [96] identified an optimal value of around 1.87 kg/batch, which differs from the

results published here because those previous models did not incorporate a minimum electrode separation as described in Section 2.6.7. It appears that imposing a minimum electrode separation drives the molten mass per batch significantly higher, as the electrodes can no longer be extremely close together. The fact that the optimal system has only one reactor also contributes to the observed increase in the molten mass per reactor.

The current per reactor (*plot e*) intuitively increases almost linearly with oxygen production level. Slight deviations from linearity occur due to the change in average current efficiency as more species are electrolyzed at higher operating temperatures, as detailed in Section 2.2.7.

The average reactor voltage (*plot f*) decreases asymptotically from a value of approximately 13V at a production level of 1,000 kg/yr to around 5V above 10,000 kg/year. This is a result of increased current through the reactor as shown in *plot e* – as the current increases, voltage can decrease while still generating enough heat to maintain the molten core. This value is slightly higher than previous work [95], which is likely due to the new requirement that the average electrode separation be greater than or equal to 2 cm, as detailed in Section 2.6.7.

The required current density (*plot g*) gradually rises with production level, from 2 A/cm² at 1,000 kg/year to 4.5 A/cm² at 10,000 kg/year. These values are sufficiently within the values obtained in experimental hardware demonstration work, as described in Section 2.6.7.

The MRE reactor design margin (*plot h* in Figure 3-6) remains relatively close to the minimum bound of 1.0. As described in Section 2.6.3, a design margin greater than 1.0 provides a measure of reactor operation flexibility by enabling a range of feasible electrode separations that satisfy the joule-heated, cold-wall constraint. Furthermore, the reactor standalone analysis presented in Section 2.7.2 indicates that a lower design margin minimizes reactor mass and power, but limits the maximum production level for a single reactor. The optimization scheme appears to favor manipulating the batch time (discussed below), rather than increasing design margin, to access higher production levels. We observe here that the optimization scheme

attempts to minimize the design margin, but this minimization hovers slightly above 1.0 to open up a range of feasible electrode separations as discussed in the next paragraph.

The feasible range of electrode separations (*plot i* in Figure 3-6) demonstrates two main points. First, the requirement that the average electrode separation be greater than or equal to 2.0 cm clearly plays a critical role in the optimal design – the optimization scheme appears to select designs with the minimal possible electrode separation. This produces designs with an average electrode separation (the thick central dashed line) directly along the lower bound of the electrode separation constraint. Furthermore, the requirement that the minimum and maximum electrode separation values are at least 0.5 cm apart (see Section 2.6.7) also drives the optimal design. The thin upper and lower black dashed lines in *plot i* of Figure 3-6 depict the minimum electrode separation bounds that satisfy this constraint, and the optimization scheme clearly follows this constraint quite closely. The required electrode separation flexibility is undoubtedly the cause behind the optimal design margin (*plot h*), in which we observe the optimization scheme selecting the lowest possible design margin that still satisfies this electrode separation range requirement.

The optimal batch time (*plot j*) appears to start around 14 hr at a production level of 1,000 kg/year. After rising to a value of 31 hr at 3,000 kg/year, the batch time generally decreases towards an asymptote of approximately 24 hr. The dramatic rise in batch time corresponds to the operating temperature crossing $\approx 2000\text{K}$, at which point the Highlands reactor can now extract more oxygen per kilogram regolith. The batch time controls the current to molten mass ratio in the reactor – larger batch times result in more molten mass per amp of current in the reactor for a given production level. A batch time of around one day appears to provide enough molten mass within the reactor to satisfy the minimum electrode separation requirement. This is an interesting relationship that falls out of the modeling – because MRE reactors extract so much oxygen from each kilogram of regolith, the batch times must be higher to allow the sufficient regolith to dwell within the reactor.

3.6 ISRU System Mass

Although the mass breakdown is shown in the upper left of Figure 3-6, it is shown in more detail in Figure 3-7. The most significant mass drivers are the oxygen liquefaction/storage system and reactor, both of which comprise 30% of the total system mass at 10,000 kg/yr. The power system also plays an important role, contributing 25% of the system mass at 10,000 kg/yr. As mentioned in Section 3.2.5, the oxygen storage system was designed to hold 6 months of oxygen production at any given time, and this requirement may be relaxed depending upon the mission needs.




Figure 3-7: The mass breakdown of the ISRU system processing Highlands regolith, compared to other models from the literature.

Figure 3-7 also compares the mass predictions from this work to several linear scaling laws in the literature. We see that the mass estimate for the reactor (blue area) compares well to the predictions from all three scaling laws at low production levels, but as production level increases the mass estimates from the MRE reactor model begin to drop away from the linear scaling laws. The linear scaling laws presented by

the FERTILE MOON study (Belachgar et al. [13]) and Taylor and Carrier III [114] are not directly cited from any previous work, but rather are described as coming from a broad literature review. Both of these mass predictions are significantly less than the results from the integrated ISRU system, but because the source of these numbers is unknown, the cause of the discrepancy is not certain. It is likely that these scaling laws do not include a power system and possible that they do not include a liquefaction and storage system. Removing these two elements from our mass predictions results in masses closer to the linear scaling laws, but we still observe a noticeable difference in the intercept and slope of the mass predictions.

The total system mass curve was fit with the following power-law curve using the `fit` function in MatlabTM:

$$M = 5.16[\pm 0.94] N^{0.622[\pm 0.021]} \quad (3.5)$$

where M is the ISRU system mass, N is the annual oxygen production level in kg/year, and 95% confidence bounds are given. The fact that the coefficient in the exponent is less than 1.0 implies that the ISRU system exhibits an economy of scale. That is, the ISRU system produces higher quantities of oxygen more efficiently.

The three most massive elements in the system are the reactor (30% at 10,000 kg O₂/year), the liquefaction and storage system (30%), and the power system (25%). The mass of the liquefaction and storage system can be reduced if some of the oxygen is used for environmental system support and does not need to be liquefied and stored. The liquefaction and storage system will also scale down as the oxygen storage duration decreases – a storage duration of 6 months was used in this study, but shorter storage times may prove viable for scenarios with higher oxygen demand.

As described in Section 3.2.6, the power system was modeled as a day-only solar array using specifications from the latest ATK MegaFlexTM arrays. This mass may go down with advances in thin-film solar arrays, but the reduction in area-specific mass must be balanced by the decrease in cell efficiency for thin-film options.

The primary drivers of reactor mass were also investigated to lend clarity to the

reactor mass estimate, as shown in Figure 3-8. The layer of refractory material lining the interior the reactor is the largest driver of reactor mass, comprising 50% of the total reactor mass at a production level of 10,000 kg O₂/year. This is because the refractory material has a significantly higher density than the insulation material (although the structural material has a higher density, it is much less thick). The second largest driver of reactor mass was the outer structural shell, comprising 24% of the total reactor mass at a production level of 10,000 kg O₂/year.

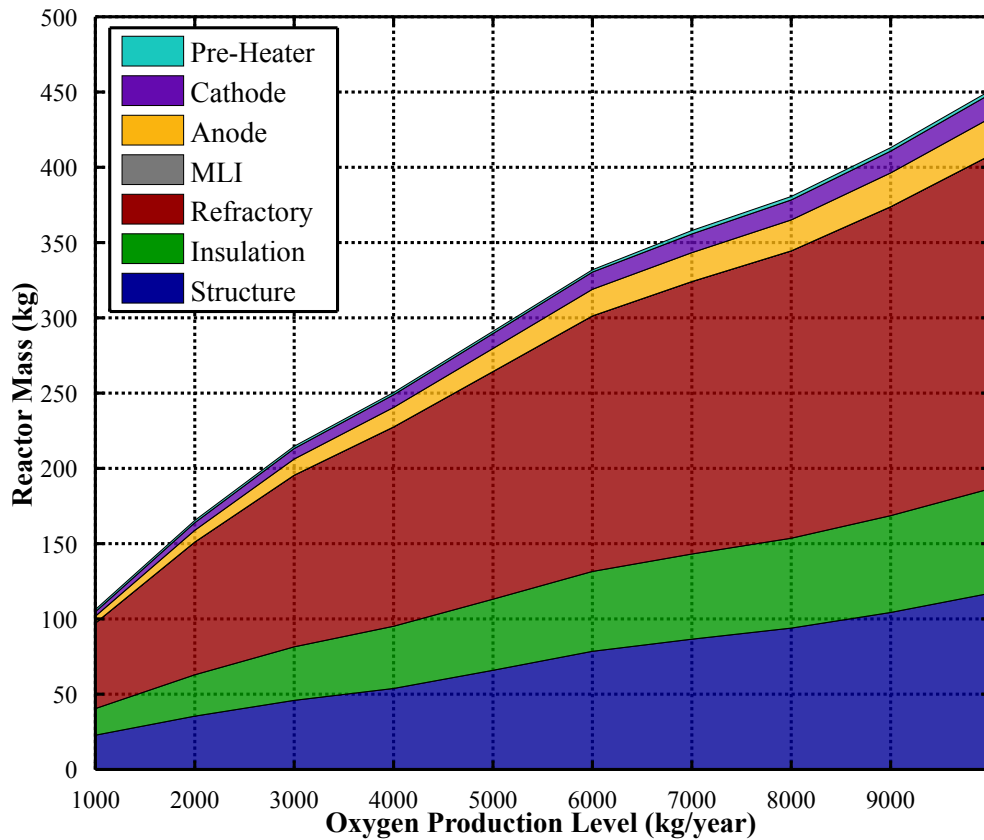


Figure 3-8: The mass breakdown of a Highlands MRE reactor over a range of oxygen production levels. The refractory material contributes $\approx 50\%$ of the reactor mass.

3.7 ISRU System Power

Although the power breakdown is given in the upper left of Figure 3-6, it is shown in more detail alongside several models from the literature in Figure 3-9. The “Chemical Electrolysis (ΔG)” section represents the power required to break the chemical bonds

in the oxides in lunar regolith. The “Regolith Heating + Phase Change” section represents the power required to heat the regolith up from the ambient temperature to the operating temperature, including the latent heat of melting in the phase change. “Radiative Heat Loss” is the heat lost from the exterior of the reactor to the lunar environment. The “Endothermic Makeup” slice depicts the amount of power required to maintain thermal equilibrium throughout the endothermic electrolysis reaction. Power demands for the “YSZ Separator”, “Feed System”, and “Liquefaction and Storage” are also shown. The comparison to models from the literature reveals a number




Figure 3-9: The power predictions from the ISRU system model compared to four linear scaling laws from the literature.

of insights. First, the MRE reactor model’s estimate of the chemical electrolysis power compares well with the linear scaling law presented in Hepp et al. [52]. The MRE reactor model estimates are slightly higher because higher operating temperatures enable the electrolysis of species that have a higher Gibbs Free Energy, increasing the

average chemical electrolysis energy. Second, when we include the expected regolith heating power our predictions compare well with Sherwood and Woodcock [98], and reasonably well with Bock et al. [15]. Christiansen et al. [26] notes that Bock et al. [15] only included the chemical electrolysis and heating power without any efficiencies, so the similarity between that scaling law and that subset of our results is unsurprising. Teeple [116] included the power to liquefy the product oxygen as well, which is why that scaling law is higher than the others.

The total power was fit with the following power-law curve using the `fit` function in MatlabTM:

$$P = 0.253[\pm 0.033] N^{0.588[\pm 0.015]}, \quad (3.6)$$

where N is the oxygen production level in kg/year, P is the power in kilowatts, and 95% confidence bounds are given. The fact that the coefficient in the exponent is less than 1.0 indicates an economy of scale – that is, higher production levels can be met with higher power efficiency. The non-linearity in the total process power raises some doubt as to the fidelity of linear scaling laws.

Most importantly, we notice that our total power estimate is significantly higher than other linear scaling laws in the literature. This is primarily due to the fact that we include an estimate of radiative heat loss and the endothermic makeup power, which together account for almost 2/3 of the total power at 10,000 kg O₂/year. Due to this discrepancy, back of the envelope calculations were carried out to provide a sanity check on the radiative heat loss numbers. Consider the current running through the molten region in the reactor. The resistance of the molten region ($R_{molten-region}$) can be approximated as:

$$R_{molten-region} \approx \frac{\bar{\rho}\Delta e}{\pi(D_{eff}/2)^2}, \quad (3.7)$$

where Δe is the electrode separation (inter-electrode distance), $\bar{\rho}$ is the average resistivity of the molten regolith, and D_{eff} is the effective diameter of the molten region. The reactor designs in this study have diameters on the order of 2 m, which with a diameter ratio of 0.15 results in an electrode diameter on the order of 0.3 m.

The diameter of the molten region can be approximated by the electrode diameter ($D_{eff} \approx 0.3\text{m}$). The average electrode separation was limited to be at least 2 cm (see Section 2.6.7), and as shown in *plot i* in Figure 3-6, optimal designs always had an average electrode separation around this limit ($\Delta e \approx 0.02\text{m}$). From Figure 2.2.5, the average resistivity $\bar{\rho} = \left(\frac{\bar{1}}{\bar{\sigma}}\right)$ can be estimated to be around 1/80 Ωm . With a current of ~ 3 kA (which gives an oxygen production level around 2500 kg/year), the resistive heating in the molten core of the reactor will be on the order of 30 kW. In fact, Colson and Haskin [29] hypothesized that *“high melt resistivities coupled with the large distance between electrodes that would seem to be required to make the approach robust might make power requirements prohibitive...”* for an MRE reactor. The results presented here are now able to quantitatively describe the power requirements due to the resistive heating in the molten region.

This is an important finding from the system modeling: the minimum electrode separation requirement results in a significant amount of radiative heat loss. Previous work [96] predicted relatively lower radiative heat loss due to the lack of a lower bound on electrode separation, but the electrode separation distances required to achieve such low heat loss are much too low for practical operation, as detailed in Section 2.6.7. However, this radiative heat loss does not necessarily have to go to waste. A few kilowatts of this resistive heating will go towards heating up fresh regolith and supplying the heat to maintain a constant temperature in light of the endothermic chemical reaction (Section 2.5.2). Future work can examine constructive uses for this extra power from the reactor.

3.8 Regolith Type Dependence

The composition (and therefore type) of regolith has wide-ranging implications in the MRE reactor model. As discussed in Section 2.2, almost all of the regolith material property models contain a dependence on the type of regolith. Molten Highlands regolith is less dense than Mare, so to accommodate the same mass of regolith a Highlands reactor will be larger. The specific heat and latent heat of melting for

Highlands regolith is around 5% more than Mare, so one can expect a slightly higher heating power per kilogram Highlands regolith when compared to Mare regolith.

The electrical conductivity of Mare regolith is higher, so one would expect Mare reactors to generate less resistive heating via the electrolysis current. Highlands regolith has less FeO, which will lead to a higher average current efficiency. However, FeO is one of the most chemically attractive electrolysis targets (low Gibbs Free Energy and Enthalpy of Formation), so this may lead to higher power requirements for Highlands regolith. The multiphysics regression trends were created for Highlands and Mare regolith separately, so the coefficients in the scaling equations for diameter, electrode separation, and wall thermal conductivity depend on the regolith type.

The MRE reactor model takes all of these considerations into account. From the standalone analysis of the reactor model in Figure 2-28, the differences in reactor mass and power based on regolith type also depend heavily on operating temperature. It is important to note that the results presented in Figure 2-28 are for a fixed (non-optimized) batch time and design margin, while the optimization scheme was allowed to set those variables in the results presented here.

The optimized mass and power estimates of the MRE-based ISRU system for Highlands and Mare regolith are presented in Figure 3-10. Quite surprisingly, the optimized mass of the entire ISRU system appears to have little dependence on regolith type although a slight decrease in power required for Mare regolith is observed.




Figure 3-10: The optimized mass and power of an MRE-based ISRU system for Highlands and High-Ti Mare regolith types.

Although the total mass and power of the Highlands and Mare ISRU systems were similar, the allocation of the mass and power amongst the subsystems displayed more variation. The mass and power breakdowns shown in Figures 3-7 and 3-9 were normalized by oxygen production level and compared side-by-side, as shown in Figure 3-11.




Figure 3-11: The specific mass and power for each subsystem in the ISRU system for systems optimized for Highlands regolith (solid line) and Mare regolith (dashed line).

For the specific mass of each subsystem (left of Figure 3-11), the Highlands system requires a slightly larger power system due to the larger power requirement for Highlands regolith seen on the right in Figure 3-10. At mid-range production levels, the Mare system requires larger excavation and hopper systems. This comes directly from the oxygen extraction ratio data shown in Figure 2-21. Between operating temperatures of 2000K to 2250K more oxygen can be extracted per kilogram regolith from Highlands regolith as compared to Mare. As shown in Figure 3-6, between production levels of 2000 kg/year and 9000 kg/year the optimal operating temperature is within this range, indicating that a Highlands reactor will have a lower regolith throughput requirement. At higher production levels this discrepancy disappears as the optimal Mare reactor exceeds the 2250K threshold.

An additional result seen in the left plot of Figure 3-11 is that most of the specific mass numbers decrease with production level. That indicates that at higher production level the mass of most subsystems should decrease relative to the production level to enable more mass-efficient oxygen production.

For the reactor specific power shown on the right of Figure 3-11, the primary difference between Highlands and Mare regolith is that Mare reactors produce less radiative heat loss than Highlands. This is due to the fact that Mare regolith has a higher electrical conductivity (see Section 2.2.5), so less heat is generated in the molten region by the electrolysis current. Mare regolith has a higher power requirement for regolith heating (green line), due to the increased regolith throughput requirements discussed two paragraphs above. Highlands regolith requires around 5% more power for chemical electrolysis (blue lines on the right of Figure 3-11). This is due to the fact that Highlands regolith has a lower concentration of iron oxide (FeO), which has the lowest Gibbs Free Energy of Formation out of the primary oxides in lunar regolith. When electrolyzing Highlands regolith, a larger portion of the electrolysis current must be spend electrolyzing higher-energy species, such as silica (SiO₂).

To better understand the effects of regolith type on the optimal ISRU system design, the trends in the optimal design variables were compared for Highlands (solid lines) and High-Ti Mare (dashed lines) regolith as shown in Figure 3-12. A number of important trends are visible in the data. First and foremost, the designs are surprisingly similar across most design variables. The diameters of the Highlands and Mare designs are almost indistinguishable and both optimal designs favor one reactor over many. The operating temperature displays similar trends for both regolith types, but Highlands regolith starts at a slightly higher temperature and at higher production levels appears to favor slightly lower operating temperatures than Mare regolith. This is likely due to the difference in liquidus temperature between the two types of regolith observed in Figure 2-10.

Mare regolith requires more current to produce the same amount of oxygen due to the lower current efficiency observed in Figure 2-21, which also results in a slightly higher required current density for optimal designs. On the other hand, Highlands




Figure 3-12: The optimal design characteristics for an MRE-based ISRU system for Highlands (solid lines) and Mare (dashed lines). Across most design variables, the two designs are reasonably similar.

regolith requires a higher voltage due to the lower electrical conductivity of Highlands regolith (see Section 2.2.5) and the imposed 2 cm minimum electrode separation (see Section 2.6.7). For both regolith types the design margin stays relatively close to 1.1 to enable the feasible electrode separation range observed in the lower left plot of Figure 3-12. The only deviation occurs at low production levels for Mare regolith, where it requires a design margin of ≈ 1.35 . This could possibly be a result of the optimization scheme being unable to locate the optimal solution with a lower design margin.

The biggest difference between Highlands and Mare reactors is the batch time. For Highlands regolith the batch time dramatically rises as the operating temperature crosses 2000K, at which point the Highlands reactor can extract significantly more oxygen per kilogram regolith (see Figure 2-21). Because no such jump occurs with the Mare regolith, the batch time remains somewhat lower and gradually rises with production level. Overall, the comparison of Highlands to Mare regolith reveals that the optimal design does not significantly change with regolith type, though minor differences are observed.

3.9 ISRU System Production Utility

With any ISRU system, it is important to compare the utility of the system to a baseline concept of simply bringing along the resources from Earth. Figure 3-13 shows the annual oxygen production normalized by the mass (left) and power (right) of the complete ISRU plant, which are measures of the plant efficiency. It is clear that at higher production levels an MRE-based ISRU system is able to produce more oxygen per unit plant mass and power. The oxygen production level normalized by system mass increases with production level, indicating that the ISRU system utilizing an MRE reactor can meet higher production levels more efficiently. Within the production levels studied in this work, the maximum efficiency of 6.2 kg oxygen per kilogram ISRU system mass was observed at the maximum production level of 10,000 kg/year. The data indicates that higher numbers should be easily attainable




Figure 3-13: The oxygen production level normalized by ISRU system mass (left) and ISRU system power (right) for both Highlands and High-Ti Mare regolith.

at higher production levels.

The following regression equation, which has a horizontal asymptote as $N \rightarrow \infty$, was fit to each data set shown in Figure 3-13:

$$\begin{aligned}
 N_M &= N_{M\infty} \exp\left(\frac{A_M}{B_M + N}\right) \\
 N_P &= N_{P\infty} \exp\left(\frac{A_P}{B_P + N}\right),
 \end{aligned}
 \tag{3.8}$$

where N_M and N_P are the mass-specific and power-specific production levels shown in Figure 3-13, N is the production level in kg O₂/year, and $N_{M\infty}$, $N_{P\infty}$, A , B are regression coefficients. The regression coefficients $N_{M\infty}$ and $N_{P\infty}$ represent the mass- and power-specific production levels at a production level of infinity (i.e. the value of the horizontal asymptote) and thus provide a good measure of the maximum performance that can be expected from the ISRU system at high production levels. The regression coefficients for Equation (3.8) are given in Table 3.2.

From the coefficient $N_{M\infty}$ in Table 3.2, one can expect to extract oxygen at a maximum rate of 8.8 and 9.2 (kg O₂/year)/(kg system mass) for Highlands and High-Ti Mare regolith, respectively. These numbers are based off of the assumption the there is a limit to the specific oxygen production numbers, which the data in Figure 3-11 does not necessarily suggest. From this work, values on the order of 9.0 (kg O₂/year)/(kg system mass) appear attainable but do not represent any sort

Table 3.2: The regression coefficients for the mass-specific and power-specific oxygen production performance of an MRE-based ISRU system.

	Highlands	High-Ti Mare
$N_{M\infty}$	8.80 ± 0.43	9.21 ± 0.71
A_M	4241 ± 671	4773 ± 1111
B_M	2405 ± 474	2653 ± 739
R_{adj}^2	<i>0.999</i>	<i>0.998</i>
$N_{P\infty}$	310.6 ± 27.4	382.6 ± 57.5
A_P	8454 ± 1707	10890 ± 3277
B_P	4815 ± 915	5878 ± 1553
R_{adj}^2	<i>0.999</i>	<i>0.998</i>

of theoretical limit on the production efficiency of lunar ISRU systems. This means an MRE-based ISRU system can produce its own mass around 9 times each year. Another way to interpret this value is to calculate the number of days until the plant produces its mass in oxygen. Using the data in Figure 3-13, it was determined that at an oxygen production level of 10,000 kg O₂/year, it takes around 59 days for the ISRU system to achieve mass pay-back and produce its mass in oxygen. At a production level of 1,000 kg O₂/yr, it will take 146 days to achieve mass pay-back.

It should be noted that this analysis does not include economic considerations, future work will investigate the price of oxygen produced and the cost of developing and emplacing the ISRU system. For this analysis, examining the mass pay-back point provides a first-order surrogate for determining the tipping point in system utility.

3.10 ISRU System Visualization

A computer-aided design (CAD) model of the ISRU system was created to visualize the optimal system designs. Figure 3-14 illustrates the CAD model of the ISRU system that will produce 1,000 kg O₂/year using Highlands regolith with a 50% operating duty cycle. The “pie”-shaped MRE reactor is pictures in the foreground with the regolith hopper and auger on the left and the YSZ separator on the right. The two oxygen storage tanks are pictured on the right side of the scene and the solar

arrays in the background. A human is included to provide a sense of scale. Recall that the hopper was sized such that it could hold 3 days worth of regolith throughput. The small size of the hopper is indicative of the high oxygen extraction efficiency, and resultant low regolith throughput requirement, of MRE reactors.




Figure 3-14: A CAD model of an ISRU system to produce 1,000 kg O₂ per year., including the solar array power system (left), MRE reactor (front), regolith hopper and auger (front left), YSZ filter (front right) and oxygen storage system (right).

Figure 3-15 displays the ISRU system designed to meet a production level of 10,000 kg O₂/year. The system is naturally larger than the 1,000 kg O₂/year system, but as the mass estimates in Figure 3-7 would suggest, the growth is less than linear with production level.

The CAD models depicted in Figures 3-14 and 3-15 serve to aid the reader in visualizing the MRE-based ISRU system. The dimensions of each component are taken from the analytical sizing models, but the relative location of each subsystem, as well as the design of the minimal support structures for the MRE reactor, solar panels, and oxygen storage tanks are not a direct result of this analysis but rather were estimated to aid in visualizing the system.




Figure 3-15: A CAD model of the ISRU system, including the solar array power system (left), MRE reactor (front), regolith hopper and auger (front left), YSZ separator (front right) and oxygen storage system (right). The 3x3 grid of solar arrays extends out of the image.

3.11 ISRU System Modeling Summary

This chapter presented estimates of the mass and power of an optimized ISRU system to extract oxygen from lunar regolith. The MRE reactor model developed in Chapter 2 was integrated with models for a power system, excavator, hopper, regolith feed system, and oxygen liquefaction and storage systems described in Section 3.2. The inputs and outputs of each model in the integrated system were interconnected to capture the coupled nature of the system, as described in Section 3.3.

This integrated model was leveraged in a hybrid genetic-algorithm/gradient-based optimization scheme, described in Section 3.4, to generate optimized system designs across a range of oxygen production levels as presented in Section 3.5. The trends in the ISRU system mass and power exhibited and were presented in Sections 3.6 and 3.7, and major drivers of mass and power were identified. The dependence on regolith type was explored in Section 3.8, and it was found that an MRE-based ISRU system functions quite similarly on Highlands and Mare regolith.

The utility of an MRE-based ISRU system was assessed to determine that at a production level of 10,000 kg O₂/year, the ISRU system can produce over 6 kg of oxygen per kilogram system mass. Higher production efficiencies, exceeding 9 kg of oxygen per kilogram system mass, can be accessed at higher production levels.

Chapter 4

Conclusions

4.1 Drivers of ISRU System Mass and Power

Several drivers of ISRU system mass are identifiable in Figure 3-7. The oxygen liquefaction and storage system and the MRE reactor comprise 30% of the total system mass each. The mass of the oxygen storage system can be reduced if the storage time of oxygen (6 months in this work) can be reduced based on mission demand. The reactor mass appears to be primarily driven by the dense refractory material that was used to line the interior of the reactor. This material needed to be thick enough to protect the reactor shell from possible random contact with the corrosive molten regolith for short periods of time. Reactor mass could be saved if this material could be thinner in areas where little contact with molten regolith is expected, or if a suitable material with a lower density ($<2700 \text{ kg/m}^3$) can be identified. Reductions in density of the insulation and structural shells would also have a significant impact on reactor mass and therefore ISRU system mass.

The largest driver of reactor power is the radiative heat loss. As shown in Figure 3-9, radiative heat loss accounts for around 50% of the reactor power at 10,000 kg O_2/year . There does not appear to be much that can be done to reduce this driver of reactor power, though it can likely be used constructively. As discussed in Section 3.6, the mental model of simply adding on more insulation to reduce heat loss is not a viable solution for an MRE reactor. With an MRE reactor, there is a heat

source in the central core of the reactor, which is caused by the electrolysis current running through the resistive molten core. The current cannot be reduced without sacrificing production level and the separation between the electrodes cannot grow too small, effectively fixing the resistance between the electrodes for a fixed electrode diameter. With a fixed current and resistance, the power generated via resistive heating cannot be reduced. When the reactor current is on the order of a few kiloamps, the heat generated can be on the order of a few tens of kilowatts. A portion of this heat can go towards heating fresh regolith and making up for the endothermic nature of the electrolysis reaction (though this is not assumed in this study). One promising avenue for mitigating the radiative heat loss are thermoelectric generators, which can be used to convert some of the thermal energy back to electric energy.

The radiative heat loss could be reduced by reducing the electrode separation of the reactor, but there are practical issues that arise when the electrodes become too close together. A more likely solution may involve increasing the reactor diameter to increase the diameter of the molten region, which will lower the resistance of the molten region and thus the radiative heat loss (though this must be done in tandem with a decrease in wall conductivity to offset the trends seen in Figure 2-20). This reduction in reactor heat loss will likely come at the cost of an increase in reactor mass. Although this tradeoff between mass and power was optimized in this model, using a higher specific mass number for the power system may drive the optimal design towards a lower radiative heat loss at the cost of higher reactor mass.

Another possible solution would be tailoring the regolith to increase the conductivity of the melt, as is often done in the metal production industry. One downside of this approach is that it necessitates bringing additional material to facilitate the reaction, and thus it may not be desirable. One promising trend in the radiative heat loss is the significantly sub-linear nature of its growth. Figure 3-9 shows how the radiative power is 24 kW at 5,000 kg/year and only 20% more (29 kW) at double the production level (10,000 kg/year). Thus, radiative heat loss will play a decreasingly important role for a single reactor at higher production levels.

4.2 Optimal MRE Reactor Design Characteristics

The optimal characteristics of an MRE-based ISRU system also yielded some interesting insights. At production levels up to at least 10,000 kg/year, a single larger reactor is preferable to multiple smaller ones, but this may or may not hold at higher production levels beyond the range in this study. For Highlands regolith, the optimal operating temperature gradually rises with production level, increasing oxygen extraction efficiency from 33% of the available oxygen (15 kg O₂/100 kg regolith) at a production level of 1,000 kg/year to 83% of the available oxygen (37.5 kg O₂/100 kg regolith) at a production level of 10,000 kg/year. Extracting 100% of the oxygen in lunar regolith requires prohibitively high temperatures (>2900K) and thus does not appear feasible at this point in time.

The current required is roughly on the order of 1 kA/mT O₂, with voltages around 10 V at low production levels and 3 V at higher production levels. Batch times around 30 hr indicate a preference towards longer residence times to satisfy the 2 cm electrode separation gap requirement. Shorter batch times result in less molten regolith in the reactor at any given time, which makes it more difficult to satisfy the minimum electrode separation requirement.

Operational flexibility for an MRE reactor is incorporated into the model by requiring a range of feasible electrode separation distances that ensure the joule-heated, cold-wall operation. This is brought about via a design margin greater than one, as described in Section 2.6.3. Figure 3-12 demonstrates that a range of electrode separations can be made feasible for both Highlands and Mare regolith. This will be critical in implementing a control system to regulate reactor temperature.

The diameter of the reactor grows from 1.8 m at 1000 kg/year to 3.5 m at 10,000 kg/year. This reactor size will fit within the 5 m Delta IV-heavy payload fairing envelope [8] and the >8 m proposed SLS payload fairing envelope [105]. The power-law equation fit to the trends in reactor diameter (Equation (3.4)) were extrapolated to determine that a reactor capable of producing 186 mT O₂ per year could fit within the proposed SLS payload fairing.

Another important finding from this study is the appropriate scaling of electrode separation with reactor diameter. Figure 2.6.4 shows how the electrode separation must increase with reactor diameter to enable the joule-heated, cold-wall operating condition. As reactor diameter increases, the larger electrode separation is required to generate enough heat to sustain the central molten core. Although an exponential fit was used in this work, perhaps a higher-order polynomial fit would also be appropriate.

4.3 MRE Reactor Feedstock Sensitivity

The comparison of the optimal MRE reactor designs between Highlands and Mare regolith provides some insight into the flexibility of an MRE reactor. The optimal design characteristics shown in Figure 3-12 demonstrate that the optimal design of an MRE reactor does not significantly vary with the type of regolith. This means that MRE reactors appear to be robust to variability in regolith composition without significant performance degradation. Although Highlands regolith results in a slightly higher power requirement, the mass of the entire system (including the power system) is not significantly higher than that of Mare reactors.

4.4 ISRU System Mass and Power Scaling

Our results yield a number of insights concerning the growth of a lunar ISRU system with production level. As discussed in Sections 3.6 and 3.7, both the mass and power of an MRE-based ISRU system exhibit an economy of scale. This means that the mass and the power of the ISRU system grows less than linearly with production level and higher production levels can therefore be achieved more efficiently.

There are a few primary drivers behind the observed economy of scale. For reactor power, the radiative heat loss grows much less than linearly (see Figure 3-9). That is, the radiative heat loss for an MRE reactor does not grow significantly with production level.

For reactor mass, the oxygen storage system is a cylindrical tank with hemisphere end caps, which inherently has a less than linear growth in mass with storage capacity. Also, as observed in Figure 3-8, the reactor mass grows less than linearly with production level, contributing to the economy of scale. This is due to the cylindrical nature of the reactor. Furthermore, because the power exhibits an economy of scale, the mass of the power system naturally exhibits an economy of scale as well.

4.5 The Utility of Lunar ISRU

The total mass of the ISRU system capable of producing 10,000 kg of oxygen each year is approximately 1,600 kg. For comparison, this is significantly less than the wet mass of the Apollo ascent stage (4,819 kg [6]). Similarly, the maximum diameter of the reactor from this study (3.5 m) is also less than the diameter of the descent module (4.2 m [6]).

On a more rigorous level, the oxygen production normalized by system mass (termed “production efficiency”) can be used to assess the utility of a lunar ISRU system in a given mission architecture. Ishimatsu et al. [55] studied the value of lunar ISRU in a Mars exploration architecture and found that production efficiencies above 5.0 (kg O₂/yr)/(kg system mass) were required to make lunar ISRU pay off in a Mars exploration scheme. They do acknowledge that their model does not take into account the time-cost of deploying the lunar ISRU infrastructure, so this number must be taken as a rough estimate rather than a defining quantity. Ho et al. [53] addressed the time cost of emplacing lunar ISRU and found that for a production efficiency of >1.0 (kg O₂/yr)/(kg system mass), lunar ISRU can be used to reduce the total mass that must be launched from Earth for a Mars exploration architecture. For production efficiencies above 5.0 (kg O₂/yr)/(kg system mass), the mass of the lunar ISRU system actually exceeded that of the Mars ISRU system. Naturally, all of these numbers depend on the assumptions made, but they provide some indication of the production efficiency that a lunar ISRU system must be capable of to reducing the total mass that must be launched from Earth for a Mars exploration campaign.

The production efficiency values calculated from the analytical hardware sizing models in this study (shown in Figure 3-13) indicate that an MRE-based ISRU system can indeed exceed these thresholds. The regression curves in Equation (3.8) predict production levels on the order of 9.0 ((kg O₂/year)/kg system mass) at high production levels for both Highlands and Mare regolith, and this value does not necessarily represent a ceiling on production efficiency. This means that an MRE-based ISRU system is expected to produce the entire system's mass around 9 times each year. **Our results, when combined with two previous analyses [55, 53], indicate that lunar ISRU has a high potential for reducing the total mass required for a Mars exploration campaign.**

4.6 MRE Reactor Future Work

Fixed Reactor Design Variables

The design variables considered in this study included reactor diameter, electrode separation, and wall thermal characteristics (thermal resistance). Future work can examine the impact of the ratio of electrode diameter to reactor diameter. A preliminary case study revealed that a value of 15% provided a suitable balance between the molten mass and operating temperature requirements. This value agrees well with range of 0.083 to 0.017 used in industrial electric slag furnaces [60]. Future work can further examine the impact of and optimal value for this design parameter.

This study also utilized a uniform thermal design around the entire exterior of the reactor – the thermal conductivity and wall emissivity were the same on the top, bottom, and sides of the MRE reactor. Future work can analyze the optimal insulation topology of the reactor, because allowing for different thermal resistances on the top, bottom, and sides of the reactor may provide an avenue for extending production range and efficiency.

An interesting avenue for future work would be to address operating pressure. For this study, the MRE model used a fixed operating pressure of 101.3 kPa, but lower pressures may improve reaction kinetics and enable a thinner structural shell

on the reactor. Additionally, a dependence upon operating pressure for the regolith material properties could also be added to the model to better understand the impact of operating pressure.

Multiphysics Simulation

The multiphysics simulation was extended from that presented in Schreiner et al. [95] to include current values of up to 3 kA, which equates to an oxygen production level of around 2500 kg O₂/year. Future work can generate addition data at higher currents to better understand the reactor scaling at higher production levels.

In terms of the multiphysics simulation, convective heat transfer was omitted from the model for two reasons. First, it would increase the simulation convergence time by several orders of magnitude, which would make the time to generate the data for the design methodology unreasonably long. Second, there is some uncertainty as to whether or not the thermal conductivity data presented in Section 2.2.4 from Eisenhüttenleute [38] contains a convective contribution as well. It is possible that this data represents both the convective and conductive heat transfer in the molten materials. Future work can further investigate this issue to ensure that the thermal modeling is accurate. Surface tension may also play a larger role in the 1/6-g lunar environment, so modeling this phenomenon may also prove useful.

Electrode Designs

More complex anode and cathode designs can also be utilized to increase reactor performance. The shaft-and-plate designs used in this study were chosen because they could be implemented in an axially-symmetric simulation for fast solver convergence, but Sibille and Dominguez [99] demonstrated that a waffle-shape anode can increase current density by around 33%.

Modeling anode and cathode life is also an area of future work. Although the MIT experiments determined that Iridium anodes experienced corrosion rates on the order of 7.7 mm/year [124], this rate is derived from a relatively short operational run at low current. As more material lifetime data becomes available, the MRE reactor

model could be updated to size the anode and cathode based on a certain design life to ensure that the anode remains intact throughout operation.

Operational Considerations

As mentioned in Section 2.5.3, the reactor model currently attempts to operate in a manner in which the voltage and current are varied inversely to achieve constant power operation. This results in approximately the same thermal performance, but the amount of power going towards breaking the oxide chemical bonds and endothermic heating will vary with the oxide species being preferentially electrolyzed. Future work can adjust the constant-power operational methodology to create a constant-heat generation power mode, such that the reactor performs similarly from a thermal perspective no matter which species is being electrolyzed.

The density of the leftover slag decreases while the density of the molten metal pool increases as more SiO_2 is reduced. This can lead to a potential density inversion issue, especially for Highlands regolith. Future work can address this issue in more detail, as this model was limited to discrete oxide species reduction, while in reality multiple oxide species will be reduced concomitantly, perhaps rendering the density issue irrelevant.

The issue of reactor start up also must be included in the reactor model. For this study, a day-only power source with no energy storage was assumed, such that the reactor is operating only during the day. This is not ideal, as it requires frequent reactor cycling for start-up and shutdown. Solar concentrators could perhaps be used to supply the start-up thermal energy, though industrial electrolysis cells are often started by striking an electric arc in a manner similar to an electric arc furnace. Future work can address the tradeoff of utilizing a solar power source without energy storage that requires a start-up/shutdown procedure versus a continuous power source that undergoes start-up and shutdown much less frequently.

From an electrochemical perspective, the electrolysis process can certainly be better understood. Utilizing proper activity coefficients for the Gibbs Free Energy and Enthalpy of Formation data, to account for the phase differences and slag composi-

tion would be a good addition to the MRE reactor model. For instance, the Gibbs Free Energy of MgO and SiO₂ become quite close around temperature of 2200K – but these would be affected by the fact that there is relatively more SiO₂ than MgO.

The electrical conductivity was assumed to remain constant throughout a batch. Although this may prove relatively true for a continuous-operation scheme, a higher fidelity simulation could also model the changing electrical conductivity in a manner similar to the liquidus temperature analysis presented in Section 2.3. Haskin et al. [45] conducted an analysis of this nature, which could be used as a starting point for this future work. The optical absorption length, as well as other regolith material properties also will depend on the changing melt composition throughout the electrolysis process – including this dependence in the model would prove useful.

The withdrawal of molten metal from an MRE reactor also was not included in this model – although the molten metal may simply be tapped out in a manner similar to that used by the steel industry, more complex techniques have been investigated [106]. Whichever technique is used, the mass associated with the hardware must be included.

4.6.1 Low Power Mode

The trends in reactor power observed in Figure 3-9 do not scale down to low production levels well. We caution against the use of this reactor model for production levels below ≈ 100 kg/year, as the model was primarily designed to understand the scaling of reactor mass and power to production levels in a full mission support scenario. Thus, further work will be required to determine how to effectively design low-power MRE reactors for early technology demonstration missions. As mentioned in Section 3.6, lower power may be achieved by somehow increasing the conductivity of the molten regolith melt (perhaps by adding a supporting electrolyte). Using closer electrode separation distances (<2 cm) will also bring down the total power if such distances can be achieved while avoiding shorting out the reactor.

4.7 ISRU System Future Work

There are a number of items that can be addressed in future work in addition to those listed for the MRE reactor model in Section 4.6. The excavator system model currently does not produce an estimate of the energy consumed by the excavator, which would be an important addition to future models. Since the model's creation, newer excavation theory and models have also been developed [92, 128, 129], which can be used to build a higher fidelity excavation model.

As mentioned in Section 3.2.4, the auger model is not yet parametrically sized to meet a given regolith insertion mass and time. Future work can dynamically size the radius and rotation rate of the auger to meet a specified insertion time that is compatible with the reactor model. This subsystem coupling would better inform an optimal reactor fill time and batch time. Furthermore, the mass of the motor to drive the auger is yet unaccounted for in the system model, although the power is calculated and taken into account.

Future design iterations can also focus on including a spare parts analysis to more accurately determine the holistic mass of a less-than-ideal ISRU system.

Naturally, the mass payback ratio used in Section 3.9 serves only as a surrogate for determining the actual feasibility and applicability of a lunar ISRU system. Blair et al. [14] noted that *“a positive mass-payback relationship does not guarantee an economic benefit.”* A proper mission-level architecture analysis is required to assess the true benefit of using lunar ISRU for a particular mission scenario. The specific production levels shown in Figure 3-13 serve as a good interface between the hardware design-level analysis presented here and the architecture study analyses.

Another avenue of future work is to integrate other reactor models, such as those for Hydrogen Reduction of Ilmenite [48] or Carbothermal Reduction of Silicates [11] discussed in Section 1.1.3. This would allow for a more quantitative comparison between these three techniques to better understand the proper applicability of each oxygen extraction method.

One clear avenue of future work is investigating constructive uses for the significant

radiative heat loss from the reactor. Although this heat must be produced due to resistive heating within the reactor molten region, it does not necessarily have to go to waste. One obvious use would be preheating the regolith prior to insertion into the reactor. Thermoelectric generators may also be able to recuperate a large portion of the radiative heat loss, dramatically reducing the total power required to produce oxygen.

Appendix A

Cutoff Line Design Justification

As mentioned in Section 2.6.1, the reactor design is largely driven by the cutoff lines between infeasible and feasible designs. Here, we provide a brief justification for designing on these cutoff lines. An examination of Figure 2-15 may cause one to wonder if perhaps designing further away from the cutoff line could result in more optimal designs.

To this end, the heat loss of the reactor was studied to better understand the impact of the cutoff lines. Figure A-1 shows the heat loss (left) and heat loss divided by the mass of molten regolith in the reactor (right) for a range of diameters (x-axis) and electrode separations (each line). From the heat loss plot (left), it becomes immediately clear that moving away from the cutoff line directly increases heat loss. From this perspective, it would appear that staying on the cutoff line results in the minimal heat loss. The plot of heat loss divided by molten mass (right) sheds additional light on the issue. It is straightforward to see that moving away from the cutoff line while staying on a line of constant electrode separation results in an increase in heat loss per molten mass.

There is a caveat to this trend, in that moving away from the cutoff line while increasing electrode separation (moving from the green to the black line) has the possibility of decreasing the heat loss per molten mass. Thus, if we have a required amount of molten mass (from the oxygen production level) it may be possible to reduce the heat loss by moving slightly off the cutoff line. Although this may result is

slightly less heat loss per kilogram molten regolith in the reactor, this effect is not very significant. That is, the slope of this line is quite shallow and moving significantly away from the cutoff line quickly encounters a dramatic increase in heat loss per kilogram regolith.




Figure A-1: The heat loss (left) and heat loss divided by molten mass (right) for an MRE reactor. Infeasible designs, where molten material has touched the wall, are crossed out in red.

Appendix B

Derivation of Nonlinear Regression Equations

When fitting nonlinear, multivariate equations to the data from the multiphysics simulation, it was necessary to determine the appropriate form of the nonlinear equations. Figure B-1 shows the radial profiles for temperature and phase (solid vs. molten) in the regolith in the MRE reactor simulated using the COMSOL simulation described in Section 2.5. The four different colored lines represent four different values of the wall thermal conductivity (all references to “conductivity” in this section refer to thermal conductivity).

A few key conclusions can be drawn from Figure B-1. First, the temperature at the center of the reactor appears to remain relatively constant with changing wall thermal conductivity. There is a slight decrease in the central temperature with higher wall conductivity, but this effect appears to be on the order of a few kelvin. Second, the exterior wall temperature remains virtually fixed irrespective of wall conductivity. This is somewhat intuitive, the reactor has to dissipate the heat generated by the electrolysis current passing through the resistive melt. The value of this required heat loss does not depend on wall conductivity and must be transferred to the environment via radiation, which is proportional to the external wall temperature (T_{out})⁴. Thus, T_{out} should not depend strongly on wall conductivity.

Figure B-1 shows that increasing the conductivity of the reactor lowers the slope

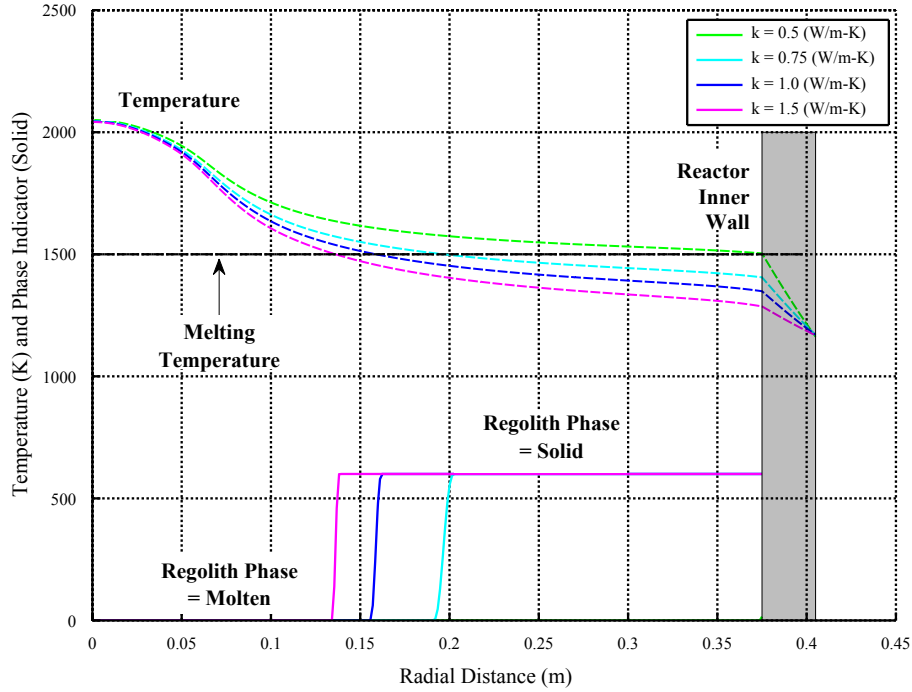


Figure B-1: The radial temperature and phase indicator profiles from the multiphysics simulation of an MRE reactor with different wall thermal conductivities.

of the temperature profile within the wall, which lowers the temperature profile in the regolith. The drop in temperature within the reactor results in a smaller molten zone. To generate an analytical relationship for this effect, a linearized model of the radial temperature profile (and phase indicator) was developed, shown in Figure B-2. This model assumes that the temperature at the outer wall and central core of the reactor remains fixed regardless of wall conductivity. Furthermore, the model assumes a linear temperature profile within the regolith and the reactor wall.

Under these assumptions, we can derive an analytical relationship between wall conductivity, diameter and molten mass within the reactor. The radial temperature distribution within the regolith, $T(r)$, can be modeled as:

$$T(r) = T_{center} - (T_{center} - T_{wall-in}) \left(\frac{r}{r_{reactor}} \right), \quad (\text{B.1})$$

where T_{center} is the temperature at the center of the reactor, $T_{wall-in}$ is the temperature at the inner wall of the reactor, and $r_{reactor}$ is the radius of the inner reactor wall.

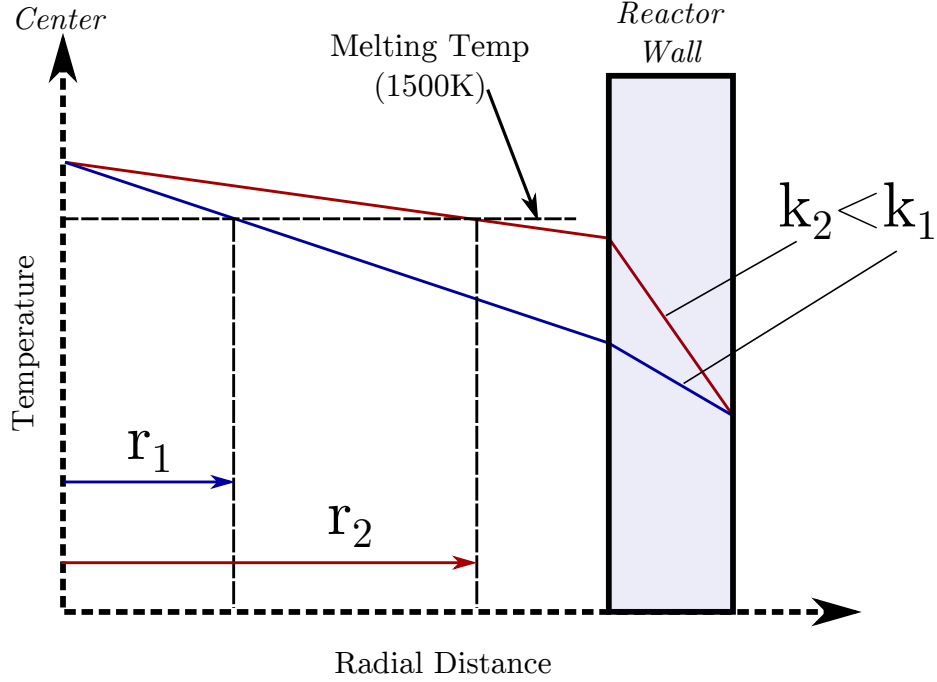


Figure B-2: A linearized representation of the radial temperature profile.

Using the conductive heat equation with a fixed outer wall temperature ($T_{wall-out}$), the temperature at the inner reactor wall can be calculated as:

$$T_{wall-in} = T_{wall-out} + \frac{\dot{Q}}{k_{wall}} \Delta x, \quad (\text{B.2})$$

where \dot{Q} is the heat loss through the wall, k_{wall} is the wall conductivity, and Δx is the wall thickness. Substituting Equation B.2 into Equation B.1, we get:

$$T(r) = T_{center} - \left((T_{core} - (T_{wall-out} + \frac{Q}{k} \Delta x)) \left(\frac{r}{r_{reactor}} \right) \right). \quad (\text{B.3})$$

Equation B.3 can then be used to solve for the radius of the phase boundary by calculating at what radius the temperature crosses the solidification temperature [$r_{solid} \equiv (T(r_{solid}) = 1500K)$]:

$$r_{solid} = \frac{(T_{center} - 1500K)}{T_{center} - (T_{wall-out} + \frac{\dot{Q}}{k_{wall}} \Delta x)} (r_{reactor}). \quad (\text{B.4})$$

The molten mass (MM) within the reactor will be proportional to the radius of

the molten region squared, resulting in the following relation:

$$MM \propto (r_{solid})^2 = \left(\frac{mm_1}{mm_2 + \frac{mm_3}{k_{wall}}} * D \right)^2, \quad (\text{B.5})$$

where mm_i are regression coefficients and D is the inner diameter of the reactor.

Here we note that molten mass will not only be proportional to the radius squared (as shown in Equation (B.5)), but also to the electrode separation. The electrode separation is also proportional to diameter, due to the geometry scaling of the reactor to stay on the cutoff line, as shown in Figure 2.6.4. Thus, Equation B.5 was not used as the regression equation, but rather was used as a starting point for developing the nonlinear, empirical regression equations presented in Section 2.6.2.

Appendix C

Regression Equation Coefficients

This appendix presents the regression coefficients to the nonlinear, multivariate equations that were fit to the multiphysics simulation data.

Table C.1: The regression coefficients for Equation 2.20, which predicts the molten mass within the reactor based off of data from the multiphysics simulation presented in Section 2.6.

T_{wall}	Regolith	m_1	m_2	m_3	m_4	m_5	m_6
1200	Highlands	20.53024	8.990361	0.4562	0.313049	-0.26814	1.546763
	Mare	12.81811	4.474273	0.648773	1.131599	0.438686	1.298745
1300	Highlands	14.39169	5.702063	0.486305	1.123273	0.149667	1.253424
	Mare	16.20177	5.417753	0.589564	0.686972	0.085608	1.555261
1400	Highlands	11.03747	5.313277	0.159837	4.987098	0.34287	0.435224
	Mare	9.731733	3.97475	0.226271	5.605015	0.499934	0.485324

Table C.2: The regression coefficients for Equation 2.21, which predicts the operating temperature within the reactor based off of data from the multiphysics simulation presented in Section 2.6.

T_{wall}	Regolith	t_1	t_2	t_3	t_4	t_5	t_6	t_7
1200	Highlands	1429.526	2209.212	0.277	-0.046	0.110	0.250	-3.298
	Mare	1427.554	4662.279	0.390	-0.027	0.114	0.194	-4.280
1300	Highlands	1443.568	11406.560	0.294	-0.022	0.015	0.145	-8.544
	Mare	1527.021	3063.317	0.483	-0.038	0.061	0.304	-2.625
1400	Highlands	1525.286	5387.139	0.351	-0.029	0.010	0.215	-5.088
	Mare	1598.214	2200.344	0.593	-0.043	0.039	0.419	-1.538

Table C.3: The regression coefficients for Equation 2.26, which predicts the required electrode separation distance to maintain thermal equilibrium in an MRE reactor based off of data from the multiphysics simulation presented in Section 2.6.

T_{wall}	Regolith	e_1	e_2	e_3	e_4	e_5	e_6	e_7	e_8
1200	Highlands	0.0011	-0.1945	2.5738	0.6487	0.4656	0.3522	0.0570	0.4348
	Mare	0.0011	-0.0540	3.1514	0.9572	0.5612	0.3292	0.0047	0.5556
1300	Highlands	0.0011	-0.0941	3.2402	0.9433	0.5211	0.3555	0.0377	0.5009
	Mare	0.0010	-0.0319	3.8635	1.0890	0.5747	0.3336	0.0088	0.5502
1400	Highlands	0.0011	-0.0526	3.9239	1.1214	0.5400	0.3491	0.0500	0.4964
	Mare	0.0008	-0.0458	4.6638	1.3147	0.5827	0.3243	0.0064	0.5857

Table C.4: The regression coefficients for Equation 2.28, which predicts the expected heat loss from an MRE reactor based off of data from the multiphysics simulation presented in Section 2.6.

T_{wall}	Regolith	h_1	h_2	h_3	h_4	h_5	h_6
1200	Highlands	1475293	-8.863	1.144	1.641	164.954	-0.033
	Mare	1239262	-8.092	1.042	1.626	158.910	-0.032
1300	Highlands	1591740	-8.890	1.161	1.307	164.021	-0.036
	Mare	1396795	-8.191	1.065	1.338	178.379	-0.034
1400	Highlands	1591221	-8.289	1.083	1.146	222.399	-0.036
	Mare	1756455	-8.843	1.165	1.135	188.927	-0.034

Table C.5: The regression coefficients for Equation 2.32, which predicts the expected operating voltage for an MRE reactor based off of data from the multiphysics simulation presented in Section 2.6.

T_{wall}	Regolith	v_1	v_2	v_3	v_4	v_5	v_6
1200	Highlands	2611.8	-11.795	1.672	1.954	0.085	-0.021
	Mare	1000.3	-6.952	0.893	1.677	0.156	-0.028
1300	Highlands	1997.4	-10.037	1.477	1.561	-0.088	-0.027
	Mare	1112.8	-6.897	0.908	1.450	0.150	-0.027
1400	Highlands	1560.3	-8.078	1.190	1.368	-0.139	-0.031
	Mare	1056.0	-6.358	0.855	1.240	0.143	-0.028

Bibliography

- [1] R. H. Aiken. *Process of making iron from the ore*. US Patent 816,142. Mar. 1906 (cit. on pp. 31, 54).
- [2] A. Allanore, L. Yin, and D. R. Sadoway. “A new anode material for oxygen evolution in molten oxide electrolysis”. In: *Nature* (2013) (cit. on pp. 30, 32, 88).
- [3] B. Altenberg. “Processing lunar in-situ resources”. In: *Technical Research and Development Project Job* (1990). 90634-002 (cit. on pp. 24, 26, 33).
- [4] M. S. Anderson, M. K. Ewert, J. F. Keener, S. A. Wagner, and L. B. CTSD. “Life Support Baseline Values and Assumptions Document”. In: (2015) (cit. on pp. 113, 114).
- [5] D Andrews, A Colaprete, J Quinn, D Chavers, and M Picard. “Introducing the Resource Prospector (RP) Mission”. In: *AIAA Space 2014 Conference*. 2014 (cit. on p. 30).
- [6] *Apollo 11 Lunar Module / EASEP*. URL: <http://nssdc.gsfc.nasa.gov/nmc/spacecraftDisplay.do?id=1969-059C> (cit. on p. 145).
- [7] C. A. Arregoitia et al. “Full Moon Storage and Delivery of Oxygen and Hydrogen”. MA thesis. International Space University, 2007 (cit. on p. 38).
- [8] *Atlas V and Delta IV: Technical Summary*. URL: http://www.ulalaunch.com/uploads/docs/Launch_Vehicles/AV_DIV_product_card.pdf (cit. on p. 143).
- [9] V. Badescu. *Moon: Prospective Energy and Material Resources*. Springer, 2012 (cit. on pp. 22, 29).
- [10] R Balasubramaniam, S Gokoglu, and U Hegde. “Determination of Chemical Kinetic Rate Constants of a Model for Carbothermal Processing of Lunar Regolith Simulant Using Methane”. In: *AIAA Proceedings* (2009) (cit. on pp. 33, 37).
- [11] R Balasubramaniam, S Gokoglu, and U Hegde. “The reduction of lunar regolith by carbothermal processing using methane”. In: *International Journal of Mineral Processing* 96.1 (2010), pp. 54–61 (cit. on pp. 33, 40, 150).
- [12] A. Basu and S. Riegsecker. “Modal mineralogic distribution in the regolith at Apollo landing sites”. In: *Journal of Geophysical Research: Planets* (1991–2012) 105.E2 (2000), pp. 4361–4368 (cit. on pp. 44, 55).

- [13] H. Belachgar et al. “FERTILE Moon: Feasibility of Extraction of Resources and Toolkit for In-situ Lunar Exploration”. MA thesis. International Space University, 2006 (cit. on pp. 30, 38, 114, 126).
- [14] B. R. Blair et al. “Space Resource Economic Analysis Toolkit: The Case for Commercial Lunar Ice Mining”. In: *Final Report to the NASA Exploration Team* (2002) (cit. on pp. 38, 150).
- [15] E. Bock et al. *Lunar resources utilization for space construction*. Tech. rep. NAS9-15560. General Dynamics Convair Division, Apr. 1979 (cit. on p. 129).
- [16] J. R. Brophy and T. Larson. “30-kW SEP spacecraft as secondary payloads for low-cost deep space science missions”. In: *International Electric Propulsion Conference 2013*. 2013, pp. 2013–405 (cit. on p. 113).
- [17] J. R. Brophy, R. Gershman, N. Strange, D. Landau, R. G. Merrill, and T. Kerslake. “300-kW solar electric propulsion system configuration for human exploration of near-earth asteroids”. In: *AIAA Paper* 5514 (2011), p. 2011 (cit. on p. 113).
- [18] G. K. Burgess and R. G. Waltenberg. *The emissivity of metals and oxides*. US Government Printing Office, 1915 (cit. on p. 65).
- [19] B. Carr. “Recovery of water or oxygen by reduction of lunar rock”. In: *AIAA Journal* 1.4 (Feb. 1963), pp. 921–924 (cit. on pp. 23, 31).
- [20] W. F. Carroll. “Research on the use of Space Resources”. In: *JPL Publication* 83-36 (Mar. 1983) (cit. on pp. 51, 88).
- [21] M. W. Chase and J. A. N. A. Force. “NIST-JANAF thermochemical tables”. In: (1998) (cit. on pp. 54, 55).
- [22] G. Z. Chen, D. J. Fray, and T. W. Farthing. “Direct electrochemical reduction of titanium dioxide to titanium in molten calcium chloride”. In: *Nature* 407.6802 (2000), pp. 361–364 (cit. on p. 54).
- [23] A. Chepko. “Technology selection and architecture optimization of in-situ resource utilization systems”. PhD thesis. Massachusetts Institute of Technology, 2009 (cit. on pp. 33, 41, 114).
- [24] A. Chepko, O. de Weck, W. Crossley, E. Santiago-Maldonado, and D. Linne. “A Modeling Framework for Applying Discrete Optimization to System Architecture Selection and Application to In-Situ Resource Utilization”. In: *12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference* (Sept. 2008) (cit. on pp. 37, 40).
- [25] A. Chepko, O. de Weck, D. Linne, E. Santiago-Maldonado, and W. Crossley. “Architecture Modeling of In-Situ Oxygen Production and its Impacts on Lunar Campaigns”. In: *AIAA SPACE 2008 Conference & Exposition*. 2008 (cit. on p. 37).
- [26] E. L. Christiansen et al. *Conceptual Design of a Lunar Oxygen Pilot Plant*. Tech. rep. NASA 9-17878. Eagle Engineering Inc., 1988 (cit. on pp. 22, 24, 36, 37, 44, 129).

- [27] A. Colaprete et al. “Detection of water in the LCROSS ejecta plume”. In: *Science* 330.6003 (2010), pp. 463–468 (cit. on p. 30).
- [28] A. J. Colozza. “Analysis of lunar regolith thermal energy storage”. In: *NASA STI/Recon Technical Report N 92* (1991), p. 14480 (cit. on pp. 48, 55).
- [29] R. O. Colson and L. A. Haskin. “Lunar oxygen and metal for use in near-Earth space: Magma electrolysis”. In: *In Arizona Univ., NASA Space Engineering Research Center for Utilization of Local Planetary Resources*. Vol. 1. 1990 (cit. on pp. 34, 50, 130).
- [30] P. Curreri et al. “Process demonstration for lunar in situ resource utilization—Molten oxide electrolysis”. In: *NASA Marshall Space Flight Center. MSFC Independent Research and Development Project* (Aug. 2006). Project No. 5-81 (cit. on pp. 25, 29, 32, 87, 88).
- [31] J. Diaz et al. *Space Transportation Architectures and Refueling for Lunar and Interplanetary Travel and Exploration (STARLITE)*. Tech. rep. NASA Grant NAG9-1535. Center for the Commercial Applications of Combustion in Space, Colorado School of Mines, June 2005 (cit. on pp. 21, 22).
- [32] J. A. Dominguez, L. Sibille, and S. Poizeau. “Modeling Joule Heating Effect on Lunar Oxygen Generation via Electrolytic Reduction”. In: *AIAA 47th Aerospace Sciences Meeting Proceedings on Disc*. 2009 (cit. on pp. 32, 34, 35).
- [33] A. Ducret, D. Khetpal, and D. R. Sadoway. “Electrical conductivity and transference number measurements of FeO-CaO-MgO-SiO₂ melts”. In: *Electrochemical Society Meeting, Philadelphia*. 2002, pp. 347–53 (cit. on pp. 50, 51).
- [34] M. B. Duke, J. Diaz, B. R. Blair, M. Oderman, and M. Vaucher. “Architecture Studies for Commercial Production of Propellants From the Lunar Poles”. In: *AIP Conference Proceedings*. AIP Publishing. 2003, pp. 1219–1226. DOI: 10.1063/1.1541422 (cit. on p. 23).
- [35] P. Eckart. “Lunar base parametric model”. In: *Journal of Aerospace Engineering* 10.2 (1997), pp. 80–90 (cit. on p. 37).
- [36] P. Eckart. *Parametric model of a lunar base for mass and cost estimates*. Vol. 1. Herbert Utz Verlag, 1996 (cit. on pp. 21, 37).
- [37] P. Eckart, B. Aldrin, A. Clarke, H. H. Schmitt, and J. Young. *The Lunar Base Handbook: an introduction to lunar base design, development, and operations (2nd Edition)*. McGraw-Hill, 2006 (cit. on pp. 23, 37).
- [38] V. D. Eisenhüttenleute. *Slag Atlas*. Verlag Stahleisen GmbH, 1995 (cit. on pp. 49, 57, 58, 147).
- [39] M. Faraday. “XXV. Experimental researches in electricity—Seventh series”. In: *The London and Edinburgh Philosophical Magazine and Journal of Science* 5.27 (1834), pp. 161–181 (cit. on pp. 53, 59).

- [40] C. A. Gallo, R. A. Wilkinson, R. P. Mueller, J. Schuler, and A. Nick. “Comparison of ISRU excavation system model blade force methodology and experimental results”. In: *American Institute of Aeronautics and Astronautics (AIAA)* (2009) (cit. on p. 109).
- [41] D. G. Gilmore and M. Bello. “Satellite thermal control handbook”. In: (1994) (cit. on p. 65).
- [42] D. G. Gilmore and M. Donabedian. *Spacecraft Thermal Control Handbook: Cryogenics*. Vol. 2. AIAA, 2003 (cit. on p. 84).
- [43] A. J. Gmitter. “The influence of inert anode material and electrolyte composition on the electrochemical production of oxygen from molten oxides”. PhD thesis. Massachusetts Institute of Technology, 2008 (cit. on pp. 29, 32, 85, 87, 88).
- [44] A. J. Hanford. “Advanced life support baseline values and assumptions document”. In: (2006) (cit. on p. 114).
- [45] L. A. Haskin, R. O. Colson, D. J. Lindstrom, R. H. Lewis, and K. W. Semkow. “Electrolytic smelting of lunar rock for oxygen, iron, and silicon”. In: *Lunar Bases and Space Activities of the 21st Century*. Vol. 1. 1992, pp. 411–422 (cit. on pp. 31, 61, 88, 149).
- [46] R. Hatakenaka, T. Miyakita, H. Sugita, M. Saitoh, and T. Hirai. “Thermal performance and practical utility of a MLI blanket using plastic pins for space use”. In: *American Institute of Aeronautics and Astronautics* (2013), pp. 1–11 (cit. on p. 84).
- [47] U Hegde, R Balasubramaniam, and S Gokoglu. “Heating-Rate-Coupled Model for Hydrogen Reduction of”. In: *AIAA Proceedings* (2010) (cit. on pp. 33, 39).
- [48] U. Hegde, R Balasubramaniam, and S. Gokoglu. “Development and Validation of a Model for Hydrogen Reduction of JSC-1A”. In: *AIAA Proceedings* (2009) (cit. on pp. 33, 37, 40, 150).
- [49] G. H. Heiken, D. T. Vaniman, and B. M. French. *The Lunar sourcebook: A user’s guide to the Moon*. CUP Archive, 1991 (cit. on pp. 44–46, 56, 70).
- [50] S Heiroth, T. Lippert, A Wokaun, et al. “Microstructure and electrical conductivity of YSZ thin films prepared by pulsed laser deposition”. In: *Applied Physics A* 93.3 (2008), pp. 639–643 (cit. on p. 108).
- [51] B. Hemingway, R. Robie, and W. Wilson. “Specific heats of lunar soils, basalt, and breccias from the Apollo 14, 15, and 16 landing sites, between 90 and 350 K”. In: *Lunar and Planetary Science Conference Proceedings*. Vol. 4. 1973, p. 2481 (cit. on pp. 46, 48).
- [52] A. F. Hepp, D. L. Linne, G. A. Landis, M. F. Wade, and J. E. Colvi. “Production and use of metals and oxygen for lunar propulsion”. In: *Journal of Propulsion and Power* 10.6 (1994), pp. 834–840 (cit. on pp. 37, 113, 128).

- [53] K. Ho, O. L. de Weck, J. A. Hoffman, and R. Shishko. “Dynamic modeling and optimization for space logistics using time-expanded networks”. In: *Acta Astronautica* 105.2 (2014), pp. 428–443 (cit. on pp. 145, 146).
- [54] D. J. Hoffman, T. W. Kerslake, A. F. Hepp, M. K. Jacobs, and D. Ponnusamy. “Thin-film photovoltaic solar array parametric assessment [space power]”. In: *Energy Conversion Engineering Conference and Exhibit, 2000.(IECEC) 35th Intersociety*. Vol. 1. IEEE. 2000, pp. 670–680 (cit. on p. 113).
- [55] T. Ishimatsu, O. De Weck, J. Hoffman, Y. Ohkami, and R. Shishko. “A Generalized Multi-Commodity Network Flow Model for Space Exploration Logistics”. In: *AIAA SPACE 2013 Conference & Exposition—San Diego, USA*. 2013 (cit. on pp. 145, 146).
- [56] N. Jarrett, S. Das, and W. Haupin. “Extraction of oxygen and metals from lunar ores”. In: *Space Sol. Power Rev.:(United States)* 1.4 (1980) (cit. on pp. 27, 31).
- [57] G. Johnson. “Recommendations for Utilization of Lunar Resources by the Working Group on Extraterrestrial Resources”. In: *Jet Propulsion Lab.(Pasadena, Calif.)* (1963) (cit. on p. 24).
- [58] B. K. Joosten and L. A. Guerra. “Early lunar resource utilization: a key to human exploration”. In: *AIAA Space Programs and Technologies Conference, AIAA*. 1993, pp. 93–4784 (cit. on p. 23).
- [59] M. Kallerud, B. Nguyen, T. Paladin, and A. Wilson. “In-Situ Resource Utilization: Investigation of Melted Lunar Regolith Simulant JSC-1A”. In: *Proceedings of the Wisconsin Space Conference*. 2011 (cit. on p. 46).
- [60] M. W. Kennedy. “Electric Slag Furnace Dimensioning”. In: *International Smelting Technology Symposium: Incorporating the 6th Advances in Sulfide Smelting Symposium*. John Wiley & Sons, Inc. 2012, pp. 279–290 (cit. on pp. 88, 146).
- [61] D. G. Kesterke. “Electrowinning oxygen from silicate rocks”. In: *NASA Special Publication 229* (1970), p. 139 (cit. on pp. 31, 88).
- [62] H. Kim, J. Paramore, A. Allanore, and D. R. Sadoway. “Electrolysis of molten iron oxide with an iridium anode: the role of electrolyte basicity”. In: *Journal of The Electrochemical Society* 158.10 (2011), E101–E105 (cit. on pp. 87–89).
- [63] H. H. Koelle and B. Jochenning. *Lunar base simulation*. TU, Inst. für Luft-und Raumfahrt, 1982 (cit. on pp. 36, 37).
- [64] H. Kojitani and M. Akaogi. “Measurement of heat of fusion of model basalt in the system Diopside-forsterite-anorthite”. In: *Geophysical research letters* 22.17 (1995), pp. 2329–2332 (cit. on p. 55).
- [65] F. Kreith, R. Manglik, and M. Bohn. *Principles of heat transfer*. Cengage learning, 2010 (cit. on pp. 83, 85).
- [66] G. A. Landis. “Materials Refining for Solar Array Production on the Moon”. In: *NASA Technical Memorandum 214014* (2005) (cit. on p. 29).

- [67] G. A. Landis. “Technology for solar array production on the Moon”. In: *Photovoltaic Specialists Conference, 2002. Conference Record of the Twenty-Ninth IEEE*. IEEE. 2002, pp. 796–798 (cit. on p. 29).
- [68] R. Lindemann and C. Voorhees. “Mars Exploration Rover mobility assembly design, test and performance”. In: *Systems, Man and Cybernetics, 2005 IEEE International Conference on*. Vol. 1. Oct. 2005, 450–455 Vol. 1. DOI: 10.1109/ICSMC.2005.1571187 (cit. on p. 110).
- [69] D. Lindstrom and L. Haskin. “Direct Electrolysis as a Means of Altering Lunar Melt Compositions”. In: *Extraterrestrial Materials Processing and Construction* NSR 09-051-001.24 (Jan. 1980), pp. 197–199 (cit. on pp. 27, 31).
- [70] D. Linne, S. Gokoglu, U. Hegde, R. Balasubramaniam, and E. Santiago-Maldonado. “Component and System Sensitivity Considerations for Design of a Lunar ISRU Oxygen Production Plant”. In: *AIAA Paper* (2009). No. 2009-1391 (cit. on pp. 33, 38–40).
- [71] D. L. Linne. “Employing ISRU models to improve hardware design”. In: *Proc., 48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. 2010 (cit. on pp. 33, 38, 39).
- [72] P. G. Lucey, D. T. Blewett, and B. L. Jolliff. “Lunar iron and titanium abundance algorithms based on final processing of Clementine ultraviolet-visible images”. In: *Journal of Geophysical Research: Planets (1991–2012)* 105.E8 (2000), pp. 20297–20305 (cit. on p. 45).
- [73] L. Mason. “Beneficiation and Comminution Circuit for the Production of Lunar Liquid Oxygen”. In: *Space* (1992), pp. 1139–1149 (cit. on p. 29).
- [74] W. W. Mendell and L. B. J. S. Center. *Lunar bases and space activities of the 21st century*. Lunar and Planetary Institute Houston, Texas, 1985 (cit. on p. 113).
- [75] R. P. Mueller, I. I. Townsend, J. G. Mantovani, and P. T. Metzger. “Evolution of Regolith Feed Systems for Lunar ISRU O₂ Production Plants”. In: *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition* (Jan. 2010) (cit. on p. 110).
- [76] R. Murray. “Water extraction from lunar rock”. In: *General Electric missile and Space Division: Internal Report* PIR 7530-1-368 (1962) (cit. on p. 23).
- [77] D. A. O’Handley, E. E. Rice, and R. J. Gustafson. “ISRU Support for a Self-Sustaining Lunar Colony (SSLC)”. In: *39th AIAA Aerospace Sciences Meeting and Exhibit* (Jan. 2001) (cit. on pp. 66, 107).
- [78] G. Olhoeft, A. Frisillo, D. Strangway, and H. Sharpe. “Temperature dependence of electrical conductivity and lunar temperatures”. In: *The moon* 9.1-2 (1974), pp. 79–87 (cit. on p. 50).
- [79] J. D. Paramore. “Candidate anode materials for iron production by molten oxide electrolysis”. PhD thesis. Massachusetts Institute of Technology, 2010 (cit. on pp. 28, 30, 32, 88).

- [80] D. Plachta and P. Kittel. “An updated zero boil-off cryogenic propellant storage analysis applied to upper stages or depots in an LEO environment”. In: *AIAA Paper* 3589 (2002), pp. 7–10 (cit. on p. 85).
- [81] D. Plachta, R. Christie, J. Jurns, and P Kittel. “ZBO cryogenic propellant storage applied to a Mars sample return mission concept”. In: *Advances in Cryogenic Engineering: Transactions of the Cryogenic Engineering Conference-CEC*. Vol. 823. 1. AIP Publishing. 2006, pp. 205–212 (cit. on p. 84).
- [82] M. Planck. “The theory of heat radiation”. In: *Masius (Blackiston, Philadelphia, 1914)*, Sec 164 (1914), p. 175 (cit. on p. 52).
- [83] S. Prasad. “Studies on the Hall-Heroult aluminum electrowinning process”. In: *Journal of the Brazilian Chemical Society* 11.3 (2000), pp. 245–251 (cit. on p. 28).
- [84] C. S. Rai. *Electrical and Elastic Properties of Basalts and Ultramafic Rocks as a Function of Saturation, Pressure and Temperature*. 1977 (cit. on p. 51).
- [85] D. Rapp. *Human Missions to Mars*. Springer, 2007 (cit. on p. 21).
- [86] P. Richet and Y. Bottinga. “Thermochemical properties of silicate glasses and liquids: a review”. In: *Reviews of Geophysics* 24.1 (1986), pp. 1–25 (cit. on p. 55).
- [87] S. W. Richter. *Experimental determination of in situ utilization of lunar regolith for thermal energy storage*. Tech. rep. SAE Technical Paper, 1992 (cit. on p. 56).
- [88] F. G. S. “Analysis of recent measurements of the viscosity of glasses”. In: *J. Am. Ceram. Soc.* 8 (1925), pp. 339–355 (cit. on pp. 50, 51).
- [89] K. R. Sacksteder and G. B. Sanders. “In-situ resource utilization for lunar and mars exploration”. In: *The 5th AIAA Aerospace Sciences Meeting and Exhibit* (Jan. 2007) (cit. on p. 21).
- [90] A. Sammells and K. Semkow. “Electrolytic Cell for Lunar Ore Refining and Electric Energy Storage”. In: *LPI Contributions* 652 (1988), p. 210 (cit. on p. 31).
- [91] G. B. Sanders and W. E. Larson. “Integration of In-Situ Resource Utilization into lunar/Mars exploration through field analogs”. In: *Advances in Space Research* 47.1 (2011), pp. 20–29 (cit. on p. 25).
- [92] G. B. Sanders and W. E. Larson. “Progress Made in Lunar In Situ Resource Utilization under NASA’s Exploration Technology and Development Program”. In: *Journal of Aerospace Engineering* 26.1 (2012), pp. 5–17 (cit. on pp. 25, 26, 150).
- [93] G. B. Sanders, W. E. Larson, K. R. Sacksteder, and C Mclemore. “NASA In-Situ Resource Utilization (ISRU) Project-Development and Implementation”. In: *AIAA Paper No 7853* (2008) (cit. on p. 25).

- [94] S. L. Schiefelbein and D. R. Sadoway. “A high-accuracy, calibration-free technique for measuring the electrical conductivity of molten oxides”. In: *Metallurgical and Materials Transactions B* 28.6 (1997), pp. 1141–1149 (cit. on pp. 50, 51).
- [95] S. Schreiner, L. Sibille, J. Dominguez, J. Hoffman, G. Sanders, and A. Sirk. “Development of A Molten Regolith Electrolysis Reactor Model for Lunar In-Situ Resource Utilization”. In: *AIAA SciTech Conference - 8th Symposium on Space Resource Utilization* (Jan. 2015) (cit. on pp. 85, 104, 114, 120, 122, 123, 147).
- [96] S. Schreiner, J. Hoffman, G. Sanders, and K. Lee. “Integrated Modeling and Optimization of Lunar In-Situ Resource Utilization Systems”. In: *IEEE Aerospace Conference*. 978-1-4799-5380-6/15. IEEE. Mar. 2015, pp. 1–11 (cit. on pp. 119, 120, 122, 130).
- [97] E. Seedhouse. “Vision for Space Exploration”. In: *Lunar Outpost: The Challenges of Establishing a Human Settlement on the Moon* (2009), pp. 1–22 (cit. on p. 25).
- [98] B. Sherwood and G. R. Woodcock. *Cost and Benefits of Lunar Oxygen: Economics, Engineering, and Operations*. 1993 (cit. on pp. 22, 36, 38, 129).
- [99] L. Sibille and J. A. Dominguez. “Joule-heated Molten Regolith Electrolysis Reactor Concepts for Oxygen and Metals Production on the Moon and Mars”. In: *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition* (2012) (cit. on pp. 32, 35, 48, 52, 53, 63, 65, 86, 88, 147).
- [100] L. Sibille et al. “Performance testing of molten regolith electrolysis with transfer of molten material for the production of oxygen and metals on the moon”. In: *AIAA: 3rd Symposium on Space Resource Utilization* (2010) (cit. on pp. 25, 32, 53, 54, 92).
- [101] L. Sibille et al. “Recent advances in scale-up development of molten regolith electrolysis for oxygen production in support of a lunar base”. In: *Proc., 47th AIAA Aerospace Sciences Meeting, American Institute of Aeronautics and Astronautics (AIAA), Reston, VA*. 2009 (cit. on p. 32).
- [102] M. C. Simon. “A parametric analysis of lunar oxygen production”. In: *Lunar Bases and Space Activities of the 21st Century*. Vol. 1. 1985, p. 531 (cit. on pp. 22, 36, 38, 113).
- [103] A. H. Sirk, D. R. Sadoway, and L. Sibille. “Direct electrolysis of molten lunar regolith for the production of oxygen and metals on the Moon”. In: *ECS Transactions* 28.6 (2010), pp. 367–373 (cit. on pp. 25, 27, 30, 32, 53, 54, 63, 88, 92).
- [104] D. Snyder, E. Gier, and I. Carmichael. “Experimental determination of the thermal conductivity of molten CaMgSi₂O₆ and the transport of heat through magmas”. In: *Journal of Geophysical Research: Solid Earth (1978–2012)* 99.B8 (1994), pp. 15503–15516 (cit. on p. 49).

- [105] *Space Launch System*. URL: http://www.nasa.gov/pdf/664158main_sls_fs_master.pdf (cit. on p. 143).
- [106] E. Standish. “Design of a molten materials handling device for support of molten regolith electrolysis”. PhD thesis. The Ohio State University, 2010 (cit. on pp. 25, 32, 149).
- [107] E. Standish, D. M. Stefanescu, and P. A. Curreri. “Ceramics for Molten Materials Containment, Transfer and Handling on the Lunar Surface”. In: *47th AIAA Aerospace Sciences Meeting* (Jan. 2009) (cit. on p. 32).
- [108] J. Stebbins, I. Carmichael, and L. Moret. “Heat capacities and entropies of silicate liquids and glasses”. In: *Contributions to mineralogy and petrology* 86.2 (1984), pp. 131–148 (cit. on pp. 45–48).
- [109] C. J. Steffen, J. E. Freeh, D. L. Linne, E. W. Faykus, C. A. Gallo, and R. D. Green. “System Modeling of Lunar Oxygen Production: Mass and Power Requirements”. In: *Proceedings of Space Nuclear Conference*. 2007 (cit. on p. 37).
- [110] K. Stehling. “Moon refueling for interplanetary vehicles”. In: *Aviation Age* 29 (1958), pp. 22–23 (cit. on p. 23).
- [111] D. Stoesser, D. Rickman, and S Wilson. *Design and specifications for the highland regolith prototype simulants NU-LHT-1M and-2M*. Citeseer, 2010 (cit. on pp. 44, 45).
- [112] F. C. Strong. “Faraday’s laws in one equation”. In: *Journal of Chemical Education* 38 (Feb. 1961), p. 98. DOI: 10.1021/ed038p98 (cit. on p. 59).
- [113] G. Tamman and W. Heese. “Die Abhängigkeit der Viscosität von der Temperatur bei unterkühlten Flüssigkeiten”. In: *Z. Anorg. Allg. Chem.* 156 (1926), pp. 245–257 (cit. on pp. 50, 51).
- [114] L. A. Taylor and W. D. Carrier III. “Oxygen production on the moon: An overview and evaluation”. In: *Resources of near earth space* (1993), p. 69 (cit. on pp. 25, 26, 32, 126).
- [115] L. A. Taylor and W. D. Carrier III. “Production of oxygen on the moon: which processes are best and why”. In: *AIAA journal* 30.12 (1992), pp. 2858–2863 (cit. on p. 24).
- [116] B. S. Teeple. “Feasibility of producing lunar liquid oxygen”. PhD thesis. Massachusetts Institute of Technology, 1994 (cit. on pp. 34, 38, 129).
- [117] *The Omegalux complete electric heaters handbook and encyclopedia*. v. 29. Omega Engineering, 1995. URL: http://books.google.com/books?id=q_i0mgEACAAJ (cit. on p. 87).
- [118] K. Tsiolkovski. *The science fiction of Konstantin Tsiolkovsky*. University Press of the Pacific, 1979 (cit. on p. 23).
- [119] M. J. Turner. *Rocket and spacecraft propulsion: principles, practice and new developments*. Springer Science & Business Media, 2008 (cit. on p. 21).

- [120] UK Space Agency. *Economic analysis to support a study on the options for the UK involvement in space exploration*. Tech. rep. London Economics, 2009 (cit. on p. 23).
- [121] A. Vai, J. Yurko, D. Wang, and D. Sadoway. *Molten oxide electrolysis for lunar oxygen generation using in situ resources*. 2010 (cit. on pp. 29, 32, 92).
- [122] J Vila et al. “Temperature dependence of the electrical conductivity in EMIM-based ionic liquids: evidence of Vogel–Tamman–Fulcher behavior”. In: *Fluid Phase Equilibria* 242.2 (2006), pp. 141–146 (cit. on p. 50).
- [123] H Vogel. “Phys. Z”. In: *J. Am. Chem. Soc* 22 (1925), pp. 645–646 (cit. on pp. 50, 51).
- [124] D. Wang, A. J. Gmitter, and D. R. Sadoway. “Production of oxygen gas and liquid metal by electrochemical decomposition of molten iron oxide”. In: *Journal of the Electrochemical Society* 158.6 (2011), E51–E54 (cit. on pp. 32, 88, 147).
- [125] J. Wang et al. “YSZ films fabricated by a spin smoothing technique and its application in solid oxide fuel cell”. In: *Journal of power sources* 163.2 (2007), pp. 957–959 (cit. on p. 109).
- [126] J. R. Wertz, D. F. Everett, and J. J. Puschell. *Space mission engineering: the new SMAD*. Microcosm Press, 2011 (cit. on p. 113).
- [127] G. R. Woodcock et al. *Robotic lunar surface operations: Engineering analysis for the design, emplacement, checkout and performance of robotic lunar surface systems*. Tech. rep. D 615-11 901. Boeing Aerospace and Electronics, Jan. 1990 (cit. on pp. 22, 37).
- [128] K. Zacny et al. “Novel approaches to drilling and excavation on the moon”. In: *AIAA SPACE Conference & Exposition*. 2009, pp. 6431–6443 (cit. on p. 150).
- [129] K. Zacny et al. “Parametric Optimization and Prediction Software for Excavation and Prospecting Tasks”. In: (2013) (cit. on p. 150).