

**IMPLEMENTATION OF CREEP-FATIGUE DAMAGE MODELING FOR INFORMED  
OPERATION AND OPTIMIZED DESIGN OF CONCENTRATED SOLAR POWER  
TOWER RECEIVERS**

by

Jacob N Wenner

A dissertation submitted in partial fulfillment of  
the requirements for the degree of

Doctor of Philosophy

(Mechanical Engineering)

at the

UNIVERSITY OF WISCONSIN–MADISON

2025

Date of oral exam: 12-15-2025

The dissertation is approved by the following members of the Final Oral Committee:

Michael J. Wagner, Assistant Professor, Mechanical Engineering

Greg Nellis, Professor, Mechanical Engineering

Douglas Reindl, Professor, Mechanical Engineering

Mark Anderson, Professor, Mechanical Engineering

Krishnan Suresh, Professor, Mechanical Engineering

*This work is dedicated to Matt and Megan. Two people who left our lives far too soon. To both of you: you were loved, you are with me in memories, and you will be missed.*

## ACKNOWLEDGMENTS

---

As I reflect on some of the best years of my life, I have a myriad of things to be grateful for.

To Mom and Dad, thanks for constant encouragement and love. Among many lessons, you taught me that there are no replacements for trusting God and working hard. These two things, like you two, have never let me down.

To Tracey, thanks for your love, support, and companionship. Above all, thank you for putting my work in eternal context and daily warming my heart in the way you care for me and everyone else in your life. And, lastly, for your legendary risotto.

To the rest of my family, thank you for being so smart and yet so goofy all at the same time.

To my advisors, Mike, thank you for leading your lab by example and deeply valuing each of your students. It is a true joy to work for someone that constantly earns respect. Greg, thanks for being a great teacher and advisor. Most of my thermodynamics and heat transfer "tools" come from you.

To my friends, you are yet to read your shoutout in my master's thesis... and I expect a similar fate for this acknowledgement. Nevertheless, I am so grateful for the life you have given me outside of work. Thanks for keeping me sane!

To my labmates, thank you for caring for those around you, always asking great technical questions, supporting me, and just being all-round fun. Your enthusiasm for all things engineering is infectious.

I view these people and my time at UW-Madison as a collection of undeserved, yet deeply valued, blessings. Nevertheless, my most treasured gift is the one that God offers to all of us — the gift of new life and forgiveness through the death of Jesus Christ.

*This is good, and pleases God our Savior, 4 who wants all people to be saved and to come to a knowledge of the truth. 5 For there is one God and one mediator*

*between God and mankind, the man Christ Jesus, 6 who gave himself as a ransom for all people.*

— 1 TIMOTHY 2:3-6A

Hopefully the present work benefits its readers and society as a whole. However, I encourage the reader to reference John 3:16, which is both infinitely more important and significantly shorter than the present document.



## CONTENTS

---

Contents iv

List of Tables vi

List of Figures vii

Abstract xvii

### 1 Introduction 1

- 1.1 *CSP Technology Background* 5
- 1.2 *Reliability Challenges of Tower Receivers* 10
- 1.3 *Statement of Research Questions* 14
- 1.4 *Existing Literature* 15
- 1.5 *Primary Contributions and Novelty* 27
- 1.6 *Document Overview* 28

### 2 Damage Modeling of Central Receivers with SRLIFE 30

- 2.1 *Introduction* 30
- 2.2 *Methodology for Creep-Fatigue Analyses* 31
- 2.3 *Damage Map for Haynes 230* 52
- 2.4 *Damage Maps for Alternative Alloys* 58
- 2.5 *Damage Maps at Other Conduction-Convection Ratios* 64
- 2.6 *FEA Versus Norton-Bailey Method for Stress Relaxation* 69
- 2.7 *The Impact of Cycle Definition on Prediction Accuracy* 75
- 2.8 *Summary of Damage Modeling Work* 79

### 3 Software Tools for Damage Analysis of Central Receiver Design and Operations 82

- 3.1 *FEA Data Availability* 82
- 3.2 *Damage Tool for Design* 84

3.3	<i>Tool Application: Lifetime Profiles of a Cylindrical Receiver</i>	90
3.4	<i>Damage Tool for Real-Time Operation</i>	98
3.5	<i>A Modified Thermal Model Tool</i>	105
3.6	<i>Summary of Damage Tool and Thermal Model Contributions</i>	127
4	<b>Reliability-Focused Aiming and Design for Central Receivers</b>	130
4.1	<i>Defining an Ideally Utilized Receiver Flux Profile</i>	131
4.2	<i>Challenges of Achieving Ideal Flux Profiles</i>	144
4.3	<i>Informed SolarPILOT Aiming Method</i>	145
4.4	<i>Development of an Informed Aiming Heuristic</i>	152
4.5	<i>Comparison of Aiming Strategies</i>	160
4.6	<i>The Impact of Informed Aiming on Design</i>	163
4.7	<i>Summary of Informed Aiming and Design</i>	194
5	<b>Summary</b>	196
	<b>References</b>	202
A	<b>Appendix: Study of Power Block Dynamics in Hybrid System Applications</b>	212

## LIST OF TABLES

---

1.1	Related literature in which methodology and results include creep-fatigue analysis and/or lifetime predictions. . . . .	19
2.1	Various alloy conductivities and the corresponding temperature difference ratio used to define boundary conditions in <i>SRLIFE</i> . . . . .	52
3.1	Fluid temperature and temperature difference ranges and data resolution for FEA data obtained in this work. . . . .	85
3.2	Lifetime results for Crescent Dunes-type cylindrical receiver. Materials include the original Haynes 230 as well as three additional ones. . . . .	97
4.1	Number of panels that fail before 30 years of operation. IA-SPT stands for informed <i>SolarPILOT</i> method, and IA-Heuristic is the informed heuristic. 162	
4.2	Thermal model parameters and simulation settings common to all cases for the aspect ratio study. . . . .	175
4.3	Minimum LCOH and the aspect ratio at which the minimum occurs for two aiming methods at two different design powers. . . . .	194

## LIST OF FIGURES

---

1.1	Global, annual trends in capacity factor and levelized cost of energy for concentrating solar power and photovoltaics. Graphic and data from IRENA (2024). . . . .	4
1.2	Rankine cycle power block, steam generation system, and nitrate salt power tower design illustration made by NREL (2018) . . . . .	6
1.3	Power tower for receiver testing with a solar field in the foreground at Plataforma Solar de Almeria in Spain. Photo taken by author. . . . .	7
1.4	A sample cylindrical receiver (1.4a) from Conroy et al. (2020) and various panel designs (1.4b) from concentrating solar research at Sandia National Lab . . . . .	8
1.5	. . . . .	11
1.6	Thermoelastic results from SRLIFE for a Haynes 230 tube at conditions of $T_f = 300^\circ\text{C}$ , $\Delta T = 220^\circ\text{C}$ . . . . .	12
1.7	Maximum wall temperature, fluid temperature, absorbed flux, and total elastic stress for a Crescent Dunes solstice clear sky thermal model simulation. Results from each of the 14 panels are shown versus solar hour. Thermal model is from Martinek et al. (2021). . . . .	13
2.1	Analysis process diagram showing how creep and fatigue damages are calculated for each thermal case and material. . . . .	32
2.2	Simulation results from thermal model showing general temperature characteristics in receiver tubes. . . . .	37
2.3	Solstice day results from thermal model in showing the range of temperature gradients observed in each panel and the respective radial/circumferential gradients. . . . .	38
2.4	Timeseries (2.4a) and circumferential (2.4b) boundary conditions used in SRLIFE. . . . .	38

2.5	Maximum crown stress (2.5a) and creep damage per cycle (2.5b) for three different temperature cases showing stress relaxation phenomena as well as stress reset in the $T_f = 475^\circ\text{C}$ , $\Delta T = 220^\circ\text{C}$ case. . . . .	45
2.6	Comparison of extrapolated 20 <sup>th</sup> day creep damage versus actual damage from extended 20-day cycles. Extrapolation inaccuracy is only prevalent in 2.6a at points with creep damage too high for viable designs, and 2.6b shows zoomed in results where accuracy is sufficient in expected design/operation region. . . . .	46
2.7	Von Mises strain range timeseries from a simulation case with $T_f = 300^\circ\text{C}$ , $\Delta T = 220^\circ\text{C}$ . . . . .	47
2.8	Yield strength (2.8a) and 100,000 hr rupture strength (2.8b) for the six receiver alloys tested in this work. Rupture strength properties for Haynes 230 are based on the Mendelson-Roberts-Manson coefficients in Eno et al. (2008). All other alloy properties are obtained using MESSNER et al. (2023). . . . .	51
2.9	Third cycle maximum elastic-plastic strain range results from SRLIFE for Haynes 230 tubes with 50.8 mm diameter and 1.25 mm thickness subject to a range of $\Delta T$ s with three different fluid temperatures. . . . .	53
2.10	Elastic stress map for all tube $T_f$ - $\Delta T$ combinations simulated in SRLIFE Haynes 230 baseline case study. . . . .	54
2.11	Day 1 (2.11a) and day 4 (2.11b) elastic-plastic stress map for all tube $T_f$ - $\Delta T$ combinations simulated in SRLIFE Alloy 230 baseline case study. Tube diameter is 50.8 mm and thickness is 1.25 mm. Note that the color scales differ. . . . .	55
2.12	Damage map for baseline case (Alloy 230), showing LTE contour lines and locations of critical operating points for receiver tubes based on operation on a 9 hr day. The stress reset region (S.R.R.) and the corrosion region (C.R.) are also shown. Lifetime units are in years, and assume that no significant stress relaxation occurs after the 20th cycle and that fatigue damage is constant after the 3rd cycle. $\Delta T$ is the total temperature difference defined in Section 2.2 and $T_f$ is the fluid temperature. . . . .	58

2.13	Damage map for Inconel 617 based on results from SRLIFE. Contour lines from 5–80 years are shown. The stress reset region (S.R.R.) and the corrosion region (C.R.) are also shown. Color levels denote regions of different ratios of fatigue to creep damage. . . . .	60
2.14	Damage map for Haynes 282 based on results from SRLIFE. Contour lines from 5–80 years are shown. The upper right shaded area is the corrosion region (C.R.). Stress reset does not occur for Haynes 282 within this parametric space, and fatigue is not significant compared to creep damage. . . . .	61
2.15	Damage map for Inconel 740H based on results from SRLIFE. Contour lines from 5–80 years are shown. The upper right shaded area is the corrosion region (C.R.). Stress reset does not occur for Inconel 740H within this parametric space. The various color levels denote regions of different ratios of fatigue to creep damage. . . . .	62
2.16	Damage map for SS 316H based on results from SRLIFE. Contour lines from 20–80 years are shown. The red shaded area represents the stress reset region. Color levels denote regions of different ratios of fatigue to creep damage. . . . .	63
2.17	Damage map for Incoloy 800H based on results from SRLIFE. Contour lines from 10–80 years are shown. The red shaded area represents the stress reset region. Color levels denote regions of different ratios of fatigue to creep damage. . . . .	64
2.18	Thermal gradients on a Haynes 230 tube cross section simulated in SRLIFE at a $\Delta T=190$ (C) and $T_f=425$ (C). Fig. 2.18a shows results from a conduction-convection ratio of 0.25, and Fig. 2.18b has a ratio of 0.75. .	66
2.19	Von Mises elastic stress plot on an Haynes 230 tube cross section simulated in SRLIFE at a $\Delta T=190$ (C) and $T_f=425$ (C). Fig. 2.19a shows results from a conduction-convection ratio of 0.25, and Fig. 2.19b has a ratio of 0.75. . . . .	66

2.20	Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Haynes 230 material. Only lifetime contours are shown, and stress reset limits are omitted. . . . .	68
2.21	Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Inconel 617 material. Only lifetime contours are shown, and stress reset limits are omitted. . . . .	69
2.22	Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Inconel 740H. Only lifetime contours are shown, and stress reset limits are omitted. . . . .	70
2.23	Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Haynes 282 material. Only lifetime contours are shown, and stress reset limits are omitted. . . . .	71
2.24	Crown stresses for a 20 day FEA simulation of a Haynes 230 cross section via SRLIFE. Every cycle experiences a maximum temperature difference of 140°C and fluid temperature is at 525°C. Dots denote the first and last day's peak operational stresses. . . . .	73
2.25	Difference between SRLIFE-predicted and Norton-Bailey-predicted stress relaxation for a range of temperatures. FEA simulation considered 20 cycles of thermal loading typical of CSP applications. Daily operation time was 12 hours. . . . .	74
2.26	A sine wave approximation of the temperature difference's temporal variation compared to tubes at the south (0), east (10) and north (20) locations around the receiver for a solstice day operation assuming ideal image size priority aiming. . . . .	79
3.1	Incident flux map for a Crescent Dunes cylindrical receiver. Heliostat field aiming with image size priority is used. The modeled material is Inconel 740H. . . . .	93
3.2	Total temperature difference at the crown at 50 axial positions per tube length for a Crescent Dunes cylindrical receiver, Heliostat field aiming uses image size priority. The modeled material is Inconel 740H. . . . .	94

3.3	Salt fluid temperature profile for a Crescent Dunes cylindrical receiver, Heliostat field aiming with image size priority. The modeled material is Inconel 740H. . . . .	94
3.4	Lifetime profile a Crescent Dunes cylindrical receiver when default <i>SolarPILOT</i> strategy is used. The thermal and damage model material is Haynes 230. Note that the colorbar is saturated at 100 years. . . . .	95
3.5	Lifetime profile a Crescent Dunes cylindrical receiver when default <i>SolarPILOT</i> strategy is used. Profile is made using an Inconel 617 damage model to interpret a thermal solution common to Figures 3.4 and 3.7. Note that the colorbar is saturated at 100 years. . . . .	95
3.6	Lifetime profile a Crescent Dunes cylindrical receiver when default <i>SolarPILOT</i> strategy is used. Profile is made using an Inconel 740H thermal and damage models. Note the log scale. . . . .	96
3.7	Lifetime profile a Crescent Dunes cylindrical receiver when default <i>SolarPILOT</i> strategy is used. Profile is made using an Haynes A282 damage model to interpret a thermal solution common to Figures 3.4 and 3.5. Note the log scale. . . . .	96
3.8	All of the evaluated thermal operating points of the receiver (red triangles) overlaid on the damage maps presented in Chapter 2 for Haynes 230 (a) and 740H (b). . . . .	97
3.9	Crown stress versus time for a Haynes 230 receiver tube simulated for 30 cycles in SRLIFE. Finite element analysis accounts for elastic and inelastic strains caused by the loading conditions. The initial 20 days and the final 5 days have identical thermal loading cycles, but cycles 21-25 had higher, more extreme loading profiles. . . . .	101
3.10	Maximum stress at the crown location for each day cycle of a 30 cycle FEA simulation. . . . .	102
3.11	Relationship between the difference of elastic and actual peak stresses versus the residual stress for a series of FEA results data. As expected, the majority of cases follow a linear trend. . . . .	103



3.12	The input total temperature ( 3.12a) and resulting actual stress ( 3.12b) for an arbitrary thermal case. . . . .	104
3.13	The resulting instantaneous damage rates ( 3.13a) and lifetimes ( 3.13b) from the thermal case shown in Figure 3.12. . . . .	104
3.14	High-level process diagram of thermal model contrasting the workflow differences between using the PySSC (dashed line) versus <i>CoPylot</i> API (solid line) methods for generating the incident flux map. . . . .	108
3.15	Demonstration of <i>CoPylot</i> integration for a flat receiver: The input user flux profile requested a high, uniform flux on the left side and low, uniform flux on the right side. . . . .	109
3.16	Thermal model results showing the effects of a heavy left side flux concentration. This simulation is for a single flow path, flat receiver case at design point. . . . .	110
3.17	Resolution study results showing peak flux (Fig. 3.17a) and incident power (Fig. 3.17b) as a function of grid resolution for a 270 m <sup>2</sup> case and a square grid scheme. . . . .	112
3.18	Resolution study results for minimum lifetimes normalized by the largest lifetime observed in the study for that tube. The x axis is the ratio of thermal model axial nodes to flux profile nodes in the axial direction. Figure 3.18a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.18b) shows results with 74 points per dimension. Both studies assumed a 270 m <sup>2</sup> case and a square grid scheme. . . . .	114
3.19	Resolution study results for absorbed thermal power. The x axis is the ratio of thermal model axial nodes to flux profile nodes in the axial direction. Figure 3.19a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.19b) shows results with 74 points per dimension. Both studies assumed a 270 m <sup>2</sup> case and a square grid scheme. . . . .	115

3.20	Resolution study results for minimum tube lifetimes in each of 18 simulated receiver panels. Each panel value is normalized by the maximum lifetime that panel had during this study. Figure 3.20a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.20b) shows results with 74 points per dimension. Both studies assumed a 270 m <sup>2</sup> case and a square grid scheme. Outer tube diameter is 50.8 mm.	116
3.21	Resolution study results for absorbed thermal power. Figure 3.21a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.21b) shows results with 74 points per dimension. Both studies assumed a 270 m <sup>2</sup> case and a square grid scheme.	116
3.22	Resolution study results for each tested panels' minimum lifetime normalized by that panel's highest magnitude value during the study (Fig. 3.22a) and absorbed thermal power (Figure 3.22b). Results are for flux profile input resolution of 74 points per dimension.	118
3.23	Resolution study results for each tested panels' minimum lifetime normalized by that panel's highest magnitude value during the study (Fig. 3.23a) and absorbed thermal power (Figure 3.23b). Results are for flux profile input resolution of 74 points per dimension.	119
3.24	13 simulated tube lifetimes normalized by each tube's highest minimum lifetime value. The modeled receiver has panel widths of 1 m and the input flux profile has 74 points per dimension.	120
3.25	Ideal fluxmap in reference to the discrete flux levels between each panel (3.25a) and tube lifetimes normalized by the highest minimum tube lifetime in that tube (3.25b).	123
4.1	Intermediate step in visualization process showing each flowpath section on a series of 14 panels. A sample location (star) has flux corresponding to the flux intensity defined by the nearest upstream section corner (dot).	137
4.2	Final visualization of generating an ideal flux map for example conditions described. The zero flux regions at the end of each flowpath represent unnecessary receiver area.	138

4.3	Flux grid profile for flux map shown in Fig. 4.2. A grid of 14 x 14 points is used. . . . .	138
4.4	Ideal flux map for a receiver with 15 m height and 1 m panel widths. The actual receiver dimensions exceed the required number of panels, so the last two panels on each side are left blank and will be treated as zero flux zones. . . . .	140
4.5	Resultant thermal operating points (4.5a) output from full model simulation with a flux grid input corresponding to the flux map shown in Fig. 4.4. Calculated lifetimes for each axial sample point are shown in (4.5b). Note that the color bar maximum is 100 years, but some lifetimes exceed that maximum. . . . .	141
4.6	Minimum area for a range of design powers, assuming a 30 year lifetime, Haynes 230 tube material, and approximately 20 parallel tubes per panel.	143
4.7	Flux profile output from <i>SolarPILOT</i> of two heliostats along the solar field midline aimed at the center of the receiver. The heliostats are 128 m (4.7a) and 858 m (4.7b) from the receiver. . . . .	145
4.8	An objective flux profile input to <i>SolarPILOT</i> ( 4.8a), and the result ( 4.8b) for a 220 MW <sub>th</sub> , 270 m <sup>2</sup> case. . . . .	147
4.9	Lifetime profile for 220 M <sub>th</sub> case shown, where receiver flux is as shown in Fig. 4.8b. . . . .	148
4.10	An objective flux profile input to <i>SolarPILOT</i> ( 4.10a), and the resulting lifetime profile ( 4.10b) for a 220 MW <sub>th</sub> , 270 m <sup>2</sup> case. . . . .	149
4.11	An objective flux profile input to <i>SolarPILOT</i> ( 4.11a) using the maximize receiver lifetime and offset functions, and the result ( 4.11b) for a 220 MW <sub>th</sub> , 270 m <sup>2</sup> case. . . . .	150
4.12	Lifetime profile for 220 M <sub>th</sub> case shown, where receiver flux is as shown in Fig. 4.11b. . . . .	151
4.13	Process diagram for informed aiming heuristic, which seeks to consecutively aim heliostats in the least damaging locations until the target outlet temperature is reached in all receiver flowpaths. . . . .	154

4.14	An example image from a default template heliostat in <i>SolarPILOT</i> , with mirror height and width of 12.2 m. Heliostat is from a generated layout capable of 140 MW <sub>th</sub> . . . . .	157
4.15	Minimum lifetimes over a range of test design powers for uniform, informed heuristic, and informed <i>SolarPILOT</i> aiming methods. . . . .	163
4.16	Informed flux grid input to <i>SolarPILOT</i> for a 0.23 aspect ratio, 170 MW <sub>th</sub> case. . . . .	172
4.17	Informed flux grid input to <i>SolarPILOT</i> for a 2.70 aspect ratio, 170 MW <sub>th</sub> case. . . . .	173
4.18	Sample layouts for 170 MW <sub>th</sub> power and a receiver with aspect ratio of 0.23 (a) and 2.70 (b). . . . .	174
4.19	Average direct normal irradiation per season at Daggett, California location, using a typical meteorological year of data. . . . .	176
4.20	Model results, curve fit, and interpolated receiver efficiencies for an example summer solstice day on a typical flat receiver subject to uniform aiming. . . . .	177
4.21	Field efficiency results from each simulated time point in the example summer solstice day for a flat receiver subject to uniform aiming as well as interpolated values at intermediate time points. . . . .	178
4.22	Model-predicted power to fluid profile for a summer solstice day for design parameter case described in text. Interpolated power to fluid is based on receiver and field efficiencies derived from the model results and the average summer DNI profile shown in Fig. 4.19. . . . .	179
4.23	Pumping power plotted against power to fluid for the summer solstice simulation. Parabolic curve fit and interpolated values at un-simulated power to fluid points are also shown. . . . .	181
4.24	Capital and total lifetime operating expenses for an aspect ratio study at 170 MW <sub>th</sub> . (a) shows results when uniform aiming is used, and (b) when informed <i>SolarPILOT</i> aiming is used. . . . .	184
4.25	Total number of required panel replacements for uniform and informed <i>SolarPILOT</i> method aiming cases when design power is 170 MW <sub>th</sub> . . .	185

4.26	Thermal operating points overlaid on the Haynes 230 damage map for uniform aiming on a receiver with 0.23 aspect ratio at 170 MW <sub>th</sub> . . . .	187
4.27	Thermal operating points overlaid on the Haynes 230 damage map for uniform aiming on a receiver with 2.70 aspect ratio at 170 MW <sub>th</sub> . . . .	188
4.28	Thermal operating points overlaid on the Haynes 230 damage map for improved <i>SolarPILOT</i> aiming on a receiver with 0.23 aspect ratio at 170 MW <sub>th</sub> . . . . .	189
4.29	Thermal operating points overlaid on the Haynes 230 damage map for improved <i>SolarPILOT</i> aiming on a receiver with 2.70 aspect ratio at 170 MW <sub>th</sub> . . . . .	190
4.30	Design point efficiency and net energy for all aspect ratios and design power of 170 MW <sub>th</sub> . Edge offset is 2 standard deviations for both aiming methods. . . . .	190
4.31	Levelised cost of heat for uniform and informed <i>SolarPILOT</i> aiming methods for 170 MW <sub>th</sub> . . . . .	191
4.32	Total number of required panel replacements for uniform and informed <i>SolarPILOT</i> method aiming cases when design power is 200 MW <sub>th</sub> . . .	192
4.33	Levelized cost of heat for two different aiming methods for 200 MW <sub>th</sub> design power. . . . .	193

## ABSTRACT

---

Concentrating Solar Power tower receivers are subject to high flux conditions during operation, resulting in creep and fatigue. Damage incurred during operation can limit the lifetime of receiver tubes in molten salt systems if not properly understood or managed. Therefore, solar receivers are faced with two options: oversize the receiver to maintain reliability in worst-case flux conditions, or risk costly panel replacements during the receiver's lifetime. Current work develops a creep-fatigue damage tool by performing large parametric studies with the FEA-based receiver structural analysis tool SRLIFE. The comprehensive methodology is novel in its in-depth modeling of creep relaxation and extensive usage of FEA, which results in a high-fidelity damage model. This damage model is visualized using damage maps that characterize receiver damage in terms of two readily determined thermal characteristics: temperature difference and fluid temperature. Large sets of FEA results are obtained, and a software tool is developed for rapid assessment of receiver lifetimes. Lifetime contours are reframed as 'contour functions' that link total allowable temperature difference to fluid temperature. Locally varying flux profiles are shown to require significantly less receiver area compared to homogenous flux profiles. Two aiming methods are presented that achieve significantly higher reliability compared to a uniform aiming baseline. Conservative estimates suggest that consideration of thermal and lifetime profiles while placing images enables achievement of at least 20 MW<sub>th</sub> higher design power compared to a simpler flux profile matching method for a flat receiver of 270 m<sup>2</sup>. A levelized cost of heat analysis is performed for two aiming methods and shows that the present work's aiming method reduces LCOH by up to 7.6% compared to uniform aiming for the 200 MW<sub>th</sub> design power case described.

## 1 INTRODUCTION

---

A person's daily life is strongly shaped by energy consumption. For example, humans depend on process heat to prepare meals, fuel or electricity to power their automobiles, and electricity to operate household appliances. Energy in the form of electricity, process heat, or fuel is considered to be a key factor of quality of life and economic development. Generally speaking, the present work aims to contribute to the field of power generation, which involves converting a heat source into usable electricity. It is certainly desirable that power is low-cost, dependable, and is generated in a way that minimally affects the surrounding environment. The present work seeks to ensure that these three desirable aspects of energy are not mutually exclusive.

Power generation in the United States is sorted into three categories, according to the Energy Information Administration: nuclear energy, fossil fuels, and renewable energy EIA (2024a). Although a detailed discussion of energy policy is outside the present scope, this work focuses on renewable energy technology development for several reasons.

Renewable energies are appealing because of their potential to achieve extremely low costs of generation. All three categories of generation have associated capital expenditures. However, fossil fuel and nuclear technologies also have an inherent fuel cost associated with generation. In contrast, every major form of renewable energy generation except bioenergy does not require fuel to generate power. These renewable technologies have high potential for low cost generation, therefore, because their cost of generation is only dependent on the cost of required operations and maintenance. Another appealing attribute of renewable energy is its low emissions during generation. Most of a renewable technology's undesirable pollutants and greenhouse gases are produced during the plant's manufacture and are independent of the quantity of generation hours the renewable plant operates. Fuel-independent generation also means that the cost of generation is unaffected by fuel volatility. As a result, renewable energy is a power generation asset that ensures resilience in regards to foreign economies and energy supply lines.

In the United States, renewable electricity generation shares are already at 21% and expected to increase EIA (2024b). Present work primarily involves solar energy, which is considered a subset of the renewables category. The massive potential of solar energy has prompted significant research and commercialization efforts for the collection and storage of solar energy. Electricity generation in the United States was around 3.8 PWh in 2021 EIA (2024a), yet the available solar resource far surpasses yearly U.S. electricity needs. As reference, M. El-Wakil estimates that the United States receives 22 EWh/year of terrestrial radiation El-Wakil (1984). Two promising solar energy technologies that already have achieved gigawatts of installed generation capacity in the United States are: photovoltaics (PV) and concentrating solar power (CSP). These two technologies have achieved large-scale development, and global solar capacity is expected to continue growing.

At the end of 2023, solar generation capacity was approximately 7.6% (91.3 GW), and solar generation accounted for 3.9% (165 TWh) of total generation in the United States EIA (2024a). Recent rapid growth of photovoltaic solar can sometimes overshadow other viable forms of solar generation, such as concentrating solar power (CSP). However, in 2023 global installed CSP capacity exceeded 7.4 GW Alami et al. (2023), with additional CSP projects underway. The installed capacity of CSP is expected to more than double in the future, with 8.1 GW of CSP capacity in development in China REN21 (2025). In addition to already achieving GW-scale capacity, CSP has the potential to further grow as power grid needs evolve.

Despite their successful large-scale implementation, it is common knowledge that solar energy is subject to intermittency problems due to variable weather conditions and limited hours of sunlight. Significantly increasing the share of solar energy generation on the grid while retaining grid reliability will require mitigation of intermittency in future solar generation. Ultimately, future solar must become *dispatchable* in order to achieve improved reliability. Another measure of intermittency is *capacity factor*, which is defined as the ratio of actual energy generation to rated energy generation for a specified amount of time. In the United States for the year 2024, PV solar achieved a capacity factor of 23.2% while thermal solar operated at 25%. Figure 1.1 shows global capacity factor trends as well as cost



of energy for both PV and CSP solar. It can be seen that CSP on the global scale generally operates at a higher capacity factor than PV, and also significantly higher than the capacity factor observed locally in the United States. The higher capacity factors of CSP can be primarily attributed to the widespread integration of CSP plants with thermal storage. For example, the Gemasolar CSP plant in Spain is designed for 15 hours of production with a nominal 55% capacity factor SENER (2025).

Integrating solar generation with energy storage is crucial for maintaining future grid reliability. PV converts solar energy directly into electricity and is naturally combined with battery energy storage systems (BESS), while CSP plants first collect solar energy in the form of thermal energy which is readily stored in the form of sensible or latent heat. If a thermal efficiency of 40% is assumed, thermal storage with molten salt technology costs approximately 75 (\$/kWhe) based on cost estimates in Zurita et al. (2020); Glatzmaier (2011). The cost of thermal energy storage (TES) is projected to be significantly lower than BESS, for which the lowest projected cost for a 4 hour unit by 2030 is 245 (\$/kWhe) according to Cole and Karmaker (2023). Energy storage systems enable more stable and efficient grids by mitigating intermittency effects and reducing price volatility Ölmez et al. (2024). With this in mind, CSP technology continues to be relevant because of its natural pairing with thermal energy storage (TES). Since CSP generation works by converting thermal energy into mechanical energy, energy storage can be readily integrated on the power block side to allow for storage of excess thermal energy.

Levelised cost of electricity (units of \$/kWh) is defined as the sum of lifetime costs divided by the sum of electrical energy produced over that lifetime. Figure 1.1 shows that the levelised cost of energy for both photovoltaic and concentrating solar has generally decreased over the last 13 years. However, the LCOE of CSP is currently higher than that of PV, which has seen a dramatic decrease in cost. LCOE therefore considers a power plant's capital cost, operation cost, and generation NREL (2024). CSP's must continue to decrease to guarantee competitiveness with other generation types, and current goals for LCOE reduction for CSP technologies are 0.05 \$/kWh by 2030 Murphy et al. (2019).




Figure 1.1: Global, annual trends in capacity factor and levelized cost of energy for concentrating solar power and photovoltaics. Graphic and data from IRENA (2024).

To summarize, renewable energy is an appealing electricity generation technology because of its inherently low cost of generation, low emissions, and independence from fuel supply chains. Solar energy is one of the most prominent forms of renewable energy, and it has enormous potential to help meet rising electricity demand in the United States. Unfortunately, solar energy is subject to natural variability and requires energy storage solutions in order to address its reliability challenges. Concentrating solar power, although less well known compared to photovoltaic solar, is readily integrated with thermal storage that is expected to be more cost effective than battery storage. A significant barrier to higher implementation of CSP is cost — the levelised cost of electricity for CSP is higher than that of photovoltaics. The present work aims to lower the cost of CSP by reducing the system's capital and operating costs for a given lifetime.

Reducing the number of unexpected tube failures is critical for reducing plant operating costs. Furthermore, accurately determining maximum allowable flux is an important step for fully utilizing existing receivers and also for reducing the

required area of potential receivers.

## 1.1 CSP Technology Background

Concentrating solar power uses mirrors to focus direct normal irradiation onto a receiver, thereby achieving a flux that is higher intensity than the intensity of unfocused sunlight. The concentrated flux heats up a fluid, which can be used in a power cycle. A variety of receiver designs and configurations exist; The two most commonly used receiver technologies are parabolic trough and power tower Alami et al. (2023). A majority of CSP project capacity currently under development plans to use power tower receivers, which can achieve higher temperatures compared to parabolic trough receivers Achkari and El Fadar (2020).

A concentrating solar power generation plant is composed of a solar field, a receiver, thermal storage, and a power block. Figure 1.2, an illustration made by NREL (2018) is included as a general concept of the typical CSP system considered in this work. A collector field of heliostats focus direct normal irradiation onto a receiver. Figure 1.2 includes a cylindrical receiver, which is a type of external receiver. Tower pumps move salt from the cold tank (orange) through the receiver, which heats the salt up to the desired outlet temperature. Hot salt (red) is pumped from the thermal storage system into the steam generation system to facilitate electricity generation. Note that the existence of salt tanks means that solar collection can be operated at independent times from operation of the steam generation system and power block. Hot salt is used to preheat, boil, superheat, and reheat steam for turbine use, and is then returned to the cold salt tank to restart the process. The power block shown uses a Rankine cycle, with steam production powering a turbine for electricity generation. The entire system's LCOE by improving the performance of each subsystem. Present work seeks to advance the operation of the collector field and tower design. Collector field and tower system performance is quantified using the levelised cost of heat (LCOH) metric, which is generally defined as the capital and operating expenditures of the collector field and the tower divided by the energy stored in the heat transfer fluid. A reduction of LCOH




Figure 1.2: Rankine cycle power block, steam generation system, and nitrate salt power tower design illustration made by NREL (2018)

results in a reduction of the plant's LCOE.

Tower receivers are typically comprised of multiple panels that each contain a large number of tubes Conroy et al. (2020), and an example of a tower receiver is shown in Fig. 1.3. Concentrated flux ultimately interfaces with some heat transfer medium in the receiver. Receivers may be used to heat up solid, liquid, or gas. Liquid tubular receivers are the most popular technology for power towers Conroy et al. (2020). We further break the category of liquid tubular receivers into internal (cavity) and external. Internal receivers focus light through an aperture, which has the potential to improve thermal efficiency of the receiver but also to incur lower image intercept efficiencies Conroy et al. (2020). The most common tube fluids used in operation are water/steam and nitrate salt, and sodium is also a tube fluid under current development. Present work focuses on external receivers that heat up nitrate salt from approximately 290 to 565 °C.

External receivers include flat (billboard) receivers, which are heated by concentrated flux from solar fields facing north or south, and cylindrical receivers,



Figure 1.3: Power tower for receiver testing with a solar field in the foreground at Plataforma Solar de Almeria in Spain. Photo taken by author.

which are heated by flux all around the receiver's circumference. In either case, the receiver is typically composed of a series of panels. The term 'panel' refers to a collection of parallel tubes attached to an upper and lower header. Figure 1.4 contains an illustration from literature showing a cylindrical receiver and a photo taken by the author that showing different test panel designs. A flowpath describes a collection of panels that are connected in series. Nitrate enters the flowpath at the temperature of cold salt and flows consecutively through each panel. The salt travels in parallel through the number tubes before getting collected and routed through the next panel.

External receiver performance is a critical aspect of the overall CSP system performance. Collector field efficiency is a function of receiver geometry and the receiver's maximum allowable flux concentrations. Receiver cost is also an important component of the total system cost. The International Renewable Energy Agency estimates that receiver costs in 2023 comprise 14% of the total system cost for power towers IRENA (2024). Receiver cost has been a significant factor in system cost for a long time — data from 2011 reports that 16% of system cost was attributed to the receiver IRENA (2024). It is crucial to decrease power tower receiver cost to enable further reduction of LCOE.

High performing receivers must be reliable, utilize their receiver area efficiently,




Figure 1.4: A sample cylindrical receiver (1.4a) from Conroy et al. (2020) and various panel designs (1.4b) from concentrating solar research at Sandia National Lab

and enable high optical efficiency of the collector field. Although there are many design and operation factors related to receiver reliability, this work specifically addresses the significant impact of panel lifetime reliability. Panel reliability depends on the reliability of its individual tubes. If a tube failure occurs, a panel replacement is required and solar collection operations must be halted until all flowpaths are complete again. Thus, a panel failure halts solar collection and therefore electricity generation as well as incurring additional operating expenditures in the form of panel replacements. A reliable receiver, as defined in this work, refers to one that

requires zero panel replacements during the lifetime of the plant. The present work provides more detail regarding how receiver tube failures occur in later sections. In general, tube lifetime is inversely related to the incident flux levels experienced during operation. Lower average incident flux on the receiver tubes typically improves the lifetime of receiver tubes. However, a larger receiver size is required to meet the same design thermal power if lower average incident flux is desired. This tradeoff leads to the second criteria of a high performing receiver, which is utilization. Incident flux levels of an ideally-utilized receiver perfectly correspond to the highest allowable flux levels that a tube can endure over the plant lifetime period. Receiver flux profiles that are lower than the maximum allowable flux levels require larger receiver areas than ideally-utilized receivers. In other words, an ideally-utilized receiver optimizes the tradeoff between reliability and required receiver area. Finally, optical efficiency of the collector field is also a function of the receiver size and allowable flux levels. Receiver performance improvements must always consider solar field efficiency because solar field costs are typically the most significant system costs IRENA (2024). Spillage occurs when heliostat image borders are not fully contained on the receiver area and the incident energy is therefore lost. Receiver size and the aiming strategy employed both impact spillage. Spillage is minimized by employing a large receiver and aiming every heliostat image at the center of the receiver. Unfortunately, this aiming and design strategy results in excessive flux levels at the center of the receiver and too low of flux levels everywhere else.

The interplay of reliability, receiver area utilization, and field efficiency forms a complex optimization challenge. Consequently, improving receiver performance requires us to: (i) quantify the relationship between reliability and incident flux (ii) establish methods to fully utilize available area for a given receiver lifetime (iii) maintain or improve solar field efficiency while achieving reliable, well-utilized receiver designs. The next section introduces the reader to the relationship between incident solar flux and tube/panel reliability.




## 1.2 Reliability Challenges of Tower Receivers


Receiver panels' tube reliability limits are directly limited to material-specific corrosion limits and creep-fatigue damage. The interaction between corrosion and creep-fatigue is an ongoing research topic that is beyond the scope of the present work. Instead, we use film temperature regions defined by literature to discuss creep-fatigue limits in context of anticipated corrosion limits. Present work investigates creep-fatigue damage in detail in Chapter 2. The purpose of the current section is to introduce the reader to the concepts of creep-fatigue damage and define the daily loading conditions of receivers that cause damage.

**Circumferential Variation of Flux on Receiver Tubes** Receiver tubes experience thermoelastic stress due to a temperature gradient induced by the absorbed flux profile around the tube circumference. The tube surface at any axial location transforms the incoming incident flux into an approximately cosinusoidal absorbed flux profile as shown in Fig. 1.5a. Absorbed flux is approximately a cosine function of the angle relative to the tube crown. Therefore, the tube crown is the only circumferential location on the tube anticipated to experience absorbed flux equal to the incoming incident flux. The absorbed flux profile causes a peak temperature difference at the tube crown which decreases to zero on the left and right sides of the tube. Tube backsides are assumed to be adiabatic. The resulting temperature profiles for two tubes in panels 0 and 7 of a modeled receiver that resembles the Crescent Dunes plant in Tonopah, NV are shown in Fig. 1.5b. Temperature profiles are determined using a detailed thermal model presented in Martinek et al. (2021). The inner and outer tube surface profiles are directly caused by the absorbed profile. We can define two temperature differences, a wall temperature difference caused by conduction ( $\Delta T_{\text{cond}}$ ) and an inner surface temperature difference caused by convection ( $\Delta T_{\text{conv}}$ ). The total temperature difference is defined as ( $\Delta T_{\text{total}}$ ) and can be used in combination with the fluid temperature to predict creep-fatigue damage by using the process described in Chapter 2. Lastly, the back wall of the tube is approximately adiabatic and therefore closely tracks the fluid temperature.





(a) Illustration of absorbed flux profile on a tube cross section. Created by Conroy et al. (2020)



(b) Inner (dashed) and outer (solid) surfaces of tube cross sections in two different receiver panels.

Figure 1.5

The resultant non-axisymmetric temperature gradients cause significant thermoelastic stresses in the receiver tubes. Typical temperature, stress, and strain profiles in a tube cross section for a single timepoint are shown in Fig. 1.6. Temperature conditions of  $T_f = 300^\circ\text{C}$  and  $\Delta T = 220^\circ\text{C}$  are input to a finite element analysis (FEA) software tool developed specifically for solar receiver tube analysis. The peak von Mises stress and strain occurs at the tube crown when only elastic conditions are considered. Spatially nonuniform receiver incident flux profiles are time-dependent and repeat daily. Therefore, the receiver tubes will experience cyclic stress and strain profiles that ultimately result in tube failures that require costly panel replacements that halt solar field production. Now we discuss how the thermoelastic stress strain varies during daily operation.




Figure 1.6: Thermoelastic results from SRLIFE for a Haynes 230 tube at conditions of  $T_f = 300^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$ .

**Temporal Variation of Flux and Stresses in Tubes** The total elastic stress experienced by a tube depends on the internal fluid pressure stress and thermoelastic stress in the tube wall. Pressure stress for molten salt receivers is typically lower in magnitude compared to the thermoelastic stress during operation. Total elastic stress, therefore, is expected to be a primary function of the incident flux profile on the receiver. For any given operational day, flux profiles mainly vary in time according to the level of direct normal irradiation available and the aiming strategy employed. Figure 1.7 shows hourly thermal results from a simulation of a Crescent Dunes type cylindrical receiver on a summer solstice clear sky day. Maximum absorbed flux on each panel varies with time of day according to the proximity to solar noon. Early morning and late evening periods are the times with the lowest DNI and incident flux on the receiver. The maximum tube stresses in each panel are also shown to be directly related to the maximum absorbed flux. Note that fluid temperature does not significantly vary throughout the day for panels near the flowpath beginning (6 & 7) or ending (0 & 13) because mass flow through each flowpath is controlled to achieve constant outlet temperature. Fluid temperature does vary throughout the day for some intermediate panels because of shifts in the receiver's flux distribution with changing solar position.

Absorbed flux cycles are generalized as gradual — starting at low magnitudes early in the day, peaking at a maximum around solar noon, and then decreasing

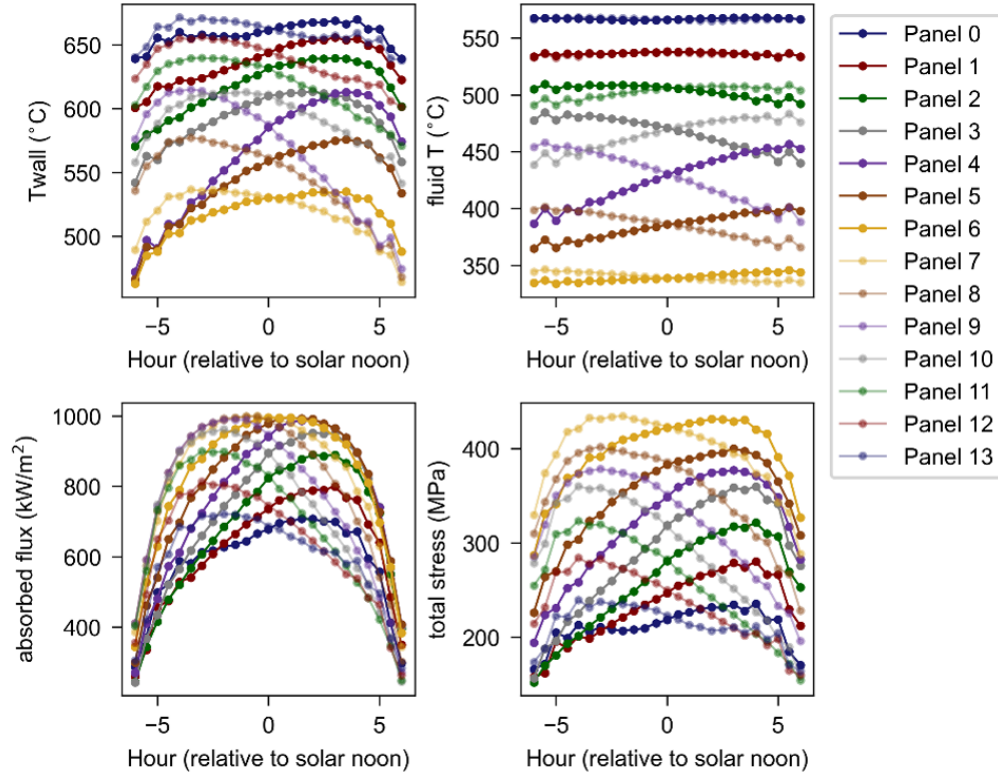


Figure 1.7: Maximum wall temperature, fluid temperature, absorbed flux, and total elastic stress for a Crescent Dunes solstice clear sky thermal model simulation. Results from each of the 14 panels are shown versus solar hour. Thermal model is from Martinek et al. (2021).

back down to zero in the evening again. The peak value is dependent on the panel, axial location on the tube, and the day of the year. It should also be noted that the temporal variance of a flux cycle is altered if aiming strategies change throughout the day. In other words, flux cycles for a given receiver location are not expected to be as gradual if the aiming strategy is constantly shifting the receiver's flux distribution throughout the day. Cloud events also cause departure from gradual flux cycles. Current work makes the problem more tractable by modeling creep-fatigue damage of clear-sky days that are defined by their peak absorbed flux, which also directly equates to a peak  $\Delta T_{\text{total}}$ .

The cycles that receiver tubes experience have long term effects on the tube

lifetimes. Incurred creep and fatigue damage both reduce the lifetime of a receiver tube, but have different causes. Fatigue damage is caused by cyclical stress and strain, while creep damage is caused by sustained stress at elevated material temperature. Thus, the magnitude of fatigue damage is dependent on the amplitude of change in strain over a given cycle, whereas creep damage increases with the time duration at a given stress and temperature. Most materials experience significant creep-fatigue interaction, meaning the combination of fatigue and creep loading is more damaging than either by itself. We develop a damage model in this work that quantifies the creep-fatigue damage associated with a given absorbed flux cycle.

### **1.3 Statement of Research Questions**

The present work's primary objective is to decrease the LCOE for tower receiver CSP systems. With that in mind, we seek to improve the reliability, utilization, and efficiency of an important subsystem: the receiver. Present work seeks to answer the following overarching research questions:

1. How do we model the creep-fatigue damage associated with a solar flux cycle that occurs on a receiver tube at a given fluid temperature and total temperature difference?
2. What are the relative magnitudes of creep and fatigue damage for six potential receiver alloys?
3. What is the relationship between alloy lifetime, fluid temperature, and total temperature difference for six potential receiver alloys?
4. For a given receiver, lifetime, and design power, what would the flux profile look like if the applied flux matched the allowable flux as closely as possible?
5. How can we inform aiming strategies to maximize tube lifetimes while meeting the design power?

6. What are the performance benefits of using an informed aiming strategy compared to a baseline uniform aiming strategy?
7. Is it possible for an aiming strategy to increase reliability and solar field efficiency simultaneously?
8. How does receiver aspect ratio impact the LCOH we can obtain if using a damage model informed aiming strategy?

The next section places the present work in the context of existing literature regarding these research questions.

## 1.4 Existing Literature

The following discussion of literature introduces work that serves as foundational knowledge for the present work and identifies existing gaps related to damage modeling, aiming, and design of tower receivers. This section reviews existing literature and organizes according to the following: (i) damage modeling of receiver tubes (ii) aiming strategies informed by tube damage estimates (iii) reducing tube damage via flowpath design.

**Damage Modeling in Literature** Tube damage modeling typically focuses on calculating the stress at the crown because it is considered to be the hottest, most stressed, and most damaged point on the tube. Literature that directly relates to the present work—determining feasible thermal conditions, and predicting tube lifetime via inelastic analysis—is discussed here, but detailed reviews of existing thermo-mechanical and damage modeling for external receiver tubes can also be found in Conroy et al. (2020); Hering et al. (2021).

Thermoelastic stress modeling of the receiver tube is a necessary first step for damage modeling, and several works have demonstrated how to apply existing techniques and evaluated the accuracy of various simplifying assumptions. Thermoelastic tube stress calculation methods are typically analytical Conroy et al.

(2020) as they drastically reduce computation time compared to FEA and can be leveraged for real-time applications, such as that in Laporte-Azcué et al. (2023). Methods for analytically determining elastic stress in solar receiver tubes can be found in or adapted from several sources including: Flores et al. (2014); Hetnarski et al. (2009); Neises et al. (2014); Ortega et al. (2015); Timoshenko and Goodier (1951); Laporte-Azcué et al. (2020); Goodier (1957); Bijlaard et al. (1968); Logie et al. (2018); Gatewood (1941); Carli (1984).

Elastic modeling does not account for relaxation due to creep or yielding, whereas the present paper accounts for these effects via inelastic material models. Inelastic modeling adds complexity and is typically only included in finite element analyses (FEA), but analytical approaches also exist in literature. When investigating high temperature molten chloride salt receivers, Barua et al. Barua et al. (2019) compared design by elastic, elastic-perfectly plastic, and inelastic analyses for 740H tubes, and Martinek et al. Martinek et al. (2021) assumed actual lifetimes to be 4-7 times longer than their elastic predictions based on previous multiday FEA results. A method of extending analytical elastic stress calculations to include the effects of yielding is presented in González-Gómez et al. (2021), where the two methods provided in Moftakhar et al. (1994); Glinka (1985) showed good agreement with FEA results for a thermal case of  $1 \text{ MW/m}^2$  flux density in A230 alloy.

Work in González-Gómez et al. (2021) also proposed using an analytical creep-strain model, the Norton-Bailey model, to predict the relaxed stress state that will result from creep. The Norton-Bailey model predicts stress relaxation for a given metal subject to an effective temperature, time, and initial stress. It was originally formulated for use in isothermal conditions at a constant stress González-Gómez et al. (2021), but work in González-Gómez et al. (2021) adapted the Norton-Bailey method by assuming each solar cycle could be represented by 1 hour at maximum temperature. Comparison to multiday FEA simulations demonstrated that implementing the Norton-Bailey model with an assumption of 30 hrs of relaxation time captures the trend of stress relaxation while being conservative González-Gómez et al. (2021), and use of the analytical Norton-Bailey model has proven useful for computationally iterative design studies such as those in Laporte-Azcué et al.

(2021a); Gentile et al. (2024). The present work also recognizes the importance of modeling stress relaxation caused by creep and therefore relies on FEA simulation of multiple cycles, which is considered a high accuracy approach.

Creep-fatigue damage modeling uses a tube's strain ranges and stresses profiles to predict the lifetime of a material. Some literature has chosen to neglect creep damage and instead solely consider fatigue damage as the primary failure mechanism Kistler (1987); Grossman et al. (1990); Liao et al. (2014); Luo et al. (2015). FEA results from Radke et al. (2014) predicted a fatigue-dominated failure mode, but recent creep-fatigue analyses have, in general, found creep to be the dominant failure mechanism. Table 1.1 provides a summary of existing literature sources that apply creep-fatigue damage modeling to solar receivers. Note that sCO<sub>2</sub> results are not considered in Table 1.1 because the pressure stress in sCO<sub>2</sub> tubes is significantly higher (20-25 MPa Neises et al. (2014)) than that seen in liquid receivers.

Understanding a receiver's ratio of creep to fatigue damage is important for operation and design, and Table 1.1 lists several sources that reported creep and fatigue magnitudes in their results. Damage modeling has been performed for several mass flow rate control, aiming strategies, and DNI condition cases as interest in power tower technology increases. The most commonly considered DNI case is clear-sky conditions, with authors finding fatigue to be negligible compared to creep in tubes made of A230 González-Gómez et al. (2021); Gentile et al. (2022); Laporte-Azcué et al. (2021a), 800H Gentile et al. (2022); Laporte-Azcué et al. (2021a), 740H Gentile et al. (2022); Laporte-Azcué et al. (2021a), and 316H Laporte-Azcué et al. (2021a). Analyses of receivers using sCO<sub>2</sub> Ortega et al. (2016); Nithyanandam and Pitchumani (2016); Neises et al. (2014) and high temperature chloride salt Barua and Messner (2023); Martinek et al. (2021); Barua et al. (2019) identified creep to be more significant than fatigue. Fatigue and creep damage in transient DNI cases depend on the mass flow control scheme. Fluctuating DNI/receiver flux will accumulate additional fatigue damage Laporte-Azcué et al. (2022). In the case of perfect mass flow control or higher than desired mass flow, temperature spikes are mitigated and the share of creep damage does not significantly increase compared to a clear sky case. These cases predict the highest fatigue damage, such

as in Laporte-Azcué et al. (2022), where, for a representative year of operation, the maximum ratio of fatigue to creep damage for a transient DNI case was 26% but only 18% for the clear sky case. In contrast, imperfect or lower than desired mass flow causes temperature spikes during unshadowing events Gentile et al. (2022); Toscani et al. (2018), thereby greatly increasing creep damage while only marginally increasing fatigue damage Gentile et al. (2022).

The present work seeks to understand the creep-fatigue relationship associated with primary operation. A thermoelastic finite element simulation in Uhlig et al. (2018) found stresses during maximum flux density operation to be significantly higher than those during preheating. Additionally, the work in Pérez-Álvarez et al. (2021) showed significant differences in tube temperature gradients for three different versions of the Vant-Hull preheating method Vant-Hull (2002). Therefore, the present work neglects preheating effects because they strongly depend on the employed preheating algorithm.



Table 1.1: Related literature in which methodology and results include creep-fatigue analysis and/or lifetime predictions.

Source	Year	Material	Type	Thermal Conditions	Structural Analysis	Creep Model	LTE/Damage Results
Berman and Rao (1983)	1983	800, 800H	H <sub>2</sub> O	clear day	FEA: elastic-creep	simulated 20 cycles	damages per cycle, LTE predictions
Fork et al. (2012)	2012	A617	air	steady, sunny, & cloudy days	FEA: includes inelastic	none	pass/fail LTE, creep & fatigue for several cases
Radke et al. (2014)	2014	-	nitrate salt	-	FEA: elastic and inelastic	included in FEA	LTE prediction, fatigue table
Conroy et al. (2018b)	2018	304SS	sodium	novel aiming strategy	analytical: elastic Logie et al. (2018)	none	LTE vs. spillage & aiming points
Conroy et al. (2019)	2019	304SS, 316H, 800H, A617, A230	sodium	flux det'd by spillage, 5 inlet-outlet ranges	analytical: elastic Logie et al. (2018)	none	allowable flux densities for 4 parameters

Table 1.1, continued

Source	Year	Material	Type	Thermal Conditions	Structural Analysis	Creep Model	LTE/Damage Results
Barua et al. (2019)	2019	740H	chloride salt	spring equinox	FEA: elastic, elastic-perfectly plastic & inelastic	calculated relaxation profile	min. tube LTEs from each modeling method
Martinek et al. (2021)	2021	740H	chloride salt	image size priority flux aiming, clear-sky days	analytical: elastic Logie et al. (2018) multiplied by adjustment factor	see note on structural analysis	flux & $\eta$ for various LTEs, creep-time plots
González-Gómez et al. (2021)	2021	A230	nitrate salt	clear sky day, flat aiming	analytical: elastic with Bijlaard et al. (1968); Laporte-Azcué et al. (2020); Goodier (1957) & plastic Mofakhar et al. (1994); Glinka (1985).	analytical: Norton-Bailey model	min. LTEs/panel and negligible fatigue

Table 1.1, continued

Source	Year	Material	Type	Thermal Conditions	Structural Analysis	Creep Model	LTE/Damage Results
Laporté-Azcué et al. (2021a)	2021	A230, 316H, A625, 740H, 800H	nitrate salt	clear-sky	analytical: see González-Gómez et al. (2021)	analytical: see González-Gómez et al. (2021)	LTE and LCOE for each material
Gentile et al. (2022)	2022	740H, 800H	nitrate salt	cloudy & clear days	analytical: see González-Gómez et al. (2021)	analytical: see González-Gómez et al. (2021)	damage & LTE per panel/case, showed critical pts.
Laporté-Azcué et al. (2022)	2022	A230	nitrate salt	cloudy & clear	analytical: see González-Gómez et al. (2021)	analytical: see González-Gómez et al. (2021)	creep & fatigue & LTEs per panel for various cases
Barua and Messner (2023)	2023	740H	chloride salt	spring equinox	FEA: elastic	three models used	creep, fatigue & LTEs results, LTE- $T_{f,out}$ plot

Table 1.1, continued

Source	Year	Material	Type	Thermal Conditions	Structural Analysis	Creep Model	LTE/Damage Results
Gentile et al. (2024)	2024	A230, 740H, 800H	nitrate salt	spring equinox, noon	analytical: see González- Gómez et al. (2021)	analytical: see González- Gómez et al. (2021)	optimization results that include LTE

**Damage-Informed Aiming in Literature** The simplest and most optically efficient aiming strategy is to aim every heliostats at the center of the receiver, but excessive peak flux results from this strategy. Therefore, the majority of aiming strategies in literature focus on homogenizing the heat flux over the receiver surface in order to ensure tube reliability Conroy et al. (2020). Aiming heuristic and optimization typically seek to minimize flux spillage while achieving acceptable reliability. It is common for aiming strategies to acknowledge tube damage by enforcing an allowable flux density, and the concept was introduced by Vant-Hull (2002) for receivers. This concept is well established, but typically is a predetermined global limit on the entire receiver. Two aiming strategies in literature that integrate thermomechanical constraints into the aiming algorithm (referred to as "closed loop" in Conroy et al. (2020)) at a local level and are discussed in more detail here.

Sánchez-González et al. (2020) developed maximum allowable flux models to advise an aiming strategy for Haynes 230 and Inconel 625 receivers. In that study, maximum allowable flux did vary locally and was dictated by the ultimate tensile strength and critical corrosion temperature of the material. Results showed that corrosion risk limited the allowable flux of high fluid temperature panels, while lower temperature panels were limited by thermal stress levels. This work demonstrated a methodology for employing local flux limits to maximize receiver power. Thermal stress limitations for a receiver are more informative if they are based on predicted creep-fatigue damage, however, because molten salt receiver tubes are typically assumed to fail due to creep-fatigue damage.

Conroy et al. (2018b) also developed a novel aiming strategy based on the simulated annealing optimization method. Results demonstrated that the aiming method is capable of limiting spillage and spreading flux evenly to achieve the design lifetimes for a sodium-cooled receiver case. Creep-fatigue damage modeling was performed to predict the receiver lifetime for each generated flux map, but stress modeling assumptions significantly differ from those used more recently in literature. Conroy et al. (2018b) assumes that a condition of zero axial force and annulled bending moment, which is shown by Logie et al. (2018) to predict lower stresses than a generalized plane strain deformation assumption. Regardless, the

authors did not investigate cases where localized flux limits are leveraged to reach the desired design power.

Lastly, an open source software titled *Heliostat Aimpoint and Layout Optimization Software* (HALOS) is available for determining the best aimpoint strategy for a CSP heliostat field Zolan et al. (2021). HALOS uses mixed-integer programming models to determine the aimpoint strategy that maximizes incident power (minimizes spillage) subject to flux constraints. For given layout, Zolan et al. (2021) designed HALOS to provide an aiming strategy that maximizes thermal power delivery to the receiver while ensuring that the flux profile obeys local flux limits. HALOS can optimize the aimpoint strategy subject to an array of flux limits that vary with receiver location, but the flux limits must be computed prior to optimization.

**Literature Related to Damage-Informed Flowpath Design** Whereas informed aiming revises the aiming of a given heliostat field on a receiver to achieve higher design power and reliability, informed design focuses on altering the receiver itself according to a given flux distribution. Table 1.1 includes literature that has incorporated damage modeling into receiver design methodology. In addition to those in Table 1.1, Binder et al. (2022) integrated consideration of interaction effects into design via so-called "interaction maps" that show the effects of safety factors on permissible strain and stress. The number of panels and the tube diameter were shown in Rodríguez-Sánchez et al. (2014) to impact several aspects of receiver design. Conroy et al. (2018a) presented a methodology to evaluate liquid sodium receiver designs through thermal performance, creep-fatigue, and pressure drop analyses. Authors performed an analysis of various receiver configurations and identified trends in key design parameters. Although the analysis framework can be generally applied, this work does not design the flowpath to leverage local flux limits.

A methodology that combines thermal, lifetime, and pressure drop analyses with a parametric optimization methodology is presented in Gentile et al. (2024). Haynes 230 and Inconel 740H are found to be the most viable materials. The optimal Haynes 230 design had a peak flux of  $1245 \text{ kW/m}^2$ , and the optimal Inconel 740H

design had a peak flux of  $1288 \text{ kW/m}^2$ . Findings from Gentile et al. (2024) favored lower receiver ratios of height to diameter to higher ones, and the optimized number of tubes per panel was in the range of 36–40 for each evaluated material. Similar to literature regarding aiming strategies, creep-fatigue damage is performed after other aspects of the system have been optimized.

Literature has previously identified general recommendations for receiver flow-paths which are often assumed to be best practice in receiver design. Rodriguez-Sanchez et al. (2015) studied 8 different flow pattern configurations for molten-salt cylindrical receivers and showed that, for the northern hemisphere conditions, configurations with flow exiting on the south side is best practice. Adding parallel flow paths are desirable in that they reduce the pressure drop for a given load. Another important observation made by Rodriguez-Sanchez et al. (2015) is that crossing flow paths is beneficial for balancing flux during uneven incident flux periods occurring in the morning and evening. Flowpath optimization work by Emerick et al. (2024) for a supercritical carbon dioxide in Spain had two interesting findings. Firstly, the authors proved that there is an optimal number of panels for the analyzed receiver, where additional panels causes excessive pressure drop but fewer panels reduces the fluid velocity and heat transfer coefficient. Secondly, the authors recommended that flowpath inlets and outlets be positioned to place low temperature fluid in lower flux and high temperature fluid in higher flux. This result highlights the importance of considering corrosion and or lifetime limits in receiver design; solely considering hydraulic and thermal performance of a receiver led the authors to a conclusion that contradicts a similar analysis by Rodriguez-Sanchez et al. (2015), where the authors considered film temperature limits and thermal stress.

**Literature Summary** In regards to damage modeling, most receiver tube creep-fatigue analyses in literature only consider elastic stress models, and finite element analyses that consider multiple cycles, inelastic strain, or stress relaxation are limited. Recent literature recognizes the importance of considering plastic strain but typically uses analytical equations to estimate the stress-relaxing effects of yielding and creep.

Many of the existing creep-fatigue analyses are case study based, only reporting lifetimes and the relative importance of fatigue and creep damage for specific operation and DNI conditions. The most commonly investigated alloys are Incoloy 800H, Inconel 617, Haynes 230 and Inconel 740H. Early literature focused on the fatigue aspect of solar tower operation, but the majority of recent receiver tube studies find that creep damage dominates fatigue damage.

A review of literature related to aiming strategies identifies a need for a closed-loop aiming methodology that dually considers spillage and reliability. The majority of aiming strategies seek to homogenize the flux over the receiver surface while minimizing spillage. Global allowable flux limits are commonly used, but these limits vary among literature, depending on the underlying factors considered. Examples of underlying factors considered in literature include corrosion limits, ratio of thermal stress to ultimate tensile strength and creep-fatigue damage. A limited number of aiming strategies implement local allowable flux limits, but underlying damage modeling methodology differs from the methodology proposed by the present work.

Literature discusses basic trends in receiver design that serve as important starting points. It is clear that increasing receiver area for a given design power generally improves reliability because the average required flux levels decrease. Also, parametric system optimization considering the cost of operation and manufacture relative to the annual heating prefers lower height-diameter ratios of 1–1.5, depending on the material. Selection of tube diameter affects salt velocity and the required pumping pressure. As a result, a tradeoff exists between improved heat transfer coefficients and higher parasitic pumping losses. Modifying the number of panels in each flowpath affects salt velocity by changing the number of tubes in parallel, total flowpath pressure drop, and the flux profiles relative to the flowpath sections. It is also common to configure flowpaths so that high flux occurs at low fluid temperatures and vice versa. Literature demonstrates the importance of performing receiver design in the context of the overall system and in consideration of as many logical flux limitations as possible. However, to the author's knowledge, literature does not typically discuss how aspect ratio and panel position relative to



the flux profile affect receiver reliability.

## 1.5 Primary Contributions and Novelty

The overarching motivation of this work is to reduce the cost of electricity associated with CSP power tower technology. We tangibly accomplish this by improving the receiver subsystem. Levelised cost of heat can be lowered by reducing capital costs, reducing operating expenses, or increasing the annual energy collected by a given receiver. This work makes a series of contributions related to these general points.

The present work improves creep-fatigue damage predictions for molten-salt power tower receivers by performing multi-cycle, parametric, FEA-based analyses for several high temperature alloys, thereby increasing reliability and better matching actual to ideal flux. These two contributions decrease required maintenance and downtime through improved reliability and reduce the required size of the receiver through improved flux distributions. This work presents a creep-fatigue damage model that is novel in both its FEA-based methodology and resulting damage maps that characterize tube lifetime over a broader range of thermal conditions than typically seen in literature for six different alloys. The damage model for Haynes 282 is especially unique, as this alloy has not been extensively studied for receiver applications.

The damage tools for design and real-time presented in this work are easy to use and provide rapid access to creep-fatigue damage values for a large range of fluid temperatures and total temperature differences. Our damage tool for design also accounts for corrosion and stress reset zones and is based on a database of more than 3700 FEA results for 6 different alloys. Novelty of the developed damage tool lies in its underlying damage model data and program versatility. Thermal operating conditions for a given receiver can be rapidly translated into lifetime predictions, which is a crucial capability for aiming and design method development. Additionally, the real-time prototype tool provides a new potential method to track state of health and lifetime even in varying conditions.

Informed aiming methods, written in Python code, define the ideal flux profile for a given receiver with an objective lifetime. This ideal flux profile definition is useful for assessing the maximum power a receiver can achieve subject to a given lifetime, determining the minimum number of panels required to reliably achieve the desired power, or providing input to aiming methods.

Two important aiming methods are presented in this work. First, an informed-image size priority method is presented that demonstrates significant improvement compared to baseline aiming. Second, an aiming heuristic is developed that integrates damage and thermal model tools to reduce damage for a required power, receiver, and heliostat field. This aiming heuristic pioneers the integration of live-updating receiver lifetime profiles to mitigate the effects of unideal optical conditions.

The present work also identifies an optimal range of aspect ratios for a given set of receiver parameters and use of the informed-image size priority method.

All study findings suggest that uniform flux profiles underutilize high flux panels and overutilize low flux panels. Therefore, results challenge the prevalent opinion in literature that it is desirable to achieve flux profile homogeneity. Instead, current work demonstrates how optical efficiency and reliability can be improved via early-process integration of creep-fatigue considerations into design and aiming.

The outcomes of this work enable improved assessment of receiver tube lifetimes, higher utilization of receiver area, higher achieved power per desired lifetime, robust aiming strategies, insights on design parameters' effects on reliability, and lower levelized system costs. Each of these outcomes directly enhance the performance of the receiver-solar field subsystem.

## 1.6 Document Overview

Chapter 2 presents the damage modeling methodology, and introduces the resulting damage map for Haynes 230, where lifetime is a function of fluid temperature and temperature difference. The presented methodology is applied to study five other receiver materials: Inconel 617, Inconel 740H, Incoloy 800H, SS 316H, and Haynes

282. Analysis is performed for all six alloys at three different ratios of conduction to convection temperature difference. The current work contributes large amounts of finite element analysis (FEA) simulation data for several different alloys to the existing literature. Damage modeling accuracy is improved and data is presented in a tractable format to enable better-informed receiver operation and design decisions.

Chapter 3 discusses software tools developed for this work as well as for future receiver design and operations. Three tools are introduced: the damage tool for design, the damage tool for real-time operation, and a modified version of the thermal model presented in Martinek et al. (2021). Sample results demonstrate good performance and highlight primary tool applications for development of tower receivers.

Chapter 4 presents a methodology for defining the ideal flux profile for a given design lifetime and panel width. A module of functions (called the **informed aiming** module) is introduced that uses the tools developed in Chapter 3 to inform aiming strategies for flat receivers. Informed aiming is executed in two aiming strategies and tested for realistic solar field-receiver scenarios to demonstrate functionality. First, we provide informed aiming input to *SolarPILOT*'s image size priority strategy and test for a range of design powers. Second, we develop a heuristic that integrates a damage tool and thermal model into its image placement algorithm. Results from testing the heuristic over a design power range show improved performance over baseline and the *SolarPILOT* image size priority method. Finally, a levelised cost of heat study is conducted over different aspect ratios while utilizing informed aiming. Aspect ratios less than 1 achieve the lowest LCOH and highest reliability for the informed aiming algorithm used.

## 2 DAMAGE MODELING OF CENTRAL RECEIVERS WITH SRLIFE

---

### 2.1 Introduction

Current work is motivated by broader questions regarding the impact of solar variability on CSP system operations, and one anticipated effect of solar variability is accelerated damage of the receiver tubes. This chapter establishes a model to evaluate creep-fatigue damage that quantifies the impact of the receiver tubes' thermal operating points on their lifetime. A thorough literature review of existing damage models is provided in Chapter 1. To summarize, most existing receiver damage model methods are analytical, only consider elastic stress models and only report lifetimes and ratios of fatigue to creep for specific operation and DNI conditions. The author's research at UW-Madison leverages high-throughput resources on campus to efficiently execute large batches of FEA simulations. Thus, the creep-fatigue damage model presented in this work is unique because it uses simulated multi-cycle loading cases with finite element analysis over a large range of molten salt operating temperatures.

The primary contributions of present damage modeling work are: (i) presenting an FEA-based methodology that leverages an existing receiver analysis tool (SRLIFE) to create a damage map – a useful decision-making tool for operations and/or design of receiver tubes that shows the impact of thermal gradient and fluid temperature on both the lifetime as well as the ratio of fatigue to creep damage for any given clear-sky thermal conditions (ii) applying this methodology to analyze Alloy 230 solar receiver tubes for a commercial molten-salt scenario located in Tonopah, Nevada. Damage is predicted via elastic-plastic-creep analysis that considers the effects of stress relaxation and temperature dependent properties (iii) conducting parametric studies that compare the lifetime performance of Inconel 617, Inconel 740H, Haynes 282, SS 316H, and Incoloy 800H when exposed to the same thermal conditions. (iv) identifying a lifetime decrease at intermediate temperature cases for Haynes 230 and SS 316H that is related to the interplay of relaxation and rupture strength. (v) showing the effects of alternative loading scenarios and thermal conditions on

damage maps.

The primary publication outcome from this work has been a journal paper titled: *Damage Modeling of Power Tower Receiver Tubes Using the SRLIFE Tool* Wenner et al. (2025), which is published in *Solar Energy*. The authors, in order, are: Jacob Wenner, Mark C. Messner (Argonne National Lab), and Michael J. Wagner. A copy of the journal paper is available upon request.

The remainder of the paper is organized as follows: A methodology for using a finite element-based solar receiver analysis tool (SRLIFE) to predict the inelastic, creep relaxed states is presented in Section 2.2. The methodology is applied for study of Crescent Dunes receiver tubes, and lifetimes as well as ratios of fatigue to creep damage are predicted in Section 2.3. The performance of five other potential receiver alloys is evaluated in Section 2.4. Receiver tube damage for different assumptions of material conductivity and convection heat transfer is investigated in Section 2.5 for the four receiver alloys that exhibit the best creep-fatigue performance in the initial study. Haynes 230 is the primary material of study for the last two sections. SRLIFE's method of predicting creep relaxation is compared to the analytical Norton-Bailey model in Section 2.6. Finally, Section 2.7 investigates the impact of cycle definition on lifetime predictions.

## 2.2 Methodology for Creep-Fatigue Analyses

The methodology objective is to characterize how creep, fatigue, and lifetime of tubes vary over the full range of thermal operating conditions encountered in a cylindrical receiver. A thermal model is used to determine the range of thermal operating conditions: the expected fluid temperatures and temperature differences as well how these conditions vary with respect to time and position. The defined thermal boundary condition cases are combined with an assumed tube geometry to serve as the input to the FEA simulation performed in SRLIFE, which is a tool for estimating the structural service life of tubular, panel solar receivers. Some of SRLIFE's modeling details will be repeated in this document for the reader's convenience, but a thorough documentation of SRLIFE's modeling methodology can be

found in Messner et al. (2022). Primary output of SRLIFE simulation includes strain range and stress timeseries profiles at the tube crown, which are then correlated to fatigue and creep damage using SRLIFE methods from the *damage* and *material* libraries. This process is applied for every fluid temperature and temperature difference within a parametric range, thereby generating a dataset of case-specific creep and fatigue damages. Each case's 3<sup>rd</sup> day fatigue damage and 20<sup>th</sup> day creep damage are assumed to represent future cycles, and the corresponding lifetime is calculated using SRLIFE *material* library methods. An overview of the process is shown in Fig. 2.1. The following subsections discuss each step in more detail.

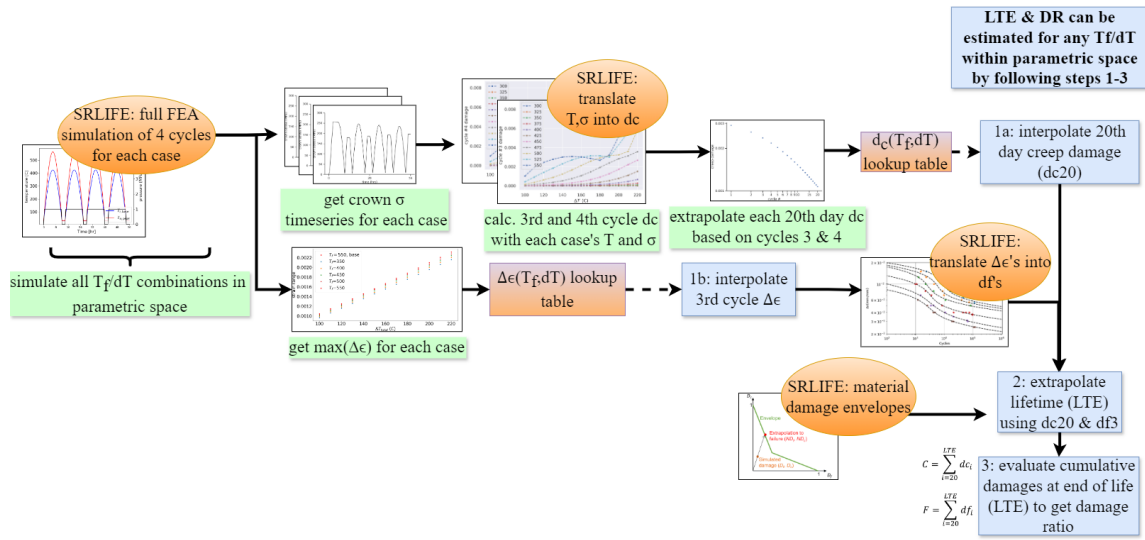


Figure 2.1: Analysis process diagram showing how creep and fatigue damages are calculated for each thermal case and material.

## Thermal Analysis

The detailed numerical thermal model discussed in Martinek et al. (2021) is implemented to formulate thermal loading cases. A design-point day is simulated, and the tube temperatures' temporal and circumferential variation serve as the basis for the thermal load assumptions made in this paper. This section discusses the key thermal results and how they are input into SRLIFE. All tube simulations in

SRLIFE use thermal boundary conditions that have an identical form in both time and position, but the magnitude of the peak flux depends on the operating point of the tube. Boundary conditions are intended to approximate a clear-sky day without any cloud events.

The thermal model predicts fluid and metal temperatures for a given flux profile input. In the present work, flux profiles on the receiver are generated with SolarPILOT Wagner and Wendelin (2018) using an image size priority aiming strategy. Image size priority is a heuristic used by SolarPILOT to achieve as uniform flux as possible while minimizing spillage losses. Images are placed on the receiver in order of image size on the receiver, and aim points are placed within an acceptable region based on their image size and positioning cutoff factor. For this work, the cutoff factor remains at the default setting of 2 standard deviations of the image being placed. The maximum simulated absorbed flux is  $1000 \text{ kW/m}^2$ . The cylindrical receiver has a diameter of 15.2 m, tube lengths of 18.59 m, and a tower optical height of 182 m. A Crescent Dunes case is simulated, with 14 panels split between a 2-path configuration that includes one stream crossover. The Haynes 230 receiver tubes have an outer diameter of 50.8 mm and thickness of 1.25 mm. The heat transfer fluid is nitrate salt ( $60\% \text{NaNO}_3$ ,  $40\% \text{KNO}_3$ ) with an inlet temperature of  $290^\circ\text{C}$ . The flux profile is from a design solstice day in Tonopah, Nevada, and salt flow rate is ideally controlled to maintain an outlet temperature of  $565^\circ\text{C}$  at each timestep. Heliostat layout generation is performed in SolarPILOT using a design power of  $565 \text{ MW}_{\text{th}}$ . The thermal model is at  $5 \times 79 \times 50$  radial, circumferential, and axial resolution per tube and is solved assuming steady state. Three tubes are simulated per panel. Radial and circumferential conduction are considered, but axial conduction is neglected. Radiative exchange considers adjacent tube reflections and the exchange of radiative energy from thermal emission in both the solar and infrared spectrum. Natural and forced convection coefficients are calculated using correlations from Siebers and Kraabel (1984). The present work assumes no significant wind, which would decrease tube temperatures. The backside of each tube is assumed to be adiabatic. A thorough discussion of modeling methodology is provided in Martinek et al. (2021). Sample results from panel 7 in Figures 2.2a

and 2.2b demonstrate the temporal and circumferential variation in inner and outer surface temperature profiles.

Figure 2.2b shows significant flux differences between the heated and unheated sides of the tube. Both the inside and outside surfaces of the tube will develop a circumferentially varying temperature profile that blends into an almost constant temperature on the back side of the tubes. There are two notable temperature differences in the tube cross section at a given axial position—the temperature difference between the outside and inside surface at the crown ( $\Delta T_{o,s}$ ) and the temperature difference between the inside surface at the crown and the backside surface ( $\Delta T_{i,s}$ ). The total temperature difference,  $\Delta T_{tot}$  is defined as  $\Delta T_{i,s} + \Delta T_{o,s}$  and illustrated in Figure 2.2b.

The tubes can be classified according to their total temperature gradient at the point of highest flux, the crown ( $90^\circ$ ), for easier implementation in parametric design studies. This classification approach is investigated in Montoya et al. (2019), which used the methodology in Montoya Sancha et al. (2018a) to propose identifying the tube with maximum stress based on the location of maximum circumferential temperature gradient. Figure 2.3a shows the maximum  $\Delta T_{tot}$  predicted by the thermal model in each of three representative tubes from the 14 panels. Panels 6 and 7 experience the highest flux and therefore have the highest  $\Delta T_{tot}$ . The plotted results are almost perfectly mirrored because panels 0–6 represent one salt flow path and 7–13 represent the second salt flow path. Results predict that typical design point maximum  $\Delta T_{tot}$ 's will be in the range 100–220°C. The conduction ratio ( $r = \Delta T_{o,s} / \Delta T_{i,s}$ ) at the design point is also calculated and shown in Fig. 2.3b on a per tube, per panel basis like in Fig. 2.3a. Results shown in Fig. 2.3b are within the range  $0.49 \pm 0.06$ . This conduction ratio depends on the salt flow's heat transfer coefficient, incident flux, conductivity of the metal, and thermal receiver efficiency. Thus, the actual conduction ratio varies with time of day, material, and aiming strategy. For Haynes 230, the averaged ratio for all tubes over all operation points of a clear sky solstice day is  $R = 0.45$  for the Crescent Dunes case.

This paper takes  $r$  to be constant in the SRLIFE boundary conditions, which enables  $\Delta T_{o,s}$  and  $\Delta T_{i,s}$  to be completely defined by  $\Delta T_{tot}$ . This resulting piecewise function



is shown in Eqns. 2.1 and 2.2, where  $T_f(t)$  is the fluid temperature,  $\Delta T_{i,s}(t)$  and  $\Delta T_{tot}(t)$  are the temperature differences at the crown. Figure 2.4b shows how SRLIFE tube surface temperatures are set to vary along the circumference of the tube for any given timestep. Note that specifying the inner and outer surface temperatures is very similar to applying a specified heat flux and assuming a constant value for the heat transfer coefficient.

$$T_i(\theta, t) = \begin{cases} T_f(t) + \Delta T_{i,s}(t)\sin(\theta) & 0^\circ \leq \theta \leq 180^\circ \\ T_f(t) & 180^\circ < \theta \leq 360^\circ \end{cases} \quad (2.1)$$

$$T_o(\theta, t) = \begin{cases} T_f(t) + \Delta T_{tot}(t)\sin(\theta) & 0^\circ \leq \theta \leq 180^\circ \\ T_f(t) & 180^\circ < \theta \leq 360^\circ \end{cases} \quad (2.2)$$


Thermal model results in Fig. 2.2a demonstrate that the temperature profiles of the heated side vary according to the operational hour. On the heated side, the overall surface temperature profile remains the same during steady operation, but the peak will gradually increase to a maximum at solar noon (hour 6) and then decrease until hour 12. Other panels in the receiver follow the same trends but may not experience time-dependent flux that is symmetrically centered about solar noon, as peak solar flux may occur before or after solar noon depending on the direction that the panel is facing.

Since the  $\Delta T$  for both surfaces is time dependent, they are defined in SRLIFE according to Eqn. 2.3, where the  $\Delta T$  in the right side of Eqn. 2.3 refers to either the total design temperature difference or the inner surface design temperature difference (see Fig. 2.2b for reference). These two variables will both reach their peak values at the same timestep. The  $\Delta T$  only activates during the operation period, when flux is gradually increasing and decreasing. The temperature of the unheated region at the inner wall is almost constant and is well-approximated using the fluid temperature. Thus, the unheated side only varies during preheat and shutdown as defined in Eqn. 2.4, where  $T_{f,op}$  is the fluid temperature during the operation period, and  $T_b$  is the base temperature of the plant during offtime,


assumed to be 30°C. Normalized inner and outer crown temperatures as well as an unheated location (270°) are shown in Fig. 2.4a to illustrate SRLIFE inputs. The temperature gradients are less severe at timesteps closer to the start of operation (hour 1) and are highest at solar noon (hour 6.5). Note that solar noon is shifted 0.5 hrs in SRLIFE because of the added preheat time.

Equation 2.3 states that the total temperature difference (and the inner surface temperature difference) variance with respect to time is approximated by a sine wave during operation. It is important to note that, for cylindrical receivers, the daily  $\Delta T$  profile experienced by a tube may differ from a sine wave approximation, depending on the tube's location. For example, some tubes will experience  $\Delta T$  timeseries profiles that more closely resemble a square wave. A single representative approximation is chosen to enable comparison solely on the basis of their fluid temperature and total temperature difference rather than the tube's individual cycle characteristics. An error analysis between model-predicted  $\Delta T$  profiles and sine-approximated profiles for Haynes 230 with the simulation conditions defined in Section 2.2 is conducted. The peak magnitudes of  $\Delta T$  and  $T_f$  considered are selected to correspond to cases with a predicted lifetime of approximately 30 years. For the cases considered, the most damaging tube profiles result in no more than 43.5% lower lifetimes when compared to a sine wave approximation. The least damaging tube  $\Delta T$  profiles result in 8.8% higher lifetimes compared to a sine wave approximation.

The operation time duration ( $t_{op}$ )—the period over which flux is used to heat molten salt—is assumed to be 12 hrs, which would represent a summer solstice design day. We choose a temperature ramping method that ensures simulated damage during startup or shutdown are minimal compared to steady operation. While startup or shutdown can pose damage risk to the tubes, the present study focuses on damage incurred due to normal on-sun operation and daily cycling. As a result, receiver startup and shutdown durations ( $t_{ph}$ ) are assumed to last 30 minutes. A preheat time of 30 min is conservative: as preheats considered are commonly less than 20 min Pérez-Álvarez et al. (2021); Laporte-Azcué et al. (2022); Vant-Hull (2002). During startup and shutdown, the tube's temperature profile



(a) Timeseries plots of panel 7 tube temperatures at three geometric locations throughout a summer solstice day. Axial position is  $z/L=0.72$ .




(b) tube temperature profile seen in panel 7 at design point. Note that the largest temperature difference is seen in the inside surface (dashed line), not the outer surface (solid line).


Figure 2.2: Simulation results from thermal model showing general temperature characteristics in receiver tubes.

isothermally and linearly ramps up or down between  $T_b$  and  $T_{f,op}$ . The cycle also includes a 1 hr period of off-time after shutdown. The total cycle period ( $t_{per}$ ), including preheat and off-time, is 14 hrs. The 24 hr day is therefore shortened because the receiver does not experience thermal loading during the nighttime hours.

Finally, the pressure cycle is specified. Internal tube pressure ramps up linearly for 30 minutes during startup, reaching a value of 1 MPa. Then, pressure is held constant until shutdown when it linearly ramps back down to 0 MPa over 30 minutes. An internal pressure of 1 MPa at the top of the tower approximately corresponds to the pumping pressure required for flow through the receiver path at Crescent Dunes. Compared to high pressure  $sCO_2$  receivers, liquid receiver tube pressures are typically low, on the order of  $10^0$  MPa Conroy et al. (2020). Consequently, liquid receiver loading is primarily caused by the thermal gradients.




(a) Maximum temperature difference at the tube's peak flux location for 3 representative tubes in each panel.




(b) Temperature difference ratios at the peak flux location. Note that  $\Delta T_{\text{cond}}$  and  $\Delta T_{\text{conv}}$  are defined in Fig. 2.2b

Figure 2.3: Solstice day results from thermal model in showing the range of temperature gradients observed in each panel and the respective radial/circumferential gradients.



(a)



(b)

Figure 2.4: Timeseries (2.4a) and circumferential (2.4b) boundary conditions used in SRLIFE.

$$\Delta T_{i,s}(t) \text{ or } \Delta T_{tot}(t) = \begin{cases} 0 & t \leq t_{ph} \\ \Delta T \sin\left(\frac{\pi(t-t_{ph})}{t_{op}}\right) & t_{ph} < t \leq t_{ph} + t_{op} \\ 0 & t > t_{ph} + t_{op} \end{cases} \quad (2.3)$$

$$T_f(t) = \begin{cases} T_b + \frac{T_{f,op}-T_b}{t_{ph}}t & t \leq t_{ph} \\ T_{f,op} & t_{ph} < t \leq t_{ph} + t_{op} \\ T_{f,op} - \frac{T_{f,op}-T_b}{t_{ph}}(t - t_{ph} - t_{op}) & t_{ph} + t_{op} < t \leq 2t_{ph} + t_{op} \\ T_b & 2t_{ph} + t_{op} < t \leq t_{per} \end{cases} \quad (2.4)$$

To summarize, boundary conditions provided as input to SRLIFE are based on actual thermal model results from a simulation of Crescent Dunes and assuming an image size priority SolarPILOT aiming strategy. Thermal loading on the receiver tubes will vary circumferentially and temporally, with large temperature gradients across the tube wall and between the heated and unheated sides of the tubes. The gradients vary gradually with time and an isothermal preheat is assumed. Ultimately, the loading cycle is dominated by the effects of thermal stress because the pressure stress is very low.

## Structural Modeling Methodology

Several structural modeling assumptions and settings are important to the application of SRLIFE to the present work. FEA is performed using the base deformation model in SRLIFE, which solves for elastic, plastic, thermal, and creep strains. Work in Berman and Rao (1983); González-Gómez et al. (2021) demonstrated the importance of using temperature dependent material properties. Consequently, all properties are temperature dependent and obtained using the Nuclear Engineering Material Library MESSNER et al. (2023).

Thermal boundary conditions and pressure at the inner and outer surfaces of the tube are computed and provided to SRLIFE which then calculates thermal gradients within the tube and subsequently performs transient FEA. The inner and outer tube temperatures vary with respect to circumferential position and time, as will be discussed later in this section. Timestep resolutions are no larger than 3 min, with the majority of simulations performed at 2 min time resolution. This timestep size is very small for transient receiver studies. For reference, timestep resolution in other recent transient creep-fatigue analyses have been in the range 5–60 min Laporte-Azcué et al. (2022); Gentile et al. (2023); González-Gómez et al. (2021). Timestep resolution was investigated in Laporte-Azcué et al. (2021b), and the authors found that larger timesteps lead to lifetime underestimation. When solving the full inelastic analysis, small timesteps prevent stackup of numerical error and improve the accuracy of the stress relaxation solution. Timesteps may be further subdivided if desired tolerances are not met when solving. Timestep convergence tolerances are set at  $1 \times 10^{-3}$  and  $1 \times 10^{-5}$  for absolute and relative tolerances, respectively. Simulation of three loading cycles, performed in the present work, shows less than 2% difference in 3rd day crown stress when comparing 1 and 2 min timesteps.

The primary drawback of FEA analysis of receiver tubes is the computational expense; therefore, several key assumptions are made to reduce simulation time. The problem is simplified from 3D to 2D by considering tube “slices” in the radial and circumferential dimensions at discrete positions along the axial direction of the tube. A 2D model only considers the thermal loading at one axial position and will neglect the effects of axial tube gradients. Previous work Messner et al. (2022) found that 2D lifetime estimates differed by less than 10% from 3D estimates as long as the tube slice was made at the axial location of maximum temperature. External receivers are comprised of panels which are made up of many parallel tubes. This work neglects additional stress/strain caused by panel-to-panel or tube-to-panel connections. Tubes are idealized as fully disconnected from one another, which means that thermal expansion of one tube does not affect stresses in adjacent tubes.

Receiver tube deformation assumptions can be assigned to one of three cate-

gories, as discussed in Laporte-Azcué et al. (2020): free bending, generalized plane strain, and clip-constrained deformation. Tube clips are commonly used in solar receivers to reduce excessive tube warpage and can significantly impact a tube's experienced stress/strain. The influence of clips on thermal stresses and deflection is studied in detail in Montoya Sancha et al. (2018b); Montoya et al. (2020), and elastic stress calculations that consider the effects of clips are shown in Montoya et al. (2020) to converge to the generalized plane strain solution as the number of clips increases. Therefore, this work assumes a receiver that uses clips and consequently enforces a generalized plane strain response, where the top is allowed to expand in the axial direction with an axial force of zero but the cross-sectional plane maintains a constant axial strain. Solar flux heating one side of receiver tubes results in compression on one side of the tube and tension in the other. As documented in Messner et al. (2022), a 2D model in SRLIFE assumes generalized plane strain. Temperature, elastic von Mises stress, and effective elastic strain on a 2D cross section are shown in Fig. 1.6 to illustrate typical gradients in a receiver of this type. This representative case is at conditions of  $T_f = 300^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$ .

Nodal convergence studies are performed to select a sufficient radial and circumferential node count. A base resolution of radial nodes and circumferential nodes is set at 3 and 72 respectively for a tube with outer diameter of 50.8 mm and tube thickness of 1.25 mm. A single cycle is simulated in SRLIFE and the strain range and crown stress are recorded for multiples of this base resolution i.e.  $6 \times 144$ ,  $9 \times 216$ , and  $12 \times 288$ . This method ensures that element aspect ratio does not change as nodal resolution increases. A resolution of  $6 \times 144$  is selected because the percent difference in the strain range predicted using the  $6 \times 144$  and  $12 \times 288$  is very low – less than 2.3%, and crown stress did not noticeably change for increasing node count beyond  $3 \times 72$ . However, the simulation time increased by almost a factor of 5 between the  $6 \times 144$  to  $12 \times 288$  resolutions. Based on this analysis,  $6 \times 144$  resolution is considered sufficient for simulating receiver tubes with reasonable computational time.

## Creep Damage

Creep damage is time-dependent material deformation caused by operation at elevated temperature and stress. It is defined in this work using the time-fraction method found in Section III, Division 5 of the ASME Boiler and Pressure Vessel Code Boiler and Code (2019) and calculated according to Eqn. 2.5:

$$D_c = \int_0^{t_p} \frac{dt}{t_r(\sigma, T)} \quad (2.5)$$

where  $\sigma$  and  $T$  are the von Mises stress and temperature. Rupture time ( $t_r$ ) correlations for all materials studied in this work are defined in Messner et al. (2022) with the exception of Alloy 230, whose rupture correlations are defined according to Eno et al. (2008). The creep, fatigue, and creep-fatigue interaction sources in Messner et al. (2022) are based on data from ASME Boiler & Pressure Vessel Code, literature data, and testing completed at Idaho National Laboratory as part of a Solar Energy Technologies Office-funded project on metallic receivers.

Note that to the authors' knowledge, there are currently no widely accepted methods for completing a statistical life analysis on metallic components. Instead, current practice is to provide a life based on lower-bound material properties for design or with median properties for remaining life evaluation, as discussed in ASME Section VIII American Society of Mechanical Engineers (2019) and ASME FFS-1 American Petroleum Institute & American Society of Mechanical Engineers. Consequently, the creep rupture properties used in this paper represent average rupture properties based on data fits discussed in Messner et al. (2022) and Eno et al. (2008).

The von Mises stress is calculated according to Eqn. 2.6. Creep damage is evaluated using the stress at the crown because maximum creep damage is anticipated to occur where stress is high and maximum tube temperature occurs.

$$\sigma_{vm} = \sqrt{\frac{(\sigma_{xx} - \sigma_{yy})^2 + (\sigma_{yy} - \sigma_{zz})^2 + (\sigma_{zz} - \sigma_{xx})^2 + 6(\sigma_{xy}^2 + \sigma_{yz}^2 + \sigma_{xz}^2)}{2}} \quad (2.6)$$



Creep results in part of the elastic strain being converted to permanent strain, thereby causing a phenomena called stress relaxation. Creep strain will accumulate at the crown as a receiver tube continues to be exposed to high operational stresses and temperatures. This accumulation of creep strain ultimately reduces crown stress, which subsequently decreases the creep strain rate. Steady thermal stresses will relax to zero over time. However, under transient conditions, if the alternating stress exceeds twice the material yield stress ( $S_y$ ) then global relaxation will not occur. Rather, the tube stresses will reset on each repetition of the thermal cycle, leading to greatly increased creep damage. Accurate prediction of a receiver tube's stress relaxation is difficult because typical CSP loading cycles apply stress/strain profiles that are gradual and constantly changing. Receiver tube stress relaxation will occur slowly and could be lower in magnitude than the ideal, global relaxation case Barua et al. (2020); Messner et al. (2022), assuming the stress reset limit is not exceeded. Since characterization of stress relaxation will have a large impact on lifetime, the FEA-based approach used in this work will now be described.

Stress relaxation caused by creep is assumed to primarily occur in the first 20 days. Recent work in González-Gómez et al. (2021) concluded that the most significant stress relaxation occurs in the first 25–35 days and this agrees with results in Berman and Rao (1983). FEA receiver tube simulations of Incoloy 800H tubes in Berman and Rao (1983) found that stress relaxation had a power law relationship with number of cycles – meaning that the logarithm of creep damage decreases linearly with the logarithm of cycles. This generalized behavior excludes a receiver's first load cycle, which is when one-time plastic yielding typically occurs. In order to understand the first-cycle yielding effect, a creep-relaxation study is performed for a Haynes 230 tube with geometry similar to that of Crescent Dunes. A series of SRLIFE simulations are run, where each case experiences 20 identical thermal loading cycles as described in Section 2.2. Each 20-cycle simulation has a different combination of  $T_f$ , ranging from 300 to 550°C, and various  $\Delta T$ s. The range of radial-direction  $\Delta T$  is 100–220°C for any cases with  $T_f$  of 300–475°C, and 100–200°C for  $T_f$  of 500–550°C. Three simulation results are shown in Fig. 2.5, with the peak daily crown stress in Fig. 2.5a and the total creep damage incurred over each cycle in

Fig. 2.5b. In general, creep damage results have two regions of behavior—a stress reset region and a stress relaxation region. Examples of both regions are shown in Fig. 2.5. In the  $T_f = 475^\circ\text{C}$ ,  $\Delta T = 190^\circ\text{C}$  case, stress relaxation is unimpeded and occurs during each cycle. The lower temperature  $T_f = 400^\circ\text{C}$ ,  $\Delta T = 210^\circ\text{C}$  case also demonstrates stress relaxation, although it occurs at a slower rate. Figure 2.5 shows that stress reset will occur when both  $\Delta T$  and  $T_f$  are large. The  $T_f = 475^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$  case experiences stress reset, and no significant relaxation occurs after approximately cycle 5. There is also some non-linearity in the log-log plots of stress relaxation, which could be due to the accumulation of numerical modeling error caused by the significant complexity of conducting fully inelastic, multi-cycle FEA simulations. All three cases in Fig. 2.5 show that the first cycle is actually when the most significant decrease in creep damage will occur due to material yielding causing stress redistribution. This observation highlights the importance of extrapolating stress relaxation and structure lifetime with data points beyond cycle 1.

This paper predicts tube life by extrapolating long-term behavior following simulation of a reduced number of cycles. Since a simulation of 20-cycles requires significant computation time, we show that it is possible to instead simulate only 4 cycles, then use a case-specific power law fit to the creep damage value after the 3<sup>rd</sup> and 4<sup>th</sup> cycle. Fitting in this way enables extrapolation of creep damage to the 20<sup>th</sup> cycle, and this creep value is assumed as the persisting long-term value used to estimate the total number of hours of operation. This approach is modified slightly in the case that stress reset is present. The stress reset region occurs where the elastic stress exceeds  $2S_y(T)$ . Stress relaxation for the 20<sup>th</sup> day is assumed to be equal to the 4<sup>th</sup> day stress relaxation if stress reset is predicted to occur in the tube.

Alternative approaches for modeling stress relaxation are presented by Gonzalez et al. González-Gómez et al. (2021) which models crown stress relaxation using the Norton-Bailey model and then predicts creep damage based on the expected crown stress. This analytical approach was proven to be computationally inexpensive and therefore useful for design purposes. However, it potentially over-predicts creep damage because small errors in predicted stress result in large differences

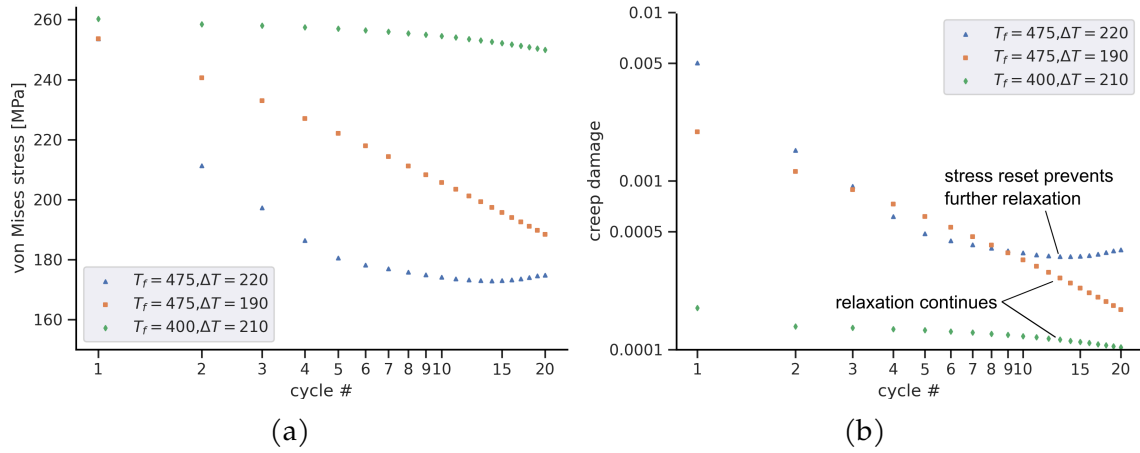


Figure 2.5: Maximum crown stress (2.5a) and creep damage per cycle (2.5b) for three different temperature cases showing stress relaxation phenomena as well as stress reset in the  $T_f = 475^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$  case.

in calculated creep damage González-Gómez et al. (2021). Instead of predicting stress relaxation and then calculating the resultant creep damage, the present work directly predicts future creep damage based on known creep damage results from simulated cycles. Thus, this FEA-based method aims to achieve higher accuracy at the expense of increased computing time compared to analytical methods. To evaluate this method, the difference between FEA simulations' 20<sup>th</sup> cycle creep damage and the predicted creep damage based on extrapolation of cycles 3 and 4, is presented in Fig. 2.6. It can be seen that extrapolating creep damage using this method is typically conservative, with the most overestimation of creep damage occurring in the 400–475°C range. The two outliers are locations where stress reset was predicted, but did not actually occur in the SRLIFE simulation. Extrapolated creep damage and FEA 20<sup>th</sup>-cycle creep damage are very similar when the magnitude of damage per cycle is less than  $2 \times 10^{-4}$ . The majority of tubes are expected to operate with damage less than  $1 \times 10^{-4}$  per cycle since solar receiver lifetime design is typically intended to exceed 10,000 cycles. In this region, the extrapolation is shown to work very well for the alloy Haynes 230.




Figure 2.6: Comparison of extrapolated 20<sup>th</sup> day creep damage versus actual damage from extended 20-day cycles. Extrapolation inaccuracy is only prevalent in 2.6a at points with creep damage too high for viable designs, and 2.6b shows zoomed in results where accuracy is sufficient in expected design/operation region.

## Fatigue Damage

Fatigue damage in receiver tubes is assumed to occur due to cyclic material strain. In this case, material strain is caused by a combination of pressure stress and thermal stress. Tube pressure during operation will stay nominally constant, but the gradually changing operating temperature and flux causes thermal stress and strain to evolve throughout the day. Because of this, the maximum strain range for a receiver tube is usually the difference between the strain at maximum operating temperature and the material strain after operation has ceased for the day and the tube has had sufficient time to equilibrate with ambient temperature. Multi-cycle simulations from SRLIFE using the base model show that the tube's maximum strain range decreases significantly after the first cycle. An example of this is shown in Fig. 2.7, where a  $T_f = 300^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$  case is simulated for four cycles. The maximum strain range decreases after the first day, and this decrease is due to yielding, which occurs once during the tube's first loading cycle. Measuring strain range after the first full cycle shows that strain range will not significantly change




Figure 2.7: Von Mises strain range timeseries from a simulation case with  $T_f = 300^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$ .

during future cycles for CSP cases, and this is also shown in Berman and Rao (1983). Strain ranges from the 3<sup>rd</sup> and 20<sup>th</sup> cycles are compared for all 20-cycle cases discussed in Section 2.2. Results show that the maximum error between the two strain ranges is less than 2% in high temperature cases. Note that the strain range calculated via SRLIFE includes creep strain. Thus, the error for the total daily strain that would contribute to fatigue damage is even lower. Therefore, the strain range for any later cycle is assumed to be the same as the 3<sup>rd</sup> cycle's maximum strain range.

The fatigue damage of a single cycle at a given equivalent mechanical strain range ( $\Delta\epsilon$ ) and maximum cycle temperature ( $T$ ) is calculated using the SRLIFE built-in fatigue curves for each material, where each material's fatigue cycles-to-failure data is fitted with a series of polylogarithmic functions. A detailed description of the fatigue data sources for each material is provided in Messner et al. (2022), and further data source details are provided in Section 2.2. In the case that strain ranges are lower than the typical cutoff strain range, extrapolation based on the corresponding polylogarithmic curve-fit is used. The total fatigue damage experienced by the tube over its entire lifetime is defined in Section III, Div. 5 of ASME

Boiler and Pressure Vessel Code Boiler and Code (2019) and shown in Eqn. 2.7:

$$D_f = \sum_{i=1}^{LTE} \frac{1}{N_{f,i}(\Delta\epsilon_{i,max}, T_{max})} \quad (2.7)$$

where  $\Delta\epsilon_{max}$  is a node's maximum strain range,  $T$  is the temperature, and  $i$  is any given cycle. The effective strain for any node is calculated according to Eqn. 2.8

$$\epsilon = \frac{\sqrt{2}}{2(1 + \nu)} \sqrt{(\epsilon_{xx} - \epsilon_{yy})^2 + (\epsilon_{yy} - \epsilon_{zz})^2 + (\epsilon_{zz} - \epsilon_{xx})^2 + \frac{3}{2}(2\epsilon_{yz}^2 + 2\epsilon_{xz}^2 + 2\epsilon_{xy}^2)} \quad (2.8)$$

This work considers the maximum strain range of any node in the tube's cross section, although it is most likely to occur at the crown, where temperature and stress will typically be the highest.

## Predicting Tube Lifetime

A single cycle's creep or fatigue damage are typically expressed in terms of the fraction of the material's total lifetime consumed during that cycle. Therefore, the definition of the tube's lifetime is the number of cycles that can occur before the cumulative damage is greater than unity. However, when a tube simultaneously experiences fatigue and creep damages, tube failure may occur before the cumulative sum of damages reaches unity. Creep-fatigue interaction is detrimental because creep-induced voids cause fatigue cracks to grow across a grain boundary Barua and Messner (2023). This impact on lifetime is typically accounted for with so-called damage envelopes. The data sources for each material's damage envelope are discussed in Messner et al. (2022). Note that the damage envelope from Inconel 617 is also used for Alloy 230 because their other mechanical properties are similar and there is currently no creep-fatigue test data to support a specific envelop for Alloy 230. In this work, the lifetime for any given tube is based on the assumption that the 20<sup>th</sup>-day fatigue and creep damages are representative of every future loading cycle. Authors also assume that the receiver experiences 365 cycles per year,

meaning that there are no assumed days of downtime in each lifetime year. With these definitions, SRLIFE's *damage* and *library* modules can be used to extrapolate the lifetime based on how many cycles will be performed before the tube's damage point lies outside of its respective damage envelope.

## Considerations for Alternative Alloys

The creep-fatigue analysis methodology discussed in Sections 2.2– 2.2 is first applied for a base case scenario that assumes Haynes 230 is used. Section 2.4 presents these results and also investigates the performance of five alternative alloys: Inconel 617, Haynes 282, Inconel 740H, SS 316H, and Incoloy 800H. The alternative alloy studies are carried out following the same methodology as for the base case study in Section 2.3. The following subsection introduces each alloy, specifies any differences in the thermal operating ranges considered, and compares the alloys' key properties. Finally, a method is introduced that accounts for the differences in thermal conductivity between the alternative materials and its effect on the thermal stresses experienced during operation. The main objective of these additional parametric studies is to identify whether each alloy can be expected to fail by fatigue or creep damage. Furthermore, the resulting damage maps also highlight differences in material performance and provide insight into expected lifetimes.

**Comparison of Alloy Properties** Each alloy considered is either a candidate for future receivers or is used in an existing receiver. The geometry is consistent with the baseline study, where  $D_o = 50.8$  mm and  $t_h = 1.25$  mm. This geometry is used so that damage maps can be readily compared, not necessarily because it represents the optimal tube design. Tube thickness can remain constant even for the lower yield strength alloys because the pressure stress is sufficiently low.

Each alloy's properties are already included in SRLIFE via MESSNER et al. (2023) and creep rupture and fatigue correlations are discussed in Messner et al. (2022). All five additional alloys are potential options for receiver tubes because of their high rupture strength and corrosion resistance at elevated temperatures. Curve

fits of temperature-dependent yield strength and 100,000 hr rupture strength are plotted in Fig. 2.8 using data obtained from MESSNER et al. (2023). Note that as previously stated, rupture time for Haynes 230 is obtained using the Mendelson-Roberts-Manson coefficients provided in Eno et al. (2008). Figure 2.8 illustrates that the six alloys (including Haynes 230) can be grouped into three primary categories according to yield strength and rupture strength: high (Haynes 282, Inconel 740H), medium (Haynes 230, Inconel 617), and low (SS 316H, Incoloy 800H). Although SS 316H and Incoloy 800H's inferior yield and rupture properties equate to generally lower lifetimes and higher cyclic damage, the lower cost of both alloys warrants an investigation into their performance. For example, recent work in Laporte-Azcué et al. (2021a) estimated SS 316H and Incoloy 800H cost to be approximately  $17.6\times$  and  $3.8\times$  cheaper respectively than Haynes 230. Tube film temperature limits, which heavily impact corrosion rates inside the receiver tubes, are assumed to be  $650^\circ\text{C}$  for Haynes 230, Inconel 617, Haynes 282, and Inconel 740H so each material's parametric range is the same as that in the baseline study discussed in Section 2.3. The film temperature corrosion limit for SS 316H is assumed to be  $600^\circ\text{C}$ , according to Miliozzi et al. (2001) and therefore the parametric simulation range is reduced to  $T_f = 300\text{--}500^\circ\text{C}$  and  $\Delta T = 100\text{--}180^\circ\text{C}$ . Incoloy 800H's parametric range is also reduced because of its low yield strength. Simulation results show that Incoloy 800H experiences stress reset conditions as early as  $\Delta T = 130\text{--}140^\circ\text{C}$ .

**Adjusting Ratio for Different Materials** Thermal conditions for each material case are unchanged except that the ratio of  $\Delta T_{o,s}$  to  $\Delta T_{i,s}$  differs slightly due to differences in thermal conductivity. The thermal boundary conditions in SRLIFE are slightly adjusted for each alloy in the following way to account for differences in conductivity. The total temperature difference, the horizontal axis in Fig. 2.12, has a ratio of  $r = \Delta T_{o,s}/\Delta T_{i,s} = 0.5$  for Haynes 230. The ratio's numerator ( $\Delta T_{o,s}$ ) is determined by the conduction heat transfer characteristics through the wall, while  $\Delta T_{i,s}$  is determined by convection heat transfer into the fluid. If all conditions other than conductivity remain the same,  $\Delta T_{i,s}$  will not change because the heat transfer coefficient is assumed to be constant. However, the wall temperature gradient will



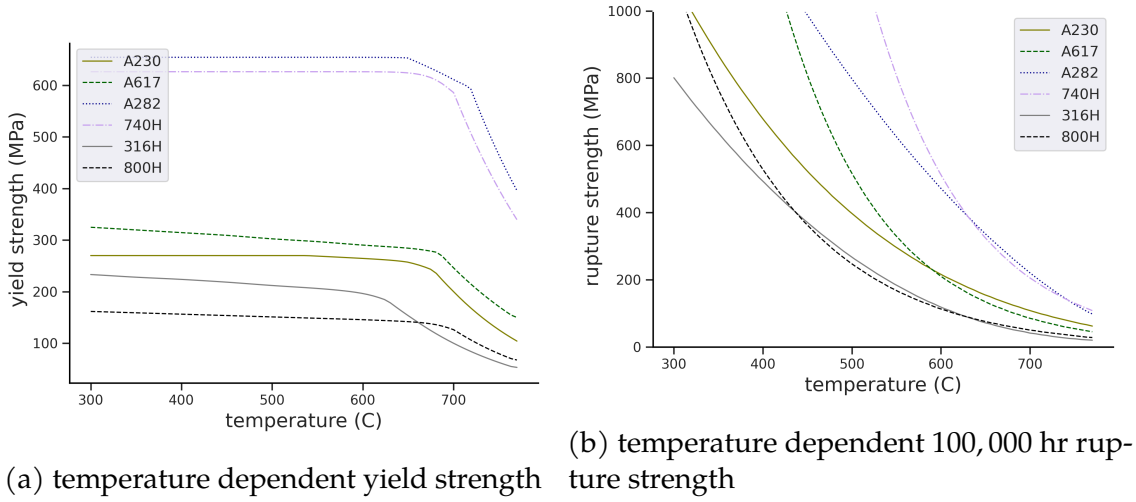


Figure 2.8: Yield strength (2.8a) and 100,000 hr rupture strength (2.8b) for the six receiver alloys tested in this work. Rupture strength properties for Haynes 230 are based on the Mendelson-Roberts-Manson coefficients in Eno et al. (2008). All other alloy properties are obtained using MESSNER et al. (2023).

change as a function of thermal conductivity as seen in Eqn. 2.9 where  $\Delta T_{o,s,new}$  is the new temperature difference of any alloy,  $k$  is conductivity, and  $\Delta T_{o,s,base}$  is the outer temperature difference for the baseline case analyzed with the thermal model in Martinek et al. (2021).

$$\Delta T_{o,s,new} = \frac{k_{base}}{k_{new}} \Delta T_{o,s,base} \quad (2.9)$$

Therefore, the temperature difference ratio for any alloy is defined using Eqn. 2.9 and assuming the  $\Delta T_{i,s}$  remains the same. Any alloy's temperature difference ratio is defined according to Eqn. 2.10, where  $r_{base}$  is the baseline case ratio. The ratio for each alloy as well as the thermal conductivity used is included in Table 2.1. Ratios do not largely differ between alloys when compared to the original temperature difference ratio for Haynes 230.

$$r_{new} = \frac{\Delta T_{o,s,new}}{\Delta T_{i,s,new}} = \frac{k_{base}}{k_{new}} \frac{\Delta T_{o,s,base}}{\Delta T_{i,s,base}} = \frac{k_{base}}{k_{new}} r_{base} \quad (2.10)$$

Table 2.1: Various alloy conductivities and the corresponding temperature difference ratio used to define boundary conditions in SRLIFE.

	A230	A617	A282	740H	316H	800H
k(T=425) [W/m-K]	16.9 International (2024a)	19.7 Metals (2024a)	19.0 International (2024b)	16.1 metals (2024)	19.9 Ho and Chu (1977)	18.3 Metals (2024b)
r [-]	0.5	0.43	0.44	0.53	0.43	0.46

## 2.3 Damage Map for Haynes 230

The discussed methodology is first used to evaluate receiver lifetime for a baseline case similar to Crescent Dunes in Tonopah, Nevada. A series of FEA simulations are performed over the defined parametric space  $T_f$ : 300–550°C and  $\Delta T$ : 100–220°C with 25°C and 10°C resolution respectively. These ranges include the maximum combinations of  $\Delta T$  and  $T_f$  of each simulated tube in the thermal model. Simulations assume a receiver sized for a design thermal input of approximately 565 MW<sub>th</sub> operating for an average of 9 hrs each day. The average cycle time is found using the System Advisor Model NREL (2023) assuming the latitude and longitude corresponding to Tonopah, Nevada. The receiver geometry and heliostat layout case for the case study as well as the alternative alloys study in Section 2.4 is the same as that described in Section 2.2. Damage modeling assumes that the first 20 days occur around the design day, and daily stress relaxation occurs according to the 12 hr operation time described in Section 2.2. All subsequent operation times are assumed to be 9 hrs and future creep damages are calculated by multiplying the 20th day creep damage by a factor of  $9/12 = 0.75$ .

The tube material is Haynes 230 ('A230' in SRLIFE) with diameter of 50.8 mm and wall thickness of 1.25 mm. As stated previously in the methodology section, each simulation includes four cycles. Average simulation time ranged from 3–5 hours and simulations were carried out using the Center for High Throughput

Computing which leverages parallelized computing resources throughout campus Center for High Throughput Computing (2006). Processing included calculation of the von Mises stress at the crown for each timestep and the maximum strain range for each cycle. A subset of strain range results from cycle 4 is shown in Fig. 2.9. It can be seen that maximum strain range shows an approximately linear relationship with  $\Delta T$ . Higher  $T_f$  values also increase the observed maximum strain range. Thus, cyclic strain range is observed to always increase with increasing  $T_f$  and  $\Delta T$  in the simulated parametric range.




Figure 2.9: Third cycle maximum elastic-plastic strain range results from SRLIFE for Haynes 230 tubes with 50.8 mm diameter and 1.25 mm thickness subject to a range of  $\Delta T$ s with three different fluid temperatures.

In contrast, the crown stress's relationship with  $T_f$  and  $\Delta T$  is more complex due to the impacts of yielding and stress relaxation due to creep. Peak crown stress results from each simulation, which are recorded at the time of peak thermal conditions, demonstrate several important points. Elastic crown stress results from all simulations are shown in Fig. 2.10, and Fig. 2.11 shows the elastic-plastic stresses from days 1 and 4. Figure 2.10 shows that elastic stress is almost solely determined by the temperature difference,  $\Delta T$ . The results in Fig. 2.11a demonstrate that crown stress is not strongly dependent on  $\Delta T$  after approximately 130°C because this is the point at which stress conditions exceed the yield strength. As a result, the crown stress range for this parametric space is relatively small—approximately 60 MPa.




Figure 2.10: Elastic stress map for all tube  $T_f$ - $\Delta T$  combinations simulated in SRLIFE Haynes 230 baseline case study.

The final simulation cycle's stress map (seen in Fig. 2.11b) shows the impact of stress relaxation. Creep that occurs in the receiver tube causes the crown stress to decrease, and is strongly dependent on operating temperature. Material at high operating temperatures, seen in the upper right-hand corner of Fig. 2.11b undergoes significant stress relaxation which ultimately reduces creep damage incurred in future cycles. However, keeping if Haynes 230 tubes at higher temperature, then higher gradient conditions is generally undesirable because of the risk of operating in the stress reset region. An example of this is shown in Fig. 2.5, where the  $T_f = 475^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$  case does not experience significant further relaxation past cycle 6. Thus, although the extreme operating conditions cause beneficial stress relaxation by the 6th simulated cycle, crown stress is not expected to further relax in subsequent cycles. It is interesting to note that in this parametric space, stress reset generally occurs after several cycles have elapsed. As a result, some permanent stress relaxation occurs during these cycles prior to the set-in of reset behavior

Using the methodology described in Section 2.2, strain and stress results for all simulated cases are analyzed to predict lifetime and the ratio of fatigue to creep damage. Additionally, SRLIFE simulations for each temperature case are also run using the elastic model, which allows for prediction of the stress reset region. The maximum ratio of fatigue to creep damage for Alloy 230 in the defined parametric




Figure 2.11: Day 1 (2.11a) and day 4 (2.11b) elastic-plastic stress map for all tube  $T_r$ - $\Delta T$  combinations simulated in SRLIFE Alloy 230 baseline case study. Tube diameter is 50.8 mm and thickness is 1.25 mm. Note that the color scales differ.

space is  $6.08 \times 10^{-4}$  and occurs at  $T_r = 550$  °C and  $\Delta T = 220$  °C. Thus, over 1,600 fatigue cycles per day would be required for the magnitudes of fatigue damage to equal that of creep damage. This observation has important implications for operations during cloud events. It should be noted that there exists a diverse range of cloud events and their thermal effects depend on how flux aiming and mass flow is managed during the event. If we assume that a cloud passage causes a decrease and subsequent increase of incident flux but mass flow is held constant, then the receiver tubes will experience a decrease and subsequent increase of thermal stress and strain that is of no larger magnitude than the thermal cycle associated with a clear sky day's operation. The present work demonstrates that the fatigue damage of a clear sky day for the baseline (Alloy 230) case is insignificant compared to the same cycle's creep damage. This fact demonstrates that creep damage is the dominant form of failure for this baseline study. Consequently, in this baseline study, tube lifetimes are not expected to be significantly decreased by cloud passages that exclusively cause fatigue damage. However, it should be noted that a day's creep damage accumulation could differ from the present work's predictions in

cases of variable solar flux. The applicability of this work's lifetime predictions to non-clear sky cases depends on the solar flux conditions and mass flow control. This work overestimates lifetime compared to days with cloud passages that cause tube temperature spikes and increased creep damage accumulation and underestimates lifetime compared to days where low solar flux results in lower tube temperatures.

All results are summarized in Fig. 2.12, where each coordinate on the damage map represents a thermal operating condition defined by  $T_f$  and  $\Delta T$ . Operation in the stress reset region (labeled "S.R.R.") shown in the upper right corner of Fig. 2.12 is not recommended, as lifetimes are severely curtailed because the tubes cannot experience stress relaxation. An additional region, the corrosion region (labeled "C.R."), is included. A critical film temperature of 650 °C is assumed based on the work in McConohy and Kruizenga (2014). The receiver tube is expected to experience excessive corrosion if its peak tube temperature exceeds this limit. The corrosion region represents all the thermal operating conditions in the damage map that result in a peak film temperature greater than 650 °C, assuming a constant heat transfer coefficient. Operation in this region is also not recommended. As expected, the lower left region of this damage map corresponds to effectively infinite life with mild conditions under which receiver tubes could operate indefinitely. In the portion of the map where lifetimes are limited, there exists a region with intermediate  $T_f$  and  $\Delta T$  where lifetimes are the lowest. This phenomena can be explained by recalling the impact of stress relaxation. Since stress relaxation is a stronger function of  $T_f$  than  $\Delta T$ , there is a region where high stress occurs but  $T_f$  is not high enough to encourage significant stress relaxation. Below and to the left of this lifetime trough, the lower  $T_f$  and  $\Delta T$  values result in increased lifetimes. Above the lifetime trough, higher  $T_f$  values encourage beneficial stress relaxation and extend the tube lifetimes. Lifetime contours in close proximity indicate a greater sensitivity to flux changes (primarily affecting  $\Delta T$ ) and mass flow rate (primarily affecting  $T_f$ ). For any  $\Delta T$  point, there is approximately 20–30°C difference in  $T_f$  between a 30 year and 80 year lifetime. This observation underscores the importance of mass flow control during operation; deviations of  $T_f$  from its expected operating point will have significant lifetime impacts.

Using the thermal model from Martinek et al. (2021), three representative tubes from each panel are simulated for a design day. Each tube is overlaid on Fig. 2.12 by defining the tube's operating conditions (i.e. location on the damage map) as the maximum  $\Delta T$  and  $T_f$  combination observed along the length of each tube. The lowest calculated tube lifetime is 13.9 yrs at  $T_f = 505.6^\circ\text{C}$  and  $\Delta T = 136.3^\circ\text{C}$ . Firstly, it should be observed that this methodology provides reasonable lifetime estimates for a real-world example. Secondly, although the magnitudes differ, both the present work and work in Laporte-Azcué et al. (2022) find that, for a clear sky analysis on a cylindrical molten salt receiver, panel lifetimes greatly differ depending on their flowpath location. The operating points shown in Fig. 2.12 indicate that high flux, lower temperature panels can operate indefinitely. Higher  $T_f$  panels have somewhat limited lifetime, however. The lifetimes of tubes experiencing high  $T_f$  could be extended without affecting receiver area by lowering their incident flux while simultaneously increasing the incident flux on lower  $T_f$  tubes. This study demonstrates the viability of the methodology presented in this work for analyzing the range of clear sky operating points receiver tubes may experience.




Figure 2.12: Damage map for baseline case (Alloy 230), showing LTE contour lines and locations of critical operating points for receiver tubes based on operation on a 9 hr day. The stress reset region (S.R.R.) and the corrosion region (C.R.) are also shown. Lifetime units are in years, and assume that no significant stress relaxation occurs after the 20th cycle and that fatigue damage is constant after the 3rd cycle.  $\Delta T$  is the total temperature difference defined in Section 2.2 and  $T_f$  is the fluid temperature.

## 2.4 Damage Maps for Alternative Alloys

Other materials are studied to gain additional insights, and the resulting damage maps for each material are presented. Fatigue damage is not found to be a significant factor for the majority of cases, although some exceptions appear. For all materials considered in this work, the end-of-life ratio of fatigue to creep damage is calculated for each operating point assuming that no further stress relaxation occurs and that strain range remains constant, i.e. the ratio of fatigue to creep damage is  $D_{f,20}/D_{c,20}$ .



**Damage Map for Inconel 617** The damage map for Inconel 617 is shown in Fig. 2.13. In contrast to Alloy 230, ratios of fatigue to creep damage are higher in magnitude. A contour plot is included in Fig. 2.13 to show regions where tubes are more sensitive to fatigue, with the damage ratio shown on a logarithmic scale. When considering tubes that operate within the 5–80 yr contour lines, fatigue damage is primarily a concern in the  $\Delta T = 180\text{--}220^\circ\text{C}$  region where tubes with 10–80 year lifetimes will experience daily fatigue that is 0.1–1% of daily creep damage. Figure 2.13 shows that tubes lifetimes at combinations of higher  $T_f$  and lower  $\Delta T$  are primarily determined by creep damage. The damage ratio increases at combinations of high  $T_f$  and high  $\Delta T$ , but tubes are not designed to be operating in this region because of proximity to the stress reset region and extremely low lifetime predictions. Lifetime steadily decreases with increasing  $T_f$  and  $\Delta T$ , and there is no beneficial stress relaxation island such as that observed in Alloy 230. Within the  $\Delta T$  range of  $130\text{--}150^\circ\text{C}$ , an elbow can be observed in each contour line which represents the  $\Delta T$  at which, for a constant  $T_f$ , increasing  $\Delta T$  results in material yielding. The stress reset region is similar in size to Alloy 230 because the yield strengths of the two materials are similar.

**Damage Map for Haynes 282** Results for Haynes 282 are shown in Fig. 2.14. Note that the fatigue-to-creep damage ratio regions are omitted because they are found to be extremely low for Haynes 282 throughout the entire parametric space, with a maximum value of  $5.65 \times 10^{-12}$  at the  $T_f = 540^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$  case. Thus, fatigue is always expected to be negligible for Haynes 282 receiver tubes operating within the present work's range of simulated conditions. The lifetime contour lines in Fig. 2.14 are approximately linear, as none of the cases cause yielding. Additionally, there are no points in the parametric space where stress reset is expected to occur. Both of these observations are explained by the high yield strength of Haynes 282. The maximum elastic stress case is  $437.1\text{ MPa}$  at the  $T_f = 550^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$ , which is less than  $2S_y(T = 770^\circ\text{C}) = 792.12\text{ MPa}$ . The damage map shows that Haynes 282 lifetimes for each temperature case are higher than any other material considered. In the regions that failure does occur, creep is the dominant damage mechanism.




Figure 2.13: Damage map for Inconel 617 based on results from SRLIFE. Contour lines from 5–80 years are shown. The stress reset region (S.R.R.) and the corrosion region (C.R.) are also shown. Color levels denote regions of different ratios of fatigue to creep damage.

**Damage Map for Inconel 740H** The results for Inconel 740H, shown in Fig. 2.15, display similarities with Haynes 282 results in that at no point is stress reset predicted to occur. Additionally, the contour lines are also linear for Inconel 740H. Again, this is due to the alloy's high yield strength. The largest elastic stress is 464.6 MPa at  $T_f = 550^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$ , and the corresponding yield strength at the crown is  $S_y(T = 770) = 679.17$  MPa. Damage ratios are higher than Haynes 282, indicating that fatigue damage is more significant in Inconel 740H. The highest ratio is 0.0162 at  $T_f = 345.5^\circ\text{C}$ ,  $\Delta T = 220^\circ\text{C}$ . However, the corresponding lifetime at this location is 2590 years. Thus, although fatigue reaches significant levels at some locations, tube lifetime at those locations is not expected to be limited.




Figure 2.14: Damage map for Haynes 282 based on results from SRLIFE. Contour lines from 5–80 years are shown. The upper right shaded area is the corrosion region (C.R.). Stress reset does not occur for Haynes 282 within this parametric space, and fatigue is not significant compared to creep damage.

**Damage Map for SS 316H** The SS 316H parametric study is performed with the same methodology as already described, but the  $T_f$  and  $\Delta T$  ranges are reduced because SS 316H has a comparatively lower film temperature limit. Although the contour lines for 30–80 yrs are linear as seen in Haynes 282 and Inconel 740H, in this case the contour lines are linear because all locations experience yielding. A reduced lifetime trough (seen as the 20 year contour line) is observed, similarly to Alloy 230. As before, this is an area where creep damage is significant due to high temperature, but stress relaxation due to creep is not significant enough to be beneficial. Over half of the contour line segments for 30–80 years lie in a region where the damage ratio is significant—where fatigue damage ( $d_f$ ) is 1–10 % of total damage. Stress reset is predicted to occur for many cases, due to SS 316H’s lower yield strength.




Figure 2.15: Damage map for Inconel 740H based on results from SRLIFE. Contour lines from 5–80 years are shown. The upper right shaded area is the corrosion region (C.R.). Stress reset does not occur for Inconel 740H within this parametric space. The various color levels denote regions of different ratios of fatigue to creep damage.

**Damage Map for Incoloy 800H** Similar to SS 316H, the parametric space considered is reduced for Incoloy 800H because of its lower creep tolerance and yield strength and the damage map is shown in Fig. 2.17. Even with the reduced parametric space, however, Incoloy 800H experiences stress reset in over half of the operating space. If the stress reset region is avoided, the majority of the contours lie at a damage ratio of  $1 \times 10^{-5}$  to  $1 \times 10^{-4}$ . Thus, creep damage is shown to dominate fatigue damage for Incoloy 800H. It can be seen that the damage ratio contour bands skip several bands in the  $T_f = 375\text{--}425^\circ\text{C}$  region, which can be explained by the nature of the fatigue curves and time-to-rupture curves. Damage properties are implemented in SRLIFE as fits of experimental data with finite temperature ranges. The damage ratios can appear to change rapidly if creep and fatigue curves




Figure 2.16: Damage map for SS 316H based on results from SRLIFE. Contour lines from 20–80 years are shown. The red shaded area represents the stress reset region. Color levels denote regions of different ratios of fatigue to creep damage.

simultaneously activate around the same temperature.




Figure 2.17: Damage map for Incoloy 800H based on results from SRLIFE. Contour lines from 10–80 years are shown. The red shaded area represents the stress reset region. Color levels denote regions of different ratios of fatigue to creep damage.

## 2.5 Damage Maps at Other Conduction-Convection Ratios

The conduction-convection ratio ( $R$ ) is defined in this work as the ratio of conduction temperature difference to the convection temperature difference. Original damage modeling uses thermal boundary conditions that assume the conduction-convection ratio is equal to 0.5. However, this ratio is dependent on the material conductivity and the heat transfer coefficient because the conduction and convection temperature differences depend on their associated thermal resistances. The initial damage model is based on a 0.5 ratio because thermal modeling of Crescent Dunes at design point found conduction-convection ratios of  $0.49 \pm 0.06$  at the highest tube flux locations. Nevertheless, cases where  $R$  differs from 0.5 can commonly occur,

especially due to the direct dependence of salt heat transfer coefficients on the flow velocity. Varying the number of tubes in parallel or adjusting the receiver's mass flow can lead to values of  $R$  greater or lesser than 0.5. This section assesses the integrity of the damage model when  $R$  is assumed to be 0.5. Ultimately, damage maps with values of  $R = 0.25$  and  $R = 0.75$  yield similar lifetime predictions to the original damage map based on  $R = 0.5$ .

As stated previously, a tube's total temperature difference can be broken up into two temperature differences. Assuming the ratio of conduction to convection enables us to simplify the thermal loading to be solely based on total temperature difference. The conduction-convection ratio contains important information about the temperature gradient throughout a tube's cross section. An  $R$  close to unity represents a case where the wall temperature difference in the radial direction is more significant than the case of  $R$  equal to zero. Figure 2.18 shows the thermal gradients for a low ratio and a high ratio. Both cross sections experience the same total temperature difference from the crown to the tube's backside and fluid temperature. Another way to interpret this is that the outer surface temperature profiles for both cases should be approximately identical. However, the hot side of the higher conduction-convection ratio case (Fig. 2.18b) displays a greater difference between the outer and inner surface temperatures. In contrast, the low conduction-convection ratio case (Fig. 2.18a) primarily only experiences a temperature that varies in the circumferential direction.

The corresponding von Mises stress profiles, which consider stress in all coordinate directions, can be seen in Fig. 2.19. It can be observed that the stress profile closely follows the temperature gradient profile in both cases. Despite having different stress distributions, the stress metric that is used to predict the tube lifetime is very similar in magnitude for both cases. The crown stress for the  $R = 0.25$  case is 257.05 MPa and 257.37 MPa for the  $R = 0.75$  case. It could be argued that failure is more likely in the  $R = 0.25$  case because more of the tube cross section is at high stress. However, similarly to SRLIFE's method, this work predicts damage based on the single most critical point in the tube. Therefore, the lifetime predictions will be similar in both cases because the crown stresses are similar.




Figure 2.18: Thermal gradients on a Haynes 230 tube cross section simulated in SRLIFE at a  $\Delta T=190$  (C) and  $T_f=425$  (C). Fig. 2.18a shows results from a conduction-convection ratio of 0.25, and Fig. 2.18b has a ratio of 0.75.




Figure 2.19: Von Mises elastic stress plot on an Haynes 230 tube cross section simulated in SRLIFE at a  $\Delta T=190$  (C) and  $T_f=425$  (C). Fig. 2.19a shows results from a conduction-convection ratio of 0.25, and Fig. 2.19b has a ratio of 0.75.



Although the elastic crown stresses are observed to be similar, alternative ratios are still studied to determine if the creep relaxation is more or less significant than in the original work. Parametric studies are performed with a methodology that is almost identical to the original methodology used at ratio  $R = 0.5$ , seen in Fig. 2.1. The only difference in methodology is how the crown stress is calculated. In the original methodology, the crown stress is calculated using the average of the four quadrature points at the crown. Now, the crown stress is equal to the maximum stress observed across all four quadrature locations. This is considered to enhance the method's accuracy, but the lifetime predictions are not expected to be significantly different than those predicted with the old method. The creep-fatigue lifetime contours for four different alloys are shown in Figures 2.20– 2.23.

Haynes 230 in Fig. 2.20 shows the most interesting differences between ratios. As discussed previously, Haynes 230 shows a lifetime "valley" where lifetime dips before rising up again on the other side of the valley where stress relaxation noticeably improves the lifetimes. On the low temperature difference and fluid temperature side, the  $R = 0.25$  contours from 40-80 years are lower than the  $R = 0.75$  contours. However, this phenomena begins to switch as thermal conditions approach the lifetime valley. The 30 year lifetime contours for both ratios are close to one another on the low temperature side of the valley, but diverge to different locations on the high temperature side of the valley. Creep relaxation becomes increasingly significant as temperature conditions approach the lifetime valley. Consequently, it can be inferred that the two ratio cases will yield similar lifetimes when creep relaxation is not prominent because their von Mises crown stresses are expected to be similar. The two ratio cases will relax differently, however, because their initial stress gradients differ, as the representative case in Fig. 2.19 depicts.

As seen in Figures 2.21 and 2.22, the lifetime contour placements, on average, differ by only a few degrees. The  $R = 0.25$  case generally predicts lower lifetimes for the same location on the damage map. Haynes 282's damage map shows virtually no difference between the two conduction-convection ratios, with the exception of the 5 year lifetime.

Results indicate that different conduction-convection ratios have observable




Figure 2.20: Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Haynes 230 material. Only lifetime contours are shown, and stress reset limits are omitted.

differences in lifetime contour placements on the damage map, with the most significant differences occurring in Haynes 230 near its lifetime valley. However, Haynes 230's 40–80 contour placements are very similar for both ratio cases, and the other three tested materials also yield similar contour placements that are within a few degrees of one another at all locations on the damage map. Thus, although noticeable, the difference in contour placements is not large enough to cause substantial lifetime prediction inaccuracies between 10–80 years in Inconel 617, Haynes 282, and Inconel 740H, and 30–80 years on the low temperature side of the lifetime valley for Haynes 230. This work implements these damage map datasets into its damage tool, which is presented in the next chapter. Nevertheless, a single lifetime damage map derived at  $R = 0.5$  is determined to be sufficiently accurate for most cases.




Figure 2.21: Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Inconel 617 material. Only lifetime contours are shown, and stress reset limits are omitted.

## 2.6 FEA Versus Norton-Bailey Method for Stress Relaxation

A key modeling aspect differentiating the present work from other damage models is how it accounts for stress relaxation. Although materials with exceptional creep tolerance at the thermal operating conditions considered will not experience meaningful decreases in stress levels, relaxation can cause significant changes in the operational stress in limiting cases where lifetime is closer to 30 years. This fact is especially true for Haynes 230, where literature has demonstrated that seemingly minor changes in stress cause significant changes in creep damage González-Gómez et al. (2021). Unfortunately, estimating the relaxed stress state of a tube is a complex task due to the loading cycles of CSP towers, which are not expected to globally relax but also do not typically reset every day Barua et al. (2020). The current section compares a method that exists in literature to the extrapolation method




Figure 2.22: Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Inconel 740H. Only lifetime contours are shown, and stress reset limits are omitted.

presented in this work in Section 2.2. Our assessment finds the FEA method to be more conservative than the Norton-Bailey model.

Section 2.2 introduces the Norton-Bailey model, which is used by recent publications to estimate the crown stress relaxation due to creep. The Norton-Bailey model, as presented in González-Gómez et al. (2021), is shown in Equation 2.11, where  $\sigma_{eq}$  is the initial von Mises stress (Pa),  $E$  is the Young's Modulus of the material (Pa),  $n$  is the stress power law,  $A$  is the stress relaxation constant,  $Q$  is the apparent activation energy,  $m$  is an equation constant,  $R$  is the gas constant (J/mol-K),  $T$  is the temperature (K), and  $t_{stab}$  is the elapsed relaxation time (s).

$$\sigma_{eq}^{relax}(\sigma_{eq}, T) = \sigma_{eq} - E \left[ (1 - n) \left\{ \frac{(\sigma_{eq}/E)^{(1-n)}}{(1 - n)} - A E^n \exp \left( -\frac{Q}{RT} \right) \left( \frac{t_{stab}^{m+1}}{m + 1} \right) \right\} \right] \quad (2.11)$$

Authors in González-Gómez et al. (2021) calculated  $A$ ,  $n$ ,  $m$ ,  $Q$  parameters of




Figure 2.23: Damage map for two different conduction-convection ratios, 0.25 (blue) and 0.75 (red) and Haynes 282 material. Only lifetime contours are shown, and stress reset limits are omitted.

2.6880e-45 (1/Pa<sup>6.6</sup>-s), 6.6, 0, and 322 (kJ/mol) respectively for Haynes 230 based on data from Boehlert and Longanbach (2011).

This model is shown to perform well for a uniaxial stress state and constant strain case. However, the cosine flux profile and daily solar cycle that receiver tubes experience causes stress states that are multiaxial with variable loading. It is important, therefore, to assess how the Norton-Bailey model-predicted relaxed stress states compare to predictions from our FEA simulations. A set of 20 day FEA simulations performed for Haynes 230 is discussed in Section 2.2, and the simulation results contain von Mises crown stress values for every simulated timestep, such as the one shown in Fig. 2.24. Simulation results include cycles with peak daily temperature differences of  $\Delta T = 100\text{--}220\text{ }^{\circ}\text{C}$  and fluid temperatures of  $T_f = 300\text{--}550\text{ }^{\circ}\text{C}$ . Note that some extreme temperature cases within this range could not be simulated due to solver limitations. FEA simulation considers elastic strain as well as plastic strain from creep and yielding. SRLIFE models time-dependent creep strain

with the Kocks-Mecking method, as stated in Messner et al. (2022), and SRLIFE uses three sources of data for the Haynes 230 parameters. Each cycle is assumed to have 12 hours of operation with isothermal preheat and shutdown, as presented in Section 2.2.

Norton-Bailey model inputs are intended to match the conditions simulated in SRLIFE as closely as possible. A relaxation time ( $t_{stab}$ ) of  $20 \times 3600$  (s) is used. This divides out to an equivalent of 1 hr of relaxation per day, which is assumed by other analyses that use Norton-Bailey for CSP González-Gómez et al. (2021); Laporte-Azcué et al. (2021a). The starting stress is assumed to be the SRLIFE-simulated peak von Mises stress during the first cycle, and it can be seen in Fig. 2.24 as the red dot. Input of  $T$  for the Norton-Bailey model is the crown temperature. For the case shown in Fig. 2.24,  $T$  would be  $665^\circ\text{C}$ .

The differences between the SRLIFE and the Norton-Bailey model's 20th day stresses are shown over the entire range of temperatures in Fig. 2.25. In general, SRLIFE predicts less stress relaxation per case compared to this implementation of the Norton-Bailey model. Stresses are on average 17.1 MPa higher for SRLIFE results than the Norton-Bailey model. The maximum difference is 63.7 MPa, and the minimum difference is -3.0 MPa. A plot in González-Gómez et al. (2021) shows that  $\pm 10\%$  error in stress estimation results in creep rupture time error in the approximate range of -50%–200% for Haynes 230. The range of 1st day von Mises stresses for SRLIFE dataset discussed here is 202–265 MPa, and 10% of this range is approximately 20–25 MPa. Therefore, an average difference of 17.1 MPa between the two methods is certainly significant enough to cause differences in estimation of creep damage.

Creep damage, or the inverse of time to rupture, is related to stress according to the equation defined in Eq. 2.5. We conclude that SRLIFE is likely to predict lower magnitudes of creep stress relaxation when considering identical thermal cycles, material properties, first day stresses and 1 hour per day of Norton-Bailey relaxation time. Consequently, we also expect our damage model's assessment of creep damage for relaxed tubes to be more conservative than an assessment that utilizes the Norton-Bailey model, assuming that both methods' elastic and plastic stress calculations output equal 1st day stresses. Note that we cannot overgeneralize




Figure 2.24: Crown stresses for a 20 day FEA simulation of a Haynes 230 cross section via SRLIFE. Every cycle experiences a maximum temperature difference of  $140^{\circ}\text{C}$  and fluid temperature is at  $525^{\circ}\text{C}$ . Dots denote the first and last day's peak operational stresses.

and conclude that the current damage model is conservative compared to other methods with regards to lifetime estimates because of the myriad of assumptions required to assess receiver lifetime. These assumptions include, but are not limited to:

- cycle definition — the thermal loading variance in time
- material properties — including properties for thermal, structural, and damage analysis
- equivalent stress definition — the most typical is von Mises stress, as defined in Eq. 2.6
- prediction of yielding — conversion of elastic strain to plastic strain




Figure 2.25: Difference between SRLIFE-predicted and Norton-Bailey-predicted stress relaxation for a range of temperatures. FEA simulation considered 20 cycles of thermal loading typical of CSP applications. Daily operation time was 12 hours.

- safety factors — typical uses include artificially increasing the creep stress or strain range to ensure conservative results
- creep relaxation assumptions

Nevertheless, this section assess an important element of the presented damage model. Additionally, the damage model in this work is considered high fidelity because it is based on FEA – which is typically the comparison standard for analytical models proposed in literature. Future work should focus on reducing uncertainty in damage modeling associated with the modeling aspects listed here. In particular, there is a need for creep testing of receiver material subject to CSP-type loading cycles.



## **2.7 The Impact of Cycle Definition on Prediction Accuracy**

Thermal loading assumptions for the damage maps derived in this work assume an isothermal preheat period followed by a  $\Delta T$  profile that, temporally speaking, gradually increases and decreases in a sinusoidal fashion. Readers may reference Section 2.2 for the fully detailed explanation of our loading cycle definition. Based on damage results presented in this work, fatigue damage accumulated from solar variability is not expected to degrade lifetime for the majority of cases. However, creep damage is dependent on time at stress during the cycle. Therefore, our assumption of how crown  $\Delta T$  varies with cycle time impacts the calculated creep damage per cycle. This section seeks to quantify how much lifetime estimates can vary when different  $\Delta T$  time variance schemes are used. A Haynes 230 case with conditions similar to that of Crescent Dunes is considered because, for the materials considered in this work, Haynes 230 is identified as one of the most sensitive alloys to creep relaxation. We simulate thermal conditions throughout a clear-sky day for a Crescent Dunes case and extract the  $\Delta T$  temporal profiles for several different positions on the receiver. Lifetimes are predicted using existing methodology and results show that different definitions of the  $\Delta T$  time variance can noticeably affect predicted lifetimes. However, the results also quantify the uncertainty associated with the current damage model and therefore guide the selection of design or operation safety factors. Finally, this section proposes a method for modifying the methodology in Section 2.2 if readers wish to derive a damage map with lower bounds on lifetimes.

### **Cycle Comparison for Haynes 230 Case**

Damage map data is generated by assuming that the temperature difference ( $\Delta T$ ) between the inner and outer surfaces progresses throughout the day according to a sine/cosine wave. Modeling the cycle's temporal variance with a sine wave is a useful approximation that enhances the applicability of the results. However,

actual flux profiles are highly dependent on heliostat field operation. As discussed in previous work, the fatigue damage for a clear sky cycle is quantified in this work, as well as in *SRLIFE*, using the cycle's maximum strain range. This means that fatigue damage is generally determined by its peak experienced  $\Delta T$ , and fatigue damage is largely independent of the stress and temperature profiles during the rest of the cycle. In contrast, creep damage is impacted by the cycle's temporal variance in  $\Delta T$  because creep damage is dependent on the time at stress. Since, therefore, the lifetime of a tube will oftentimes be determined by the creep damage, it is worthwhile to consider the impacts of how we define the transient thermal loading profile. The transient loading profile of a cycle refers to how the tube's boundary conditions change with respect to cycle time. In this analysis, the authors focus on investigating the creep damage-related error introduced by a sine wave approximation of a receiver tube's transient loading profile.

Error is quantified by inputting tube  $\Delta T$  timeseries data from the full thermal model into *SRLIFE* and then comparing the creep damage results to those obtained when using a sine wave approximation. The peak inner and outer surface timeseries temperatures are obtained by simulating a solstice day for a CSP plant resembling Crescent Dunes. A sine wave approximation typically overpredicts lifetime compared to a solstice day operation with image size priority aiming. For the thermal operating points considered, the most damaging tube  $\Delta T$  profiles result in no more than 43.5% lower lifetimes when compared to a sine wave approximation. The least damaging tube  $\Delta T$  profiles actually result in higher lifetime, but lifetimes remain less than 8.8% higher than results from a sine wave  $\Delta T$  approximation. A sine wave approximation's overestimation of lifetime is due to either underestimating the time at stress during operation or overestimating the stress relaxation that occurs.

The time spent near the maximum stress condition for a given tube in a cylindrical receiver depends on where that tube is located around the circumference. To better understand this, tube  $\Delta T$  profiles from various receiver circumferential positions were simulated at various fluid temperatures and  $\Delta T$ s. Each  $\Delta T$  profile had the same magnitude, but their timeseries profile differed according to location — meaning that some profiles would experience more time at higher thermal stress

despite all the profiles having equivalent maximum thermal stress. This idea is shown in Fig. 2.26, which shows the timeseries  $\Delta T$ s for southward facing tube 0, eastward facing tube 10, and northward facing tube 20 compared to a sine wave assumption. A westward facing tube would have a mirrored profile compared to tube 10. Profiles show that the  $\Delta T$  reaches 0.4–0.8 of its maximum quickly — this is due to the plant operating with a pump with minimum allowable speed. The pump speed is adjusted to maintain salt outlet temperature and will vary directly with DNI. There exists a minimum flux level, therefore, that corresponds to the minimum required flux for plant operation, which is referred to as the cutoff DNI. Results showed that in regards to  $\Delta T$  profile shape, the tubes facing north and south experience the  $\Delta T$  profiles with the most time at near-maximum thermal stress. In contrast, tubes facing east and west experience profiles that more closely match the sine wave approximation, but exhibit a skewed distribution about solar noon.

Authors also investigated use of alternate  $\Delta T$  profile approximations, such as  $\sin^A(x)$  where  $A$  is some fractional exponent. As  $A$  decreases, the function increasingly matches a square wave. This approximation is also found to be prone to limitations, resulting in underestimations or overestimations, depending on the  $\Delta T$  timeseries plot it was compared to and the  $A$  value selected. Interestingly, square wave approximations are found to underestimate the damage because of their increased stress relaxation, which ultimately reduces creep damage. Thus, it is the authors' interpretation that both limits of  $\Delta T$  approximations experience limitations – either by underestimating the time at stress or by overestimating the creep relaxation. Therefore, the authors have elected to use a sine wave approximation and the associated error is now characterized for a cylindrical receiver case. There are two important outcomes from this analysis.

Firstly, this analysis offers insights on further improvement of damage methodology. Results demonstrate that any cycle approximation is expected to introduce error into lifetime predictions due to either overestimation of stress relaxation or underestimation of the time at peak stress. Therefore, the most generalizable cycle approximation is the conservative one. Authors suggest the following procedure for identifying the worst case temporal variation:

1. perform thermal analysis for a range of tube locations on a representative day. It is important to select the most representative aiming strategy.
2. calculate the  $\Delta T$  for each timestep, and normalize it by the peak  $\Delta T$  — this enables comparison of different tube cycles despite their different thermal operating points.
3. use the resulting  $\Delta T$  timeseries profiles as input to SRLIFE and perform a multiday parametric study of all the profiles at a few select thermal operating points.
4. extrapolate each cycle's future creep damage using the method described in previous reports
5. identify the cycle with the highest 20<sup>th</sup> day creep damage and define the normalized temporal variance for each thermal operating point based on this cycle

It is important to distinguish that authors are *only* suggesting that each tube's thermal cycle's *normalized* temporal variation be defined by the worst case cycle observed. This ensures that the resulting damage model's lifetime estimation is still dependent on the critical thermal operating point of the tube.

Secondly, this analysis quantifies uncertainty in the current damage model. For Haynes 230 in a Crescent Dunes scenario, the highest lifetime prediction error along the 30 year contour is at  $T_f = 475^\circ\text{C}, \Delta T = 135^\circ\text{C}$  with a lifetime that is 43.5% lower than the lifetime obtained assuming a temporal sine approximation. Reliability of the current damage map is maintained if a designer or operator uses a higher lifetime contour in the damage model to ensure the same actual lifetime. For example, in the case of a 30 year desired lifetime, an operator or designer can ensure tube lifetimes above 30 years if they use the damage model's approximate 69–70 year contour. In reality, this is a conservative safety factor that could be relaxed at thermal operating points other than the worst case point of  $T_f = 475^\circ\text{C}, \Delta T = 135^\circ\text{C}$ .




Figure 2.26: A sine wave approximation of the temperature difference's temporal variation compared to tubes at the south (0), east (10) and north (20) locations around the receiver for a solstice day operation assuming ideal image size priority aiming.

## 2.8 Summary of Damage Modeling Work

The present work both provides insight into suggested operating practices and can also serve as a useful receiver design tool. A FEA-based methodology for building damage maps of any given alloy is established. This methodology leverages full inelastic FEA analyses in SRLIFE and considers the important effect of stress relaxation at the tube crown. The power law extrapolation techniques predicts 20<sup>th</sup> day creep damage based on FEA simulation of the first four thermal cycles and proves to be sufficiently accurate for expected operating conditions and conservative in cases that experience extreme levels of stress relaxation. The Haynes 230 parametric study, using the methodology developed here, analyzes a receiver with baseline case geometry and assumes daily clear-sky flux profiles. Classification of baseline case tubes based on their critical temperature points and overlaying those operating points on the damage map shows that low  $T_f$  tubes will be operating with extremely high lifetimes, while higher  $T_f$  tubes display limited lifetime. Based on these results,

it would be advantageous to further concentrate the flux on low temperature panels and reduce the flux concentration on high temperature panels to achieve a more uniform lifetime prediction for all receiver tubes. Fatigue damage is never found to be significant when compared to the corresponding creep damage incurred for any temperature in the baseline case. The large difference in magnitude between incurred fatigue and creep damages for a clear sky case implies that additional cloud events causing smaller amplitude cyclic fatigue damage will not significantly affect a receiver tube's state of health. Note that this statement only applies to cloud effects related to strain range cycles. Operation that causes extreme temperature spikes that further aggravate creep damage should be avoided whenever possible.

Parametric analyses are performed for other candidate alloys, including Inconel 617, Haynes 282, Inconel 740H, SS 316H, and Incoloy 800H. Results show that only Inconel 617, SS 316H, and possibly Inconel 740H tubes experience significant fatigue damage. Lifetime estimates show that Haynes 282 and Inconel 740H exhibit excellent creep resistance for the simulated operating conditions and more extreme fluid temperatures and fluxes can be employed compared to the benchmark Alloy 230 case. The low yield strength of Incoloy 800H results in a highly reduced operational space unless the receiver tubes are operated in the stress reset region. Thus, with the primary exceptions of Inconel 617 and SS 316H, receiver tube fatigue damage is found to be negligible when compared to creep damage. Therefore, aiming strategies and mass flow control should focus on mitigation of creep, and such control strategies could increase cycling (if necessary) without significantly affecting tube lifetime.

Additional FEA campaigns at conduction-convection ratios of 0.25 and 0.75 are performed for Haynes 230, Haynes 282, Inconel 740H, and Inconel 617. Haynes 230 shows minimal differences for the 40–80 year lifetime contours, but some significant differences exist near the alloy's lifetime "valley" at intermediate thermal operating points. The other three alloys produce highly similar contours for each conduction-convection ratio, and the temperature points are shifted by only small magnitudes for 10–80 year contours. In general, a single lifetime damage map derived at  $R = 0.5$  is sufficiently accurate for most applications.

Cycle definition is shown to significantly affect lifetime predictions. Thermal and structural simulations are performed over a range of peak  $T_f, \Delta T$  magnitudes for cycle profiles with temporal variability corresponding to different locations on a cylindrical receiver. For Haynes 230 in a Crescent Dunes scenario, the highest lifetime prediction error along the 30 year contour is at  $T_f = 475^\circ\text{C}, \Delta T = 135^\circ\text{C}$  with a lifetime that is 43.5% lower than the lifetime obtained assuming a temporal sine approximation. A lifetime safety factor offers a simple method of accounting for uncertainty associated with cycle definition. Ultimately, a modified methodology is provided for development of plant-specific damage models if desired in future research.

The damage model developed in this work is distinct from other available models and is considered high-fidelity for several reasons. This damage model's most unique feature is that it is based on multi-cycle, finite element analyses results instead of analytical models with lumped assumptions. Utilizing FEA is expected to improve modeling accuracy of creep relaxation and, to a lesser degree, yielding effects. Simulation in SRLIFE also predicts lower magnitudes of stress relaxation relative to the Norton-Bailey model used in literature. Additionally, many creep-fatigue analyses apply safety factors to provide conservative design margin. This model is not developed with excessive conservatism in mind — instead it aims to estimate creep and fatigue damage as accurately as possible. The modeled cycle is generalizable and not dependent on preheat procedures, which allows it to offer insight on general operation and design. For this reason, the lifetime contours in Fig. 2.12 can be used to calculate maximum allowable receiver flux at any fluid temperature from 275–575 °C. The next chapter integrates this damage model into a multilevel simulation format that includes thermal and damage simulation tools.

**Future Work** Future damage modeling work should investigate model adjustments for different cycle lengths, simulate multiple thermal loading cycles at as high a cycle count as possible, perform experimental work related to material creep, and investigate alternative tube geometry.

### 3 SOFTWARE TOOLS FOR DAMAGE ANALYSIS OF CENTRAL RECEIVER DESIGN AND OPERATIONS

---

The primary focus of the current chapter is to introduce and demonstrate important tools developed over the duration of this work and how they might be applicable to future research. Present work builds on existing thermal and damage analysis tools to improve computation speed and accuracy for design and real-time applications. The existing receiver lifetime analysis tool, *SRLIFE*, is an invaluable software tool. Our damage tool relies on pre-computed receiver results from *SRLIFE* and cycle assumptions to make quickly-accessible lookup tables. The result is a design tool that is capable of providing lifetimes for large arrays of thermal operating points in less time than a single case can be simulated in *SRLIFE*. Also, the present work develops a real-time prototype tool that is focused on operations applications and a highly customized version of the Martinek et al. (2021) thermal model. The current Chapter provides an overview of the obtained FEA data in Section 3.1. Section 3.2 discusses how the damage tool is programmed in Python and explains logic for the core class methods. Section 3.3 analyzes a cylindrical receiver sized for Crescent Dunes conditions in cases of Inconel 740H, Inconel 617, Haynes 282, and Haynes 230 tubes. Section 3.4 introduces the real-time damage tool and the FEA basis for assumptions made. Modifications to the original thermal model are discussed in Section 3.5, as well as best practices for resolution parameters.

#### 3.1 FEA Data Availability

Both damage tools discussed in this chapter rely on the datasets obtained from extensive FEA simulation campaigns for all six alloys introduced in Chapter 2. Primary datasets are in the form of lifetime tables, which are visualized in this work as damage maps. However, our lifetime tables are built with several assumptions that are already described in the modeling methodology. We also make the foundational cycle and timeseries data available for the sake of repeatability and



support of future research that may wish to investigate other lifetime prediction assumptions. Timeseries data is defined as the unprocessed strain and stress data obtained at every timestep of a given case. Cycle data is defined as a processed product of timeseries data and refers to total creep or fatigue damages incurred by the simulated cycle. Additional data that is unused by present work is also made available – maximum stress observed over the entire tube cross section. As discussed previously, it is common practice to focus analysis on the damage occurring at the crown because this is the hottest location. However, future work may use maximum stress occurring at locations elsewhere than the crown for alternative analyses if desired. A brief overview of the cycle and timeseries data collected is provided in this section. Additionally, Table 3.1 defines the range of thermal conditions over which data is obtained. Present work conservatively estimates at least 3700 case results are available over the specified ranges.

**Cycle Data** Present work collected the following FEA cycle data for each temperature case:

- total cycle creep damage for cycles 1–4 of operation
- extrapolated 20<sup>th</sup> day creep damage, based on the method discussed in Chapter 2.
- Maximum observed strain range observed over the cycle at the crown location for days 1–4 of operation
- the damage ratio, if finite. Defined as the ratio of cumulative total fatigue to cumulative total creep damage over the tube’s lifetime.
- extrapolated lifetime, following methodology in Chapter 2.

**Timeseries Data** Present work also collected the following FEA timeseries data for each temperature case:

- von Mises stress at the outer surface crown

- maximum von Mises strain range at the outer surface crown
- von Mises strain at the outer surface crown
- temperature at the outer crown
- maximum von Mises stress observed in the tube cross section at that timestep
- maximum von Mises strain observed in the tube cross section at that timestep
- locations of the maximum stresses and strains
- temperatures of the maximum stresses and strains

These timeseries values include yielding and creep relaxation. Timestep resolution varied in the range of 1.5–3 minutes, with the most common resolution being 2 minutes. Recall that standard simulation length in Chapter 2 is 4 cycles of 14 hr periods.

Note that available data for material ratios other than 0.25, 0.50, 0.75 as seen in Table 3.1 is slightly more limited because they are from the initial study made by the author. It should also be noted that Haynes 230 at 0.5, Inconel 617 at 0.5, SS 316H at 0.43, and Incoloy 800H at 0.46 also have datasets containing elastic stress and strain timeseries results. Elastic stress data is used by this work to predict stress reset conditions, as presented in Chapter 2. Finally a set of 20 cycle simulation timeseries results are available for both Haynes 230 and Inconel 617.

The next section introduces damage tool code that uses sets of preprocessed lifetime data to analyze reliability of any given receiver-solar field case within dataset range.

## 3.2 Damage Tool for Design

Work in Chapter 2 discusses the full methodology for making creep-fatigue damage models in 6 candidate receiver alloys. The damage tool discussed here consolidates the damage model results into a usable package for development of receiver designs

Table 3.1: Fluid temperature and temperature difference ranges and data resolution for FEA data obtained in this work.

Material	Ratio	$T_f$ range (C)	$T_f$ res. (C)	$\Delta T$ range (C)	$\Delta T$ res. (C)
A230	0.25	275–575	25	50–300	10
	0.50	275–575	25	40–370	10
	0.75	275–575	25	50–300	10
A617	0.25	275–575	25	50–300	10
	0.50	275–575	25	50–300	10
	0.75	275–575	25	50–300	10
740H	0.25	275–575	25	50–300	10
	0.53	300–550	25	100–220	10
	0.75	275–575	25	50–300	10
A282	0.25	275–575	25	50–300	10
	0.44	300–550	25	100–220	10
	0.75	275–575	25	50–300	10
316H	0.43	300–500	12 or 13	100–180	5
800H	0.46	300–500	25	100–180	10

and aiming strategies. Key features and accessible data are discussed here, as well as the impact of interpolation on accuracy. The common functionality in all damage tool features is that they link fluid temperature and allowable flux/temperature difference to a lifetime value. The damage tool is made in Python code, and key features are referred to by their method names.

In general, each material has a series of accessible files that the damage tool uses as data sources. The file data is obtained using the methodology discussed in Chapter 2. The lifetime table used by the damage tool is specified when initializing the instance of the damage tool class — users specify the alloy, and the instance will load lookup tables based on its specified material. Users can also specify what ratio of conduction to convection temperature difference to assume and the desired interpolation method. Readers should reference Section 2.5 for more information

on the impacts of ratio assumption. By default, the damage tool assumes the most conservative ratio available. Ratio selection should be based on the range of ratios observed via thermal analysis of receiver conditions.

As a reminder, lifetime tables used by the damage tool are built on the following underlying assumptions: (i) a tube section's thermal operating point refers to the maximum temperature difference at solar noon and the corresponding fluid temperature at that location in the flowpath (ii) all cycles are identical and are defined in Section 2.2 (iii) the tube has already undergone its break-in phase, meaning that creep relaxation and yielding have already occurred.

## Overview of Available Class Methods

**Lifetime Assessment Methods** The most commonly used damage tool methods are the `get_LTEs` methods, which either use a radial basis function interpolator or 2-dimensional linear interpolator. The process for making the interpolator is introduced here before the `get_LTEs` methods are discussed. First, the lifetime lookup tables are loaded, which contains lifetimes as a function of  $\Delta T$  and  $T_f$ . The lookup tables have resolution of data according to Table 3.1. Some of the less extreme thermal cases have infinite lifetimes predicted. To improve the accuracy of the interpolator, a subarray is created that only contains the finite lifetime data points from the original lookup table. Once this is done, the actual interpolator is generated using either the `RBFInterpolator` or the `LinearNDInterpolator` function from SciPy's `interpolate` package Virtanen et al. (2020). A significant benefit of using these functions is that the input data can have irregular intervals. Additionally, required interpolation times are found to be relatively short — on the order of 20 ms for 100 sets of  $\Delta T, T_f$ . The linear interpolator uses a fill value of 100,000, which will primarily apply to data points that are predicted by `SRLIFE` to be infinite. These infinite life operating points can be pictured as located left or below the contours on a damage map. Note that the linear interpolator also may predict extremely high lifetimes for the extreme top right points of a damage map because `SRLIFE` simulation typically failed there. However, this region is a highly improbable

operating location and the lifetime assessment methods that consider stress reset or corrosion would return zero lifetime here, regardless. The radial basis function uses 100 nearest neighbors and a linear kernel.

There are three lifetime prediction methods that use the lifetime interpolator. The output is a lifetime prediction for every thermal operating point set ( $\Delta T, T_f, R$ ) input to the method. First, **get\_LTEs\_no\_limits** returns the estimated tube lifetime in years and neglects the corrosion and stress reset regions. Lifetime is solely predicted with the expected creep-fatigue lifetimes presented in Chapter 2. Second, **get\_LTEs** is similar to the first method, but **get\_LTEs** also estimates the film temperature and proximity to the stress reset region. Detailed corrosion and stress reset models are outside of the scope of this work, so the method returns a predicted lifetime of 0 years for every thermal operating point that is in either the corrosion region or the stress reset region. The regions are discussed in more detail in Chapter 2. Subfunctions for estimating the film temperature and proximity to the stress reset region are discussed later in this section. Third, the **get\_LTEs\_w\_penalties** method exists for the purpose of optimizing or using a heuristic with the stress reset and corrosion regions as constraints. This method has the base functionality of **get\_LTEs** but heavily penalizes thermal operating points that enter the stress reset region and/or the corrosion region. The calculated penalty from incurring stress reset is added to the calculated penalty of operating in a high corrosion region. This penalization method provides a soft limit with a sharp gradient and removes incentive for a thermal operating point to proceed farther into corrosion or stress reset conditions. Heuristic testing presented later in this work finds this method to be useful for cases of over-utilized receivers, where it is virtually impossible to avoid flux violations due to insufficient receiver area.

A helper class called **materialLimits** is also written to provide validity bounds on the damage model. The damage model in Chapter 2 does not consider the interactions of corrosion and creep-fatigue damage. In fact, this is an ongoing research question. The general strategy in literature is to identify critical film temperatures between the salt and the inner tube surface, above which excessive corrosion is experimentally observed. A given alloy's critical film temperature is therefore

regarded as a thermal constraint. For example, the critical film temperature limit is used as one of the flux density limits in Sánchez-González et al. (2020). Thus, critical film temperatures are enforced for each alloy as listed in Chapter 2. An **estimate\_Tfilm** method is written that conservatively approximates the maximum film temperature of any thermal operating point as the tube wall temperature at the inner crown. Equation 3.1 formally states this, where  $T_{\text{film}}$  is the film temperature,  $R$  is the ratio of conduction temperature difference to convection temperature difference,  $T_f$  is the fluid temperature, and  $\Delta T_{\text{total}}$  is the temperature difference between the crown and the tube's backside/ $T_f$ .

$$T_{\text{film}} = T_f + \frac{\Delta T_{\text{total}}}{1 + R} \quad (3.1)$$

Finally, lifetime assessments of a receiver panel are based on sub-assessments on each of the panel's tube lifetimes. Tube lifetime is predicted by performing a lifetime assessment of each its axial location's thermal operating points. The axial location with the lowest predicted lifetime determines the entire tube's life. In a similar way, the lowest expected tube lifetime in a panel is assumed to limit the panel's lifetime. This means that the authors are assuming a single tube failure warrants an entire panel replacement. Assuming the tube with the lowest lifetime determines the entire panel's lifetime is commonly seen in literature, such as in Gentile et al. (2024). The damage tool method **calc\_minimum\_panel\_LTEs** takes an array of lifetimes as input and, based on the number of tubes per panel and panels, predicts the lifetime of each panel in the receiver.

**Design Contour Generation Method** A procedural code for translating damage map data into accessible lifetime contours is presented here. The method name is **make\_contour\_function\_from\_interpolator**. Users input a desired lifetime in years, and the output is a 1-dimensional interpolator. The interpolator returns the maximum allowable temperature difference for any fluid temperature. Chapter 4 describes how this method is used for receiver design. A brief overview of its procedure is provided here, however, because it is a damage tool class method. The

general steps to producing a design contour are: (a) building a damage map with a lifetime contour level specified by the user (b) extracting the fluid temperature and total temperature difference vertices and (c) making a linear 1D interpolation function via **SciPy** methods. In regards to procedure methodology, it is important to build the damage map with an interpolation method consistent with that used to interpolate lifetimes (i.e. **get\_LTEs**, etc). Otherwise, the design contour does not exactly lie on the intended lifetime level. Users should also note that the method does not create contour functions for lifetimes less than 30 years because contour vertices for some materials are problematic for making a monotonic contour function. However, design for lifetimes below 30 years is not recommended nor expected to be frequently desirable.

**Material Limits Class** The helper class **materialLimits** aids in locating the thermal operating point's proximity to corrosion and stress reset regions as well as assigning a penalty value to operation in these regions.

The method **check\_SR** returns a boolean value based on if the ratio of elastic to peak plastic stress for the cycle is greater or less than 2. Visually, this means that the method returns "True" if the thermal operating point is in the solid-red colored region of the damage map. To accomplish this, the method consults a precomputed lookup table of elastic to peak plastic stress ratios specific to the specified material. A radial basis function interpolates this lookup table with a setting of 20 neighbors and a linear kernel. Another method, **penalize\_SR**, is responsible for returning penalties for thermal operating points in the stress reset region. The penalty for any thermal operating point determined to be in the stress reset zone is determined using Eq. 3.2, where  $\Delta T_{SR}$  is the total temperature difference at which the stress ratio becomes greater than 2 for the thermal operating point's  $T_f$ , and  $\Delta T$  is the thermal operating point's total temperature difference.

$$LTE = -100 \times (\Delta T_{SR} - \Delta T) \quad (3.2)$$

**penalize\_corrosion** has a similar purpose as **penalize\_SR**. This method uses a

similar equation, shown in Eq. 3.3, where  $\Delta T_{\text{corr}}$  is the maximum critical temperature difference for a given fluid temperature, above which excessive corrosion is expected to occur.

$$\text{LTE} = -100 \times (\Delta T_{\text{corr}}(T_f) - \Delta T) \quad (3.3)$$

The variable  $\Delta T_{\text{corr}}$  is calculated by rearranging Eq. 3.1 into Eq. 3.4, where  $T_{\text{crit,film}}$  is a fixed critical film temperature for a given material, and  $T_f$  and  $R$  are the fluid temperature and ratio of conduction to convection for a given thermal operating point.

$$\Delta T_{\text{corr}} = (T_{\text{crit,film}} - T_f) \times (1 + R) \quad (3.4)$$

Error and lifetime prediction accuracy in the stress reset zone are expected to be significantly lower than elsewhere on the damage maps shown.

### 3.3 Tool Application: Lifetime Profiles of a Cylindrical Receiver

It is common for creep-fatigue analyses in literature to analyze either Gemasolar, located in Spain, or Crescent Dunes, located in the United States. Analyzing a Crescent Dunes case, therefore, serves as a comparison benchmark for this tool. A preliminary damage analysis of a commercial plant resembling Crescent Dunes is performed in Chapter 2, and this section extends that analysis for several reasons. First, we demonstrate the design damage tool's capabilities to predict lifetimes in four alloy cases: Inconel 617, Haynes 230, Haynes 282, and Inconel 704H. Second, the majority of the present work focuses on flat receivers. Although it is necessary to focus on simplified flat receiver case, it is also important to demonstrate that the scope of the damage tool is not limited to flat receivers. Third, the current section illustrates entire lifetime profiles of the receiver, which is significantly more detail than the analysis in Section 2.2 provides. Results show that for the default



aiming strategy assumed, the most limiting lifetimes are expected to occur in the final panel of each flowpath. Haynes 230 is the exception to this, however, as the closest intersection of thermal operating points and the lifetime contours occurs at an intermediate fluid temperature. Inconel 740H and Haynes 282 demonstrate superior performance and are capable of handling even higher flux densities than the magnitudes seen in this case.

A flux profile is simulated for a solar noon position on summer solstice, which is often used as a design basis. Receiver geometry, heliostat field details, and aiming strategy remain consistent with what is presented in Section 2.2's thermal analysis. The flux profile shown in Fig. 3.1 is obtained via *SolarPILOT*. Aiming strategies for flat receivers are further discussed in the next chapter. It is interesting to note that a main difference between each receivers' flux profiles is that there are only offset regions at the top and bottom of each panel on a cylindrical receiver, while a two flowpath flat receiver typically also has a sizable low flux region on the outermost panels to avoid spillage losses.

A full thermal model analysis is performed, and the post-processed results yield Fig. 3.2 and Fig. 3.3, which are the total temperature difference and the fluid temperatures respectively. Two thermal model analyses are performed, one for tube conductivity corresponding to Haynes 230 material, and one for conductivity of Inconel 740H. Temperature-dependent conductivity models for the other two alloys were not readily available at the time of simulation. Three tubes are simulated per panel. Only the fluid temperature and total temperature difference profiles of the Inconel 740H case are shown because the two cases only slightly differ — the maximum observed temperature difference is 220°C for 740H and 215°C for Haynes 230.

The damage tool presented in this Chapter takes the calculated fluid temperatures and total temperature difference results and predicts lifetimes at each simulation node. As a result, so-called "lifetime profiles" show the predicted lifetime of each axial subsection of the tube. Thus, lifetime resolution matches that of the thermal model's settings used in Martinek et al. (2021): 3 tubes  $\times$  14 panels in the horizontal direction, and a vertical resolution of 50 points. A tube's minimum

lifetime is the lowest subsection lifetime along its axial length. Using the material original to Crescent Dunes, Haynes 230, yields a damage profile as seen in Fig. 3.4. 740H's lifetime profile is shown in Fig. 3.6. The thermal solution is assumed to be sufficiently representative of the thermal solutions obtained if Inconel 617 or Haynes 282 temperature-dependent conductivity models were available. Table 2.1 reports higher conductivities for Inconel 617 and Haynes 282 than that of Haynes 230, so we can assume this analysis will yield conservatively higher estimates of total temperature differences per location. Therefore, lifetime profiles for Inconel 617 and Haynes 282 are inferred by using the thermal solution obtained from assuming Haynes 230 metal conductivity. Inconel 617's lifetime profile is shown in Fig. 3.5, and Haynes 282's profile is shown in Fig. 3.7.

Table 3.2 summarizes results from all four alloys. Columns of the table include: the global minimum lifetime on the receiver ( $LTE_{min}$ ), the critical total temperature difference ( $\Delta T_c$ ) it occurs at, the critical fluid ( $T_f$ ) temperature it occurs at, and the number of panels that will fail prior to 30 years ( $N_{sub30}$ ). Similar to what is discussed in Chapter 2, Haynes 230 and Inconel 617 have significantly lower creep-fatigue endurance compared to Haynes 282 and Inconel 740H. The Inconel 617 receiver has the lowest minimum lifetime of 9.2 years, and Haynes 230 yields a similar prediction of 14.5 years. In contrast, the other two alloys are not expected to have any failures before 30 years. It can be observed in Table 3.2 that the fluid temperature corresponding to the receiver's minimum lifetime location occurs at the same point for every alloy except Haynes 230. Also, the difference in critical fluid temperature point for Haynes 230 versus the other alloys is too large to be explained by the difference in conductivity models used. Instead, one must consider the difference in general shape of the damage map contours between Haynes 230 and the other alloys. Haynes 230 has a significant elbow in its 30–80 year contour lines, where the slope changes significantly. The critical point in the Haynes 230 receiver is approximately where the elbow is located on the damage map.

To illustrate this, Figure 3.8 shows the thermal operating point of every calculation node on the receiver. As far as the operating points are concerned, we see that the seven panel's operating points on both flowpaths collapse on top of each




Figure 3.1: Incident flux map for a Crescent Dunes cylindrical receiver. Heliostat field aiming with image size priority is used. The modeled material is Inconel 740H.

other due to symmetry in the flux profile. The peak flux locations are located at the highest temperature difference coordinates, and the the long "tails" of each panel are formed by the low flux regions at the beginning and end of each panel. The current aiming strategy gradually decreases the flux on on each panel of the flowpath. It is helpful to visualize the highest temperature difference points as forming a "front", with the right-most points on each panel connected to form an approximately linear line with negative slope. For this style of aiming, a receiver's *most* critical point occurs where the lifetime line slope intersects with the slope of the receiver's critical points. Damage maps with fairly linear lifetime contours have their lowest lifetimes in the panels with higher fluid temperatures because this is where the line of critical panel points are closest to intercepting the lifetime contours. This explains why the three alloys other than Haynes 230 have the same critical lifetime location. Haynes 230's lifetime contours form a more prominent elbow that more closely intersects the thermal operating points at an intermediate fluid temperature compared to the other alloys.

To summarize, the current section demonstrates the design damage tool's capabilities by analyzing thermal results from a Crescent Dunes case. The present work is capable of producing lifetime profiles, which provide helpful visualization and




Figure 3.2: Total temperature difference at the crown at 50 axial positions per tube length for a Crescent Dunes cylindrical receiver, Heliostat field aiming uses image size priority. The modeled material is Inconel 740H.




Figure 3.3: Salt fluid temperature profile for a Crescent Dunes cylindrical receiver, Heliostat field aiming with image size priority. The modeled material is Inconel 740H.




Figure 3.4: Lifetime profile a Crescent Dunes cylindrical receiver when default *SolarPILOT* strategy is used. The thermal and damage model material is Haynes 230. Note that the colorbar is saturated at 100 years.




Figure 3.5: Lifetime profile a Crescent Dunes cylindrical receiver when default *SolarPILOT* strategy is used. Profile is made using an Inconel 617 damage model to interpret a thermal solution common to Figures 3.4 and 3.7. Note that the colorbar is saturated at 100 years.




Figure 3.6: Lifetime profile a Crescent Dunes cylindrical receiver when default *SolarPILOT* strategy is used. Profile is made using an Inconel 740H thermal and damage models. Note the log scale.




Figure 3.7: Lifetime profile a Crescent Dunes cylindrical receiver when default *SolarPILOT* strategy is used. Profile is made using an Haynes A282 damage model to interpret a thermal solution common to Figures 3.4 and 3.5. Note the log scale.

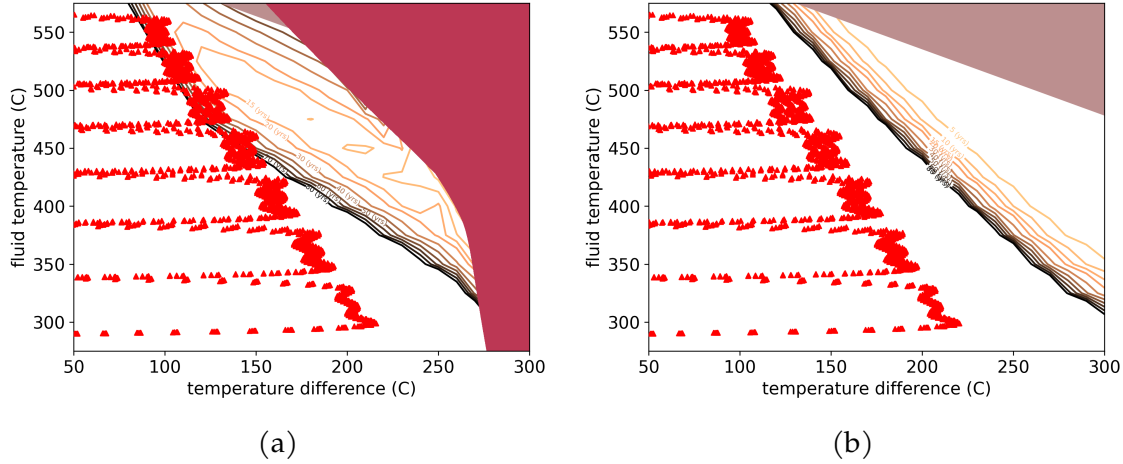


Figure 3.8: All of the evaluated thermal operating points of the receiver (red triangles) overlaid on the damage maps presented in Chapter 2 for Haynes 230 (a) and 740H (b).

Table 3.2: Lifetime results for Crescent Dunes-type cylindrical receiver. Materials include the original Haynes 230 as well as three additional ones.

material	$LTE_{\min}$	$\Delta T_c$	$T_{f,c}$	$N_{\text{sub30}}$
A230	14.5	132.2	500.0	6
A617	9.2	102.0	560.5	8
740H	572.6	105.6	560.5	0
A282	1531.3	102.0	560.5	0

data that could be used for flowpath optimization focused on modifying aspect ratio, panel width, or tube geometry to achieve improved lifetimes. We leverage our damage tool to test alternative alloy performance. The four analyzed alloys can be ordered in terms of reliability from most to least required panel replacements: Inconel 617, Haynes 230, Inconel 740H, Haynes 282. If used, the last two alloys would not be expected to limit the lifetime of the receiver. For the the default aiming strategy in *SolarPILOT* on cylindrical receivers, we expect the final panels of each flowpath to yield the lowest lifetimes for Haynes 282, Inconel 617, and Inconel 740H.

For Haynes 230, however, the limiting lifetime occurs at an intermediate panel.

Findings illustrate how the intersection of thermal operating points and lifetime contour on a damage map identify where the receiver reliability is most limited. Effective aiming strategies and or receiver design should focus on aligning thermal operating points to account for the slope of the lifetime contours. The next section introduces a prototype that uses the same underlying datasets as the design damage tool currently being discussed but for a real-time application.

### 3.4 Damage Tool for Real-Time Operation

A real time decision tool that takes thermal input and outputs instantaneous damage rates is created by using the damage model datasets obtained from Chapter 2. This section provides an overview of the procedure, presents the underlying database, and provides results from a sample case. Further details regarding this model can be found in a supporting document titled: *Instantaneous Damage Model for Solar Power Tower Receiver* by Frankie Iovinelli, who performed undergraduate research under the supervision of the present author as well as Assistant Professor Mike Wagner. The real time damage tool presented here demonstrates a valuable prototype for dealing with creep damage during variable operating conditions.

**Overview of Methodology** This real time damage tool's functionality and methodology is distinct from damage tool presented in Section 3.2, although it relies on the same underlying FEA data. These two damage tools have different intended uses — design versus real time operation. Output from the design damage tool is broader in scope and includes lifetimes, penalized lifetimes, design contours, and damage maps. In contrast, this real time tool output is focused on tracking state of health and calculating instantaneous creep damage rate. Each model predicts lifetime differently. The design damage tool assumes all cycles are identical with the  $\Delta T$ ,  $T_f$ ,  $R$  input defining conditions at the peak of a gradually increasing and decreasing (cosinusoidally) cycle. The real time damage tool's lifetime output is simply based on the instantaneous damage rate of the thermal operating point input provided



and the current state of health. Real time damage tool usage is currently restricted to the Alloy 230 case and does not consider the effects of stress reset or corrosion, although additional material options could be added if desired.

The real time decision tool assumes several things. Tube lifetime is based on the loading conditions that occur at the crown, which is the tube location subject to the highest absorbed flux. This tool assumes that creep damage dominates fatigue damage. Stress values and relaxation behavior is based on SRLIFE-generated data. Authors assume that the receiver tubes have completed a 20 day thermal break-in period, during which the thermal stresses have relaxed according to 20 cycles at their nominal thermal operating point. If the tube's thermal operating point shifts during real-time, this tool assumes that the receiver accordingly relaxes to that higher thermal loading condition instantly. While this assumption is a simplification of the actual relaxation problem, the resultant error is expected to become insignificant after a sufficient number of cycles, thereby justifying the assumption. Finally, note that this real time damage tool predicts remaining lifetime by assuming that future hours of operation are identical to the conditions currently being experienced.

A primary challenge for real-time instantaneous creep damage modeling is determining what the actual tube stress is for a given thermal operating point. The most conservative option for predicting the tube's stress is to always calculate the elastic stress, but stress relaxation due to creep has been shown by this work and other related literature to be non-negligible in some alloys. Actual stress in a given tube cross section is likely to be lower than the predicted elastic stress because of relaxation that occurs from yielding and material creep. Current work refers to the initial tube stress relaxation as the "break-in" period, and Chapter 2 discusses this concept in greater detail. The issue of stress relaxation is further complicated if in variable conditions, a tube encounters higher flux conditions than occurred during its break-in period. Damage modeling in Chapter 2 is capable of predicting the 20<sup>th</sup> day residual stresses for a tube subject to 20 identical cycles at a given thermal operating point peak. If we assume that a tube will relax to the new break-in state of the higher flux conditions, we can use existing FEA data to predict the new stress

state of the tube.

Relaxed stress is calculated based on an established relationship defined in Eq. 3.5, where  $\sigma_{\text{actual}}$  is peak elastic-plastic-creep stress, and  $\sigma_{\text{residual}}$  is the residual stress for that case after all stress relaxation has occurred. The slope,  $m$ , is determined via analysis of a series of 20 day FEA simulations performed in SRLIFE, which are presented in Chapter 2.

$$\sigma_{\text{actual}} = \sigma_{\text{elastic}} - m * \sigma_{\text{residual}} \quad (3.5)$$

It is important to note that the slope  $m$  is a best fit line of data that excludes FEA cases that experienced stress reset. This is not expected to impact the tool's utility because the stress reset region is not a desirable operation region. The resultant  $m$  value and its fit for the data considered are shown in Fig. 3.11. It can be observed that the difference between elastic and actual, relaxed stress for any case is approximately equal to the residual stress times some fitting constant. The fitting constant is selected to be consistent with FEA data, and it accounts for numerical error and/or limitations of underlying assumptions.

Assuming a tube always relaxes to its most extreme thermal loading case is shown to approximate FEA results well for the demonstrative case seen in Fig. 3.9 and 3.10. The shown FEA simulation is a multicycle case of 30 days/cycles in total. The first 20 days are typical cycles with increasing and decreasing loading profiles defined by a peak fluid temperature of 300 (°C) and temperature difference of 100 (°C). The next five days have similar load profiles, but more extreme magnitudes: 450 (°C) fluid temperature and 160 (°C) temperature difference. Lastly, the loading cycles return to 300 (°C) / 100 (°C) for the final 5 cycles. Figure 3.9 shows the daily crown stress cycles. Negligible residual stress is seen in the first 20 days, but the extreme cycles from 21–25 cause residual stress to accumulate. Residual stress does not decrease, even during the milder cycles at the end of the test. Therefore, we use the residual stress as a means to track the tube's relaxation history and predict stress for any thermal operating point of lower fluid temperature and/or temperature difference than the tube has already experienced. Figure 3.10 demonstrates the




Figure 3.9: Crown stress versus time for a Haynes 230 receiver tube simulated for 30 cycles in *SRLIFE*. Finite element analysis accounts for elastic and inelastic strains caused by the loading conditions. The initial 20 days and the final 5 days have identical thermal loading cycles, but cycles 21-25 had higher, more extreme loading profiles.

important impact of relaxation history on the tube cross section — the maximum crown stress is lower than initial cycles because of the accumulation of residual stress that opposes the stress induced during operation hours.

**Workflow Overview** The damage tool uses the following procedure to predict lifetime and instantaneous damage rate. First, the fluid temperature and incident solar flux are obtained to provide model input. Flux can be readily converted to a wall-to-fluid temperature difference by solving the lightweight thermal model discussed in Section 3.5. Next, elastic stress is determined by referencing a pre-computed lookup table with the given  $T_f$  and  $\Delta T$  operating points. Alternatively, elastic stress can be calculated directly by using relations from Logie et al. (2018). Then, the residual stress, which is caused by plastic yielding and creep relaxation,




Figure 3.10: Maximum stress at the crown location for each day cycle of a 30 cycle FEA simulation.

is obtained from another precomputed lookup table.

Once the tube's actual stress is predicted, two useful outputs can be obtained: instantaneous damage and remaining lifetime. Instantaneous damage represents the damage incurred per second at the current operating point. This value is calculated by correlating the thermal operating point and the actual stress to a rupture time with a Mendelson-Roberts-Manson parametrization. The inverse of rupture time can be interpreted as the instantaneous damage rate. Integrating the damage rate over time results in the total damage to the tube that is incurred during a specified period. Remaining lifetime is calculated by tracking the total damage that has already occurred in the tube and also considering the instantaneous damage rate. This relationship is defined in Eq. 3.6, where  $LTE_r$  is the remaining lifetime in seconds,  $D_{c,inc}$  is the damage that has already been incurred, and  $d_{c,i}$  is the instantaneous damage rate.




Figure 3.11: Relationship between the difference of elastic and actual peak stresses versus the residual stress for a series of FEA results data. As expected, the majority of cases follow a linear trend.

$$\text{LTE}_r = \frac{1 - D_{c,inc}}{d_{c,i}} \quad (3.6)$$

The numerator quantifies the tube's state of health, with a tube beginning its life at 0 incurred damage and ending its life when the incurred damage equals 1.

**Application Example** A demonstration of this plot is shown in Figures 3.12 and 3.13. Fluid temperature ( $T_f$ ) and temperature difference ( $\Delta T$ ) are set for a series of points on the receiver, which are represented by the black dots. Figure 3.12a shows the maximum temperature seen on the outside surface, which is the sum of each point's fluid temperature and temperature difference. The demonstration is for a generic simple aiming scheme, where heliostats are all aimed at the center of the receiver to maximize optical efficiency. Figure 3.12b shows the actual stress for all receiver points. Higher temperature locations experience more stress relaxation and therefore have a lower predicted actual stress. The resulting instantaneous damage rates and remaining lifetimes are shown in Figures 3.13a and 3.13b respectively.




Figure 3.12: The input total temperature ( 3.12a) and resulting actual stress ( 3.12b) for an arbitrary thermal case.




Figure 3.13: The resulting instantaneous damage rates ( 3.13a) and lifetimes ( 3.13b) from the thermal case shown in Figure 3.12.

**Summary and Future Work** In summary, the real time damage tool discussed in this section is capable of quickly translating thermal data into lifetime and damage predictions for any given receiver with available creep data. Implementation of pre-computed lookup tables allow for low computation time. When paired with a computationally efficient thermal model, this damage tool is expected to be extremely useful for helping operators make decisions about the thermal operating points of the receiver during real time.

Note that the model is proven as a promising prototype, but has not been rigorously tested against large numbers of multicycle FEA simulations. Recall that  $m$ 's slope is derived from Chapter 2 campaign of 20 cycle FEA simulations for Haynes 230, which all have maximum temperature differences during the cycle of 100–220 °C and exclude cases approaching stress reset conditions. Thus, suggestions for future work include additional FEA simulations similar to that seen in Fig. 3.9 but with different combinations of cycles with varying periods, temperature ranges, and cycle counts. Whenever possible, extended cycle FEA simulations provide valuable insight into tube material behavior after repeated loading. If residual stress measurement methods are available, experimental testing would also be extremely informative. Finally, note that this real-time model focuses on quantifying creep damage rates. This focus is valid for materials that experience significantly higher levels of creep damage compared to fatigue damage during their lifetime. A future iteration of this real-time damage model could incorporate a strain ranges lookup table into the work flow, where the table contains data for each thermal operating point simulated in Chapter 2.

### 3.5 A Modified Thermal Model Tool

The present work uses a previously developed thermal model to predict wall temperatures and fluid temperatures along a given receiver's flowpath. A thermal model was developed in Martinek et al. (2021) to perform sensitivity analysis and optimization of receiver design parameters and flux density for central receivers using high temperature chloride salts with inlet/outlet temperatures of 500°C/735°C.

Model methodology is described in Martinek et al. (2021), including the various radiation calculations, energy balances, and correlations. Although the present work considers receivers that use nitrate salt instead of chloride salt, the thermal model is generalizable for any liquid heat transfer fluid. In general, the present work uses the thermal model to output the receiver tubes' thermal states for a given flux profile. Several modifications were made to the original model, including an adaptation from circular to flat receiver geometry, integration with *CoPylot*, improvement in the numerical discretization scheme, and model simplification for reduced computation time. This section discusses these modifications in more detail. Results demonstrate that the thermal tool can be readily integrated with the damage tool to advise and assess aiming strategies.

## **Modifying the Thermal Model for Flat Receivers**

The thermal model was originally built to solve thermal problems for cylindrical receivers. Only one minor change is made to adapt this model to flat receivers: the panel width calculation. Previously, each panel width was calculated by dividing the receiver's circumference input value by the number of panels input value. Now, the panel width is simply calculated by dividing the receiver's width input value by the number of panels input value. It should be noted that the radiation heat transfer calculations still assume every tube is adjacent to two others — one on each side. Although this is true for cylindrical receivers, flat receivers' left and right-most end tubes are only in radiation contact with one other adjacent tube. On the tubes' other side there is increased exposure to the ambient. This small difference only affects 2 tubes and is not expected to significantly impact the thermal solutions. Results are presented later in this chapter that show agreement between the predicted power from the optical model and the incident power calculated by the thermal model.

## **Integration with *CoPylot***

A more flexible *SolarPILOT* interface is now integrated with the existing thermal model. Previously, thermal model flux inputs were limited to generic flux pro-



files. Authors can now test custom receiver flux profiles, allowing for design and evaluation of damage informed aiming schemes. Additionally, the heliostat field layout can be set to a specific file, allowing for consistent fields to be used during simulation campaigns.

Thermal model workflow begins by obtaining a flux profile as input, with the model performing an optical simulation with *SolarPILOT* to predict the flux profile for a given receiver geometry. *SolarPILOT* is a tool developed at NREL used for CSP design. Users can perform a variety of tasks related to the design and optimization of heliostat fields and receivers as well as simulate receiver flux profiles (Wagner and Wendelin (2018)). *SolarPILOT*'s graphical user interface, which has full functionality, is not callable by other programs and therefore cannot be efficiently implemented with the thermal model. Instead, the thermal model uses an application programming interface (API) to access *SolarPILOT*. Originally, the thermal model used 'PySSC', which accesses the System Advisor Model Simulation Core (SSC) compiled library version of *SolarPILOT*, and only has access to basic functionality of *SolarPILOT*. Users provide key inputs, such as receiver and heliostat geometry, design thermal power, and simulation time of day, then *SolarPILOT* generates a heliostat layout and returns the optical efficiencies and relative fluxes at the desired simulation points. Two key features lacking from this library version of *SolarPILOT* are the ability to assign heliostat layouts from existing files and to provide user flux profiles for controlling the aiming algorithm. Firstly, the authors desire the ability to manually assign *SolarPILOT* layouts, thereby keeping solar resources constant to ensure fair comparison between receiver designs. Secondly, the flux profiles currently under evaluation by this project must be requested manually through the "user\_flux\_profile" feature in *SolarPILOT*.

For these reasons, another existing API called *CoPylot* is implemented that uses a compiled *SolarPILOT* library with significantly more functionality compared to 'PySSC'. A basic workflow overview of the thermal model is shown in Fig. 3.14, where users provide inputs, a version of *SolarPILOT* generates an incident flux map, then the thermal model imposes that flux distribution on the receiver and solves for the complete temperature profile of the receiver. It can be seen that the change in

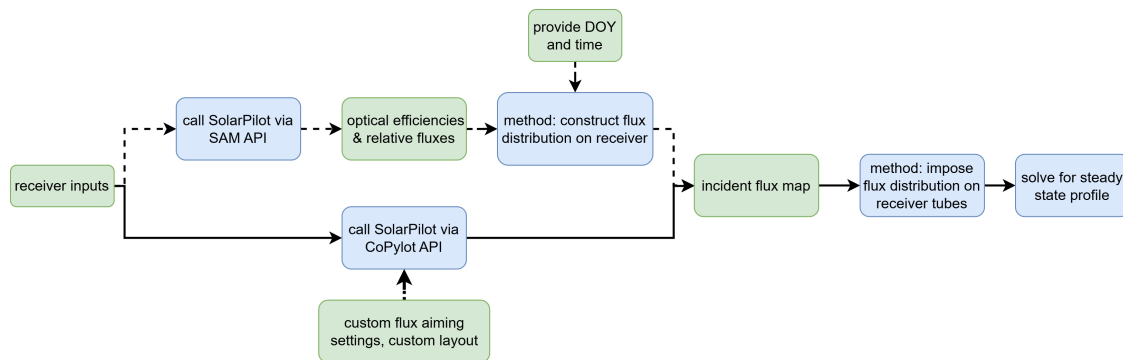


Figure 3.14: High-level process diagram of thermal model contrasting the workflow differences between using the PySSC (dashed line) versus *CoPilot* API (solid line) methods for generating the incident flux map.

*SolarPILOT* API versions only affects the beginning of the model's workflow. Users now have total control of the flux simulation, including the ability to provide a pre-generated heliostat field and request custom flux patterns. A cylindrical receiver case is simulated using the PySSC and *CoPilot* API workflows to verify successful integration of *CoPilot* and the NREL thermal model. Comparison of solution results show only negligible differences in receiver power, incident power, and mass flow.

The thermal model's enhanced functionality is demonstrated in Fig. 3.15 and Fig. 3.16. A flat receiver case is simulated with thermal design power of  $180 \text{ MW}_{\text{th}}$ , height of 15 m, width of 18 m, tower height of 170 m, 9 panels, and a single flowpath. A heliostat layout capable of achieving approximately  $200 \text{ MW}_{\text{th}}$  of incident power on the selected receiver is generated, and a custom aiming profile is specified. The custom aiming profile requests a flux level of 100% on the left half and 10% on the right half. Results show that this flux boundary condition input directly translates to the thermal model, as the absorbed flux results plotted in Fig. 3.16 show the left and right halves of the receiver with contrasting flux levels. The fluid temperature accordingly quickly increases in the first half but rises noticeably more slowly on the low flux side. As a result of this functionality addition to the thermal model, custom profiles can now be specified as an input.




Figure 3.15: Demonstration of *CoPylot* integration for a flat receiver: The input user flux profile requested a high, uniform flux on the left side and low, uniform flux on the right side.

## Numerical Discretization Methodology

Damage model-informed design and aiming for central receivers requires the integration of three model stages: an optical model, a thermal model, and a damage model. Each model's output is fed as input to the next simulation stage. For a given heliostat field and aiming strategy, the primary output from the optical model is a receiver flux profile, which typically is a 2 dimensional array of discretized values. That flux profile output is then translated to tube fluxes inputs in the thermal model. The thermal model is numerical, which means each simulated tube has a finite set of computation nodes, according to the selected radial, circumferential, and axial resolution. Additionally, only a reduced number of tubes are actually modeled in each receiver panel to reduce computation time. Thermal model output is primarily an array of fluid temperatures and corresponding total temperature differences; the array dimensions correspond to the number of simulated tubes and the axial points on each tube. Lastly, the damage model outputs a lifetime array whose dimensions match those of the thermal model. Consequently, it is crucial to understand how




Figure 3.16: Thermal model results showing the effects of a heavy left side flux concentration. This simulation is for a single flow path, flat receiver case at design point.

the various model resolutions impact the model accuracy and how the various model resolutions impact the accuracy of later simulation stage accuracy.

The current section investigates the following sensitivities related to numerical resolution: i) optical model's incident power versus output array resolution ii) thermal model output versus its four discretization parameters iii) thermal model output versus optical model output resolution. In reference to the third point, this work identifies an important impact of optical model output array resolution on thermal and damage model results. The author selects resolution parameters based on the results presented here.

**Impact of Optical Model's Array Resolution** The first simulation layer discretization is the square grid that characterizes the flux output from *SolarPILOT* simu-

lations. Optical simulation in *SolarPILOT* uses analytical hermite expansions to predict each heliostat's image projection on the receiver Wagner and Wendelin (2018) which is then evaluated at the user-selected number of grid points along the vertical and horizontal dimensions of the receiver. It can be expected, then, that a finer grid will result in higher fidelity output from *SolarPILOT*, while a coarser grid will result in some inaccuracies. A study is conducted to determine the minimum acceptable grid resolution.

The heliostat field generated for this study is sized with a maximum field radius of 13 times the tower height, with is 170 m. Mirror parameters are consistent with *SolarPILOT* default settings. A layout is generated in *SolarPILOT* that is capable of 200 MW absorbed power when a uniform aiming strategy is used. Receiver dimensions are 15 m height and 18 m width. Solar conditions are assumed to correspond to solar noon on a summer solstice day. A custom flux profile that represents a typical desired aiming strategy is input to *SolarPILOT* via the *CoPyLOT* interface, and the simulated flux profile output is processed to determine the peak flux and total incident power on the receiver. This process is repeated at different grid point resolutions. The results are shown for maximum flux (Fig. 3.17a) and incident power (Fig. 3.17b).

Results in Fig. 3.17 show that peak flux decreases with increasing grid resolution. Resolution greater than approximately 10 points per square meter will yield similar peak flux values. Peak flux varies slightly as resolution increases, and this may be explained by the change in aiming strategy itself: as the evaluation resolution increases, the image size priority algorithm may produce different results even for the same desired flux profile input. Thus, the evaluation locations and the algorithm are simultaneously changing, although only by a minor amount for the latter. Incident power also decreases with increasing grid resolution and converges more slowly than peak flux. Similarly to peak flux, more than 10 pts/m<sup>2</sup> should be used to maintain fidelity. 10-20 pts/m<sup>2</sup> is assumed to be adequate for expressing optical model output. Even at 20 pts/m<sup>2</sup>, the slope of the incident power is nonzero, however. Therefore, although this resolution is assumed to adequate, readers should keep in mind that the actual incident power may be lower than the simulated

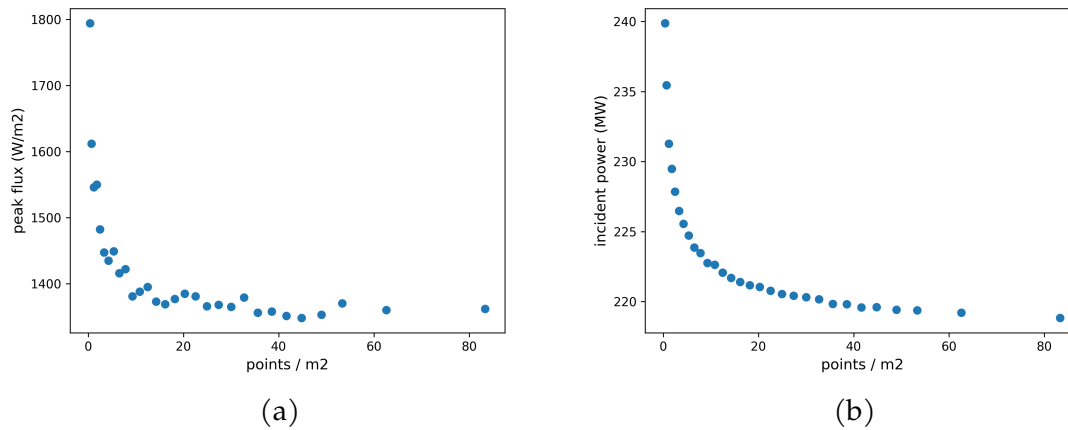


Figure 3.17: Resolution study results showing peak flux (Fig. 3.17a) and incident power (Fig. 3.17b) as a function of grid resolution for a 270 m<sup>2</sup> case and a square grid scheme.

value. Higher resolutions are possible but deemed excessive from a computational resource perspective. This study cannot be interpreted separately from its impact on other systems, however. The next sections provide further information on the impact of optical result resolution on thermal and damage results.

**Considerations for Thermal Model Numerical Discretization** This work investigates four simulation parameters: the circumferential resolution, the radial resolution, the axial resolution, and the number of simulated tubes per panel. Nodal resolution for each parameter is held to a default of 79, 5, 50, and 1 respectively unless it is the parameter of interest. For each parameter, a range of values are tested in the same thermal model case, which is identical to the case described for the optical model convergence study. Thermal model input, the flux profile from the optical model, is provided in two resolutions corresponding to the range identified in the convergence study. Using two different optical model resolutions, 10 and 20 pts/m<sup>2</sup> allows us to identify any interactions between the optical output resolution and thermal model output magnitudes. An optical model resolution of 10 pts/m<sup>2</sup>

on the 270 m<sup>2</sup> receiver corresponds to a square grid of 52 pts per dimension, and the higher resolution of 20 pts/m<sup>2</sup> corresponds to 74 pts per dimension. The ultimate goal of the entire simulation suite is to convert flux profiles into expected tube lifetimes. Therefore, main convergence metric for the thermal model parameters is the final output, tube lifetime. More specifically, the minimum tube lifetime from each panel. The secondary convergence metric for thermal model parameters is the absorbed thermal power, which serves to place tube lifetime in context. These two metrics are evaluated as a function of each parameter resolution.

Note that mass flow is controlled in this study to maintain a constant outlet temperature from each flowpath. Thus, the heat transfer coefficient will increase with increased absorbed power. An increased heat transfer coefficient will yield a decreased total temperature difference in the tube, and potentially yield a higher predicted lifetime.

**Axial Resolution Study** Axial resolution is defined as the number of nodes used by the thermal model along the direction of the tubes. This direction visually equates to the vertical direction when looking at a receiver. Axial resolution is tested in the range of 10–200 nodes in increments of 5. The thermal model solves the full heat transfer problem, and lifetime at each and absorbed thermal power are subsequently calculated for each case. Results for both optical input flux profile resolutions are shown in Figures 3.18 and 3.19. The change in absorbed thermal power is less than 2 MW<sub>th</sub>, so the mass flow and resultant heat transfer coefficient does not largely vary for the axial resolution study.

Tube lifetimes vary in the range of 30-1500 years in this case, depending on the panel. Consequently, Fig. 3.18 shows all lifetimes normalized by their largest observed value, and it can be seen that there is a range of nodal sensitivities to axial node count. The most sensitive tubes experience oscillation as resolution increases. This effect can be explained by understanding how the thermal model interprets flux inputs. Flux inputs are interpreted from the supplied flux distribution to thermal model nodes in the z direction via 1 dimensional interpolation. Increasing resolution shifts the thermal model's axial node positions relative to

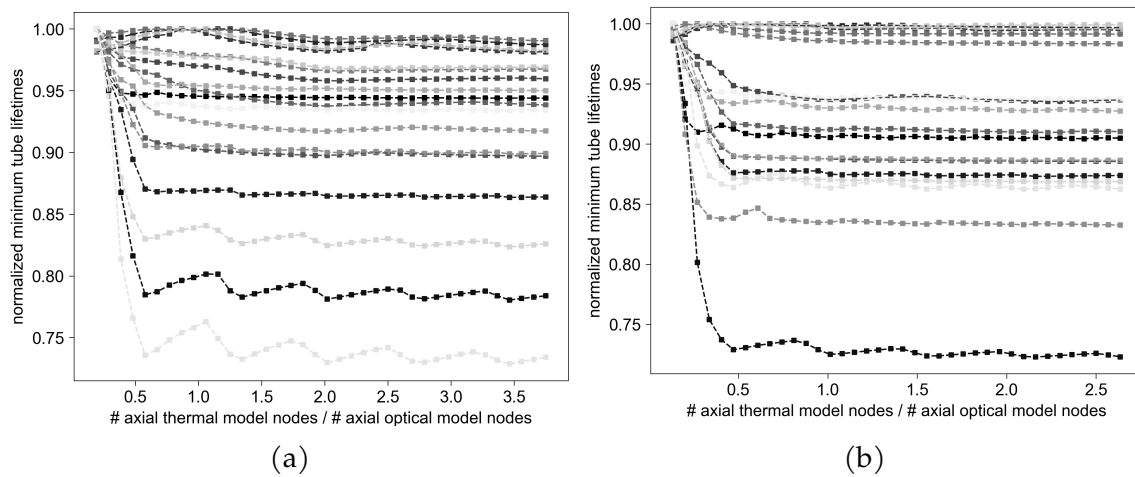


Figure 3.18: Resolution study results for minimum lifetimes normalized by the largest lifetime observed in the study for that tube. The x axis is the ratio of thermal model axial nodes to flux profile nodes in the axial direction. Figure 3.18a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.18b) shows results with 74 points per dimension. Both studies assumed a  $270 \text{ m}^2$  case and a square grid scheme.

their nearest flux profile input point. Oscillation is reduced in Fig. 3.18b when compared to Fig. 3.18a because the difference between flux points decreases. Both of the minimum tube lifetime plots show that lifetimes most significantly vary with resolution for ratios of axial thermal model nodes to flux profile points less than 0.5. Absorbed thermal power shows a similar sensitivity to resolution, with the largest slope between points at ratios less than 0.5. Therefore, it is recommended to use 1–2 times the axial thermal model nodes relative to the optical model’s flux profile resolution.

**Circumferential Resolution Study** Circumferential resolution, in the context of this thermal model, refers to the number of computational nodes around each tubes’ circumference for a given radial distance from the tube’s center. A study analogous to the axial resolution study is conducted, with circumferential resolution ranging from 30–300 nodes. Figures 3.20 and 3.21 show the minimum panel lifetimes and






Figure 3.19: Resolution study results for absorbed thermal power. The x axis is the ratio of thermal model axial nodes to flux profile nodes in the axial direction. Figure 3.19a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.19b) shows results with 74 points per dimension. Both studies assumed a 270 m<sup>2</sup> case and a square grid scheme.

absorbed thermal power results, respectively. Note that each figures' x-axis is not a ratio of nodes to flux profile nodes, as seen in Fig. 3.18 and 3.19. Consideration of the node ratio is assumed to not be relevant in the discussion of circumferential resolution because this resolution pertains to scaling the nearest single incident flux value from the optical model to all the other circumferential locations on the tube. This scaling task is in contrast to axial 1D interpolation, which interpolates between the two nearest incident flux points.

Figure 3.20 shows that, similarly to the axial study findings, increasing circumferential resolution leads to lower lifetimes, with the predicted lifetimes' rate of change significantly decreasing with node counts greater than approximately 100. The reason for tube lifetime decline with increasing node count must be interpreted in the context of the significant decrease in absorbed thermal power, seen in Fig. 3.21, which results in a lower system mass flow and salt-side heat transfer coefficient. A decrease in thermal power will degrade the salt-side heat transfer coefficient and cause higher total temperature differences across the tube cross section. The sys-




Figure 3.20: Resolution study results for minimum tube lifetimes in each of 18 simulated receiver panels. Each panel value is normalized by the maximum lifetime that panel had during this study. Figure 3.20a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.20b) shows results with 74 points per dimension. Both studies assumed a 270 m<sup>2</sup> case and a square grid scheme. Outer tube diameter is 50.8 mm.




Figure 3.21: Resolution study results for absorbed thermal power. Figure 3.21a) shows results for flux profile input resolution of 52 points per dimension, and Fig. 3.21b) shows results with 74 points per dimension. Both studies assumed a 270 m<sup>2</sup> case and a square grid scheme.

tem of circumferential nodes approximates the co-sinusoidal incident flux profile typical for solar receiver tubes. Since the incident flux varies with the angle of the tube's surface relative to the flux, it decreases along the tube surface for locations further away from the tube crown. In fact, the tube crown is the only location that experiences the full magnitude of the optical model's predicted incident receiver flux — all other tube locations experience a flux that is a fraction of the crown's flux. Figure 3.21 shows that thermal power decreased by almost  $25 \text{ MW}_{\text{th}}$  in the node count range 30–300. A low number of circumferential points does not accurately approximate each tube surface's incident flux profile, especially at the crown; this misrepresentation of the tube flux causes a significant overestimation of the absorbed thermal power. A circumferential node count of at least 100 is recommended, with higher node counts yielding thermal powers that differ by less than approximately 1–2%. The reader should note that the studied tubes have a diameter of 50.8 mm, so a node count of 100 would correspond to approximately 1.6 mm of circumference between each node.

Conclusions are found to be largely independent of input flux profile resolution, because the 52 and 74 nodes/dimension plots show only minor differences in terms of behavior and magnitude. This finding is consistent with the author's expectations because the optical resolution will typically be significantly coarser than the tube spacing. Thus, the same flux profile point likely covers all circumferential positions on a tube at a given axial position.

**Radial Resolution Study** Radial resolution refers to the number of nodes dispersed through the thickness of the tube for any given circumferential point on the tube. Radial resolutions are tested in the range of 2–30. It can be seen in Fig. 3.22 that minimum tube lifetimes and absorbed thermal power are not significantly sensitive to radial node count. A node count of 5 or more is considered sufficient. Tube thickness for this study is set to 1.25 mm, so a count of 5 nodes is equivalent to 4 nodes per mm.




Figure 3.22: Resolution study results for each tested panels' minimum lifetime normalized by that panel's highest magnitude value during the study (Fig. 3.22a) and absorbed thermal power (Figure 3.22b). Results are for flux profile input resolution of 74 points per dimension.

**Impact of Number of Simulated Tubes Per Panel** Typical receiver design uses a series of 50.8 mm diameters tubes to span distances several meters wide. For example, an 18 m wide flat receiver would be comprised of approximately 350 tubes. Simulation of all 350 tubes is not computationally efficient. Instead, a representative number of tubes in each panel are simulated and then the tubes' absorbed power is scaled up to represent the total absorbed energy in each panel. The location of the simulated tube in the panel determines what incident flux it experiences. If a single tube per panel is simulated, then the tube locations are in the center of each panel. For 3 simulated tubes per panel, the tube locations are on the left panel edge, center, and right panel edge. Higher numbers of simulated tubes per panel are distributed between these general bounds so that the set of simulated tubes are situated equidistantly along the panel width. The impact of simulated tube count on minimum panel lifetimes and the receiver's absorbed thermal power is studied by varying the simulated tube count between 1–13 tubes per panel, in steps of 2 tubes. Results are shown in Figures 3.23 and 3.24 for an input flux profile resolution




Figure 3.23: Resolution study results for each tested panels' minimum lifetime normalized by that panel's highest magnitude value during the study (Fig. 3.23a) and absorbed thermal power (Figure 3.23b). Results are for flux profile input resolution of 74 points per dimension.

of 74 points per dimension.

Figure 3.23 demonstrates that there is little difference between 3 or higher tube counts in regards to tube lifetimes. We can conclude that the most important modeling detail to capture is that the left and right edge tubes are modeled, because these tube lifetimes typically bound the lifetime range of the panel. Minimum tube lifetimes for each of 13 simulated tubes in each panel are shown in Fig. 3.24, with left edge and right edge tube lifetimes highlighted by unique plot markers. It can be observed that on the receiver's left side, the right-most tubes are the lowest lifetimes, and vice versa on the right side. Panel 8 is the only exception to this observation. Recall that the modeled input flux profile and flowpath configuration is typical of the present work — two paths of molten salt start in the panel's center and serpentine out to the left and right receiver edges. The informed flux profile attempts to fully utilize the receiver by concentrating high flux in the low fluid temperature center and decreasing flux as the fluid temperature increases towards the receiver edges. Thus, flux is generally decreasing in the direction of the center




Figure 3.24: 13 simulated tube lifetimes normalized by each tube's highest minimum lifetime value. The modeled receiver has panel widths of 1 m and the input flux profile has 74 points per dimension.

out to the edge. For any given panel, its tube that is closest to the center will experience higher flux, and will therefore be the most limited in lifetime. This logic breaks down at central panels, however, where the flux is actually closer to a constant value instead of varying with horizontal position. Therefore, it is crucial to consider the left and right edge tubes when evaluating panel lifetimes.

A critical number of 3 simulation tubes per panel will also ensure accurate thermal power calculations — Figure 3.23b demonstrates that increasing the number of simulated tubes per panel does not have a significant effect on the receiver's absorbed thermal power, as the total power differs by less than  $1 \text{ MW}_{\text{th}}$ . Simulating 3 tubes per panel is assumed to be adequate unless the panel width significantly differs from the conditions simulated here.

**Resolution Interaction between Optical Model and Thermal Model** As discussed previously, the left and right edge tube locations of each panel are oftentimes the determiners of panel lifetime. Informed aiming strategies seek to achieve discrete flux levels for adjacent panels, based on the local fluid temperatures. Work in this section shows that the existing thermal model requires a code modification to be able to capture flux differences between adjacent panels. Additionally, a minimum input flux resolution in the horizontal dimension proportional to receiver width is required to prevent interpolation-induced discrepancies between the simulated flux and the thermal model's interpreted flux.

The original thermal model, as presented in Martinek et al. (2021), takes a discrete grid of flux profile points as input. Model code assigns incident flux values to each tube by determining its horizontal position and then interpolating between the grid points for every axial position on the tube. Each tube will have unique incident flux values if one tube per panel is simulated because the tube positions are in the centers of each panel. However, the present work's simulated tubes per panel study demonstrates the importance of considering left and right edge tubes. Original thermal model code approximated panel edge tube locations by assuming they were at the exact extents of the panel. As a result, two panels' adjacent edge tubes would actually be assigned the same incident flux because they are technically at the same position. The aiming goal of the present work is to achieve as discrete or 'stepped' flux profiles as possible per panel so that allowable flux closely follows the local panel's salt temperature. A code modification is necessary, then, to avoid adjacent edge tubes having the same incident flux profile because having equivalent edge tube incident flux profiles undermines our aiming strategy. The code solution is to slightly offset edge tubes from the exact edges of the panels. The new code assigns tubes on panel edges approximately one tube diameter away from the nearest panel edge. Thus, the neighboring edge tubes of two adjacent panels are located a distance of 2 diameters offset from each other. It is important to account for this offset when deciding the resolution of the input flux profile. A flux profile with grid points coarser than this neighboring tube gap will negate the code modification. Fortunately, the magnitude of error will decrease with increasing resolution. In

the present work, where tube diameter is typically 50.8 mm, a grid resolution of approximately 10 points/m is required. This critical value can be approximated with the following relation:  $1 \text{ pt}/(2 \times D_o) \text{ m}$ , where the denominator is the distance between two neighboring edge tubes.

In the ideal optical case, where stepped flux is achieved for each panel with no image spillover to the adjacent panels, critical resolution is the number of points required by the model to assign edge tubes incident flux that is not impacted by interpolation. This is illustrated in Fig. 3.25. A flux map is designed for a receiver with 15 m height and 18 m width, and consisting of 24 panels. Figure 3.25b shows that the tube lifetimes are sensitive to the number of horizontal flux grid points used until more than 10 points/m are used. At this point, each panel has enough flux grid points to no longer require the adjacent tube's flux grid points to interpolate flux for its own edge tubes. The critical resolution of 10 points/m holds true for other panel widths, also. An identical test is run with panels of 1 m wide and 10 points/m is still the critical resolution. Therefore, the present work concludes that 10 points/m are required for the 50.8 diameter tube geometry being considered. It should be noted that 10 points/m translates to significantly higher optical grid resolution than is typically used with *SolarPILOT*.

The critical resolution only ensures that flux maps with no image spillage to adjacent panels maintain their fidelity in the thermal model. Actual flux simulations have gradual flux changes between panels, and thermal model output and resulting lifetimes are therefore more dependent on the optical accuracy rather than optical resolution. The present work does not find increasing optical flux grid resolution to yield lifetime convergence for actual flux simulations. Instead, tube lifetimes are found to fluctuate between high and low lifetimes with increasing optical flux grid resolution. This may be explained by recalling that the image size priority method used in *SolarPILOT* for obtaining actual flux maps produces slightly different flux profiles when different grid point resolutions are used. Regardless, the lifetime fluctuations from *SolarPILOT*-produced flux profiles at different grid point resolutions decrease with decreasing lifetime. Horizontal grid resolution in the range of 72–720 points results in lifetimes that differ by no more than 4.5 years for tubes



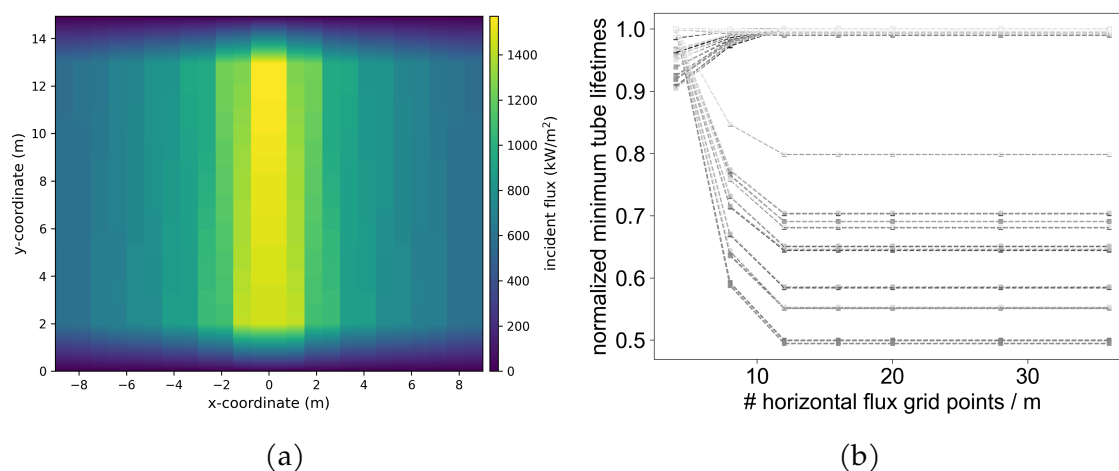


Figure 3.25: Ideal fluxmap in reference to the discrete flux levels between each panel (3.25a) and tube lifetimes normalized by the highest minimum tube lifetime in that tube (3.25b).

that are predicted to have sub-40 year lifetimes.

*SolarPILOT* is used in this work as a quick assessment tool rather than a detailed simulation. If ultra-high accuracy is desired, a ray tracing program such as *SolTrace* may be used. High input flux resolutions are computationally expensive and do not yield significantly different values. Additionally, the input flux profiles should not be such coarse resolution that ideal optical input flux profiles will be misconstrued to be unideal optical inputs. Therefore, if low flux input resolution is provided to the thermal model, a "blocking" technique is implemented that repeats elements. For example, an input flux profile resolution of 5 points/m is blocked by a factor of 2, which increases the necessary resolution above the critical value. This method artificially increases the resolution after *SolarPILOT* has simulated, thereby reducing computational cost while still ensuring that the flux profile output from *SolarPILOT* is identically interpreted by the thermal model.

**Summary of Recommended Simulation Settings** The present work identifies five simulation settings that can potentially affect lifetime estimates. First, the optical

model simulates an incident flux profile based on a given aiming strategy. Results show that setting *SolarPILOT* to use 10-20 evaluation pts per  $\text{m}^2$  of receiver area is recommended to reliably calculate peak flux and total incident power. A code modification is implemented because the thermal model's previous method of interpolating flux input onto each receiver tube is found to cause discrepancies between the input flux profile and the flux profiles on each tube. Ideal profiles, with discrete fluxes on each panel, require at least 10 points per horizontal meter, to prevent flux interpolation-induced discrepancies. A blocking function is introduced that artificially increases input flux resolution without requiring an ultra-high resolution optical simulation. Second, the number of axial nodes on each of the thermal model's simulated tubes is shown to be positively related to the absorbed thermal power and negatively related to minimum lifetime. Repeating the same axial node study for two different flux input resolutions also demonstrates the two settings are somewhat coupled. A minimum ratio of 0.5 axial nodes per vertical flux points is recommended. Third, the number of circumferential nodes is shown to be significantly and negatively correlated to absorbed thermal power as well as tube lifetimes, although readers should note that, in this study, the heat transfer coefficients degrade with decreasing thermal power. This decrease in salt-side convective heat transfer may be driving the significant decrease in lifetime with increasing number of circumferential nodes. At least 80 circumferential nodes are recommended to accurately model the absorbed thermal power. Fourth, the number of radial nodes is shown to not largely affect the lifetime or power results. Fifth, the number of simulated tubes per panel does not largely affect absorbed power, but minimum panel lifetimes are typically limited by the edge tubes' lifetimes. As a result, high fidelity lifetime predictions for a receiver include at least 3 tubes per panel — one tube on each edge and one in the middle.

These simulation settings are adopted as best practice and used when assessing aiming strategy design performance. However, reducing computational expense of the thermal model is useful for fast assessments such as in cases of heuristic, real-time, or optimization applications. Modifying the model to be lightweight and fast is the topic of the next section.

## Lightweight Model for Faster Computation

The full thermal model presented in Martinek et al. (2021) is configured to thoroughly simulate the entire collector system, including generating heliostat fields, performing optical simulation, calculating thermo-elastic stresses, and iteratively solving for the mass flow that yields the target salt outlet temperature for a given incident power. Solution time varies, but is typically at least several seconds. For instance, a 50 axial node, 5 radial node, 79 circumferential node, 3 tubes per panel, flux profile resolution 52 x 52 case requires 6.3 s to solve at one time point. This section discusses how the full model is configured to be a less resource intense, "lightweight" thermal model for the purpose of high speed applications, where the thermal model may need to be called on the order of 100–1000 times per run.

The lightweight model is a version of the full model and is configured to receive an input flux array and output steady state thermal results for a single case. The full model guesses a mass flow and iteratively solves for the mass flow and coupled steady state solution within their specified tolerances. In contrast, the lightweight mode assigns a mass flow and solves the resultant steady state solution within specified tolerances. Thus, the full model's solution will have varying mass flow to achieve the target salt outlet temperature. Oppositely, the lightweight model will have varying target salt outlet temperature for a fixed mass flow case. Solving the lightweight model's thermal problem requires fewer iterations and ensures that all comparisons are on the same mass flow basis and therefore have the same salt side heat transfer coefficients. Two additional modifications cut simulation time: not performing an optical simulation each time the model solves and instantiating a receiver object without instantiating its steady state analysis parent object. Directly solving the receiver object's thermal problem also removes post processing procedures performed by the parent object which were not being fully utilized for lightweight applications. These modifications yield a lightweight model typically capable of solving in less than 1 s. Note that this time does not include the time to setup the model. However, it is assumed that the model will only require a single initialization for high speed applications and therefore the simple thermal solving

time is the most relevant metric.

Further model simplifications also reduce speed without sacrificing significant accuracy. The lightweight model does not require an initial guess to accurately solve, which reduces the number of required solution iterations to reach the specified tolerance. Test cases show that the combination of approximating the heat transfer through the tube wall as 1D and assuming an average conductivity only results in 2.5% lifetime error while reducing the execution time by a factor of 3x. Note that a 1D dimensional assumption in this model implies that axial and circumferential heat transfer is neglected. One reason this assumption does not largely affect accuracy is due to the thin walls of the tube when compared to the overall circumference: 1.25 mm versus 50.8 mm. However, reducing the number of circumferential nodes, as discussed previously, is not worth the significant tradeoff in accuracy that it incurs. Finally, lightweight model solve time almost linearly scales with the number of axial nodes used — reducing axial node count by half also reduced the solve time by a factor of 2x. The result of these modifications is a lightweight thermal model that can solve a thermal problem on the order of 100–500ms, depending on the axial resolution and number of simulated tubes per panel.

## Post Processing Thermal Model Output

Thermal model output from the steady state solution is stored as attributes of the receiver object. Each simulated tube contains an axial x circumferential x radial array of wall temperature as well as a bulk salt temperature at each axial position. This section briefly elaborates on how total temperature difference  $\Delta T$ , fluid temperature ( $T_f$ ), and  $\Delta T$  ratio ( $R$ ) are extracted from the full solution array.

The maximum outer wall temperature ( $T_{ow,max}$ ) at each axial position is identified by finding the maximum over all circumferential nodes at the outermost radial node. Maximum inner wall temperatures ( $T_{iw,max}$ ) are identified by finding the maximum circumferential temperature at the innermost radial node layer. Both the maximum inner and outer wall temperatures are assumed to occur at the crown position where the highest absorbed flux occurs. The minimum wall temperature

$(T_{w,\min})$ , which is assumed to be at the back of the tube, is identified at each axial position by finding the global minimum over the circumferential and radial dimensions at that axial node.

Total temperature difference, which is the x-axis used for damage maps in the present work, is stated in Eq. 3.7. Stated simply, total temperature difference is the measured between the front and back outer surfaces of the tube.

$$\Delta T = T_{ow,\max} - T_{w,\min} \quad (3.7)$$

The temperature difference ratio slightly impacts lifetime predictions as discussed in Section 2.5. This ratio represents the relative magnitudes of the outer wall-inner wall difference and the inner wall-fluid temperature difference, and is shown in Eq. 3.8

$$R = \frac{T_{ow,\max} - T_{iw,\max}}{T_{iw,\max} - T_{w,\min}} \quad (3.8)$$

Note that the minimum wall temperature difference on the backside is observed to closely match the fluid temperature at a given axial node. Therefore,  $T_{w,\min}$  in Eq. 3.8 could be replaced by  $T_f$  if desired.

All postprocessing discussed here exists in receiver class methods and a thermal model helper module. Postprocessing yields an array of  $\Delta T$ ,  $R$ , and  $T_f$  values for every axial position. The damage model can then predict lifetime for each axial position using these thermal model outputs.

## 3.6 Summary of Damage Tool and Thermal Model Contributions

The current chapter presents several necessary tools, algorithms, and practices that directly enable improvement of receiver reliability. Code is designed to be data base flexible and readily versatile. We use an open-source environment and intend the code serve as a starting point for future research.

Present work has produced over 3700 FEA results by implementing the methodology proposed in Chapter 2 to simulate six different alloys over a meaningful domain of fluid temperatures and temperature differences. Data is available in the form of unprocessed timeseries data, processed cycle data, and estimated lifetime tables. In this Chapter, we present Python code to utilize lifetime tables for design and operation settings. The damage tool class can be used in central receiver analysis to identify expected low lifetime panels and inform aiming strategies. Extrapolation methods are sophisticated, accounting for corrosion and stress reset regions if desired. Additionally, design and aiming routines can use a modified extrapolation method that heavily penalizes operating points located in the corrosion and stress reset zones. The next chapter utilizes the introduced method for generating design contours from given lifetime tables.

This damage tool is applied to analyze lifetime profiles for a cylindrical receiver. Default aiming in *SolarPILOT* for a cylindrical receiver with Crescent Dunes geometry and 565 MW<sub>th</sub> design power produces flux levels that gradually step down from high flux panels to low flux panels. The lifetime profiles for receivers made of Haynes 230 and Inconel 617 predict 6 and 8 required panel replacements before 30 years of constant operation. In contrast, the panel lifetimes for Haynes 282 and Inconel 740H receivers far exceed 30 years. These results are consistent with the damage maps presented in Chapter 2. Interestingly, the final panel is predicted to fail for all alloys except Haynes 230, which encounters failure farther upstream in the flowpath.

The damage tool prototype for real-time operation presented here offers a solution for variable cycle conditions by proposing a method to track material relaxation. Extended 30 cycle FEA simulation in *SRLIFE* demonstrates the impact of accumulated residual stress on a tube's future stress state. Accurate knowledge of a tube's current stress state allows for state of health tracking by adding the real-time damage to the already accumulated tube damage.

Current work integrates an existing model from Martinek et al. (2021) with *CoPylot*, which is an application programming interface for *SolarPILOT*. Now, users can provide custom desired flux profiles to *CoPylot* and the results feed directly

into the thermal model. Additionally, minor modifications made allow the thermal model to simulate flat receivers, solve quickly for high-speed applications, and interpolate independent flux values for adjacent tubes of neighboring panels. Best practices are established for thermal modeling regarding node count and coupling to optical and damage models. Numerous best practices are established via analysis of resolution effects on optical, thermal, and damage models. This work investigates the coupled impact of a single model's resolution on the other models it is integrated with, which is not frequently discussed in literature. The number of tubes simulated per panel is shown to have a significant impact on lifetime evaluations, and the location of tube failure in each panel is very predictable for the aiming strategy employed in Section 3.5.

The next chapter leverages the tools developed here to define ideal flux profiles, formulate damage tool-informed aiming methods, and perform system cost analyses for different receiver designs.

## 4 RELIABILITY-FOCUSED AIMING AND DESIGN FOR CENTRAL RECEIVERS

---

The detailed damage models already developed provide for lifetime estimation of a tube experiencing a given temperature difference at fluid temperatures ranging from 275-575°C. Thus, a series of maximum allowable flux limits have been established for a range of lifetimes from 10-80 years. Each flux limit is a constant lifetime contour. Inspection of any given lifetime contour shows that the significant impact of thermal stress on receiver tube lifetime leads to large changes in maximum allowable flux along a salt receiver's flowpath. Design of a solar receiver and its incident flux profile often assumes a constant flux limit, which results in under-utilization of some tubes and over-utilization of other tubes.

Analysis of critical thermal points on the Crescent Dunes receiver, performed in Section 2.3 shows an undesirable distribution of tube lifetimes. Tube lifetimes would ideally be uniformly at or slightly above 30 years. Instead, a large number of tubes have unnecessarily long lifetimes and a significant number of tubes are predicted to fail before 30 years of operation. Design power and receiver reliability could be improved if flux is increased on the underutilized tubes and decreased on the low lifetime tubes. Present work uses damage model information to more closely match local tube temperature constraints to flux conditions. The author's general class of concepts related to this idea is so-called 'informed aiming' because flux aiming is informed by a lifetime contour. This important concept is the motivation for work presented in the current chapter.

Key contributions of the present chapter are: (i) a definition of the "ideal flux profile" that maximizes receiver area utilization for a given lifetime (ii) a module of functions to generate and visualize ideal flux profiles for a receiver of desired dimensions (iii) a computationally-fast method for approximating ideal flux profiles that is well suited for design or real-time applications (iv) a novel heuristic that directly integrates damage modeling with heliostat aiming decisions and further improves performance gains (v) levelized cost of heat results across a wide range



of height to width aspect ratios for two aiming strategies.

Section 4.1 begins this discussion by defining ideally utilized receiver flux profiles based on damage map lifetime contours, presenting the underlying thermal assumptions, introducing the **Informed Aiming** module, verifying the process with a preliminary thermal analysis, and showing the minimum required area for a design power. Section 4.2 introduces the primary barrier to implementation of ideal flux profiles: image spreading. Section 4.1 introduces a method of using *SolarPILOT* built-in image placement heuristics to achieve a pseudo-ideal flux profile, and Section 4.4 presents a new heuristic capable of considering flux violations during image placement. Section 4.5 demonstrates the performance of both developed methods. Section 4.6 presents an implementation of a literature cost model and applies it to study system cost trends over a range of aspect ratios for both uniform and informed *SolarPILOT* method aiming methods.

## 4.1 Defining an Ideally Utilized Receiver Flux Profile

The allowable absorbed flux for a given receiver dictates the maximum allowable quantity of heating that can be achieved for a given receiver area. Present work defines an "ideal" flux profile as referring to a flux profile that matches the incident flux as closely as possible to the allowable incident flux to utilize a given receiver area as fully as possible. A receiver with a given height and width, therefore, has a maximum power that is limited by the maximum allowable flux at every point along the flowpath. Present work develops a method to assess the power that a fully utilized receiver with given dimensions can achieve and presents an implementation of a two flowpath flat receiver case. Key code implementations of procedural methods and visualizations are also presented.

### Translating Design Contours into Flux Profiles

Damage maps, such as that shown in Fig. 2.12, are populated by lifetime contours. The contours are shown over a range of thermal operating points defined

by the maximum total temperature difference ( $\Delta T_{\text{total}}$ ) and corresponding fluid temperature ( $T_f$ ) experienced by a tube at design point. Each of these lifetime contours can be recast as a design contour – a guideline dictating the limit of thermal conditions for the tubes along the flowpath. The damage tool method **make\_contour\_function\_from\_interpolator** is described in Section 3.2 and outputs a design contour corresponding to the desired lifetime. Output from this function is called a contour function, and it describes the relationship between  $T_f$  and  $\Delta T$  for any given case. A receiver is fully utilized when the total temperature difference is on or below a selected contour at each fluid temperature point along the flowpath. Conservative design can select a high lifetime design contour with the tradeoff of increased receiver area, and more aggressive designs may desire lower lifetimes that allow for higher flux densities.

Design contours are in terms of fluid temperature and total temperature difference for the purpose of being generalizable. However, incident flux is typically the metric for aiming and design. Therefore, the basic procedure used to translate total temperature difference into incident flux is presented here. We assume that heat transfer from the crown to the molten salt is only in the radial direction. Note that use of this assumption is limited to determination of informed aiming/flowpath strategies and not common practice in the thermal model unless explicitly stated. This assumption allows for the tube wall to be approximated as a plane wall with tube thickness as the length of conduction. Neglecting circumferential conduction is justified by observing that molten salt receiver tubes are typically very thin compared to their diameter. For example, Crescent Dunes uses tubes with outer diameter of 50.8 mm and thickness of 1.25 mm. Applying a series resistance formula at the crown leads to Eq. 4.1, where  $R_{\text{cond}}$  is the conduction resistance at the crown,  $R_{\text{conv}}$  is the convection resistance to the salt inside the tube,  $\dot{q}_{\text{ab}}''$  is the absorbed flux at the crown, and  $A$  is the heat transfer area considered at the crown.

$$\dot{q}_{\text{ab}}'' A = \frac{\Delta T}{R_{\text{cond}} + R_{\text{conv}}} \quad (4.1)$$

Expanding the resistance terms and solving for  $\Delta T$  leads to Eq. 4.2, where  $k$  is

the effective wall conductivity,  $t_h$  is the tube thickness, and  $\bar{h}$  is the heat transfer coefficient. Conductivity and heat transfer coefficient values are location dependent and determined using the local  $T_f$  and  $\Delta T$ . Conductivity is material dependent and also is dependent on the temperature at the inner crown surface. The heat transfer coefficient is determined with the same correlation that the full thermal model in Chapter 3 uses, which is the Gnielinski correlation for smooth pipes.

$$\dot{q}_{ab}'' A = \frac{\Delta T}{\left(\frac{t_h}{kA} + \frac{1}{\bar{h}A}\right)} \quad (4.2)$$

Heat transfer area,  $A$ , can be canceled from the resistance terms and the left hand side of Eq. 4.2 to result in Eq. 4.3.

$$\dot{q}_{ab}'' = \frac{\Delta T}{\left(\frac{t_h}{k} + \frac{1}{\bar{h}}\right)} \quad (4.3)$$

Incident flux is obtained from absorbed flux via the tube absorptivity by the relation  $\dot{q}_{inc}'' = \dot{q}_{ab}''/\alpha$ , with  $\alpha$  assumed to be 0.96. Ultimately, Eq. 4.3 yields a relation for  $\dot{q}_{ab}''$  that is a function of mass flow rate via  $\bar{h}$ , and fluid temperature via  $\Delta T$  whose relation to  $T_f$  can be defined by a design contour. Note that  $k$  is also fluid temperature dependent.

The total area required to meet a flowpath's design power is determined with the following procedure. First, the allowable incident flux for every  $T_f$  in the fluid temperature range (usually 290–565 °C) is determined using a selected design contour, which provides a  $\Delta T(T_f)$  function. Typical resolution for the difference between  $T_f$  points ( $\Delta T_f$ ) is 1 °C, so a typical molten nitrate salt flowpath would have at least 275 points where  $\Delta T$  is determined for the given  $T_f$ . The required heating to raise the temperature of the fluid is defined as  $\dot{Q}_\Delta$  and determined using the energy balance in Eq. 4.4, where  $\dot{m}_{path}$  is the mass flow per flowpath,  $c_p$  is the fluid temperature dependent specific heat capacity.

$$\dot{Q}_\Delta = \dot{m}_{path} c_p \Delta T_f \quad (4.4)$$

Required area ( $A_\Delta$ ) for any flowpath section is calculated with the relation

shown in Eq. 4.5. Each flowpath section's outlet  $T_f$  dictates the maximum allowable  $\Delta T$ , which is translated to the incident flux  $\dot{q}_{inc,\Delta}''$  through Eq. 4.3.

$$A_{\Delta} = \frac{\dot{Q}_{\Delta}}{\dot{q}_{inc,\Delta}'' \alpha} \quad (4.5)$$

The process described here is performed for all fluid temperatures so that each flowpath section has a required area and corresponding incident flux value. Total area required by the flowpath to reach the objective outlet temperature is the sum of all flowpath sections. Total required receiver area is calculated by multiplying the required flowpath area by the number of receiver flowpaths. Functions in a python module named **informed\_aiming** automate the process of determining ideal flux for a receiver flowpath. Next, a logical overview of this python module is provided.

## Informed Aiming Module

Key informed aiming module functions are described here and are organized by expected order of use.

An ideal flux map refers to a pseudo-continuous map of flux values for a given receiver design. The function **generate\_ideal\_fluxmap** generates the ideal flux profile in the form of a python dictionary. Required receiver design inputs include an instance of a typical tube thermal object, the desired panel lifetime, the mass flow, the panel width, the panel height, and the flowpath configuration. Tube diameter and thickness can be set as object attributes, and more details about the tube object can be found in the accompanying python code for Martinek et al. (2021). Present work assumes an outer tube diameter of 50.8 mm and thickness of 1.25 mm. The ideal flux dictionary output contains a number of sub-dictionaries and each one represents a panel in the flowpath. Each flowpath dictionary contains a list of y coordinates and corresponding incident flux values. Each y coordinate is the upstream coordinate, and all panel locations downstream of that y coordinate are assigned the corresponding incident flux value. Incident flux is updated at the next

y coordinate. The distance between y coordinates is the flowpath section length and can be determined by dividing the required square section area calculated from Eq. 4.5 by the panel width. An ideal flux dictionary can be thought of as a collection of piecewise functions, with each function output being the incident flux and the domain of each function being the y coordinates that bound the section. Generating ideal flux maps are only dependent on one resolution — the specified  $\Delta T_f$  of each flowpath section. Larger  $\Delta T_f$ s will result in "blockier" ideal flux maps, whereas small  $\Delta T_f$ s result in smaller flowpath sections and more continuously varying incident flux. There is also an alternative version of this function called **generate\_ideal\_fluxmap\_with\_offset**, which is discussed in a later section.

The ideal flux map is generated by first calculating each flowpath section's incident flux and section length, and then translating that information into a ideal flux map dictionary via the helper function **assign\_fluxes\_to\_panels**. This helper function is currently configured to create a dictionary for the '2\_ctr' flowpath configuration, where a flat receiver has two flowpaths that start at the top, side by side, in the center of the receiver. Flow through the panels is assumed to move in a serpentine fashion, with the flow alternating between with and against gravity. Although the algorithm developed here is configuration specific, it is intended to be a reference algorithm in the case that other configurations are desired. The algorithm considers the direction of flow for each panel and completes any partially finished section on the next previous panel. Sections are assigned to each panel until all sections have been assigned. If less than a full panel is required at the end of the flowpath, the algorithm assigns zero incident flux to the remaining portion of the panel.

Ideal flux map dictionaries are created with the intent of being used to inform aiming strategies and flowpath designs. Flux map visualization is important tool for both of these tasks. The highest fidelity visualization of an ideal flux map is one that does not distort each flowpath section's dimensions by interpolating at regular intervals. A function named **plot\_ideal\_fluxmap** is written that translates an ideal flux map dictionary into a series of rectangular boxes with color corresponding to its incident flux intensity and coordinates determined by the panel and y coordinates.

Figure 4.1 shows an intermediate step in generating an ideal flux map visualization to illustrate the process. A receiver with height of 10 m, total width of 14 m, and panel width of 1 m is considered. Molten salt flows through two paths at 200 kg/s, and both paths begin at the top center. Design power is approximately 150 MW<sub>th</sub> for this case. These inputs are provided to **generate\_ideal\_fluxmap**, which calculates the lengths and fluxes for a series of flow path sections required to raise the fluid up to 565°C. The helper function **assign\_fluxes\_to\_panels** arranges these sections on a receiver according to the flowpath configuration. Finally, **plot\_ideal\_fluxmap** creates a figure for visualization. The small rectangles seen in Fig. 4.1 each represent a flowpath section over which the fluid heats up by 1°C. Section lengths, which are the heights of each rectangle, are shown to increase as fluid progresses through the flowpath. This is because higher  $T_f$  sections are constrained to lower incident fluxes, in accordance to the 30 year design contour used. Lower flux intensity sections require area to achieve the same increase in fluid temperature. Figure 4.2 shows the final visualization that is obtained by coloring each flowpath section by its allowable incident flux intensity. Note that the end of each flowpath has a zero flux region. Presence of a zero flux region means that some of the final panels' length is not required in the ideal case. In this case, mass flow or lifetime could be increased to more fully utilize the receiver.

Lastly, the **informed\_aiming** model contains functions to construct flux grids from flux maps. Flux maps are shown to be collections of piecewise incident flux functions with irregularly spaced x and y domains over the receiver. In contrast, the term 'flux grids' is used in this work to define an array of regularly spaced grid interpolations of a flux map. Flux grids are compatible with other aspects of this work, such as design, aiming, and analysis. For example, thermal model input is sometimes in terms of a flux grid, which is then imposed as boundary conditions on each simulated receiver tube. The function **build\_ideal\_fluxgrid** interpolates a flux map at discrete points specified by the user. Each flux grid point's incident flux level is determined by determining the panel based on x position and then taking the incident flux level corresponding to the section its y position is in. This process is illustrated in Fig. 4.1, where the sample grid point (star) is located in a




Figure 4.1: Intermediate step in visualization process showing each flowpath section on a series of 14 panels. A sample location (star) has flux corresponding to the flux intensity defined by the nearest upstream section corner (dot).

section defined by the panel number and the nearest upstream y section coordinate (circle). For the shown case, where flow is in the direction of gravity, the top corner y coordinate will determine the incident flux section that the sample lies in. In the case the flow is opposing gravity, the bottom corner y coordinate determines the incident flux section that the sample lies in. Function output's resolution is partially determined by the user — a function parameter is the y resolution. The resolution in the x direction is equal to the number of receiver panels. An x resolution lower than this will result in poor performance because flux levels of adjacent panels will be projected onto neighboring panels. A helper function **plot\_ideal\_fluxgrid** is used to plot the resulting fluxgrid, which is shown in Fig. 4.3. The dots represent each location at which the flux map is interpolated.

All **informed\_aiming** functions that pertain to designing and visualizing ideal flux maps as well as flux grids are presented in this section. Other modules function




Figure 4.2: Final visualization of generating an ideal flux map for example conditions described. The zero flux regions at the end of each flowpath represent unnecessary receiver area.




Figure 4.3: Flux grid profile for flux map shown in Fig. 4.2. A grid of 14 x 14 points is used.



exist that are used to deal with real world situations including edge offset algorithms to improve optical efficiency, and methods to fit ideal flux maps to fixed receiver dimensions. These additional functions are presented in upcoming sections.

## Process Verification

It is important to verify that ideal aiming schemes developed with the procedures presented here produce lifetimes that are consistent with our objective lifetime. A full test case is run to evaluate the process. First, the **informed\_aiming** module functions generate an ideal flux map and then build a flux grid. Then, the flux grid is input to the thermal model. Output from the thermal model is an array of  $T_f$ s and  $\Delta T$ s at axial positions along the flowpaths. Finally, these thermal operating points are input to the damage model and lifetimes are predicted at each point. The test receiver has a height of 18 m, a total width of 15 m, a panel width of 1 m, and a nominal design power of 180 MW<sub>th</sub>. Our objective lifetime is 30 years for all panels. Receiver tubes are made of Haynes 230 alloy. The thermal model simulates three tubes in every panel.

Figure 4.4 shows the ideal flux map for this case when the radial basis function interpolator is used to make the design contour. It can be seen that at least 4 out of the total 18 panels are unnecessary if flux can be applied ideally to this receiver. Figure 4.5a clearly shows that all thermal operating points are below and to the left of the 30 year contour. This is expected, because the ideal flux map is built using similar modeling concepts to what the full thermal model uses. It can be observed that thermal operating points slowly diverge from the contour line as the fluid temperature increases. This divergence is a result of the ideal flux map being built without considering convection losses along the flowpath. Damage model output, shown in Fig. 4.5b, shows lifetimes are uniformly near 30 years of lifetime with two exception regions. The design contour does not actually follow the damage model data's 30 year lifetime contour at low fluid temperatures because the 30 year lifetime contour intercepts the stress reset region around approximately  $\Delta T = 270^\circ\text{C}$ . A hard cutoff (vertical line) is used here to avoid operation in the stress reset region. The




Figure 4.4: Ideal flux map for a receiver with 15 m height and 1 m panel widths. The actual receiver dimensions exceed the required number of panels, so the last two panels on each side are left blank and will be treated as zero flux zones.

vertical line causes thermal operating points in the low fluid temperature range to shift to higher lifetimes. Additionally, the zero flux region in the last 2 panels of each flowpath also yields very high lifetime regions on the receiver.

Results from this process evaluation demonstrate that the **aiming\_informer** algorithms for producing ideal flux maps and flux grids can successfully create a flux profile capable of maintaining the desired lifetime of 30 years, according to our thermal and damage models presented in the previous chapters. Also, the presented process of generating an ideal flux map for a given receiver is shown to have sufficient performance without requiring detailed or iterative thermal modeling techniques.

## Minimum Required Area Study

The minimum required area for a receiver is the sum of all flow path section areas in a flux map and dependent on the design power. This metric is useful for assessing receiver utilization and several key parameters that impact minimum required area are: desired lifetime, material, panel width, and mass flow. Desired lifetime is




Figure 4.5: Resultant thermal operating points (4.5a) output from full model simulation with a flux grid input corresponding to the flux map shown in Fig. 4.4. Calculated lifetimes for each axial sample point are shown in (4.5b). Note that the color bar maximum is 100 years, but some lifetimes exceed that maximum.

inversely related to minimum required area, and is typically targeted to be 30 years. It may be desired to design for a higher lifetime in order to deal with uncertainties in the expected thermal conditions, creep or fatigue properties, or aiming precision. Material properties, namely creep and fatigue limits, strongly affect allowable flux and therefore impact the minimum required area for a given design power. Alloys with improved creep tolerance such as 740H and A282 are expected to have lower required areas compared to Haynes 230 and Alloy 617 for the same design power. Panel width determines the number of tubes placed in parallel for the molten salt to flow through. This fact stems from the assumption that tubes are attached to the panel as densely as possible and therefore the number of tubes per panel is approximately equal to the panel width divided by the outer tube diameter. A larger number of parallel tubes along the flow path results in lower mass flows and velocities per individual tube for a given design power and fixed mass flow rate. Consequently, the predicted heat transfer coefficient is also inversely related to the panel width or tubes per panel. Finally, design power and mass flow are directly proportional to one another because the target outlet temperature and inlet

temperature remains constant. A higher design power is associated with higher mass flow rates and therefore improved heat transfer coefficients, which lead to decreased receiver area. The additional benefit of higher heat transfer coefficients is that they yield lower total temperature differences across the tube compared to lower coefficients. Thus, higher heat transfer coefficients allow for higher allowable flux values for the same desired lifetime.

Authors use the function **generate\_ideal\_fluxmap** to calculate minimum required areas for a range of design powers between 100–400 MW<sub>th</sub>. The resulting required area-design power plot is presented in Figure 4.6 for a receiver with 1 m width panels, 50.8 mm outer diameter tubes with 1.25 mm wall thickness and Haynes 230 material. This curve assumes default flowpath configuration of 2 paths starting in the receiver's center and uses an objective lifetime of 30 years. Mass flow is related to each design power by assuming a constant value for specific heat and a total flowpath temperature rise of 275°C between inlet and outlet. Note that design power is defined as the power to the heat transfer fluid, and is distinct from absorbed and incident powers, which are related to design power through their efficiencies.

Results show a fairly linear relationship between required area and design power, which is expected because greater area equates to larger net heat transfer into the fluid. Figure 4.6 provides a useful relationship between required area and design area for the case described, but it should be interpreted in context of two important limits that the 30 year lifetime design contour does not account for. Firstly, increasing mass flow requires increased pumping power. Therefore, an optimum may exist that maximizes the net power, which is design power minus the pumping power. Secondly, increasing salt velocity can also contribute to accelerated rates of material degradation. For reference, work in Gentile et al. (2024) enforces a salt velocity limit of 5 m/s. It can be observed that minimum required areas range from approximately 150–285 m<sup>2</sup> for the 100–400 MW<sub>th</sub> design power range in the flat receiver case considered here.

Future work should leverage the functions presented in this section to explore the impact of design parameters on the ideal power limits of a receiver. The slope




Figure 4.6: Minimum area for a range of design powers, assuming a 30 year lifetime, Haynes 230 tube material, and approximately 20 parallel tubes per panel.

of Fig. 4.6 is expected to change with respect to material, panel width, inner tube diameter, design lifetime, and tube length. Finally, a more complete picture of the ideal design space requires consideration of parasitic pumping losses and salt velocity limits.

## Summary of Ideal Flux Development

Work in Section 4.1 accomplishes several important objectives related to reducing required areas of central receivers. Ideal flux profiles are defined as maximum allowable flux profiles for a chosen design lifetime. These profiles are distinct from global flux limits that are commonly seen in present operation and design schemes. Instead, the maximum allowable flux varies according to fluid temperature. This

section also documents the theory of translating damage map lifetime contours into design contours, which provide a convenient relation between fluid temperature and allowable flux. We program and present several important procedures for creating flux maps and synthesizing them into flux grids to be used by other modeling suite tools. Additional functions are written and presented for visualizing the flux maps because visualization is believed to be crucial for process development. Present work develops code for a two flowpath, center-starting receiver, but future work can build on the presented workflows to develop processes for other desired flowpath configurations. Testing verifies the viability of the overall workflow by evaluating its performance in a specific 180 MW<sub>th</sub> scenario. The results yield a field of lifetimes at 30 years or higher for all evaluated points on the receiver. Minimum required area across a range of 100–400 MW<sub>th</sub> design powers is determined to be in the range of 150–285 m<sup>2</sup> for the given receiver parameters. Overall, the present work establishes and visualizes a target flux profile for aiming strategies to strive towards. The target flux profile shows excellent potential to reduce required receiver area. Subsequent sections will assess the viability of ideal flux profiles in real scenarios and will present methods for attaining approximate versions of ideal flux profiles.

## 4.2 Challenges of Achieving Ideal Flux Profiles

Solar images are ideally reflected off of heliostats and retain the intensity distribution of the original image of the sun. In reality, however, the heliostat-reflected images spread into a lower magnitude, larger area due to atmospheric attenuation and/or heliostat imperfections. The effect of image spreading is demonstrated in Fig. 4.7, which shows two heliostat images on a generic flat receiver in *SolarPILOT*. Both heliostats are positioned along the midline of the heliostat field and are aiming at the receiver's center. Figure 4.7a shows the flux distribution from a heliostat that is significantly closer to the receiver compared to the heliostat used to make the flux distribution in Fig. 4.7b. It can be seen that although the heliostats dimensions are the same, the flux intensity is significantly lower for the more distant heliostat and covers a larger area. The discrete flux levels on each panel that are required to

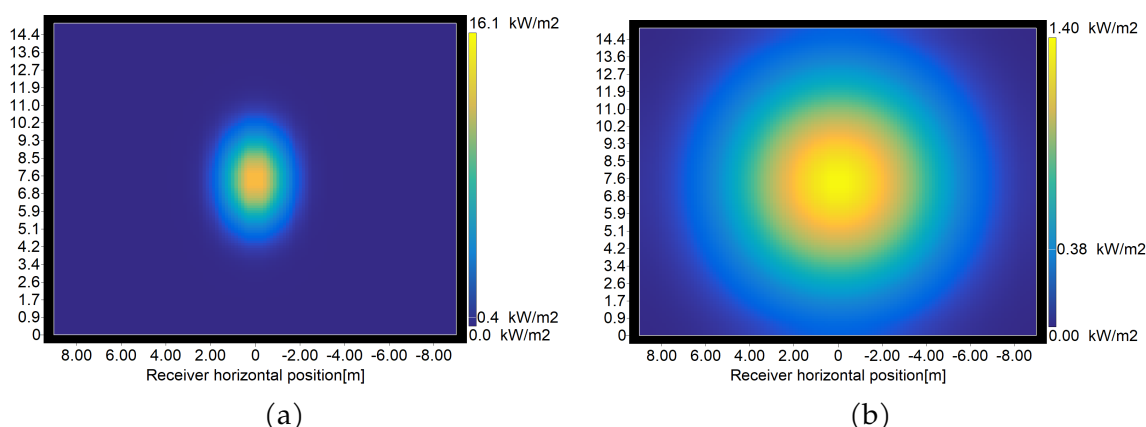


Figure 4.7: Flux profile output from *SolarPILOT* of two heliostats along the solar field midline aimed at the center of the receiver. The heliostats are 128 m (4.7a) and 858 m (4.7b) from the receiver.

achieve informed aiming are therefore more easily accomplished in smaller CSP systems, where the heliostats are all located in close proximity to the receiver. The rated design power of a receiver ultimately determines the average heliostat distance from the receiver — as design power increases, additional heliostats must be added in the back and sides of the field that have larger flux distributions and are therefore less desirable. *SolarPILOT* is currently addressing this challenge by using an aiming method called ‘Image Size Priority’. As explained in the *SolarPILOT* documentation (Wagner (2018)), this aiming method places heliostats on the receiver to achieve the desired flux profile distribution, and heliostats with the largest image sizes are placed first. To the author’s knowledge, *SolarPILOT* constantly evaluates the flux distribution during heliostat placement but does not enforce any exact flux magnitude limits.

### 4.3 Informed *SolarPILOT* Aiming Method

The first developed aiming method modifies a heuristic that already exists in *SolarPILOT*— the image size priority method. Present work introduces the image size

priority heuristic in Section 2.2. For flat receiver cases, *SolarPILOT* users may input an objective flux profile that the image size priority heuristic uses as an objective profile. A demonstration of this capability is shown in Section 3.5, where a custom flux profile is input to *SolarPILOT* with high flux on the left and low flux on the right (Fig. 3.15). We combine this capability with our generated ideal flux profiles introduced in Section 4.1. The resulting method is named the 'informed *SolarPILOT* method' because our objective flux profiles are serving as input to the existing *SolarPILOT* aiming method to achieve improved performance. Present work finds that the informed *SolarPILOT* method's performance from a panel reliability perspective strongly depends on the quality of the objective flux profile input. This section presents procedures that modify the ideal flux profile to enable higher fidelity between the objective profile and the *SolarPILOT* result.

**Generating Objective Flux Profiles for *SolarPILOT*** Ideal flux grids generated with the **informed\_aiming** module can be supplied directly into *SolarPILOT*. Present work uses the API *CoPylot* to read saved .csv flux grid profiles as a matrix and then set the custom user profile using that matrix. Users should note that reading the saved file as a matrix sometimes "flips" the grid, so code developed here typically saves a flux grid in an inverted state to address this issue. An example of input user grid is shown in Fig. 4.8a for a 220 MW<sub>th</sub> flat receiver with dimensions of 18 m width and 15 m height. Heliostats are assumed to be 12.2 m x 12.2 m with a total reflectivity error of 3.07 mrad. The ideal flux profile does not require the full area of the receiver to reach target salt temperature, so we say that this receiver is not being fully utilized. Figure 4.8b shows the *SolarPILOT* output when Fig. 4.8a is used as an objective flux profile. An extremely high flux density is achieved at the bottom center of the receiver.

The lifetime profile shown in Fig. 4.9 shows that lifetimes are not limited in the center of the receiver. However, the intermediate regions of the flowpath do experience undesirably low lifetimes, with the minimum lifetime being 17.82 years and the receiver requiring 6 total panel replacements over its lifetime. Factors that cause the discrepancy between the objective profile and the *SolarPILOT* are due






Figure 4.8: An objective flux profile input to *SolarPILOT* ( 4.8a), and the result ( 4.8b) for a  $220 \text{ MW}_{\text{th}}$ ,  $270 \text{ m}^2$  case.

to image spreading and *SolarPILOT*'s enforcement of an offset region around the perimeter of the receiver. Ultimately, a number of factors can cause divergence of the flowpath's thermal points from the contour that the objective profile assumes. These factors can be mitigated by offsetting the design contour — that is, strategic underutilization of the receiver to include error margin for aiming complications.

**Accounting for *SolarPILOT*'s Enforced Offset** *SolarPILOT* reduces image spillage by enforcing an edge offset. Each image's aimpoint options are restricted to locations that are at least a set number of image deviations away from the receiver edges. In the cases shown for this section, an edge offset in the x and y directions of 2 is assumed. Increasing the edge offset value improves the field's intercept efficiency on the receiver, but the usable receiver surface is also decreased. After all, the limit of maximized optical efficiency for a given receiver is aiming all the images on the center and results in unsustainable flux concentration there.

Custom flux output from *SolarPILOT* will more closely resemble the objective output if the edge offset is accounted for when building the objective profile. The `informed_aiming` module's `generate_ideal_fluxmap_with_offset` function




Figure 4.9: Lifetime profile for 220  $M_{th}$  case shown, where receiver flux is as shown in Fig. 4.8b.

is designed to account for edge offset by anticipating the dropoff in flux levels for panel sections within an offset region. This function all the input required by **generate\_ideal\_fluxmap\_with\_offset** (see Section 4.1) and also an assumed offset distance from the top and bottom receiver edges. Note that the input is specified in meters, not standard image deviations. The current work assumes an offset region that is approximately 2 meters in length. Our **generate\_ideal\_fluxmap\_with\_offset** function determines the maximum allowable flux for every subsection, but assumes that flux linearly must decrease from the previous, non-offset region's flux level down to zero when ending a panel. The opposite assumption is made for panel starts — the function iteratively solves for the maximum possible temperature difference in the offset region, assuming that flux must start at zero on the edge and then linearly ramp up to full value by the end of the offset length.




Figure 4.10: An objective flux profile input to *SolarPILOT* ( 4.10a), and the resulting lifetime profile ( 4.10b) for a 220 MW<sub>th</sub>, 270 m<sup>2</sup> case.

Figure 4.10a shows the offset grid objective profile resulting from the process described for the same design case as in Fig. 4.8a. It can be seen that the regions from 13–15 m and from 0–2 m taper linearly from full scale flux values down to zero at the panel edges. This mimics the typical effects of *SolarPILOT*'s enforced edge offset. The resulting lifetime profile is shown in Fig. 4.10b. Minimum lifetime increases to 21.25 years, but the total number of required panel replacements remains at 6. Nevertheless, 4 years of improvement in minimum lifetime is significant. Additionally, this offset-considering method becomes even more impactful for shorter panels, where the offset regions occupy increased fractions of each panel.

**Maximize Lifetime for a Given Receiver** Objective profiles built on ideal or ideal with offset methods can be observed to have zero flux flowpath area at the end of both flowpaths. For a set receiver size, we can leverage this unused flux area to reduce flux levels at points further upstream in the flowpath. An **informed\_aiming** function called **fully\_utilize\_receiver\_maximize\_lifetime** is developed to accomplish this task. This function builds a fluxmap using previously introduced methods and checks if all of the receiver area is being utilized.




Figure 4.11: An objective flux profile input to *SolarPILOT* ( 4.11a) using the maximize receiver lifetime and offset functions, and the result ( 4.11b) for a  $220 \text{ MW}_{\text{th}}$ ,  $270 \text{ m}^2$  case.

In the case that unused area exists, the function regenerates the fluxmap with a higher lifetime contour. This process is iteratively performed until approximately all of the receiver area is used. The result is an objective profile with lower overall flux magnitudes, but the entire available area is being used for heat transfer. Figure 4.11a shows the objective flux profile for our  $220 \text{ MW}_{\text{th}}$  example case when the **fully\_utilize\_receiver\_maximize\_lifetime** function is used. *SolarPILOT* flux output shown in Fig. 4.11b has a noticeably different shape compared to Fig. 4.8b, and Fig. 4.12 shows significant improvement in panel lifetimes. The minimum lifetime is 28.85 years and the total required panel replacements are only 3. Overall location of critical panels has not changed for any of these cases, but the intensity at those locations decreases as our methods develop.

**Summary of Informed *SolarPILOT* Method** The informed *SolarPILOT* method is presented in this section as a means to accomplish ideal aiming in actual solar field scenarios. Method speed is fast — on the scale of seconds — and leverages an existing placement heuristic in *SolarPILOT* via Python code that uses *CoPylot*. Best practices include accounting for offset and relaxing the design contour whenever




Figure 4.12: Lifetime profile for 220  $M_{th}$  case shown, where receiver flux is as shown in Fig. 4.11b.

possible. Implementing these practices improves the minimum panel lifetime from 17.8 to 28.9 years. Detailed results are not shown, but the same thermal and lifetime analysis using a uniform aiming strategy yielded a significantly lower minimum lifetime of 12 years and 10 total panel replacements over the same 30 year lifetime.

Achieving near-30 year lifetime on a 220  $MW_{th}$ , 270  $m^2$  receiver is considered by the present work to be extremely difficult. These design parameters result in an overall average incident flux of approximately 943  $kW/m^2$ . For reference, the default cylindrical receiver in *SolarPILOT* has an average incident flux of approximately 648  $kW/m^2$ .

Section 4.5 compares this *SolarPILOT* aiming method to uniform aiming for a range of design powers. Results show that the informed *SolarPILOT* method is a significant improvement over the uniform strategy and well suited for iterative

design applications or real-time cases. However, the informed *SolarPILOT* method does not integrate damage analysis into its aiming method and therefore is not considered as robust as the method presented in the next section.

## 4.4 Development of an Informed Aiming Heuristic

Informed aiming, when implemented with *SolarPILOT*'s image size priority algorithm, demonstrates creep-fatigue reliability advantages over uniform aiming. However, the image size priority method is only focused on achieving the shape of the requested flux profile. This aiming algorithm is initially informed by an ideal flux profile input, but the damage model is not consulted during the procedure. Problems arise in real systems because images borders are not discrete — inevitable image spillage onto adjacent panels causes undesirable and unexpected trouble spots, where incident flux exceeds the local allowable value. Since informed aiming in *SolarPILOT* does not actually consult a damage model while placing images, it places images irrespective of the effect their borders may have on neighboring panels. Work in the current section discusses the development of an informed aiming heuristic that places images according to damage tool information that updates as images are placed. The heuristic is theoretically more likely to produce a less damaging aiming scheme compared to a simple image size priority method because image placement decisions anticipate the effect that the next-placed image's borders may have on neighboring panels.

It should be noted that this so-called heuristic is not an actual optimization routine such as in Zolan et al. (2021). In its current form, the heuristic progresses in one direction through the list of images and only performs a single pass. A series of single image placement optimizations are performed in succession, and the result is not independent of image order. Nevertheless, this heuristic is shown to be robust in that it provides high quality aiming schemes capable of meeting the objective power. At the present, HALOS optimizes the receiver power subject to upper limits on incident flux. Meanwhile, the heuristic in the present work simply mitigates incident flux violations while being required to meet a given design power. The two

aiming methods differ, therefore, in their ability to produce solutions for receivers with high degrees of utilization. A HALOS optimization becomes over constrained by the rigid, spatially varying flux limits. In contrast, the heuristic guarantees a solution because it does not rigidly enforce flux limits.

The primary contribution of this heuristic lies in producing aiming schemes that are expected to outperform both uniform aiming and simple image size priority methods. This section presents the heuristic workflow, shows intermediate steps of the process, and discusses performance outcomes compared to other methods.

**Heuristic Workflow** There are three inputs to the heuristic: receiver parameters, a corresponding ideal flux profile, and a set of heliostat images. Important receiver parameters include the design power, the number and width of panels, tube material, tube geometry, the number and orientation of flowpaths, the target outlet temperatures, and the accompanying heliostat field layout. An ideal flux map is generated that is specific to these receiver parameters, and a flux profile in the form of an array of gridpoints is provided as a guiding input to the heuristic. The gridpoint array shape must match that of heuristic's grid of aimpoint options. Finally, there are a set of individual heliostat flux images as input for the heuristic to place on the receiver. Output from the heuristic is a list of aimpoints for each heliostat and the predicted flux profile.

An overview of the heuristic workflow is shown in Fig. 4.13. Starting at the top of Fig. 4.13, an image is selected from the input ordered set of individual heliostat images. This set of images is composed of receiver flux profiles based on individual optical simulation in *SolarPILOT*. Each profile, then, represents the flux contribution of a single image. The resultant set size corresponds to the number of heliostats in the field layout. Authors use an existing method in the HALOS software called `get_single_helio_flux_dict` to build the set. Single heliostat optics are assumed to stay relatively constant for all aimpoints on the receiver. Consequently, we assume that the originally simulated flux array can be translated to different coordinates on the receiver without the need to re-simulate true optical performance at every aimpoint. Pre-simulating individual heliostat flux profiles reduces heuristic run-

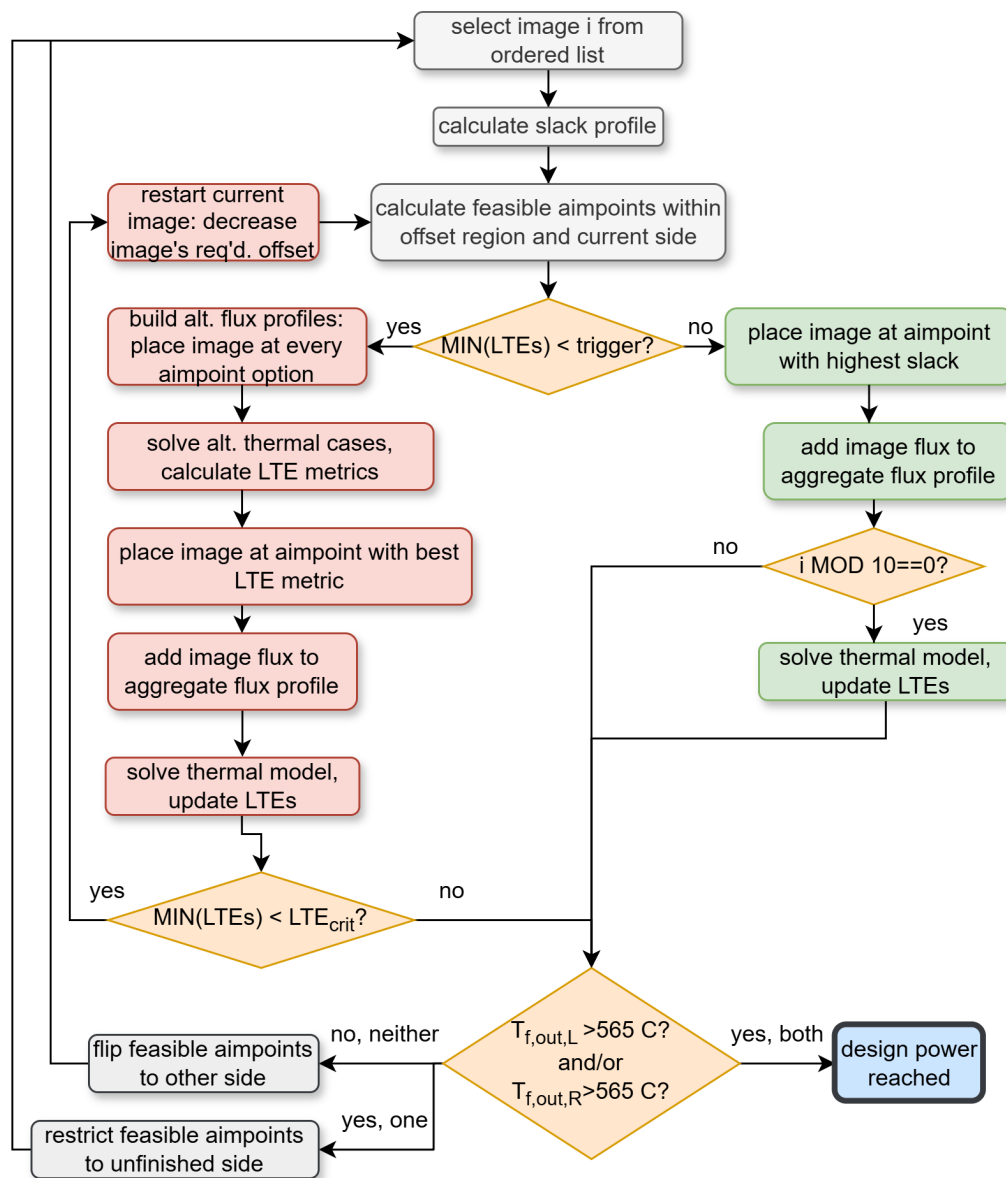


Figure 4.13: Process diagram for informed aiming heuristic, which seeks to consecutively aim heliostats in the least damaging locations until the target outlet temperature is reached in all receiver flowpaths.



time. Note, however, that this method introduces some inaccuracy for images that are larger than the receiver when originally placed at the center. Any portion of an image's flux profile that is not on the original profile is not captured and would therefore not be considered for any other image placements. The set is sorted prior to being used in the heuristic. In its current state of development, the set is sorted in descending order of image size. Thus, the heuristic places large images before small images. Large images have a more limited subset of viable aimpoints and generally need to be aimed closer to the receiver center when compared to small images.

The next step in the heuristic process is to calculate the current slack profile. Present work defines a slack profile as the difference between the ideal flux profile and the existing flux profile. At the beginning of the heuristic process, the slack profile is simply a negatively signed ideal flux profile because no heliostat fluxes are placed on the receiver yet. If the design power were met with ideally sized and placed images, the slack profile would be zero at all receiver locations. The slack profile is referenced by the heuristic in its initial placements before flux violations occur, and this is a time saving technique that imitates the simple image size priority method used by *SolarPILOT* to approximate user defined flux profiles Wagner and Wendelin (2018). Section 4.3 discusses the simple image size priority method in greater detail. Images are not placed according to the slack profile if flux violations are detected on the receiver. The two subprocesses of image placement will be further discussed later in this subsection.

After calculating the slack profile, the feasible aimpoints are determined based on the image's size. Selection of feasible aimpoints relies on precomputed estimates of the image's height and width. The border of the image is assumed to be where flux is less than 1% of the maximum value. Image height is defined as the maximum vertical distance between image borders, and image width is the maximum horizontal distance between borders. An example of a single heliostat simulation is shown in Fig. 4.14. The maximum flux is approximately  $\text{kW}/\text{m}^2$  and so the image border is where the flux is less than  $0.16 \text{ kW}/\text{m}^2$ . For this specific heliostat, the determined height is approximately 8.33 m and the width is 4.66 m. Feasible

aimpoints for a given image are a subset of all aimpoint options on the receiver. An edge offset is enforced for each image to reduce image spillage. A larger edge offset limits the an image's aimpoint options and increases the image intercept efficiency. *SolarPILOT* enforces edge offset as a number of standard deviations of the flux distribution. This heuristic enforces offset in terms of image standard deviations to be consistent with *SolarPILOT*. The image height and width are assumed to correspond to  $6\sigma_x$  or  $6\sigma_y$ , where  $\sigma$  is the standard deviation of the image in a given direction. This assumption is based on the vertical and horizontal flux distributions being approximately gaussian. Present work initially enforces an edge offset of  $1\sigma$ . Consequently, feasible aimpoints are restricted to those with y coordinates more than  $1/6$  of the image height and x coordinates more than  $1/6$  of the image width away from the receiver edges. Feasible aimpoints are further restricted if multiple flowpaths exist in the receiver. In that case, the subspace's horizontal bounds also depend on which flowpath is allowed to accept the image. For example, consider the two flowpath case described in this work. If the left flowpath is selected, then the horizontal bounds span from the left offset edge to the receiver's center. The outlet temperature decision diamond decides which flowpath to select. Finally, the initial edge offset requirement is sometimes relaxed if the aimpoint options are too restrictive, as also discussed later in this section.

The first yellow diamond in Fig. 4.13 represents the point in each image placement's process where the heuristic assesses what aimpoint selection method should be used for the current image. This decision is made by consulting the damage tool with the most current thermal operating points of the receiver. Minimum lifetimes for every panel are predicted by the damage model via methods discussed in Chapter 3. The minimum panel lifetimes are checked against a lifetime trigger ( $LTE_{tr}$ ) parameter. Selection of the  $LTE_{tr}$  value has a significant effect on heuristic performance and runtime. If panel lifetimes are all higher than the trigger value, the heuristic elects to place the image using the green method, which is significantly faster than the red method. However, the heuristic uses the red process if any panel lifetimes are below the trigger value. In general, the green subprocess is used to select aimpoints for the majority of images, and the red process is used in the later




Figure 4.14: An example image from a default template heliostat in *SolarPILOT*, with mirror height and width of 12.2 m. Heliostat is from a generated layout capable of 140 MW<sub>th</sub>.

stages when each placement has more significant effects on lifetime.

The green aimpoint selection subprocess, shown in Fig. 4.13, applies a strategy similar to that of image size priority in *SolarPILOT*. Aimpoints are selected with the objective of matching the overall shape of the ideal flux profile as closely as possible. Each image aimpoint is selected to be the location with the highest slack value and that is located within the determined subspace. That image's flux profile is translated to the corresponding aimpoint location, and then added to the overall flux profile. Lifetimes are updated every 10 image placements by: i) solving the lightweight thermal model with the overall flux profile as input ii) extracting temperature differences and fluid temperature results at each analysis point iii) calculating the corresponding lifetimes. The overall lifetime array is further simplified into a set of minimum panel lifetimes, which then are used to decide what aimpoint

selection subprocess to use for the next image placement. Selecting the number of image placements that occur before the lifetimes predictions are updated is a balance of speed and performance. Updating the lifetimes less frequently saves time but the process is less tightly controlled.

The flowpath outlet temperature diamond (yellow, bottom in Fig. 4.13) determines what flowpath side to allow image placement on. Logic for a two flowpath case is shown in the diagram. If the left flowpath's objective temperature is reached, image placement should shift focus to solely heating the right flowpath. The opposite logic is true if the right flowpath's temperature is reached. In the case that both outlet temperatures are reached, no more flux is necessary and the heuristic task is finished. When neither flowpaths are finished, the selected flowpath switches after every placement. Constant switching is necessary because many images are large enough to impact both flowpaths' thermal operating points. The coupled nature of the two flowpaths necessitates building each flowpath's flux profile in parallel. It is undesirable to add image flux on one side if doing so also increases flux levels on a finished flowpath side; All the finished flowpath's thermal operating points shift, causing damage to also increase as a consequence. Note that future versions of this heuristic should investigate utilizing field symmetry and simultaneously placing pairs of images on the receiver. This improvement could save time and ensures that each flowpath's flux profile is being developed simultaneously.

Aggregate flux profile magnitudes as well as the flowpath temperatures increase as more images are placed. Both of these factors combine to decrease predicted lifetimes of each panel. The heuristic begins using the red subprocess to select aimpoints after the first panel lifetime drops below the trigger lifetime value. Generally speaking, every image placement degrades lifetime predictions. We design the red subprocess to select the aimpoint that degrades panel lifetimes the least. Making this evaluation requires analysis of every aimpoint option within the feasible space. First, a potential aggregate flux profile must be constructed for every aimpoint option being considered. Each option's potential flux profile is calculated by translating the individual image to the potential aimpoint and then superimposing it on the existing aggregate profile. Second, thermal model cases are solved with

the potential aggregate flux profiles as inputs. These cases can be run in parallel to save time if large number of aimpoints are being considered. The resulting thermal model solutions are fed directly into the damage tool, which then outputs predicted lifetime profiles for every aimpoint option. The lifetime profiles are further processed to find the points on each panel that most limit its lifetime. Panel lifetimes are defined by those minimum points and other lifetime points are neglected. Third, the aimpoint with the best lifetime metric is selected for image placement. A lifetime metric (LM) for this heuristic is defined to be a function of the panel lifetimes ( $LTE_{\min,p}$ ), the trigger lifetime ( $LTE_{tr}$ ), and the number of panels ( $N_p$ ) as shown in Eq. 4.6.

$$LM = \frac{\sum \min(LTE_{tr}, LTE_{\min,p})}{N_p} \quad (4.6)$$

This lifetime metric is a single average value of the panel set's minimum lifetimes. Additionally, however, the numerator bounds any panel lifetimes that are unnecessarily high by limiting them to be no more than the trigger value. Bounding the panel lifetimes ensures that panels with high lifetimes do not obscure any panels with critically low lifetimes. Also, a large degradation of a large lifetime is not as significant as degradations of lifetimes that are similar to the magnitude of the design life. For the example of a 30 year design, an image placement that causes a panel lifetime decrease from 900 to 700 years does not impact the number of panel replacements a 30 year plant would require. However, a smaller decrease of panel lifetime from 40 to 20 years requires a panel replacement during the plant's lifetime. The lifetime metric therefore is designed to only react to relevant lifetime degradation rather than lifetime changes that will not impact the plant life. Nevertheless, it is important to consider all panels in the lifetime metric so that the heuristic designs a flux profile with the lowest number of required replacements. In other words, this definition of a lifetime metric defines reliability as the total number of required replacements rather than the time elapsed before the first panel's failure.

After the best aimpoint is selected, the image flux is translated from the center to that aimpoint and then added to the aggregate flux profile. An instance of the

thermal model with the proposed flux profile is solved and the lifetime profile checked. A final decision diamond is at the end of the red subprocess in cases of extreme receiver utilization. There exists a minimum acceptable lifetime threshold, and the edge offset must be relaxed if the least degrading aimpoint still causes lifetimes to fall below the lifetime threshold. Decreasing the edge offset allows for additional feasible aimpoints to be considered that, although less desirable in terms of spillage, allow the heuristic to maintain panel lifetimes above the threshold. In this case, the current image placement restarts and its contribution to the aggregate flux profile is removed. The subprocess is rerun to consider the additional points. Edge offset continues relaxing until the image can be placed and is reset to the default value for the next image placement.

This subsection explains the basic logic of the heuristic. It is expected that consulting the damage tool while placing images leads to a reduction in flux violations. The next section compares the two aiming strategies discussed in this chapter to uniform aiming.

## 4.5 Comparison of Aiming Strategies

Three aiming strategies are compared on their ability to preserve panel lifetime over a range of design powers: Uniform aiming represents the baseline strategy and will result in approximately even flux distribution everywhere on the receiver excluding the offset region. The other two aiming strategies are those developed in Sections 4.3 and 4.4: the informed *SolarPILOT* aiming and the informed aiming heuristic.

The objective is that each case's design power closely matches for each nominal test point to ensure similar mass flow rate, heat transfer coefficient between all three strategies. Several actions are performed to enforce comparable design powers. Layouts are specific to each design power and generated to be large enough that the lowest efficiency aiming method produces the design power  $\pm 1 \text{ MW}_{\text{th}}$ . Heliostats are removed from the layout in cases that yield power to fluids higher than the design power. Finally, informed *SolarPILOT* method uses 2 standard deviation

offset from the edge. Uniform aiming uses 2.25 standard deviation offset for the 140 MW<sub>th</sub> case and 2.15 for all others.

The informed *SolarPILOT* aiming settings used in this study are preliminary and use an intermediate set of methods compared to what is presented in Section 4.3. An ideal flux profile for 80 years lifetime is generated as an flux profile objective input for each design power and offset region is not considered.

Heuristic settings are as follows. Ideal flux profiles that serve as heuristic reference are identical to those used as input to the informed *SolarPILOT* aiming method. The lifetime trigger parameter is set to 240 years, and the threshold lifetime is 15 years. A grid of 54 x 54 points represents all of the aimpoint options on the receiver.

These settings listed here are common to all simulations. Receiver height is 15 m and width is 18 m. The material is Haynes 230. Heliostats are 12.2 m by 12.2 m and assumed to have a total error of 3.07 mrad. Flux resolution is 54 horizontal and vertical points over the receiver.

The optical simulation method for uniform and *SolarPILOT* simulations is via *CoPylot* as described previously. Heuristic simulation differs slightly. We simulate each heliostat individually on the center of the receiver and then translate those images as desired during the heuristic procedure. Negligible differences between these two optical simulation methods are assumed, as the image aimpoint change's effects on optics is assumed to be minor over small distances. Thermal simulation differs for the heuristic method in its assumption of 1D conduction. Note that previous work in the present paper has assessed maximum error from this assumption to be 1.3 K.

Each aiming method positions the heliostats and the resulting thermal profiles are calculated. Post-processing is performed using the same damage tool version to analyze case's thermal operating points. This fact implies that the lifetime comparison basis is consistent for all three methods. Therefore, even if damage tool interpolation changes result in different lifetime magnitudes, the relationship between the three aiming method's performance is not expected to alter unless lifetime prediction methodology significantly changes.

Resulting lifetime profiles are further processed to estimate minimum panel lifetimes for each case. Minimum panel lifetimes are plotted in Figure 4.15. Note that the lifetime axis is limited to 0–50 years, but several less critical panels have extremely high lifetimes and are hidden. It can be seen that both of the methods developed in this work significantly outperform the uniform aiming baseline case. When comparing the informed heuristic to the informed *SolarPILOT* method, the informed heuristic achieves higher panel reliability until approximately 200 MW<sub>th</sub>. This work shows that there exists a critical design power for a receiver area, above which achieving reliable power is expected to be infeasible regardless of the aiming strategy employed. However, informed heuristic is estimated, based on these results, to be capable of achieving 10–20 MW<sub>th</sub> of additional power compared to the informed *SolarPILOT* method before sub-30 year lifetimes are observed. Table 4.1 clearly shows that even in the high power design regions where failure occurs, the informed heuristic succeeds at minimizing the number of failures that occur. The value of the informed heuristic is its ability to update and consult lifetime profiles during image placement, which leads to better mitigation of flux violations caused by unideal image conditions. Uniform aiming performance suffers from failing to match flux levels according to local fluid temperature — allowing flux limits to vary locally is shown in this work to be extremely beneficial.

Table 4.1: Number of panels that fail before 30 years of operation. IA-SPT stands for informed *SolarPILOT* method, and IA-Heuristic is the informed heuristic.

Power MW <sub>th</sub>	Uniform	IA-SPT	IA-Heuristic
140	4	0	0
160	4	0	0
180	5	2	0
200	8	5	1
220	8	6	4






Figure 4.15: Minimum lifetimes over a range of test design powers for uniform, informed heuristic, and informed *SolarPILOT* aiming methods.

## 4.6 The Impact of Informed Aiming on Design

Present work now investigates how improved aiming strategies impact receiver design. We choose the LCOH metric to quantify designs because it considers panel reliability, system cost, and yearly net energy simultaneously. LCOH is evaluated across a series of aspect ratios for uniform and informed *SolarPILOT* aiming cases. Findings provide insights on overall trends in optical efficiency, reliability, and minimum cost versus aspect ratio. First, a cost model from literature is introduced. Second, the corresponding performance analysis is presented for evaluating net energy over off-design conditions. Third, the overall methodology is described, which includes damage, thermal, optical and cost analysis tools. Fourth, results are presented over height to width aspect ratios of 0.23–2.70. This LCOH study

places the reliability and optical efficiency improvements of informed aiming into an economic context and demonstrates that significant cost reductions are possible.

## Model for Levelised Cost of Heat

Broadly speaking, Levelised Cost of Heat (LCOH) refers to the ratio of expenditures to collected energy deposited in the heat transfer fluid. This metric is chosen because it is specific to the collection side of a CSP plant and neglects costs of the plant that do not change with respect to design aspects of the receiver or solarfield. Reductions in LCOH are expected to directly contribute to reductions in LCOE as long as our communicating variables to other parts of the system remain consistent. The hot and cold salt tank temperature requirements remain constant for every evaluated receiver/field design so that power block operation is unaffected. We select the existing LCOH model proposed in Gentile et al. (2024) to use as a model basis because it is recently published, developed specifically for power towers, and considers all of the design parameters and inputs relevant to the present work.

Present work builds a Python code implementation of the LCOH model based on its published description. The current paper focuses on the application of this model rather than its development, and therefore does not fully repeat the cost model discussion of the original work in Gentile et al. (2024). Instead, the objective of the current section is to state any assumptions or interpretations the present author makes to fully implement the literature model. A brief overview of the model is provided for context and so that the reader has a high level understanding of the model, but readers should refer to Gentile et al. (2024) for a detailed discussion of the cost model.

The cost model defines LCOH ( $\text{€}/\text{MWh}_{\text{th}}$ ) in annual terms as defined in Eq. 4.7, where CAPEX is the capital expenditures, CRF is the capital recovery factor, OPEX is the yearly operating expenditures, and  $Q_{\text{HTF}}^y$  is the yearly energy collected by the heat transfer fluid.

$$\frac{\text{CRF} \times \text{CAPEX} + \text{OPEX}}{Q_{\text{HTF}}^y} \quad (4.7)$$

The capital recovery factor incorporates the effects of discount  $\tilde{d}$  and inflation  $\tilde{i}$  and depends on plant lifetime  $N_{life}$  via Eq. 4.8,

$$CRF = \frac{\left( \frac{\tilde{d}-\tilde{i}}{1+\tilde{i}} \left( 1 + \frac{\tilde{d}-\tilde{i}}{1+\tilde{i}} \right)^{N_{life}} \right)}{\left( 1 + \frac{\tilde{d}-\tilde{i}}{1+\tilde{i}} \right)^{N_{life}} - 1} \quad (4.8)$$

Capital expenditures are based on total equipment cost and adjusted based on contingency and engineering and procurement costs. The total equipment cost considers receiver, solar field, tower, piping, pump, control system, and spare parts costs. Of these sub-categories, tower, piping, pump, control system, and parts costs are scaled using methods from Turchi and Heath (2013). Solar field cost is dependent on the required solar field area ( $A_{SF}$ ) and how much total land area is required. Solar field area is equal to the number of heliostats multiplied by the mirror area of each heliostat. The required solar field land area is dependent on the solar field area as well as how much space is between each heliostat and its neighbor. Finally, the total required land area is proportional to how much solar field land area is required according to an assumed solar land multiplier. Present work calculates total required land ( $A_{L,tot}$ ) based on an assumed heliostat field area to required land factor ( $F_{SF-L}$ ) of 5.8, and a solar land multiplier  $S_{LM}$  of 1.3, according to Eq 4.9,

$$A_{L,tot} = A_{SF} F_{SF-L} S_{LM} \quad (4.9)$$

The present work assumes a  $S_{LM}$  of 1.3 because this is a default value in *SolarPILOT*. Our assumed value for  $F_{SF-L}$  is conservative based on simulations in *SolarPILOT*. Solar field land area is layout specific, which means that design parameters such as tower height, thermal design power, and receiver size impact how compactly heliostats are spaced. Several layouts are generated for a flat receiver area of 270 m<sup>2</sup> and tower height of 170 m. Layout settings do not restrict the maximum field radius and the design power ranges from 160–220 MW<sub>th</sub>. Calculated ratios of solar field area to solar field land area were in the range of 4.8–5.8. As expected,

increasing design power and decreasing tower height lead to larger  $F_{SF-L}$  factors.

Receiver cost is defined in Eq. 4.10, where  $C_{rec,eng}$  is the engineering cost,  $C_{rec,ref}$  is the reference receiver cost,  $D_{rec}$  is the diameter,  $D_{rec,ref}$  is the reference diameter,  $H_{rec}$  is the height,  $H_{rec,ref}$  is the reference height,  $A_{rec}$  is the receiver area,  $c_{compl}$  is the complexity cost,  $c_{compl,ref}$  is the reference complexity cost,  $A_{coat}$  is the area required to be coated, and  $c_{coat}$  is the cost of application and purchase of Pyromark 2500 coating on the tubes. It can be seen in Eq. 4.10 that the model is scaling input geometry by reference receiver geometry, and this equation is originally derived for cylindrical receiver cases. The present work is analyzing LCOH of flat receivers, so we convert receiver width into an equivalent diameter by dividing by  $\pi$ . Additionally, we assume the that the entire tubes' surface area is being coated.

$$C_{rec} = C_{rec,eng} + C_{rec,ref} \left( \frac{D_{rec}}{D_{rec,ref}} \right) \left( \frac{H_{rec}}{H_{rec,ref}} \right)^{0.6} + A_{rec} (c_{compl} - c_{compl,ref}) + A_{coat} (c_{coat,purch} + c_{coat,app}) \quad (4.10)$$

Complexity cost accounts for material and labor costs of the receiver according to Eq. 4.11, where  $m_{metal}$  is the total mass of metal required,  $c_{metal}$  is the alloy cost,  $N_t$  is the number of tubes per panel,  $N_p$  is the *total* number of panels,  $c_{lab}$  is the labor cost/hour,  $N_{h,nozzles}$ ,  $N_{h,welds}$  &  $N_{h,clips}$  are the labor hours for nozzles welds & clips.

$$c_{compl} = \frac{m_{metal}c_{metal} + N_t N_p c_{lab} (2N_{h,nozzles} + 2N_{h,welds} + N_{clips} N_{h,clips})}{A_{rec}} \quad (4.11)$$

We calculate the mass of metal according to Eq. 4.12, where  $m_t$  is the mass of a single tube and  $m_{mfd}$  is the mass of the manifold. Headers connect the flux-experiencing absorber tubes to the manifolds, and are assumed to add additional 5% of the total tube mass, based on page 91 of Kelly (2010). Manifold mass is calculated based on the required diameter and assumed wall thickness. The manifolds, which

distribute and collect salt flow at the top and bottom of the panel, are assumed to have a T-shape, with half of total flowpath mass flow traveling in either direction. The required manifold diameter is assumed to be that which maintains constant velocity. Thus, each side of the manifold should have roughly the same flow area as half of the total tube flow area of the panel. This assumption enables calculation of the manifold diameter. Thickness is assumed to be 9.53 mm, which corresponds to NPS 14, schedule 30. For simplicity, we assume that the headers and manifolds are constructed from the same material as the tubes, and that the number of tubes per panel is calculated by dividing the panel width by outer tube diameter.

$$m_{\text{metal}} = 1.05m_t N_p N_t + m_{\text{mfd}} \quad (4.12)$$

In general, a receiver design's cost is modeled as a function of its height, width, number of panels, number of tubes, and material.

Annual operating expenditures include two terms: a fixed fractional cost of the receiver, assumed to be 1.4% of the receiver cost, and panel replacement costs. Replacement cost for a single panel is calculated by dividing the receiver cost by the number of panels and correcting with the contingency and engineering and procurement costs used to adjust capital expenditures. The annualized operating expenditures, therefore, are defined in Eq. 4.13, where  $N_{\text{repl}}$  is the total number of panel replacements during plant lifetime, and  $c_{\text{repl}}$  is the cost of a single replacement.

$$\text{OPEX} = \frac{N_{\text{repl}} c_{\text{repl}}}{N_{\text{life}}} + 0.014 \text{ CAPEX} \quad (4.13)$$

Annual energy ( $Q_{\text{HTF}}^y$ ) is calculated according to Eq. 4.14, where  $N_{\text{repl}}$  is the total number of replacements,  $N_{\text{dy,o}}$  is the number outage days required to replace a panel,  $\dot{Q}_{\text{HTF},i}$  is the hourly receiver power to fluid,  $P_{p,i}$  is the hourly electric pump power, and  $\Delta t_i$  is the hourly timestep, and  $\eta_{\text{PB}}$  is the estimated power block efficiency.

$$Q_{\text{HTF}}^y = \frac{1 - N_{\text{repl}} N_{\text{dy,o}}}{365 N_{\text{life}}} \sum_{i=1}^{8760} \left( \dot{Q}_{\text{HTF},i} - \frac{P_{p,i}}{\eta_{\text{PB}}} \right) \Delta t_i \quad (4.14)$$

Dividing the required pumping power by power block efficiency quantifies the power lost to pumping on equal terms as the thermal power of that timestep. Each panel replacement is assumed to require 10 days to complete.

The present work's Python implementation is compared to reported receiver costs and overall capital expenditures in Gentile et al. (2024). Operating expenditures are not reported in Gentile et al. (2024) and annual energy is calculated in a simplified manner compared to the reference. Regarding receiver and capital expenditure calculations, initial evaluation finds that estimated receiver costs differ by an approximate factor of two. However, present work cross-references the source of the receiver cost formula in Kelly (2010) and calculates similar values for the case described in that paper's Table 25. Present work notes that Gentile et al. (2024) calculate significantly higher receiver costs than the reference receiver yet with smaller dimensions, which suggests there may be a discrepancy between the model and reported results. Receiver costs from both the present work and Gentile et al. (2024) align well when the present work artificially decreases the reference receiver geometry from 20 m to 10 m. Also, capital expenditures, which include receiver cost, are higher for the present model than that in Gentile et al. (2024). This discrepancy may be due to differences in estimating the total land area — the present work applies factors as described, but the literature model description does not discuss the required land area calculation. Values align well when the present work assumes that land area is simply equal to solar field area. Thus, there are important discrepancies between the present work's model implementation and the reported output from Gentile et al. (2024) in terms of reported receiver cost and how land area is calculated. 1–2 sample outputs from Gentile et al. (2024) are drawn from Tables 10–12 of that work and total capital expenditures are found to be within 3% of each other if we modify our Python implementation to assume no scaling between solar field and total land area. We assume that the developed model is satisfactory, therefore, and use a reference diameter of 20 m as described in the source paper Kelly (2010) and use the land multiplication factors determined in this section.

To summarize, the LCOH model adopted by the present work captures sufficient

detail regarding the solar field and receiver costs. Inputs include: design power, tower height, receiver height, receiver width, panel width, tube diameter, tube thickness, material, heliostat area, total land area, plant design lifetime, required panel replacements, annual energy to HTF, and pumping power. The next section discusses how the present work estimates the energy to HTF, pumping power, and reliability inputs to the LCOH model.

## Performance Analysis & Modeling Methodology

As discussed in Section 4.6, calculating the LCOH requires estimates of power to HTF fluid, pumping power, and required number of panel replacements. Calculating the LCOH of a receiver-field system is a complicated task that requires integration of four different models already discussed in this work: (i) the damage model, accessed via the tool developed in Chapter 3 (ii) the optical model, which is *SolarPILOT* accessed via the *CoPylot* interface (iii) the thermal model from Martinek et al. (2021), adapted for flat receivers in Section 3.5 (iv) the cost model from Gentile et al. (2024), implemented as discussed in Section 4.6. Present work uses a simplified method for estimating LCOH performance inputs that considers off-design operation based on field efficiency, receiver efficiency, and pumping power results from representative clear-sky simulations. Yearly values are obtained by interpolating performance curves at average DNI levels that vary with season. The current section discusses simulation methods, provides an overview of each performance analysis, and summarizes the analysis process into a series of steps.

**Simulation Settings and Parameter Values** Each model requires a large number of inputs to simulate a given case, so this discussion aims to inform the reader of key input values provided to each model during the conducted aspect ratio study. Inputs are organized according to model, and unspecified inputs should be assumed equal to simulation default settings.

In general, each optical simulation is defined by a heliostat layout, selected aiming method, and heliostat attributes. Solar field layouts are generated in *SolarPILOT*

using the procedure titled *Heliostat Field Layout Procedure* in the *SolarPILOT* user manual, and the process is briefly summarized as follows. *SolarPILOT* generates a field of possible heliostats, simulates their individual performance over a series of representative days, and selects the highest performing units. Heliostats are selected in descending order of performance until the nominal design power input is met. The resulting solar field's actual performance may be lower than expected, depending on the aiming strategy and solar position used during actual simulation. Therefore, the present work standardizes its process of layout generation by requiring that the layout achieve the requested *power to fluid* under the following conditions:

1. uniform aiming as executed by *SolarPILOT* with its built-in image size priority method
2. two standard deviations of offset from the edges for all images
3. simulated day, month, and hour of 6, 21, and 12 respectively.
4. Overall receiver area of  $270 \text{ m}^2$  for the current study, and receiver height and width according to the case's aspect ratio.
5. A tower height of 170 m
6. Assuming a nominal receiver efficiency of 0.91, which accounts for absorption and thermal losses

Note that the assumed heliostat attributes also impact layouts, and we present these fixed parameters later in this discussion. These assumptions are implemented in code to generate all layouts. For the objective power to fluid, the code inputs a nominal power to the *CoPylot* instance, the *CoPylot* instance generates a layout, and then the code simulates the flux profile and estimates the actual power to fluid. This code iteratively modifies the nominal power input and restarts the process until the estimated power to fluid is within  $1 \text{ MW}_{\text{th}}$  of the objective power to fluid. Two layout examples are shown in Fig. 4.18 for a desired power to fluid of  $200 \text{ MW}_{\text{th}}$



and aspect ratios of 0.23 and 2.70. It can be observed that the receiver's aspect ratio noticeably impacts the layout generation, with wider receivers resulting in wider layouts and vice versa. This is expected, as heliostats on the sides of the field have more significant horizontal image 'smearing' when they impact the receiver. Spillage associated with image smearing decreases for increasing receiver widths and thus the side heliostats become more viable options.

The heliostats themselves are assumed to have a height and width of 12.2 and reflectivity area ratio of 0.97, which results in an approximate reflective area of 144.4 m<sup>2</sup> per heliostat. Heliostat optical error, which includes pointing, slope, and beam error, will decrease the efficacy of informed aiming due to the resultant image spreading. Image error is defined in *SolarPILOT* to be the standard deviation of the reflected image from the heliostat according to Eqn. 4.15, where  $\sigma_{el}$  is the elevation pointing error,  $\sigma_{az}$  the azimuthal pointing error,  $\sigma_s$  the slope error in the X/Y directions, and  $\sigma_r$  is the reflected beam error in the X/Y directions. Slope error pertains to the waviness of the mirror surface, while reflected beam error models the specularly of the mirror surface.

$$\sigma_{total} = \sqrt{4(\sigma_{el}^2 + \sigma_{az}^2 + \sigma_{s,x}^2 + \sigma_{s,y}^2) + \sigma_{r,x}^2 + \sigma_{r,y}^2} \quad (4.15)$$

*SolarPILOT* settings are 0 for azimuthal and elevation pointing error,  $1.53 \times 10^{-3}$  rad for slope error in X and Y directions, and  $2 \times 10^{-4}$  rad for reflected beam error.

*CoPylot* directs *SolarPILOT* use its image size priority aiming heuristic to achieve one of the following flux profiles: either a uniform profile with a fixed edge offset or a custom profile defined by the present work's aiming informer. The latter option is introduced in Section 4.3 and referred to as 'informed *SolarPILOT* aiming'. Both strategies' edge offset parameters in *SolarPILOT* are fixed at 2 standard deviations.

The informed *SolarPILOT* aiming method's reliability performance is shown in Section 4.3 to closely depend on the input flux grid. Grid inputs to *SolarPILOT* are specific to each receiver geometry and generated by using the informed aiming module's **fully\_utilize\_receiver\_maximize\_lifetime** function, with a specified vertical edge offset of 2 m and assuming design point mass flow. Informed input grid




Figure 4.16: Informed flux grid input to *SolarPILOT* for a 0.23 aspect ratio, 170 MW<sub>th</sub> case.

resolution and optical simulation grid resolution are held equal for each case and defined on a per m basis. The resolutions are 3.6 points/m in the vertical direction and 3 points/m in the horizontal direction, which corresponds to approximately 10 points/m<sup>2</sup>. Example flux grids for aspect ratios of 0.23 and 2.70 are shown in Fig 4.16 and Fig. 4.17. The x and y axes proportions accurately represent the proportionality of the receiver heights and widths. Maximum fluxes of the 0.23 and 2.70 aspect ratio flux grids are 1380 and 1213 kW/m<sup>2</sup> respectively. Wider, shorter receivers are expected to require higher flux densities to achieve the same power because the amount of usable receiver area is lower than for taller, narrower receivers. Shorter flat receivers' low flux vertical offset regions take up a more significant fraction of each panel. Present work's aiming method does not account for horizontal offset, which, if considered, would likely cause flux densities for all aspect ratios to be more similar than that of current work.

Table 4.2 defines fixed thermal model parameters for this study and resolutions used. Axial resolution of each simulated tube ( $n_{z_{th}}$ ) varies according to tube length i.e. receiver height. The names  $n_r$  and  $n_\theta$  refer to radial and circumferential resolution of each tube. Both the horizontal and the vertical resolutions of the thermal model's flux profile interpolation definition are set to be equal to the optical model's output flux profile resolution. These resolutions also vary proportionately to the receiver dimensions.  $f_{block}$  refers to the blocking factor, which is introduced




Figure 4.17: Informed flux grid input to *SolarPILOT* for a 2.70 aspect ratio, 170  $\text{MW}_{\text{th}}$  case.

in Section 3.5 to prevent adjacent simulated tubes from artificially interpolating the same flux value.

The aspect ratio and receiver width vary with case. Ranges are 0.23–2.70 and 10–34 for aspect ratio and width respectively. Height for each case is calculated using the definition of aspect ratio:  $H = \text{AR} \times W$ .

Cost model parameters are defined in the section introducing the model. Recall that total land is calculated by multiplying the solar field area by a factor of 7.5. Also, second law efficiency of the power block, which is used to equate pump power to an equivalent amount of power to fluid, is assumed to 0.45.




Figure 4.18: Sample layouts for 170 MW<sub>th</sub> power and a receiver with aspect ratio of 0.23 (a) and 2.70 (b).

The damage tool instance for each case is defined with Haynes 230 material, conduction-convection ratio of 0.75, and a linear N-dimensional interpolator for obtaining lifetimes.

**Estimating Yearly Solar Resources** Power to fluid is directly dependent on the available solar resource at the given location. CSP systems use the direct normal normal radiation (DNI) metric, which is defined as the intensity of direct rays from the sun on a surface constantly held normal to the incoming ray direction. DNI varies with solar position and seasonal atmospheric conditions. We assume a location of Daggett, California, with latitude and longitude of 34.87,-116.78. Hourly DNI data is obtained via a typical meteorological year file from Sengupta et al. (2018).

Present work seeks to simplify the analysis process so that a reduced number of optical simulations can be performed. To accomplish this, yearly DNI data is summarized into four average DNI hourly profiles according to the four seasons. Each season's months are grouped according to the month before, of, and after the representative solar day. For example, the summer season considers all daily solar

Table 4.2: Thermal model parameters and simulation settings common to all cases for the aspect ratio study.

fixed parameter	value	model resolution	value
panel width	1 m	$n_{z_{th}}$	3.6 pts/m
tube outer diameter	50.8 mm	$n_r$	5
tube wall thickness	1.25 mm	$n_\theta$	99
tube absorptivity	0.96	$n_{x_{flux}}$	3 pts/m
tube emissivity	0.87	$n_{z_{flux}}$	3.6 pts/m
pump efficiency	0.85	$f_{block}$	9
flowpaths	2	sim. tubes	3/panel
timestep	1 hr		
ambient temp.	25°C		
relative humidity	25%		
wind speed	0 m/s		
tube material	Haynes 230		

profiles of May, June, and July. At each hour, hourly DNI from every solar profile is averaged to create a 24-hour average profile. Results are shown in Fig. 4.19 for every season. It can be seen that fall and summer are the most productive solar seasons for CSP operation at this location. Spring and fall are expected to have seasonal rainfall and humidity differences, which would explain the difference in average DNI profiles. Next, power to fluid is estimated via thermal model simulation.

**Calculating Power to Fluid** Generally speaking, power to fluid for a given receiver and layout is assumed to be a function of the available solar resource (DNI), receiver efficiency  $\eta_r$ , and overall field efficiency  $\eta_f$ . Calculating hourly power to fluid requires us to characterize how these variables change throughout the year and period of operation. DNI resource as a function of season and hour of day is characterized as already presented. Receiver efficiency is defined as  $\eta_r = \dot{Q}_f / \dot{Q}_{inc}$ , where  $\dot{Q}_f$  is power to fluid and  $\dot{Q}_{inc}$  is incident power on the receiver. This definition




Figure 4.19: Average direct normal irradiation per season at Daggett, California location, using a typical meteorological year of data.

means that receiver efficiency accounts for absorptivity losses, radiation losses, and convection losses. Present work characterizes receiver efficiency as primary function of incident power, and uses a parabolic curve fit to interpolate model data. Curve fits are individual for every simulated day, and an example of a curve fit for a 200 MW<sub>th</sub> design power, 18 m width, 15 m height flat receiver subject to uniform aiming on a summer solstice day is shown in Fig. 4.20. Model results show that receiver efficiency improves with higher incident power. Present work assumes for this analysis that mass flow is perfectly controlled so that outlet temperature remains constant at 565°C. Consequently, the general trend seen in Fig. 4.20 is expected, as an increase in incident power corresponds to a higher mass flow and salt side heat transfer coefficient.

Field efficiency for a given day, layout, receiver, and tower height is characterized




Figure 4.20: Model results, curve fit, and interpolated receiver efficiencies for an example summer solstice day on a typical flat receiver subject to uniform aiming.

according to its relationship with simulation time of day, which is directly related to solar position. The function `np.interpolate` is used to 1D interpolate at desired points between simulation points. Similar to field efficiency characterization, each interpolation is specific to the day being simulated. Figure 4.21 shows an example of model field efficiencies from the same 200 MW<sub>th</sub> uniform aiming, summer solstice simulation. It can be observed that model results only exist for several hours around solar noon. The thermal model is set to only solve for hours daylight hours with DNI greater than a cutoff value of 250 W/m<sup>2</sup>. Field efficiencies outside the range of model data are assumed to be equal to the nearest data point, as shown in Fig. 4.21. This extrapolation is not expected to affect results because, as discussed later, power to fluid for a given timestep is automatically set to zero if DNI is below the model cutoff. Thus, the field efficiency may be artificially high but it is irrelevant because




Figure 4.21: Field efficiency results from each simulated time point in the example summer solstice day for a flat receiver subject to uniform aiming as well as interpolated values at intermediate time points.

interpolations at sub-cutoff DNIs are set to zero.

Characterization of field efficiency and receiver efficiency enables us to calculate the power to fluid for any hour  $i$  according to Eq. 4.16, which multiplies the total available energy from the solar field by the overall optical efficiency for that hour and the receiver efficiency as a function of the incident flux power of that hour.

$$\dot{Q}_{\text{HTF},i} = \text{DNI}_i \times A_{\text{SF}} \times \eta_{f,i} \times \eta_{r,i}(\dot{Q}_{\text{inc},i}) \quad (4.16)$$

Figure 4.22 shows the calculated power to fluid hourly profile for a given day compared to the simulated clear-sky model output that its performance characterizations are based on. The clear-sky solstice day has a peak DNI of  $950 \text{ W/m}^2$ , while the average summer DNI peak is approximately  $830 \text{ W/m}^2$ . The lower values






Figure 4.22: Model-predicted power to fluid profile for a summer solstice day for design parameter case described in text. Interpolated power to fluid is based on receiver and field efficiencies derived from the model results and the average summer DNI profile shown in Fig. 4.19.

of average DNI cause the overall summer interpolated power to fluid profile to be lower in magnitude than that of the simulated solstice day.

**Calculating Required Pumping Power** Power is required to move salt from the cold tank up the tower and through each flowpath in the receiver, where the current analysis assumes two parallel receiver flowpaths. An existing pressure loss model from Martinek et al. (2021) is implemented, and the reader is directed to that publication for additional details. A brief overview of the model is provided here to convey what the model accounts for. The pumping pressure model accounts for losses due to tower height, straight pipe section frictional loss, changes in height as salt traverses the receiver flowpath, and tube bends where headers attach to the

panel manifolds. Equation 4.17 shows the calculation for loss along a differential element of length  $dz$  of tube, where  $f_d$  is the Darcy friction factor,  $\rho$  is the salt density at that element,  $u_{\text{HTF}}^2$  is the salt velocity,  $D_i$  is internal tube diameter, and  $g$  is acceleration due to gravity. Note that Eq. 4.17 adds a positive increase in pressure if salt is ascending the tube path and negative if salt flow is descending.

$$dp = \frac{1}{2} f_d \rho u_{\text{HTF}}^2 \frac{dz}{D_i} \pm \rho g dz \quad (4.17)$$

Each panel also has pressure loss associated with bends near connection points according to Eq. 4.18, where  $N_b$  is the assumed number of 4 effective bends for each tube in a panel, and  $(L/D)_e$  is the assumed effective length to diameter ratio of 30. Note that an effective number of bends is assumed — tube paths can include various bends depending on how they are designed. Readers should refer to Rodríguez-Sánchez et al. (2014) for an illustration of tube bends.

$$\Delta P_{\text{bend}} = \frac{1}{2} N_b f_d \rho u_{\text{HTF}}^2 (L/D)_e \quad (4.18)$$

Each panel's pressure loss is calculated by integrating over the straight section of tube length with Eq. 4.17 and adding to it the pressure loss due to bends from Eq. 4.18. Total flowpath pressure loss is the sum of tower pressure loss and the sum of all panel pressure losses.

Thus, pressure loss at any given timestep is dependent on several design parameters: the design power, tower height, number of panels in each series, total length of straight tube, tube inner diameter, and each panel's number of tubes in parallel. Pressure loss is also dependent on the salt velocity and therefore is coupled to the power to fluid in the case that mass flow is controlled to maintain constant outlet temperature. Pressure loss is related to each flowpath's electric pumping power according to Eq. 4.19, where  $\dot{m}_{\text{fp}}$  is the mass flow per flowpath,  $\Delta P_c$  total pressure loss through the flowpath, and  $\eta_p$  is the pump efficiency.

$$P_{\text{e,fp}} = \frac{\dot{m}_{\text{fp}}}{\rho} \frac{\Delta P_c}{\eta_p} \quad (4.19)$$




Figure 4.23: Pumping power plotted against power to fluid for the summer solstice simulation. Parabolic curve fit and interpolated values at un-simulated power to fluid points are also shown.

Present work characterizes pumping power as a function of power to fluid for each simulated day. Figure 4.23 shows the relationship between pumping power and power to fluid for the summer case. Note that the pumping power for hours of the day with no power to fluid is not interpolated from the curve and instead enforced to be zero. The curve fit does not include a zero point because times of day with no operation are simply omitted from the daily simulations.

**Reliability Assessment** Reliability refers to the creep-fatigue lifetimes of each panel in the receiver. Lifetime is assessed in a manner consistent with the cycle assumptions made in Chapter 2. We input the tube crown temperature differences and corresponding fluid temperatures at every solution node of the thermal model into the damage tool's `get_LTEs` method. Thermal operating points are

obtained by post-processing results from a simulation of design point, which is summer solstice at solar noon. The resulting lifetime profile is post-processed using **calc\_minimum\_panel\_LTEs** to determine the number years each panel can operate before failure. The total number of replacements  $N_{\text{repl}}$  is calculated according to Eq. 4.20, where  $\text{LTE}_{p,i}$  is the lifetime of a panel, and  $N_p$  is the total number of panels in the receiver. Each panel's calculation is floored to result in integer values that exclude each position's original panel from consideration.

$$N_{\text{repl}} = \sum_{i=1}^{N_p} \left\lfloor \frac{N_{\text{life}}}{\text{LTE}_{p,i}} \right\rfloor \quad (4.20)$$

**Overview of Performance Analysis Procedure** The following steps are performed in order to analyze the yearly performance of the each receiver case.

1. A clear-sky thermal simulation is run for a summer solstice day.
2. Performance curves for  $P_{p,i}$ , and  $\dot{Q}_{\text{HTF},i}$  are obtained by characterizing the system variables as described in this section. These curves are ultimately a function of DNI and time of day. An average hourly profile is therefore built by assuming average summer DNI levels and interpolating performance curves at regular time intervals.
3. The resulting 24-hour profiles are repeated for as many days as are in the season to form a sequence of daily profiles.
4. Steps 1,2 & 3 are repeated for fall equinox, winter equinox, and spring solstice to create season sequences for fall, winter, and spring seasons.
5. Seasonal sequences are combined to yearly performance arrays for pump power and power to fluid that each have a size of 8760 elements, which corresponds to the number of hours in a year.
6. The design point simulation is performed and input to the damage model to yield the total number of required panel replacements.

7. LCOH is calculated with the receiver parameters, solar field parameters, and yearly performance arrays as input.

The next section presents detailed results from the parametric study and compares the effects of aiming strategy on LCOH.

## Parametric Aspect Ratio Study

In this section, LCOH is evaluated for two power scenarios, 170 and 200 MW<sub>th</sub>, and at each power with two different aiming strategies, uniform and informed *SolarPILOT* aiming. For each scenario and case, the height to width aspect ratio of the flat receiver is varied from 0.23 to 2.70. We also show and discuss the trend of panel replacements versus aspect ratios for both aiming methods and explain them in terms of the Haynes 230 damage map. Results show that optical efficiency primarily drives minimum LCOH, especially when operating costs are low compared to capital costs. However, the informed *SolarPILOT* aiming method achieves the lowest overall LCOH.

**Results at 170 MW<sub>th</sub> Design Power** We select a baseline design power rating with the objective of equating the average flux density to that of the default cylindrical receiver in *SolarPILOT*. Accounting for all losses yields an estimated 627 kW/m<sup>2</sup> average flux density. For our baseline receiver area of 270 m<sup>2</sup>, this average flux density would achieve approximately 170 MW<sub>th</sub>. Performance modeling is conducted using uniform and informed *SolarPILOT* aiming strategies. Figure 4.24b shows the capital expenditures and total lifetime operating expenses over the 30 year lifetime of the receiver and solar field for a range of aspect ratios. Capital expenditures for both aiming methods exhibit a minima at an intermediate aspect ratio. The minimum capital cost for informed and uniform aiming is 154 and 155.9 M€ respectively at 1.05 aspect ratio. The difference in capital cost is explained by the informed aiming method requiring a slightly smaller subset of the layout required by the uniform aiming method due to higher optical design efficiency. Total operating costs are dependent on capital expenditures and well the number of

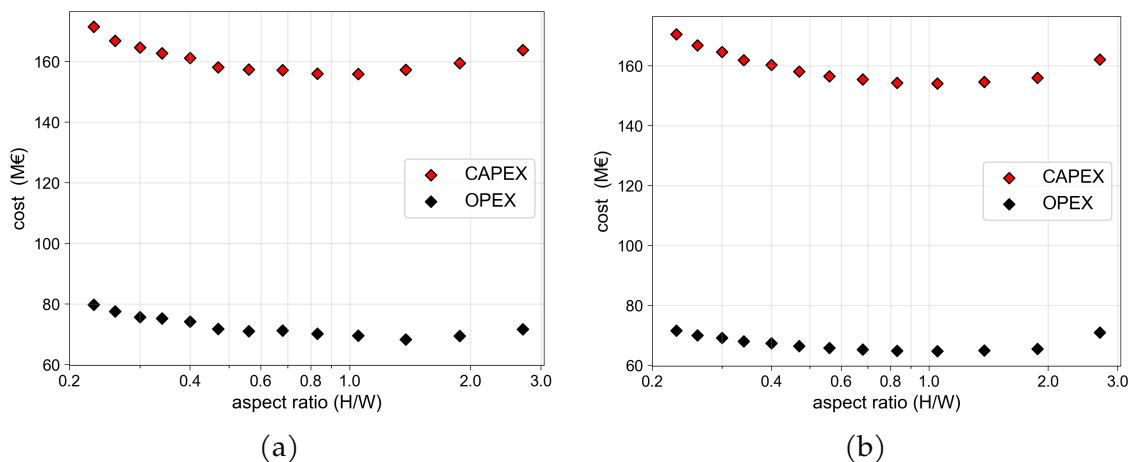


Figure 4.24: Capital and total lifetime operating expenses for an aspect ratio study at  $170 \text{ MW}_{\text{th}}$ . (a) shows results when uniform aiming is used, and (b) when informed *SolarPILOT* aiming is used.

panel replacements, as defined in Eq. 4.13. Minimum operating cost for informed aiming is  $64.7 \text{ M€}$  at 1.05 aspect ratio. Uniform aiming yields a minimum operating cost of  $68.3 \text{ M€}$  at 1.38. Thus, we see that there is an offset between the location of the minimum capital cost and the minimum operating cost for uniform aiming.

The average difference in operating costs between the two aiming methods for each is calculated to be  $5.5 \text{ M€}$ . If we consider that the average difference in capital costs for each aspect ratio is only  $1.26 \text{ M€}$ , it can be concluded that panel replacements are the driving factor behind uniform aiming's higher operating expenses. Figure 4.25 shows that replacements are required for every aspect ratio case if uniform aiming is used. By contrast, informed aiming achieves 30 year reliability in every panel until the extreme aspect ratio of 2.7. In fact, 2.7 is the only aspect ratio at which these two aiming strategies achieve the same level of panel reliability. Interestingly, the number of required replacements for uniform aiming is inversely proportional to aspect ratio — the minimum lifetime for 0.23 aspect ratio is 12.2 years, and for 2.70 aspect ratio is 20.3 years.

Figures 4.26 and 4.27 show the difference in thermal operating points in the context of the Haynes 230 damage map. The thermal power of design point for




Figure 4.25: Total number of required panel replacements for uniform and informed *SolarPILOT* method aiming cases when design power is  $170 \text{ MW}_{\text{th}}$ .

these two cases differs by less than 1%, so they are comparable in terms of average flux. It can be seen that a taller, narrower receiver consolidates the undesirable thermal operating points onto fewer panels compared to the wider, shorter receiver case.

The improved placement of operating points achieved by informed *SolarPILOT* aiming is shown in Figures 4.28 and 4.29. Reliability increases for lower aspect ratio receivers compared to higher aspect ratio receivers when using informed *SolarPILOT* aiming because the shorter panels break the flowpath into smaller subsections, which results in lower increases in fluid temperature per panel. Undesirable image spreading that causes flux violation is less damaging because the allowable temperature difference does not largely vary between panels. Figure 4.29 demonstrates the undesirable effects of images that are larger than the extents of the panel they are

being aimed at. Starting at a low fluid temperature in either one of the flowpaths, we see that thermal points start at a high temperature difference and gradually decrease along the flowpath. Each set of panel thermal operating points is visually identified by observing the locations of rapid decreases in temperature difference, which correspond to the low flux offset regions at the top and bottom of a panel. As a result, the incident flux is intentionally high at the top center of the receiver where salt flow enters. The flow finishes the first panel at the bottom of the receiver and begins traversing back up the second panel. Unfortunately, image spreading from the intended flux on the first panel at the top of the receiver causes flux violations in the second panel. Notice that only some of the panel tubes are experiencing flux violations — the thermal points split into three general lines according to different temperature differences. A similar thing happens at every subsequent panel turn. Performance is still improved over uniform aiming, however. At 0.23 aspect ratio, the minimum lifetime is 90.9 years. For 2.70 aspect ratio, the minimum lifetime decreases to 22.0 years.

Figure 4.30 shows design point efficiencies and yearly net energy for both aiming methods. As expected, informed *SolarPILOT* aiming achieves higher efficiencies at every aspect ratio. There is an offset in the location of peak efficiencies between the two aiming methods, however. Net energy considers annual power to fluid, pumping power, and outages as defined in Eq. 4.14. Informed aiming net energy is more consistent across all aspect ratios and, on average, 16.1 GWh/year higher per aspect ratio case.

The LCOH results are shown in Fig. 4.31. Both aiming methods' trends in LCOH follow a parabolic shape on a log x-scale of aspect ratio. A minimum LCOH of 26.0 €/MWh exists for the informed case at an aspect ratio of 1.05. Uniform aiming's minimum LCOH is 27.8 €/MWh, occurring at an aspect ratio of 1.38. There is some numerical uncertainty in the uniform case due to the fact that each aspect ratio has its own layout. This uncertainty would not affect the informed case because the informed aiming is more efficient and is not always using all of the layout heliostats. The lower efficiency uniform method is subject to the exact performance limitations of the layout itself. Ultimately, the informed *SolarPILOT*






Figure 4.26: Thermal operating points overlaid on the Haynes 230 damage map for uniform aiming on a receiver with 0.23 aspect ratio at  $170 \text{ MW}_{\text{th}}$ .

method achieves a minimum LCOH that is 6.5% lower than the minimum achieved by uniform aiming method.

**Results at  $200 \text{ MW}_{\text{th}}$  Design Power** Design power is increased on the same receiver area by  $30 \text{ MW}_{\text{th}}$ . The same performance analysis is then applied over the same range of aspect ratios. Results demonstrate that the difference between aiming methods becomes more evident as average flux on the receiver increases.

The difference in capital and operating costs show a similar trend to that seen at  $170 \text{ MW}_{\text{th}}$ , but at different magnitudes. Referring to the  $200 \text{ MW}_{\text{th}}$  design power, informed *SolarPILOT* aiming cases' capital costs are on average 2 M€ lower than those of the corresponding uniform aiming cases. Total operating expenses are on average 9.1 M€ lower for informed *SolarPILOT* aiming cases. This is almost twice as high a difference as compared to results for operating expenses at  $170 \text{ MW}_{\text{th}}$  design power.




Figure 4.27: Thermal operating points overlaid on the Haynes 230 damage map for uniform aiming on a receiver with 2.70 aspect ratio at  $170 \text{ MW}_{\text{th}}$ .

The number of required replacements for uniform aiming cases increases significantly compared to  $170 \text{ MW}_{\text{th}}$  design power, as shown in Fig. 4.32. Overall trend of panel replacements follow the same trend as for  $170 \text{ MW}_{\text{th}}$  except that now uniform aiming requires more replacements per aspect ratio. Informed *SolarPILOT* aiming sustains high panel reliability until an aspect ratio of 1.38.

Yearly net energy for the informed *SolarPILOT* method is on average 21.0 GWh/yr greater than for uniform method at the same aspect ratios. Design point field efficiency results demonstrate similar findings to that of  $170 \text{ MW}_{\text{th}}$ . The design point efficiency of the informed *SolarPILOT* method is on average 2.2% higher than the uniform method over all tested aspect ratios.

The resulting levelized cost of heat is shown in Fig. 4.33. Minimum LCOH for uniform aiming is  $27.8 \text{ €/MWh}$  and  $25.5 \text{ €/MWh}$  for informed *SolarPILOT* aiming at 0.68 aspect ratio. The minimum LCOH using informed aiming is 7.6% lower than that of uniform aiming.




Figure 4.28: Thermal operating points overlaid on the Haynes 230 damage map for improved *SolarPILOT* aiming on a receiver with 0.23 aspect ratio at  $170 \text{ MW}_{\text{th}}$ .

**Aspect Ratio Study Summary** Aspect ratio study results show aspect ratio trends for reliability and LCOH, and results also compare the two aiming methods on a technoeconomic basis. Findings are based on performance analysis and cost modeling on an annual basis and accounting for temporal and seasonal variability in solar resource. Present work shows that informed *SolarPILOT* aiming outperforms uniform aiming by all metrics.

Both design powers share similar qualitative trends in reliability, efficiency, and LCOH. Required replacements decreases with aspect ratio for uniform aiming, but increases for informed *SolarPILOT* aiming. In the former aiming case, higher aspect ratios are favorable because they consolidate high damage regions into fewer panels. In the latter case, higher aspect ratios worsen the effect of image spreading between neighboring panels and increase the magnitude of fluid temperature increase from panel to panel. Design point efficiency is optimal at intermediate aspect ratios for both aiming methods because area is constant and extreme aspect ratios cut off




Figure 4.29: Thermal operating points overlaid on the Haynes 230 damage map for improved *SolarPILOT* aiming on a receiver with 2.70 aspect ratio at 170 MW<sub>th</sub>.

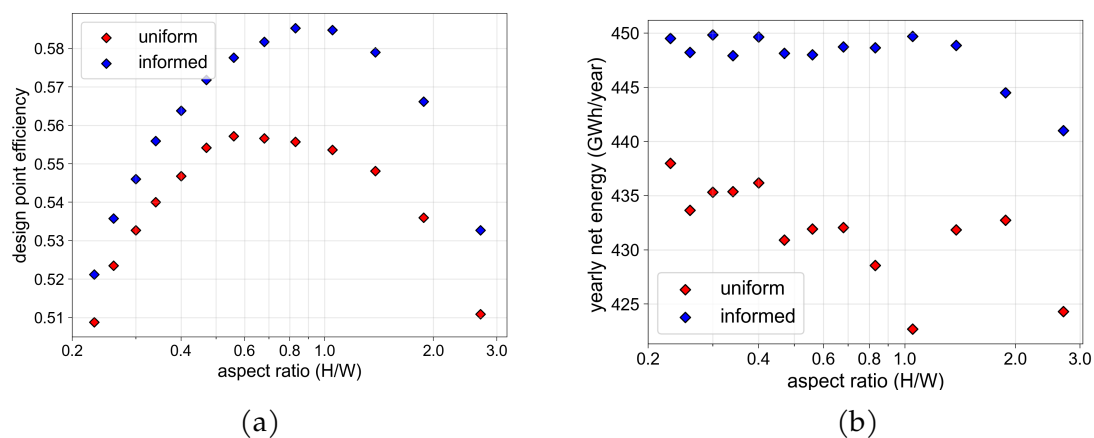


Figure 4.30: Design point efficiency and net energy for all aspect ratios and design power of 170 MW<sub>th</sub>. Edge offset is 2 standard deviations for both aiming methods.




Figure 4.31: Levelised cost of heat for uniform and informed *SolarPILOT* aiming methods for 170 MW<sub>th</sub>.

image borders in either the vertical or horizontal direction. As expected, LCOH trends generally follow the optical efficiency trends because solar field and land costs are a very significant factor in the total equipment cost.

Quantitative results from this study recommend height-width aspect ratios in the range of 0.68–1.05 for flat receivers in the design range presented here. Table 4.3 shows minimum LCOH and the corresponding aspect ratio for each aiming strategy and design power tested. Informed aiming consistently achieves a minimum LCOH at 1.05 while uniform aiming has minima on the extents of this range. Visual evaluation of Figures 4.31 and 4.33 show that the difference between minimum LCOH and neighboring aspect ratio LCOH values is small, thus authors suggest a range of aspect ratios instead of the single points seen in Table 4.3. Lower aspect ratios for informed aiming are less desirable from an LCOH perspective because




Figure 4.32: Total number of required panel replacements for uniform and informed *SolarPILOT* method aiming cases when design power is 200 MW<sub>th</sub>.

the increasing cost of capital expenditures does not yield any gains in reliability at these tested design powers. We can conclude that considering informed aiming improves levelized system cost but does not significantly affect the design's aspect ratio selection unless reliability is a severe challenge.

The difference in overall performance increases for higher design powers. Average differences between LCOH increases in magnitude from -1.62 to -1.92 €/MWh from 170–200 MW<sub>th</sub>. Yearly net energy average differences also increase from 16.1 to 21.0 GWh/year. Average differences in operating expenses over the receiver's lifetime change from -5.5 to -9.1 M€. Optical efficiency average differences marginally increase from 2.1% to 2.2%, where informed aiming is always more efficient than uniform aiming. These metrics all indicate that, when increasing design power and holding receiver area constant, informed *SolarPILOT* aiming becomes even more




Figure 4.33: Levelized cost of heat for two different aiming methods for 200 MW<sub>th</sub> design power.

appealing.

Informed aiming improves panel reliability and also outperforms uniform aiming in terms of efficiency. This finding is novel and due to the strategic arrangement of flowpaths and use of local flux limits. Allowing flux to vary according to fluid temperature better utilizes each panel in the receiver compared to global flux limits, which causes underutilization in some panels to reduce damage in critical panels. Placing both flowpaths at the center of the receiver is complementary to the local flux limits, as the center of the receiver has higher flux limits than the sides. Spillage, then, is naturally reduced compared to uniform aiming because more images are being aimed at the receiver's center. The current study shows that designs over a large of aspect ratios and layouts experience performance improvements compared to uniform aiming. Performance improvements, when placed in the economic

Table 4.3: Minimum LCOH and the aspect ratio at which the minimum occurs for two aiming methods at two different design powers.

aiming	power $\text{MW}_{\text{th}}$	Min(LCOH) €/MWh	$\text{AR}_{\text{min,LCOH}}$
uniform	170	27.8	1.38
	200	27.6	0.68
informed	170	26.0	1.05
	200	25.5	1.05

context of LCOH, are shown to be significant at a system level — a reduction of 6.5%–7.6% LCOH for 170 and 200  $\text{MW}_{\text{th}}$  design powers.

## 4.7 Summary of Informed Aiming and Design

This chapter builds on the previous chapters’ novel damage model and tools to directly improve panel creep-fatigue reliability and overall system cost.

Present work develops definitions and visualizations of ideal flux profiles that fully utilize receiver area at a given power and lifetime. We create several important procedures for creating flux maps and synthesizing them into flux grids to be used by other modeling scenarios. Testing results show that ideal flux aiming, if achieved, strategically places all thermal operating points on the desired contour, resulting in optimal flux levels with acceptable panel reliability. Minimum required area across a range of 100–400  $\text{MW}_{\text{th}}$  design powers are in the range of 150–285  $\text{m}^2$  for the given receiver parameters.

Two methods are presented as methods of achieving ideal flux aiming in unideal optical scenarios. The first method is computationally fast and achieves noticeable reliability improvements when compared to uniform aiming, but lacks robust handling of flux violations. The second method, although slower, further improves panel reliability by updating the thermal and lifetime profiles before selecting aim-points. This more robust method is shown to achieve fewer panel failures and markedly increase the achievable design power by 10–20  $\text{MW}_{\text{th}}$  without incur-



ring panel failures at less than 30 years. Both methods are significantly higher performing when compared to uniform aiming in terms of achieved panel lifetimes.

The economic and design effects of different aiming assumptions is investigated for a range of height-width aspect ratios. A cost model from literature is integrated with the previously presented optical, thermal, and damage modeling software tools to perform a detailed analysis of levelized cost of heat. Results from 170 and 200 MW<sub>th</sub> design power scenarios show that informed aiming's improvement in panel lifetimes translates to economic benefits. Additionally, LCOH of informed aiming cases further improves thanks to increased optical efficiency when compared to the baseline uniform aiming case. LCOH for systems using informed *SolarPILOT* aiming are always lower than uniform aiming. Minimum LCOHs are 6.5% and 7.6% lower than uniform aiming cases for 170 and 200 MW<sub>th</sub> scenarios. Interestingly, lower aspect ratios appear to be better suited for informed aiming methods while higher aspect ratios result in fewer panel replacements for uniform aiming methods. However, the analyzed power levels do not challenge the informed aiming method and therefore no significant replacement costs are incurred for informed aiming methods. As a result, aspect ratios close to one yield the lowest LCOH values in these two power scenarios. Future work should consider higher power scenarios to determine if the aspect ratio of minimum LCOH shifts due to more impactful panel replacement costs.

## 5 SUMMARY

---

The present work directly contributes towards operation and design of CSP systems by (a) developing a novel model and tool to characterize tube creep-fatigue lifetime as a function of thermal conditions (b) presenting methods of improved receiver utilization that also improve panel reliability and optical efficiency for central receivers. In addition to directly benefiting the CSP industry, planned work is expected to be generally applicable to other thermal problems. The novel flowpath analysis and design methods presented in this work can apply to any heat exchanger design task that must meet challenging lifetime requirements while maximizing thermal power. More specific conclusions are stated in order of chapter.

**Damage Modeling Findings** The present work both provides insight into suggested operating practices and can also serve as a useful receiver design tool. A FEA-based methodology for building damage maps of any given alloy is established. This methodology leverages full inelastic FEA analyses in *SRLIFE* and considers the important effect of stress relaxation at the tube crown. The power law extrapolation techniques predicts 20<sup>th</sup> day creep damage based on FEA simulation of the first four thermal cycles and proves to be sufficiently accurate for expected operating conditions and conservative in cases that experience extreme levels of stress relaxation. The Haynes 230 parametric study, using the methodology developed here, analyzes a receiver with baseline case geometry and assumes daily clear-sky flux profiles. Classification of baseline case tubes based on their critical temperature points and overlaying those operating points on the damage map shows that low  $T_f$  tubes will be operating with extremely high lifetimes, while higher  $T_f$  tubes display limited lifetime. Based on these results, it would be advantageous to further concentrate the flux on low temperature panels and reduce the flux concentration on high temperature panels to achieve a more uniform lifetime prediction for all receiver tubes. Fatigue damage is never found to be significant when compared to the corresponding creep damage incurred for any temperature in the baseline case. The large difference in magnitude between incurred fatigue and creep damages

for a clear sky case implies that additional cloud events causing smaller amplitude cyclic fatigue damage will not significantly affect a receiver tube's state of health. Note that this statement only applies to cloud effects related to strain range cycles. Operation that causes extreme temperature spikes that further aggravate creep damage should be avoided whenever possible.

Parametric analyses are performed for other candidate alloys, including Inconel 617, Haynes 282, Inconel 740H, SS 316H, and Incoloy 800H. Results show that only Inconel 617, SS 316H, and possibly Inconel 740H tubes experience significant fatigue damage. Lifetime estimates show that Haynes 282 and Inconel 740H exhibit excellent creep resistance for the simulated operating conditions and more extreme fluid temperatures and fluxes can be employed compared to the benchmark Alloy 230 case. The low yield strength of Incoloy 800H results in a highly reduced operational space unless the receiver tubes are operated in the stress reset region. Thus, with the primary exceptions of Inconel 617 and SS 316H, receiver tube fatigue damage is found to be negligible when compared to creep damage. Therefore, aiming strategies and mass flow control should focus on mitigation of creep, and such control strategies could increase cycling (if necessary) without significantly affecting tube lifetime.

Additional FEA campaigns at conduction-convection ratios of 0.25 and 0.75 are performed for Haynes 230, Haynes 282, Inconel 740H, and Inconel 617. Haynes 230 shows minimal differences for the 40–80 year lifetime contours, but some significant differences exist near the alloy's lifetime "valley" at intermediate thermal operating points. The other three alloys produce highly similar contours for each conduction-convection ratio, and the temperature points are shifted by only small magnitudes for 10–80 year contours. In general, a single lifetime damage map derived at  $R = 0.5$  is sufficiently accurate for most applications.

Cycle definition is shown to significantly affect lifetime predictions. Thermal and structural simulations are performed over a range of peak  $T_f, \Delta T$  magnitudes for cycle profiles with temporal variability corresponding to different locations on a cylindrical receiver. For Haynes 230 in a Crescent Dunes scenario, the highest lifetime prediction error along the 30 year contour is at  $T_f = 475^\circ\text{C}, \Delta T = 135^\circ\text{C}$  with

a lifetime that is 43.5% lower than the lifetime obtained assuming a temporal sine approximation. A lifetime safety factor offers a simple method of accounting for uncertainty associated with cycle definition. Ultimately, a modified methodology is provided for development of plant-specific damage models if desired in future research.

The damage model developed in this work is distinct from other available models and is considered high-fidelity for several reasons. This damage model's most unique feature is that it is based on multi-cycle, finite element analyses results instead of analytical models with lumped assumptions. Utilizing FEA is expected to improve modeling accuracy of creep relaxation and, to a lesser degree, yielding effects. Simulation in *SRLIFE* also predicts lower magnitudes of stress relaxation relative to the Norton-Bailey model used in literature. Additionally, many creep-fatigue analyses apply safety factors to provide conservative design margin. This model is not developed with excessive conservatism in mind — instead it aims to estimate creep and fatigue damage as accurately as possible. The modeled cycle is generalizable and not dependent on preheat procedures, which allows it to offer insight on general operation and design. For this reason, the lifetime contours in Fig. 2.12 can be used to calculate maximum allowable receiver flux at any fluid temperature from 275–575 °C.

**Software Tool Contributions and Capabilities** The present work presents several necessary tools, algorithms, and practices that directly enable improvement of receiver reliability. Code is designed to be data base flexible and readily versatile. We use an open-source environment and intend the code serve as a starting point for future research.

Present work has produced over 3700 FEA results by implementing the methodology proposed in Chapter 2 to simulate six different alloys over a meaningful domain of fluid temperatures and temperature differences. Data is available in the form of unprocessed timeseries data, processed cycle data, and estimated lifetime tables. In this Chapter, we present Python code to utilize lifetime tables for design and operation settings. The damage tool class can be used in central receiver analysis to

identify expected low lifetime panels and inform aiming strategies. Extrapolation methods are sophisticated, accounting for corrosion and stress reset regions if desired. Additionally, design and aiming routines can use a modified extrapolation method that heavily penalizes operating points located in the corrosion and stress reset zones. The next chapter utilizes the introduced method for generating design contours from given lifetime tables.

This damage tool is applied to analyze lifetime profiles for a cylindrical receiver. Default aiming in *SolarPILOT* for a cylindrical receiver with Crescent Dunes geometry and 565 MW<sub>th</sub> design power produces flux levels that gradually step down from high flux panels to low flux panels. The lifetime profiles for receivers made of Haynes 230 and Inconel 617 predict 6 and 8 required panel replacements before 30 years of constant operation. In contrast, the panel lifetimes for Haynes 282 and Inconel 740H receivers far exceed 30 years. These results are consistent with the damage maps presented in Chapter 2. Interestingly, the final panel is predicted to fail for all alloys except Haynes 230, which encounters failure farther upstream in the flowpath.

The damage tool prototype for real-time operation presented here offers a solution for variable cycle conditions by proposing a method to track material relaxation. Extended 30 cycle FEA simulation in *SRLIFE* demonstrates the impact of accumulated residual stress on a tube's future stress state. Accurate knowledge of a tube's current stress state allows for state of health tracking by adding the real-time damage to the already accumulated tube damage.

Current work integrates an existing model from Martinek et al. (2021) with *CoPylot*, which is an application programming interface for *SolarPILOT*. Now, users can provide custom desired flux profiles to *CoPylot* and the results feed directly into the thermal model. Additionally, minor modifications made allow the thermal model to simulate flat receivers, solve quickly for high-speed applications, and interpolate independent flux values for adjacent tubes of neighboring panels. Best practices are established for thermal modeling regarding node count and coupling to optical and damage models. Numerous best practices are established via analysis of resolution effects on optical, thermal, and damage models. This work investigates

the coupled impact of a single model's resolution on the other models it is integrated with, which is not frequently discussed in literature. The number of tubes simulated per panel is shown to have a significant impact on lifetime evaluations, and the location of tube failure in each panel is very predictable for the aiming strategy employed in Section 3.5.

**Informed Aiming and Design Study Conclusions** Final work builds on the previous chapters' novel damage model and tools to directly improve panel creep-fatigue reliability and overall system cost.

Present work develops definitions and visualizations of ideal flux profiles that fully utilize receiver area at a given power and lifetime. We create several important procedures for creating flux maps and synthesizing them into flux grids to be used by other modeling scenarios. Testing results show that ideal flux aiming, if achieved, strategically places all thermal operating points on the desired contour, resulting in optimal flux levels with acceptable panel reliability. Minimum required area across a range of 100–400 MW<sub>th</sub> design powers are in the range of 150–285 m<sup>2</sup> for the given receiver parameters.

Two methods are presented as methods of achieving ideal flux aiming in unideal optical scenarios. The first method is computationally fast and achieves noticeable reliability improvements when compared to uniform aiming, but lacks robust handling of flux violations. The second method, although slower, further improves panel reliability by updating the thermal and lifetime profiles before selecting aim-points. This more robust method is shown to achieve fewer panel failures and markedly increase the achievable design power by 10–20 MW<sub>th</sub> without incurring panel failures at less than 30 years. Both methods are significantly higher performing when compared to uniform aiming in terms of achieved panel lifetimes.

The economic and design effects of different aiming assumptions is investigated for a range of height-width aspect ratios. A cost model from literature is integrated with the previously presented optical, thermal, and damage modeling software tools to perform a detailed analysis of levelized cost of heat. Results from 170 and 200 MW<sub>th</sub> design power scenarios show that informed aiming's improvement in

panel lifetimes translates to economic benefits. Additionally, LCOH of informed aiming cases further improves thanks to increased optical efficiency when compared to the baseline uniform aiming case. LCOH for systems using informed *SolarPILOT* aiming are always lower than uniform aiming. Minimum LCOHs are 6.5% and 7.6% lower than uniform aiming cases for 170 and 200 MW<sub>th</sub> scenarios. Interestingly, lower aspect ratios appear to be better suited for informed aiming methods while higher aspect ratios result in fewer panel replacements for uniform aiming methods. However, the analyzed power levels do not challenge the informed aiming method and therefore no significant replacement costs are incurred for informed aiming methods. As a result, aspect ratios close to one yield the lowest LCOH values in these two power scenarios. Future work should consider higher power scenarios to determine if the aspect ratio of minimum LCOH shifts due to more impactful panel replacement costs.

## REFERENCES

---

Achkari, O, and A El Fadar. 2020. Latest developments on tes and csp technologies—energy and environmental issues, applications and research trends. *Applied Thermal Engineering* 167:114806.

Alami, Abdul Hai, Abdul Ghani Olabi, Ayman Mdallal, Ahmed Rezk, Ali Radwan, Shek Mohammad Atiqure Rahman, Sheikh Khaleduzzaman Shah, and Mohammad Ali Abdelkareem. 2023. Concentrating solar power (csp) technologies: Status and analysis. *International Journal of Thermofluids* 18:100340.

American Petroleum Institute & American Society of Mechanical Engineers. API 579-1/ASME FFS-1 Fitness-for-Service. Standard, API & ASME.

American Society of Mechanical Engineers. 2019. Boiler and Pressure Vessel Code — Section VIII Rules for Construction of Pressure Vessels . Standard, ASME.

Barua, B, MC Messner, and MD McMurtrey. 2019. Comparison and assessment of the creep-fatigue and ratcheting design methods for a reference gen3 molten salt concentrated solar power receiver. In *Pressure vessels and piping conference*, vol. 58943, V003T03A100. American Society of Mechanical Engineers.

Barua, Bipul, Michael McMurtrey, Ryann E Rupp, and Mark C Messner. 2020. Design guidance for high temperature concentrating solar power components. Tech. Rep. ANL-20/03, Argonne National Lab.(ANL), Argonne, IL (United States).

Barua, Bipul, and Mark C Messner. 2023. Structural design challenges and implications for high temperature concentrating solar power receivers. *Solar Energy* 251: 119–133.

Berman, I, and MSM Rao. 1983. Sensitivity of estimated tube life to material property variation. *J. Pressure Vessel Technology*.



Bijlaard, Paul P, RJ Dohrmann, and JM Duke. 1968. Thermal stress analysis of nonuniformly heated cylindrical shell and its application to a steam generator membrane wall. *Journal of Eng. Power*.

Binder, Matthias, Christian Schuhbauer, Ralf Uhlig, Peter Schwarzbözl, Ruth Schwaiger, and Robert Pitz-Paal. 2022. Comparison of different safety concepts for evaluation of molten salt receivers. *Solar Energy* 234:119–127.

Boehlert, CJ, and SC Longanbach. 2011. A comparison of the microstructure and creep behavior of cold rolled haynes® 230 alloy™ and haynes® 282 alloy™. *Materials Science and Engineering: A* 528(15):4888–4898.

Boiler, ASME, and Pressure Vessel Code. 2019. Asme section iii, division 5, high temperature reactors. Tech. Rep., American Society of Mechanical Engineers.

Carli, G. 1984. Molten salt solar receiver subsystem research experiment. phase 1, final report. volume 1. technical. Tech. Rep., Foster Wheeler Solar Development Corp., Livingston, NJ (USA).

Center for High Throughput Computing. 2006. Center for high throughput computing.

Cole, Wesley, and Akash Karmaker. 2023. Cost projections for utility-scale battery storage: 2023 update. Tech. Rep., National Renewable Energy Lab.(NREL), Golden, CO (United States).

Conroy, Tim, Maurice N Collins, James Fisher, and Ronan Grimes. 2018a. Levelized cost of electricity evaluation of liquid sodium receiver designs through a thermal performance, mechanical reliability, and pressure drop analysis. *Solar Energy* 166: 472–485.

———. 2018b. Thermal and mechanical analysis of a sodium-cooled solar receiver operating under a novel heliostat aiming point strategy. *Applied energy* 230:590–614.

Conroy, Tim, Maurice N Collins, and Ronan Grimes. 2019. Sodium receiver designs for integration with high temperature power cycles. *Energy* 187:115994.

———. 2020. A review of steady-state thermal and mechanical modelling on tubular solar receivers. *Renewable and Sustainable Energy Reviews* 119:109591.

EIA. 2024a. Electricity explained.

———. 2024b. Short-term energy outlook.

El-Wakil, Mohamed Mohamed. 1984. *Powerplant technology*. McGraw-Hill Companies.

Emerick, Bruno S, Allan R Starke, and Alexandre K da Silva. 2024. Pseudo-transient flow path optimization of a solar tower receiver operating with supercritical carbon dioxide. *Applied Thermal Engineering* 252:123472.

Eno, Daniel R, George A Young, and T-L Sham. 2008. A unified view of engineering creep parameters. In *Asme pressure vessels and piping conference*, vol. 48296, 777–792.

Flores, O, C Marugán-Cruz, D Santana, and M García-Villalba. 2014. Thermal stresses analysis of a circular tube in a central receiver. *Energy Procedia* 49:354–362.

Fork, David K, John Fitch, Shawn Ziaei, and Robert I Jetter. 2012. Life estimation of pressurized-air solar-thermal receiver tubes. *Journal of solar energy engineering* 134(4):041016.

Gatewood, Buford Echols. 1941. Xxiii. thermal stresses in long cylindrical bodies. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 32(213):282–301.

Gentile, Giancarlo, Giovanni Picotti, Marco Binotti, Michael E Cholette, and Giampaolo Manzolini. 2022. Dynamic thermal analysis and creep-fatigue lifetime assessment of solar tower external receivers. *Solar Energy* 247:408–431.

———. 2024. A comprehensive methodology for the design of solar tower external receivers. *Renewable and Sustainable Energy Reviews* 193:114153.

Gentile, Giancarlo, Giovanni Picotti, Marco Binotti, Michael E Cholette, Theodore A Steinberg, and Giampaolo Manzolini. 2023. A methodology for the creep-fatigue analysis of external receivers for solar tower plants using real weather data. In *Aip conference proceedings*, vol. 2815, 100004. AIP Publishing.

Glatzmaier, Greg. 2011. Developing a cost model and methodology to estimate capital costs for thermal energy storage. Tech. Rep., National Renewable Energy Lab.(NREL), Golden, CO (United States).

Glinka, Grzegorz. 1985. Calculation of inelastic notch-tip strain-stress histories under cyclic loading. *Engineering Fracture Mechanics* 22(5):839–854.

González-Gómez, PA, María de los Reyes Rodríguez-Sánchez, Marta Laporte-Azcué, and D Santana. 2021. Calculating molten-salt central-receiver lifetime under creep-fatigue damage. *Solar Energy* 213:180–197.

Goodier, JN. 1957. Thermal stress and deformation. *Journal of Applied Mechanics*.

Grossman, James W, Wendell B Jones, and Paul S Veers. 1990. Evaluation of thermal cycling creep-fatigue damage for a molten salt receiver. Tech. Rep., Sandia National Lab.(SNL-NM), Albuquerque, NM (United States).

Hering, Dominik, Matthias Binder, Peter Schwarzbözl, Ruth Schwaiger, and Robert Pitz-Paal. 2021. Monitoring of service life consumption for tubular solar receivers: Review of contemporary thermomechanical and damage modeling approaches. *Solar Energy* 226:427–445.

Hetnarski, Richard B, M Reza Eslami, and GML Gladwell. 2009. *Thermal stresses: advanced theory and applications*, vol. 41. Springer.

Ho, Cho Yen, and TK Chu. 1977. Electrical resistivity and thermal conductivity of nine selected aisi stainless steels. Tech. Rep., Cindas Report 45.

International, Haynes. 2024a. Haynes 230 alloy data sheet.

———. 2024b. Haynes 282 alloy.

IRENA. 2024.

Kelly, Bruce D. 2010. Advanced thermal storage for central receivers with supercritical coolants. Tech. Rep., Abengoa Solar Inc.

Kistler, BL. 1987. Fatigue analysis of a solar central receiver design using measured weather data. Tech. Rep., Sandia National Lab.(SNL-CA), Livermore, CA (United States).

Laporte-Azcué, Marta, Antonio Acosta-Iborra, Todd P Otanicar, and D Santana. 2023. Real-time estimation of the transient thermomechanical behaviour of solar central receivers. *Thermal Science and Engineering Progress* 41:101834.

Laporte-Azcué, Marta, PA González-Gómez, María de los Reyes Rodríguez-Sánchez, and D Santana. 2020. Deflection and stresses in solar central receivers. *Solar Energy* 195:355–368.

———. 2021a. Material selection for solar central receiver tubes. *Solar Energy Materials and Solar Cells* 231:111317.

———. 2022. A procedure to predict solar receiver damage during transient conditions. *Renewable and Sustainable Energy Reviews* 154:111905.

Laporte-Azcué, Marta, María de los Reyes Rodríguez-Sánchez, PA González-Gómez, and D Santana. 2021b. Assessment of the time resolution used to estimate the central solar receiver lifetime. *Applied Energy* 301:117451.

Liao, Zhirong, Xin Li, Chao Xu, Chun Chang, and Zhifeng Wang. 2014. Allowable flux density on a solar central receiver. *Renewable Energy* 62:747–753.

Logie, William R, John D Pye, and Joe Coventry. 2018. Thermoelastic stress in concentrating solar receiver tubes: A retrospect on stress analysis methodology, and comparison of salt and sodium. *Solar Energy* 160:368–379.

Luo, Y, X Du, L Yang, C Xu, and Y Yang. 2015. Study on the allowable flux density for a solar central dual-receiver. *Energy Procedia* 69:138–147.

Martinek, Janna, Sameer Jape, and Craig S Turchi. 2021. Evaluation of external tubular configurations for a high-temperature chloride molten salt solar receiver operating above 700° c. *Solar Energy* 222:115–128.

McConohy, Geoff, and Alan Kruizenga. 2014. Molten nitrate salts at 600 and 680 c: Thermophysical property changes and corrosion of high-temperature nickel alloys. *Solar Energy* 103:242–252.

MESSNER, MARK, TIANCHEN HU, and US DOE NE-NEAMS. 2023. Neml2 - the new engineering material model library, version 2.

Messner, Mark C., Bipul Barua, and Michael McMurtrey. 2022. srlife: A fast tool for high temperature receiver design and analysis. Tech. Rep. ANL-22/29, Argonne National Lab. (ANL), Argonne, IL (United States).

metals, Special. 2024. A developmental precipitation-hardenable ni-cr-co superalloy for high temperature service in the automotive and power industries.

Metals, Special. 2024a. Inconel alloy 617.

———. 2024b. Inconel alloy 800h.

Miliozzi, A, GM Giannuzzi, P Tarquini, and A La Barbera. 2001. Fluido termovettore: dati di base della miscela di nitrati di sodio e potassio. In *Enea/sol/rd/2001/07*, rev. 0.0, 1. ENEA.

Moftakhar, A, A Buczynski, and G Glinka. 1994. Calculation of elasto-plastic strains and stresses in notches under multiaxial loading. *International journal of fracture* 70:357–373.

Montoya, Andrés, MdIR Rodríguez-Sánchez, Jorge López-Puente, and Domingo Santana. 2020. Influence of longitudinal clips in thermal stresses and deflection in solar tubular receivers. *Solar Energy* 198:224–238.

Montoya, Andrés, MR Rodríguez-Sánchez, Jorge López-Puente, and Domingo Santana. 2019. Thermal stress variation in a solar central receiver during daily operation. In *Aip conference proceedings*, vol. 2126, 030038. AIP Publishing.

Montoya Sancha, Andrés, María de los Reyes Rodríguez Sánchez, Jorge López Puente, and Domingo José Santana Santana. 2018a. Numerical model of solar external receiver tubes: Influence of mechanical boundary conditions and temperature variation in thermoelastic stresses. *Solar Energy*.

Montoya Sancha, Andrés, M.R. Rodriguez-Sanchez, Jorge Puente, and Domingo Santana. 2018b. Thermal and mechanical stresses in a solar central receiver. *Renewable Energy and Power Quality Journal* 1:528–532.

Murphy, Caitlin, Yinong Sun, Wesley J Cole, Galen J Maclaurin, Mark S Mehos, and Craig S Turchi. 2019. The potential role of concentrating solar power within the context of doe's 2030 solar cost targets. Tech. Rep., National Renewable Energy Lab.(NREL), Golden, CO (United States).

Neises, Ty W, Michael J Wagner, and Allison K Gray. 2014. *Structural design considerations for tubular power tower receivers operating at 650 c*, vol. 45868. American Society of Mechanical Engineers.

Nithyanandam, K, and R Pitchumani. 2016. Thermal and structural investigation of tubular supercritical carbon dioxide power tower receivers. *Solar Energy* 135: 374–385.

NREL. 2018. Concentrating solar power.

———. 2023. System advisor model version 2023.12.17.

NREL. 2024. Electricity Definitions. <https://atb.nrel.gov/electricity/2024/about>.

Ölmez, Mehmet Emre, Ibrahim Ari, and Gülfem Tuzkaya. 2024. A comprehensive review of the impacts of energy storage on power markets. *Journal of Energy Storage* 91:111935.

Ortega, Jesus, Sagar Khivsara, Joshua Christian, Clifford Ho, and Pradip Dutta. 2016. Coupled modeling of a directly heated tubular solar receiver for supercritical carbon dioxide brayton cycle: Structural and creep-fatigue evaluation. *Applied Thermal Engineering* 109:979–987.

Ortega, Jesus D, Joshua M Christian, and Clifford K Ho. 2015. Structural analysis of a direct heated tubular solar receiver for supercritical co2 brayton cycle. In *Energy sustainability*, vol. 56840, V001T05A015. American Society of Mechanical Engineers.

Pérez-Álvarez, Rafael, Pedro Ángel González-Gómez, Antonio Acosta-Iborra, and Domingo Santana. 2021. Thermal stress and fatigue damage of central receiver tubes during their preheating. *Applied Thermal Engineering* 195:117115.

Radke, Edward F, David T Wasyluk, David J Dewees, and James M Tanzosh. 2014. Creep-fatigue design applied to molten salt solar receivers. *Pressure Technology* 40740:232–247.

REN21. 2025. Renewables 2025 global status report.

Rodriguez-Sanchez, María de los Reyes, Alberto Sanchez-Gonzalez, Carolina Marugan-Cruz, and D Santana. 2015. Flow patterns of external solar receivers. *Solar Energy* 122:940–953.

Rodríguez-Sánchez, María Reyes, Antonio Soria-Verdugo, José Antonio Almendros-Ibáñez, Antonio Acosta-Iborra, and Domingo Santana. 2014. Thermal design guidelines of solar power towers. *Applied Thermal Engineering* 63(1): 428–438.

Sánchez-González, Alberto, María Reyes Rodríguez-Sánchez, and Domingo Santana. 2020. Allowable solar flux densities for molten-salt receivers: Input to the aiming strategy. *Results in Engineering* 5:100074.

SENER. 2025.

- Sengupta, Manajit, Yu Xie, Anthony Lopez, Aron Habte, Galen Maclaurin, and James Shelby. 2018. The national solar radiation data base (nsrdb). *Renewable and sustainable energy reviews* 89:51–60.
- Siebers, Dennis L, and JS Kraabel. 1984. Estimating convective energy losses from solar central receivers. *United States*.
- Timoshenko, Stephen P, and James Norman Goodier. 1951. *Theory of elasticity*. McGraw-hill New York.
- Toscani, Andrea, Francesco Crespi, David Sánchez, Marco Binotti, and Giampaolo Manzolini. 2018. Assessment of different control strategies to manage cloud-induced transients in central receiver systems using molten salts. In *Aip conference proceedings*, vol. 2033, 040038. AIP Publishing.
- Turchi, Craig S, and Garvin A Heath. 2013. Molten salt power tower cost model for the system advisor model (sam). Tech. Rep., National Renewable Energy Laboratory (NREL), Golden, CO (United States).
- Uhlig, Ralf, Cathy Frantz, Robert Flesch, and Andreas Fritsch. 2018. Stress analysis of external molten salt receiver. In *Aip conference proceedings*, vol. 2033, 040040. AIP Publishing.
- Vant-Hull, Lorin L. 2002. The role of “allowable flux density” in the design and operation of molten-salt solar central receivers. *J. Sol. Energy Eng.* 124(2):165–169.
- Virtanen, Pauli, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17:261–272.



Wagner. 2018. *Solarpilot user's manual*.

Wagner, Michael J, and Tim Wendelin. 2018. Solarpilot: A power tower solar field layout and characterization tool. *Solar Energy* 171:185–196.

Wenner, Jacob, Mark C Messner, and Michael J Wagner. 2025. Damage modeling of power tower receiver tubes using the srlife tool. *Solar Energy* 299:113627.

Zolan, Alexander, William Hamilton, Kashif Liaqat, Michael Wagner, USDOE Office of Energy Efficiency, and Renewable Energy. 2021. Halos (heliostat aimpoint and layout optimization software).

Zurita, Adriana, Carlos Mata-Torres, José M Cardemil, and Rodrigo A Escobar. 2020. Assessment of time resolution impact on the modeling of a hybrid csp-pv plant: A case of study in chile. *Solar Energy* 202:553–570.

## A APPENDIX: STUDY OF POWER BLOCK DYNAMICS IN HYBRID SYSTEM APPLICATIONS

---

The author's current work focuses improving the collection side of CSP — its receiver and solar field. Previous work by the author investigated the generation side of CSP, including overall power block research and component specific simulation. This chapter lists work by the author in the field of dynamic and real-time simulation of CSP power blocks. The overall research outcome was a Power-Hardware-in-the-Loop physical testbed for studying hybrid solar plants composed of photovoltaic and CSP generation. An overview of the resulting publications and presentations is provided here. The major publications are: a published journal paper, a successfully defended Master's Thesis, a published conference paper, and a prepared journal manuscript.

### Superheater Modeling Publication

A dynamic heat exchanger model was developed to analyze dynamic operating changes. A journal article discussing the methodology and results from simulating several mass flow change events has been published in the ASME Journal of Solar Energy Engineering in 2023 titled: *Dynamic Simulation of a Molten Salt-Steam Superheater for Load-Following Applications* by Jacob Wenner and Mike Wagner. The paper is available here: <https://doi.org/10.1115/1.4065768>.

### PHIL Simulation of Hybrid Solar Plants

The author successfully defended a Master's Thesis titled: *A Power-Hardware-In-The-Loop Testbed for Dynamic Analysis and Control of PV-CSP Plants* under the supervision of Assistant Professor Mike Wagner, Professor Mark Anderson, and Professor Bulent Sarlioglu in 2022. The master's thesis document can be accessed here: <https://esolab.engr.wisc.edu/publications/>. The primary outcome was

a physical testbed capable of simulating dynamic CSP power block models in real-time communication with physical power electronics. A dynamic model was created in Simulink capable of capturing and controlling general steam generation transients. Data from real-time tests during cloud events demonstrated the potential synergy of photovoltaic and CSP generation in real time. A followup journal manuscript that extends the plant control development and studies the effects of PV production setpoints on system dynamics is prepared for submission. The proposed title is: *PHiL Evaluation of the Steam-Rankine Dynamic Response During Real-Time Cloud Events in a PV-CSP Hybrid Energy System*, and the authors are Jacob Wenner, Andrew Fernandez, Eric Haag, and Mike Wagner. A copy of the manuscript is available from the author upon request.

Various aspects of this work have been presented in several poster sessions and one oral presentation. Posters have been presented at the Wisconsin Electric Machines and Power Electronics Consortium Annual review in 2022, 2023, and 2024. One additional poster session was accepted for and presented at the SolarPACES 2022 conference. The subsequent conference paper titled "*Development of a PHiL Real-Time Simulation Testbed for Optimization of Hybrid Power Plant Generation*" by Jacob Wenner, Ben Bates, and Mike Wagner is available here: <https://doi.org/10.52825/solarpaces.v1i.639>. A second topic was accepted and selected for oral presentation at the SolarPACES 2024 conference, titled "*Real-Time PHiL Simulation of PV-CSP Plant Operation During Cloud Events*" by Jacob Wenner, Andrew Fernandez, and Mike Wagner.