

Thermal Energy Grid Storage: Liquid Containment and Pumping

by
Caleb Amy

BS, Mechanical Engineering (2015)
University of Central Florida

MS, Mechanical Engineering (2017)
Georgia Institute of Technology

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy in Mechanical Engineering
at the
Massachusetts Institute of Technology
September 2020

© 2020 Massachusetts Institute of Technology
All rights reserved

Signature of Author: _____

Caleb Amy
Department of Mechanical Engineering
July 29, 2020

Certified by: _____

Asegun Henry
Department of Mechanical Engineering
Thesis Supervisor

Accepted by: _____

Nicolas Hadjiconstantinou
Department of Mechanical Engineering
Chair, Committee for Undergraduate Students

Thermal Energy Grid Storage: Liquid Containment and Pumping

by
Caleb Amy

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the Degree of
Doctor of Philosophy in Mechanical Engineering

Abstract:

As the cost of renewable energy falls below fossil fuels, the key barrier to widespread sustainable electricity has become availability on demand. Energy storage can enable dispatchable renewables, but only with drastic cost reductions compared to current batteries. In this thesis, I investigate an electricity storage concept that stores electricity as sensible heat in an extremely hot liquid ($>2000^{\circ}\text{C}$) and uses multi-junction photovoltaics (MPV) as a heat engine to convert it back to electricity on demand hours, or days, later. In addition to a technoeconomic analysis, this thesis focuses experimentally on heating, liquid containment, and pumping. The transfer of the storage liquid is key because it enables conversion to and from electricity and compact, efficient heat transfer. However, operating at these extreme temperatures introduces many practical challenges, so several novel solutions related to containment and pumping are investigated including high-performance heaters, sealing a large multi-part tank with affordable materials, and pumping above 2000°C . The key result is that although affordable silicon can be contained in affordable graphite and pumped at these temperatures, temperature variation in the system causes it the graphite infrastructure to rapidly dissolve and ultimately fail in a matter of hours. Alternative embodiments are proposed with recommendations on areas of future work. The key takeaway from the technoeconomic modeling is that integrating low-cost thermal storage with an inexpensive heat engine can enable an economical approach to electricity storage, even without high round trip efficiencies. Thus, despite the challenges, future work is warranted.

Thesis Committee:

Asegun Henry (Advisor & Chair)

Robert N. Noyce Career Development Professor

Gang Chen

Carl Richard Soderberg Professor of Power Engineering

Ahmed Ghoniem

Ronald C. Crane (1972) Professor

Evelyn Wang

Department Head; Gail E. Kendall Professor

Acknowledgments:

I have been very fortunate to have Professor Asegun Henry as an advisor and close mentor over the last five years. From discussions about personal and career goals to brainstorming solutions to the hundreds of crazy experimental problems we have encountered, his advice has been invaluable. Auspiciously, I was able to transfer from Georgia Tech to MIT with Dr. Henry, and through this process he entrusted me to lead the design and construction of our new lab, which was a rare and valuable experience.

I also must thank my committee for their constant feedback throughout the process. Professor Evelyn Wang, as a co-PI on our ARPA-E project which this thesis is related to, has been involved from even before I formed my committee. Dr. Wang was also critical to making our lab construction possible and timely, working with me and the construction crew every week for months. Professor Gang Chen has also been a critical mentor, especially helping me to understand mechanisms governing the behavior of electrons and photons at extreme temperatures—both from his lectures and personal mentoring. Professor Ahmed Ghoniem sparked some of my interest and ideas in his class on advanced energy conversion and helped me evaluate them through class projects and feedback on my thesis work.

I also owe much gratitude to mentors beyond my committee. Since many of the fundamental questions in this thesis are rooted in material science, I have been privileged to be mentored by Professor Ken Sandhage and Dr. Mehdi Pishahang. Dr. Sandhage of Purdue University has been a PI or co-PI on several projects I have been involved in and is an expert in high-temperature materials. He has provided critical insight into the governing and rate-limiting interactions of the materials in our high-temperature systems. Similarly, Dr. Pishahang joined our lab as a research scientist at MIT and has been a vital mentor for all my materials questions, and more. He was always open to working with me in the lab and always encouraged me to pursue the fundamental cause of results (good or bad). Professor Shannon Yee and Sonja Brankovic of Georgia Tech helped with the design of high-temperature experiments and made FTIR measurements of insulation reported here.

I also must thank my labmates—especially Colin Kelsall, Alina LaPotin, and Daniel Budenstein—for all their help with experiments, brainstorming, and checking the units. And lastly, I must thank sponsors who helped to fund my research and experiments over the years, most importantly the ARPA-E and EERE of the US DOE.

[blank page]

Table of Contents

Chapter 1: Introduction	15
1.1 Motivation.....	16
1.2 TEGS-MPV Concept.....	17
1.3 Prior Work.....	21
1.4 Alternative Embodiments	22
Chapter 2: System Efficiency and Technoeconomics	26
2.1 Introduction	27
2.2 Losses	30
2.2.1 Converting Electricity to Heat	32
2.2.2 Converting Heat back to Electricity.....	32
2.2.3 Heat Loss	34
2.2.4 Parasitic Losses.....	36
2.3 Round Trip Efficiency	37
2.4 Technoeconomics and Value	37
2.4.1 Storage Cost	38
2.4.2 MPV Cost	42
2.4.3 Heater Cost.....	44
2.4.4 Other Costs.....	45
2.4.5 Overall Cost	46
2.4.6 Value.....	47
2.4.7 Comparison	50
Chapter 3: Joule Heating Input	56
3.1 Introduction	57
3.2 Electrical Isolation in Solids.....	58
3.3 Electrical Isolation in gases	63
3.4 Heater Material.....	71
3.5 Thermal Insulation	72
3.6 Temperature Measurement	77
Chapter 4: Liquid Containment.....	80
4.1 Silicon Carbide.....	81
4.2 Proof of Concept	84
4.3 Sealing	88
4.4 Effect of Graphite Pores.....	91
4.5 The Effect of Silicon Purity	96
4.6 Thermal Cycling.....	104
Chapter 5: Pumping	115
5.1 Introduction	116
5.2 Mechanical.....	116
5.3 Thermal Considerations	123
5.4 Chemical Considerations.....	127
5.5 Results.....	132
Chapter 6: Conclusions and Future Directions	139
6.1 Thesis Contributions	139
6.2 State of TEGS-Si.....	141
6.3 Future Work.....	142

List of Figures:

Figure 1.1: Schematic overview of the TEGS-MPV concept. Excess electricity (left) is converted to heat and stored at a peak temperature above 2000°C until needed. Then, it is pumped through a multi-junction photovoltaic heat engine, where electricity is produced. The cold fluid is then stored until excess electricity is available again. 18

Figure 1.2: Efficiency limited by entropy as a function of temperature. A lower limit applies to turbines where heat is added over a range of lower temperatures. 19

Figure 1.3: Considerations for selecting a liquid. The materials shown in red have been recently of interest in the research community, while the others have been investigated previously, but are generally avoided for safety reasons. Cost is given in \$/mol since molar heat capacity, and thus energy stored, is nearly constant across the elements. Each of the fluids shown has strengths and weaknesses, since the ideal fluid would remain liquid over a very wide range, but also be inexpensive..... 20

Figure 1.4: Prior work pumping Sn above 1400°C. A gear pump made from AlN was used to circulate liquid Sn from a tank and through graphite piping. A hole was cut in the insulation to observe the flowing liquid metal in situ. 22

Figure 2.1: Full Scale overview of the TEGS-MPV concept. Part A shows the full system with a car and a person shown for scale. Part B shows the heater, where horizontal graphite pipes transport silicon exposed radiatively to vertical graphite heaters. Part C shows the tank wall and insulation while Part D shows the tank wall looking down—which is made from parts bolted together. Part E shows the centrifugal pump. Part F shows the multi-junction photovoltaic heat engine, with a zoomed in view in Part G. Here, a water-cooled rod covered in PV cells is lowered into a cavity where it is radiated by tungsten fins heated by the hot silicon. 29

Figure 2.2: Losses in thermal storage of electricity. The black lines indicate a molten salt energy storage system operating at 565°C, such as existing concentrating solar power plants. The solid lines track the energy stored throughout the cycle, where the dominant loss is in the heat engine. The dotted lines track the exergy or availability of the energy stored throughout the cycle, which show that for TEGS-MPV the conversion to heat generates much less entropy than for molten salt, due to the relatively high temperature of the TEGS-MPV system..... 31

Figure 2.3: MPV design (A) and efficiency model results (B). The MPV has two junctions, enables electrons above 1.2eV to be captured at that voltage, while those between 1.0-1.2eV can be harnessed at 1.0V. This utilizes more of the energy of the higher energy photons, while still collecting lower energy photons. The back reflector is key, and enables photons below 1.0eV to be reflected back to the emitter, to be remitted at the original spectrum. The

efficiencies shown in part B show a peak efficiency of 57% for an isothermal 2-junction emitter at 2400°C, however the actual efficiency is lower since heat is supplied between 1900-2400°C..... 33

Figure 2.4: Graphite insulation. On the left a flexible graphite felt is shown while on the right a rigid graphite board is shown. The rigid material is most expensive and slightly more conductive, so it is often used only where compressive strength is needed..... 35

Figure 2.5: Effect of scale on the cost of storing thermal energy. Most (~90%) of the benefits of large scale are realized at the GWh-e scale..... 41

Figure 2.6: Storage cost breakdown for 100% Si base case 42

Figure 2.7: MPV Heat Engine cost breakdown for the base case..... 44

Figure 2.8: Heater cost breakdown 45

Figure 2.9: Capital cost summary for TEGS-MPV System. On the left, CPE is shown where the dominant cost is Si, except in the “Fe medium” case, which shows major savings if Fe could be used. On the right, the CPP costs are shown, where the cell cost is relatively small, especially if the substrates used in processing can be reused many times as show on the far right. 47

Figure 2.10: Value of Arbitrage in 2007 PJM as a function of RTE. There is no value below 36% RTE because the cost of buying so much electricity exceeds the revenue from sales. Also, in this market/time period durations longer than 10 hours were rarely beneficial. 49

Figure 2.11: CPE and CPP (white shapes) of three competitive energy storage technologies. Arrows and black shapes indicate maximum CPP to break even. Arrow direction and length indicates NPV. The colored contour represents the RTE for zero NPV, assuming 10% IRR, 30 year system, and 10 hours of storage. 53

Figure 2.12: Levelized cost of electricity of energy storage versus duration. Batteries become expensive for durations longer than five hours, because their power and energy components are coupled. Other technologies reach a minimum LCOS where $CPE=CPP \cdot h$ where h =hours..... 54

Figure 3.1: Electrical breakdown of candidate high temperature insulators. All commercially available insulators other than hafnia (HfO₂) degrade below 2000°C. *SiC-B4C data is based on the patent mentioned [85]..... 60

Figure 3.2: Ellingham diagram depicting the thermal limit of oxide stability in presence of carbon [101]. The vertical red line indicates the peak system temperature and the angled redline represents the stability of carbon monoxide (CO). Above 2000°C, CO is more stable than all other oxides, including hafnia. So, carbon reduces oxides above 2000°C. 61

Figure 3.3: Hafnium oxide electrical insulators reduced by carbon. The white discs made from hafnia (HfO_2) were reduced to hafnium when heated to 2000°C in the presence of solid carbon beneath them, and gaseous evaporated carbon surrounding them. 62

Figure 3.4: Final heater is suspended by W threaded rod. The W rod serves as a mechanical mount, the electrical supply, and the thermal bridge from 2500°C to 200°C where the tungsten rod is bolted to copper wire. 63

Figure 3.5: Graphite heater arcing through argon gas to the surrounding conductive graphite insulation. Part A shows the serpentine coil of the graphite heater before arcing. Part B shows the initiation of the arc, which may actually be similar to Part C except that the camera exposure had not adjusted yet. In Part C, the arc can be seen travelling up from the anode, through the graphite insulation and then back down through another arc to the cathode. 64

Figure 3.6: Electrical breakdown of noble gases. Helium resists arc discharge to about twice the voltage of argon between 1900 K and 2500 K. 69

Figure 3.7: Heater resistance remains for three hours indicating no arcing at 2200°C and 70V. 70

Figure 3.8: Heater electric field reduced by smooth corners A: model show electric field peaks in line with B experimental observation. 70

Figure 3.9: A: Baseline heater electric field is highest at the cathode. B: By grounding the center of the heater, the peak field was reduced by half. C: The highlight rod was used for grounding. 71

Figure 3.10: A - Electrical resistivity of graphite and tungsten vs. temperature [60, 90, 102, 113]. B – Vapor pressure of graphite and tungsten vs. temperature [47]. 72

Figure 3.11: Contribution of the three modes of heat transfer to losses in insulation. Below 0°C in vacuum solid conduction becomes important. Near room temperature, the dominant loss is convection. Above 1000°C radiate becomes the main form of heat loss. Adapted from [115]. 74

Figure 3.12: Thermal conductivity of graphite insulation in various gases vs temperature. Dotted lines are estimated by adding effective conductivity in vacuum to gas stationary conductivity. 75

Figure 3.13: Heater geometry (left) and thermal model (right). The heater is suspended above the insulation by tungsten conductors. Electrical isolation is by solid separation until the bottom of the insulation where an AlN heat sink serves the dual purpose of passively cooling the tungsten conductors and electrically isolating them. The thermal model takes the heater temperature of 2150°C as an input and includes convective and radiative external boundary conditions. Internal radiation in the heater cavity is computed with ray tracing. 77

Figure 4.1: Si C Phase diagram. The horizontal red line shows the peak TEGS temperature. 82

Figure 4.2: Diffusion coefficients of silicon and carbon in silicon carbide [27]. The straight lines are Arrhenius plots because of the log y axis and the inverse x axis according to Equation 3.1. Carbon diffuses two orders of magnitude faster than silicon, so most SiC growth is due to carbon diffusion toward the liquid silicon..... 83

Figure 4.3: General experimental setup for small (~50 mm diameter) tanks heated by induction 85

Figure 4.4: Graphical induction heating experimental process. A 100mm tall tank and ¼-20 CFC bolt were heated to 2200°C for four hours and both formed a protective SiC layer as shown on the right..... 86

Figure 4.5: Two gallon tank successfully contained 553 Si at 2200°C for four hours 87

Figure 4.6: Evaporated silicon reached the graphite insulation, reacting to form porous SiC 87

Figure 4.7: Mechanical failure of thick graphite seal due to doubling of the seal thickness from reaction with infiltrated silicon..... 90

Figure 4.8: Self-sealing behavior of SiC reaction. After leaking 1 mm along the seal path, the leak was blocked by the formation of SiC. 91

Figure 4.9: Illustration of how pore size depends on initial particles and processing. The dark gray particles on the right indicate the effect of increased time, temperature, and/or pressure during the sintering process. 92

Figure 4.10: Conceptual illustration of pore closure by SiC reaction competing with leak rate. As the blue SiC layer grows, the leak rate decreases and eventually stops when the SiC fully closes the pore..... 93

Figure 4.11: Extruded graphite grades mechanically fail (A) and/or leak (B) due to large pore size. 95

Figure 4.12: Various pressed graphite grades successfully contain silicon above 2000°C..... 96

Figure 4.13: Fe-Si-C ternary phase diagram at 1450°C and approximate projection to 2400°C.. 98

Figure 4.14: FeSi dissolves tank unless pre-melted to form homogenous alloy; %-mass. In each of these tests the FeSi dissolved the tank wall due to locally high concentrations of Fe in direct contact with C..... 101

Figure 4.15: Graphite crucible precoated in SiC in preparation for FeSi containment testing. . 102

Figure 4.16: FeSi ingots pre-melted in alumina, heated inside graphite for induction heating 102

Figure 4.17: Effect of Fe concentration on protective SiC barrier. An SEM image is shown below each image of the C-FeSi interface. The crucible graphite grade and concentration of iron (%-mass) is shown below each image. The remaining fluid balance was 553 silicon. Only test A with 67% Fe successfully formed a protective SiC layer, while the other tests with more iron did not show any SiC formation and/or significant graphite loss. 103

- Figure 4.18: Thermal and fluid model of dissolution experiment. On the left is a black graphite crucible containing silicon and surrounded by insulation. It is heated from the bottom inductively, causing the thermal gradient shown on the right. This thermal gradient induces natural convection which enhances the mass transfer, so that dissolution can be measured in a matter of hours. 105
- Figure 4.19: Concentration of carbon in silicon for dissolution experiment. The snapshots in time show that large natural convection eddies carry away significant amounts of silicon. They also show that quasi-steady mass transfer has been established. The model predicts that the hottest region at the bottom will dissolve at 0.30 mm/hour. 107
- Figure 4.20: Temperature of spatial thermal cycling experiment over time. The tank was heated from the middle until the lid which was measured reach above the melting point of silicon (1414°C), as shown in the first 2.5 hours of the plot. Then the heating was moved to the bottom of the tank, inducing natural convection for the four hour duration of the experiment..... 108
- Figure 4.21: Results of four hour spatial thermal cycling experiment. Dissolution material loss of about 0.8mm in the hot 2100°C zone at the bottom was observed. This dissolved graphite was precipitated and formed porous SiC at the top of the liquid region. The darker particles in the middle and right figures are SiC..... 110
- Figure 4.22: Temperature of transient thermal cycling experiment over time. The temperature was ramped between 1600-2100°C 22 times over 8 hours to check for thermal stress induced SiC delamination. 113
- Figure 4.23: Results of transient thermal cycling experiment. Part A shows a tank which was cycled only once for reference. Part B shows the tank that was cycled 22 times, resulting in up to 3mm (the entire wall thickness) material loss. Part C shows a cut perpendicular to Part B, where the entire wall thickness was lost as indicated by the arrow and detail view in Part D. 114
- Figure 5.1: Impeller design (left) with backward curved blades to increase pressure at low speed. Velocity is shown on the right with streamlines indicating minimal flow separation. .. 117
- Figure 5.2: Mass flow meter. Liquid is pumping into the top tank (flow meter) which drains at a rate proportional to the height of the liquid in it. So, as the pump flowrate increases, the height and weight of the top tank increases until the outlet flowrate equals the inlet flowrate..... 118
- Figure 5.3: Flow path in combined pump / flow meter. The fluid is first accelerated to the outer diameter of the pump, then travels up the pump outlet path and into the tank (flowmeter) above the pump. The fluid level inside the flowmeter increases until the height is

sufficient to drive the fluid out of the small outlet at the same flowrate as the flowmeter inlet from the pump..... 120

Figure 5.4: The tank rests on the insulation, which rests on the scale. When the pump is turned on, the flowmeter above the pump fills, making the tank lighter. Then, the flowrate can be measured by measuring the weight of the tank. 121

Figure 5.5: Water drain curves (left) used to calculate flow rate (right) from measured mass. The drain curves are measured by filling the mass flow meter from Figure 5.3 and then stopping the inlet flow. So, the weight is recorded over time while the flow meter drains as shown on the left. This curve is then numerically differentiated with respect to time to get the mass flow rate vs time. Lastly, the mass flow rate is plotted vs the mass as shown on the right. 122

Figure 5.6: Mechanical design of high temperature pump. The pump is mounted asymmetrically to the motor with respect to the shaft to avoid thermal expansion induced shaft misalignment. The pump is connected to the motor with a low thermal conductivity ceramic shaft to thermally separate the hot and cold components. To accommodate the thermal expansion of the pump mounts relative to the internal bolts, springs are used in the cold region. The pump is located inside the tank so that a shaft seal is not needed since the small amount of fluid that leaks along the shaft remains in the tank. The motor is mounted to a water cooled plate to thermally protect it from the hot pump..... 123

Figure 5.7: Temperature distribution of pump system (left) and pump (right). The model uses a fixed temperature input to the base of the tank, based on the heater temperature of 2100°C. External convection and radiation boundary conditions are applied, and the motor and pump mounts are actively water cooled. The figure on the right shows that temperature varies by 100°C within the pump. 124

Figure 5.8: Temperature distribution of pump motor from side (left) and top (right) view. The motor is kept cold by the water cooling which not only cools the motor/pump mount plate, but also geometrically separates the hot pump mounts from the cold motor—effectively breaking the thermal circuit and protecting the motor from overheating. . 125

Figure 5.9: Thermal expansion of pump supports relative to internal CFC thread rod. The graphite and ceramic pump supports expanded 5 mm vertically, while the CFC thread rod expanded less than 1 mm. The springs at the top accommodated this mismatch, visibly compressing 4 mm..... 126

Figure 5.10: Full prototype pump system. The pump is lowered into the tank once the silicon is melted by linear actuators mounted to the motor plate. The tank is heated from below and is supported by the insulation which rests on a scale. The scale reports the weight

change of the tank, and by extension, the weight of the flow meter above the pump, enabling the calculation of flow rate. 127

Figure 5.11: Vapor (left) and liquid (right) silicon source for SiC layer formation. The path for vapor diffusion is so long and narrow that a SiC layer would take days to form compared to only minutes when submerged as shown on the right. 128

Figure 5.12: Dynamic SiC surfaces contact at the top of the shaft, causing binding which stopped the pump from rotating. The thick, rough SiC formation was caused by the rapid introduction of silicon to the top of the sleeve bearing when the pump speed was increased in a step change. 130

Figure 5.13: SiC shaft surfaces separated by silicon. The shaft did not bind in this case because the SiC formed slowly and smoothly since the pump was slowly lowered into the silicon. 131

Figure 5.14: Silicon drain curves (left) stacked (middle) and used to calculate flow rate (right). The drain curves are measured by filling the mass flow meter from Figure 5.3 and then stopping the inlet flow. So, the weight is recorded over time while the flow meter drains as shown on the left. This curve is then numerically differentiated with respect to time to get the mass flow rate vs time. Lastly, the mass flow rate is plotted vs the mass as shown on the right. 133

Figure 5.15: Pumping duration and temperature in three experiments. The top shaded region shows the different phases of the experiment: melting, pumping, and freezing. In the main plot, a horizontal temperature plateau is visible during phase change due to the high latent heat of silicon. The 0.1 hour test was stopped when the pump shaft became bound to the pump by SiC. The 1 hour test ended when the tank leaked. The 10 hour test was stopped because there was no flow rate signal. Images of the pump cross sections after each test are shown in Figure 5.18. 133

Figure 5.16: Part A shows the pump cross section, which suffered major dissolution after 10 hours of pumping. The regions closest to the pump entrance at the bottom center experienced the most thickness loss. Part B shows the pump outlet channel that is blocked by porous SiC and Part C shows the outer surfaces of the pump which are also covered in SiC. Part D shows the bottom of the pump, where macroscopic (~1mm) porosity has developed from the dissolution of carbon. Part E shows the thick SiC present is poorly adhered, and was deposited on top of silicon, rather than forming at the surface because when the SiC scale was chipped off, revealing the silicon beneath. 135

Figure 5.17: The pump dissolves over time. These images are from three separate experiments of varying length, but otherwise similar conditions. The buildup of precipitated SiC can be

seen after one hour, but it is much more obvious after 10 hours, where the entire pump outlet is clogged..... 137

Figure 6.1: TEGS-Sn concept where energy is stored in solid graphite and heat is transferred with liquid tin. During charging the left side of the piping is activated and tin transfers heat from the heater to the graphite. During discharging the right side of the piping is activated and tin transfers heat from the graphite to the MPV. 143

Figure 6.2: Regeneration of dissolved graphite by precipitation in the same location. Part A shows a graphite pipe with tin flowing through it before any dissolution and without heat transfer to drive dissolution. Part B shows an exaggerated geometry change of the pipe, after cold tin has been flowing in from the left and heated by the pipe, dissolving the wall mostly on the hot (right) side. Part C shows the pipe after hot tin has been flowing in from the right and cooled by the pipe, causing precipitation mostly on the right since the solubility gradient is larger at higher temperatures. 144

Figure 6.3: TEGS-Sn concept with combined Heater and MPV section, which can enable natural regeneration of carbon dissolution and precipitation cycling since the hot side of the Heater/MPV piping changes between charging and discharging..... 147

List of Tables:

Table 2.1: Losses throughout the storage cycle resulting in roundtrip efficiency (RTE)	37
Table 2.2: Material Costs and Properties	47
Table 2.3: Effect of Rate on Maximum CPP	52
Table 2.4: Value comparison of leading energy storage technologies.....	53
Table 3.1: Mesh convergence of heater thermal model	77

Chapter 1:

Introduction

Sustainable energy from renewable sources has become cost competitive with traditional fossil resources [1, 2], but its intermittent availability will limit deployment until affordable energy storage is available [3, 4]. Existing technologies are either too expensive (batteries) [5, 6] or rely on limited geography (pumped hydro and compressed air) [7, 8]. Thus, in this chapter a new energy storage concept is introduced that is geographically independent and is estimated to be very inexpensive. This concept appears feasible based on prior work in several areas of science and engineering which are considered in this chapter. Lastly, alternative embodiments related to the main concept are introduced, several of which are currently under development.

1.1 Motivation

Over the last decade the cost of electricity derived from intermittent renewables, i.e., solar photovoltaics (PV) and wind, has fallen drastically [1, 2] making renewables cheaper than fossil-derived electricity in many locations. However, this is only part of the story. While the levelized cost per unit energy (LCOE), which is the total cost divided by the lifetime electricity output, is low, it does not account for the intermittent nature of these energy sources. Thus, fossil-based technologies remain the main source of on-demand energy. As a result, Denholm [4] and others [4, 5, 9] have shown that renewable penetration onto the grid will be limited to < 10-15% without long duration (10+ hours) grid-level storage. Thus, how to store energy at the grid-scale cheaply, has emerged as one of the most important technological barriers to the decarbonization of the grid—a critical step toward mitigating anthropogenic climate change [10].

The cheapest mature grid storage technology is pumped hydroelectric storage (PHS), which has a high roundtrip efficiency (RTE) near 80%, as well as a low cost per unit energy stored (CPE) ~\$60/kWh-e and cost per unit power generated (CPP) ~\$1/W-e [11]. CPE includes costs that scale with the amount of storage (e.g. the size of the water reservoir), while CPP represents costs that scale with the input/output power (e.g. water pump/turbine). The issue with PHS, and also compressed air energy storage (CAES), is that they are geographically limited, and in the case of PHS the prime locations have already been exploited [7, 8]. Electrochemical batteries, on the other hand, have promising new chemistries [12] and architectures, but it is unclear if any will displace Li-ion batteries whose prices continue to drop from \$300-400/kWh-e down to a predicted asymptote ~\$150/kWh-e [7, 13]. However, even this lower asymptote for Li-ion is still not cheap enough to enable the eventual 100% penetration of renewables which may require a storage CPE as low as \$20/kWh-e [5, 6].

Therefore, alternative solutions to the storage problem are needed, as costs below \$50/kWh-e CPE and \$0.50/W-e [14-17] will be required. This low-cost requirement arises from a logical upper limit of the storage cost of \$0.06/kWh—the current average electricity price [18], and that 10+ hours [19] of storage are needed to reliably and cost-effectively supply the grid. So, given that a new cost regime is required to enable a renewable grid, the goal of this research was to investigate a new regime of energy storage. Namely, storing electricity as extremely high-

temperature heat and converting back to electricity using photovoltaics—thus completely rethinking the two traditionally high costs of storing electricity: storage and conversion.

1.2 TEGS-MPV Concept

To find a cost-effective storage solution, one approach is to search broadly for inexpensive materials and methods that require minimal processing. Similar to PHS, thermal storage fundamentally works with cheap (literally all) materials and can require little to no processing. However, unlike storing electricity in a gravitational potential, storing electricity as heat has fundamental losses due to the entropy associated with the heat, which also must be discharged when the electricity is outputted. So, the conversion of heat back to electricity tends to occur at relatively low efficiency (~35-40%) and high cost (~\$1/W-e) for conventional turbine-based heat engines. However, even though the low efficiency is undesirable, the critical metric is overall profitability, which can prove quite attractive if new embodiments achieve somewhat >40% efficiency, coupled with low cost of storage and conversion components [20], as explored in Chapter 2 on System Efficiency and Technoeconomics. While counterintuitive, the most profitable energy storage solutions may not be the most efficient. This is not only because low efficiency can enable low cost, but also because the energy that is stored is relatively inexpensive excess of cheap renewable sources.

The Thermal Energy Grid Storage (TEGS) concept is based on storing excess electricity as heat, and then converting back to electricity when it is needed. This excess electricity (for example, solar power in the middle of the day) is converted to heat using resistance heaters. This heat is stored in a cheap liquid at high temperature to preserve its usefulness. Then, when more electricity is needed, for example at night, the liquid is pumped through a heat engine that generates electricity. A conceptual view of the TEGS concept is shown in Figure 1.1.




Figure 1.1: Schematic overview of the TEGS-MPV concept. Excess electricity (left) is converted to heat and stored at a peak temperature above 2000°C until needed. Then, it is pumped through a multi-junction photovoltaic heat engine, where electricity is produced. The cold fluid is then stored until excess electricity is available again.

The cost of storage must be especially low because it provides a service that was not previously part of the electricity system. Not only must electricity be generated from primary sources, but when the sources are largely intermittent renewables, much of it must also be stored. In the same way, efficiency is critical since it is an additional loss mechanism, on top of pre-existing losses. These constraints on cost and efficiency drove this TEGS concept to employ a multi-junction photovoltaic (MPV) heat engine, which is estimated to have an efficiency similar to conventional turbomachinery (~50%) while being significantly cheaper (~\$0.30/W-e), and with a much faster response time, on the order of seconds. However, high efficiency and low cost are only enabled by an *extremely* hot heat source, which enables high power density (low specific cost) and high efficiency for the MPV cells. It is primarily this reason that temperatures as high as 2400°C are targeted for storage. As shown in Figure 1.2, most of the thermodynamic benefits of high temperatures are reached by about 1500°C, but in this case another 1000°C is warranted for its specific effect on the power density and superior photon emission spectrum for a PV heat engine. It is also important to note that Carnot efficiency can only be reached by thermal cycles that intake all heat at a peak temperature, while turbines are limited to a lower efficiency since

heat is added to a working fluid over a range of lower temperatures {Henry, 2014 #35}. This is a distinct advantage of a photovoltaic heat engine, which extracts heat at the peak system temperature. While the heat engine itself presents its own set of risks and challenges, this thesis focuses on heating, storage, and pumping.




Figure 1.2: Efficiency limited by entropy as a function of temperature. A lower limit applies to turbines where heat is added over a range of lower temperatures.

The options for storage materials become relatively narrow with these extreme conditions. The material needs to be inexpensive (<\$50/kWh-e [6, 14-17]), liquid around 2000°C, and able to be contained and pumped. At lower temperatures, salts are attractive because of their low cost and melting point. Metallurgical grade Silicon, on the other hand, is inexpensive and has high thermal stability. Tin (Sn) has a very wide liquid range and high thermal conductivity but is too expensive for bulk storage, as shown in Figure 1.3. Although there are a few materials that meet the first two requirements (Si, Fe, Al, SiO₂), containment turns out to be the most stringent requirement. In fact, *no* affordable material pair exists that is thermodynamically stable under these conditions. For example, although C-Sn, HfO₂-Al₂O₃, and other pairs are stable, Sn and HfO₂ are cost-prohibitive as bulk materials. Thus, the only apparent option for containment is graphite, since it is affordable and retains mechanical integrity at 2400°C. All the candidate fluids react with graphite, but prior experience in the electronics industry has shown that under

the right circumstances, Si may form a protective SiC barrier between itself and graphite [21, 22]. On the other hand, it is also known that Fe and Al continuously react without the formation of such a protective barrier [23, 24].




Figure 1.3: Considerations for selecting a liquid. The materials shown in red have been recently of interest in the research community, while the others have been investigated previously, but are generally avoided for safety reasons. Cost is given in \$/mol since molar heat capacity, and thus energy stored, is nearly constant across the elements. Each of the fluids shown has strengths and weaknesses, since the ideal fluid would remain liquid over a very wide range, but also be inexpensive.

For this reason, silicon contained in a graphite tank was selected as the primary material pair for storage. Although there is vast experience with molten silicon containment due to its use in the electronics industry, this application has a significantly different set of requirements. Namely, TEGS-MPV is 1000°C hotter, requires a longer lifetime, requires low cost of the storage medium and containment, and involves the use of a single batch of storage medium over the life of the tank. For example, at lower temperatures, quartz crucibles and coated graphite are often used for containment. The last difference is also important because any impurities in the silicon and the gas environment only enter once, unlike silicon processes where a tank is subject to many cycles of new impurities both from the silicon itself and the gas environment, such as silicon processing for the electronics industry [25]. TEGS-MPV also experiences different thermal cycling since it always remains above the melting point of the silicon.

1.3 Prior Work

The challenges in this project are only approachable because of prior discoveries and advances in science and engineering. Each aspect of the TEGS concept, from technoeconomics to liquid containment to pumping to the photovoltaic heat engine attempts to take a large step forward from decades of prior work in the respective areas.

Electrical heating above 2000°C has been available for decades for industrial furnaces[26]. However, the techniques employed here that enable reaching such high temperatures are fundamentally incompatible with energy storage. That is, these furnaces use minimal insulation and actively water cooled components to avoid overheating. This thermal design results in nearly all of the heat generated being lost as waste heat. While this is acceptable for industrial processes, where a physical product is produced and sold, for energy storage it is the energy itself that must be sold, so the losses must be much lower (e.g. less than 10% instead of more than 90%). For this reason, an important contribution of this thesis is in the development of an efficient 2000°C heater.

Containment of liquid silicon has been demonstrated by the electronics industry for decades [21, 22, 25], albeit under less extreme conditions. The work in this industry has shown that silicon carbide (SiC) *can* form a thin mass diffusion barrier between silicon and graphite, drastically slowing further reaction to about 1mm growth over 30 years [27, 28]. That is, a well adhered SiC layer limits the reaction rate to the speed at which silicon atoms diffuse through the SiC lattice. This discovery hints at the possibility that such a protective layer *may* also form and remain stable at the higher temperatures envisioned here. So, this prior work serves as a rationale to evaluate whether SiC remains protective with cheaper graphite, cheaper silicon alloys, much higher temperatures, much longer life, much larger tanks, and a less controlled gas environment. Since its possible the SiC layer may dissolve, detach, or otherwise become unprotective under this range of new conditions, a major focus of this thesis is to identify, model, and test each new and risky condition the SiC layer is subject to, including thermal cycling and pumping.

Similarly, high-temperature pumps have been engineered for industrial and energy applications, albeit until recently at moderate temperatures below 750°C [29-31]. In 2017, a

breakthrough was made by the author and researchers at Georgia Tech, raising this temperature limit to above 1400°C as depicted in [32]. Many learnings from this work have been applied to the current work.




Figure 1.4: Prior work pumping Sn above 1400°C. A gear pump made from AlN was used to circulate liquid Sn from a tank and through graphite piping. A hole was cut in the insulation to observe the flowing liquid metal in situ.

1.4 Alternative Embodiments

While the focus of this work is on the TEGS embodiment of a silicon bearing fluid contained in a graphite infrastructure, other concepts are being developed within our research group, across MIT, and throughout the world.

Several embodiments [33, 34] are under development involving the conversion of electricity to heat, which is then stored and later converted back on demand, such that they may be grouped into a class of technologies termed thermal energy grid storage (TEGS). What these various incarnations share is the storage of heat, which is exploited to be as inexpensive as possible and can be 1-2 orders of magnitude cheaper than electrochemical batteries [13]. They

vary mainly in the form of storage (solid, liquid, or latent), temperature, and method of energy conversion. The simplest embodiment that is arguably closest to commercialization, is to use molten salt, as is currently done in concentrated solar power (CSP) plants [35], except that one would need to replace the solar heat input with joule heating. With this approach, one can today achieve a CPE < \$100/kWh-e [36], but the problem would be the low RTE (~35-40%) and significant CPP (~\$1/W-e). A more clever approach introduced by Laughlin [33], is under development by multiple companies including Malta, Inc. and Brayton Energy. Their approach involves the use of a heat pump instead of joule heating, which can in theory almost double the RTE to ~ 72%, perhaps at similar CPP, and makes this embodiment of TEGS a very attractive option.

Even at ultra-high temperatures above 1000°C, several storage systems are being developed. For example, researchers at the Technical University of Madrid are experimentally developing a technology that stores energy in the melting of silicon at 1414°C [37]. The field is expanding such that a textbook is being released to address the various considerations in these ultra-high temperature systems [38]. At even higher temperatures, scientists and engineers at Antora Energy are working on storage in solid graphite around 2000°C, using photovoltaics to directly convert heat back to electricity, where radiation is the sole method of heat exchange.

Other interesting and potentially attractive embodiments also exist, but to determine the best option, the value of RTE must be assessed with respect to CPE and CPP. It is therefore important to have a framework for quantitatively evaluating the tradeoffs between RTE, CPE, and CPP, which ultimately dictates the economics and value to the grid. In the next chapter, a simple framework for assessing such tradeoffs is introduced, followed by a plausible embodiment of TEGS-MPV, which the analysis shows may be one of the few solutions to the storage problem that is inexpensive enough to eventually enable a fully renewable grid.

Among the concepts mentioned above are those that include storage in a material that is solid (or freezes), and there are important tradeoffs to consider here. On one hand, solids are in many ways lower risk since they hold their shape. On the other hand, liquids can be crucially effective at changing between charging, storage, and discharge modes of a storage system. With a liquid, these modes can be selected using pumps and/or valves, nearly infinitely variable surface

area changes between modes and negligible conduction length scales—and thus small thermal losses.

However, the use of a liquid storage medium requires pumping, which could be avoided if a solid storage medium were used. This work focuses specifically on liquids/metals because of the heat transfer issues that arise from using a solid storage medium. Notably, even though gases could be transported, they generally do not have sufficiently high energy density to offer a competitive embodiment, since the density of gases is generally 2-3 orders of magnitude lower than condensed phases. This assumes atmospheric pressure since high pressures combined with high temperatures necessitate an extremely thick-walled pressure vessel and become cost-prohibitive at the large grid scales of interest. This also assumes energy content based on sensible heat and not a chemical reaction.

Thus, to avoid pumping and containment, and enable storage in low-cost solids, one may consider a solid medium. Since a pump is not used, the mode of heat transfer is conduction through the solid, while energy can potentially be exchanged during charging and discharging by radiation. However, for storage durations more than 10 hours, the conduction length scale would be ~ 1 m, which is a key drawback since heat inherently must be conducted along long paths at the high flux (~ 100 kW/m²) required to make the system efficient and cost-effective. This induces very significant thermal gradients and transients in the solid mass, especially near the surfaces. Such a large temperature drop reduces efficiency and power density and may cause mechanical failure through thermal stress and shock.

A hybrid of these techniques exists and has been proposed in several embodiments. In general, this approach involves storing energy in a stationary material such as graphite, rocks, or even the Earth itself and exchanging that heat with a fluid. For example, researchers at MIT have proposed using air, oil, or even liquid metals as a heat transfer fluid to add and remove heat from cheap bulk storage materials [39-41].

Given the need for inexpensive energy storage to enable a sustainable future, and the challenges to achieving such a solution, this thesis answers the following questions central to the feasibility of the described TEGS-MPV concept. The first question on economics is addressed in Chapter 2 where a framework is developed to directly consider the effect of efficiency on

profitability. The estimated cost of TEGS-MPV is also developed in this chapter. The issues of electrical isolation which is required to convert excess electricity into heat is covered in Chapter 3 on Joule Heating Input. This challenge is novel in that this is the first application where extremely high temperature heat needs to be generated while greatly reducing losses compared to existing process furnaces. Question 3 and its parts address the key issue of whether silicon can actually be contained and pumped, and is addressed by Chapters 4 and 5 on Liquid Containment and Pumping, respectively.

1. What is the impact of roundtrip efficiency on the cost competitiveness of electricity storage technologies?
2. What are the issues with and engineering solutions to electrical isolation above 2000°C?
3. What materials compatibility issues arise when containing and pumping silicon alloys above 2000°C?
 - 3.1. What is the effect of graphite pores on the reaction between silicon and graphite?
 - 3.2. What is the effect of Fe, Al, and Ca impurities in metallurgical grade silicon on the reaction between silicon and graphite?
 - 3.3. Does the SiC layer stay adhered to the graphite during thermal cycling? Does the SiC layer delaminate, abrade, or cause binding during mechanical pumping?

Chapter 2:

System Efficiency and Technoeconomics

To determine the required cost of energy storage to enable affordable sustainable energy, a framework was developed to compare relative effects of efficiency, energy costs, and power costs on the overall economics of an energy storage system. This is in contrast to prior work which largely focused on cost, considering efficiency as a constant near 80%. The key result of this framework is that lower efficiency technologies can actually be very attractive for energy storage, particularly if their cost is also low.

Then, the efficiency and cost of the TEGS-MPV concept was estimated and compared to existing storage technologies. Roundtrip efficiency includes the losses throughout the system, including the converting electricity to heat, heat loss, several parasitic losses, and finally converting back to electricity. The cost includes all of the energy and power components, such as the heater, storage tanks and medium, pipes and pumps, insulation, and the multi-junction photovoltaic system. Also included in the cost considerations is construction and financing costs.

With the cost estimated, and a framework for comparison in place, the TEGS-MPV system is compared to existing technologies. This comparison shows that TEGS-MPV is similarly as profitable as pumped hydro, but without geographic constraints. Thus, these results suggest that TEGS-MPV or similar thermal storage concepts may be the most (if not the only) affordable options for energy storage.

2.1 Introduction

A very important distinction between TEGS and the vast majority of energy technologies is in the pursuit of efficiency. While the overall goal in all energy system is to maximize profitability, the economics of fossil fuel power generation and storing excess renewable energy work out very differently. For example, while recent progress in fossil generation focuses on tweaks to improve performance [42], the solar PV industry has primarily focused on silicon (and even cheaper [43]) cells, despite the efficiency benefits of multi-junction and direct bandgap materials which are only economical in niche cases [44]. Part of the reason for this dichotomy is fundamental, in that fossil energy involves many upstream costs before conversion to electricity, while renewable energy need only be converted. This difference means energy storage technologies can be optimized farther from peak efficiency. These technologies not only store energy from cheap renewables but by design store the cheapest portion of it. Just as low efficiency peaking simple cycle gas turbines buffer the energy grid of today, cheap moderate efficiency storage can buffer the grid of tomorrow.

To evaluate the economic prospects for TEGS-MPV and answer the question “if it works, will it matter?”, the basic concept was expanded to a full-scale model. As part of this model, key mechanical and thermal considerations were applied, such as the strength of the tank, thickness of insulation, and size of pumps. The result of this model is conceptually illustrated in Figure 2.1 and complemented by the basic concept illustration in Figure 1.1. Part A of Figure 2.1 depicts a full 1 GWh-e system, with a 100 MW-e MPV heat engine (back) and heater (front) between the hot and cold storage tanks with a person and car shown for scale. The tanks contain and insulate a low-cost thermal storage fluid, nominally 553 metallurgical grade (98.5% pure) silicon, which costs \$1.6/kg at high volume. These tanks are envisaged to be made from meter-sized plates bolted together with carbon fiber composite (CFC) threaded rod and sealed with a soft graphite foil seal.

To charge the system, 1900°C “cold” silicon is pumped, using an all graphite seal-less sump pump shown in part E, through a series of horizontal graphite pipes which connect to headers that are externally irradiated by vertical graphite rods acting as resistive heaters, drawing electricity from the grid. In this heater sub-system, part B, the temperature of the Si is nominally

raised to $\sim 2400^{\circ}\text{C}$ as it is pumped into the “hot” tank, where it is stored. The tanks are large, with diameters on the order of 10 m, which allows the surface area to volume ratio to be small enough that it is feasible for less than 1% of the energy stored to be lost each day, which is similar in size and heat loss to existing concentration solar power (CSP) plants using molten salt thermal energy storage (TES) [45]. In this extreme temperature case, the insulation is more expensive as detailed in the section on Heater Cost, but heat loss can still be minimal. Assuming such a storage resource were to be discharged once a day, this leads to an almost negligible penalty on the RTE.

When electricity is needed, the 2400°C Si is pumped out of the hot tank and through the MPV heat engine, shown in part F. The MPV power cycle consists of an array of PV cells that are vertically actuated into cavities between graphite pipes that are covered in tungsten (W) foil. The W foil acts as a lower vapor pressure barrier between the graphite pipes and the MPV cells, which are mounted to an actively cooled block that keeps their temperature near the ambient temperature. The W foil, therefore, serves as a photon emitter, almost identical to an incandescent lightbulb [46], that emits light to the MPV cells which convert a fraction to electricity. As the Si passes through the power cycle piping network it cools and returns to the “cold” tank. Notably, at these extreme temperatures, more than 25% of the light being converted is visible, and materials with bandgaps even higher than typical silicon PV are envisioned. Therefore, these cells are arguably just PV cells, as opposed to thermophotovoltaic (TPV) cells. This is why the term multi-junction PV (MPV) instead of TPV is used herein, to highlight the fact that the envisioned cells bear resemblance to, and use many of the advances that have been made for MPV, in the context of concentrated PV (CPV) [47].




Figure 2.1: Full Scale overview of the TEGS-MPV concept. Part A shows the full system with a car and a person shown for scale. Part B shows the heater, where horizontal graphite pipes transport silicon exposed radiatively to vertical graphite heaters. Part C shows the tank wall and insulation while Part D shows the tank wall looking down—which is made from parts bolted together. Part E shows the centrifugal pump. Part F shows the multi-junction photovoltaic heat engine, with a zoomed in view in Part G. Here, a water-cooled rod covered in PV cells is lowered into a cavity where it is radiated by tungsten fins heated by the hot silicon.

It should be appreciated that the temperature regime chosen here represents an upper limit for industrially manufactured materials, namely graphite and tungsten. Although both materials remain solid at even higher temperatures, at 2400°C a substantial vapor pressure develops in graphite [48] which can lead to deposition on the MPV cells—degrading their optical

performance, similarly tungsten is very soft (below 50 MPa) at these temperatures [49]. The most extreme temperatures possible are employed to achieve the highest possible RTE.

To date, the efficiency of PV that converts light from a terrestrial heat source has reached 29%, which was achieved using single-junction Si cells and a 2000°C emitter [50]. However, more recent work using an InGaAs cell achieved almost the same efficiency (28.8%) [51], but only required an emitter temperature of 1250°C [51]. This was enabled by the use of relatively high voltage direct bandgap cells [52]. Direct bandgap cells have relatively low voltage loss, but given that this loss is roughly constant, an important pathway to reaching higher efficiencies (> 40%) with a terrestrial light source is to use higher bandgaps. However, higher bandgaps require higher temperatures so that a substantial portion of the emitted spectrum is above the bandgap and can be converted. This is important because achieving a low CPP requires that the cells be operated at high power density, so that their cost, which scales with the total cell area can remain low. Furthermore, although photons below the bandgap can be returned to the emitter by mounting the cells on a back surface reflector (BSR), this recycling is imperfect, so the proportion of photons above the bandgap must be substantial to outweigh the parasitic below bandgap absorption. It is for these reasons that the most extreme temperatures possible are considered in this TEGS embodiment.

2.2 Losses

Although all losses are included in the efficiency calculations, it turns out that the only important exergetic losses are in the heater and heat engine. Exergy, or availability, can be defined as the usefulness of energy. Exergy and energy are equivalent for electrical energy (kinetic and potential energy too). This is demonstrated by Figure 2.2, where at the beginning of the storage cycle both exergy and energy have the same value. Once the electricity is converted to heat, it starts to degrade, and continually degrades either by heat loss or temperature decrease throughout the storage cycle. While the next sections will walk through the various loss mechanisms in detail, Figure 2.2 sums up the results, as it traces the usefulness of one unit of energy traveling through the storage system. As shown by the dotted red line representing TEGS, only about 10% of the usefulness is lost by conversion to heat, which is remarkably low—owing

to the extreme temperatures present. This 10% comes from Carnot efficiency where TEGS $\sim 3000\text{K}$ temperature is compared to the 300K room temperature. This can be contrasted with molten nitrate salt energy storage, at a peak temperature of 565°C [53], where 36% of usefulness is lost as soon as the heat is stored. For both systems, heat leakage and parasitic losses are relatively minor, due in large part to the large size of the storage tanks which minimize the surface area to volume ratio. The losses also differ in the heat engine. For TEGS, this is where most of the losses occur—in the conversion of radiated heat to electricity by the PV cells. This plot shows that the overall efficiency of the system with a low cost heat engine is truly made possible by operating at extreme temperatures, even if the heat engine can be outperformed on an exergetic basis by turbomachinery.

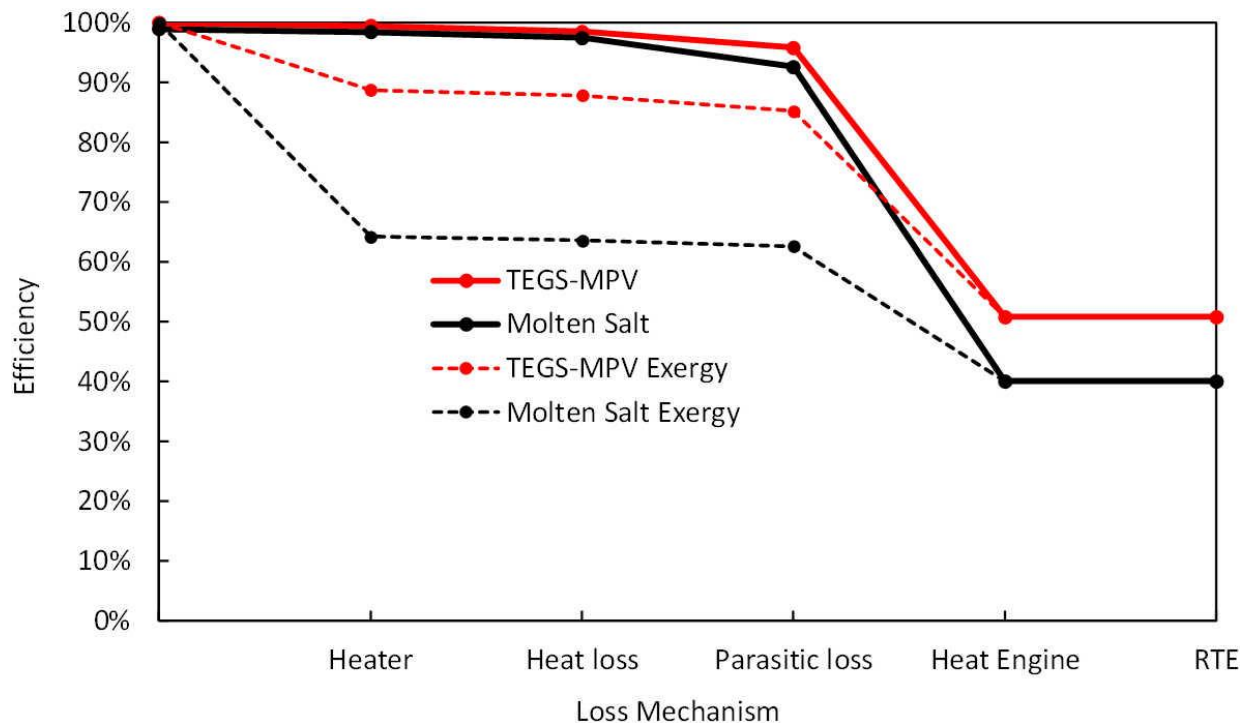


Figure 2.2: Losses in thermal storage of electricity. The black lines indicate a molten salt energy storage system operating at 565°C , such as existing concentrating solar power plants. The solid lines track the energy stored throughout the cycle, where the dominant loss is in the heat engine. The dotted lines track the exergy or availability of the energy stored throughout the cycle, which show that for TEGS-MPV the conversion to heat generates much less entropy than for molten salt, due to the relatively high temperature of the TEGS-MPV system.

2.2.1 Converting Electricity to Heat

While little effort is required to convert electricity into heat, controlling where the heat is generated is more of a challenge at extreme temperatures. Heaters are used very widely industrially and even residentially, but the constraints on large scale, extreme temperature energy applications are more stringent. Although commercial technologies exist to generate heat from electricity above 2000°C, they suffer major heat loss from the active cooling of the electrical insulators and joule heating in the wires and electronics supplying the heater. The incoming electricity must be conditioned to the required voltage and also modulated to adjust to the desired charge rate at any given time. These electronics are often expensive and lossy, but by using thyristor based power supplies with silicon controlled rectifiers (SCR) to modulate power by rapidly switching, efficiencies as high as 99.5% [54] can be achieved. Thyristors are simple, compact, solid state switches which are cheaper and more efficient than transistors, and can avoid the use of transformers. Although transistors can switch much more rapidly (10^9 compared to 10^2 cycles per second) and can be controlled linearly, thyristors have much lower losses [55]. Similarly, by using only refractory materials around the heater, active cooling can be avoided, thus minimizing parasitic heat loss. Thus, less than 1% of the electricity absorbed from the grid would be lost in the heating process. The actual realization of high-performance heating is explored in detail in Chapter 3 on Joule Heating Input.

2.2.2 Converting Heat back to Electricity

For the MPV power cycle, some of the initial important system-level considerations were addressed in work by Seyf and Henry [41], such as the need for the power cycle to be MW scale to overcome the losses associated with heat leakage to the environment by minimizing the ratio of surface area to volume. Their prior work also identified the BSR reflectivity or more specifically the net amount of below bandgap cell absorption, as the most critical parameter. Recent work by Amy et al. have shown that their initial predictions for the cell efficiency do not fully capture the realistic voltage losses that tend to occur in real cells [16], which drive the system towards operation at much higher temperatures [41]. A model of the cell efficiency, developed in collaboration between MIT and NREL is shown in Figure 2.3 and detailed in a recent publication [16]. Here, part A shows the portions of the incident spectrum that are converted by each

junction of the cell, and the portion that is reflected and recycled (gray). Part B shows the modeled efficiencies of 1- and 2-junction PV cells for 1900°C–2400°C emitter temperatures as a function of (bottom) junction bandgap. For 2-junction cells, the top-junction bandgap is selected to give the highest efficiency for given bottom-junction bandgap. The key result is that the average efficiency for the 1900-2400°C spectrum is expected to be 53%.

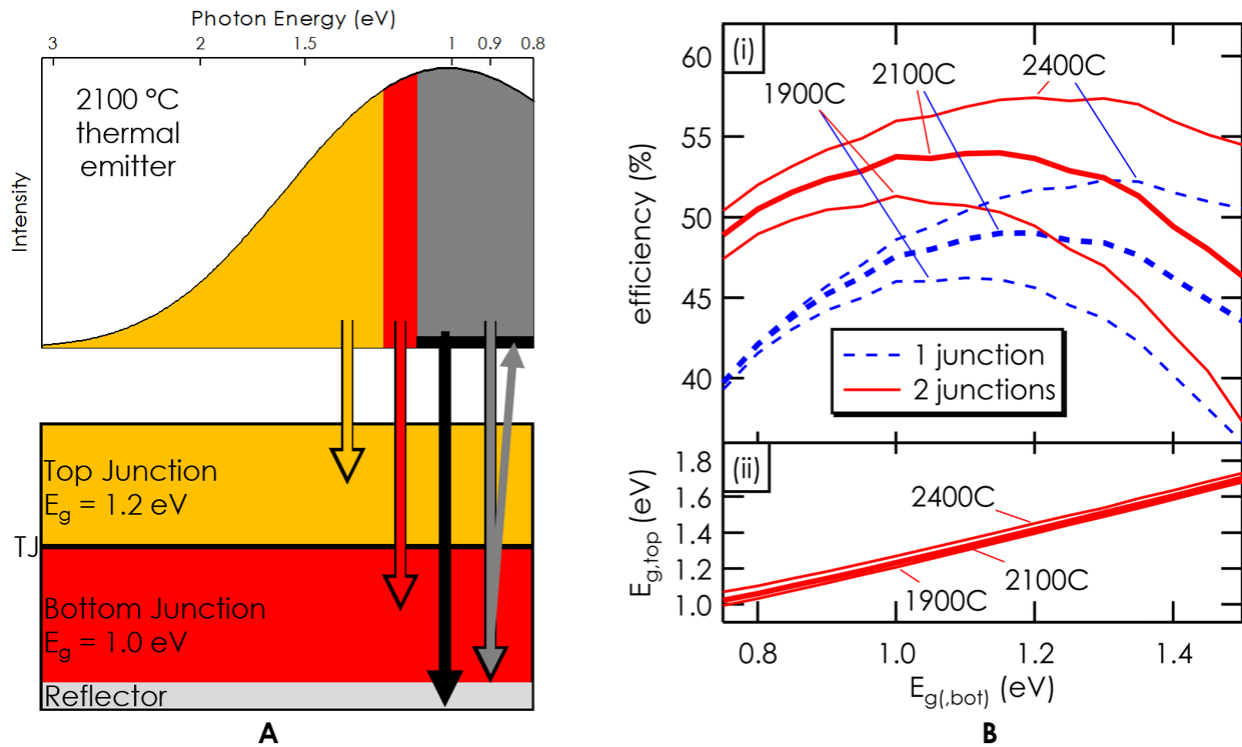


Figure 2.3: MPV design (A) and efficiency model results (B). The MPV has two junctions, enables electrons above 1.2eV to be captured at that voltage, while those between 1.0-1.2eV can be harnessed at 1.0V. This utilizes more of the energy of the higher energy photons, while still collecting lower energy photons. The back reflector is key, and enables photons below 1.0eV to be reflected back to the emitter, to be remitted at the original spectrum. The efficiencies shown in part B show a peak efficiency of 57% for an isothermal 2-junction emitter at 2400°C, however the actual efficiency is lower since heat is supplied between 1900-2400°C

While this efficiency would be revolutionary for terrestrial PV, it begs the question: why PV is chosen for the heat engine instead of turbomachinery, which could likely be more efficient, even at lower temperatures. There are three main reasons for this choice. First, turbines that take external heat input (instead of combusting the working fluid) and operate at high efficiencies (> 50%) do not exist. Although it may be possible to develop such a system, a large barrier to commercial deployment exists, since a large OEM must undertake an expensive (> \$100M)

development effort for a high-risk application. On the other hand, existing PV cell manufacturers exist to aid the deployment of these cells with greater interest and much less investment. Next, the cost of the proposed MPV system can be much lower than that of a turbine. Lastly, the ramping speed of PV can be orders of magnitude faster than turbomachinery—seconds instead of minutes. This is possible because, as shown in Figure 2.1G, the MPV modules can be actuated in and out of the light very quickly, which offers greater value to the grid, since TEGS-MPV can provide load following.

Notably, the electricity generated by the MPV heat engine is DC but must be converted to AC for use on the grid. While the cost of inverters is very significant and included in the MPV Cost section, their efficiency is very high at 98% [1] so the losses due to the inverter are 1% of system efficiency since efficiencies are multiplied (i.e., 51% → 50% RTE).

2.2.3 Heat Loss

The ideal rate of heat loss from thermal storage is a tradeoff between the cost of insulation, storage duration, cost of input electricity, and capital cost of other systems components. For this model, a storage duration of 24 hours was assumed, and the value of storage was estimated as described in the section on Value. With these inputs, the optimum (profit-maximizing) heat loss rate was determined to be approximately 1% of the energy stored per day. Although this seems like a challenging target, the sheer size of the tanks makes this feasible. This is because volume increases faster than surface area (L^3 vs L^2), so the ratio of volume to surface area increases linearly with the length scale. Thus, a 10 m tank loses heat at $1/10^{\text{th}}$ the rate of a 1 m tank, on a % loss per unit time basis. The effect of scale is also shown in Figure 2.5 in the section on Storage Cost.

For these large tanks, this is a heat flux of $\sim 400 \text{ W/m}^2$, which is notably more than two orders of magnitude lower than commercial ovens [26], which are actively cooled. This flux reduction is in addition to the effect of these tanks being $\sim 100\text{X}$ the diameter of commercial ultra-high temperature ovens, which are $\sim 0.2 \text{ m}$. Thus, taken together, the heat loss in terms of % per time is more than *four* orders of magnitude lower than commercial ovens. The graphite tank is insulated with multiple layers of insulation, with the insulation thickness calculated using Equation 2.1 and the effective thermal conductivities explored in detail in Chapter 3 on Joule

Heating Input. Here, r_i is the inner radius of the insulation layer, H is the tank height, k is thermal conductivity, $(T_i - T_o)$ is the temperature decrease across the layer, and Q is the heat loss.

$$t_{ins} = r_i \times \exp\left(\frac{2\pi Hk(T_i - T_o)}{Q}\right) \quad 2.1$$

Immediately outside the graphite tank is graphite felt insulation. This material is used until the temperature decreases to 1,350°C. At this point, a lower cost aluminum silicate ceramic fiber is used. Even lower in cost and thermal conductivity is fiberglass insulation so this material is used as an outer insulation layer below 540°C. The insulation cost is dominated by the graphite felt layer which is shown in Figure 2.4 and explored in the following cost section.



Figure 2.4: Graphite insulation. On the left a flexible graphite felt is shown while on the right a rigid graphite board is shown. The rigid material is most expensive and slightly more conductive, so it is often used only where compressive strength is needed.

The graphite tanks rest directly on rigid graphite insulation board, shown on the right of Figure 2.4. Below 1,700°C, this board rests on a calcium aluminate based castable cement (WAM ALII HD). This material is used for its compressive strength in a cinder block geometry, and the cavities are filled with aluminum silicate insulation to minimize cost, radiative heat loss, and natural convection. The castable cement then rests on a concrete foundation that can be cooled by forced air or water, as is the case in current molten salt CSP plants [56]. These tanks reside inside an inert atmosphere, achieved with a cool steel shell, which is explored in detail in the following cost analysis [57].

2.2.4 Parasitic Losses

There are several other small losses in the system that were quantified and included in efficiency calculations. Some losses not discussed below are included in the subsystems above (e.g. power conditioning electronics). These include the energy to pump silicon through the system, to pump water to cool the PV cells, to blow air to cool that water, to blow gas over the surface of the cells to keep them clean, and even the energy “lost” from the slow sublimation of the tank wall. The most significant of these losses is a result of selecting a dry-cooling scheme for the heat engine. While this is not currently the industrial standard, it is sustainable and enables TEGS-MPV to be located independent of water supply.

The energy required to pump silicon was estimated and included in efficiency calculations based on the head required to lift the silicon out of the tank and viscous losses based on 100 m of 250 mm diameter pipe. The pump was conservatively assumed to operate at 50% efficiency. To output 100MW-e of power, a flow rate of 3000 gpm is required, which is similar to a firehose. This is actually quite low for the rate of energy transfer, and therefore the pumping power turns out to be only 0.04% of the electricity produced.

The water used to cool the PV cells is under similar conditions (flow rate and viscosity) to the silicon itself and removes the same amount of energy. The main difference is that the water is limited to a narrower temperature range, so the viscous pressure drop is higher since smaller passages are needed for adequate heat transfer and higher flow rates are needed due to the smaller temperature rise. Here, the losses from cooling water pumping are 0.1% of the electricity produced and this loss was included in the efficiency calculations.

To reject heat from the water to air requires even more power since air has a low density and thermal conductivity. The losses here could be calculated directly from available data on dry cooling systems, assuming an 8°C temperature difference between the external environment and the water temperature. The losses from air cooling are 0.5% [58]. Notably, the dry cooling penalty in this system is less than for steam turbines because PV cells are less sensitive to temperature. That is, a Rankine cycle will not work well unless the water condenses at low pressure, whereas PV cells have a moderate efficiency dependence on temperature ($\sim 0.1\%/^{\circ}\text{C}$) [47].

2.3 Round Trip Efficiency

The roundtrip efficiency (RTE) of an energy storage system can be calculated by *multiplying* the efficiency of each process in series. This is a very important distinction since it yields a different (better) result than arithmetically subtracting each loss. For example, if state and federal income taxes are calculated arithmetically (and independently), but the total tax would be lower if it was multiplicative. That is, $(1 - x)(1 - y) > 1 - x - y$. The reason for this effect is that once losses occur in one process, the absolute loss in the next process is less significant. For example, because half of the energy is lost in the heat engine, a 1% loss elsewhere in the system only has a 0.5% impact on the RTE. So, for this system, the efficiency is calculated as shown in Table 2.1.

Table 2.1: Losses throughout the storage cycle resulting in roundtrip efficiency (RTE)

Loss Mechanism	Equation	Process Efficiency	RTE
Heater	η_{heater}	99.50%	99.50%
Storage heat loss	$(1 - \%Heat_loss_per_cycle)$	99.00%	98.51%
Si Pump	$(1 - P_{pump,Si} / P_{out})$	99.96%	98.47%
MPV air cooling	$(1 - P_{air_cool} / P_{out})$	99.50%	97.98%
MPV water cooling	$(1 - P_{pump,H_2O} / P_{out})$	99.75%	97.73%
Inverter	$\eta_{inverter}$	98.00%	95.78%
MPV	η_{MPV}	53.00%	50.76%

2.4 Technoeconomics and Value

While estimating the cost of a new energy system is difficult, in the case of energy storage estimating the value is even more challenging because it involves how it compares to other alternatives and future projections, which yields compounded uncertainty. This is in large part because the biggest value of energy storage has yet to manifest itself on the grid—balancing a grid dominated by intermittent renewables [4, 59]. This then results in a problem where neither can exist without the other. In this section, the cost of each component is addressed, followed by an estimate of the value (revenue streams), and finally a comparison between TEGS-MPV and established storage technologies.

For a given storage technology, the total capital expenditure (CAPEX) can be thought of as a sum of two main components, $CapEx = CPE + CPP / t$ where t is the time that the resource can be discharged at maximum power, CPE is cost per energy, and CPP is cost per power.

The major advantages of TEGS-MPV over other grid-level energy storage technologies are its expected low cost, geographically flexibility, and fast response time. Thus, it is important to demonstrate the basis of the cost estimates provided, as summarized in Figure 2.9. As a nominal design point, we considered a 100 MW-e output system with 10 hours of storage. The CPE includes the storage medium, tank, insulation, auxiliary components, and construction, using a similar procedure to Glatzmaier [20] and Wilk *et. al.* [60]. The CPP includes the heater, MPV cells, inverter, emitter, insulation, construction, and cooling system.

A model was developed which compares storage technologies based on three key parameters: cost per unit energy stored (CPE), cost per unit power produced (CPP), and round trip efficiency (RTE). To determine the minimum cost of TEGS, the model explores a wide range of variables in search of a global minimum cost, including tank size, insulation thickness, temperature range, and charge/discharge schedule and duration. This model also takes into account the lifetime of the components and the current and predicted future value of energy storage to determine the constraints on the cost and efficiency of the TEGS-MPV concept for renewables + storage to be the cheapest form of energy.

2.4.1 Storage Cost

In the nominal design case, the storage material is 553 grade (98.5% pure) Si at a market price of \$1.60/kg. This liquid silicon is stored in a tank made from isostatically molded graphite of density 1.8 g/cm^3 , at a cost of \$7/kg based on multiple quotes from large suppliers. This graphite grade has multiple trade names, such as KYM-20, AS-TJ, AR-06, and G330. The common features are a density greater than 1.750 g/cm^3 , with particle size below $10 \text{ }\mu\text{m}$. These large tanks can be built in sections as shown in Figure 2.1D, approximately one meter in size. The tank has two layers, which reduces the likelihood of leaks. The units are connected by flanges on all edges with high strength (120 MPa fracture under tensile loading) carbon fiber composite (CFC) threaded rod and nuts, as shown in Figure 2.1C. The units are sealed using a soft graphite seal as described in the section on Sealing. The CFC threaded rod which compresses the seal is

importantly much stronger than graphite (200 MPa vs 30 MPa tensile strength). This high strength ratio enables the total rod cross-section area to be smaller than the tank wall cross-section, which enables a compact system that takes full advantage of the mechanical strength of both materials. If, on the other hand, graphite threaded rod was used, the cross-section would need to be equal to that of the tank to get full mechanical use out of the expensive tank wall. This is impractical because the rods must be spaced away from each other, and so a very thick flange with many rows of rods would have been required.

Mechanically, the graphite tank wall thickness is designed to resist hoop stress with a minimum factor of safety (SF) of two, based on the tensile strength of isostatic graphite at room temperature (50 MPa), as shown in Equation 2.2. Here, t_g is the graphite thickness, ρ is the fluid density, g is the acceleration of gravity, H is the tank height, and r is the tank radius. This SF increases with temperature, as the graphite strength increases with temperature up to 2600°C [61]. The hoop stress in the wall decreases linearly with height, as the internal pressure arises from the gravitational force acting on the fluid. This fact can be leveraged to reduce the graphite required by up to a factor of two.

$$t_g = \frac{(\rho g H) r}{(\sigma_{fracture} / SF)} \quad 2.2$$

Immediately outside the graphite tank is flexible graphite felt insulation, at \$7,000/m³ based on multiple quotes. This material is used up to the radial location that the temperature decreases to 1,350°C. At this point, a lower cost aluminum silicate ceramic fiber blanket is used at \$400/m³. The cost of this material is low in part because it is widely used in ovens. Even lower in cost and thermal conductivity is fiberglass insulation, at \$85/m³, so this material is used as an outer insulation layer below 540°C. The insulation cost is dominated by the graphite felt layer, which bridges a 1,000°C temperature decrease in the hot tank, but only 550°C in the cold tank. For this reason, the cold tank cost is \$10/kWh-e cheaper than the hot tank. The insulative properties of the graphite insulation are very important and are explored as a function of the temperature and gas environment in Chapter 3 on Joule Heating Input.

The graphite tanks rest directly on rigid graphite insulation board and cost \$13,000/m³. Below 1,700°C, this board rests on a calcium aluminate based castable cement (WAM ALII HD),

which costs \$6,000/m³. The castable cement then rests on a concrete foundation that can be cooled by forced inert gas or water, similar to current molten salt CSP plants [56]. The concrete cost is \$200/m³ and cooling cost is estimated \$60/m² based on a designed heat flux of 400 W/m² and a cost of recirculated cooling [58] of \$80/kW. These tanks reside inside an inert atmosphere, achieved with a cold steel shell [57].

Creating a large inert atmosphere can be expensive, depending on the exact requirements. In the case of TEGS, the leaking of oxygen into the system causes oxidation of the graphite felt insulation and storage medium. Almost all the oxidation should happen in the insulation though because it has a very high surface area, and also reacts with oxygen before it can reach hotter components inside the insulation. This fact has been observed experimentally, where only the outer (colder, ~1000°C) graphite insulation forms gaseous oxides CO and CO₂. With the degradation mechanism known, the acceptable oxygen leak rate into the system can be estimated. Here, 10% of the insulation was allowed to degrade over a 30 year period, which is equivalent to 30 m³ of air leaking *into* the system per day, which is quite large. One way to minimize the rate of leaking into the system is to have a positive gage pressure so that mechanical pressure drives any leaks *out* of the system— but the oxygen concentration gradient will still drive some diffusion against the convective mass transfer. In the case of positive gage pressure, another leak rate limit emerges, which comes from the cost of replenishing the inert gas. This limit turns out to be very similar to the first, since at 30 m³ of air leaking *out* per day corresponds to an effective capital cost of inert gas ~ \$0.7/kW-year as an operational expense (OpEx) basis.

A similar commercially available process for controlled atmospheres is that used for extended fruit storage [62, 63]. From discussions with vendors, an effective typical leak geometry was determined based on actual blower door tests [64] they have conducted on fruit storage buildings. Here, the building is pressurized to 200 Pa gauge pressure and the leak rate is measured. To normalize for the size of the building, and for simplicity of measurement, the leak rate is measured in pressure change over time, where the limit for a tolerable leak rate is 10 mPa/s. Since a 1 GWh-e TEGS system would require a 100x40x20m building, this amounts to a volume of 60,000m³. So, if the building was pressurized to 200 Pa, the leak rate of 10mPa/s would amount to 10g/s or 300 m³/day. Since the maximum tolerable leak rate for TEGS was calculated

above to be 30 m³/day, 200 Pa is too high. However, it is not necessary for the building pressure to remain this high, this is pressurized test to quantify the leak rate. Once the rate is known, and since flowrate in pressure driven flow is linearly related to pressure, the maximum building gauge pressure can be calculated, resulting in a maximum gauge pressure of ± 20 Pa (or 0.003 PSI). This means that if the system pressure can be modulated between ± 20Pa, the leak rate with commercially deployed technology (essentially silicone gaskets) is tolerable.

The cost of this inert containment based on controlled atmosphere technology from three studies [65-67] over the last 50 years is \$750-\$900/m² or \$1.3-1.6/kWh-e for the TEGS system. This is about 3% of the total CapEx. Notably this cost scales like insulation cost, so the larger the system is, the less the specific inert containment cost. The full-scale size is selected to approach the asymptotic minimum cost while remaining at a scale that could be feasibly manufactured (1 GWh-e), as shown in Figure 2.5.




Figure 2.5: Effect of scale on the cost of storing thermal energy. Most (~90%) of the benefits of large scale are realized at the GWh-e scale.

Construction costs are estimated based on the cost of constructing molten salt CSP tanks [20], along with the cost to assemble additional components. For example, the cost to layout and bolt together the graphite tanks is estimated assuming that each section takes five minutes to position and one minute to install each nut. With an estimated worker salary of \$50,000 per year,

the total tank construction cost is \$3.14/kWh-e. Given all these considerations, the overall cost of the storage system, which makes up most of the energy costs (CPE) is shown in Figure 2.6.




Figure 2.6: Storage cost breakdown for 100% Si base case

2.4.2 MPV Cost

The hot components in the MPV power block are inexpensive relative to the rest of the infrastructure, and are very similar to those in the heater described next. The cost is dominated by the \$0.10/W-e MPV cell cost, \$0.08/W-e inverter cost [1], and the \$0.07/W-e dry cooling cost. This cell cost is based on a power density of 100 kW/m² and a cell cost [41] of \$10,000/m². In reality, this cost may be much lower if metalorganic vapor-phase epitaxy (MVPE) can be used instead of metalorganic chemical vapor deposition (MOCVD)[16], along with reusable substrates [68]. Under this scenario, the cost of the cell manufacturing may reduce by about a factor of five, to \$2,000/m² or \$0.02/W-e at the high-power density present in this system. If this advancement were achieved, the cell cost would be nearly negligible, and periodic cell replacements could easily be tolerated (e.g., as improvements are made).

Thus, the major variable in the power block cost is the MPV cell cost. It is expected that this cost will fall somewhere between the price of silicon PV cells [1] at \$50/m² and the current cost of manufacturing GaAs cells at \$10,000/m². Since the power density of this high-temperature system is ~ 100 kW/m², the resulting cell cost will likely be between \$0.0005/W-e and \$0.10/W-e. Nonetheless, a conservative upper limit on cost and lower limit on power density is used in the primary cost model.

Another important cost is that of the inverters to convert the DC power to AC. These are priced at \$0.08/W-e based on national averages [1] for central inverters in utility-scale PV. Cooling of the MPV cells is also estimated at \$0.08/W based on industrial prices [58] for

recirculated dry cooling 8°C above ambient, and the required cooling power per unit of electrical power generated is calculated as $(1 - \text{RTE})/\text{RTE}$. The numerator of this equation accounts for the waste heat generated, while the denominator normalizes it to electricity output, rather than heat input. Since the RTE is about 50%, this means approximately 1W of cooling is needed per watt of electricity output.

Tungsten (W) foil, which is used as a vapor pressure barrier to suppress the evaporation of graphite infrastructure was quoted at \$700/m², or \$0.035/W-e. The graphite piping and insulation is similar to the previously described systems and has a small effect on cost (\$0.014/W-e), but is included. The cost of constructing the MPV power block and the heater is estimated by including CPP construction costs based on previous analysis and adding the estimated labor cost to assemble additional components. For example, the time to install each pipe in the heater and MPV systems is estimated to be 10 minutes (essentially tightening 1 nut), and 30 minutes is estimated to install each unit of tungsten foil. Based on the labor rates discussed in the section on Storage Cost, the construction cost of components that scale with power is estimated at \$0.03/W-e. A breakdown of the MPV costs is given in Figure 2.7. An overview and detail view of the MPV layout is shown in Figure 2.1F and G.

The cost of pumps and piping is also included in this section since these components scale with the power capacity of the TEGS-MPV system. Because of the energy density of silicon over the nominal 500°C storage temperature range (i.e., 1900-2400°C), to discharge at 100MW-e, a flow rate of only 0.2 m³/s (~3000 gpm) is needed. For perspective, this is similar in size to the water pumps found on fire engines. The pressure required is mostly needed for gravitational head and will vary between 0.1-0.4 MPa (1-4 atm), which can easily be met with a centrifugal pump consuming ~ 40 kW of power. This flow rate and pressure can be met with a 330 mm (~ 1 ft) diameter centrifugal pump (SAE size 8x10-13), and the power requirement is negligible compared to the 100 MW-e power output, which a key advantage of using liquid Si/metal. The pump material is graphite, and the cost will be dominated by the 10 m shaft required to locate the pump in the bottom of the tank with the motor above the tank, similar to CSP plants. The total mass of a pump is estimated at 2,000 kg and a pump is needed for each of the two tanks. The primary piping network between the tanks, heater, and MPV has a nominal diameter of 250

mm to minimize dynamic head loss. The cost of pumps is based on isostatically molded graphite (e.g. KYM-20) of density 1.8 g/cm^3 , at $\$7/\text{kg}$.




Figure 2.7: MPV Heat Engine cost breakdown for the base case

2.4.3 Heater Cost

The cost of the heater includes electronics, graphite heating elements, graphite pipes and headers to distribute the silicon to be heated, insulation, and inert containment. Thyristor based power supplies for the heaters were quoted at $\$5/\text{kW}$. These supplies use silicon controlled rectifiers (SCR) to modulate power by rapidly switching, with efficiencies as high as 99.5%[54]. In a large heater such as this one, heating elements can be arranged in a series-parallel configuration to match the overall heater resistance to supplied voltage, thus reducing or eliminating the need for voltage transformers. In this specific design, the heater consists of a 37×15 array of graphite pipes of 10 mm inner diameter and 20 mm outer diameter. Between columns of these pipes are graphite rods that are used as electric resistance heating elements as shown in Figure 2.1F. The graphite components are made from isostatically molded graphite (e.g. KYM-20) of density 1.8g/cm^3 , at $\$7/\text{kg}$ based on multiple quotes from large suppliers. The heater is designed so that the peak heater temperature is 2500°C . The pricing for pipes and rods is from graphitestore.com, at $\$100/\text{m}$ for pipes, and $\$18/\text{m}$ for rods. A quote was obtained for these custom headers, at a cost of $\$0.25/\text{kW}$. The heater is insulated using the same approach as the tanks, by restraining the heat flux loss to be the same 400 W/m^2 as the tanks. The cost of the insulation and inert containment follows the same methods explained in the section on Storage Cost, but is much lower because the heater is smaller. Overall, the total cost of the heater is about $\$0.02/\text{W-e}$, which is less than 10% of the total cost of the power components, since the heat MPV system dominates the cost. A breakdown of the heater cost is given in Figure 2.8.




Figure 2.8: Heater cost breakdown

2.4.4 Other Costs

Several minor costs have been grouped with the sections above where they best fit. For example, pumping cost is discussed and included with the MPV cost, since pumps are a power component and the MPV is the main power system. Construction costs are broken up between the storage and MPV cost sections above. These components of the cost can be identified by the legends of Figure 2.6 and Figure 2.7.

Two important remaining costs are associated with the operations and maintenance (O&M) and also the financing. One interesting aspect of this system is that it is designed to remain (very) hot during its entire lifetime, which is primarily a drawback because it limits access to the hot portion of the system and any maintenance that can be done. Nonetheless, O&M costs are estimated based on industry standards for the fixed O&M of combined-cycle natural gas power plants, at \$11/kW-year [69]. This cost does not appear on a CapEx basis but is instead included as a detractor from revenue in the following value and comparison modeling.

Another critically important cost is financing. In general, the cost of financing often exceeds the amount financed. For example, for a 30-year loan at 5% annual percentage rate (APR), the financing charge is approximately equal to the principal borrowed. For first-of-a-kind systems, interest rates are even higher because of the risk with trying an unproven technology, so several rates are explored in Table 2.3 to show their effect on the break-even system cost. This table is in the section on Comparison because it includes five other existing technologies. For later figures, a rate of 10% is used to compare all technologies, although riskier technologies may be subject to higher rates (or lower if government-guaranteed loans can be secured).

2.4.5 Overall Cost

The total capital cost of the TEGS-MPV system is shown in Figure 2.9, broken into energy and power components. The first (left) bar for each component is the baseline which has been described above, resulting in a CPE of \$36/kWh-e and CPP of \$0.34/kW-e which would be economical for long durations [14-17]. The second bar takes into account less conservative, but plausible cost reductions to present a lower possible cost. Table 2.2 shows costs and sources for key materials.

For energy-related costs (CPE), an alternative embodiment of interest is to use iron (Fe) to partially replace Si as the storage medium. In this scenario, the cost of the medium becomes much lower, especially if scrap steel is used, and the other tank costs, especially insulation and construction, dominate. This effect is particularly dramatic because the volumetric energy density of Fe is 3X that of silicon, so the rest of the system becomes more compact. Generally the volumetric energy density of elements decreases with atomic number since specific heat per mole generally approaches a constant value of $3k_b$ [70] and since the atoms get larger, atomic density generally decreases. However, Iron's molar specific heat reaches almost $5k_b$ due to contributions from electrons [71]. Also, despite the larger size of iron atoms, they pack 70% more densely than silicon atoms. In this less conservative lower cost case, it is also assumed that a lower grade extruded (\$2/kg) graphite can be used for the tank and a higher heat loss of 2% per day, instead of 1%. The effect of these changes is shown in the "Fe" case of Figure 2.9.

Similarly, a lower cost scenario for the power components is considered. Here, the only change is the manufacturing method of the PV cells. If MVPE can be used instead of MOCVD, along with high substrate reuse, the cost of the cell manufacturing may reduce by about a factor of five, to \$2,000/m² or \$0.02/W-e at the high power density present in this system. Notably, the overall cost of the power cycle remains relatively similar to the base case because the cell cost did not dominate to begin with. Other potential cost reductions (~\$0.12/W-e; not shown) would involve using wet cooling and avoiding purchasing inverters by collocating with a PV farm or along a DC line.




Figure 2.9: Capital cost summary for TEGS-MPV System. On the left, CPE is shown where the dominant cost is Si, except in the “Fe medium” case, which shows major savings if Fe could be used. On the right, the CPP costs are shown, where the cell cost is relatively small, especially if the substrates used in processing can be reused many times as show on the far right.

Table 2.2: Material Costs and Properties

Material	Density (kg/m ³)	Cost (\$/kg)	Thermal Cond. (W/m-K)	Temperature Limit (°C)	Source
553 Silicon	2400	1.6	25	3250	Quotes[72]
Isostatic Graphite	1850	7.0	30	3600	Quotes
Rigid Graphite Insulation	24	540	0.5 (0.15-1.4)	2800	Quotes[73]
Graphite Felt	14	500	0.5 (0.15-1.4)	2800	Quotes[73]
Aluminum Silicate	100	4.0	0.2	1350	Quotes[74]
Fiberglass Blanket	12	7.1	0.05	540	[75]
WAM ALII	2700	2.2	1.5	1700	Quotes[76]
Scrap steel	7000	0.1	30	2862	[77]
Tungsten Foil (0.1mm thick)	19000	350	100	3400	Quotes

2.4.6 Value

There are many ways to estimate the highly uncertain value of long-duration storage. Broadly, one can demonstrate past value or estimate future value. Similarly, the revenue streams of an energy storage asset on the grid can be analyzed, or a renewable source firmed by energy storage can be viewed as a package. In the future grid, it is unlikely that storage will be used to

balance individual grid elements, because of the economic benefits of diversity [78]. Because of the complexity and uncertainty in future value modeling, and the excellent ongoing research in these areas by experts [3, 5, 6, 9], this work conservatively bases the value of energy storage on demonstrated past value. This is conservative because the need and value of storage are expected to increase as the share of cheap intermittent renewable increases.

The two primary existing unsaturated sources of revenue for energy storage are capacity payments and arbitrage. Capacity payments essentially exist to enable peaking gas turbines to cover their CapEx in a market (effectively all US markets) with electricity price caps [79, 80]. It is paid annually with respect to the power output promised/supplied (i.e., \$/kW). Arbitrage often has negative connotations in the financial industry, but in the case of energy systems, it is synonymous with balancing. That is, in a free market, prices will be low when supply exceeds demand, and high when demand exceeds supply. Energy storage can efficiently buffer this imbalance, earning revenue by buying low and selling high. Arbitrage value is defined as the summation of the annual revenue that could be earned by a device, minus the cost to purchase energy at off-peak times [3]. Arbitrage is earned on the free market, but can be analyzed in the same units as capacity payments (\$/kW) earned annually, and it is here that RTE plays a critical role.

Notably, this value modeling assumes no intervention from governments or otherwise to incentivize sustainable energy despite the positive externalities. Thus, the results here do not rely on or assume a value for avoided emissions, but instead are based on free market value alone. Sioshansi *et al.* [3] have quantified how much value a storage resource would have received from arbitrage, as a function of the RTE and storage duration (t , in hours, in the following equations), by using the Pennsylvania New Jersey Maryland (PJM) grid as an example. Their work showed that there would be a diminishing increase in value for long-duration resources on the 2007 PJM grid and they quantified how the value of storage changes with RTE, which is the basis for Figure 2.10. This plot shows that a storage technology with $RTE \leq 36\%$ would not have generated any value from arbitrage on the 2007 PJM grid. Fundamentally, this is because the input energy must be purchased and therefore the ratio between on-peak and off-peak pricing sets a lower efficiency limit, η_{\min} , to earn arbitrage profit, as shown in Eq. 2.3. That is, if a technology must

buy three times as much energy as it sells, it must sell that energy for at least three times the purchase price to derive positive value from arbitrage. Nonetheless, capacity payments enable a storage system to earn revenue for even lower efficiency, if it supplies the promised capacity when scheduled by the grid system operator. Because these devices cannot charge or discharge instantaneously, the closer their efficiency is to η_{\min} , the less frequently they can profitably engage in arbitrage.

$$\eta_{\min} \approx \frac{P_{\text{offpeak}}}{P_{\text{peak}}}$$

2.3

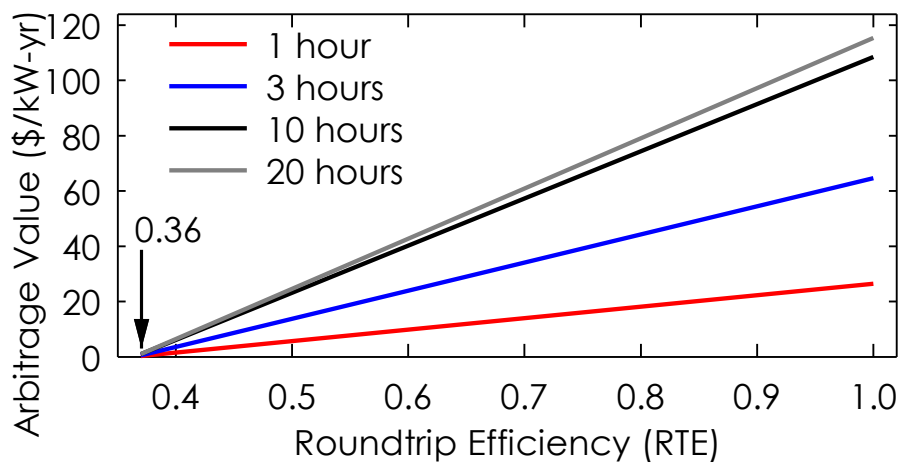


Figure 2.10: Value of Arbitrage in 2007 PJM as a function of RTE. There is no value below 36% RTE because the cost of buying so much electricity exceeds the revenue from sales. Also, in this market/time period durations longer than 10 hours were rarely beneficial.

There are other forms of revenue for energy storage that small resources such as lithium-ion are currently exploiting. These include payments for short duration services, such as primary frequency response, regulation, and contingency [81]. These revenue streams are not included here because they are expected to be saturated quickly and are best served by short-duration sources—namely those with a high CPE but low CPP. Notably, if long-duration energy storage technologies like TEGS are successful, arbitrage will also be saturated, which will set a price ceiling at the cost of the best available technology. This is a natural effect, and needs to be considered when projecting the future revenue streams of storage—the arbitrage profit will always be limited by the most competitive participant.

Capacity payments (CP) exist because of an electricity market failure that arises from electricity price limits that are intended to protect consumers [79, 80]. Since prices cannot exceed a set limit, but electricity demand only reaches peak levels for a few hours per year, peaking natural gas plants which only operate during peak periods are not profitable. Hence, with no market correction, these plants would not be built, and the grid would experience periodic blackouts. To avoid this, and incentivize new peaking generation capacity, CPs are set. Thus, these payments vary widely and are high in markets short on capacity and low in markets with excess capacity [80]. CPs can be thought of like an insurance policy paid by the grid operator to producers, in exchange for a guarantee of a maximum electricity price, much like how a driver purchases car insurance to avoid unexpected large expenses. Ideally, at the end of the day, everything works out the same, with the buyer paying a price for avoided risk.

Thus, instead of taking an average of actual CPs, which range \$0-500/kW-yr [4], a more fundamental method is used to estimate the long term average CP in a market with price caps. That is, the reason a CP is offered in the first place, is that in regulated markets, grid resources that only operate during peak times, namely peaking gas turbines, do not earn enough revenue from energy sales to be profitable. Then, logically, a CP should be the subsidy needed to allow a balancing resource to have zero net present value (NPV, a metric which includes profit in terms of an internal rate of return (IRR)). NPV is the key financial metric used to decide whether to make an investment. If it is positive, the investment is profitable in that it exceeds one's required rate of return. It is like the breakeven point, except it includes a return on investment in addition to return of investment.

CPs are actually determined based on this required additional revenue, and surveys are routinely conducted by grid operators to determine the net cost of new entry (Net CONE) [82] parameter, so it is widely reported. It is calculated as the total cost of a resource minus the revenue it earns from energy sales and other ancillary revenues. Thus, it is this net loss that needs to be compensated by a CP. This value averaged \$95/kW-yr in the PJM market in 2018 [83].

2.4.7 Comparison

Prior analysis and comparison of energy storage systems has been largely based on the cost of storage (CPE) alone, or using an all-encompassing metric like levelized cost of storage

(LCOS). Neither of these approaches was ideal. That is, while technologies in a given class can be compared on an energy cost basis (e.g. batteries), this approach neglects the important effect and interplay of roundtrip efficiency (RTE) and power costs (CPP). On the other hand, LCOS includes the effect of CPE, CPP, and RTE but only outputs a single metric which does not elucidate which are the most important factors. LCOS is defined as the total cost of a storage system (CPP and CPE) divided by the lifetime energy output, as shown in Equation 2.4. This is analogous to levelized cost of electricity (LCOE) for electricity which is directly consumed. Here, P is the price of electricity that is stored, OM is the annual operations and maintenance cost (O&M), and H is the hours of storage duration. Future revenue and expenses are discounted with the factor α as shown in Equation 2.5, which assumes revenue is accrued uniformly over time. Here, L is the life of the system in years and r is the internal rate of return (IRR). Here, electricity is purchased over time and O&M is paid overtime, so those values are discounted. On the other hand, capital costs CPE and CPP are paid upfront so that are not discounted. Discounting enables all these expenses to be compared in terms of today's dollars. For example, if the IRR is 10% and the term is one year, a revenue of \$1 one year from now is worth \$0.90 compared to having \$1 today.

$$LCOS = \frac{(1/RTE - 1)P\alpha + OM\alpha + CPE/RTE + CPP/H}{N\alpha} \quad 2.4$$

$$\alpha = \frac{rLe^{rL}}{e^{rL} - 1} \quad 2.5$$

Once the energy costs (CPE), power costs (CPP), roundtrip efficiency (RTE), and lifetime of an energy storage system are known, it can be compared to others. This comparison not only can show which technology is best suited for a given set of conditions (e.g. duration), but also the show gaps in existing technologies. In this section, prominent energy storage technologies are compared to TEGS-MPV in two ways: by looking at the conditions and level of profitability, and by comparing the technologies on a levelized cost of storage (LCOS) basis.

Existing storage technologies have high efficiencies (>80%) but are too expensive [14-17] or are geographically limited (PHS, CAES) [84]. Thus, TEGS-MPV aims to address the latter challenges by sacrificing the former. That is, low cost and geographical independence at the expense of efficiency. To assess the value of RTE relative to CPE and CPP, the simple relation in

Equation 2.6 is used. Here, the CPP for zero net present value (NPV) is evaluated where the total cost is equal to total revenue earned during the system's life, discounted with an internal rate of return (IRR) of 10%, denoted by r . The result is the highest amount an investor would be willing to pay for power costs (e.g. the heat engine) in order to maintain a desired rate or return, given the CPE and RTE. This 10% rate is based on typical interest rates of energy storage systems [85], although the effect of this variable is explored in Table 2.3. For example, Li-ion can only tolerate a rate of 4% for the CPE and CPP given. The other variables in Equation 2.6 include L which is the lifetime in years and $V_{arb}(RTE)$ is arbitrage value in $\$ \text{ kW}^{-1} \text{ yr}^{-1}$, which is a function of RTE, as shown in Figure 2.10. The capacity payment (CP) is estimated based on the average net cost of new entry (Net CONE) of peaking gas turbines as described in the section on Value. Future revenue is discounted with the factor α as discussed above and in Equation 2.5.

$$CPP = \frac{L}{\alpha} (V_{arb}(RTE) + CP) - t \times CPE \quad 2.6$$

Table 2.3: Effect of Rate on Maximum CPP

	CPE	RTE	CPP	Life	Max CPP ($\$/\text{W-e}$)			
					0%	4%	15%	20%
TEGS-MPV	\$36	50%	\$0.34	30	\$3.06	\$1.62	\$0.37	\$0.18
PHS	\$60	80%	\$0.75	30	\$4.90	\$2.60	\$0.61	\$0.31
CAES	\$27	75%	\$0.60	30	\$4.46	\$2.49	\$0.77	\$0.52
Li-ion	\$150	90%	\$0.08	10	\$0.33	\$0.01	-\$0.55	-\$0.71
Lead-acid	\$300	80%	\$0.45	10	-\$1.34	-\$1.63	-\$2.14	-\$2.28
Flywheel	\$2900	60%	\$0.30	30	-\$25.03	-\$26.69	-\$28.13	-\$28.34

Using this simple relationship, the maximum CPP value for zero NPV was estimated for various technologies, assuming the values [13, 84] given in Table 2.4, while alternative IRR's are considered in Table 2.3. Using this framework, any storage technology can be evaluated by knowing its RTE, CPE, and CPP. By using its actual CPE (horizontal axis) and corresponding RTE (color) from Figure 2.11, the maximum allowable CPP is indicated by the vertical axis. If a given technology's actual CPP is lower than the corresponding maximum CPP in Figure 2.11, then it would be profitable under the stated assumptions. Thus, the maximum CPP and actual CPP values for different technologies are then indicated in Figure 2.11, as well as the estimated values for

the TEGS-MPV, which could be profitable based on the relatively conservative baseline results of the following Overall Cost section.

Table 2.4: Value comparison of leading energy storage technologies
 (* indicates not profitable for > 10-hour storage duration)

	CPE	RTE	CPP	Life	IRR	α	Max CPP
TEGS-MPV	\$36	50%	\$0.34	30yr	10%	3.2	\$0.70
PHS	\$60	80%	\$0.75	30yr	10%	3.2	\$1.14
CAES	\$27	65%	\$0.60	30yr	10%	3.2	\$1.23
Li-ion*	\$150	85%	\$0.08	10yr	10%	1.6	-\$0.34
Lead-acid*	\$300	80%	\$0.45	10yr	10%	1.6	-\$1.95
Flywheel*	\$2900	60%	\$0.30	30yr	10%	3.2	-\$27.74

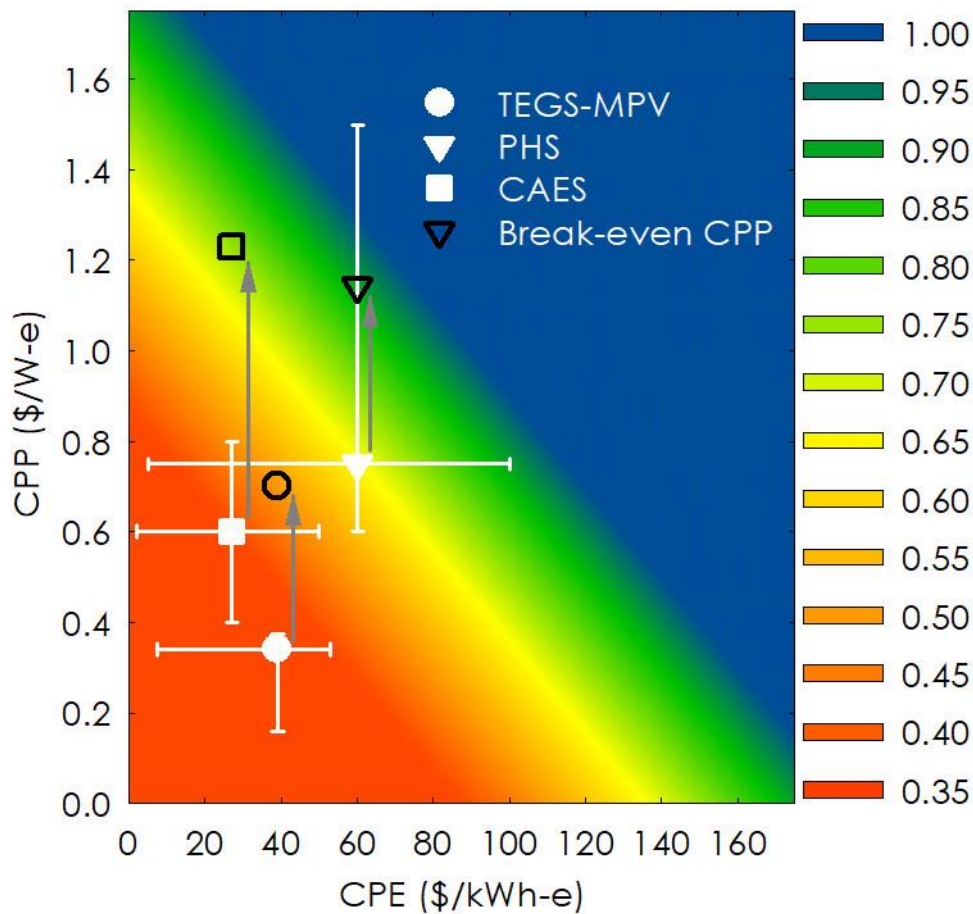


Figure 2.11: CPE and CPP (white shapes) of three competitive energy storage technologies. Arrows and black shapes indicate maximum CPP to break even. Arrow direction and length indicates NPV. The colored contour represents the RTE for zero NPV, assuming 10% IRR, 30 year system, and 10 hours of storage.

The results in Figure 2.11 show that technologies that have very low CPE and CPP values can be competitive even with low efficiency. For example, a storage technology with an RTE of

50%, $CPE < \$50/\text{kWh-e}$, and $CPP < \$0.5/\text{W-e}$ can be profitable like PH, while batteries would not be profitable under the stated assumptions in Table 2.4 [14-17]. This is largely due to the limited cycle-life of batteries compared to the storage of energy via bulk mechanical/thermal methods. To consider the coordinates of batteries on Figure 2.11 assuming their life is ten years (which is not shown because the maximum CPP is negative) their cost must be multiplied by $3\alpha_{10}/\alpha_{30} \cong 1.5$. This results in an effective $CPE = \$230/\text{kWh-e}$ and $CPP = \$0.12/\text{kW-e}$, where discount factor, α , accounts for the shorter loan/payback period of batteries as shown in Equation 2.5.

The second method used to compare energy storage technologies is levelized cost of storage (LCOS) [85]. This method of comparison is powerful, in part because of its wide use, but it does not contain the same information contained in the analysis above—so both are discussed herein. For example, while CPP and CPE are independent inputs into each model, their individual effects are less obvious on an LCOS basis. With Table 2.4 as inputs and a purchase price of $\$0.025/\text{kWh}$ (LCOE) of electricity, Figure 2.12 shows the LCOS of various energy storage technologies as a function of duration. The horizontal gray line indicates the average cost of electricity in the United States (LCOE), which can be considered as something of an upper limit on cost since the cost of stored energy is LCOS plus the original cost of the off-peak electricity. This is not a hard limit though, because not all energy needs to be stored, and thus the price of electricity can be above average for the times when energy must be stored versus used directly.

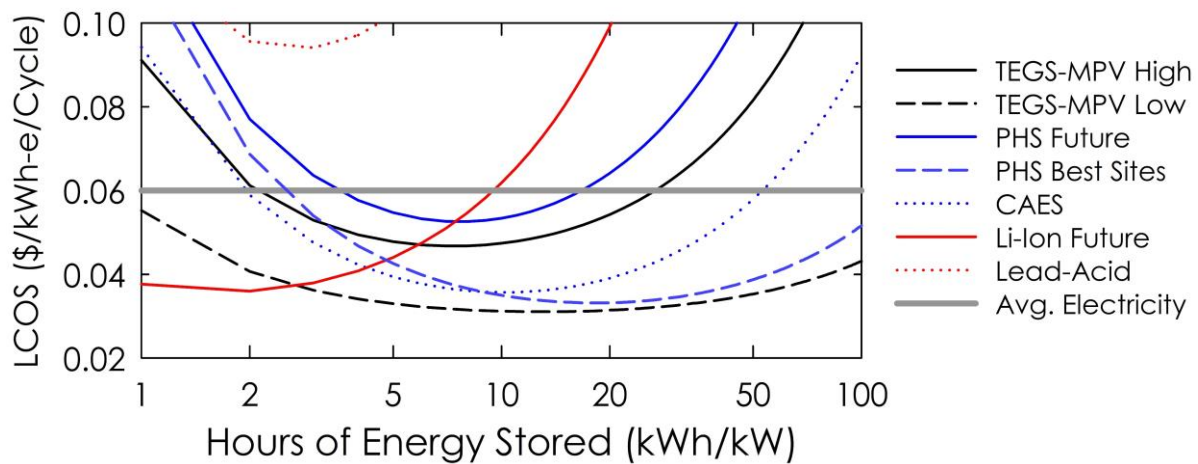


Figure 2.12: Levelized cost of electricity of energy storage versus duration. Batteries become expensive for durations longer than five hours, because their power and energy components are coupled. Other technologies reach a minimum LCOS where $CPE = CPP \cdot h$ where $h = \text{hours}$.

Notably, lithium-ion (red) has a low LCOS for short durations of a few hours—it is not a coincidence that it has been deployed for these precise durations. On the other hand, technologies with lower energy cost (CPE) reach a cost minimum at longer durations, once the cost of their power components (CPP; e.g., turbomachinery) is drowned out. In fact, the cost minimum occurs approximately at the point where $h = CPP/CPE$. In a new way, this plot shows that the baseline TEGS-MPV embodiment (black) is affordable for durations between 2 and 20 hours. Thus, based on this economic motivation, the TEGS-MPV concept that employs ultra-high temperatures and multi-junction photovoltaics (MPV) to achieve a profitable combination of CPE, CPP, and RTE has competitive economics. With this approach, a storage technology that is not geographically limited, yet has similar cost-effectiveness to PH could be realized and could become the most cost-effective embodiment of energy storage.

Thus, given the low CPE and CPP of TEGS-MPV, it can remain profitable down to efficiencies as low as 40%. Given the estimated efficiency of 50% and its position as a profitable, geographically independent energy storage concept in contrast to all existing technologies, a deeper study into its feasibility is warranted. In the next chapters, some key risks to implementing this technology are reviewed, with proposed solutions modelled and experimentally tested to determine if affordable silicon can be contained and pumped by an affordable graphite infrastructure.

Chapter 3:

Joule Heating Input

In this chapter, challenges and solutions to high temperature, high efficiency heating are presented. Although electrical heating above 2000°C has been available for decades for industrial furnaces [26], the techniques employed that enable reaching such high temperatures are fundamentally incompatible with energy storage. That is, these furnaces use minimal insulation and actively water-cooled components to avoid overheating. This thermal design results in nearly all of the heat generated being lost as waste heat. While this is acceptable for industrial processes, where a physical product is produced and sold, for energy storage it is the energy itself that must be sold, so the losses must be much lower (e.g. less than 10% instead of more than 90%). For this reason, an important contribution of this thesis is in the development of a more efficient 2000°C heater. There are significant material and design challenges to this, primarily because the heater exceeded the temperature limit of conventional electrical insulators. This chapter also includes considerations for thermal insulation and temperature measurement.

3.1 Introduction

Although heaters are used very widely in industries and even residentially, the requirements for large scale, extreme temperature energy applications are more stringent. Commercial technologies exist to generate heat from electricity at temperatures above 2000°C, but they suffer major heat loss from active cooling of the electrical insulators and joule heating in the wires and electronics supplying the heater. The problem is that above 2000°C the only affordable insulation is graphite, which is electrically conductive. So, the heater cannot be supported directly by the insulation. Then, the question becomes, how can the heater be supported by an electrical insulator? At lower temperatures, many methods exist but they do not generally transfer to extreme temperatures—most prominently, polymer, silicone, and even common ceramic electrical insulators degrade. On the other hand, existing methods of power conditioning in the cold zone can be applied. That is, by using thyristor based power supplies with silicon controlled rectifiers (SCR) to modulate power by rapidly switching, efficiencies as high as 99.5% [54] can be achieved outside the hot zone. The benefits of thyristor control are discussed in Chapter 2 on System Efficiency and Technoeconomics in the Heater Cost section. So, the focus in this chapter is on reducing or avoiding potential losses and degradation mechanisms in the hot zone.

Here, two methods are employed to minimize heat loss, namely increasing voltage, and removing active cooling. Industrial heaters operating above 2000°C generally operate below 20V and lose heat at high fluxes above 100kW/m² through water-cooled electrical and mechanical heater connections [26]. Notably, the TEGS prototype system would in concept achieve a 20X reduction in this flux, and the full-scale system is designed to have a 300X reduction, as previously discussed in Chapter 2. This relatively low heat flux of 400W/m² is especially effective in combination with the very large size of the system, such that the surface area is relatively small compared to the volume or amount of energy stored. The key challenge to avoiding active cooling is that no conventional electrical insulators can withstand temperatures as high as the 2500°C heater. So, to enable these improvements, an array of geometries and refractory materials were

designed and evaluated. A search for higher temperature electrical insulators was conducted, and some were tested experimentally. Considerations for physically separating the electrically charged heater and wires from the rest of the system are then presented, followed by a similar discussion for avoiding arcing through gasses. Lastly, thermal properties and modeling of thermal insulation are included.

3.2 Electrical Isolation in Solids

The simplest method to reduce heat loss relative to existing heaters above 2000°C is to avoid active cooling all together. However, since the heater has to be electrically separated from the electrically conductive thermal insulation, without active cooling this insulator would have to withstand the peak temperature direct contact with potentially reactive graphite. For example, if the power in a toaster oven was increased until the elements reached 2000°C (and if the heating elements didn't fail first), the supports for the heating elements themselves would melt and the oven would short circuit. So, the first approach was to select and test high-temperature electrical insulators including diamond, hafnium oxide, and silicon carbide. None of these have previously been shown to survive as high as 2500°C, which the heater must reach to heat the storage medium to 2400°C at a sufficient rate. An alternative approach was to employ novel heat transfer and geometry methods to hold the electrical insulators at a lower temperature while minimizing heat loss.

As a first step in the search for an electrical insulator, a literature review was conducted which revealed several candidates, including diamond, oxide ceramics, and a specially doped compound of silicon carbide [86]. Figure 3.1 shows the electric conductivity of these materials as a function of temperature [87-91]. Graphite is shown for comparison, since it is a candidate for the heater itself. Thus, the conductivity of the insulator must be far lower than graphite because the path through the insulation is three orders of magnitude wider than the graphite heater. Based on this, an upper limit of 1 S/cm was considered for the electrical insulator.

There are two key takeaways from Figure 3.1: all insulators degrade with temperature and hafnia is the only insulator with data near the target temperature of 2500°C. The reason for the uniform degradation of all these different materials is that defect density fundamentally

increases with temperature due to entropy, and especially so for compounds where any of the constituents would have an appreciable vapor pressure on their own, such as oxides, nitrides, borides, carbides, and silicides. In the case of ceramic oxide insulators, the vapor pressure of oxygen is of course high even at room temperature and as the temperature increases oxides begin to reduce since most of these materials support some level of off-stoichiometry. The vacancies left behind facilitate electronic conduction similar to doping [92] according to Equation 3.1. Furthermore, the Fermi level increases with temperature providing more positions for electrons, which monotonically increases electrical conductivity as well. This occurs because atoms move faster at high temperatures and, in a solid, more frequently have enough kinetic energy to break their bonds with neighboring atoms, generating a vacancy [93]. For this reason, the same trends are also observed for both thermal and electrical insulators and conductors—conductivity properties approach mediocre values as disorder increases with temperature.

$$\sigma = \sigma_0 \exp(-Q/kT) \quad \mathbf{3.1}$$

The reason data is not available for any material other than hafnia above 2000°C is largely because the other materials melt or decompose at lower temperatures. For example, while diamond is an excellent insulator at low temperatures, it is not thermodynamically stable even at room temperature—it is in fact favorable for it to graphitize albeit so slowly it is in essence kinetically stabilized. However, as a diamond is heated, the rate of graphitization increases exponentially—also via Arrhenius rate dependence [94-96]. Other materials including Silica and alumina melt below 2100°C. Zirconia remains solid but undergoes mechanically destructive phase transformations around 1200°C and 2000°C, and while doping (e.g., with Yttrium) can stabilize the higher temperature cubic phase, these dopants also increase electrical conductivity [97, 98]. One candidate that isn't shown on the plot because of insufficient data is B₄C doped SiC. This material was investigated in the 1980s as an electrical insulator and is documented in a patent, but no recent information was found [86]. Nonetheless, this material warrants future investigation because the patent claims that it remains stable and resistive (10⁻⁸ ohm-cm or 10⁸ S/cm) above 2500°C. Thus, hafnia is the only commercially available (despite cost >\$20,000/kg) electrical insulator above 2000°C that the author is aware of.




Figure 3.1: Electrical breakdown of candidate high temperature insulators. All commercially available insulators other than hafnia (HfO_2) degrade below 2000°C . *SiC-B4C data is based on the patent mentioned [85].

A further reason to consider hafnia as a strong candidate is its commercial use for electrical insulation in the highest temperature thermocouples. Here, hafnia is used at up to 2300°C to physically separate two wires made from a tungsten alloy and enclosed in a sealed tungsten sheath filled with argon gas. On the other hand, the electrical insulator in the TEGS system may also come into contact with silicon and/or carbon, among other possible contaminants. Thus, it was warranted to investigate the effect of these additional elements on the stability of hafnia.

An Ellingham diagram, shown in Figure 3.2, is a tool that can be used to quickly check the stability of compounds (oxides in this case) under various conditions [99]. Many insights can be gained from this diagram. For example, since the y-axis represents the Gibbs free energy, the lower a compound is on the diagram, the more stable. Thus, on this diagram calcium oxide is the most stable compound at moderate temperatures. This means that calcium would tend to reduce aluminum oxide. For the present purpose, one can read along the x-axis to the desired temperature, indicated by the vertical red line. Then, following this red line upward it first intersects carbon monoxide (CO), and then hafnia. This indicates that CO is more stable than HfO_2 , and therefore carbon will tend to reduce hafnia (and all other oxides) above 2000°C . Despite this information, with hafnia as the only remaining candidate for electrical isolation, an

experiment was conducted to determine if this reduction by carbon indeed occurred at a sufficient rate to degrade the hafnia.




Figure 3.2: Ellingham diagram depicting the thermal limit of oxide stability in presence of carbon [101]. The vertical red line indicates the peak system temperature and the angled redline represents the stability of carbon monoxide (CO). Above 2000°C, CO is more stable than all other oxides, including hafnia. So, carbon reduces oxides above 2000°C.

In this experiment shown in Figure 3.3, a tungsten (W) heating element was placed on many small hafnia spacers, which in turn rested on a graphite plate. The heater was powered by

two tungsten leads in the middle of the heater. Unfortunately, as shown in the upper right part of the figure, the hafnia was reduced by the graphite plate beneath it, and the remaining hafnium was both conductive and liquid, which shorted the electrical isolation of the heater. While some improvements could be made to isolate the hafnia from carbon, due to the appreciable vapor pressure (5 Pa [48]) of the graphite, the use of oxides co-located with carbon above 2000°C was avoided based on these results.



Figure 3.3: Hafnium oxide electrical insulators reduced by carbon. The white discs made from hafnia (HfO_2) were reduced to hafnium when heated to 2000°C in the presence of solid carbon beneath them, and gaseous evaporated carbon surrounding them.

In the end, no material was found to be both electrically and chemically stable in the given chemical and thermal environment, except for (potentially) SiC-B₄C mentioned above, which warrants future work [86]. Thus, the alternative approach of using lower temperature electrical isolation, while still avoiding active cooling, was employed. As shown in Figure 3.4, this was achieved by suspending the heater by rigid tungsten electrical leads that are mechanically supported and electrically insulated 250 mm below the heater, where the temperature is much lower around 1000°C. This method effectively isolated the heater from other solids in the hot zone while minimizing the heat loss. Notably, this heater was made from graphite because tungsten is too conductive to be made thick enough to support its weight in this experiment. The thickness can be calculated since the required resistance to generate 10kW at 100V is 1 Ω , the maximum length coil that can practically fit in the 200x200mm space is 3 m, the resulting width

is 5 mm, and tungsten's conductivity at 2000°C is 0.5 μΩ-m. With these values, the thickness can be calculated according to Equation 3.2, resulting in a thickness of only 0.3 mm. At this thickness, the tungsten is approximately as flexible as a sheet of paper, and thus requires complete support. The mechanical strength of tungsten also degrades by at least two orders of magnitude between room temperature and 2500°C [100]. Nonetheless, future, and especially full scale, heaters could employ tungsten due to its lower vapor pressure (0.6 mPa [48]) if alternative geometries are available.

$$R = \frac{\sigma L}{wt} \quad 3.2$$




Figure 3.4: Final heater is suspended by W threaded rod. The W rod serves as a mechanical mount, the electrical supply, and the thermal bridge from 2500°C to 200°C where the tungsten rod is bolted to copper wire.

3.3 Electrical Isolation in gases

Although infeasible at temperatures even as high as 1000C, another path for electricity to travel when temperatures exceed ~1500C, is through the inert gas via arc discharge. In this case, an electrical arc can develop between a region of high temperature and electric field and the surrounding (conductive) graphite insulation, bypassing the heater. Statistically, some gas discharge always happens, but the rate only becomes appreciable at extreme temperatures and/or voltages. Here, due to increases in temperature (i.e. atomic kinetic energy) and electric field, a point is reached whereby the repulsive forces generated during gas molecule collisions exceed the attraction of the electrons to their nuclei and as a result, electrons can be ejected to

form an inert gas plasma. The degree of ionization at which a gas is called a plasma varies, so herein a plasma will refer to a gas that contains enough ions to be a good electrical conductor (in this case the plasma resistance is about 0.5Ω) which can occur even when the vast majority of atoms remain neutral—that is, it isn't a necessary condition for *all* electrons in the gas near the arc to be stripped from nuclei [101]. An image of arc discharge through plasma is shown in Figure 3.5, where Part C shows the arc travelling up the left side of the heater, through the insulation, then down the right side.




Figure 3.5: Graphite heater arcing through argon gas to the surrounding conductive graphite insulation. Part A shows the serpentine coil of the graphite heater before arcing. Part B shows the initiation of the arc, which may actually be similar to Part C except that the camera exposure had not adjusted yet. In Part C, the arc can be seen travelling up from the anode, through the graphite insulation and then back down through another arc to the cathode.

In this case, arc discharge is initiated by thermionic emission which is governed by Richardson's law as shown in Equation 3.3. This equation is valid for a condensed phase emitter in a vacuum without the presence of an electric field [102]. Here, J is the current density or flux that can be thermionically emitted. A_G is a material parameter that generally only varies slightly from a combination of fundamental constants including electron charge and Planck's constant. W is the work function of the material which is the energy required to emit an electron. T is temperature, and k is Boltzmann's constant. Thus, all the variables in this equation are known, and given the $\sim 0.1 \text{ m}^2$ surface area of the heater, the thermal driven emission would be 1200 Amps (A)—absent other factors. This enormous value serves to show that thermionic emission will not only be significant, but also will not be rate-limited by thermally driven emission alone.

$$J = A_G T^2 \exp(-W/kT) \quad \mathbf{3.3}$$

Notably, thermionic emission can occur even if no current is supplied, but once the first electrons are emitted into the gas, they are attracted back to the emitter since it becomes positively charged. However, in this system a current is continuously supplied so this effect is negated up to the applied current, which in this case is limited to 200A by the hardware in the lab. Importantly, Richardson's result does not depend on the voltage of the heater and neglects any electric field present. Here, the thermionic emission is instead driven by the thermal or atomic kinetic energy of the heater. Thus, the addition of voltage provides an additional driving force for arcing. Richardson's law also neglects the space charge effect, which works to reduce thermionic emission. Here, as electrons are emitted locally they repel new electrons, reducing the current. The effect of space charge effect can be modelled using Equation 3.4, Child's Law [103]. Here, A_c is a group of constants, V is voltage, and d is the distance between the emitter and collector. Based on this equation, current would be limited in the TEGS heater to only 100mA. In practice, the current is larger due to the developing plasma which aids the transport of current across the gap.

$$J = A_c V^{3/2} / d^2 \quad \mathbf{3.4}$$

Thus, in addition to the thermal driving force for arcing, there is also a force from the electric field. This can act to increase the current if it is an applied voltage, or decrease the current

via the space charge effect if no plasma forms. Thus, the electric and thermal forces are separately strong enough to cause arc discharge. When the forces are combined, they work together to effectively lower the effective work function of the material via the Schottky effect, inducing more emission [104].

It should be noted that arc discharge is intentionally used in heating applications above 3000°C since no other method exists to heat beyond this point, since nearly all solids would melt. Even the adiabatic flame temperature of fuels is exceeded by electric arc heaters, also known as plasmatrons [105]. Applications for these plasmatrons include reduction of metal oxide ores and cutting/melting metals (including tungsten at 3422°C). Unfortunately, several factors make the current application incompatible with arc heating. First, the erosion rate of the electrode material would necessitate periodic heater replacement, approximately weekly [106], which would be infeasible economically. Also, carefully controlled and forced gas flow is required to direct the arc away from solids. This results in significant heat loss and complexity. Lastly, when sustained arcing develops, the resistance through the arc is very low, such that the voltage drops, and the current must be increased to maintain the heating power—increasing cost and thermal losses. Thus, even though plasmatrons are capable of reaching extremely high temperatures, they are not particularly energy efficient, and therefore are not a candidate for the TEGS-MPV application. Interestingly, in these plasmatrons, tungsten and graphite are used for their high electron emission and arcing is induced by electric field gradients above 10 V/cm. Unfortunately, these materials which arc readily are also the only material candidates for the heating element in the TEGS system.

Given the powerful thermal and electrical driving forces for arcing, which are fundamentally necessary for the function of the heater (supplying electricity to generate high-temperature heat), arcing seems inevitable. However, the above considerations neglect the presence of a gas separating the cathode and anode. In practice, for example in plasma vapor deposition processes [101], gases at moderate pressure often act to enhance arcing by their own ionization. Nonetheless, under the right conditions, gases can also suppress arcing. Fundamentally, this comes down to what happens when an emitted electron collides with a gas molecule. If it is absorbed, the molecule becomes charged, but since its mass is least three orders

of magnitude larger than the absorbed electron, its acceleration and effective speed are essentially negligible—the electron is effectively out of play. That is, the ion moves so slowly that an absorbed electron contributes orders of magnitude less current than a free electron. On the other hand, if instead of being absorbed, the electron breaks one or more electrons free from the molecule, the arc current continues and can even multiply if more than one electron breaks free and joins the arc discharge. This multiplication is known as cascading or avalanche impact ionization [107].

Thus, given that the heater fundamentally generates conditions favorable to arc discharge, and that arc discharge would reduce the ability to heat, as well as the life of the heater, the question becomes: what materials and engineering solutions can be applied to prevent arcing? There are several variables to consider including temperature, gas pressure, gas composition, voltage and heater geometry. Temperature could be reduced, but that would reduce system performance, and was therefore not pursued. Pressure could be modulated, but a premise for the economical deployment of TEGS-MPV is the operation at atmospheric pressure so that a giant pressure vessel to prevent oxygen penetration is avoided. Then, the remaining variables are gas composition, and modifications of the electric field by changing voltage or geometry. It is economically preferred to operate at high voltages because current-carrying conductors can be smaller and have lower losses—so the solution approach was to optimize the gas and geometry, then use the highest tolerable voltage.

At low temperatures, arcing in strong electric fields (e.g. power grid electronics) is suppressed by the use of gases with large electronegative molecules such as SF₆, C₄F₈, CF₂ClCF₂Cl, and other chlorofluorocarbons [108]. These gases are especially effective at absorbing electrons, because the addition of one or more electrons fills their valence electron shell, and thus they act as arc suppressants. Unfortunately, all studied arc suppressant gases decompose at high temperatures—and their constituent elements are reactive with other elements present in the TEGS-MPV system. Based on the behavior of these gases, it was theorized that carbon monoxide (CO) may serve as an arc suppressant, which also happens to be the only non-inert gas that is chemically inert and thermally stable in the TEGS system. So, an experiment was run with 10% CO/Ar balance, but no significant difference in arcing behavior from pure Argon was observed.

Then, since no high performing arc suppressant gases could be used, gases that may at least exhibit a lower rate of avalanche ionization were evaluated. Here, the search was limited to noble gases, since others would react or decompose in this extreme thermal and chemical environment. For example, even nitrogen, which is generally very inert, forms various cyanides when exposed to carbon at high temperatures ($> 1000^{\circ}\text{C}$). Among the noble gases, helium (He) is the most resistant to ionization—that is, it has the highest ionization energy (25eV compared to argon’s 15eV) [109]. Neon has an intermediate ionization energy, but is more expensive than He, so He was pursued instead. Thus, while argon had been considered as the primary gas in the TEGS-MPV system, helium emerged as a candidate to reduce arcing, at a similar cost. However, it should be noted that helium has very high thermal conductivity especially at high temperature (an order of magnitude higher than argon) and this has important implications for thermal insulation, as discussed in the section on Thermal Insulation.

Nonetheless, by having selected helium as a candidate arc suppressant, experiments were conducted using the same geometry but changing the gas between helium and argon. As expected, the temperature-voltage relationship was shifted to the right as shown in Figure 3.6. The key effect was that about twice the voltage could be applied without arcing when helium was used. This result is also in good agreement with similar prior work [110]. So, while some improvement was gained by changing the gas composition, arcing remained a voltage limiter.




Figure 3.6: Electrical breakdown of noble gases. Helium resists arc discharge to about twice the voltage of argon between 1900 K and 2500 K.

To further optimize the heater, its geometry was re-evaluated, and two key changes were made. First, based on the findings from a 3D COMSOL electrical model, the path spacing and curvature were optimized—reducing the peak electric field by more than 25%. These changes essentially involved rounding corners which were arc initiation points (evident in both the model and experiment as shown in Figure 3.8) and increasing the clearance between the heater and the surrounding insulation. A more important (and apparently novel) change was to ground the center of the heater to the insulation as shown in Figure 3.9C. In this way, the maximum voltage difference between the insulation and a part of the heater—and thus the electric field—was reduced by a factor of two. Here, half of the heater was below the voltage of the insulation and half above it. This turned out to have the biggest effect on the breakdown voltage. Overall, by modifying the gas environment, the heater geometry, and the relative insulation voltage, the breakdown voltage was increased by more than a factor of three, which effectively reduced the amount of material needed for current carrying wiring and electronics on the cold side by a factor of three, a significant cost savings. That is, compared to existing heaters above 2000°C which operate below 20V, this heater can operate above 60V which means for a given power only a

third of the current is required. This not only reduces the cost of copper wiring, but also the power conditioning electronics including the thyristors. With all these changes the heater was operated at 70V for four hours at up to 2200°C without arcing or otherwise degrading. The resistance was measured during this experiment to check for degradation and is plotted in Figure 3.7. Here, once the temperature and contact resistance between the heater and the tungsten leads stabilized, the resistance varied less than 1% for more than three hours.




Figure 3.7: Heater resistance remains for three hours indicating no arcing at 2200°C and 70V.




Figure 3.8: Heater electric field reduced by smooth corners **A**: model show electric field peaks in line with **B** experimental observation.




Figure 3.9: **A:** Baseline heater electric field is highest at the cathode. **B:** By grounding the center of the heater, the peak field was reduced by half. **C:** The highlight rod was used for grounding.

3.4 Heater Material

Given the thermal, chemical, mechanical, and electrical, and cost requirements for the heater, the only two candidates are graphite and tungsten. The advantages of graphite are its strength at high temperatures and its relatively high and constant resistance across temperatures, as shown in Figure 3.10 [61, 88, 100, 111]. Tungsten, on the other hand, has a much lower vapor pressure (0.6 mPa compared to 5000 mPa at 2500°C [48]) which can prolong the heater’s life. In these experiments, both were used, but because of the geometry constraints

of this small system and the need to suspend the heater, graphite proved to be a more convenient material. Nonetheless, both remain cost effective and competitive options for full-scale deployment of TEGs. Notably, these are also the only two materials used commercially for heating above 2000°C.

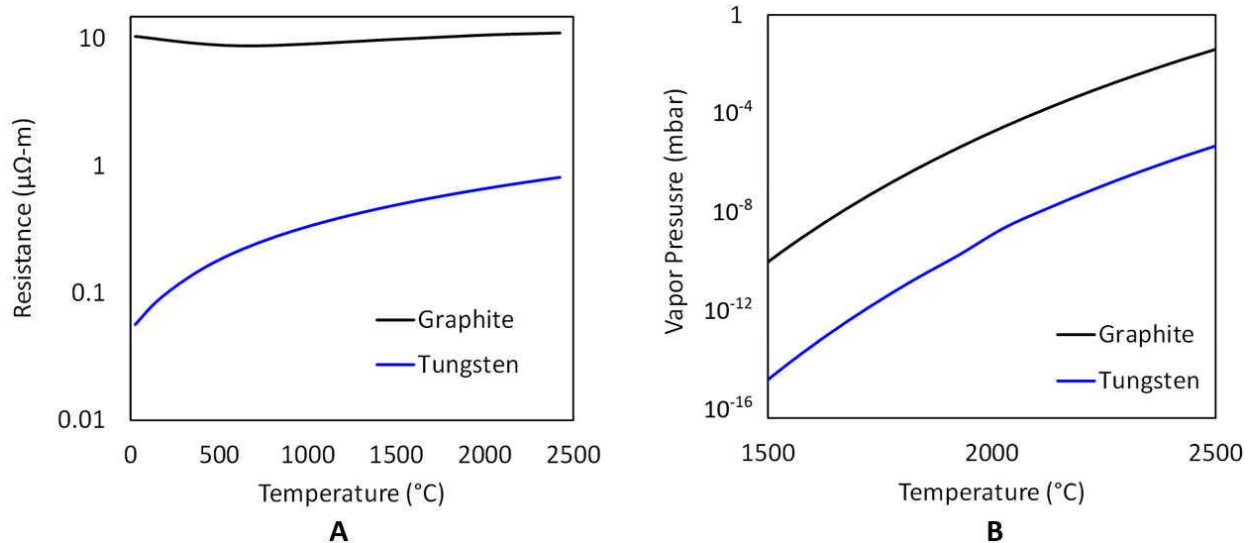


Figure 3.10: A - Electrical resistivity of graphite and tungsten vs. temperature [60, 90, 102, 113]. **B** – Vapor pressure of graphite and tungsten vs. temperature [47].

3.5 Thermal Insulation

At temperatures above 2000°C, the only commercially available thermal insulation is made from highly porous carbon. One alternative is tungsten radiation shields, but these are generally used where heat loss only needs to be marginally mitigated and carbon vapor is chemically harmful [112]. However, since graphite insulation must be used, the temperature limit of oxide insulations is correspondingly reduced by the reducing carbon and carbon monoxide vapor formed from the graphite insulation, according to the Ellingham diagram in Figure 3.2 in the section on Electrical Isolation in Solids. Thus, the presence of carbon limits the thermochemical stability of silica-bearing insulations to about 1500°C. Nearly all high-temperature insulations use at least some silica to bond alumina fibers, so for experiments and cost models herein, graphite is used above 1400°C.

This thick region bridging from 1400C to 2400C (1000°C difference) using the poor performance and high cost of graphite insulation, results in it dominating the volume and cost of

the thermal insulation system. The high effective thermal conductivity of graphite insulation is largely due to radiation being the major contributor [113, 114] in insulations above 1000°C, as shown in Figure 3.11. Due to this, all thermal insulation above 1000°C is effectively an arrangement of radiation shielding [115]. Here, each cell, fiber, or particle in the insulation acts as an optical barrier to photons. This is related to actual planar shields used to block radiation, such as those made of tungsten for use in high temperature ovens, but in this case the “shields” are actually small particles that locally block incoming photons [113]. This is in contrast to the traditional, lower temperature purpose of insulation which is to block convection within the cells. Furthermore, at very low temperatures where radiation is negligible, the best insulation is a vacuum with no solid infrastructure, which is the exact opposite of what is needed at high temperature.

Thus, in this radiation dominated regime, the two most important parameters become the emissivity and number of shields as indicated by Equation 3.5 [116] where q is the heat flux, σ is the Stefan-Boltzmann constant, T is the temperature, N is the number of shields, and ε is the emissivity. Unfortunately the emissivity of graphite is high, around 0.9. Extreme temperature thermal insulation could, therefore, be improved by increasing the density of shields or reducing the emissivity by using alternative materials or coatings, such as tungsten. One area of active research is in high-temperature aerogels, which are very porous and even begin to defy the trends shown in Figure 3.11, by suppressing radiation and even gas conduction due to the pores being smaller than the mean free path of gases [117]. The cost of carbon aerogels is currently far too high for large scale applications and their mechanical strength is very low, although that is also being improved upon by carbon fiber reinforcement [118]. In the future, carbon aerogels could enable compact, effective thermal isolation an order of magnitude better than current graphite insulation because of their higher surface area, which is analogous to a higher density of radiation shields.

$$q = \frac{\sigma(T_1^4 - T_2^4)}{(N+1)\left(\frac{2}{\varepsilon} - 1\right)} \quad 3.5$$

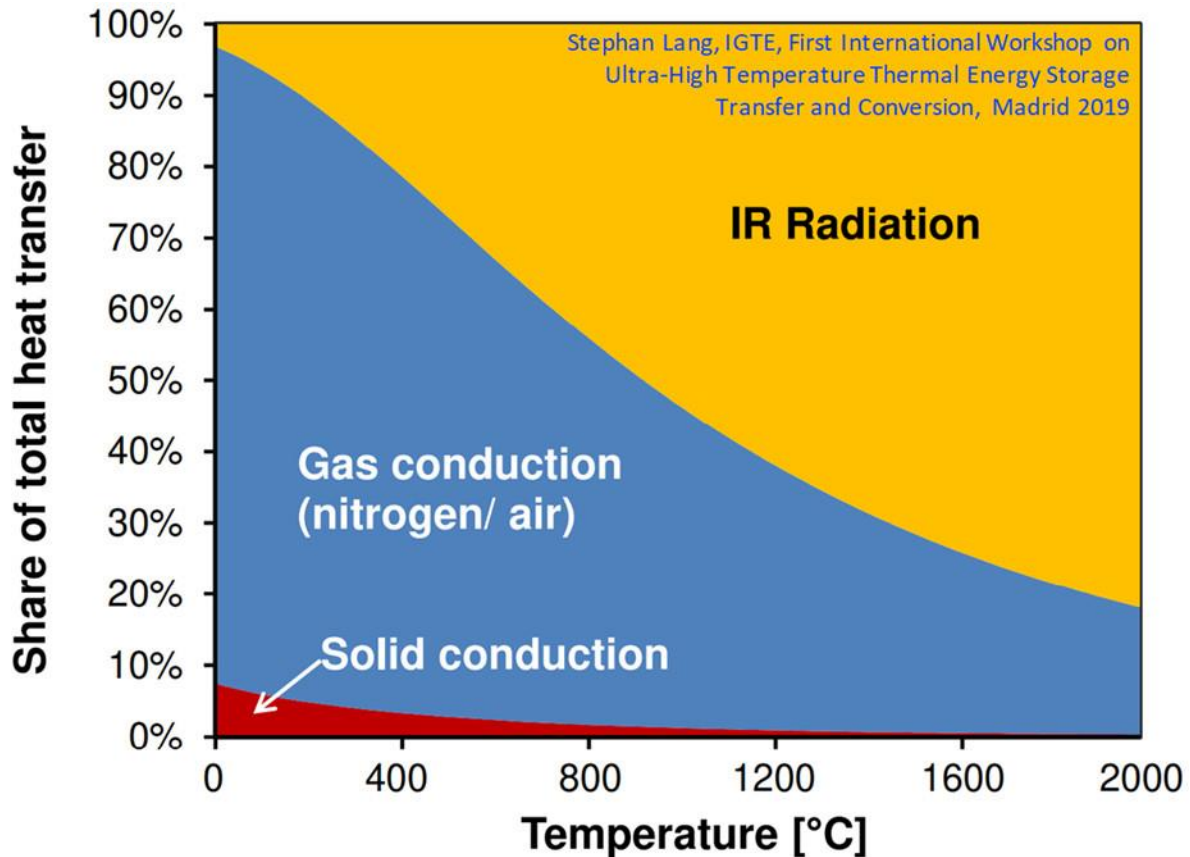


Figure 3.11: Contribution of the three modes of heat transfer to losses in insulation. Below 0°C in vacuum solid conduction becomes important. Near room temperature, the dominant loss is convection. Above 1000°C radiate becomes the main form of heat loss. Adapted from [115].

The heat loss in the TEGS system can be especially high if helium is used to help suppress arcing, as discussed in the section on Electrical Isolation in gases. Unfortunately, the effective thermal conductivity of graphite insulation in helium gas has not been reported, so it was estimated here. The idea was that since insulation suppresses natural convection in gases and the two methods of heat transfer (here, dominantly radiation and gas conduction) operate nearly independently, perhaps the effective thermal conductivity could be estimated by adding the conductivity of helium to that of graphite insulation in a vacuum. This simplistic model turns out to be accurate when tested against available data for graphite insulation in nitrogen. That is, the effective thermal conductivity of graphite insulation has been measured in vacuum and nitrogen yet the difference between these two is almost exactly equal to the thermal conductivity of nitrogen, as shown by the agreement between the solid and dotted gray lines in Figure 3.12. There are two key takeaways from this plot. First, the thermal conductivity of insulation increases

drastically (roughly by a factor of four) from 500°C to 2000°C. This trend is not unique to graphite but also follows for other insulations including the common aluminum silicate [119]. Second, helium approximately doubles the effective thermal conductivity, essentially doubling the cost of thermal insulation. For this reason, the use of helium may be optimally constrained to the heater section or avoided altogether. To verify that the insulation used in experiments matched the available data a limited temperature range of diffusivity measurements were made using laser flash (NETZSCH LFA 457) by collaborators in STEEL Lab at Georgia Tech. These results, shown as blue points, agree well with previously reported values. The error bars show the standard deviation where three samples were measured at each temperature.

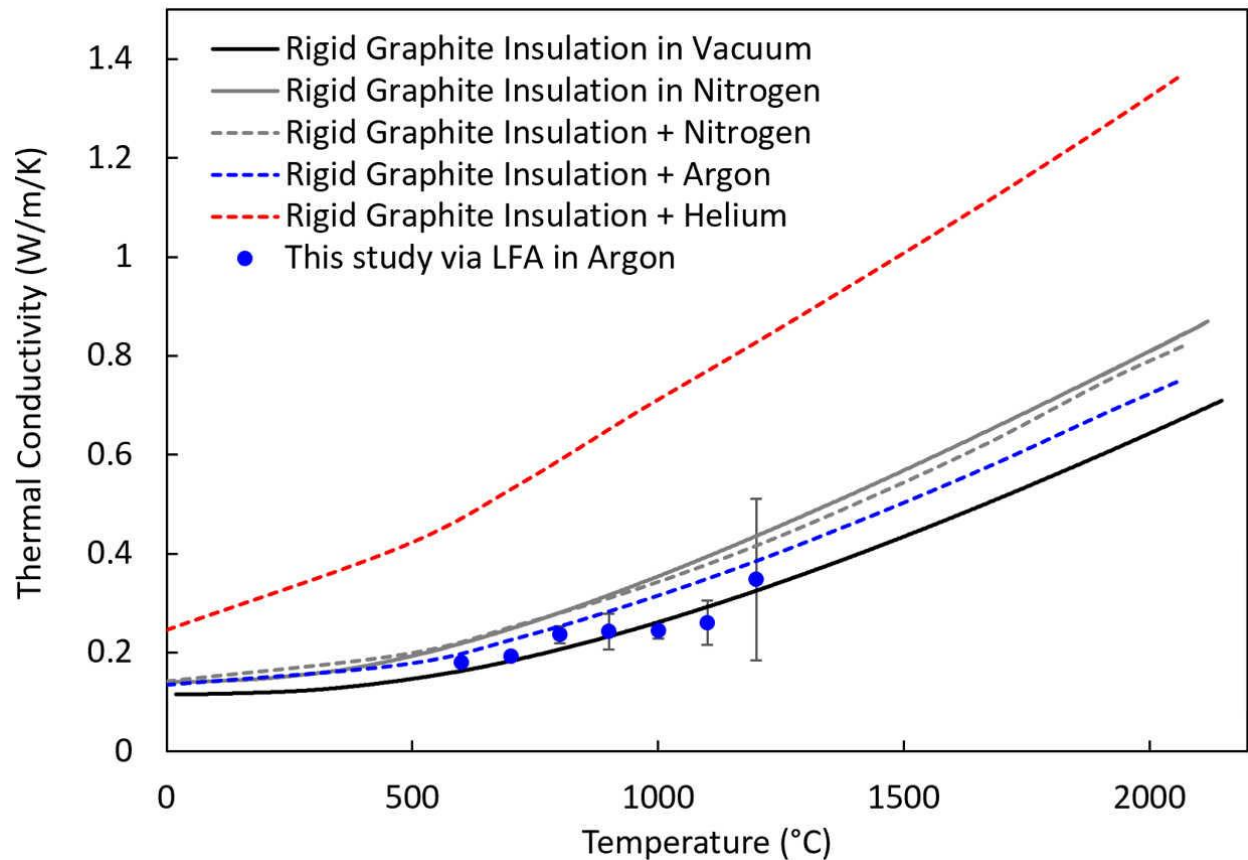


Figure 3.12: Thermal conductivity of graphite insulation in various gases vs temperature. Dotted lines are estimated by adding effective conductivity in vacuum to gas stationary conductivity.

Taking these insulation properties into account, a thermal system was designed to test the heater, which was later used as the basis for the large liquid containment and pumping experiments. This design is built around a 200 x 200 mm heater, which was selected based on

the minimum practical size a pump could be made and the necessary heater length (~3 m) to achieve the needed resistance. The pump size was set based on the desire to operate at low speed (<500 rpm) for this prototype and the required pressure head of 100 mm. With these variables the radius can be calculated based on inviscid limit of Bernoulli's Equation 5.1 in Chapter 5 in the Mechanical section. Then a first analytical model was made based on spherical symmetry to estimate the thickness needed for graphite and oxide insulation layers, based on a target of 5 kW heat loss. This model was based on the steady heat equation in spherical coordinates (r is the radius, T is the temperature) without heat generation shown in Equation 3.6. Here, a fixed internal temperature boundary condition is used for a sphere with 200mm diameter and parallel radiative and convective external boundary conditions are applied. This model resulted in the need of 200 mm of insulation if the average thermal conductivity is 0.6 W/mK—which is conservative based on Figure 3.12.

$$\frac{1}{r^2} \frac{d}{dr} \left(r^2 \frac{dT}{dr} \right) = 0 \rightarrow q = \frac{4\pi k (T_1 - T_2)}{\left(1/r_1\right) - \left(1/r_2\right)} \quad 3.6$$

Next, a 3D finite element model was developed in COMSOL. This model used more accurate geometry and temperature-dependent material properties. It was checked for mesh convergence by decreasing element size until the solution became independent, which in this case occurred as long as elements were smaller than about 1 inch as shown in Table 3.1. Here, the heater power was held constant while the steady state temperature was compared as the independent variable, for the chosen mesh. The results varied by less than 1% from the result with many more elements, and thus a maximum mesh element size of 0.75 in was considered converged, and subsequently used. A transient version of the model was run to estimate the time required to reach the 2000°C, which turned out to be several hours for 5-10 kW heat input. The model was also used to inform the design of the heater electrical connections, since at the base of the insulation the tungsten conductors interfaced with copper wires, as shown in Figure 3.13. Initially, the base of the heater was predicted to be as high as 1000°C, due to the conduction of heat down the tungsten conductors, and the radiative transmission in the gaps surrounding these leads. Notably, despite this temperature, the heat loss through these leads was less than 100W, because of their small cross-section area. Nonetheless, to avoid overheating the copper wires, a

thermally conductive, electrically insulating plate (i.e., a heat spreader) was used to secure and passively cool the base of the tungsten rods and maintain them below 400°C. For these tests, a thermocouple was embedded in a thermal slug above the heater as discussed in the section on Temperature Measurement.




Figure 3.13: Heater geometry (left) and thermal model (right). The heater is suspended above the insulation by tungsten conductors. Electrical isolation is by solid separation until the bottom of the insulation where an AlN heat sink serves the dual purpose of passively cooling the tungsten conductors and electrically isolating them. The thermal model takes the heater temperature of 2150°C as an input and includes convective and radiative external boundary conditions. Internal radiation in the heater cavity is computed with ray tracing.

Table 3.1: Mesh convergence of heater thermal model

Mesh Quality	Maximum element size (inch)	Minimum element size (inch)	Number of elements	Heater Temperature (°C)	Error
Very Coarse	4	0.75	48,600	2275.49	-9%
Coarser	2.75	0.5	61,025	2251.49	-7%
Coarse	2	0.701	90,986	2126.71	-1%
Normal	1.5	0.45	132,450	2086.11	0%
Fine	1.375	0.25	191,815	2106.45	-1%
Finer	0.75	0.1	391,446	2098.5	0%
Extra Fine	0.25	0.0375	1,122,864	2093.01	0%
Extremely Fine	0.05	0.005	4,112,308	2095.79	0%

3.6 Temperature Measurement

In this work, thermocouples and optical pyrometers were used to measure temperature. Each method has strengths and weaknesses. Thermocouples are widely used for many

temperature ranges, where copper-based alloys are typically used below 500 °C, while nickel alloys are used up to 1300 °C, then platinum alloys are used to 1700 °C and finally, tungsten alloys are used at even higher temperatures. A thermocouple works by generating a voltage when dissimilar metals are held at different temperatures via the Seebeck effect [120]. Due to the electrical nature of thermocouples, and as discussed in the section on Electrical Isolation in Solids, the temperature limit of tungsten-based thermocouples (usually referred to as Type-C) is actually set by the electrical insulator and not the conductor. Most high-temperature thermocouples use alumina insulators, but at higher temperatures, hafnia is used for its thermal stability to 2300°C. Although hafnia is stable above this temperature, owing to its increasing conductivity it acts as a partial short circuit and also starts to decompose in the reducing atmosphere sealed inside the thermocouple to protect the tungsten. Some other ceramics including beryllium oxide and thorium oxide have been used at even high temperatures, up to 2850°C, but they aren't commercially available owing to their high cost and toxicity [121]. Given the 2300°C limit of commercially available thermocouples, and the fact that they can be degraded by interaction with the rest of the high-temperature system, pyrometers were used as an alternative for even higher temperatures.

Optical pyrometers function by measuring the rate and wavelength of photons emitted by a surface to determine the temperature of a gray emitter that would produce that spectrum. The simplest pyrometers are known as IR thermometers, and they are very cheap and convenient non-contact methods to measure the forehead of patients to check for a fever. As temperature increases, the difference between IR and visible (solar and artificial light) become less distinguished, and it becomes especially important to not only sum the photons, but also measure their wavelength to accurately predict temperature. The best optical pyrometers use multiple sensors with a narrow bandpass filter to measure emission at two specific wavelengths relevant to the temperature range of interest [122]. Optical pyrometers were used for the highest temperature measurements in this work, and are capable of measurements above 3000°C. The main drawbacks of this method are first the need for line of sight, which necessitates a relatively large (~10 mm diameter) path through the insulation and chamber wall to the pyrometer. This

increases heat loss and can become blocked. Secondly pyrometers have a lower limit of operating temperature, in this case $\sim 1000^{\circ}\text{C}$.

Chapter 4:

Liquid Containment

Typically, material pairs for liquids and containers are selected to avoid reactivity. For example, polymers hold water, and steel contains acetone, but not vice versa. However, silicon was not selected for its inertness, it was selected despite its reactivity. That is, above 2000°C, there does not exist a low-cost material pair (liquid & container) without reaction or dissolution. Nonetheless, prior work has shown a possible path for affordable silicon containment via a protective SiC layer.

Viewed from another perspective, graphite is the only container candidate under these conditions. Among cheap materials (i.e. < \$10/kg), graphite is the only one that retains its mechanical strength above 2000°C—and its strength actually increases to 2600°C [61, 111]. Other refractories soften and eventually melt at lower temperatures, except for some exotic, cost-prohibitive, and/or toxic compounds like thorium oxide (thoria, ThO₂) and hafnium carbide (HfC) [123]. Graphite could be used to contain liquid tin without reaction, but tin is too expensive for bulk storage. Quantitatively, tin costs about \$750/kWh-e compared to \$23/kWh-e for silicon for a temperature change of 500°C, or several times the cost of Li-ion [38]. Exotic ceramics would also be too expensive, and the narrow availability/use brings into question if the ~1,000 tons of material needed could even be produced [123, 124]. Thus, in this chapter, the graphite-silicon material pair is investigated to determine if affordable grades can withstand the containment conditions present in the TEGS-MPV system.

4.1 Silicon Carbide

Silicon quickly reacts with carbon to make silicon carbide, even near its melting point (1414°C) [21]. However, under some conditions this reaction forms a thin layer of SiC which acts as a mass diffusion barrier, impeding further reaction. The electronics industry takes advantage of this, by using graphite to contain molten silicon during the formation of single-crystal silicon ingots, known as the Czochralski method [21, 22, 25]. While this is very encouraging, the conditions present in the TEGS-MPV system are much more extreme. So, a major effort in this work was to evaluate whether SiC can remain protective with cheaper graphite, less pure silicon, higher temperatures, longer life, multi-component tanks, and a less controlled gas environment.

To begin, the experimentally determined phase diagrams were reviewed as shown in Figure 4.1, which indicate two important facts [125-127]. First, SiC itself is thermodynamically stable, up to about 2700°C, which is above the maximum temperature of interest here (2400°C) [128]. Fundamentally, this means that the enthalpic driving force (ΔH) for the bonds between Si and C atoms exceeds the entropic driving force ($T\Delta S$) for dissociation up to this temperature. Thermodynamically, the Gibbs free energy (ΔG) shown in Equation 4.1 for SiC at 2700°C changes from negative to positive, indicating thermodynamic equilibrium for dissociated Si and C.

$$\Delta G = \Delta H - T\Delta S \quad 4.1$$

$$x = \exp \frac{\Delta H_f}{R} \left(\frac{1}{T_f} - \frac{1}{T} \right) \quad 4.2$$

Second, there is significant solubility of carbon in silicon at 2400°C [129], about 5 at% as indicated by the intersection between the red and black lines on the right of Figure 4.1. Solubility can be estimated using Equation 4.2, which is only strictly valid for ideal solutions but the trend is nonetheless applicable. Here, x is the solubility in terms of the mole fraction, ΔH_f is the enthalpy of formation, R is the gas constant, T_f is the melting or dissociation temperature, and T is the operation temperature [130]. This solubility increases rapidly with temperature until it approaches the melting point, where it becomes concave down and asymptotically approaches 100% solubility at the melting or dissociation temperature. Since the phase diagram only represents the result after infinite time at a given temperature and pressure, it does not teach us

whether dissolution of graphite will actually occur in the TEGS system. This dissolution may be mitigated by the SiC layer, or by other resistances. On the other hand, if carbon does dissolve, it may remain dissolved indefinitely in the system, or precipitate out in the colder regions. Because of this uncertainty, a set of experiments was devised to check for dissolution, which is described in the section on Thermal Cycling.




Figure 4.1: Si C Phase diagram. The horizontal red line shows the peak TEGS temperature.

Another important question is how quickly the SiC layer grows since it could clog pipes or cause leaks on the decades time scale. The rate of layer growth is governed by either Si or C diffusion through SiC, whichever is faster. Mass diffusion is governed by Fick's law as shown in Equation 4.3, which is analogous to Fourier's law for heat conduction. For a fixed diffusion coefficient, D , and mass concentrations at the boundaries, this concentration gradient ($\bar{\nabla}m$), and thus diffusion rate (\bar{j}_1), decays parabolically with layer growth. The diffusion coefficient, however, increases with temperature according to the Arrhenius relation in Equation 3.1, described in the section on Electrical Isolation in Solids in Chapter 3 on Joule Heating Input. The

diffusion coefficients for both C and Si in SiC have been experimentally determined [27, 28] up to 2200°C and follow an Arrhenius curve well, and their slight extrapolation to 2400 °C is reasonable if the diffusion mechanism remains the same. The diffusion of each element are relatively similar (although carbon diffuses about 100 times faster than silicon) and on the order of 1×10^{-14} m²/s as shown below in Figure 4.2. This value is similar to the rate of metal/oxygen diffusion through copper [131] and aluminum [132] oxides at lower temperatures—which are important protective/ passivation layers—and results in uniform thickness growth of about 1 mm in 30 years. Notably, this degradation could be dramatically accelerated if the SiC layer delaminates or spalls off during thermal cycling, so this was a focus of the Thermal Cycling experiments. Nonetheless, since it appeared *possible* that silicon could be contained in graphite above 2000°C based on this initial theoretical analysis, a set of initial experiments were conducted to see if affordable grades of silicon and graphite could survive.




Figure 4.2: Diffusion coefficients of silicon and carbon in silicon carbide [27]. The straight lines are Arrhenius plots because of the log y axis and the inverse x axis according to Equation 3.1. Carbon diffuses two orders of magnitude faster than silicon, so most SiC growth is due to carbon diffusion toward the liquid silicon.

$$\vec{j}_1 = -\rho D \vec{\nabla} m$$

4.3

4.2 Proof of Concept

An early baseline experiment was performed to determine if affordable silicon could be contained in affordable graphite made from multiple sections above 2000°C. In this test, a dense (1.85 g/cc) graphite (KYM-20) small tank was filled with 553 grade Si was heated above 2,000°C for 60 minutes. The tank was made from two sections and sealed with a thin graphite foil face seal that was compressed by carbon fiber composite (CFC) threaded rod and nuts, as shown in Figure 4.3. Initially a thicker seal was used, but as described in the next section on sealing, it was fully penetrated by silicon and expanded until the bolts broke. The tank was insulated with graphite felt and aluminum silicate insulation inside a quartz tube, under high purity Argon gas ($< 1 \times 10^{-6}$ atm O₂). The tank was heated by induction and its temperature was measured using a C-type thermocouple.




Figure 4.3: General experimental setup for small (~50 mm diameter) tanks heated by induction

Here, similar to prior work at lower temperatures, Si penetrated the graphite tank approximately 400 μm before creating a dense 20 μm thick SiC layer at the surface, preventing further penetration as shown in Figure 4.4. In this test, a carbon fiber composite (CFC; graphite with ~5 mm fiber whiskers) threaded rod was submerged in the silicon to check for degradation since in the full-scale design some rods are inside the tank. A thin SiC layer formed here as well. Later experiments with similar materials exceeded 2300°C and 4-hour duration, with no significant degradation or differences from this initial experiment.




Figure 4.4: Graphical induction heating experimental process. A 100mm tall tank and ¼-20 CFC bolt were heated to 2200°C for four hours and both formed a protective SiC layer as shown on the right.

Based on this successful containment, a larger tank designed to fit a pump was also tested using the same 553 silicon. This tank was nominally 200 x 200 x 250 mm tall and was made from 6 graphite plates bolted together and sealed with the same graphite foil sheet as the smaller experiments. This two-gallon tank was heated by resistive heating described in Chapter 3 on Joule Heating Input to 2200°C for four hours. The tank did not leak during the test, although it did crack on freezing due to the expansion of silicon on freezing, similar to how a glass containing water can crack in a freezer. Similarly, a dense 20 μm thick SiC layer formed at the surface and even on the lid of the tank, as shown in Figure 4.5.

This coating of regions that were not submerged occurs mainly because of the significant vapor pressure (4 mbar) of silicon at 2200°C [48]. Estimating a lower bound on the mass flux from the liquid silicon surface to the 100 mm away, Fick’s law was used from equation 4.3, neglecting advection. The diffusion coefficient for dilute silicon in helium at 2200°C is $\sim 10^{-4}$ m²/s based on Chapman–Enskog theory [133], so given the boundary conditions above the mass flux is about 1 mg/m²s, forming a 20 μm layer in a few hours. This rate was accelerated by significant natural convection in the tank caused by heating from the bottom such that the formation of the SiC layer shown in is expected on all interior tank surfaces, submerged or not. In later tests the lid had holes to support a pump, evaporated silicon travelled through these holes, reacting with the

nearby graphite insulation as shown in Figure 4.6. This test also demonstrated for the first time, a situation where the mass of silicon significantly exceeded that of graphite. These tests served as initial evidence that a path existed to the containment of liquid silicon in graphite above 2000°C, but several important questions remained, which are addressed in the following sections.




Figure 4.5: Two gallon tank successfully contained 553 Si at 2200°C for four hours




Figure 4.6: Evaporated silicon reached the graphite insulation, reacting to form porous SiC

4.3 Sealing

Sealing is one of the biggest challenges at high temperatures since seals generally require a soft material which conforms to the imperfections between two harder materials. Typically, polymers are used as seals because they are composed of strong polymer chains which are only weakly attached to each other, enabling them to slide. However, polymers decompose at much lower temperatures (generally below 300°C), so they couldn't be used here [31]. Instead, sealing was achieved through the use of flexible 'flake' graphite. This graphite is flexible because it consists of largely aligned graphene layers, unlike the isotropic graphite used for the tank itself. These layers behave like polymers mechanically because they are only bound to each other with weak Van der Waals forces such that they can slide and deform to match the shape required, similar to how polymer chains can slide.

There are two general classes of seals: dynamic and static. Dynamic sealing, which is more difficult to achieve, is used between surfaces that move with respect to each other. A dynamic seal is usually used between a pump or valve shaft and a pump or valve body to prevent leaks. However, we avoid the need for a dynamic seal in this work by using a sump pump. Here, the pump is located *inside* the tank, submerged in the liquid, so that the small amount of liquid leaking along of the pump shaft remains inside the tank. The full-scale embodiment of TEGS would use the same type of simple centrifugal pump, since the required pressure is low unlike existing molten salt energy storage tanks today which are coupled with ~100m tall concentrating solar power towers [134].

Nonetheless, using a liquid storage medium requires that it be impermeable, so static seals are needed throughout the system. Seals are needed between each section of the tank, which is too large at scale to be made from one piece. The gasket was compressed using CFC threaded rod and nuts, which are importantly much stronger than graphite (200MPa vs 30MPa tensile strength). This high strength ratio enables the total rod cross-section area to be smaller than the tank wall cross-section. That is, since the tank wall is loaded in circumferential tension, the bolts and the wall are subject to an equal tensile *force*. Thus, the required area to withstand the applied force depends on the strength (which has units of pressure) of each material since $Area = Force/Strength$. If the bolts were the same graphite material, the same cross-section

area of bolts would be needed as the tank wall itself, but since these two materials overlap, this would result in very thick flanges (4X wall thickness to clear bolts and nuts) and many more bolts. Instead, the strong CFC bolts enable a compact system that takes full advantage of the mechanical strength of both materials—an important cost savings for the tank material as discussed in the section on Storage Cost.

Related to the high strength of CFC is its low rate of thermal expansion. These properties are linked, just as stiffness and strength are generally linked, fundamentally because the intermolecular bonds holding carbon atoms in place in carbon fiber reside in a deep, narrow potential well [135, 136]. Here, strength is related to the depth of the well, stiffness is related to the width, and thermal expansion is related to the slope of the walls of the potential well. Since the thermal expansion coefficient of the CFC is low compared to graphite (1 $\mu\text{m}/\text{mK}$ compared to 5 $\mu\text{m}/\text{mK}$), the stress on the rods becomes tighter as the tank is heated. This can be used as an advantage, since the bolts do not have to be fully tightened at room temperature but must be taken into account in the design of graphite/CFC bolted systems.

Initial experiments used a thick (~5 mm) gasket shaped similar to an O-ring and fit into a matching groove, which was made from compacted graphite foil in the form of a rope. This material is commercially available and even used in some industries as a dynamic seal, not for high-temperature capability, but for chemical resistance to solvents that would damage polymers. This material has recently been successfully used to seal pumps at up to 1400°C, but here the liquid was Sn and thus did not react with the seal [32]. The O-ring like geometry is often preferred for sealing because it enables a higher sealing pressure by reducing the seal area. This increased pressure can provide additional force to squeeze the seal into surface scratches, reducing the leak rate. Here, the idea was that when the seal was compressed, it would form a dense, conformed barrier to silicon penetration, with only a shallow surface reaction—as had occurred with the bulk graphite.

However, when reactive silicon was melted in a tank sealed with this material, it penetrated deeply into it, reacting to form SiC. This reaction is volume expansive, despite the higher density of SiC than either Si or C. The expansion occurs essentially because silicon is added to the original carbon. In this case, expansion or contraction depends essentially on the choice of

control volume and naming conventions. That is, if silicon and carbon are considered together before the reaction, a net volume decrease is observed [137]. However, here carbon is considered to be the initial system, since as silicon is added, the net solid volume increases, which is the relevant parameter for sealing and strain. The amount of expansion can be calculated using Equation 4.4, where V_f is the final volume, V_i is the initial volume, and ρ is density. The factor $40/12$ is the ratio of the molar mass of SiC to C. Since SiC is less than 3.3X the density of C (its only 1.6X), the volume increases by a factor of two. For a 5 mm thick seal, this amounts to 5 mm of expansion, far exceeding the strain limit of the threaded rod holding the tank together. The large forces that were generated by this volume expansion not only broke the threaded rod, but even the tank wall itself, as shown in Figure 4.7.

$$V_f = V_i \frac{40}{12} \frac{\rho_{SiC}}{\rho_C} \quad 4.4$$




Figure 4.7: Mechanical failure of thick graphite seal due to doubling of the seal thickness from reaction with infiltrated silicon.

As a result of the expansive failure of a thick seal, future seals were made thin (0.1 mm), avoiding this mechanical issue, at the expense of reduced seal pressure due to the larger seal area. The reason a thin seal leads to an increase in seal area is that the thickness change is accompanied by a material property change, and the thinner seal is less compressible and conformal. Since the thin seal cannot be forced to deform into scratches as well as a thicker, softer seal, in this case, the leak rate is minimized by spanning a large area and completely covering smaller scratches.

Fortunately, in this case, the SiC expansion was actually advantageous. That is, the expansion is very effective at closing off leak paths. This can be seen directly in Figure 4.8, where

silicon leaked along the sealing surface until the conversion to SiC stopped further leaking. Thus, SiC not only acts as a self-healing reaction barrier, but also a self-sealing leak preventer. The ability to reliably seal a graphite tank made from multiple components is an important step toward realizing the full TEGS system since a large (~10 m) tank must be made from multiple parts.



Figure 4.8: Self-sealing behavior of SiC reaction. After leaking 1 mm along the seal path, the leak was blocked by the formation of SiC.

4.4 Effect of Graphite Pores

It has been established that high-quality graphite exposed to molten silicon can form a thin SiC layer slowing further reaction, but given the large volume and low cost requirements to build the TEGS system [16], it was important to determine precisely how “high quality” the graphite must be. The microstructure of the graphite is a key factor in whether this thin layer forms, best represented by the distribution and size of pores—with average pore size, particle size, grain size, and porosity as approximate representations. The porosity and density are also mathematically related, but the size is more important than the quantity, so neither is an ideal representation. The degree of interconnection of pores is also important. That is, closed-cell porosity is not a problem, because silicon cannot reach those pores to continue penetrating. Just as the thick seal in the previous section failed because silicon penetrated too far into it and converted too much of it to SiC, graphite itself can mechanically fail if it is too porous. This has

been observed in prior work [21, 22], and was also demonstrated here as part of the effort to determine the maximum tolerable pore size.

Isotropic (i.e. rigid) graphite has three broad classifications based on the manufacturing method: extruded, uniaxially pressed (also called 'molded'), and isostatically pressed (in order of increasing quality/cost). Notably, on a molecular scale graphite is fundamentally anisotropic and it only becomes relatively isotropic when many grains are sintered together in a random pattern. In all of these, initial particles are pressed together and heated above 2500°C, forming a sintered bulk solid. The pore size is then strongly related to the initial particle size (and by extension, grain size) since pores are the gaps between particles/grains. However, the final grain size and shape depend on processing, as shown in Figure 4.9, so even grain size is an imperfect indicator of pore size. Pores can even be closed by grain growth during longer periods of heating, especially under compression [138].

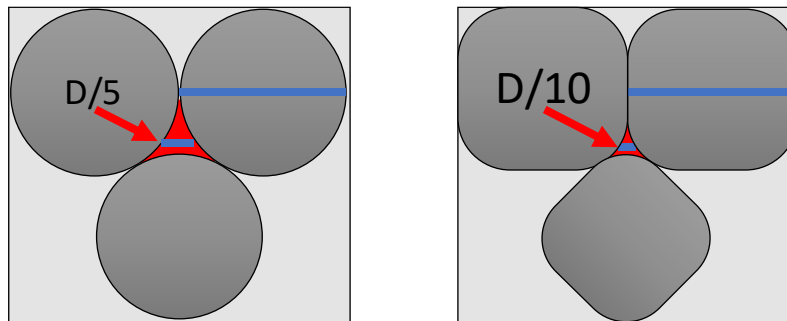


Figure 4.9: Illustration of how pore size depends on initial particles and processing. The dark gray particles on the right indicate the effect of increased time, temperature, and/or pressure during the sintering process.

For extruded graphite, the grains are about 100-1000 μm in diameter, while for pressed graphite they are in the range of 1-10 μm . White and Israel [21, 22] showed that some grades of extruded graphite mechanically fail under Si infiltration due to their macroscopic interconnected pores and the internal stresses developed during reaction. In this thesis work, it has also been shown that for some extruded structures, silicon may leak directly through the wall without causing mechanical failure but failing to contain silicon nonetheless, as shown in Figure 4.11. On the other hand, White and Israel [21, 22] also showed that high-quality isostatic graphite with grain sizes smaller than 15 μm form dense SiC layers under a limited range of conditions. However, these tests were done with 6-9N (6N is 1 part per million impurity), pure silicon

(~\$500/kg), 6-9N pure gas environments, and 1000°C colder than TEGS. Notably, even though these prior experiments were at much lower temperatures, SiC formation occurs just after melting, so the results apply to TEGS—at least for the tanks. Even if the initial exterior layer delaminates occasionally at higher temperatures, the internal pores are likely to remain closed by SiC while a new layer forms. Essentially, a certain combination of size and interconnectedness of pores is required to generate a sufficiently tortuous path, such that the silicon reacts and seals itself before it can leak or overstress the bulk graphite infrastructure.

To determine the critical pore size, a simple model was developed which was compared to more detailed reactive infiltration models and experiments above 2000°C. The basic model treats a pore as a cylindrical hole with radius as the critical variable, as illustrated in Figure 4.10. There are two fundamental reasons why small pores clog at short depths. First, leak velocity decreases with radius since resistance to flow is proportional to r^2 , which is a stronger function of radius than the capillary driving force for flow, which is proportional to r . Second, the time to close the pore also decreases with radius.




Figure 4.10: Conceptual illustration of pore closure by SiC reaction competing with leak rate. As the blue SiC layer grows, the leak rate decreases and eventually stops when the SiC fully closes the pore.

A comparison is made between the leak rate at the sealing rate. The leak rate is capillary driven lubrication flow as shown in Equation 4.5 since $L \gg r$ and $\sigma/r \gg \rho gh$, where V_{Si} is the leak velocity, L is the length of the hole (~10mm), r is the radius (~1 μ m), σ is the surface tension between silicon on graphite in helium (0.7 J/m [137]), μ is the viscosity of silicon (like water 10^{-3} Pa-s), ρ is the density of liquid silicon (2300 kg/m³), g is the acceleration due to

gravity, h is the hydrostatic height of silicon (only $\sim 0.1\text{m}$ above reaction plane). The reaction is rate limited by mass transfer through the SiC layer, since the diffusion coefficient, $D_{Si,SiC}$ is so low ($10^{-14}\text{ m}^2/\text{s}$). The metric for successful pore closure is a penetration depth of less than 10 mm, which is selected to avoid leaks and retain most of the tank wall as graphite (not SiC) for mechanical stability. So, when these equations are solved for r_{cr} , it is found to be $0.3\ \mu\text{m}$. While this value is quite small, it represents the largest acceptable average pore radius. Thus, when considering that pores are generally an order of magnitude or more smaller than the particles (see illustration in Figure 4.9), this suggests a maximum particle/grain size of $6\ \mu\text{m}$, which is in the middle range for pressed graphite grades.

$$L \approx V_{Si}t \quad V_{Si} \approx \frac{r\sigma \cos \theta}{4\mu L} \quad t \approx \frac{V_{SiC}}{r} \quad V_{SiC} \approx \frac{D_{Si,SiC}}{r} \quad 4.5$$

$$r_{cr} \approx \left(\frac{4\mu D_{Si,SiC} L_{\max}^2}{\sigma \cos \theta} \right)^{1/3} = 0.3\ \mu\text{m} \quad 4.6$$

This result is in strikingly good agreement with a more detailed model proposed by Chiang and others for the design of graphite preforms for SiC production via reactive infiltration [137]. Here, Messner and Chiang modified the century-old Washburn equation for non-reactive infiltration [139], using more accurate geometry parameters, accounting for the non-linear acceleration of fluid velocity and layer growth, and including kinetics in addition to mass transfer. Nonetheless, their result is within 10% of the more basic approach presented above. This result also agrees with more recent work by Israel and others within 30%, who predicted a critical pore diameter of $1\ \mu\text{m}$ [22]. Notably, a single large pore is unlikely to cause failure, so long as it does not fully penetrate the wall, or at least not for long, because the surrounding graphite can suppress the crack propagation from a small number of holes. Also, this result is conservative because it assumes straight cylindrical paths and a free supply of silicon, but actual paths are more tortuous, and the supply of silicon is usually during melting, where two phase flow may add resistance.

Given that theory predicts an intermediate grade of graphite can be used, an array of isostatic and molded graphite materials were tested under TEGS conditions (above 2000°C , in commercially available Ar gas, in contact with 553 grade Si) to determine the most affordable

graphite grade that can be used for TEGs. Some experiments were also performed with extruded grades (GR008G and ASCX) that confirm the previous work showing their failure due to large (10-100 μm estimated based on 200-1000 μm particle size) pores as shown in Figure 4.11.




Figure 4.11: Extruded graphite grades mechanically fail (A) and/or leak (B) due to large pore size.

More than ten grades of pressed graphite were tested, ranging from uniaxially pressed with particles as large as 15 μm and costs as low as \$3/kg to isostatically pressed with particles as small as 1 μm and costs as high as \$100/kg. A selection of these results is shown in Figure 4.12. All of these pressed grades successfully contained silicon above 2000°C without leaking or mechanical failure. This is due not only to the smaller initial particle sizes compared to extruded graphite, but also to the longer, high-pressure sintering process, which further reduces pore size. Testing conditions ranged from 2000-2300°C and 60-240 minutes, with no significant difference observed between conditions or materials.

The primary grade considered for use in the full-scale system has a cost at a large scale of \$7/kg and has multiple trade names, such as KYM-20, AS-TJ, AR-06, and G330. The common features are a density greater than 1.75 g/cm^3 , with particle size below 10 μm , and thus a pore size approximately below 1 μm . The results of these experiments, i.e., that 1 μm pores are tolerable but 10 μm pores are not, is in very good agreement with the modeled critical pore size above, and past work [137, 139]. Here, the tolerable porosity from experiments was slightly above the predicted critical value (a few μm instead of tenths of μm) but this was expected since the model conservatively assumed a worst-case deep cylindrical pore and quasi-steady pore closure while the actual pores are more tortuous and exhibit superlinear closure.




Figure 4.12: Various pressed graphite grades successfully contain silicon above 2000°C

4.5 The Effect of Silicon Purity

The price of silicon varies from \$1-1000/kg as purity increases from 96% to 9N [140, 141], so it was critical to determine the lowest purity silicon that could be used for TEGS. The major impurities in silicon, which is initially made by the carbothermal reduction of sand in an electric arc furnace, are other metals and residual oxygen: Fe (0.5-3%), Al (0.5-1%), Ca (0.3%), and O (0.1%) [141-143]. Notably, the cost of silicon can be reduced even below \$1/kg if it is intentionally diluted by Fe, which has a price as low as \$0.1/kg for scrap steel [16].

The effect of these impurities depends in part on whether stable compounds form between them, Si, or C at the temperatures, pressure, and gas environment present in the TEGS system. Each contaminant is first considered individually. As discussed in the section on Electrical Isolation in Solids, the most stable oxide above 2000°C is carbon monoxide (CO). CO is so stable at high temperatures because unlike to other oxides, it is a gas and thus has relatively high entropy due to the relatively high disorder of gases compared to condensed states. At high temperatures, entropy becomes a larger factor in the stability of a substance because it is

multiplied by temperature in the Gibbs free energy Equation 4.1. Thus, given the very low boiling point of oxygen and the presence of carbon in the solid and vapor phases, oxygen reacts and leaves the hot zone as gaseous CO.

Calcium boils at 1480°C but may form CaC₂ if exposed to carbon before boiling [144]. This carbide melts at 2100°C, so it would remain solid in the “cold” tank, while in the hot tank it would melt and likely deposit in the cold tank. This contaminant only forms silicides below 1300°C and does not form refractory aluminum intermetallics. Similarly, aluminum forms a carbide, Al₄C₃, which melts at 2180°C [145]. Thus, the aluminum that does not evaporate ($P_v=0.7\text{atm}$ at 2400°C [48]) is likely to collect in the cold tank as well. Another dilute contaminant in metallurgical grade silicon is phosphorus, which does not form a stable carbide. Its silicide, SiP, melts at 1160°C [146], and P itself boils at 280°C, so it will either deposit in the outer insulation around the system or remain dissolved in the silicon melt.

Iron is of special importance not only because it is the dominant impurity, but also because of interest in further increasing its concentration to reduce the overall storage cost. Fe is especially cost-effective because of its high volumetric energy density, which enables the entire TEGS storage system to shrink for the same amount of energy storage, reducing insulation and balance of plant costs [16]. Fe does not form refractory carbides (all melt below 1600°C), but carbon has great solubility in iron—up to 30% at 2400°C [147]. While this solubility is high and would surely inhibit containment if pure iron were used, if a SiC barrier forms before Fe significantly dissolves the carbon infrastructure, the further reaction may be drastically slowed by mass diffusion through this SiC layer. Iron diffuses through SiC about 100 times slower than carbon or silicon, at $\sim 10^{-16}$ m²/s compared to 10^{-14} m²/s for Si [148]. Iron forms several compounds with the other impurities, but they all melt below 1600°C [149]. In fact, iron is even used as a getter to clean silicon due to the relative stability of ferrous alloys and compounds [149], the effect does not occur here though, unless the cold tank is held below 1600°C to collect these compounds.

Therefore, once again liquid containment in the TEGS system relies on a protective SiC layer not only to prevent entropic dissolution into hot silicon, but even more so into hot iron. In fact, this effect can be visualized especially well from a ternary Fe-Si-C phase diagram. Although

higher temperature diagrams are not available, one at 1450°C [150] has been extrapolated to 2400°C based on the shape at 1450°C and the known intercepts based on binary phase diagrams [125, 126, 147]. This indicates that as the concentration of iron increases, the solubility of carbon in the melt increases, roughly linearly, from 5% for pure silicon to 30% for pure iron. Notably, 5% and 30% have a similar consequence for TEGS, in that they are far too high since the volume of the tank wall is only 2% the volume of the liquid, and the solubility is temperature dependent so if dissolution occurred, precipitation would be likely in the colder cold tank. This underscores the importance of SiC as a barrier to prevent the dissolution of the graphite infrastructure.

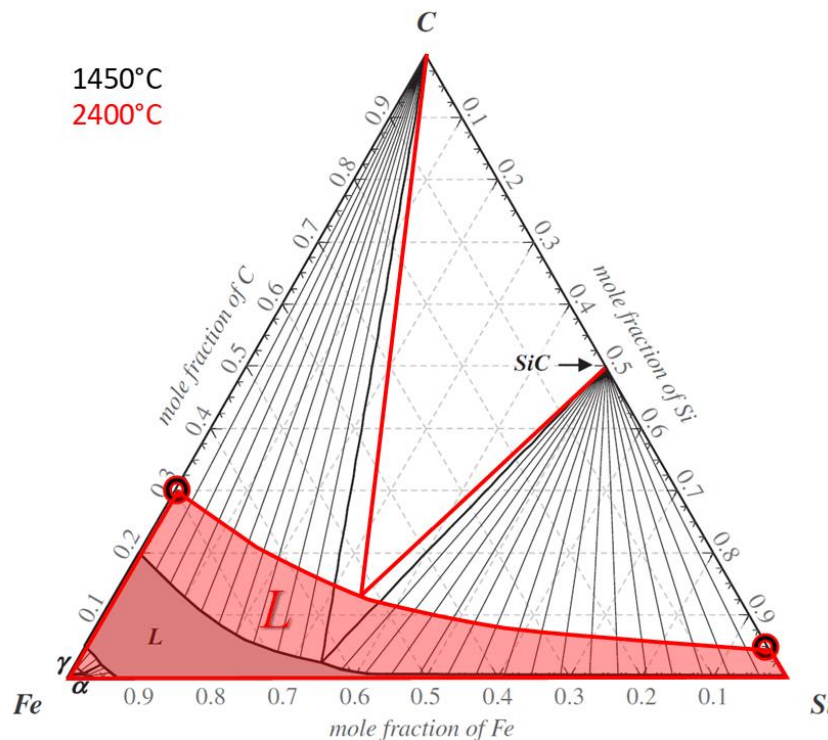


Figure 4.13: Fe-Si-C ternary phase diagram at 1450°C and approximate projection to 2400°C

Given these thermodynamic predictions, various alloys of silicon were tested to verify their effect. In most of these tests, KYM-20 graphite was used for the tank, which is an affordable, compatible type described in the section on the Effect of Graphite Pores. As a baseline, a few initial experiments were performed with 99.95% pure Si (see labels in Figure 4.12), which formed a stable SiC barrier as expected. Then an affordable grade of silicon known as 553, due to the impurity percentages (0.5% Fe, 0.5% Al, 0.3% Ca, 0.2% other for 98.5% purity), was used. This grade was selected because it is the least pure well-characterized grade of silicon. That is, while

some metallurgical grade silicon has 95-98% purity, the concentration of contaminants is inconsistent, so while they may be used for significant cost savings in the future, the main metallurgical grade used here is 553. Two of the experiments shown in Figure 4.12 from the section on the Effect of Graphite Pores used 553 silicon, and its impurities did not cause any noticeable degradation or change in the SiC layer compared to 99.95% pure silicon. This is a very important result because this grade of silicon is affordable, at \$1.6/kg or \$22/kWh-e on a capital cost per energy (CPE) stored basis. This is also the grade that is used for cost modeling in Chapter 2 on System Efficiency and Technoeconomics.

The purpose of the next set of purity tests was to determine the effect of higher concentrations of iron. The effect of adding iron has the most beneficial effect if at least 60-90%-mass (40-80%-at) of iron can be used because in this case the energy storage cost of the entire system is reduced by a factor of 2-5. In other words, if less than 50%-mass Fe can be used, the cost reduction is less than 50%, which while important is not as revolutionary. On the other hand, once 90%-mass is reached, the remaining savings to be gained become small, only about 5% based on this initial cost, or 15% on the new lower cost. Thus, the key interest was to see if $\sim 2/3$ iron could be used. From a materials perspective, %-at is the preferred unit, but since material costs are usually reported on a mass basis, %-mass also becomes an important unit.

According to the ternary phase diagram of Fe-Si-C in Figure 4.13, SiC is the only stable compound and it remains stable through 2700°C [127] so, theoretically, only enough silicon is needed to form a thin SiC layer, which would be less than 1%. However, in practice, the first layer of SiC must form quickly, before significant dissolution occurs, and that layer may delaminate partially or occasionally, and would need to be formed again. That is, while the graphite could potentially undergo initial processing where a SiC layer is formed without the use of iron, to retain the self-healing capability, it must have enough silicon to outrun the dissolution caused by iron. This issue could be especially important in the pump, where SiC may be continually worn away mechanically.

Given the uncertainty in just how much Si would be necessary, with bounds between 1% and 98.5% based on previous experiments and thermodynamic predictions, an array of concentrations was tested. The first test was with pure (99.98%) Fe to establish a baseline on the

other extreme. Due to the small size of these prototype tests, the volume of the tank wall is large enough compared to the fluid such that it becomes saturated before leaking. Nonetheless, the effect of dissolution is visible in the 99% Fe case in Figure 4.17. Here, although containment was technically achieved, it is clear that significant dissolution occurred, with more than 1 mm of wall thickness dissolved, and significant precipitated graphite and cementite (Fe_3C) visible in the SEM micrograph below. Thus, as predicted and previously reported, graphite dissolves until the iron is saturated if no barrier prevents it.

In later tests, it was observed that the way iron is added to the crucible was very important to the outcome. Three methods were used to add ferrosilicon. In the first approach, particles of solid iron and silicon were mixed, then added to a graphite tank, and melted. Here, significant dissolution was always observed, as shown in Figure 4.14, where an array of five concentrations was used. Another interesting result of this test demonstrated a minimum size for these tests, since the effect of solid-liquid-vapor surface tension was greater than gravity, causing non-uniform melting in these very small (~10 mm) tanks.




Figure 4.14: FeSi dissolves tank unless pre-melted to form homogenous alloy; %-mass. In each of these tests the FeSi dissolved the tank wall due to locally high concentrations of Fe in direct contact with C.

The second approach was to first melt 553 silicon in graphite crucibles to coat them with SiC, then remove the Si and add the appropriate alloy. The results of this were good, except that it was very cumbersome to fully coat a crucible due to the low packing density of particles. Here, the crucibles were filled, melted, filled, melted, then flipped to drain the silicon. Then, the SiC coated crucible was filled with an FeSi alloy and melted one last time, as shown in Figure 4.15. Finally, a third method was used where the premixed alloy was melted at a relatively low temperature in an alumina crucible and quenched to retain the dissolved mixture, as shown in Figure 4.16. This alloyed ingot was then heated in a graphite tank. This method of premixing had similar results to the previous method of precoating, so it was adopted as the main method. It is also important that this method works because it indicates that no preprocessing of the graphite would be required in the full-scale system, as long as the composition is uniform in the tank.




Figure 4.15: Graphite crucible precoated in SiC in preparation for FeSi containment testing.



Figure 4.16: FeSi ingots pre-melted in alumina, heated inside graphite for induction heating

The highest concentration of Fe found not to inhibit the growth of a SiC layer was 67%-mass (50%-atom). This suggests that drastic cost savings can be made by replacing two-thirds of the silicon with iron in the TEGS system, reducing the total capital cost per energy (CPE) by nearly a factor of three. This result is shown on the left of Figure 4.17, where silicon penetrated a short distance into the graphite, forming a well adhered SiC layer. This layer even remained during freezing, while the FeSi alloy shrank and separated from the SiC layer. This bodes very well for layer adhesion with thermal cycling. Notably, this effect is not observed for pure silicon tests because silicon expands on freezing.




Figure 4.17: Effect of Fe concentration on protective SiC barrier. An SEM image is shown below each image of the C-FeSi interface. The crucible graphite grade and concentration of iron (%-mass) is shown below each image. The remaining fluid balance was 553 silicon. Only test A with 67% Fe successfully formed a protective SiC layer, while the other tests with more iron did not show any SiC formation and/or significant graphite loss.

The middle two images in Figure 4.17 show that higher concentrations of iron prevented the formation of a SiC layer, although these results do not rule out the possibility of a concentration above 67%-mass Fe surviving under more ideal conditions. For example, precoating crucibles with SiC, or other changes to the melting process (e.g. thermal-temporal conditions) may enable the initial SiC layer to form before significant dissolution occurs. Another important caveat is that all these tests were largely spatially and temporally isothermal, two conditions that would be different at full-scale. That is, the temperature gradients in the heater and the heat engine may cause dissolution-precipitation cycles. Similarly, these sections of the system will experience transient thermal cycles through charging, storage, and discharging.

Additionally, pumping from the hot tank to the cold tank may cause mechanical wear enhanced dissolution. These issues are investigated in the next section and next chapter, respectively.

4.6 Thermal Cycling

While it has been shown in the previous sections that affordable silicon alloys can be contained in affordable graphite tanks up to 2300°C, it remained to be seen whether this is still true under transient or spatial thermal cycling. Transient thermal cycling involves changing the temperature of a component uniformly and this could cause delamination because of thermal or solid-solid phase change stresses that develop between the SiC and graphite. On the other hand, spatial thermal cycling refers to circulating silicon between two regions of different temperatures. This may accelerate dissolution, especially if SiC itself has significant temperature dependent solubility in silicon.

Spatial thermal cycling refers to the circulation of silicon between regions of graphite infrastructure at steady but different temperatures. For example, in the TEGS system, silicon is pumped from the cold tank, through the heater, and to the hot tank. Here, although the hot tank remains hot and the cold tank remains cold, silicon is pumped between the tanks at varying levels of carbon saturation. To check whether this varying carbon solubility resulted in dissolution of the graphite infrastructure, or if the SiC or other resistances mitigated the dissolution, the objective of this first thermal cycling experiment was to generate the ~500°C temperature gradient that will be present between the hot and cold tanks of the TEGS system. The reason this is a concern is that the solubility of C is temperature dependent such that thermodynamics shows up to 3.25% solubility of C in Si at 2400°C, but only 1% at 1900°C [129]. The geometry of this experiment is similar to previous small induction furnace tests described in this Chapter, except that this tank was much taller, 200 mm compared to 75 mm which enabled a significant temperature gradient within a single tank—especially when heated from the bottom with the top uninsulated. Then, via natural convection, the hotter silicon at the bottom of the tank, via buoyancy, circulated to the top of the tank. Then, if carbon dissolved in the hot zone, it may precipitate and deposit in the cold zone at the top.

A 3D COMSOL heat transfer, mass transfer, and fluid dynamics multiphysics model was made to inform the heating/insulation scheme required to generate the desired temperature profile, which is strongly influenced by the convective heat transfer in the silicon, which is driven by the temperature dependent density of silicon. Thus, the fluid mechanics and heat transfer physics were solved simultaneously. The model was also used to predict the flow rate between the hot and cold zones. In the actual experiment temperature is measured by a pyrometer focused on the tank lid as shown in Figure 4.18, so the model was adapted to match the temperature attained in the experiment, of 1070°C. Although this temperature is lower than the melting point of silicon, the lid is 50 mm above the silicon, and 200 mm from the heated zone. Here, the base of the tank was found to be 2100°C based on the insulation, external convective and radiative boundary conditions, and the measured and matched lid temperature. As shown on the right of Figure 4.18, the sharp thermal gradient generates buoyancy driven flow on the order of 30 mm/s or a 1 g/s (assuming 80% of the flow recirculates in smaller eddies).




Figure 4.18: Thermal and fluid model of dissolution experiment. On the left is a black graphite crucible containing silicon and surrounded by insulation. It is heated from the bottom inductively, causing the thermal gradient shown on the right. This thermal gradient induces natural convection which enhances the mass transfer, so that dissolution can be measured in a matter of hours.

If the convective mass transfer resistance were negligible, every 70 grams of silicon would carry a gram of carbon out of the hot zone, since the solubility in the hot zone is about 1.5% [129]. resulting in 1 mm of wall thickness loss in only ten minutes. However, convective mass transfer is, in fact, a significant resistance in this case. Conservatively modeling the flow as laminar pipe flow based on Hagen-Poiseuille flow and using an average ($5 \times 10^{-9} \text{m}^2/\text{s}$) of the closest diffusion coefficients available which range less than an order of magnitude—for carbon in molten iron [151], germanium in silicon [152], and other non-reactive liquids [153]—the predicted time to 1 mm wall loss is ten hours. This is conservative in that the actual mass transfer would be faster because of the complex and not fully developed flow profile, so if the experiment duration is on the order of several hours, it is expected there would be enough time for observable dissolution—if it indeed occurs based on the model above. This mass transfer model uses the mass-heat transfer analogy to calculate the effectiveness of the crucible as a mass exchanger using the ϵ -NTU method [154]. Due to the low surface area, diffusion coefficient, and residence time—the mass transfer is especially poor, and the effectiveness is about 0.5%. In other words, the mass transfer rate would be 200X if the silicon in the hot zone became saturated. So, considering saturated mass transfer and Hagen-Poiseuille convective mass transfer as upper and lower analytical limits, respectively, dissolution in the range of 10-600 minutes was expected for 1mm of thickness loss.

With this regime estimated, physics for mass transfer in dilute species was added to the thermofluids model from Figure 4.18. Here, the wall concentration of carbon was set based on the temperature dependent solubility of carbon in silicon, and the diffusion coefficient was also temperature dependent based on the average described above. This model predicts 1 mm thickness loss in 200 minutes based on the average rate of net dissolution from the hot zone over an hour long transient study, well in line with the range estimated analytically. The results of this model are shown in Figure 4.19, where the concentration of carbon is shown at various points in time. This diffusion is clearly enhanced by the naturally driven fluid flow as evidenced by the concentration being dominated by the natural convection cells.




Figure 4.19: Concentration of carbon in silicon for dissolution experiment. The snapshots in time show that large natural convection eddies carry away significant amounts of silicon. They also show that quasi-steady mass transfer has been established. The model predicts that the hottest region at the bottom will dissolve at 0.30 mm/hour.

Now that an estimate was in place for the rate of dissolution for this specific geometry and thermal conditions, an experiment was conducted to validate this model. Here, the 200 mm crucible was filled with silicon and first heated from the middle to melt all of the silicon, as shown in the first 2.5 hours of Figure 4.20. Then the heater was moved to the bottom of the tank to induce natural convection. Here, although the top of the tank decreased to only 1070°C, based on the COMSOL modelling this corresponds to a peak temperature of the bottom of the tank of 2100°C.




Figure 4.20: Temperature of spatial thermal cycling experiment over time. The tank was heated from the middle until the lid which was measured reach above the melting point of silicon (1414°C), as shown in the first 2.5 hours of the plot. Then the heating was moved to the bottom of the tank, inducing natural convection for the four hour duration of the experiment.

This experiment showed dissolution nearly as rapid as modelled, suggesting that the SiC layer did not prevent saturation of carbon in silicon near the surface of the SiC layer. That is, carbon from the SiC itself clearly dissolved into the silicon, since a measurable material loss was observed after the test of 0.82 ± 0.16 mm as shown in Figure 4.21. This uncertainty is based on the standard deviation of six measurements taken each of the dissolved bottom region and the top region that was not submerged. The variation in measurements is mainly due to the nonuniform dissolution in the hot zone, while the measurements themselves were made with 0.05 mm accuracy. This amount of material loss is in fairly good agreement with the COMSOL model that predicted 1.18 mm of material loss, 30% more dissolution that was observed.

Notably, the precipitated carbon formed highly porous SiC in the top of tank which interrupted the flow profile, slowing the rate of dissolution. Here, as the SiC built up, the effective height of the circulating region decreased. This caused the temperature range of the circulating region to reduce, perhaps from 500°C to about 250°C. This estimate is based on the ~50 mm height of SiC deposition visible in Figure 4.21 compared to the temperature profile in Figure 4.18. Since solubility is temperature dependent, this reduced the variation in solubility between the

hot and 'cold' zone. Then, since the concentration is determined by solubility, the concentration gradient was reduced. Mass transfer is driven by the concentration gradient, so the driving force—and thus the rate—of mass transfer was reduced. Similarly, the smaller temperature range also reduced the velocity of the fluid since the buoyant affects were suppressed by the smaller changes in density that accompany a small change in temperature. Thus, the driving force for diffusion and the advective enhancement of natural convection were both reduced as SiC continued to deposit in the top of the tank. This is likely the main reason for the overprediction of the COMSOL model.

Two other variables introduce uncertainty to the model. First, the actual temperature profile may be different than predicted due to gaps in the insulation or inaccuracies in the boundary conditions. For example, the vertical induction furnace generated external natural convection along its silica tube, causing accelerated and non-uniform convective cooling while a constant convection coefficient of $10 \text{ W/m}^2\text{K}$ was applied in the model. Also, there is significant uncertainty in the diffusion coefficient of carbon in silicon, since it is estimated based on values ranging almost an order of magnitude from different authors for carbon in molten iron [151], germanium in silicon [152], and other non-reactive liquids [153]. Although the effect of this uncertainty is suppressed by the advective enhancement compared to pure diffusion, the limits of the diffusion coefficient uncertainty change the model prediction by up to $\pm 0.09 \text{ mm}$ dissolution over the 4 hour experiment. Overall, these uncertainties combined with the effect of carbon deposition over time can account for the disagreement between the model and the experiment. Thus, there is a reasonably good agreement between the two, suggesting the physics driving the material loss are well understood. Essentially, carbon dissolves from SiC to locally saturate silicon near the SiC surface just as it does from graphite itself. In other words, even though SiC is protective against the reaction between carbon and silicon, it is not protective against the dissolution of carbon. Thus, in any molten silicon system contained by SiC or graphite above $\sim 1700^\circ\text{C}$ that is not isothermal, significant and rapid dissolution can be expected. This is a very important result because it implies that the TEGS energy storage system based on silicon above 2000°C may be infeasible due to dissolution of the graphite infrastructure in a matter of hours.




Figure 4.21: Results of four hour spatial thermal cycling experiment. Dissolution material loss of about 0.8mm in the hot 2100°C zone at the bottom was observed. This dissolved graphite was precipitated and formed porous SiC at the top of the liquid region. The darker particles in the middle and right figures are SiC.

Another form of thermal cycling also occurs in the TEGS system—transient thermal cycling. For example, between heating and storage phases the pipes in the heater section of the TEGS system will cool by 100-600°C. This could cause delamination because of thermal or solid-solid phase change stresses that develop between the SiC and graphite. The thermal stress that develops from the thermal expansion mismatch between SiC and graphite can be calculated since the thermal expansion of each material is known and the magnitude of thermal cycling is nominally $\sim 500^\circ\text{C}$. The thermal expansion coefficient of graphite [61] is relatively constant, at about $5 \mu\text{m/mK}$, while for SiC [155] it is $4 \mu\text{m/mK}$. Since graphite expands more than SiC, when the temperature is increased the SiC will be stressed in tension. Then, if it fractures and more SiC forms, it would be loaded in compression on cooling. Because the thickness of the graphite is

more than 1000 times thicker than the SiC, essentially all of the strain is expected to be resolved in the SiC, generating stress according to Equation 4.7. Here, σ is stress, E is the elastic modulus of SiC, ε is the strain, α is the thermal expansion coefficient, and ΔT is the temperature change. This stress exceeds the tensile strength of SiC along grain boundaries, although it is below the compressive strength along grain boundaries and the tensile strength of single-crystal SiC [155]. This suggests the SiC layer may initially crack on expansion along grain boundaries, but once these cracks are filled with more SiC, the layer may eventually be continuously loaded with only compressive loads.

$$\sigma = E\varepsilon = E(\alpha_c - \alpha_{SiC})\Delta T = (400 \cdot 10^9) \left((5-4)10^{-6} \right) (500) = 200 \text{MPa} \quad 4.7$$

Another potential driving force for the delamination of SiC during transient thermal cycling is the varying stability of SiC polytypes at different temperatures. Polytype refers to the crystal structure of a solid compound. Despite the hundreds of possible SiC polytypes [156], for the TEGS system the two most expected to be present are α -SiC (called 6H-SiC since it is hexagonal with six-layer periodicity) and β -SiC (3C-SiC; cubic, three-layer periodicity). These polytypes are very well studied because of their hardness, electrical properties, and stability [157]. SiC used for high-temperature applications is cubic because of its high-temperature stability [158]. However, above 1600-1900°C the hexagonal polytypes actually become more stable again. That is, cubic SiC is only stable between roughly 300-1700°C, while hexagonal phases are thermodynamically preferred above $\sim 1700^\circ\text{C}$ [159]. The reason for this transition goes back to Gibbs free energy: since the hexagonal configuration is less ordered, the transformation from cubic to hexagonal generates entropy, so as the temperature increases the driving force for transformation increases—as indicated by Equation 4.1. This change in stability around 1800°C could be a cause of delamination during thermal cycling because of the $\sim 1.5\%$ volume change accompanying the transition. However, it is not clear if the transformation will occur, or whether its effect will cause delamination when a thin SiC layer is bonded to a thick graphite substrate. While the thermal expansion of SiC depends on the present phase, the mismatch between phases is almost an order of magnitude smaller than the mismatch between SiC and graphite (β - α mismatch is 1-2%, but SiC-C is 10%) [160], so a representative value was selected in the thermal stress model above. Another interrelated phenomenon that could occur is the thermal stress

induced phase transformation of SiC. While this has been observed to only happen under GPa scale isostatic pressure [161], stress developed by the thermal expansion mismatch could contribute a phase transformation. Fundamentally, this is because when a compound is put under pressure, there is an increased driving force for the phase or polytype with higher density. This was an additional reason to investigate the effect of thermal cycling experimentally.

Due to the range of conditions and potential failure modes that could degrade the SiC layer, an experiment was devised to simulate/approximate potential issues in the full-scale system at the lab scale. For example, some potential issues may only become apparent after many thermal cycles. To address this, the experiment was designed to rapidly thermal cycle. Here, the existing experimental setup shown in Figure 4.3 was thermally cycled between 1600°C and 2100°C, a slightly different range than the planned storage cycle, but the widest range experimentally feasible such that the silicon did not freeze, and the peak temperature reached by the induction heater. Here, the number of thermal cycles was selected such that the degradation would be measurable. Since the SiC layer grows to ~20 µm thick in about 30 minutes, 25 cycles could cause 0.5 mm of wall thickness loss, which would be very obvious since the wall is only 3mm thick to start.

With this prediction, the transient thermal cycling experiment was conducted, varying the temperature from 1600-2100°C 22 times over 8 hours, as shown in Figure 4.22. The experiment stopped short of 25 thermal cycles because the optical path to the temperature measuring pyrometer became blocked near the end of the experiment, as indicated by the trailing peak temperature for the last three cycles.




Figure 4.22: Temperature of transient thermal cycling experiment over time. The temperature was ramped between 1600-2100°C 22 times over 8 hours to check for thermal stress induced SiC delamination.

The results of the experiment are shown in Figure 4.23. Here, material loss ranged from nearly zero, at the top and to bottom, to up to 3 mm (the entire wall thickness) near the middle. The remaining flakes of SiC above the SiC surface indicate that delamination did indeed occur, but it is difficult to fully decouple the delamination effects from the dissolution effects. It is also not completely clear why the material loss varies so greatly spatially. The most material loss occurred at the top surface of the silicon, and then the loss trailed off lower down the tank wall. This was roughly the hottest part of the crucible, although the heating was spread over at least 50 mm vertically. Thus, although this experiment proves transient thermal cycling *can* cause delamination, more analysis and experiments would be needed to determine under which specific conditions it occurs. For example, the rate of thermal cycling and the dwell time at the high and low temperature may affect whether delamination occurs. However, increasing the length of the experiment may exacerbate the effects of dissolution, further mixing the two effects.




Figure 4.23: Results of transient thermal cycling experiment. Part A shows a tank which was cycled only once for reference. Part B shows the tank that was cycled 22 times, resulting in up to 3mm (the entire wall thickness) material loss. Part C shows a cut perpendicular to Part B, where the entire wall thickness was lost as indicated by the arrow and detail view in Part D.

Chapter 5:

Pumping

The ability to pump is central to the TEGS-MPV concept because the liquid storage medium must be transferred to enable charging and discharging. However, the highest temperature commercial pumps operate only near 600°C [134], while this pump must operate at 2400°C, more than three times the absolute temperature. This new temperature regime introduces new challenges and uncertainties that are addressed in this chapter. Material selection has been limited to graphite by mechanical, thermal, and chemical constraints described in previous chapters, but now graphite is first used as a dynamic component and mechanically connected to low temperature metals and polymers in the motor. So, new challenges in thermal isolation despite mechanical connection arise, in addition to thermal stress from expansion with temperature. There are also critical chemical concerns given that the graphite pump could react with or dissolve in the silicon fluid. A key question is whether the formation of SiC will cause the shaft to seize or become rough, accelerating wear. On the other hand, SiC is known to be an excellent wear-resistant material [162], so if a smooth (or self-polishing) layer forms, it may actually enhance the life of the pump.

5.1 Introduction

In theory, the spatial and transient thermal cycling discussed in Chapter 4 on Liquid Containment would not be present in the pump because it is relatively isothermal. However, in this small scale experiment, the insulation and thermal isolation length scales in the pump are so short that the coldest (top) part of the pump is actually more than 100°C colder than the hottest silicon according to the thermal model described later in the Thermal section of this chapter. This temperature variation may be critically important if the SiC layer does not prevent the dissolution of the graphite since the change in solubility over 100°C in this temperature regime is about 1%. This potential for dissolution is modeled in the following Chemical section.

To address these uncertainties, a graphite pump was designed, modeled, and experimentally tested to circulate silicon. This builds on earlier work by the author [31, 32] where a ceramic pump for molten tin was designed and tested [32] that holds the world record for pumping temperature of 1400°C. Although the temperatures here are much higher, from a mechanical standpoint graphite behaves similarly between 1400°C and 2400°C, in fact, its mechanical strength actually increases with temperature up to 2600°C [61]. This experiment tests several potential failure modes together, including convection enhanced reactions, wear enhanced reactions, SiC binding the pump shaft to the pump body, and the thermal isolation challenge of keeping the electronics cool while in contact with a pump above 2000°C.

In the following sections, these risks and issues are addressed through design and modeling, followed by experiments where the pump performed as designed from a thermal and mechanical perspective, but suffered rapid material loss due to chemical dissolution which severely limited its life.

5.2 Mechanical

The hydrodynamic requirements of this high-temperature centrifugal sump pump are very similar to existing low pressure water pumps that have been very well studied [163]. Here, the efficiency at the lab and full scale is actually even less important than usual, due to the very high energy density of the fluid being pumped. The effect of pump efficiency on the cost and efficiency of the overall TEGS system is explored in Chapter 2 on System Efficiency and

Technoeconomics. For the prototype experiment, the pump was designed to operate at low speed so that harmonics and vibration issues could be decoupled from the primary interests—thermal and chemical performance. To operate at low speed with sufficient head, two parameters were altered from water pumps designed for similar pressure and flow rate: diameter and blade shape. Specifically, a large diameter of 125 mm was used, despite the low 1 gpm design point so that sufficient pressure (2 kPa) could be generated at low speed. The key change in blade shape was the use of forward, as opposed to backward, curved blades as shown in Figure 5.1. Forward curved blades cause the fluid to accelerate circumferentially in addition to radially, generating a higher absolute fluid speed compared to backward curved blades [164].




Figure 5.1: Impeller design (left) with backward curved blades to increase pressure at low speed. Velocity is shown on the right with streamlines indicating minimal flow separation.

The specific blade shape was selected based on empirically established design guidelines intended to avoid flow separation and maximize efficiency [165]. After varying the inlet and outlet blade angles in 2D ANSYS CFD model, it was determined that six blades with an inlet angle of 25° relative to radial and an outlet angle of 40° relative to tangential provided the highest head, with minimal flow separation, shown on the right of Figure 5.1. Here, the ANSYS model predicted a shut-off (no flow) pressure of 4kPa (equivalent to 0.2 m silicon head) at only 300 rpm, which is twice as high as a similar design using backward curved blades. Notably, even without considering losses in the pump outlet, this model predicted an efficiency of only 15%, which is largely due to the low speed—so low that viscous losses are not drowned out by the inertia imparted by centrifugal force. In the ideal inviscid limit, Bernoulli’s equation shown in Equation 5.1, describes the relationship between pressure and velocity across the pump. Here, B is a

constant, P is pressure, ρ is density, v is velocity which is the radius times the angular speed, h is head, and g is gravity. When solved for head, the equation takes the form of $h = (r\omega)^2 / 2g$, where ω is angular speed. For the geometry of this pump, Bernoulli's equation predicts a head of 0.19 m, in very good agreement with the ANSYS result. The reason the pump can outperform this Bernoulli calculation is that the blades are forward curved, such that the fluid velocity can exceed $r\omega$.

$$B = P + \frac{\rho v^2}{2} + \rho gh \quad 5.1$$

With the pump geometry designed, the next challenge was to experimentally verify its performance. Due to the high-temperature operation, and the fact that this pump is located inside a tank, no visual inspection was possible during testing. Furthermore, no commercial method of measuring pressure or flow rate could be used, since they would melt and react in this corrosive, extreme temperature environment. Instead, the pump was designed to have a constant head and the flow rate was measured with a simple mass based flow meter, conceptually shown in Figure 5.2.




Figure 5.2: Mass flow meter. Liquid is pumping into the top tank (flow meter) which drains at a rate proportional to the height of the liquid in it. So, as the pump flowrate increases, the height and weight of the top tank increases until the outlet flowrate equals the inlet flowrate.

Here, the fluid is pumped into a separate tank with a hole in the bottom. Then, as the flow rate increases, the pressure required to force the fluid out of the hole also increases. This pressure is generated by the height of the fluid in the tank since the pressure on all the external surfaces is atmospheric. Since the outlet hole is thin, viscous effects are minimal and this process can also be modeled using Bernoulli's equation from before in Equation 5.1. When solved for velocity, a relationship is found between the height of the fluid in the tank and the flowrate— $v = \sqrt{2gh}$. The variables in this equation can be exchanged for mass flow rate and the mass of the fluid in the tank, since mass flow rate is directly proportional to velocity and mass is directly proportional to height, yielding $\dot{m} = c\sqrt{m}$. With this approach, if the proportional constant and the mass of the fluid in the flow meter is known, the flow rate is also known.

This mass flow meter is an incredible tool where a static device and measurement can determine a dynamic variable—flow rate. It is this simplicity that was so useful in these high-temperature experiments. Here, the tank can be weighed to determine the mass flow rate. Now, given that the pipes and multiple tanks shown in Figure 5.2 would result in a large volume and heat loss, another innovation was made to simplify the system. Here, the pump and the flow meter were merged into a single component, as shown in Figure 5.3. Essentially, it is a pump with a tank on top of it. Since the pump and flow meter tank are integral, no external piping was needed. The constant head was generated by the use of an internal passageway up to the top of the tank, shown on the left of Figure 5.3. As the pump speed increased, the fluid would build up in the tank, since more pressure was needed to drive it out of the hole on the right. With this design, now only the pump, or equivalently the tank, needed to be weighed to know the flow rate. For practicality, in these experiments the tank was weighed as shown in Figure 5.4. Here, the tank rested on the insulation, which in turn rested on a thick stainless-steel plate with force sensors on each corner, outputting the total weight of the tank and the insulation on top of it. The overall weight of the system was approximately 50 kg, and the signal from pumping was about 1kg. This was easily measured using sensors designed for measuring postal packages, with a 10g resolution.




Figure 5.3: Flow path in combined pump / flow meter. The fluid is first accelerated to the outer diameter of the pump, then travels up the pump outlet path and into the tank (flowmeter) above the pump. The fluid level inside the flowmeter increases until the height is sufficient to drive the fluid out of the small outlet at the same flowrate as the flowmeter inlet from the pump.




Figure 5.4: The tank rests on the insulation, which rests on the scale. When the pump is turned on, the flowmeter above the pump fills, making the tank lighter. Then, the flowrate can be measured by measuring the weight of the tank.

This combined pump and flow meter performed very well and was first tested and calibrated by pumping water at room temperature. The calibration procedure was used to determine the constant of proportionality discussed above. This can be calculated for ideal flow, but is usually 30-50% lower due to minor losses [166]. Here, the pump is turned on to fill the flow meter and then turned off while it drains. The mass decrease over time during draining is shown on the left of Figure 5.5 and follows very closely to the parabolic rate predicted by inviscid flow. When this mass measurement is differentiated in time, the rate of mass change is found—which is equivalent to the mass flow rate. Once the mass and flow rate are known with respect to the same parameter i.e., time, they can be plotted against each other, as shown on the right of Figure 5.5. As shown, there is excellent agreement between the experimental measurement and the theoretical shape.




Figure 5.5: Water drain curves (left) used to calculate flow rate (right) from measured mass. The drain curves are measured by filling the mass flow meter from Figure 5.3 and then stopping the inlet flow. So, the weight is recorded over time while the flow meter drains as shown on the left. This curve is then numerically differentiated with respect to time to get the mass flow rate vs time. Lastly, the mass flow rate is plotted vs the mass as shown on the right.

Now that the pump and flow meter were designed, a system was made to mount and actuate the pump, which is shown conceptually in Figure 5.6. This design builds on the author's previous work [31] with important improvements. Most notable is the axisymmetric mounting of the pump. That is, the vertical rods that mechanically attached the pump to the same plate as the motor, preventing translation and rotation, are symmetric to the pump shaft. This is important because the pump expands with temperature, and any other type of mounting system would cause it to be misaligned with respect to the motor. These mounts also include springs in series with the hollow tubes that are compressed with CFC threaded rod to accommodate thermal expansion mismatch. Here, when the system is heated, the graphite tubes expand vertically up to 5 mm, while the CFC expands only 1 mm due to the very low thermal expansion of carbon fiber. So, without these springs very large strain would be developed—which actually happened in an early experiment causing mechanical failure. The importance of these springs is underscored in Figure 5.9 in the Thermal section.

The pump is actuated by an insulated shaft connected to the motor. The low thermal conductivity of this shaft is important because it is not actively cooled before directly contacting the polymer and aluminum motor. To further protect the motor, the plate which secures both it and the pump is water cooled, as further described in the Thermal section.




Figure 5.6: Mechanical design of high temperature pump. The pump is mounted asymmetrically to the motor with respect to the shaft to avoid thermal expansion induced shaft misalignment. The pump is connected to the motor with a low thermal conductivity ceramic shaft to thermally separate the hot and cold components. To accommodate the thermal expansion of the pump mounts relative to the internal bolts, springs are used in the cold region. The pump is located inside the tank so that a shaft seal is not needed since the small amount of fluid that leaks along the shaft remains in the tank. The motor is mounted to a water cooled plate to thermally protect it from the hot pump.

5.3 Thermal Considerations

The thermal design of the pump system builds on the work described in Chapter 3 on Joule Heating Input, where the heater and insulation materials and geometry were selected. For example, it was shown that 200 mm insulation thickness was needed for a tolerable heat loss of ~ 5 kW. This set the required length of the pump shaft and supports. A steady 3D thermal model of the pump system was made in COMSOL, as shown in Figure 5.7. In this model, the base of the tank is set to 2100°C and external convective and radiative boundary conditions are applied. Temperature dependent thermal conductivity is used for the insulation layers including the effect of the gas. The inner insulation layer is rigid graphite foam, also known as sliced rayon foam, which is surrounded by microporous insulation.

Although the temperature profile is similar to the heater system previously analyzed, the larger size of the system to accommodate the two-gallon silicon tank and the heat loss in the

pump shaft and supports increased the heat loss to nearly 8 kW. This heat loss could only be mitigated by significantly increasing the graphite insulation thickness, since increasing the thickness of the surrounding ceramic insulation alone would cause it to overheat. Another important result of this thermal model is that the pump is up to 100°C colder than the silicon. While pumping helps to make this temperature more uniform, this nonetheless underscores the fact that the pump is not isothermal.




Figure 5.7: Temperature distribution of pump system (left) and pump (right). The model uses a fixed temperature input to the base of the tank, based on the heater temperature of 2100°C. External convection and radiation boundary conditions are applied, and the motor and pump mounts are actively water cooled. The figure on the right shows that temperature varies by 100°C within the pump.

The pump and motor both mount to the same water cooled stainless steel plate, as shown in Figure 5.8. This design incorporates several heat transfer techniques to protect low temperature components while keeping the system simple and minimizing the risk of coolant leaks. The most difficult component to thermally protect is the motor since it is mechanically attached to the pump. Two main approaches are used to achieve this: thermal isolation and geometrically isolating active cooling. The thermal isolation is mainly employed by the long pump shaft, a section of which is made from partially stabilized zirconia, with a relatively very low thermal conductivity of 2 W/mK. The pump mounts and the motor are also geometrically isolated by locating water cooled channels between them. Through this arrangement, the coolant water

essentially breaks the heat transfer circuit between the pump mounts and the motor, even with large (12 mm) channels and moderate heat transfer coefficients ($\sim 2000 \text{ W/mK}$) as shown in Figure 5.8.




Figure 5.8: Temperature distribution of pump motor from side (left) and top (right) view. The motor is kept cold by the water cooling which not only cools the motor/pump mount plate, but also geometrically separates the hot pump mounts from the cold motor—effectively breaking the thermal circuit and protecting the motor from overheating.

The pump mounts are installed mechanically in series with springs to accommodate the relatively large expansion rate of graphite compared to the CFC bolts securing them. This net thermal expansion is explored in detail in the section on Sealing and can be calculated using Equation 4.7 resulting in 4 mm, which would cause a stress of 450 MPa—far exceeding the 120 MPa tensile strength of the CFC—if not relieved by the springs. On the other hand, the springs have a force constant of 15 N/mm and are installed with a preloaded displacement of 5 mm for a maximum displacement of 10 mm and 150 N maximum force. Taking the active cross-section area of the CFC threaded rod into account, this amounts to a tensile stress of only 2 MPa, well below the fracture strength. Figure 5.8 shows how drastically these springs are compressed when the pump is heated.




Figure 5.9: Thermal expansion of pump supports relative to internal CFC thread rod. The graphite and ceramic pump supports expanded 5 mm vertically, while the CFC thread rod expanded less than 1 mm. The springs at the top accommodated this mismatch, visibly compressing 4 mm.

With all these constraints and considerations in mind, the full pump system was designed. This system builds on the conceptual system shown in Figure 5.6, by adding vertical actuators, weight sensors, and the heater as shown in Figure 5.10. The actuators are needed because the silicon is frozen during assembly, so the pump is held above the silicon and then lowered once it melts. These actuators are known as non-captive linear actuators and are powered by stepper motors which rotate a nut to climb up and down a fixed lead screw. The weight sensors shown suspend the tank and insulation above the metal pump infrastructure and are used to measure the flow rate of the silicon as described in the Mechanical section. Here, it was geometrically simpler to weigh the tank instead of the pump, so the measurement was simply the negative of if the pump was weighed. These sensors are aluminum bars with strain gauges attached and are commonly used in shipping for weighing packages. The heater is also shown beneath the tank, and the tank is supported by 4 rods on the corners (not shown) so that it remains at least 20 mm away from the heater to avoid short circuiting by arc discharge, as described in the section on Electrical Isolation in gases in Chapter 3 on Joule Heating Input.




Figure 5.10: Full prototype pump system. The pump is lowered into the tank once the silicon is melted by linear actuators mounted to the motor plate. The tank is heated from below and is supported by the insulation which rests on a scale. The scale reports the weight change of the tank, and by extension, the weight of the flow meter above the pump, enabling the calculation of flow rate.

5.4 Chemical Considerations

It was shown in Chapter 4 on Liquid Containment that SiC can form a protective layer on graphite greatly slowing further reaction under isothermal conditions, but it was unclear how this layer would behave between dynamic surfaces and under advective enhancement of any

degradation mechanism. Therefore, in these pump tests, the key goal was to learn the effect of pumping on the SiC layer. Two main failure modes were anticipated: seizure and wear.

SiC can be formed by silicon penetration in the gas or liquid phase, as shown in Figure 5.11. This was observed in the first large tank test, where silicon formed on the lid of the tank from condensed silicon vapor in addition to the sides of the tank that were submerged, as described in Chapter 4 on Liquid Containment. However, the vapor process is much slower than the liquid reaction, especially here, because of three factors: the low concentration of the vapor in the gas, the narrow path compared to the surface area of the shaft/pump gap to be coated, and the lack of capillary pressure to drive silicon into the pores. The time required for a 10 μm thick layer of SiC to form on the pump shaft can be estimated based on the length and cross section of the path silicon vapor must travel. Following the same analysis method for the coating of the tank lid described in Chapter 4 on Liquid Containment, it would take nearly one year to form a protective layer on the pump shaft from silicon vapor. This is in stark contrast to the layer growth for graphite submerged in molten silicon which forms a 10 μm layer in two hours, since the reactive source is fully concentrated at both sides of the SiC layer. This means that the pump shaft will not be coated with SiC until it is submerged, an important fact in the initial operation of the pump before SiC has formed.




Figure 5.11: Vapor (left) and liquid (right) silicon source for SiC layer formation. The path for vapor diffusion is so long and narrow that a SiC layer would take days to form compared to only minutes when submerged as shown on the right.

With it established that the SiC layer will form only when liquid reaches a given location, another question emerged regarding the rate of formation. Here, the concern was that if the layer formed rapidly, too much SiC would simultaneously bridge the gap between the shaft and the pump body (which can fluctuate from 0-200 μm as the shaft moves) and the shaft would seize. Since the SiC formation is fast, the rate of formation can only be limited by controlling the rate at which silicon is put in contact with the graphite. Here, the pump was slowly lowered at 2mm/minute into the silicon using the actuators shown in Figure 5.10, slowly submerging the bottom half of the pump.

The pump was therefore lowered slowly into the silicon in all experiments, but in the first test the sleeve bearing (i.e. the static part of the pump surrounding the shaft) was too long to be fully submerged. During this test, the pump was slowly lowered to the bottom of the tank while rotating slowly, over the course of 30 minutes. The torque required to rotate that shaft remained roughly constant during this time, but once the pump speed was increased to overcome the head to actually start pumping, the shaft seized very quickly after only a few minutes. While this caused the test to fail, this test turned out to be an excellent validation of the theory that slow immersion was necessary. As shown in Figure 5.12, at the bottom of the shaft where immersion was slow, the surfaces remained smooth and separated by silicon as desired. On the other hand, near the top of the shaft where silicon was rapidly introduced by the increase in pressure from the change in pump speed, the surfaces are very rough and the SiC is thick and even bridges the gap in some areas. In fact, the roughness exceeded the average gap width, such that there was interference to the rotation, which caused the pump to seize. Although this issue could be resolved by changing the pump speed slowly, the sleeve bearing was simply made shorter in future experiments to avoid the issue altogether.




Figure 5.12: Dynamic SiC surfaces contact at the top of the shaft, causing binding which stopped the pump from rotating. The thick, rough SiC formation was caused by the rapid introduction of silicon to the top of the sleeve bearing when the pump speed was increased in a step change.

As shown in Figure 5.13, future experiments used a shorter sleeve bearing that was fully submerged when the pump was lowered to the bottom of the tank. This resulted in a smooth SiC layer even at the top of the sleeve as shown on the right of the figure. The SiC is much smoother here than in the earlier tank experiments, which suggests the layer is indeed self-polishing. That is, large or poorly attached SiC particles were likely worn away, leaving a surface with smaller asperities [162]. In these microscopy images, the cracks are not an issue, but an artifact of the shrinking of the system on cooling. Notably, even though silicon expands on freezing, the dimensions here are dominated by the graphite components since the gap is small compared to the shaft diameter.




Figure 5.13: SiC shaft surfaces separated by silicon. The shaft did not bind in this case because the SiC formed slowly and smoothly since the pump was slowly lowered into the silicon.

With the mechanical, thermal, and SiC formation matters addressed, the pump was almost ready for long-duration testing. However, given the results of the pump thermal model in Figure 5.7, which show up to 100°C temperature variation between the pump surfaces and the silicon, the spatial thermal gradient driven dissolution analysis conducted in the section on Thermal Cycling in Chapter 4 on Liquid Containment was applied to the pump as well. Here, the concern was that if the SiC layer did not prevent the dissolution of carbon from the pump, it may dissolve away during the experiment. Although the temperature variation is relatively small in the pump, the solubility of carbon in silicon increases by about 1% per 100°C in the 2000-2400°C range according to the phase diagram in Figure 4.1. When this fact is coupled with the very small diffusion length scales and high velocities in the small gaps in the pump, the pump can actually dissolve very quickly if the SiC does not prevent it.

To estimate this rate, the method used in the Thermal Cycling section was adjusted to model the gap between the pump impeller blade and the surface beneath it, which is a gap 0.1 mm thick and 50 mm wide. This model considered Poiseuille flow in this wide rectangular gap where the recently heated but potentially unsaturated bulk silicon has a concentration of carbon 1% lower than the pump surfaces, where the concentration was governed by the solubility at silicon's temperature. This model predicts 1 mm of material loss from the exposed surfaces in

only four hours, which would drastically reduce the performance of the pump and cause deposition of the carbon elsewhere in the system. Notably, in a larger system, the pump could be held in closer thermal equilibrium with the silicon and this rate could be reduced. For example, if the temperature difference was only 10°C, the concentration gradient would also be reduced by approximately an order of magnitude, corresponding to an order of magnitude increase in pump life—since mass diffusion and advection (just like conduction and convection) depend linearly on the concentration (or temperature) gradient. This concept can be extrapolated to predict the maximum temperature variation to allow the pump to last for one year, yielding 0.05°C—an incredibly stringent requirement. This further underscores the importance of the SiC layer to not only slow the reaction of carbon with silicon but also prevent the dissolution of carbon.

5.5 Results

The pump operated for up to 10 hours and at temperatures as high as 2000°C during three experiments. In two experiments, the flow rate was measured which is shown in Figure 5.14. As described in the Mechanical section, by measuring the mass of the tank over time, the flow rate could be determined. This relationship between the pump speed setting and measured mass was the primary way the performance of the pump was confirmed in these tests. The temperature and pumping histories of these experiments are shown in Figure 5.15, where the duration of pumping and phase change (melting/freezing) are shown with horizontal lines at the top of the figure. In all three of these tests, the pump did not fail mechanically, nor did any components overheat or suffer excessive thermal stress. In the 1-hour test, no discernable freezing plateau exists because most of the silicon leaked out of the tank during that test, likely due to thermal stress developed in the bolts, as discussed in the Mechanical section. The tank did not leak in the other two tests.




Figure 5.14: Silicon drain curves (left) stacked (middle) and used to calculate flow rate (right). The drain curves are measured by filling the mass flow meter from Figure 5.3 and then stopping the inlet flow. So, the weight is recorded over time while the flow meter drains as shown on the left. This curve is then numerically differentiated with respect to time to get the mass flow rate vs time. Lastly, the mass flow rate is plotted vs the mass as shown on the right.

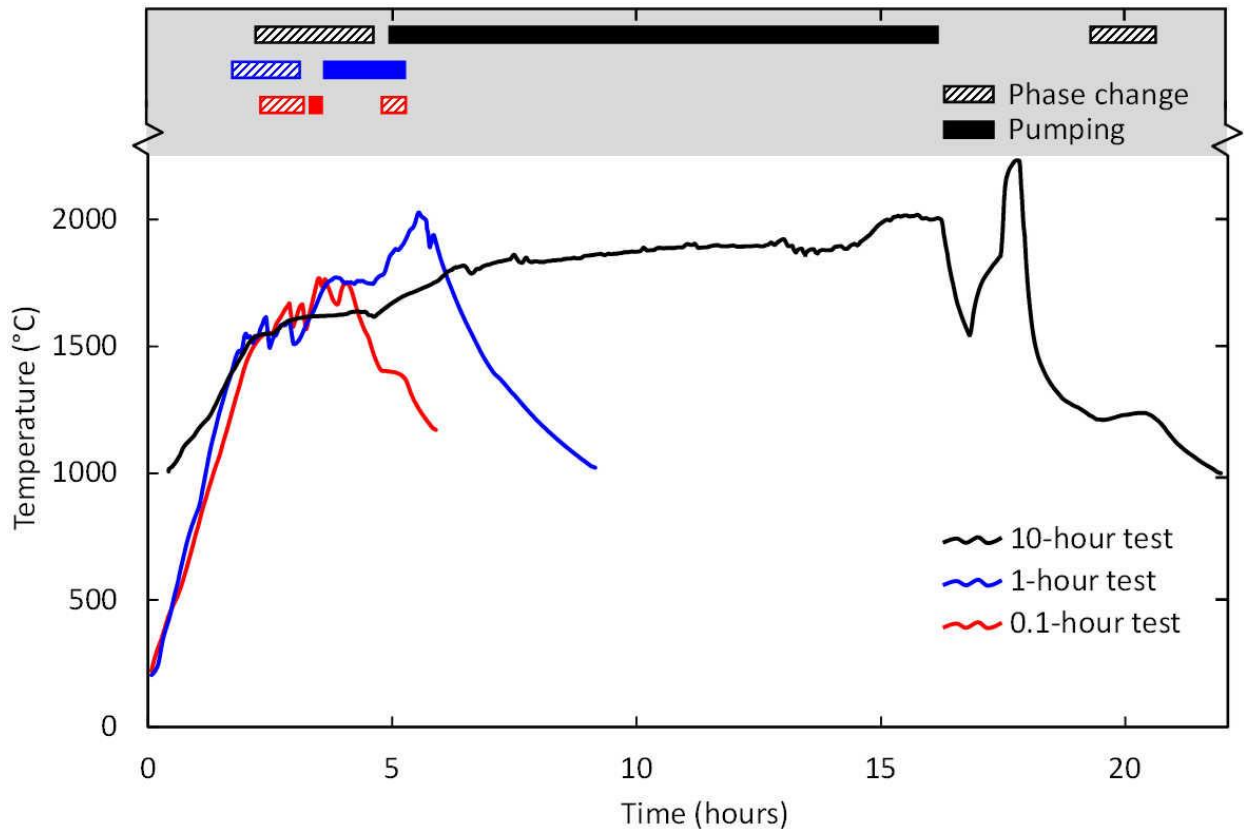


Figure 5.15: Pumping duration and temperature in three experiments. The top shaded region shows the different phases of the experiment: melting, pumping, and freezing. In the main plot, a horizontal temperature plateau is visible during phase change due to the high latent heat of silicon. The 0.1 hour test was stopped when the pump shaft became bound to the pump by SiC. The 1 hour test ended when the tank leaked. The 10 hour test was stopped because there was no flow rate signal. Images of the pump cross sections after each test are shown in Figure 5.18.

However, during the 10-hour experiment, there was no indication of flow after the first few minutes of the pump rotating. The pump continued to rotate throughout the test as observed visually and by the shaft encoder measuring its speed, and there was no mechanical fracture of the pump during the entire test. After ten hours of rotating, though, the torque required gradually increased until the shaft seized. The cause of this lack of mass signal and shaft seizure was clear after the experiment—the pump outlet and shaft were covered in SiC. The source of this SiC was also clear because of the large amount of missing material from the pump components, as shown in Figure 5.16. That is, carbon from the hottest parts of the pump, including the impeller and surrounding pump body, had dissolved into the silicon and precipitated out in the colder areas, including the pump outlet, the top of the pump, and the top of the tank. Notably, even if the flow path was not clogged, so much SiC was generated (~500 ml of highly porous SiC) that it even bridged the 10 mm gap between the pump and the tank so that neither could be weighed independently.




Figure 5.16: Part A shows the pump cross section, which suffered major dissolution after 10 hours of pumping. The regions closest to the pump entrance at the bottom center experienced the most thickness loss. Part B shows the pump outlet channel that is blocked by porous SiC and Part C shows the outer surfaces of the pump which are also covered in SiC. Part D shows the bottom of the pump, where macroscopic (~1mm) porosity has developed from the dissolution of carbon. Part E shows the thick SiC present is poorly adhered, and was deposited on top of silicon, rather than forming at the surface because when the SiC scale was chipped off, revealing the silicon beneath.

Prior modeling in this and previous chapters predicted that SiC protection was necessary to prevent excessive carbon dissolution, and the rate of dissolution is indeed similar to that which was predicted in the case that SiC was not protective. For example, in the Chemical section above, the impeller surfaces were estimated to dissolve at a maximum rate of 0.25 mm per hour of pumping, and in this experiment where the pump was rotated for 10 hours, approximately 1 mm of material loss was observed, even where there was no dynamic solid-solid contact (for example, the center of the pump body below the impeller in Figure 5.16A). That this experimental rate is lower may be due to the blockage of the pump outlet, such that the silicon inside the pump became saturated, slowing the dissolution process. In the experiments specifically investigating

the effect of Thermal Cycling in Chapter 4 on Liquid Containment, a more controlled comparison was able to be made.

Given that the SiC layer has been shown not to prevent dissolution, and the approximate agreement with the mass diffusion model, the critical combination of solubility and liquid diffusion coefficient can be estimated to inform the design of alternative embodiments of the TEGS system, or related systems. In an energy storage system with significant solubility variation with temperature, the first component to fail may be the pump or the pipes where heat transfer occurs (e.g., the heater and heat engine). In the pump, velocities are high and diffusion length scales are low, while in the regions of heat exchange the temperature variation is large. The three variables concentration, diffusion coefficient, and length scale are the key factors present in Fick's law, Equation 4.3, which governs diffusive mass transfer. This equation has been solved for the rate of material loss (or deposition) in thickness per time units and is reformatted below in Equation 5.2 to emphasize these variables. Convective effects are considered by adding in Sh , the Sherwood number, which analogous to the Nusselt number for heat transfer, and accounts for the enhancement of convection such that the equation reduces to solely diffusive mass transfer if $Sh = 1$.

$$i \approx Sh \cdot D_{C-M} \frac{(m_s - m_b)}{L} \quad 5.2$$

Here, i is the rate of wall thickness loss (e.g. of a pipe wall or the impeller) in m/s, D_{C-M} is the liquid diffusion coefficient of carbon in the melt, m_s is the mass fraction concentration set by the solubility of carbon in the melt, which is essentially equal to the wall concentration [167]. The concentration of carbon in the bulk fluid is denoted by m_b , which can be nearly zero if a cold sink is in the circulation loop. The length scale for diffusion is denoted by L , for example the diameter of a pipe. Although other schemes may be employed to prolong life, a general limit on the thickness loss could be on the order of 1 mm per year for an industrial-scale system, considering pipe wall thicknesses and impeller blades are generally on the order of a few mm. This equation cannot be further simplified in general, but by inserting typical values for all but the solubility, a rough estimate can be found for the targeted solubility, which can be a starting point in the search for materials and temperature limits. For this example, the diffusion

coefficient in many high-temperature liquids is on the order of 10^{-8} m²/s is used [151-153]. The Sherwood number for laminar internal pipe flow in the limit of dilute mass transfer is 48/11 based on the analogous Nusselt number, and a pipe size of 50 mm is common in large energy systems, so these values are used to estimate the solubility limit. This yields a maximum solubility mass fraction of 3×10^{-5} , which is roughly similar to the mole fraction, depending on the molar mass and density of the liquid selected relative to carbon. While this is more than three orders of magnitude lower than the solubility of carbon in silicon at 2400°C, some materials approach this low value, especially at moderately lower temperatures. One very interesting candidate is tin (Sn), which is considered as an alternative in the section on Future Work.

The series of experiments presented above are shown sequentially in Figure 5.17, where the time dependent degradation becomes very clear. For example, the vertical gap above the impeller is much larger in the 10-hour test than in the shorter duration tests. Another notable result in that the 1-hour test the pump outlet labeled on the left of the figure is partially clogged with SiC and SiC has also deposited at various levels on the inside of the flow meter. The fact that this deposition/precipitation is not present in the shorter test, and far more is present in the longer test clearly represents a time dependent degradation mechanism that turned out to cause the pump to fail after only hours of use. In other words, without the protection of SiC, this chemical situation is by analogy actually quite similar to a pump made of sugar pumping water. That is, the solubility of sucrose in water at 20°C is within a factor of two of the solubility of carbon in silicon at 2400°C [168].




Figure 5.17: The pump dissolves over time. These images are from three separate experiments of varying length, but otherwise similar conditions. The buildup of precipitated SiC can be seen after one hour, but it is much more obvious after 10 hours, where the entire pump outlet is clogged.

By applying all of the design methodology and modelling in the previous chapter and sections, the pump performed as expected from a thermal, fluids, and mechanical standpoint but nonetheless its life was severely limited by chemical dissolution. While the aspirational goal of these experiments was to demonstrate pumping an affordable storage material above 2000°C without degradation, instead what was demonstrated is that the key principle that made it plausible—the protection provided by a SiC layer—does not prevent dissolution when a thermal gradient is present, despite its ability to greatly slow the rate of reaction under isothermal conditions.

Chapter 6:

Conclusions and Future Directions

6.1 Thesis Contributions

The most central result of this thesis is that of all the challenges to the use of silicon to store electricity, only one stands out as unsurmountable—dissolution. That is, above 1600°C the solubility of carbon in silicon becomes appreciable, including the carbon in an otherwise protective silicon carbide layer. Several other important results were also generated as part of this thesis work that can inform the design of related thermal storage concepts, and even other forms of energy storage more broadly based on the economic concepts considered.

In Chapter 2, a new economic framework was developed for evaluating and comparing storage technologies. Central to this method is the consideration of roundtrip efficiency as a tunable variable in the development of an energy storage technology. This is important because it shows that efficiencies as low as 30% may be profitable if accompanied by a corresponding decrease in cost compared to existing technologies. A technoeconomic evaluation of the full scale TEGS-MPV concept was also conducted, which not only resulted in an estimate of the system cost, but also detailed designs of each system component. Here, the cost of extreme temperature insulation, inert containment, and the graphite infrastructure was found to be affordable using existing materials and methods.

In Chapter 3, the challenges to efficient and affordable joule heating to store electricity were evaluated and addressed. Cost effective and efficient power conditioning can be achieved through the use of existing thyristor based controls. However, avoiding electrical short circuit

through the materials supporting the heater and even the noble gas environment is a significant challenge. Several methods were employed to reduce or avoid this degradation, and in the end three methods were employed which enabled heating at four times (80V vs 20V) the industrial limit of heating above 2000°C. These included optimizing the heater geometry to minimize the electric field gradient, grounding the center of the heater to reduce the peak voltage relative the insulation by a factor of two, and the use of helium as opposed to argon to suppress arcing.

In Chapter 4, the conditions under which silicon can be contained were evaluated. It was found that 553 metallurgical grade silicon, which is affordable at \$1.60/kg (\$22/kWh-e for 500°C storage temperature range), can be contained in affordable isostatically pressed graphite grades which are as inexpensive as \$3/kg. In fact, as much as two thirds iron can be added to silicon by mass without degrading the protective silicon carbide layer that silicon forms on the graphite tank. This dilution of silicon can reduce the total capital cost per energy (CPE) of the TEGS-MPV system by 65%. The main property of commercially available graphite grades that determines whether it can be used to contain silicon is pore size. A simple model was developed and validated experimentally that can predict the critical pore size within an order of magnitude. This is very useful because pore size ranges four orders of magnitude for graphite, and this result showed that the pores of extruded graphite are too large but those of essentially all pressed grades are small enough.

The effect of thermal cycling was also evaluated in Chapter 4. Here, two types of temperature variation were considered: spatial and transient. In the TEGS-MPV system, the graphite tanks remain at a largely constant temperature for the life of the system, but the silicon may enter these tanks with more or less than the equilibrium concentration of dissolved carbon. Furthermore, in the sections of the system where heat is exchanged—namely the heater and heat engine—the temperature profile varies spatially and temporally. For spatial thermal cycling, an experiment was designed to measure the rate of dissolution and compared to a COMSOL 3D heat transfer, mass transfer, and fluid dynamics multiphysics model. The model predicted 1.18 mm of material loss during the four hour experiment and 0.82 ± 0.16 mm was observed. While the experiment material loss in the experiment was 30% less than predicted, this can be explained by the large amount of dissolved carbon precipitating as silicon carbide and partially

blocking the flow path, reducing the rate of further dissolution. Thus, overall, this experiment suggests the underlying mechanisms for carbon dissolution from graphite coated in SiC is well understood, and that it proceeds at a rapid rate above about 1700°C.

The second thermal cycling experiment involved heating a small tank from 1600°C to 2100°C repeatedly, to check for thermal stress induced delamination. The predicted rate of delamination if it occurred with each cycle was 20µm/cycle since that is approximately the thickness of the SiC layer. However, in the experiment, rates ranging from 20-150µm/cycle were observed. The highest rates were likely enabled by the delamination occurring deeper in the weaker graphite wall, such that the SiC removed significant amounts of graphite with each delamination. Thus, these experiments show that delamination is *possible*, but more work is warranted to investigate how it might be avoided. For example, reducing the rate of temperature change, or the high and low temperatures, may affect the delamination.

In Chapter 5, the effect of pumping was investigated on the graphite infrastructure. Here, the pump eventually performed as designed and modeled from a thermal, fluid, and mechanical perspective at temperatures up to 2000°C and 10 hours. However, since the pump experienced a significant thermal gradient, it was also subject to major dissolution that was also observed in the spatial thermal cycling experiment. So, although a pump could theoretically survive in a truly isothermal environment, the practical gradients present at the prototyping scale lead to about 1 mm of material loss in just a few hours. It also turns out that the more rapid degradation mechanism is the deposition of this carbon as highly porous silicon carbide in the pump outlet, which quickly clogs the outlet piping.

6.2 State of TEGS-Si

The results of this work show that a major risk in the use of silicon significantly above its melting point has been confirmed, due to the high solubility of graphite containment in the liquid silicon. In fact, since silicon was investigated for its unique position as an affordable, potentially containable liquid, these results suggest that the entire idea of using a single liquid for sensible energy storage and heat transfer above 1500°C may be impractical. However, the performance benefits of *moving* heat with a liquid are so significant, further work is warranted on the

containment and pumping of liquids above the recently set 1400°C world record [32]. Similarly, the cost and performance benefits of operating a photovoltaic heat engine at extreme temperatures near 2000°C warrant the investigation of alternative extreme temperature storage materials. In the following section, a simple alternative embodiment of TEGS is introduced to motivate future work, along with some unresolved issues with the TEGS-MPV system that would translate to this new embodiment and similar alternatives.

As described in Chapters 1 and 2, the use of a liquid to move heat between storage, charging, and discharging components of a storage system is very important since the requirements in each component are different, and liquids (unlike solids) deform continuously to meet the shape of their tanks, pipes and heat exchanges. Similarly, liquids can have very high volumetric energy density (unlike gases), enabling a compact efficient heat transfer system.

In the same way, given that the renewable energy problem has become largely an energy storage problem, it remains critical to find a low cost way to buffer intermittent renewables. Since there is no clear path to affordable long duration grid storage from batteries or otherwise [17], it remains valuable to develop thermal technology which can be fundamentally low cost, if it works. High temperature thermal systems have a special possibility to be inexpensive and relatively efficient. These systems can store energy over a wide temperature range, potentially more than 1000°C, which could make them more energy dense than existing oil and salt based storage systems. Their efficiency can be higher based on Carnot efficiency, and the heat can be converted back to electricity cheaply, without the use of turbomachinery, only if the temperature is high enough to enable heat fluxes that drown out much of the cost and losses of a photovoltaic heat engine.

6.3 Future Work

An alternative embodiment of TEGS which merits further consideration is the separation of the storage and heat transfer systems. One specific material pair that appears ideal is graphite storage and tin heat transfer. Graphite is actually a much cheaper storage material than silicon because low grade graphite can cost as little as \$0.50/kg, and its mass based specific heat

continues to increase up to 2 kJ/kg/K at 2000°C, approximately twice as high as silicon. However, the challenge becomes, how can thermal energy be exchanged with the graphite.




Figure 6.1: TEGS-Sn concept where energy is stored in solid graphite and heat is transferred with liquid tin. During charging the left side of the piping is activated and tin transfers heat from the heater to the graphite. During discharging the right side of the piping is activated and tin transfers heat from the graphite to the MPV.

Liquid tin could be used to transfer thermal energy from the heater to the graphite, and from the graphite to the heat engine as shown in Figure 6.1. However, the question arises, will tin have the same solubility problem as silicon? Thermally enhanced solubility is fundamental to all materials [130], and always becomes appreciable as the solid approaches its melting or dissociation temperature, as described by Equation 4.2 in the section on Silicon Carbide [130]. Graphite is actually very far from decomposing at 2400°C, unlike SiC which decomposes at 2700°C. In fact, graphite does not melt at atmospheric pressure, and only sublimates at 3600°C. Since graphite does not react with tin, the only driving force for dissolution is mixing entropy, and it has been measured at up to 2300°C [169]. At this temperature, carbon has 0.035% at solubility in tin. While this is two orders of magnitude lower than silicon, it is still not negligible.

Nonetheless, this much lower solubility shifts the system to a regime where mitigation is far more practical.

As explored in the Results section Chapter 5, the dissolution can be reduced by reducing temperature, convective effects, the solubility variation, and the length scale for diffusion. In addition to these four tools, a regenerative, or reversible dissolution-precipitation scheme could also be employed. In this case, the direction of fluid flow and or thermal/solubility gradient could be periodically reversed in an attempt to redeposit lost material in the same location. A simple diagram showing how this could work is given in Figure 6.2. Here, the direction of flow is reversed axially, and the direction of heat transfer is reversed radially in order to reverse the effects of dissolution locally. Lastly, a refractory coating could be applied to reduce dissolution. This would be similar to the idea of SiC protecting graphite, but an even more refractory coating such as WC or ZrC which would have lower solubility due to its higher dissociation temperature (due fundamentally to the stronger covalent bonds compared to SiC).

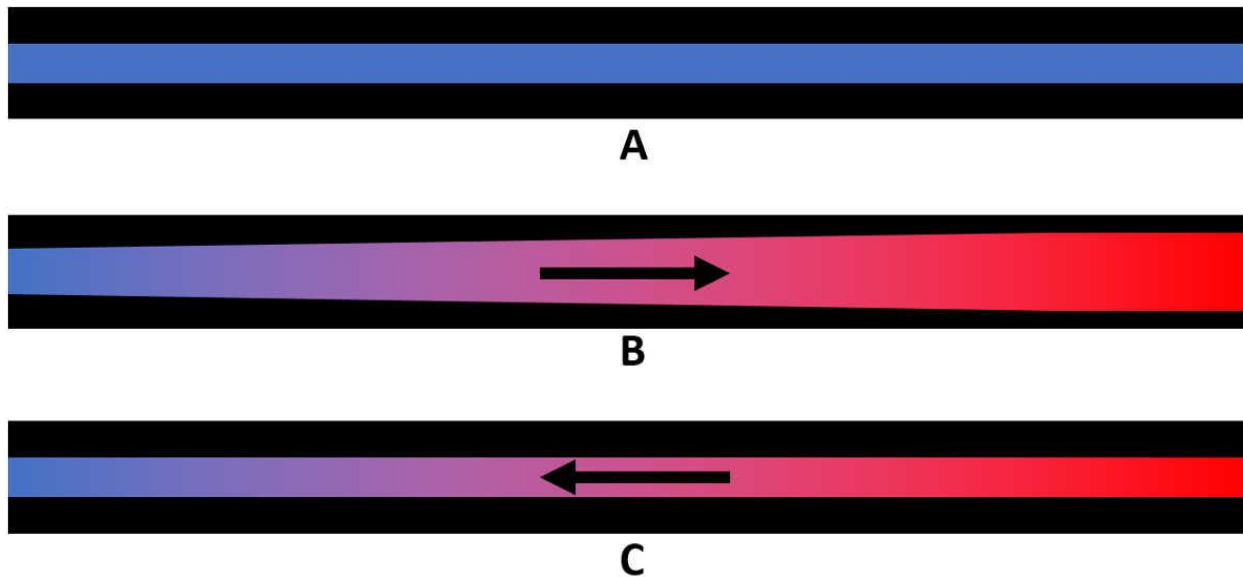


Figure 6.2: Regeneration of dissolved graphite by precipitation in the same location. Part A shows a graphite pipe with tin flowing through it before any dissolution and without heat transfer to drive dissolution. Part B shows an exaggerated geometry change of the pipe, after cold tin has been flowing in from the left and heated by the pipe, dissolving the wall mostly on the hot (right) side. Part C shows the pipe after hot tin has been flowing in from the right and cooled by the pipe, causing precipitation mostly on the right since the solubility gradient is larger at higher temperatures.

Before quantifying the effects of each of these tools, it is important to establish a baseline for a typical design where they are not implemented. This baseline was considered in the previous chapter, where a solubility of 0.003% was estimated to be the critical solubility before mitigation would be needed to prevent excessive material loss. This is exactly one order of magnitude lower than the solubility of C in Sn at 2300°C. Thus, some mitigation is required.

The most obvious way to reduce dissolution rates is to reduce temperature. Here, the 10X reduction needed could be achieved by reducing the peak temperature to 2000°C [169]. This is a simple and low risk solution, but it would significantly reduce system efficiency and increase cost. Here, the multi-junction photovoltaic efficiency is reduced by approximately 2% per 100°C, although it is nonlinear [16]. This alone increases the cost of storage, but the lower temperature also reduces heat flux by 70% due to the T^4 dependence of radiation. This lower heat flux means more MPV cells and emitter area are needed and the effect of parasitic heat loss becomes more important. Nonetheless, with the attractive cost of this system, lowering temperature is still a feasible option.

Reducing convective effects can reduce the mass transfer rate if the system is in a regime where the mass transfer is slow compared to the heat transfer. The Lewis number was specifically formulated for this comparison, which is defined as the thermal diffusivity divided by the mass diffusivity. For carbon in tin, it is approximately 5000 assuming the diffusion coefficient for C in tin is $5 \times 10^{-19} \text{ m}^2/\text{s}$, which is typical for hot liquids [151-153]. This indicates that heat transfer is much faster than mass transfer, which is very encouraging. However, realizing the benefits of this disparity are somewhat limited by other system constraints. First of all, if the flow is laminar, then changing the velocity has no effect on the convective enhancement. Furthermore, the velocity is not truly a free variable since an energy balance must be maintained between the heat flux from the pipes to the MPV (or from the heater to the pipes). Large pipes could be used, thus increasing the length scale for diffusion, but the cost of tin to fill these pipes quickly becomes exorbitant and the size of the heat exchanger also becomes large if the pipes are large, even with the good heat transfer inside the pipes. Despite these limitations, it is feasible for mass flux to vary by an order of magnitude or more by varying the Sherwood number and pipe radius. Notably, in the example used for the baseline, the flow is laminar, and the pipes are 50 mm diameter, so the

only way to reduce flux by a factor of ten would be to increase diameter by a factor of 10, which is not practical.

Another parameter that could be targeted is the solubility variation. Here, the upper limit can be considered fixed by temperature, but the lower limit could be intentionally *increased* to reduce the difference between them. In this way, the tin could maintain a significant amount of dissolved carbon throughout circulation, reducing the driving force and capacity for further dissolution. The problem here is that solubility increases exponentially with temperature, such that the only way to have a significant “cold” solubility would be for the hot and cold temperatures to be with about 100°C of each other. For example, the solubility of C in Sn at 2220°C is 0.015%, which is half the solubility at 2300°C [169]. This narrow temperature range is possible but would greatly reduce the range of sensible storage temperatures making the system much more expensive. This approach may also greatly increase the required velocity of the tin to avoid it cooling too much, which could counteract any attempt at reducing the convective enhancement of the mass transfer as discussed about.

A final, and very elegant, approach to mitigating the effects of dissolution is to simply reverse them periodically as shown in Figure 6.2. This may actually be uniquely plausible and convenient in an energy storage system, which already reverses state on a periodic basis during charge/discharge cycles. For example, during discharging, cold tin flows through the storage graphite and dissolves away graphite as it is heated. However, during charging, hot tin is cooled in the same region, precipitating out graphite, reversing the direction of the initial mass transfer. It may also be possible to cause a similar ‘natural regeneration’ in the heater and MPV system as well if they are combined. As shown in Figure 6.3, if the same pipes are used to receive heat from the heaters as to reject heat to the MPV, then regeneration can occur automatically. Here, the change from charging to discharging alone applies a reversed temperature profile which can cause precipitation which dissolution was occurring in the previous state. It is clear that some amount of regeneration will occur naturally, but the extent to which this happens, and its uniformity warrants further analysis—and likely experimentation.

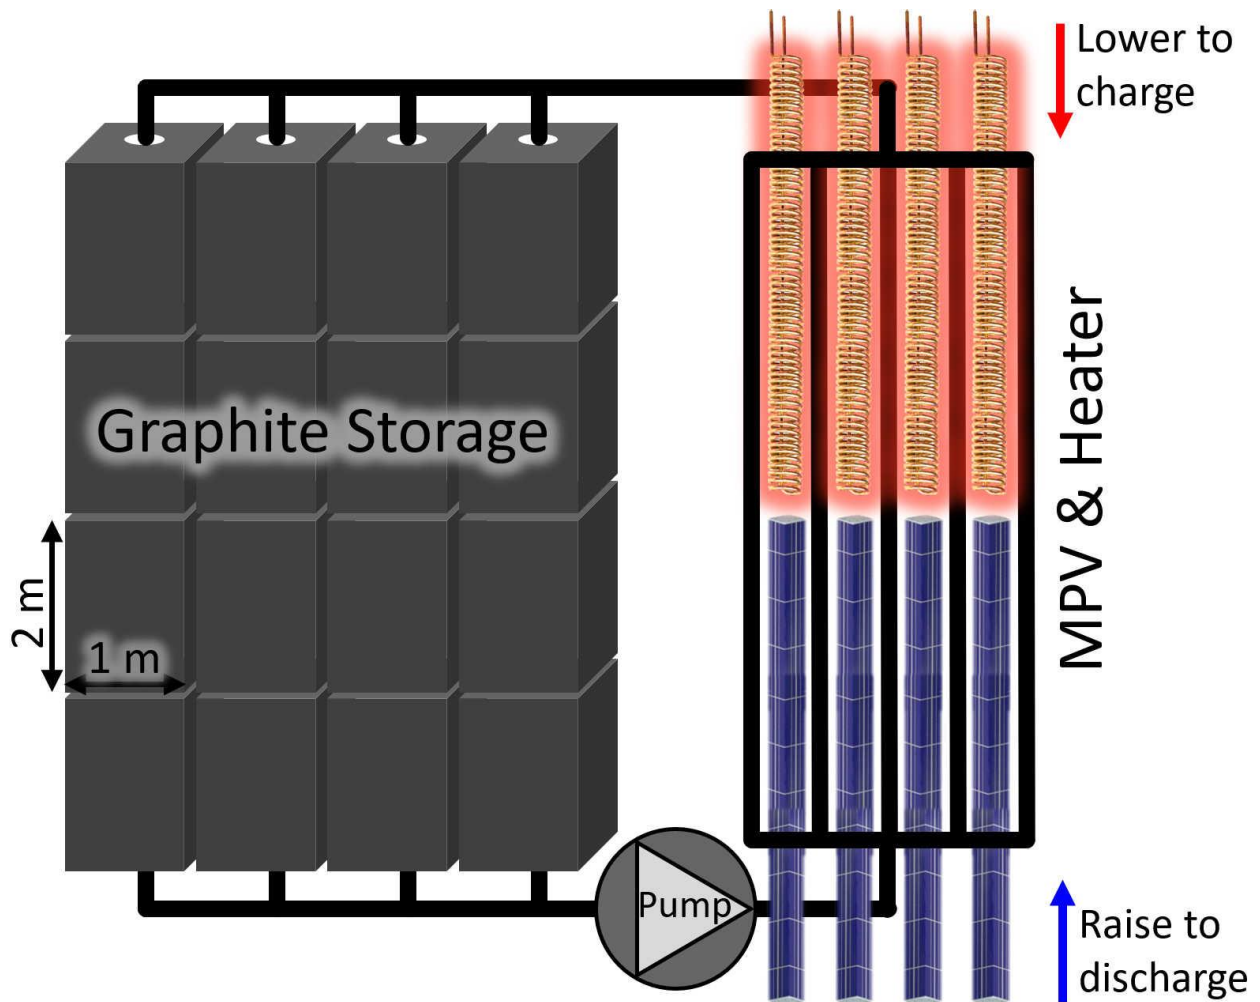


Figure 6.3: TEGS-Sn concept with combined Heater and MPV section, which can enable natural regeneration of carbon dissolution and precipitation cycling since the hot side of the Heater/MPV piping changes between charging and discharging.

Another area for future work involves high temperature solid electrical insulators. Since none were found to survive above 2300°C, or even 2000°C in a carbon rich gas environment, new material alloys or composites may need to be considered. One is B₄C doped SiC. This material was investigated in the 1980s as an electrical insulator and its fabrication and properties are documented in a patent, but no recent information was found [86]. Nonetheless, this material warrants future investigation because the patent claims that it retains stable and resistive above 2500°C.

Further work is also warranted to investigate the detailed effects of transient thermal cycling of silicon contained in graphite. This was investigated in Chapter 4, where it was shown

that SiC layer delamination during thermal cycling is a possibility due to the large thermal stress that can develop since the thermal expansion of graphite and SiC is different. Experimentally, it was shown that this delamination, likely in parallel with dissolution, can cause material loss at rates at least as high as 150 μm /cycle. However, this rate was spatially dependent, warranting future work to determine under which conditions the delamination can be avoided. It has been shown [37, 170, 171] at lower temperatures around the melting point (1414°C) of silicon (and fewer cycles (~4)) that delamination does not seem to occur. However, for energy applications hundreds, if not thousands, of cycles will be needed.

One last unaddressed issue with this extremely hot infrastructure that warrants future work is creep. Here, although the tensile strength of graphite increases to 2600°C [61], creep becomes an important degradation mechanism above 2000°C, especially under stresses near the fracture stress [111]. This problem is largely avoided by the use of solid containment and thick walled pipes, such that the stress remains only a few percent of the fracture stress. However, providing insulation below that can bear the weight of the storage material remains a challenge. Here, rigid graphite felt has sufficient compressive strength to hold ~50 m graphite tall (or silicon) above it, but it will likely experience devastating creep. Although data is not available, it is expected that this porous, partially sintered material will degrade much more quickly than dense graphite. One (perhaps the only) solution to this is to use at least some dense graphite as insulation beneath the tank. This can resolve the creep issue at the tradeoff of heat loss and the cost/size of this graphite. Nonetheless, further work to develop this concept is warranted.

References:

- [1] R. Fu, D. Feldman, R. Margolis, M. Woodhouse, and K. Ardani, "U.S. Solar Photovoltaic System Cost Benchmark: Q1 2017," National Renewable Energy Laboratory NREL/TP-6A20-68925 2017, Available: <https://www.nrel.gov/docs/fy17osti/68925.pdf>, Accessed on: March 9th, 2018.
- [2] T. Stehly, D. Heimiller, and G. Scott, "2016 Cost of Wind Energy Review," National Renewable Energy Laboratory NREL/TP-6A20-70363, 2017, Available: <https://www.nrel.gov/docs/fy18osti/70363.pdf>.
- [3] R. Sioshansi, P. Denholm, T. Jenkin, and J. Weiss, "Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects," *Energy Economics*, vol. 31, no. 2, pp. 269-277, 3// 2009.
- [4] P. Denholm, J. Jorgenson, M. Hummon, T. Jenkin, and D. Palchak, "The Value of Energy Storage for Grid Applications," National Renewable Energy Laboratory NREL/TP-6A20-58465, 2013, Available: <https://www.nrel.gov/docs/fy13osti/58465.pdf>.
- [5] F. J. de Sisternes, J. D. Jenkins, and A. Botterud, "The value of energy storage in decarbonizing the electricity sector," *Applied Energy*, vol. 175, pp. 368-379, 2016/08/01/ 2016.
- [6] M. S. Ziegler *et al.*, "Storage Requirements and Costs of Shaping Renewable Energy Toward Grid Decarbonization," *Joule*, vol. 3, no. 9, pp. 2134-2153, 2019/09/18/ 2019.
- [7] O. Schmidt, A. Hawkes, A. Gambhir, and I. Staffell, "The future cost of electrical energy storage based on experience rates," *Analysis* vol. 6, p. 17110, 07/10/online 2017.
- [8] R. Bowers, "Energy storage and renewables beyond wind, hydro, solar make up 4% of U.S. power capacity," U.S. Energy Information Administration 2017, Available: <https://www.eia.gov/todayinenergy/detail.php?id=31372>.
- [9] W. A. Braff, J. M. Mueller, and J. E. Trancik, "Value of storage technologies for wind and solar energy," *Nature Climate Change*, Article vol. 6, p. 964, 06/13/online 2016.
- [10] C. D. Thomas *et al.*, "Extinction risk from climate change," *Nature*, vol. 427, no. 6970, pp. 145-148, 2004/01/01 2004.
- [11] S. Schoenung, "Energy Storage Systems Cost Update " Sandia National Laboratories SAND2011-2730 2011.
- [12] M. Park, J. Ryu, W. Wang, and J. Cho, "Material design and engineering of next-generation flow-battery technologies," *Nature Reviews Materials*, Review Article vol. 2, p. 16080, 11/08/online 2016.
- [13] W. J. Cole and A. Frazier, "Cost projections for utility-scale battery storage," National Renewable Energy Lab.(NREL), Golden, CO (United States) 2019.
- [14] Y.-M. Chiang, L. Su, M. S. Pan, and Z. Li, "Lowering the Bar on Battery Cost," *Joule*, vol. 1, no. 2, pp. 212-219, 2017/10/11/ 2017.
- [15] Z. Li *et al.*, "Air-Breathing Aqueous Sulfur Flow Battery for Ultralow-Cost Long-Duration Electrical Storage," *Joule*, vol. 1, no. 2, pp. 306-327, 2017/10/11/ 2017.
- [16] C. Amy, H. R. Seyf, M. A. Steiner, D. J. Friedman, and A. Henry, "Thermal energy grid storage using multi-junction photovoltaics," *Energy & Environmental Science*, 10.1039/C8EE02341G 2019.

- [17] P. Albertus, J. S. Manser, and S. Litzelman, "Long-Duration Electricity Storage Applications, Economics, and Technologies," *Joule*, vol. 4, no. 1, pp. 21-32, 2020/01/15/ 2020.
- [18] S. M. Schoenung and W. V. Hassenzahl, "Long- vs. Short-Term Energy Storage Technologies Analysis," in "A Life-Cycle Cost Study, A Study for the DOE Energy Storage Systems Program " Sandia National Laboratories Albuquerque, New Mexico SAND2003-2783 2003.
- [19] C. Budischak, D. Sewell, H. Thomson, L. Mach, D. E. Veron, and W. Kempton, "Cost-minimized combinations of wind power, solar power and electrochemical storage, powering the grid up to 99.9% of the time," *Journal of Power Sources*, vol. 225, pp. 60-74, 2013/03/01/ 2013.
- [20] G. Glatzmaier, "Developing a Cost Model and Methodology to Estimate Capital Costs for Thermal Energy Storage," National Renewable Energy Laboratory 2011.
- [21] J. F. White, L. Ma, K. Forwald, and D. Sichen, "Reactions Between Silicon and Graphite Substrates at High Temperature: In Situ Observations," *Metallurgical and Materials Transactions B*, vol. 45, no. 1, pp. 150-160, 2014/02/01 2014.
- [22] R. Israel, R. Voytovych, P. Protsenko, B. Drevet, D. Camel, and N. Eustathopoulos, "Capillary interactions between molten silicon and porous graphite," *Journal of Materials Science*, vol. 45, no. 8, pp. 2210-2217, 2010/04/01 2010.
- [23] Y. Zhang *et al.*, "Containment materials for liquid tin at 1350 °C as a heat transfer fluid for high temperature concentrated solar power," *Solar Energy*, vol. 164, pp. 47-57, 2018/04/01/ 2018.
- [24] V. I. Nizhenko and L. I. Floka, "Contact reaction of graphite with liquid iron and iron base melts," *Soviet Powder Metallurgy and Metal Ceramics*, vol. 13, no. 6, pp. 487-492, 1974/06/01 1974.
- [25] G. Fiegl and W. Torbet, "Continuous replenishment of molten semiconductor in a Czochralski-process, single-crystal-growing furnace," ed: Google Patents, 1981.
- [26] (2020). *Carbolite Laboratory High temperature Furnace*. Available: <https://www.carbolite-gero.com/products/vacuum-furnace/laboratory-furnaces/lhtg/function-features/>
- [27] R. N. Ghoshtagore and R. Coble, "Self-diffusion in silicon carbide," *Physical Review*, vol. 143, no. 2, p. 623, 1966.
- [28] J. D. Hong, M. H. Hon, and R. F. Davis, "Self-diffusion in alpha and beta silicon carbide," *Ceramurgia International*, vol. 5, no. 4, pp. 155-160, 1979/10/01/ 1979.
- [29] A. Fritsch, J. Flesch, V. Geza, C. Singer, R. Uhlig, and B. Hoffschmidt, "Conceptual Study of Central Receiver Systems with Liquid Metals as Efficient Heat Transfer Fluids," *Energy Procedia*, vol. 69, pp. 644-653, 2015/05/01 2015.
- [30] X. Dong, X. Huang, L. Liu, L. He, and P. Li, "A liquid aluminum alloy electromagnetic transport process for high pressure die casting," *Journal of Materials Processing Technology*, vol. 234, pp. 217-227, 8// 2016.
- [31] C. Amy, "Liquid Metal Pumps for Enabling Heat Transfer at Extreme Temperatures," M.S., Mechanical Engineering, Georgia Institute of Technology, 2017.
- [32] C. Amy *et al.*, "Pumping liquid metal at high temperatures up to 1,673 kelvin," *Nature*, Article vol. 550, p. 199, 10/11/online 2017.

- [33] R. B. Laughlin, "Pumped thermal grid storage with heat exchange," *Journal of Renewable and Sustainable Energy*, vol. 9, no. 4, p. 044103, 2017/07/01 2017.
- [34] A. Dietrich, "Assessment of Pumped Heat Electricity Storage Systems through Exergoeconomic Analyses," Mechanical Engineering, University of Darmstadt, Germany, 2017.
- [35] U. Pelay, L. Luo, Y. Fan, D. Stitou, and M. Rood, "Thermal energy storage systems for concentrated solar power plants," *Renewable and Sustainable Energy Reviews*, vol. 79, pp. 82-100, 2017/11/01/ 2017.
- [36] C. S. Turchi, J. Vidal, and M. Bauer, "Molten salt power towers operating at 600–650 °C: Salt selection and cost benefits," *Solar Energy*, vol. 164, pp. 38-46, 2018/04/01/ 2018.
- [37] A. Datas *et al.*, "AMADEUS: Next generation materials and solid state devices for ultra high temperature energy storage and conversion," *AIP Conference Proceedings*, vol. 2033, no. 1, p. 170004, 2018/11/08 2018.
- [38] C. Amy, C. C. Kelsall, A. LaPotin, M. Pishahang, and A. Henry, "Ultra High Temperature Sensible Thermal Storage and Heat Transfer Fluids," in *Ultra-High Temperature Energy Storage, Transfer and Conversion*, A. Datas, Ed.: Elsevier, 2020.
- [39] C. W. Forsberg, D. C. Stack, D. Curtis, G. Haratyk, and N. A. Sepulveda, "Converting excess low-price electricity into high-temperature stored heat for industry and high-value electricity production," *The Electricity Journal*, vol. 30, no. 6, pp. 42-52, 2017/07/01/ 2017.
- [40] C. Forsberg, "Hybrid systems to address seasonal mismatches between electricity production and demand in nuclear renewable electrical grids," *Energy Policy*, vol. 62, pp. 333-341, 11/01 2013.
- [41] H. R. Seyf and A. Henry, "Thermophotovoltaics: a potential pathway to high efficiency concentrated solar power," *Energy & Environmental Science*, 10.1039/C6EE01372D vol. 9, no. 8, pp. 2654-2665, 2016.
- [42] J. P. Clinch and J. D. Healy, "Cost-benefit analysis of domestic energy efficiency," *Energy Policy*, vol. 29, no. 2, pp. 113-124, 2001.
- [43] M. Cai, Y. Wu, H. Chen, X. Yang, Y. Qiang, and L. Han, "Cost-performance analysis of perovskite solar modules," *Advanced Science*, vol. 4, no. 1, p. 1600269, 2017.
- [44] E. L. Ralph, "High efficiency solar cell arrays system trade-offs," in *Proceedings of 1994 IEEE 1st World Conference on Photovoltaic Energy Conversion - WCPEC (A Joint Conference of PVSC, PVSEC and PSEC)*, 1994, vol. 2, pp. 1998-2001 vol.2.
- [45] R. Sioshansi and P. Denholm, "The Value of Concentrating Solar Power and Thermal Energy Storage," *IEEE Transactions on Sustainable Energy*, vol. 1, no. 3, pp. 173-183, 2010.
- [46] A. Leroy, B. Bhatia, K. Wilke, O. Ilic, M. Soljačić, and E. N. Wang, "Combined selective emitter and filter for high performance incandescent lighting," *Applied Physics Letters*, vol. 111, no. 9, p. 094103, 2017/08/28 2017.
- [47] P. J. Verlinden and J. B. Lasich, "Energy rating of Concentrator PV systems using multi-junction III–V solar cells," in *2008 33rd IEEE Photovoltaic Specialists Conference*, 2008, pp. 1-6.
- [48] C. B. Alcock, V. P. Itkin, and M. K. Horrigan, "Vapour Pressure Equations for the Metallic Elements: 298–2500K," *Canadian Metallurgical Quarterly*, vol. 23, no. 3, pp. 309-313, 1984/07/01 1984.

- [49] F. F. Schmidt and H. R. Ogden, "The engineering properties of tungsten and tungsten alloys," DTIC Document 1963.
- [50] R. M. Swanson, "Silicon photovoltaic cells in thermophotovoltaic energy conversion," in *1978 International Electron Devices Meeting*, 1978, vol. 24, pp. 70-73.
- [51] Z. Omair *et al.*, "Experimental Demonstration of 28.2% Thermophotovoltaic Conversion Efficiency," in *Conference on Lasers and Electro-Optics*, San Jose, California, 2018, p. AW30.7: Optical Society of America.
- [52] R. R. King *et al.*, "40% efficient metamorphic GaInP/GaInAs/Ge multijunction solar cells," *Applied Physics Letters*, vol. 90, no. 18, p. 183516, 2007/04/30 2007.
- [53] V. M. B. Nunes, C. S. Queirós, M. J. V. Lourenço, F. J. V. Santos, and C. A. Nieto de Castro, "Molten salts as engineering fluids," *Applied Energy*, vol. 183, pp. 603-611, 12/1/ 2016.
- [54] T. Nagel, S. Kosik, and D. Fuhlbohm, "How SCR Controllers Help Boost Energy Efficiency," Advanced Energy Industries, Inc. ENG-SCR Controllers-Energy Efficiency-270-01 11.15, 2015, Available: https://www.advanced-energy.com/upload/File/White_Papers/ENG-SCRControllers-EnergyEfficiency-270-01.pdf.
- [55] M. Adler and S. Westbrook, "Power semiconductor switching devices—A comparison based on inductive switching," *IEEE Transactions on Electron Devices*, vol. 29, no. 6, pp. 947-952, 1982.
- [56] R. Gabrielli and C. Zamparelli, "Optimal Design of a Molten Salt Thermal Storage Tank for Parabolic Trough Solar Power Plants," *Journal of Solar Energy Engineering*, vol. 131, no. 4, pp. 041001-041001-10, 2009.
- [57] G. Wilk, "Liquid metal based high temperature concentrated solar power: Cost considerations," M.S., Mechanical Engineering, Georgia Institute of Technology, 2016.
- [58] H. Zhai and E. S. Rubin, "Performance and cost of wet and dry cooling systems for pulverized coal power plants with and without carbon capture and storage," *Energy Policy*, vol. 38, no. 10, pp. 5653-5660, 2010/10/01/ 2010.
- [59] A. Edington, "The role of long duration energy storage in decarbonizing power systems," 2019.
- [60] G. Wilk, A. DeAngelis, and A. Henry, "Estimating the cost of high temperature liquid metal based concentrated solar power," *Journal of Renewable and Sustainable Energy*, vol. 10, no. 2, p. 023705, 2018/03/01 2018.
- [61] C. Malmstrom, R. Keen, and L. Green, "Some Mechanical Properties of Graphite at Elevated Temperatures," *Journal of Applied Physics*, vol. 22, no. 5, pp. 593-600, 1951/05/01 1951.
- [62] S. Bodbodak and M. Moshfeghifar, *Advances in controlled atmosphere storage of fruits and vegetables* (Eco-Friendly Technology for Postharvest Produce Quality). Academic Press, 2016, pp. 39-76.
- [63] A. K. Thompson and D. Bishop, *Controlled Atmosphere Technology* (Reference Module in Food Science). Elsevier, 2016.
- [64] D. Keefe, "Blower door testing," *Journal of Light Construction*, pp. 1-7, 2010.
- [65] G. R. Griffith, "The Profitability of Controlled Atmosphere Storage of Apples in the Orange Area of New South Wales: A Benefit-Cost Study," New South Wales Department of Primary Industries Research Economists Aug 1974, Available: <https://ideas.repec.org/p/ags/nswarc/260382.html>.

- [66] M. Ormeci kart and V. Demircan, "Analysis of investment cost of apple cold storage facilities," *Custos e Agronegocio*, vol. 11, pp. 53-70, 06/26 2015.
- [67] M. Laing and R. L. Sheppard, "The economics of controlled atmosphere storage and transport for nectarines, apples and kiwifruit," 1984.
- [68] C.-W. Cheng, K.-T. Shiu, N. Li, S.-J. Han, L. Shi, and D. K. Sadana, "Epitaxial lift-off process for gallium arsenide substrate reuse and flexible electronics," *Nature communications*, vol. 4, no. 1, pp. 1-7, 2013.
- [69] "NREL Annual Technology Baseline," in "Natural Gas Plants," NREL2018, Available: atb.nrel.gov.
- [70] H.-J. Hoffmann, "Reasons for Melting of Chemical Elements and some Consequences," *Materialwissenschaft und Werkstofftechnik*, vol. 34, no. 6, pp. 571-582, 2003.
- [71] M. Beutl, G. Pottlacher, and H. Jäger, "Thermophysical properties of liquid iron," *International Journal of Thermophysics*, vol. 15, no. 6, pp. 1323-1331, 1994/11/01 1994.
- [72] R. Endo, Y. Fujihara, and M. Susa, *Calculation of the density and heat capacity of silicon by molecular dynamics simulation*. 2003, pp. 505-511.
- [73] (April 22, 2018). *Carbon Fiber Products - Soft felt*. Available: http://www.carbon.co.jp/english/products/fine_carbon/
- [74] (April 22, 2018). *Alumina-Silica Blanket - Type ASB*. Available: <https://www.zircarceramics.com/product/asb/>
- [75] D. R. Smith, J. G. Hust, and L. J. V. Poolen, "Effective Thermal Conductivity of a glass fiberblanket standard reference material," Thermophysical Properties Division, National Engineering Laboratory, National Bureau of StandardsNBSIR 81-1640, 1981, Available: <https://www.gpo.gov/fdsys/pkg/GOVPUB-C13-da95d64f577af8468f325f311d3f5c7b/pdf/GOVPUB-C13-da95d64f577af8468f325f311d3f5c7b.pdf>.
- [76] (April 22, 2018). *ALII Products*. Available: <http://westadmat.com/aliiproducts/>
- [77] T. Nishi, H. Shibata, Y. Waseda, and H. Ohta, "Thermal conductivities of molten iron, cobalt, and nickel by laser flash method," *Metallurgical and Materials Transactions A*, vol. 34, no. 12, pp. 2801-2807, 2003/12/01 2003.
- [78] C. Amy, J. Barber, and C. Kelsall, "Paths to Scalable Carbon Neutrality for MIT," *arXiv preprint arXiv:1904.06296*, 2019.
- [79] C. Batlle and I. J. Pérez-Arriaga, "Design criteria for implementing a capacity mechanism in deregulated electricity markets," *Utilities Policy*, vol. 16, no. 3, pp. 184-193, 2008.
- [80] T. Jenkin, P. Beiter, and R. Margolis, "Capacity Payments in Restructured Markets under Low and High Penetration Levels of Renewable Energy," National Renewable Energy LaboratoryNREL/TP-6A20-65491, 2016, Available: <https://www.nrel.gov/docs/fy16osti/65491.pdf>.
- [81] P. Denholm, E. Ela, B. Kirby, and M. Milligan, "Role of energy storage with renewable electricity generation," National Renewable Energy Lab.(NREL), Golden, CO (United States)2010.
- [82] "State of the Market Report for PJM," Monitoring Analytics, LLC2015, Available: http://www.monitoringanalytics.com/reports/PJM_State_of_the_Market/2015.shtml.

- [83] "2018/2019 RPM Base Residual Auction Planning Period Parameters," PJM2015, Available: <http://pjm.com/~media/markets-ops/rpm/rpm-auction-info/2018-2019-planning-parameters-report.ashx>.
- [84] "Challenges and Opportunities For New Pumped Storage Development," NHA's Pumped Storage Development Council2012, Available: http://www.hydro.org/wp-content/uploads/2012/07/NHA_PumpedStorage_071212b.pdf.
- [85] "Lazard's Levelized Cost of Storage Analysis - Version 3.0," Lazard and Enovation Partners2017, Available: <https://www.lazard.com/media/450338/lazard-levelized-cost-of-storage-version-30.pdf>.
- [86] W. D. G. Boecker, "Sintered silicon carbide ceramic body of high electrical resistivity," EU, 1985.
- [87] J. Pappis and W. D. Kingery, "Electrical Properties of Single-Crystal and Polycrystalline Alumina at High Temperatures," *Journal of the American Ceramic Society*, vol. 44, no. 9, pp. 459-464, 1961/09/01 1961.
- [88] R. Powell and F. Schofield, "The thermal and electrical conductivities of carbon and graphite to high temperatures," *Proceedings of the physical society*, vol. 51, no. 1, p. 153, 1939.
- [89] V. A. Nemchinsky and M. S. Showalter, "Cathode erosion in high-current high-pressure arc," *Journal of Physics D: Applied Physics*, vol. 36, no. 6, pp. 704-712, 2003/02/27 2003.
- [90] W. Gajewski *et al.*, "Electronic and optical properties of boron-doped nanocrystalline diamond films," *Physical Review B*, vol. 79, no. 4, p. 045206, 01/15/ 2009.
- [91] J. Dixon, L. LaGrange, U. Merten, C. Miller, and J. Porter II, "Electrical resistivity of stabilized zirconia at elevated temperatures," *Journal of the Electrochemical Society*, vol. 110, no. 4, p. 276, 1963.
- [92] D. R. Gaskell and D. E. Laughlin, "The Third Law of Thermodynamics," in *Introduction to the Thermodynamics of Materials*: CRC press, 2017, p. 148.
- [93] J. Hong *et al.*, "Exploring atomic defects in molybdenum disulphide monolayers," (in eng), *Nature communications*, vol. 6, pp. 6293-6293, 2015.
- [94] F. P. Bundy, H. P. Bovenkerk, H. M. Strong, and R. H. Wentorf Jr, "Diamond-graphite equilibrium line from growth and graphitization of diamond," *The Journal of Chemical Physics*, vol. 35, no. 2, pp. 383-391, 1961.
- [95] M. Seal, "Graphitization of Diamond," *Nature*, vol. 185, no. 4712, pp. 522-523, 1960/02/01 1960.
- [96] G. Davies, T. Evans, and R. W. Ditchburn, "Graphitization of diamond at zero pressure and at a high pressure," *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, vol. 328, no. 1574, pp. 413-427, 1972.
- [97] H.-T. Kim, J.-S. Han, J.-H. Yang, J.-B. Lee, and S.-H. Kim, "The effect of low temperature aging on the mechanical property & phase stability of Y-TZP ceramics," *The Journal of Advanced Prosthodontics*, vol. 1, no. 3, pp. 113-117, 2009 2009.
- [98] J. Wang and R. Stevens, "Zirconia-toughened alumina (ZTA) ceramics," *Journal of Materials Science*, vol. 24, no. 10, pp. 3421-3440, 1989// 1989.
- [99] DerSilberspiegel, "Ellingham Richardson-diagram," in *DerSilberspiegel [CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0>)]*, ed, 2016.

- [100] "Temperature and Resistivity for Tungsten," Pasco Scientific, Available: http://physics.usask.ca/~bzulkosk/Lab_Manuals/EP354/ep354_thermal-radiation-lab_Tungsten-Temp-Resistivity.pdf.
- [101] D. M. Mattox, "Chapter 5 - The Low Pressure Plasma Processing Environment," in *Handbook of Physical Vapor Deposition (PVD) Processing (Second Edition)*, D. M. Mattox, Ed. Boston: William Andrew Publishing, 2010, pp. 157-193.
- [102] O. W. Richardson, "Electron Emission from Metals as a Function of Temperature," *Physical Review*, vol. 23, no. 2, pp. 153-155, 02/01/ 1924.
- [103] C. D. Child, "Discharge From Hot CaO," *Physical Review (Series I)*, vol. 32, no. 5, pp. 492-511, 05/01/ 1911.
- [104] J. G. Simmons, "Poole-Frenkel effect and Schottky effect in metal-insulator-metal systems," *Physical Review*, vol. 155, no. 3, p. 657, 1967.
- [105] "Electric arc heater," Thermopedia2011, Available: <http://www.thermopedia.com/content/718/>.
- [106] V. Messerle, A. Ustimenko, V. Lukyashchenko, V. Z. Ushanov, and E. Karpenko, "Long life DC arc plasmatron with nanocarbon coating of electrodes," in *Proceedings of the 20th International Symposium on Plasma Chemistry (ISPC20)-Philadelphia, USA, 2011*, pp. 24-29.
- [107] A. N. Bhoj and M. J. Kushner, "Avalanche process in an idealized lamp: II. Modelling of breakdown in Ar/Xe electric discharges," *Journal of Physics D: Applied Physics*, vol. 37, no. 18, p. 2510, 2004.
- [108] A. D. A. Mazli, N. A. M. Jamail, and N. A. Othman, "Electric Field and Current Density Performance Analysis of Sf6, C4f8 and CO2 Gases As An Insulation," *IOP Conference Series: Materials Science and Engineering*, vol. 226, p. 012132, 2017/08 2017.
- [109] G. Frenking, "Another noble gas conquered," *Nature*, vol. 406, no. 6798, pp. 836-837, 2000/08/01 2000.
- [110] D. George and H. Messerle, "Electrical breakdown in high-temperature gases," *Nature*, vol. 197, no. 4867, pp. 584-585, 1963.
- [111] P. Wagner, A. R. Driesner, and L. A. Haskin, "High-Temperature Properties of Graphite. II. Creep in Tension," *Journal of Applied Physics*, vol. 30, no. 2, pp. 152-154, 1959.
- [112] B. Ebihara, "Refractory metal shielding/insulation/increases operating range of induction furnace," 1965.
- [113] S. Lang, D. Bestenlehner, R. Marx, and H. Drück, "Thermal insulation of an ultra-high temperature thermal energy store for concentrated solar power," in *AIP Conference Proceedings*, 2018, vol. 2033, no. 1, p. 090020: AIP Publishing LLC.
- [114] L. Stephan, *Ultra-high temperature thermal insulation*. 2020.
- [115] V. A. Petrov, "Combined radiation and conduction heat transfer in high temperature fiber thermal insulation," *International Journal of Heat and Mass Transfer*, vol. 40, no. 9, pp. 2241-2247, 1997/06/01/ 1997.
- [116] F. P. Incropera, A. S. Lavine, T. L. Bergman, and D. P. DeWitt, *Fundamentals of heat and mass transfer*. Wiley, 2007.
- [117] M. Wiener, G. Reichenauer, S. Braxmeier, F. Hemberger, and H. P. Ebert, "Carbon Aerogel-Based High-Temperature Thermal Insulation," *International Journal of Thermophysics*, vol. 30, no. 4, pp. 1372-1385, 2009/08/01 2009.

- [118] J. Feng, C. Zhang, and J. Feng, "Carbon fiber reinforced carbon aerogel composites for thermal insulation prepared by soft reinforcement," *Materials Letters*, vol. 67, no. 1, pp. 266-268, 2012/01/15/ 2012.
- [119] W. C. Turner and J. F. Malloy, "Thermal insulation handbook.[Includes glossary]," 1981.
- [120] T. J. Seebeck, *Magnetische polarisation der metalle und erze durch temperatur-differenz* (no. 70). W. Engelmann, 1895.
- [121] R. R. Asamoto and P. E. Novak, "Tungsten-Rhenium Thermocouples for Use at High Temperatures," *Review of Scientific Instruments*, vol. 38, no. 8, pp. 1047-1052, 1967/08/01 1967.
- [122] L. Michalski, K. Eckersdorf, J. Kucharski, and J. McGhee, "Automatic Pyrometers," in *Temperature Measurement*: Wiley, 2002, pp. 177-208.
- [123] W. G. Fahrenholtz and G. E. Hilmas, "Ultra-high temperature ceramics: Materials for extreme environments," *Scripta Materialia*, vol. 129, pp. 94-99, 3/1/ 2017.
- [124] J. Schemel, *ASTM manual on zirconium and hafnium* (no. 639). Astm International, 1977.
- [125] W. V. WRIGHT Jr and F. T. BARTELS, "CONSIDERATIONS OF THE TERNARY SYSTEM CARBON-SILICON-GERMANIUM," in *Silicon Carbide High Temperature Semiconductor: Proceedings of the Conference on Silicon Carbide, Boston, Mass. April, 1959*, 1960, p. 31: Pergamon.
- [126] H. Kleykamp and G. Schumacher, "The constitution of the silicon-carbon system," *Berichte der Bunsengesellschaft für physikalische Chemie*, vol. 97, no. 6, pp. 799-804, 1993.
- [127] H. Kleykamp, "Gibbs energy of formation of SiC: A contribution to the thermodynamic stability of the modifications," *Berichte der Bunsengesellschaft für physikalische Chemie*, vol. 102, no. 9, pp. 1231-1234, 1998.
- [128] O. Ruff, "The Formation and dissociation of silicon carbide," *Transactions of the Electrochemical Society*, vol. 68, no. 1, p. 87, 1935.
- [129] R. I. Scace and G. A. Slack, "Solubility of Carbon in Silicon and Germanium," *The Journal of Chemical Physics*, vol. 30, no. 6, pp. 1551-1555, 1959/06/01 1959.
- [130] P. W. Atkins, J. De Paula, and J. Keeler, *Atkins' physical chemistry*. Oxford university press, 2018.
- [131] W. J. Moore and B. Selikson, "The Diffusion of Copper in Cuprous Oxide," *The Journal of Chemical Physics*, vol. 19, no. 12, pp. 1539-1543, 1951/12/01 1951.
- [132] A. Heuer, "Oxygen and Aluminum Diffusion in α -Al₂O₃: How Much Do We Really Understand?," *Journal of the European Ceramic Society*, vol. 28, pp. 1495-1507, 12/31 2008.
- [133] S. Chapman, T. G. Cowling, and D. Burnett, *The mathematical theory of non-uniform gases: an account of the kinetic theory of viscosity, thermal conduction and diffusion in gases*. Cambridge university press, 1990.
- [134] J. Vidal, "Molten Salt Technology," DOE CSP Gen 3SAND2016-8201PE, 2016, Available: <https://www.osti.gov/servlets/purl/1378067>.
- [135] K. Schulgasser, "The relationship between stiffness and thermal expansion of carbon fibres," *Materials Science and Engineering: A*, vol. 110, pp. L35-L37, 1989/03/01/ 1989.
- [136] E. Le Bourhis, "Appendix 3: Thermal Expansion and Elasticity," in *Glass*: Wiley-VCH Verlag GmbH & Co. KGaA, 2007, pp. 271-274.

- [137] R. P. Messner and Y.-M. Chiang, "Liquid-Phase Reaction-Bonding of Silicon Carbide Using Alloyed Silicon-Molybdenum Melts," *Journal of the American Ceramic Society*, vol. 73, no. 5, pp. 1193-1200, 1990.
- [138] R. S. Averback, H. J. Höfler, H. Hahn, and J. C. Logas, "Sintering and grain growth in nanocrystalline ceramics," *Nanostructured Materials*, vol. 1, no. 2, pp. 173-178, 1992/03/01/ 1992.
- [139] E. W. Washburn, "The dynamics of capillary flow," *Physical review*, vol. 17, no. 3, p. 273, 1921.
- [140] R. E. Thomas, "25 - SILICON SOLAR CELLS (2) — PRACTICAL ASPECTS," in *Solar Energy Conversion*, A. E. Dixon and J. D. Leslie, Eds.: Pergamon, 1979, pp. 805-830.
- [141] N. Yuge *et al.*, "Purification of metallurgical-grade silicon up to solar grade," *Progress in Photovoltaics: Research and Applications*, vol. 9, no. 3, pp. 203-209, 2001/05/01 2001.
- [142] N. N. Greenwood and A. Earnshaw, *Chemistry of the Elements*. Elsevier, 2012.
- [143] A. M. Kuhlmann, "Silicon metal production," ed: Google Patents, 1965.
- [144] E. Schürmann and H. Jacke, "Schmelzgleichgewichte im System Fe—C—Ca," *Steel Research*, vol. 58, no. 9, pp. 399-405, 1987.
- [145] H. Ohtani, M. Yamano, and M. Hasebe, "Thermodynamic analysis of the Fe-Al-C ternary system by incorporating ab initio energetic calculations into the CALPHAD approach," *ISIJ international*, vol. 44, no. 10, pp. 1738-1747, 2004.
- [146] R. Olesinski, N. Kanani, and G. Abbaschian, "The P- Si (Phosphorus-Silicon) system," *Bulletin of Alloy Phase Diagrams*, vol. 6, no. 2, pp. 130-133, 1985.
- [147] H. Okamoto, "The C-Fe (carbon-iron) system," *Journal of Phase Equilibria*, vol. 13, no. 5, pp. 543-565, 1992.
- [148] K. Takano *et al.*, "Volume and dislocation diffusion of iron, chromium and cobalt in CVD β -SiC," *Science and Technology of Advanced Materials*, vol. 2, no. 2, pp. 381-388, 2001/01/01 2001.
- [149] S. Esfahani and M. Barati, "Purification of metallurgical silicon using iron as impurity getter, part II: Extent of silicon purification," *Metals and Materials International*, vol. 17, no. 6, pp. 1009-1015, 2011/12/01 2011.
- [150] S. Kawanishi, T. Yoshikawa, and T. Tanaka, "Equilibrium Phase Relationship between SiC and a Liquid Phase in the Fe-Si-C System at 1523–1723 K," *MATERIALS TRANSACTIONS*, vol. 50, no. 4, pp. 806-813, 2009.
- [151] T. SAITO, Y. KAWAI, K. MARUYA, and M. Michiyoshi, "Diffusion of some alloying elements in liquid iron," *Science reports of the Research Institutes, Tohoku University. Ser. A, Physics, chemistry and metallurgy*, vol. 11, pp. 401-410, 1959.
- [152] W. Yu, Z. Wang, and D. Stroud, "Empirical molecular-dynamics study of diffusion in liquid semiconductors," *Physical Review B*, vol. 54, no. 19, p. 13946, 1996.
- [153] J. Qin, X. Li, J. Wang, and S. Pan, "The self-diffusion coefficients of liquid binary M-Si (M=Al, Fe, Mg and Au) alloy systems by first principles molecular dynamics simulation," *AIP Advances*, vol. 9, no. 3, p. 035328, 2019/03/01 2019.
- [154] L. A. Sphaier and W. M. Worek, "Parametric analysis of heat and mass transfer regenerators using a generalized effectiveness-NTU method," *International Journal of Heat and Mass Transfer*, vol. 52, no. 9, pp. 2265-2272, 2009/04/01/ 2009.
- [155] G. L. Harris, *Properties of silicon carbide* (no. 13). let, 1995.

- [156] J. F. Kelly, G. R. Fisher, and P. Barnes, "Correlation between layer thickness and periodicity of long polytypes in silicon carbide," *Materials Research Bulletin*, vol. 40, no. 2, pp. 249-255, 2005/02/15/ 2005.
- [157] J. M. Bind, "Contribution of the ionic bonding to the thermal stability of SiC polytypes," *Materials Research Bulletin*, vol. 14, no. 9, pp. 1173-1179, 1979/09/01/ 1979.
- [158] L. N. Hailey, "Method of producing composite ceramic articles," ed: Google Patents, 1983.
- [159] E. Scalise, A. Marzegalli, F. Montalenti, and L. Miglio, *Temperature-dependent stability of polytypes and stacking faults in SiC: reconciling theory and experiments*. 2019.
- [160] D. N. Talwar and J. C. Sherbondy, "Thermal expansion coefficient of 3C-SiC," *Applied Physics Letters*, vol. 67, no. 22, pp. 3301-3303, 1995.
- [161] D. Varshney, S. Shriya, M. Varshney, N. Singh, and R. Khenata, "Elastic and thermodynamical properties of cubic (3C) silicon carbide under high pressure and high temperature," *Journal of Theoretical and Applied Physics*, vol. 9, no. 3, pp. 221-249, 2015/09/01 2015.
- [162] I. Hutchings and P. Shipway, *Friction and Wear of Engineering Materials* (Tribology (2nd)). Butterworth-Heinemann, 2017, pp. 1-5.
- [163] W. E. Forsthoffer, *Pump Best Practices* (Forsthoffer's Best Practice Handbook for Rotating Machinery). Boston: Butterworth-Heinemann, 2011, pp. 25-91.
- [164] T. Wang, F. Kong, B. Xia, Y. Bai, and C. Wang, "The method for determining blade inlet angle of special impeller using in turbine mode of centrifugal pump as turbine," *Renewable Energy*, vol. 109, pp. 518-528, 2017.
- [165] J. Tuzson, *Centrifugal pump design*. John Wiley & Sons, 2000.
- [166] R. Spikes and G. Pennington, "Discharge coefficient of small submerged orifices," *Proceedings of the Institution of Mechanical Engineers*, vol. 173, no. 1, pp. 661-674, 1959.
- [167] A. Dokoumetzidis and P. Macheras, "A century of dissolution research: From Noyes and Whitney to the Biopharmaceutics Classification System," *International Journal of Pharmaceutics*, vol. 321, no. 1, pp. 1-11, 2006/09/14/ 2006.
- [168] C. E. Crestani, A. Bernardo, C. B. B. Costa, and M. Giuliotti, "Experimental data and estimation of sucrose solubility in impure solutions," *Journal of Food Engineering*, vol. 218, pp. 14-23, 2018/02/01/ 2018.
- [169] L. L. Oden and N. A. Gokcen, "Sn-C and Al-Sn-C phase diagrams and thermodynamic properties of C in the alloys: 1550 °C to 2300 °C," *Metallurgical Transactions B*, vol. 24, no. 1, pp. 53-58, 1993// 1993.
- [170] B. Grorud, "Interaction of Eutectic Fe-Si-B Alloy with Graphite Crucibles," NTNU, 2018.
- [171] J. Jiao *et al.*, "The use of eutectic Fe-Si-B alloy as a phase change material in thermal energy storage systems," *Materials*, vol. 12, no. 14, p. 2312, 2019.