



KANDIDAT

Dugan William (10034)

PRØVE

FY3900 2 Masteroppgave i fysikk

Emnekode	FY3900
Vurderingsform	Oppgave
Starttid	13.04.2026 10:00
Sluttid	15.05.2026 12:00
Sensurfrist	--
PDF opprettet	16.06.2026 06:59
Opprettet av	Jon Andreas Støvneng

Masteroppgave

NTNU
Norges teknisk-naturvitenskapelige universitet
Det humanvitenskapelige fakultet
Institutt for historiske studier

Student Studentsen

Min oppgave

Masteroppgave i Mitt studieprogram

Veileder: Veileder Veilederen

Januar 2020



i**Levering av masteroppgave ved NTNU**

Gratulerer! Du skal nå levere din masteroppgave.

Forarbeid - lage PDF med NTNU mal

Masteroppgaven skal leveres som en PDF-fil. Før du leverer, anbefaler vi at PDF-en genereres via [NTNU Grafisk senter \(lenke\)](#). Dette gjøres for at oppgaven din får et NTNU-uttrykk. I tillegg blir bildene i oppgaven kvalitetssjekket, slik at de er gode nok dersom oppgaven skal trykkes. [Utfyllende veiledning finner du ved å følge denne lenken.](#)

Når du nå leverer oppgaven din må du også:

- Velge om du vil publisere oppgaven din i Nasjonalt vitenarkiv (NVA), den nasjonale tjenesten som samler publikasjoner og forskningsresultater.
- Bestemme om eventuelle vedlegg skal publiseres i NVA
- Legge inn nødvendig metadata, som brukes i NVA og på vitnemålet ditt

[Utfyllende informasjon, samt brukerveiledninger, finner du ved å følge denne lenken.](#)

Merk: Inneholder oppgaven sensitive opplysninger? I så fall må den ikke publiseres eller leveres i Inspira Assessment. Se «**Fortrolig, eller strengt fortrolig oppgave (oppgave med sensitive opplysninger)**» på nettsiden [Publisering og båndlegging av oppgaver](#)

☒ **Vi anbefaler at du kopierer tittelen fra en word-fil, eller skrive det direkte inn i ruta.**

Kopiering fra PDF fører ofte til en kodefeil i NVA, derfor anbefaler vi at du kopierer fra en word-fil, eller skriver direkte inn i feltene her.

Tittel på oppgaven din

Fyll inn tittelen på oppgaven din. Tittelen skal være på språket oppgaven er skrevet på. Denne tittelen vil automatisk legges på vitnemålet ditt, vil ligge i metadata i NVA, og vil være søkbar der.

Ord: 0

Oppgave på et annet språk enn engelsk? Da skal du også legge inn engelsk tittel her:

Den engelske tittelen vil legges inn i det engelske tillegget til vitnemålet ditt, samt i NVA. Du trenger ikke å skrive noe i dette feltet hvis du har skrevet masteroppgaven din på engelsk.

Ord: 0

Språket masteroppgaven din er skrevet på

- ☐ Norsk - bokmål
- ☐ Norsk - nynorsk
- ☐ Engelsk
- ☐ Svensk
- ☐ Dansk
- ☐ Ingen av disse

Besvart.

**OBLIGATORISK****Sammendrag**

Et sammendrag av oppgaven skal skrives både på norsk, eller et annet skandinavisk språk, og engelsk (ref. [Studieforskriften §5-8](#)). Disse sammendragene skal være like de sammendragene som står i PDF-versjonen av masteroppgaven din.

Vi anbefaler at du kopierer sammendraget fra en word-fil, eller skrive det direkte inn i ruta.

Kopiering fra PDF fører ofte til en kodefeil i NVA, derfor anbefaler vi at du kopierer fra en word-fil, eller skriver direkte inn i feltene her.

Sammendraget vil være tilgjengelig og søkbart i Nasjonalt vitenarkiv (NVA), uavhengig om du samtykker til publisering eller ikke.

Dersom du ber om publisering av oppgaven på et senere tidspunkt ([båndlegging](#)), vil sammendraget ikke publiseres før båndleggingsdato er passert. Istedenfor sammendrag, vil følgende tekst vises: *"Abstract will be available on YYYY-mm-dd"*.

Sammendrag på norsk, eller et annet skandinavisk språk

Ord: 0

Sammendrag på engelsk

Ord: 0

Besvart.

EKSEMPEL

Oppgi navn på forfatter, eller forfattere, av oppgaven: Ditt navn. Dersom du har skrevet oppgaven sammen med andre, skriv inn alle navnene. Informasjonen du legger inn her bli synlig, og søkbar, i NVA. Dette er nyttig, fordi du kan finne igjen oppgaven din. Du kan da vise den frem til framtidige arbeidsgivere, og andre interesserte.

Forfatter(e)s navn skrives på denne formen: Etternavn, Fornavn. Er det flere forfattere, skal disse separeres med linjeskift

Normann, Kari
Doe, Jane Elizabeth
Medel-Svensson, Filip

Oppgi veileders navn. Har du flere veiledere, skal også disse oppgis.

Navn på veileder(e) skrives på formen: Etternavn, Fornavn. Dersom det er flere veiledere, skal disse separeres med linjeskift.

Nilssen, Otto
Modal, Lise Müller

**OBLIGATORISK**

Oppgi navn på forfatter, eller forfattere, av oppgaven: Ditt navn. Dersom du har skrevet oppgaven sammen med andre, skriv inn alle navnene. Informasjonen du legger inn her blir synlig og søkbar i NVA. Dette er nyttig, fordi du kan finne igjen oppgaven din. Du kan da vise den frem til framtidige arbeidsgivere, og andre interesserte.

Forfatter(e)s navn skrives på denne formen: Etternavn, Fornavn. Er det flere forfattere, skal disse separeres med linjeskift

Ord: 0

Oppgi veileders navn. Har du flere veiledere, skal også disse oppgis.

Navn på veileder(e) skrives på formen: Etternavn, Fornavn. Dersom det er flere veiledere, skal disse separeres med linjeskift.

Ord: 0

Besvart.

1

OBLIGATORISK

Leverer masteroppgaven:



Din fil ble lastet opp og lagret i besvarelsen din.

Last ned

Fjern

Erstatt

Filnavn:	thesis.pdf
Filtype:	application/pdf
Filstørrelse:	38.48 MB
Opplastingstidspunkt:	14.05.2026 22:04
Status:	Lagret

Besvart.

2 Lever eventuelle vedlegg:



Last opp filen her. Maks én fil.

Følgende filtyper er tillatt: **.zip** Maksimal filstørrelse er **50 GB**.

 Velg fil for opplasting

Er du usikker på hvordan du lager en ZIP-fil? [Følg denne lenkja til veiledning på Innsida](#).

Ubesvart.

3 KI-deklarasjon

Vennligst last opp ferdig utfylt og signert KI-deklarasjon.

[Du kan laste ned KI-deklarasjonen som skal fylles ut her \(lenke\)](#)



Din fil ble lastet opp og lagret i besvarelsen din.



Last ned



Fjern



Erstatt

Filnavn:

ki_skjema.pdf

Filtype:

application/pdf

Filstørrelse:

118.23 KB

Opplastingstidspunkt:

12.05.2026 20:24

Status:

Lagret

Besvart.

**OBLIGATORISK****Publisering av masteroppgaver**

Når du leverer inn masteroppgaven din må du ta stilling til om oppgaven skal gjøres åpent tilgjengelig, ha utsatt publiseringstidspunkt eller holdes lukket.

NTNU oppfordrer til publisering av alle oppgaver, men valget er opp til deg/dere som forfatter(e).

Åpen publisering

Velger du å gjøre oppgaven din tilgjengelig for allmennheten, kreves godkjenning i henhold til gjeldende publiseringsavtale (versjon 1.0). Oppgaven vil først bli publisert etter sensuren er endelig og alle oppgaver publiseres i [NVA \(lenke\)](#).

Du finner gjeldende publiseringsavtale, samt mer informasjon [ved å følge denne lenken til brukerveiledning om publisering og båndlegging av oppgaver på Innsida](#)

Reservasjon mot publisering/utsatt publisering

Dersom du ikke ønsker å publisere oppgaven, vil kun metadata (tittel, forfatter, veileder) og sammendrag bli gjort tilgjengelig. Ved avtale om utsatt publisering (båndlegging) vil verken metadata, sammendrag eller oppgaven være tilgjengelig før den angitte dato i avtalen.

Dersom du ikke har levert masteravtale med informasjon om utsatt publisering, men ønsker dette likevel, tar du kontakt med instituttet ditt etter at du har levert.

Masteroppgave med flere forfattere

Dersom oppgaven som leveres har flere forfattere, så gjelder publiseringsavtalen for samtlige forfattere. Alle forfatterne MÅ samtykke til publisering. Dersom en av forfatterne i gruppa ikke ønsker publisering, skal svaret på hvorvidt oppgaven skal publiseres være "Nei".

Tillater du publisering av din oppgave?

- ☐ Ja, jeg tillater publisering av masteroppgaven
- ☐ Nei, jeg tillater ikke publisering av masteroppgaven

Publisering av vedlegg

Dersom du leverer vedlegg som en del av masteroppgaven din, kan du velge om du også ønsker å gjøre disse tilgjengelig.

Tillater du publisering av vedlegg?

- ☐ Ja, jeg tillater publisering av vedleggene
- ☐ Nei, jeg tillater ikke publisering av vedleggene
- ☐ Jeg har ikke levert vedlegg

Besvart.

William Dugan

Thermal Analysis of Regeneratively Cooled Additively Manufactured Rocket Engines

Master's thesis in Physics

Supervisor: Jon Andreas Støvneng

Co-supervisor: Even Drugli

May 2026



Norwegian University of
Science and Technology

William Dugan

Thermal Analysis of Regeneratively Cooled Additively Manufactured Rocket Engines

Master's thesis in Physics
Supervisor: Jon Andreas Støvneng
Co-supervisor: Even Drugli
May 2026

Norwegian University of Science and Technology
Faculty of Natural Sciences
Department of Physics



Abstract

This thesis presents and applies a reduced-order model for thermal and hydraulic analysis of regeneratively cooled liquid rocket thrust chambers. The objective was to establish a fast and transparent engineering tool that connects chamber geometry, cooling-channel design, heat transfer, pressure loss, material behaviour, and additive-manufacturing effects, without replacing detailed CFD or conjugate heat-transfer analysis.

Verification is performed against chemical-equilibrium and forced-convection reference cases, followed by benchmarking against a published RL10A-3-3A cooling-jacket model. For the RL10 benchmark, the predicted peak heat flux, throat gas-side wall temperature, coolant temperature rise, coolant pressure drop, and total heat-transfer rate differ from the reference values by -6.8% , -2.2% , $+4.7\%$, -43.9% , and $+13.0\%$, respectively, indicating reasonable agreement for the main thermal quantities while identifying pressure-loss prediction as the largest discrepancy.

The framework is applied to the Propulse NTNU P25 engine using available firing data and hydraulic measurements. The analysis shows that roughness, effective channel constriction, and coolant-side boiling are important for interpreting the measured pressure and temperature response. A design study is also performed for the P26 engine, where a bifurcated cooling-channel layout is analysed and compared with water-flow testing. For this case, the predicted pressure drop agrees with the measured value within $+3.9\%$.

Overall, the results demonstrate that the model captures the main axial trends in heat flux, wall temperature, coolant temperature, and pressure loss across benchmark and engine-relevant cases. The framework provides a practical foundation for cooling-channel design, interpretation of engine-test data, and further development of additively manufactured liquid rocket engines.

Sammendrag

Denne oppgaven presenterer og anvender en lavordensmodell for termisk og hydraulisk analyse av regenerativt kjølte rakettbrennkamre for flytende drivstoff. Målet er å etablere et raskt og etterprøvbart ingeniørverktøy som knytter sammen brennkammergeometri, kjølekanaldesign, varmeoverføring, trykktap, materialoppførsel og effekter knyttet til additiv tilvirkning, uten å erstatte detaljerte CFD-beregninger eller konjugert varmeoverføringsanalyse.

Modellen verifiseres mot referansecaser for kjemisk likevekt og tvungen konveksjon, etterfulgt av en sammenligning med en publisert modell av RL10A-3-3A-kjølesystemet. For RL10-sammenligningen avviker predikert maksimal varmekraft, veggtemperatur på gassiden i halsområdet, temperaturøkning i kjølevæsken, trykktap i kjølekretsen og total varmeoverføringsrate fra referanseverdiene med henholdsvis -6.8% , -2.2% , $+4.7\%$, -43.9% og $+13.0\%$. Dette indikerer rimelig samsvar for de viktigste termiske størrelsene, samtidig som prediksjon av trykktap identifiseres som det største avviket.

Rammeverket anvendes videre på Propulse NTNU P25-motoren ved bruk av tilgjengelige fyringsdata og hydrauliske målinger. Analysen viser at overflateruhet, effektiv kanalinnsvævring og koking på kjølevæskesiden er sentrale for å tolke den målte trykk- og temperaturresponsen. I tillegg gjennomføres en designstudie av P26-motoren, der et avgrenset kjølekanaloppsett analyseres og sammenlignes med vannstrømningstesting. For dette tilfellet stemmer predikert trykktap med målt verdi innen $+3.9\%$.

Samlet viser resultatene at modellen gjengir de viktigste aksiale trendene i varmekraft, veggtemperatur, kjølemediets temperatur og trykktap, både i referansecaser og i anvendelser på relevante motorgeometrier. Rammeverket gir et praktisk grunnlag for kjølekanaldesign, tolkning av motortestdata og videre utvikling av additivt produserte rakettmotorer med flytende drivstoff.

Acknowledgements

I would like to thank my NTNU supervisor, Jon Andreas Støvneng, for his guidance and support during the thesis work.

I am especially grateful to Even Drugli for his close supervision, technical guidance, and continued involvement throughout the project.

I would also like to thank Fieldmade AS for providing engineering data, test prints, engine prints, and access to optical microscopy equipment.

Finally, I would like to thank Propulse NTNU for enabling the engine testing that forms an important part of this work, and the thrust chamber team for providing CAD models, simulation data, and technical discussions related to the P25 and P26 engines.

List of Abbreviations

AM	Additive manufacturing
CAD	Computer-aided design
CEA	NASA Chemical equilibrium with applications
CFD	Computational fluid dynamics
CHF	Critical heat flux
CHT	Conjugate heat transfer
DoE	Design of experiments
HTC	Heat-transfer coefficient
IN625	Inconel 625
SS347	Stainless steel 347
LPBF	Laser powder-bed fusion
O/F	Oxidizer-to-fuel mass flow ratio
TBC	Thermal barrier coating
TEOS	Tetraethyl orthosilicate (chemical)
RL10	Pratt & Whitney RL10 rocket engine
P25	Propulse NTNU P25 rocket engine
P26	Propulse NTNU P26 rocket engine

List of Symbols

Latin symbols

Symbol	Description	Unit
A	Area	m^2
b	Rib thickness	m
c_p	Specific heat capacity	$\text{J}/(\text{kg K})$
D_h	Hydraulic diameter	m
f	Darcy friction factor	–
G	Mass flux	$\text{kg}/(\text{m}^2 \text{ s})$
h	Heat-transfer coefficient	$\text{W}/(\text{m}^2 \text{ K})$
k	Thermal conductivity	$\text{W}/(\text{m K})$
$\dot{m}$	Mass flow rate	kg/s
p	Pressure	Pa
Δp	Pressure drop	Pa
q''	Heat flux	W/m^2
R	Thermal resistance per unit area	$\text{m}^2 \text{ K}/\text{W}$
R	Combustion chamber radius	m
R	Combustion chamber radius of curvature	m
T	Temperature	K
t	Thickness	m
v	Velocity	m/s
w	Channel width	m
x	Axial coordinate	m
I_{sp}	Specific impulse	s
L^*	Characteristic length of combustion chamber	m
R_a	Arithmetic mean profile roughness	m
R_z	Maximum height profile roughness	m
S_a	Arithmetic mean areal roughness	m
S_z	Maximum height areal roughness	m

Greek symbols

Symbol	Description	Unit
η	Fin efficiency	–
ε	Roughness height	m
μ	Dynamic viscosity	Pa s
ρ	Density	kg/m^3
θ	Overhang angle	deg

Dimensionless numbers

Symbol	Description
M	Mach number
Nu	Nusselt number
Pr	Prandtl number
Re	Reynolds number

Subscripts

Subscript	Description
c	Coolant
ch	Chamber
e	Nozzle exit
g	Gas side
i	Inner surface / coating surface
in	Inlet
out	Outlet
sat	Saturation condition
t	Throat
TBC	Thermal barrier coating
w	Wall
wc	Coolant-side wall
wg	Gas-side wall

Table of Contents

List of Abbreviations	iv
List of Symbols	v
List of Figures	ix
List of Tables	x
1 Introduction	1
1.1 Liquid rocket thrust chambers	2
1.1.1 Combustion chambers and nozzles	2
1.1.2 Injectors	2
1.1.3 Cooling methods	3
1.2 Manufacturing and material considerations	4
1.3 Modelling approach	6
1.4 Thesis goals, scope, and structure	7
2 One-dimensional model	8
2.1 Model formulation	8
2.2 Axial marching solution procedure	9
2.3 Chemical equilibrium model	10
2.4 Heat transfer models	11
2.4.1 Gas-side forced convection correlation	11
2.4.2 Coolant-side forced convection correlations	12
2.4.3 Fin efficiency	13
2.4.4 Nucleate boiling heat-transfer coefficient	14
3 Implementation	16
3.1 General software structure	16
3.2 Solver workflow	16
3.3 Software architecture	16
3.4 External libraries and dependencies	18
3.5 Representation of geometry and operating conditions	18
3.6 Gas-state and thermodynamic property evaluation	18
3.7 Evaluation of heat transfer and pressure loss	19
4 Verification and benchmark cases	20

4.1	Chemical equilibrium solver verification	20
4.2	Single-phase forced-convection channel validation	23
4.3	Pratt & Whitney RL10A-3-3A cooling-jacket benchmark	24
4.3.1	Geometry reconstruction	27
4.3.2	Modelling decisions	29
4.3.3	Operating conditions	30
4.3.4	Results	31
4.3.5	Conclusion	34
5	Development of a surface roughness model	35
5.1	A brief introduction to surface roughness	36
5.2	Design of experiment	37
5.3	Surface roughness characterization	38
5.4	Formulation of the roughness model	40
5.5	Comparison with channel flow measurements	42
6	Case study: Propulse NTNU P25	43
6.1	Geometry	46
6.2	Operating conditions	47
6.3	Modelling decisions	49
6.4	Results	50
6.5	Discussion	52
6.5.1	Pressure model	52
6.5.2	Thermal model	54
6.5.3	Observed chamber-wall crack	55
6.6	Conclusion	56
7	Design study: Propulse NTNU P26	58
7.1	Hot-wall geometry	59
7.2	Bifurcated cooling channels	60
7.3	P26 cooling-channel design	62
7.4	Operating conditions	64
7.5	Modelling decisions	65
7.6	Results	66
7.7	Discussion	68
7.8	Conclusion	68

8	Discussion	69
8.1	Model purpose and validation basis	69
8.2	Performance of the reduced-order thermal model	69
8.3	Hydraulic modelling and surface roughness	70
8.4	Engineering lessons from P25 and P26	71
8.5	Main model limitations	72
8.6	Future improvements	74
9	Conclusion	75
	Bibliography	77
	Appendix	80
A	Surface-roughness data	80
B	Inconel 625 properties	81
C	347 Stainless steel properties	82
D	Ethanol–water mixture properties	83
E	Reconstructed RL10 geometry with source data	85
F	Cantera input file for cryogenics at normal boiling point	86
G	SUNDIALS KINSOL residual function	87

List of Figures

1.1	P26 static ground test	1
1.2	Combustion chamber hot-wall geometry parametrization.	2
1.3	Regeneratively cooled thrust-chamber wall cross section	3
1.4	Examples of thermal-barrier-coating configurations used in related rocket-engine studies.	4
1.5	Laser powder-bed fusion process	5
1.6	Definition of the overhang angle used in this work.	6
2.1	One-dimensional thermal-resistance schematic	8
3.1	Flowchart of the axial marching cooling-channel solver	17
4.1	Forced-convection validation at two Reynolds numbers	23
4.2	RL10 test stand assembly	24
4.3	RL10 tubular cooling configuration schematic	25
4.4	RL10 inlet manifold region	26
4.5	RL10 cooling jacket cut-away views	26
4.6	RL10 hot-wall geometry	27
4.7	RL10 per-channel areas	28
4.8	RL10 channel hydraulic diameter	28
4.9	RL10 gas-side wall temperature	31
4.10	RL10 heat flux	32
4.11	RL10 coolant temperature	32
4.12	RL10 coolant pressure	33
5.1	Workflow used to develop and assess the roughness model	35
5.2	Roughness-coupon specimens used for surface roughness characterization.	37
5.3	Cooling-channel test specimen manufactured for hydraulic assessment of the model.	38
5.4	Surface roughness data for 0°, 15°, 36°, 45° overhang angles.	39
5.5	Angle-dependent equivalent sand-grain roughness computed from the measured roughness data using the Adams conversion.	40
6.1	Propulse NTNU P25 engine assembly	43
6.2	P25 static ground test	44
6.3	Launch of the Heimdall flight vehicle	45
6.4	P25 hot-wall geometry	46
6.5	P25 cooling channel flow area	47
6.6	Pressure sensor data from the flight-test of <i>Heimdall</i> , powered by the P25 engine.	48
6.7	P25 flight-test temperature data	48

6.8	P25 coolant pressure and temperature	50
6.9	P25 heat flux	51
6.10	P25 gas-side wall temperature	51
6.11	P25 chamber wall crack	52
6.12	P25 wall superheat	54
7.1	Propulse NTNU P26 combustion chamber	58
7.2	P26 hot-wall geometry	60
7.3	Effect of cooling-channel aspect ratio on hot-wall structural response	60
7.4	Historical examples of non-bifurcated and bifurcated cooling-channel concepts . . .	61
7.5	P26 cooling-channel flow area	63
7.6	Internal cooling channel geometry of the P26 combustion chamber	64
7.7	P26 coolant pressure and temperature	66
7.8	P26 heat flux	67
7.9	P26 gas-side wall temperature	67
A.1	Surface roughness data for different overhang angles, separated by upskin/downskin and wall-thickness.	80
B.1	Inconel 625 annealed material properties	81
C.1	Stainless steel 347 annealed material properties	82
D.1	Thermophysical properties as a function of temperature for a binary mixture of 90 wt% ethanol and 10 wt% water at $p = 2.5$ MPa.	83
D.2	Surface tension of a binary mixture of 90 wt% ethanol and 10 wt% water as a function of saturation pressure.	84
E.1	Reconstructed RL10 cooling-jacket geometry	85

List of Tables

4.1	Constant enthalpy-pressure equilibrium results for methane-air at 1 atm for lean, stoichiometric, and rich mixtures.	21
4.2	Chemical equilibrium results using constant enthalpy-pressure for chamber and con- stant entropy-pressure for throat and nozzle exit, for SSME parameters.	21
4.3	Chemical equilibrium and performance parameters for Case 3.	22
4.4	RL10 hot-wall geometry parameters used for the benchmark.	27
4.5	RL10A-3-3A operating conditions used for the benchmark.	30
4.6	Summary of key RL10 benchmark results compared with the New System Model of Binder et al. [7].	33
5.1	Design of the roughness-coupon specimens. Each specimen corresponds to one over- hang angle and contains three wall-thickness regions.	38
5.2	Comparison of predicted and measured pressure drops for the P25 cooling-channel geometry under different modelling approaches.	42

6.1	P25 hot-wall geometry parameters.	46
6.2	P25 operating and simulation parameters used in the thermal analysis.	47
7.1	Principal P26 hot-wall geometry parameters.	59
7.2	P26 operating and simulation parameters used in the thermal analysis.	65
7.3	Comparison of predicted and measured pressure drop for the P26 cooling-channel water flow-test geometry. The Reynolds number in the throat region is $1.4 \cdot 10^4$. .	66

1 Introduction

A liquid-propellant rocket engine generates thrust by injecting fuel and oxidizer into a thrust chamber, burning the propellants at high pressure, and accelerating the resulting combustion products through a nozzle. The thrust chamber is therefore both the main energy-conversion component and one of the most thermally loaded parts of the engine. Gas temperatures can exceed the allowable temperature of the chamber material by a large margin, so the chamber wall must be cooled while still carrying pressure loads, preserving injector and nozzle alignment, and maintaining acceptable pressure losses in the propellant feed system.

Liquid rocket engines are normally developed through an iterative process. Vehicle-level requirements first define approximate engine requirements such as thrust, chamber pressure, propellant combination, mixture ratio, envelope, and burn duration. These requirements are translated into a conceptual engine design, after which reduced-order models are used to size the chamber, nozzle, injector, cooling system, and feed-system pressure budget. Promising concepts can then be refined using higher-fidelity analysis, manufactured, tested, and redesigned based on the results. The reduced-order model in this thesis fits into this early-to-intermediate design loop, where many cooling-channel concepts and operating assumptions must be evaluated before committing to detailed analysis or hardware.

Propulse NTNU develops liquid rocket engines in which the fuel is routed through internal cooling channels before injection. Figure 1.1 shows a static fire test of the P26 engine and illustrates the type of hardware that motivates this work: a compact regeneratively cooled thrust chamber operating under high heat flux, limited pressure budget, and manufacturing constraints. For this class of engine, the cooling-system design is not a secondary detail. Wall temperature, coolant pressure drop, available injector pressure, manufacturability, and material limits are directly coupled.



Figure 1.1: Static fire test of the P26 engine in a horizontal firing configuration. Photo by Philip Draget, Propulse NTNU.

Regenerative cooling design is difficult because the dominant design drivers are strongly coupled. The largest heat loads are typically found near the throat, where the gas velocity and convective heat-transfer coefficient are high and the chamber wall is geometrically constrained. Narrow cooling channels increase coolant velocity and can improve heat transfer, but they also increase pressure loss and become more sensitive to surface roughness, local blockage, and manufacturing tolerances. Additive manufacturing enables integrated channel geometries that would be difficult to machine conventionally, but it also introduces rough internal surfaces, dimensional uncertainty,

and material-property variations. In addition, local wall-temperature measurements are rarely available during firing tests, which makes thermal validation difficult. A useful engineering model must therefore be simple enough for repeated design studies while still retaining the main thermal and hydraulic mechanisms that determine cooling performance.

1.1 Liquid rocket thrust chambers

The thrust-chamber assembly consists of an injector, combustion chamber, converging-diverging nozzle, and the surrounding structural and cooling features. Propellants are injected into the chamber, where they mix, react, and form high-temperature combustion products. The gas is then accelerated through the nozzle, converting thermal energy into directed kinetic energy. The throat is the minimum-area section of the nozzle and is usually one of the most thermally demanding regions, because high gas-side heat-transfer coefficients coincide with a small wall radius and limited space for cooling features.

1.1.1 Combustion chambers and nozzles

The hot-wall geometry determines both the gas expansion and the available surface through which heat is transferred to the coolant. The chamber radius, throat radius, contraction contour, expansion contour, and axial station spacing therefore enter directly into a one-dimensional regenerative cooling model. Figure 1.2 illustrates the type of hot-wall parametrization used to represent the chamber and nozzle contour in the reduced-order analysis. The figure is not intended as a complete geometric description of a specific engine, but as a compact definition of the main geometric quantities that connect thrust-chamber design to the thermal model.




Figure 1.2: Combustion chamber hot-wall geometry parametrization.

1.1.2 Injectors

The injector controls how the propellants enter the chamber and therefore influences atomization, mixing, combustion efficiency, stability, and the spatial distribution of heat release. Non-uniform injection can also affect the local gas-side wall heat flux and the distribution of film cooling or coating deposition. These effects are important for real thrust chambers, but they are not resolved directly in this thesis. The injector is instead represented indirectly through operating conditions, mixture ratio, chamber pressure, and any assumed gas-side correction such as film-cooling effectiveness or thermal-barrier-coating behaviour. This keeps the model focused on regenerative cooling performance rather than on combustion-device design.

1.1.3 Cooling methods

Rocket thrust chambers can be cooled in several ways. The simplest distinction is between passive or consumable methods, where the wall or a protective layer absorbs heat without actively circulating coolant through the structure, and active methods, where a fluid removes heat during operation. Heat-sink cooling relies on the thermal mass of the chamber material and is mainly suitable for short-duration firings. Ablative cooling uses a material that decomposes, chars, melts, or erodes in a controlled manner, carrying heat away from the wall. Radiation cooling rejects heat from the external surface by thermal radiation and is most useful where the heat flux is moderate, the wall can tolerate high temperature, and the external surface has sufficient view to the surroundings. Film cooling injects a portion of propellant along the hot wall to create a cooler protective layer between the combustion gas and the chamber surface. Transpiration cooling is a related concept in which coolant passes through a porous wall, but this requires a specialised wall structure and is not considered in this thesis.

Regenerative cooling is the main cooling method considered in this work. Figure 1.3 shows a representative cross section of a regeneratively cooled thrust-chamber wall. In a regeneratively cooled engine, one of the propellants is routed through channels in or around the chamber wall before it enters the injector. Both fuel and oxidizer can be used as the coolant, but the fuel is the desired choice in the majority of engines and in the engines considered in this thesis. The coolant absorbs heat from the wall and is simultaneously preheated before combustion. This allows the chamber to survive repeated firings without consuming a protective liner and can also recover part of the heat that would otherwise be lost from the engine.




Figure 1.3: Representative regenerative-cooling layout with axial gas flow, counter-flowing coolant passages, and a chamber-wall cross section.

This coupling creates the central regenerative cooling tradeoff. Smaller channels, larger wetted perimeter, or higher coolant velocity can improve heat transfer and reduce wall temperature, but they also increase pressure loss. Excessive pressure loss reduces injector pressure drop and feed-system margin. Conversely, large channels can reduce pressure loss, but may provide insufficient heat transfer near the throat. The cooling-channel design must therefore provide adequate thermal margin without consuming too much of the available pressure budget.

Film cooling is relevant to this thesis because it changes the gas-side boundary condition seen by the wall. A film-cooled wall may experience a lower effective driving temperature and lower heat flux than an unprotected wall. However, film-cooling performance depends on injector-hole layout, injection angle, momentum ratio, evaporation, combustion, turbulent mixing, and axial film decay. These are inherently three-dimensional processes and are not resolved in the present model. Film cooling is therefore treated as a model limitation or an imposed correction to the gas-side thermal boundary condition rather than as a separate flow simulation.

Thermal barrier coatings (TBCs) reduce the heat flux into the metal wall by adding a low-conductivity layer on the hot-gas side of the chamber. This can lower the metal temperature and increase thermal margin, but it also introduces uncertainty because coating thickness, conductivity, adhesion, erosion resistance, and spatial uniformity are difficult to know precisely. Figure 1.4 illustrates two different coating concepts. The left image shows a comparatively thick, deliberately applied ceramic coating system, while the right image shows an in-situ silicon-dioxide coating pattern that is more representative of the simplified coating assumption used in this thesis. The visible streaking in Fig. 1.4b highlights that an in-situ coating may not form a perfectly uniform annular layer. In the reduced-order model, the coating is therefore represented as a one-dimensional thermal resistance with prescribed thickness and conductivity. This is appropriate for design studies and sensitivity analysis, but it should not be interpreted as a detailed coating-deposition model.



(a) Example of a TBC configuration dissimilar to the one used in this study. A 200 μm to 500 μm layer of ThermoHold[®] is applied by atmospheric plasma spraying over a 50 μm nickel-based bond coat [1]. (b) Example of a TBC configuration closer to the one assumed in this study. A fine layer of silicon dioxide (SiO_2) is formed in situ by combusting a fuel blend containing tetraethyl orthosilicate (TEOS), as demonstrated by Copenhagen Suborbitals [2]. The visible streak patterns result from the discrete injector orifices, rather than from a continuous annular distribution.

Figure 1.4: Examples of thermal-barrier-coating configurations used in related rocket-engine studies.

1.2 Manufacturing and material considerations

The P25 and P26 combustion chambers considered in this thesis are produced in Inconel 625 using laser powder-bed fusion (LPBF). LPBF is an additive manufacturing process in which a component is built layer by layer from metal powder: a thin powder layer is spread across the build plate, selected regions are melted by a laser according to the part geometry, the build plate is lowered, and a new powder layer is spread. Repeating this sequence creates the three-dimensional part, after which unfused powder is removed and functional surfaces may be machined or otherwise post-processed. The basic process is shown schematically in Fig. 1.5. Unlike conventional machining, where material is removed from a larger billet, LPBF creates internal features during the build. This is attractive for liquid rocket thrust chambers because cooling channels, manifolds, thin ribs,

and complex contours can be integrated into a single printed part. The same capability also creates constraints: enclosed channels must allow powder removal, downskin surfaces must be printable without excessive sagging or droop formation, and thin walls must remain sufficiently accurate and leak tight after printing.




Figure 1.5: Schematic illustration of the laser powder-bed fusion process. Reproduced from De Leon et al. [3], licensed under CC BY 4.0.

Internal surface quality is particularly important for regeneratively cooled chambers. As-printed channels can have significant roughness due to partially fused powder particles, layer stair-stepping, melt-pool features, and degraded downskin surfaces [3]. This roughness affects hydraulic performance by increasing frictional pressure loss. It may also influence coolant-side heat transfer, although the present work primarily uses roughness to improve pressure-loss prediction. Since the roughness is not spatially uniform, a single constant roughness value is often a poor representation of additively manufactured channels. In particular, roughness depends strongly on local build orientation and overhang angle. This motivates the spatially varying roughness model developed later in the thesis.

The upskin and downskin terminology used for build-orientation effects is defined in Fig. 1.6. This distinction is important because the lower side of an overhanging feature is generally more difficult to print cleanly than the upper side. In a cooling channel, this means that the hydraulic roughness can vary around the perimeter and along the engine axis. The roughness model used in this thesis does not attempt to resolve every local surface feature, but introduces a practical way to include build-orientation effects in a one-dimensional pressure-loss calculation.

Material behaviour is also relevant. Inconel 625 (IN625) is attractive for additively manufactured thrust chambers because of its high-temperature capability and corrosion resistance, but both thermal and mechanical properties vary with temperature. The thermal conductivity affects the temperature gradient through the chamber wall, while high-temperature strength, residual stress, and cyclic thermal loading influence failure risk. The present thesis includes temperature-dependent thermal properties in the wall-conduction model, using the material data summarized in Appendix B. Stainless steel 347 is also used in RL10 the benchmark case, with the corresponding property data given in Appendix C. Structural and fatigue analysis are not part of the cooling solver, but the predicted wall-temperature distributions provide input that would be required for such assessments.




Figure 1.6: Definition of the overhang angle used in this work. The angle θ is measured between the build direction and the local tangent to the printed surface.

1.3 Modelling approach

High-fidelity approaches could in principle be used for regenerative-cooling analysis. A quasi-two-dimensional wall model or a finite-difference solution through the wall would resolve more of the local temperature field, while a full three-dimensional conjugate heat-transfer simulation would capture effects such as three-dimensional gas flow, channel curvature, secondary flow, non-uniform wall heating, and local geometric details. These methods are valuable for detailed verification of mature designs, but they require substantial geometry preparation, meshing effort, boundary-condition specification, computational cost, and case-specific modelling decisions. They are therefore less suitable as the primary tool for early design iteration, where many channel geometries and operating assumptions may need to be evaluated quickly.

The approach taken in this thesis is to use a quasi-one-dimensional regenerative-cooling model. The chamber and cooling channels are discretized into stations along the coolant flow path. At each station, the local heat flux and wall temperatures are calculated from a thermal-resistance network between the combustion gas and the coolant. The coolant bulk temperature and pressure are then marched from station to station. This formulation prioritizes rapid design iteration and direct inspection of the coupling between chamber geometry, cooling-channel layout, heat-transfer correlations, pressure loss, material properties, and surface roughness.

Heat transfer through the chamber wall is represented as one-dimensional and wall-normal. Axial and circumferential conduction in the wall are therefore neglected. This approximation is appropriate for the present application because the dominant temperature gradients in regeneratively cooled thrust chambers are expected to occur through the wall thickness rather than along the axial direction. The gas-side state is supplied by a chemical-equilibrium thermochemistry model, which provides gas temperature, pressure, composition, and transport properties for the heat-transfer calculation. The coolant side is primarily modelled using internal-flow forced-convection correlations, with a nucleate-boiling correction included when the local wall temperature exceeds the coolant saturation temperature.

Some physical effects are outside the present implementation. Gas-side radiation is not included, so the imposed wall heat load is determined by convective heat transfer alone. Film cooling is also not modelled explicitly, although it could be introduced in a future version through an effective adiabatic wall temperature and, depending on the formulation, an effective gas-side heat-transfer coefficient. Since no film-cooling correction is applied here, the predicted gas-side heat flux and wall temperature correspond to an unfilmed boundary condition. Thermal-barrier coatings are included as prescribed thermal resistances rather than as simulated coating-deposition processes.

1.4 Thesis goals, scope, and structure

The main goal of this thesis is to establish and assess a reduced-order model for regeneratively cooled liquid rocket thrust chambers. The specific objectives are:

1. to implement a modular and extensible quasi-one-dimensional regenerative cooling solver;
2. to verify and benchmark the gas-state, heat-transfer, and coolant-flow submodels against literature and reference cases;
3. to develop and apply a pressure-loss model that accounts for surface roughness in laser powder-bed-fused cooling channels;
4. to apply the model to the Propulse NTNU P25 chamber as the main hardware comparison case;
5. to apply the model to the Propulse NTNU P26 chamber as a design study for a newer bifurcated cooling-channel concept.

The main contribution of this thesis is the integration of established regenerative-cooling correlations, chemical-equilibrium gas-property evaluation, additive-manufacturing surface-roughness characterization, and Propulse-specific thrust-chamber geometries into a transparent engineering analysis framework. The novelty is therefore not the development of a new gas-side or coolant-side heat-transfer correlation, but the combined application of existing models to additively manufactured student rocket engines, including the use of measured surface-roughness data to inform hydraulic modelling.

The scope is limited to steady-state cooling analysis for prescribed geometry and operating conditions. The model predicts axial distributions of heat flux, wall temperature, coolant temperature, and coolant pressure. It does not resolve injector atomization, turbulent combustion, detailed two-phase flow structure, transient startup or shutdown behaviour, structural stress, fatigue, or crack propagation. P25 provides the main comparison with available firing-test data, while P26 is treated primarily as an application and design-evaluation case supported by hydraulic flow-test data.

2 One-dimensional model

Classical thrust-chamber design methods commonly combine quasi-one-dimensional gas properties, empirical gas-side and coolant-side heat-transfer correlations, and one-dimensional coolant energy and pressure balances [4–6]. Similar reduced-order regenerative-cooling models have continued to be used in later system-level and engine-cycle studies, including the RL10 cooling-jacket model of Binder et al. [7] and recent liquid-rocket cooling analyses [8].

The resulting model structure is illustrated in Fig. 2.1. The combustion gas provides the gas-side thermal boundary condition, heat is transferred through the thermal barrier coating and chamber wall, and the coolant bulk state is updated from one axial station to the next.




Figure 2.1: Schematic representation of the thermal-resistance formulation used in the present one-dimensional regenerative-cooling model. The figure illustrates the physical layers, axial coolant update, and wall-normal heat-transfer path at one cooling-channel station.

2.1 Model formulation

At each axial station, the thrust-chamber wall is represented as a one-dimensional thermal-resistance network between the combustion gas and the coolant. The thermal path consists of gas-side convection, optional thermal-barrier-coating conduction, metal-wall conduction, and coolant-side convection. The gas-side boundary condition is expressed using a reference temperature, T_{ref} , which is typically the adiabatic wall temperature supplied by the gas-state model. The coolant is represented by its local bulk temperature, T_c .

The unknown temperature vector at each station is written as

$$\mathbf{x} = [T_i \quad T_{wg} \quad T_{wc} \quad T_c]^T, \quad (1)$$

where T_i is the temperature of the inside surface exposed to the combustion gas, T_{wg} is the gas-side metal wall temperature, T_{wc} is the coolant-side metal wall temperature, and T_c is the coolant bulk temperature after heat addition over the current station. If no thermal barrier coating is included, the coating resistance is set to zero and T_i coincides with the gas-side wall temperature.

For a given estimate of the unknown temperatures, the heat-transfer coefficients and conductive

resistances are evaluated as

$$h_g = h_g(T_i), \quad h_c = h_c(T_{wc}, T_c), \quad (2)$$

and

$$R_{\text{TBC}} = R_{\text{TBC}}(T_i, T_{wg}), \quad R_w = R_w(T_{wg}, T_{wc}), \quad (3)$$

where the conductive resistances are calculated using mean wall temperatures,

$$R_{\text{TBC}} = \frac{t_{\text{TBC}}}{k \left(\frac{T_i + T_{wg}}{2} \right)}, \quad R_w = \frac{t_w}{k \left(\frac{T_{wg} + T_{wc}}{2} \right)}. \quad (4)$$

The local heat flux is then calculated from the total thermal resistance,

$$q'' = \frac{T_{\text{ref}} - T_c}{\frac{1}{h_g} + R_{\text{TBC}} + R_w + \frac{1}{h_c}}. \quad (5)$$

The local solution is obtained by requiring the same heat flux to pass through each part of the thermal network, while also enforcing conservation of energy for the coolant over the axial control volume. This gives the nonlinear residual system

$$\mathbf{F}(\mathbf{x}) = \begin{bmatrix} h_g(T_{\text{ref}} - T_i) - q'' \\ \frac{T_{wg} - T_{wc}}{R_w} - q'' \\ h_c(T_{wc} - T_c) - q'' \\ \dot{m}c_p(T_c - T_{c,\text{in}}) - q''A_c \end{bmatrix} = \mathbf{0}. \quad (6)$$

Here $\dot{m}$ is the coolant mass flow rate through the considered channel or channel group, c_p is the coolant specific heat capacity, $T_{c,\text{in}}$ is the coolant bulk temperature entering the current station, and A_c is the local coolant-wall contact area associated with the axial control volume. The last equation states that the heat transferred through the wall over the current station is absorbed by the coolant. The heat flux is not treated as an independent unknown in the nonlinear system. Instead, it is evaluated directly from the total thermal resistance using Eq. (5) for the current temperature iterate. The residual equations then enforce consistency between the interface temperatures, conductive and convective heat-transfer relations, and the coolant energy balance.

This formulation makes the coupling between heat transfer and coolant state explicit. The heat-transfer coefficients depend on the unknown wall and coolant temperatures, while the coolant temperature depends on the heat flux. The local problem is therefore nonlinear and is solved iteratively at each axial station. In the implementation, this nonlinear system is solved using KINSOL from the SUNDIALS library; the implementation details are described in Section 3.

2.2 Axial marching solution procedure

The cooling-channel solution is obtained by marching along the coolant flow direction. For the engines considered in this work, the coolant enters near the nozzle end and flows toward the injector end. The inlet coolant temperature and pressure are prescribed, and the solver advances from one channel station to the next.

At each station, the local geometry is supplied to the model. This includes the chamber radius or area, cooling-channel hydraulic diameter, wall thickness, contact area, roughness parameters, and any thermal-barrier-coating properties. The gas-state model provides the local reference temperature and gas properties, while the coolant property model evaluates liquid properties from the current coolant state.

The local nonlinear system in Eq. (6) is then solved for the wall temperatures and outlet coolant bulk temperature. Once the station has converged, the outlet coolant temperature becomes the inlet temperature for the next station,

$$T_{c,\text{in}}^{i+1} = T_c^i. \quad (7)$$

The coolant pressure is updated separately using a one-dimensional pressure-loss relation. In the present model, the pressure drop over one station is calculated using the Darcy–Weisbach form,

$$\Delta p_i = f_i \frac{\Delta x_i}{D_{h,i}} \frac{\rho_i u_i^2}{2}, \quad (8)$$

where f_i is the local Darcy friction factor, Δx_i is the station length, $D_{h,i}$ is the hydraulic diameter, ρ_i is the coolant density, and u_i is the mean coolant velocity. The pressure entering the next station is then

$$p_c^{i+1} = p_c^i - \Delta p_i. \quad (9)$$

The marching procedure can be summarized as follows:

Algorithm 1 Axial marching procedure for the one-dimensional cooling-channel model

- 1: Prescribe inlet coolant state $T_{c,\text{in}}^1$, p_c^1 and mass flow rate $\dot{m}$
 - 2: **for** each axial station $i = 1, \dots, N$ **do**
 - 3: Read local geometry: A_c^i , D_h^i , Δx_i , wall thickness, roughness, and TBC properties
 - 4: Evaluate gas-side state and reference temperature T_{ref}^i
 - 5: Evaluate coolant properties from $T_{c,\text{in}}^i$ and p_c^i
 - 6: Solve $\mathbf{F}(\mathbf{x}^i) = \mathbf{0}$ for T_c^i , T_{wg}^i , T_{wc}^i , and T_c^i
 - 7: Store converged local solution quantities and derived properties
 - 8: Compute pressure drop Δp_i
 - 9: Set $T_{c,\text{in}}^{i+1} = T_c^i$
 - 10: Set $p_c^{i+1} = p_c^i - \Delta p_i$
 - 11: **end for**
-

The complete solution is therefore built sequentially. Each station uses the converged coolant state from the previous station, solves the local nonlinear thermal-resistance problem, stores the resulting heat flux and wall temperatures, and updates the coolant temperature and pressure before continuing downstream. The result is a set of axial distributions for heat flux, wall temperature, coolant temperature, coolant pressure, heat-transfer coefficients, Reynolds number, and related derived quantities.

2.3 Chemical equilibrium model

The chamber gas state is calculated as an enthalpy-pressure equilibrium state. The combustor is treated as adiabatic and constant-pressure, and the reactants are assumed to reach chemical equilibrium at the chamber pressure. The equilibrium composition is found by minimizing the Gibbs free energy of the mixture subject to elemental conservation. The corresponding equilibrium temperature is constrained by the enthalpy balance

$$h_r(T_r, p_0, Y_r) = h_p(T_0, p_0, Y_{\text{eq}}), \quad (10)$$

where T_r and Y_r denote the reactant inlet temperature and composition, p_0 is the chamber stagnation pressure, and T_0 and Y_{eq} are the resulting equilibrium product temperature and composition. This is commonly referred to as an enthalpy-pressure, or HP, equilibrium calculation.

After the chamber equilibrium state has been determined, the chamber and nozzle flow field is represented using a quasi-one-dimensional isentropic approximation. At each axial station, the local Mach number is obtained from the area–Mach relation, and the local static pressure is calculated from the chamber stagnation pressure using the isentropic pressure relation. The gas state is then updated at the chamber entropy and the local static pressure.

In the default model used in this thesis, the combustion-gas mixture is re-equilibrated at each axial station using a constant-entropy, constant-pressure equilibrium calculation. The gas composition and transport properties therefore vary along the chamber and nozzle, corresponding to a shifting-equilibrium approximation. The implementation also supports freezing the composition from a

prescribed axial station onward, but this was not used as the default setting for the results presented here.

The local Mach number is obtained from the isentropic area–Mach relation,

$$\frac{A}{A_t} = \frac{1}{M} \left[\frac{2}{\gamma+1} \left(1 + \frac{\gamma-1}{2} M^2 \right) \right]^{\frac{\gamma+1}{2(\gamma-1)}}. \quad (11)$$

The subsonic solution is used upstream of the throat, while the supersonic solution is used downstream of the throat.

The local static gas temperature and pressure are then calculated from the isentropic relations

$$\frac{T_0}{T_g} = 1 + \frac{\gamma-1}{2} M^2 \quad (12)$$

and

$$\frac{p_0}{p_g} = \left(1 + \frac{\gamma-1}{2} M^2 \right)^{\gamma/(\gamma-1)}. \quad (13)$$

The local gas state is therefore determined by the chamber equilibrium state, the local area ratio, and the assumed value of γ .

Although the composition is frozen during the nozzle expansion, the gas properties used in the heat-transfer model are evaluated at the local static temperature and pressure. This allows the viscosity, thermal conductivity, specific heat, and Prandtl number to vary along the chamber and nozzle. These properties are then used in the gas-side heat transfer correlation described in Section 2.4.1.

2.4 Heat transfer models

2.4.1 Gas-side forced convection correlation

The gas-side convective heat-transfer coefficient is evaluated using the semi-empirical Bartz correlation [5]. The correlation relates the local gas-side heat-transfer coefficient to the chamber pressure, characteristic velocity, local nozzle geometry, gas properties, and Mach number, and is widely used for preliminary thermal analysis of rocket thrust chambers. In the present model, it provides a practical closure for the gas-side boundary condition using quantities obtained from the quasi-one-dimensional gas-state calculation. The corresponding gas-side heat flux is written as

$$q_g'' = h_g (T_{aw} - T_s), \quad (14)$$

where T_{aw} is the adiabatic wall temperature and T_s is the temperature of the surface exposed to the combustion gas. If no thermal barrier coating is included, T_s is the hot-gas-side metal wall temperature. If a thermal barrier coating is included, T_s is the gas-side thermal barrier coating temperature.

The adiabatic wall temperature is the temperature the gas-side surface would approach if the wall were adiabatic, such that the convective heat flux at the wall vanished. It is higher than the local static gas temperature because part of the kinetic energy in the boundary layer is recovered near the wall. In the present model, it is approximated using the recovery-temperature relation [9]

$$\frac{T_{aw}}{T_g} = 1 + \text{Pr}^{1/3} \frac{\gamma-1}{2} M^2 \quad (15)$$

where T_g is the local static gas temperature, M is the local Mach number and γ is the ratio of specific heats. The use of T_{aw} rather than the chamber stagnation temperature gives a gas-side driving temperature that is consistent with the local quasi-one-dimensional gas state.

The Bartz correlation is used in the form

$$h_g = \left[\frac{C_B}{D_t^{0.2}} \left(\frac{\mu^{0.2} c_p}{\text{Pr}^{0.6}} \right)_0 \left(\frac{P_0}{c^*} \right)^{0.8} \right] \left(\frac{A_t}{A} \right)^{0.9} \sigma. \quad (16)$$

Here, D_t is the throat diameter, A_t is the throat area, A is the local chamber or nozzle area, P_0 is the chamber stagnation pressure, and c^* is the characteristic velocity. The gas properties inside the square brackets are evaluated at the stagnation state. This term is therefore constant for a given operating point, while the local variation in h_g comes from the area ratio and the correction factor σ .

The coefficient C_B is the empirical Bartz coefficient. Bartz gives the commonly used value $C_B = 0.026$, but also notes that the coefficient may require adjustment for a given engine or test configuration [5]. In this thesis, any scalar correction to the gas-side heat transfer coefficient is applied by changing C_B directly. For example, reducing the Bartz prediction by 25% is equivalent to using $C_B = 0.0195$ instead of $C_B = 0.026$. This avoids introducing a separate correction factor in the default Bartz model.

The correction factor σ accounts for property variation through the gas-side boundary layer and is given by

$$\sigma = \left[\frac{1}{2} \frac{T_s}{T_0} \left(1 + \frac{\gamma - 1}{2} M^2 \right) + \frac{1}{2} \right]^{-0.68} \left[1 + \frac{\gamma - 1}{2} M^2 \right]^{-0.12}. \quad (17)$$

In this expression, T_0 is the stagnation temperature, M is the local Mach number, and γ is the ratio of specific heats. The use of T_s in Eq. (17) is important when a TBC is present, since the combustion gas is then in contact with the coating surface rather than the metal wall.

Some forms of the Bartz equation also include a weak throat-radius correction term which is not included in the default implementation used here, since the exponent is small and the gas-side model is already empirical. Any scalar correction to the gas-side heat transfer level is instead applied through the empirical coefficient C_B .

2.4.2 Coolant-side forced convection correlations

When the coolant is single-phase, the liquid-side heat transfer coefficient is calculated from internal forced-convection correlations. The heat transfer coefficient is related to the Nusselt number by

$$h_l = \frac{\text{Nu} k_l}{D_h}, \quad (18)$$

where k_l is the coolant thermal conductivity and D_h is the hydraulic diameter of the cooling channel. The Reynolds and Prandtl numbers are defined as

$$\text{Re}_{D_h} = \frac{\rho_l u_l D_h}{\mu_l}, \quad \text{Pr} = \frac{c_{p,l} \mu_l}{k_l}. \quad (19)$$

Unless otherwise stated, the coolant properties are evaluated at the bulk liquid temperature.

A standard correlation for turbulent flow in smooth tubes is the Dittus–Boelter correlation [10],

$$\text{Nu} = 0.023 \text{Re}_{D_h}^{0.8} \text{Pr}^n. \quad (20)$$

The exponent is normally taken as $n = 0.4$ when the fluid is being heated and $n = 0.3$ when the fluid is being cooled [10]. Since the coolant absorbs heat from the chamber wall, $n = 0.4$ is used here. The correlation is mainly intended for fully developed turbulent flow in smooth tubes, and is commonly applied for $\text{Re}_{D_h} > 10^4$ and $0.7 \leq \text{Pr} \leq 160$ [10]. It is therefore a simple baseline model, but it does not include any direct correction for wall roughness. When the Reynolds number is expressed in terms of coolant mass flow rate rather than velocity, and the fluid properties are evaluated at a film temperature defined as the average of the bulk coolant and wall temperatures, this formulation is often referred to as a Colburn-type correlation in rocket cooling-jacket modelling [7].

If the wall-to-bulk temperature difference is large enough that the viscosity changes significantly across the boundary layer, the Sieder–Tate correlation may be used instead,

$$\text{Nu} = 0.027 \text{Re}_{D_h}^{0.8} \text{Pr}^{1/3} \left(\frac{\mu_l}{\mu_w} \right)^{0.14}, \quad (21)$$

where μ_l is evaluated at the bulk coolant temperature and μ_w is evaluated at the liquid-side wall temperature [10]. The correction accounts only for the effect of temperature-dependent viscosity within a single-phase liquid boundary layer. It does not account for nucleate boiling, bubble formation, two-phase flow, or any enhancement of heat transfer caused by phase change. Consequently, eq. (21) is only applied while the coolant-side wall temperature remains below the local saturation temperature of the coolant at the channel pressure. If the predicted wall temperature exceeds this limit, the liquid-side heat transfer can no longer be treated as purely non-boiling forced convection, and a boiling heat transfer model is required instead. The same limitation applies to the other non-boiling correlations in this section.

A more general turbulent-flow correlation is the Gnielinski correlation [10],

$$\text{Nu} = \frac{(f/8)(\text{Re}_{D_h} - 1000)\text{Pr}}{1 + 12.7(f/8)^{1/2}(\text{Pr}^{2/3} - 1)}, \quad (22)$$

where f is the Darcy friction factor. The correlation is commonly used for turbulent internal flow and has a wider range of applicability than the Dittus–Boelter correlation, especially close to the lower end of the turbulent regime [10]. For smooth tubes, the friction factor can be estimated from the first Petukhov equation

$$f = (0.790 \ln \text{Re}_{D_h} - 1.64)^{-2}. \quad (23)$$

For rough channels, the friction factor can instead be found from a rough-wall relation such as the Colebrook equation,

$$\frac{1}{\sqrt{f}} = -2 \log_{10} \left[\frac{\varepsilon/D_h}{3.7} + \frac{2.51}{\text{Re}_{D_h} \sqrt{f}} \right], \quad (24)$$

where ε is the equivalent sand-grain roughness [10]. This is useful for additively manufactured cooling channels, since the same friction factor can be used both in the pressure-loss model and in the Gnielinski heat transfer correlation. For this reason, the Gnielinski correlation is used as the preferred single-phase forced-convection model in this thesis.

2.4.3 Fin efficiency

The coolant-side heat-transfer coefficient gives the heat flux per unit coolant-side surface area, but the one-dimensional thermal-resistance model does not resolve the two-dimensional temperature field in the ribs between adjacent cooling channels. The effective coolant-side contact area A_c in Eq. 6 therefore depends on how much of the channel surface is assumed to participate in heat transfer to the coolant. Using only the hot-wall channel width gives a lower-limit area, while including the channel side walls accounts approximately for lateral heat conduction into the ribs.

To account for partial heat removal through the channel side walls, the ribs are treated as straight fins of height H and thickness b . For a fin with an adiabatic tip and uniform heat-transfer coefficient, the fin efficiency is written as [10]

$$\eta = \frac{\tanh(mH)}{mH}, \quad m = \sqrt{\frac{2h_c}{k_w b}}, \quad (25)$$

where h_c is the coolant-side heat-transfer coefficient, k_w is the wall thermal conductivity, H is the channel height, and b is the rib thickness between adjacent channels. The factor of two in the expression for m accounts for heat transfer from both sides of the rib to the coolant.

For one control volume of length Δx , the effective coolant-side contact area is then approximated as

$$A_c = (w + 2\eta H) \Delta x, \quad (26)$$

where w is the hot-wall channel width exposed to the coolant. The first term represents heat transfer through the hot-wall channel surface, while the second term represents the effective contribution from the two channel side walls. If fin effects are neglected, only the hot-wall channel width is included. If the side walls are nearly isothermal with the hot wall, η approaches unity and the full side-wall contribution is recovered.

2.4.4 Nucleate boiling heat-transfer coefficient

When the coolant-side wall temperature exceeds the local saturation temperature, vapour formation at the wall can increase the coolant-side heat-transfer coefficient. This effect is included using a Chen-type flow-boiling superposition model, where the heat transfer is written as the sum of a forced-convection contribution and a nucleate-boiling contribution [11]. Similar Chen-type formulations have also been applied in one-dimensional regenerative-cooling models for student LOX/ethanol rocket engines [12].

The boiling model is used as an engineering correction to the coolant-side heat-transfer coefficient rather than as a resolved two-phase-flow model. The model is limited to nucleate boiling; critical heat flux prediction, dryout, and post-critical-heat-flux regimes are not included. Vapour slip, void fraction, bubble dynamics, dryout, and detailed onset of nucleate boiling are not modelled. The coolant flow is still represented by a one-dimensional bulk state, while saturated liquid and vapour properties are evaluated at the local coolant pressure.

In the Chen model, the forced-convection contribution is modified by a convective enhancement factor F , while the nucleate-boiling contribution is modified by a suppression factor S . Written as an equivalent coolant-side heat-transfer coefficient, the model becomes

$$h_l = F h_{fc} + S \frac{T_{wc} - T_{sat}}{T_{wc} - T_b} h_{nb}, \quad (27)$$

where h_{fc} is the single-phase forced-convection coefficient, h_{nb} is the nucleate-boiling coefficient, T_{wc} is the coolant-side wall temperature, T_{sat} is the saturation temperature at the local coolant pressure, and T_b is the bulk coolant temperature. The equivalent-coefficient form follows from applying the forced-convection term to $T_{wc} - T_b$ and the nucleate-boiling term to $T_{wc} - T_{sat}$ [12].

The enhancement factor F accounts for the increase in convective heat transfer caused by two-phase flow. It is evaluated from the reciprocal Martinelli parameter, X_{tt}^{-1} , which depends on the local vapour quality and the saturated liquid and vapour properties. The local equilibrium vapour quality is first estimated from the bulk coolant enthalpy as

$$x = \frac{h_b - h_{f,sat}}{h_{fg}}, \quad (28)$$

where h_b is the bulk coolant enthalpy, $h_{f,sat}$ is the saturated-liquid enthalpy at the local pressure, and h_{fg} is the latent heat of vaporisation.

The reciprocal Martinelli parameter is evaluated as [12, 13]

$$X_{tt}^{-1} = \left(\frac{x}{1-x} \right)^{0.9} \left(\frac{\rho_l}{\rho_v} \right)^{0.5} \left(\frac{\mu_v}{\mu_l} \right)^{0.1}, \quad (29)$$

where ρ_l , ρ_v , μ_l , and μ_v are the saturated liquid and vapour densities and dynamic viscosities, respectively.

Chen originally introduced the enhancement and suppression factors as part of the flow-boiling superposition model [11]. In the present work, the analytical expressions for these factors follow the formulation summarized by Chen and Fang [13]. The enhancement factor is evaluated as

$$F = \begin{cases} 1, & X_{tt}^{-1} \leq 0.1, \\ 2.35 (X_{tt}^{-1} + 0.213)^{0.736}, & X_{tt}^{-1} > 0.1, \end{cases} \quad (30)$$

and the suppression factor is evaluated as

$$S = \frac{1}{1 + 2.53 \times 10^{-6} (Re_l F^{1.25})^{1.17}}. \quad (31)$$

The Reynolds number used in the suppression factor is evaluated as

$$Re_l = \frac{(1-x)GD_h}{\mu_l}, \quad (32)$$

where G is the total coolant mass flux and D_h is the coolant-channel hydraulic diameter.

For subcooled liquid, x is negative. Since the Martinelli parameter used in the Chen enhancement factor is defined for non-negative vapour quality, the value used in the two-phase correction is clipped to the interval $0 \leq x < 1$. This means that the enhancement factor remains $F = 1$ while the bulk coolant is subcooled. Local wall boiling may still contribute through the nucleate-boiling term when $T_{wc} > T_{sat}$, but no additional bulk two-phase convective enhancement is introduced in the subcooled region.

The nucleate-boiling coefficient h_{nb} is evaluated using the Forster–Zuber correlation [14],

$$h_{nb} = 0.00122 \frac{\Delta T_e^{0.24} k_l^{0.79} c_{p,l}^{0.45} \rho_l^{0.49} \Delta p_e^{0.75}}{h_{fg}^{0.24} \sigma^{0.5} \rho_v^{0.24} \mu_l^{0.29}}, \quad (33)$$

where $\Delta T_e = T_{wc} - T_{sat}$ is the wall superheat, Δp_e is the saturation-pressure difference associated with the wall superheat, k_l is the saturated-liquid thermal conductivity, $c_{p,l}$ is the saturated-liquid specific heat capacity, and σ is the liquid–vapour surface tension.

In the present implementation, boiling is activated only when the coolant-side wall temperature exceeds the local saturation temperature by more than 2 K. Below this threshold, the single-phase forced-convection coefficient is used. This avoids numerical activation of the boiling model for very small wall superheats. A detailed onset-of-nucleate-boiling model is not included.

The Forster–Zuber correlation requires the liquid–vapour surface tension at the local saturated coolant state. For pure fluids, this property is obtained directly from the thermophysical-property backend. For mixtures, direct surface-tension evaluation is not generally available for all property backends, and a parachor-based estimate is therefore used.

The parachor method relates surface tension to the difference between the coexisting liquid and vapour densities. Hugill and Van Welsenes present the classical Macleod–Sugden parachor relation and describe the Weinaug–Katz method as a widely used extension for mixture surface-tension prediction [15]. For a pure component, the parachor is evaluated as

$$P_i = \frac{\sigma_i^{1/m}}{\rho_{\text{mol},l,i} - \rho_{\text{mol},v,i}}, \quad (34)$$

where σ_i is the pure-component surface tension, $\rho_{\text{mol},l,i}$ and $\rho_{\text{mol},v,i}$ are the saturated liquid and vapour molar densities, and $m = 4$ is the classical parachor exponent.

For mixtures, Hugill and Van Welsenes write the Weinaug–Katz method in terms of liquid- and vapour-phase mixture parachors, and further show that the mixing rule can be generalized by introducing binary interaction parameters [15]. Following this notation, the liquid- and vapour-phase mixture parachors are evaluated as

$$P_l = \sum_i \sum_j z_{l,i} z_{l,j} \lambda_{ij} \frac{P_i + P_j}{2}, \quad (35)$$

and

$$P_v = \sum_i \sum_j z_{v,i} z_{v,j} \lambda_{ij} \frac{P_i + P_j}{2}, \quad (36)$$

where $z_{l,i}$ and $z_{v,i}$ are the liquid- and vapour-phase mole fractions of component i , respectively, and λ_{ij} is an interaction factor.

The mixture surface tension is then calculated from

$$\sigma_{\text{mix}} = (P_l \rho_{\text{mol},l} - P_v \rho_{\text{mol},v})^m, \quad (37)$$

where $\rho_{\text{mol},l}$ and $\rho_{\text{mol},v}$ are the saturated liquid and vapour molar densities of the mixture.

In the present work, no mixture-specific surface-tension data were available for fitting binary interaction parameters. The interaction factor was therefore set to its default value, $\lambda_{ij} = 1$, which reduces the generalized mixing rule to the original Weinaug–Katz form. In the code this is implemented through $\lambda_{ij} = 1 - \delta_{ij}$, with $\delta_{ij} = 0$ by default. The resulting surface-tension estimate is used only in the Forster–Zuber nucleate-boiling contribution.

3 Implementation

The numerical model described in Section 2 was implemented as a modular C++ code for one-dimensional regenerative cooling analysis. The implementation follows the axial marching procedure introduced in Algorithm 1. The coolant state is advanced from one station to the next, while each station solves the local non-linear thermal-resistance problem using the gas-side state, coolant properties, heat-transfer correlations, and pressure-loss model.

3.1 General software structure

The code is organized around a cooling-channel solver that operates on a prescribed sequence of channel stations, while the physical sub-models are supplied through separate interfaces. The ordering of the stations is defined by the cooling-channel representation itself. This is important for geometries where the coolant path does not follow a simple monotonic axial coordinate, such as multi-pass channels or bifurcated channel layouts. The solver therefore iterates over the station sequence it is given, solves the local non-linear thermal problem at each station, updates the coolant state, and stores the resulting station quantities.

This separation is deliberate. The solver does not hard-code a specific engine geometry, coolant path, propellant combination, heat-transfer correlation, or friction model. These quantities are instead provided by case definitions, channel objects, and model interfaces. The same numerical machinery can therefore be reused for different channel concepts and operating conditions without changing the core station-solve procedure.

Individual simulation cases are defined as small executable programs. A case program specifies the thrust-chamber contour, cooling-channel geometry, material properties, coolant inlet state, chamber operating point, and selected heat-transfer and pressure-loss models. The same solver can therefore be used for the verification cases in Section 4, the roughness-model comparisons in Chapter 5, and the P25 and P26 engine studies in Chapters 6 and 7.

The solver produces axial distributions of heat flux, wall temperatures, coolant temperature, coolant pressure, heat-transfer coefficients, Reynolds number, friction factor, and related derived quantities. These outputs are written to file and processed separately to generate the figures and tables used in the later chapters. The implementation is therefore intended primarily as an analysis library rather than as a stand-alone graphical design tool.

3.2 Solver workflow

The solver workflow is summarized in Fig. 3.1. The calculation consists of an outer axial march and an inner station solve. The outer loop advances the coolant bulk state through the prescribed channel-station sequence, while the inner solve enforces local consistency between wall temperature, heat flux, heat-transfer coefficients, and pressure loss. This structure separates the global coolant-state propagation from the local non-linear thermal problem.

3.3 Software architecture

The solver implementation follows a modular architecture in which thermodynamic property evaluation, heat-transfer models, pressure-loss models, geometry definitions, and numerical solution procedures are separated into independent components. The simulation case provides the geometry and operating conditions. The gas-state model provides the combustion-gas reference temperature and transport properties. The coolant-property model evaluates the local liquid state from temperature and pressure. The heat-transfer and friction models then use these states together with the local geometry to compute the quantities required by the station solve. The cooling-channel solver coordinates these components during the axial march.




Figure 3.1: Flowchart of the axial marching cooling-channel solver. The outer loop (blue) advances the coolant state from one axial station to the next. The local station solve (orange, dashed panel) evaluates the temperature-dependent heat-transfer coefficients and conductive resistances, solves the nonlinear thermal-resistance residual system, and returns the converged wall temperatures, heat flux, and heat-transfer coefficients. The pressure loss is then evaluated separately before the updated coolant state is passed to the next station.

3.4 External libraries and dependencies

The implementation was written in C++ and interfaces directly with the C or C++ APIs provided by Cantera, CoolProp, and SUNDIALS. This allowed thermodynamic-property evaluation and nonlinear-solver routines to be integrated in the same compiled codebase as the cooling-channel solver.

Cantera is used for chemical-equilibrium calculations and combustion-gas property evaluation. The chamber gas state is calculated as an enthalpy-pressure equilibrium state, and the resulting composition is used by the gas-state model to evaluate thermodynamic and transport properties along the chamber and nozzle. Cantera is accessed through its C++ interface.

Coolant properties are evaluated using CoolProp. The CoolProp interface is accessed through a thin C++ wrapper that provides a convenient representation of coolant states and collects the property calls required by the heat-transfer models. This layer also allows model-specific extensions to be added where the underlying property backend does not provide all required quantities directly. In particular, the parachor-based surface-tension estimate described in Section 2.4 is implemented in this wrapper and is used by the nucleate-boiling model when direct surface-tension data are unavailable.

To improve computational efficiency, the wrapper includes caching of coolant-state calculations, since the nonlinear station solve may request the same or nearly identical thermodynamic states repeatedly.

The nonlinear station problem at each station is solved using KINSOL from the SUNDIALS library, which solves the residual system defined in Eq. (6). An example implementation of this residual function is provided in Listing 2.

3.5 Representation of geometry and operating conditions

The solver operates on a discretized representation of the combustion chamber and cooling channel geometry. The axial coordinate is divided into a sequence of stations that correspond to control volumes along the chamber and nozzle wall. Each station contains the geometric parameters required for evaluation of the thermal system and pressure losses.

The most important geometric quantities are the local chamber cross-sectional area, the hydraulic diameter of the cooling channel, and the contact area between the coolant and the wall. The axial discretization is typically chosen such that the geometry changes smoothly between adjacent stations, allowing the quasi-one-dimensional assumptions of the model to remain valid.

For regenerative cooling channels produced by additive manufacturing, the local channel orientation relative to the build direction may vary along the chamber contour. This variation affects the surface roughness of the channel walls. To accommodate this effect, the solver stores the local overhang angle of each channel wall. These angles are later used to evaluate the roughness model developed in Chapter 5.

Operating conditions are provided as a set of global parameters describing the engine and coolant system. These include the chamber stagnation pressure, mixture ratio, mass flow rate, coolant inlet temperature and pressure, and geometric parameters such as throat diameter and nozzle expansion ratio. From these quantities the solver determines the local flow state along the chamber and cooling channel.

3.6 Gas-state and thermodynamic property evaluation

The gas-state model provides the combustion-gas quantities required by the heat-transfer calculation. The chamber state is first calculated using Cantera [16] as a constant-enthalpy, constant-pressure equilibrium state. This calculation gives the chamber stagnation temperature, equilibrium

composition, and stagnation transport properties for the specified propellant combination, mixture ratio, chamber pressure, and reactant inlet conditions.

Cantera provides the thermodynamic and transport-property evaluation routines, while the species data are supplied through an external YAML mechanism file. In this work, the San Diego combustion mechanism is used as the Cantera source file for the gas-phase equilibrium and transport-property calculations [17]. The YAML file specifies the species included in the mixture, the thermodynamic polynomial data used to evaluate enthalpy, entropy, and heat capacity, the reaction set available to the equilibrium solver, and the transport data used for viscosity, thermal conductivity, and Prandtl number evaluation. This choice is part of the numerical model rather than only a software detail, because changing the mechanism file can change the equilibrium composition, flame temperature, molecular weight, and gas transport properties, and therefore also the gas-side heat-transfer coefficient.

The nozzle and chamber flow field is then represented by a quasi-one-dimensional gas-state provider. For each axial station, the local area ratio is used to determine the Mach number from the area–Mach relation. The corresponding static temperature and pressure are then calculated from the isentropic relations. In the default implementation, the combustion-gas mixture is re-equilibrated locally at the station pressure and temperature. The gas composition and transport properties therefore vary along the chamber and nozzle. This corresponds to a shifting-equilibrium approximation, rather than a frozen-composition expansion. A frozen-composition gas model may also be used within the same interface, but was not used as the default model in this work.

The equilibrium calculation is performed once for a given operating point, while the gas-state provider supplies station-wise quantities during the axial march. The cooling solver therefore only depends on the gas-state interface and does not need to know whether the supplied gas properties come from Cantera, precomputed tables, or another gas model. The later verification against CEA and RPA in Section 4.1 should therefore be interpreted as a check of the complete Cantera-based gas-state setup used in this thesis, including both the implementation and the selected source file.

3.7 Evaluation of heat transfer and pressure loss

At each station, the heat-transfer model receives the local geometry, gas-side state, coolant state, and material data. The gas-side heat-transfer coefficient is evaluated using the selected gas-side correlation, which is the Bartz correlation for the default model. The coolant-side heat-transfer coefficient is evaluated from the selected liquid-side model, using local coolant properties and channel geometry. For single-phase calculations this is typically based on the Gnielinski correlation, while cases including wall boiling use the nucleate-boiling correction described in Section 2.4.

The station solve then evaluates the thermal-resistance network introduced in Chapter 2. For each non-linear temperature iterate, the heat-transfer coefficients and temperature-dependent conductive resistances are updated, and the heat flux is calculated from the total resistance between the gas-side reference temperature and the coolant bulk temperature. The residual equations enforce consistency between this heat flux, the wall-interface temperatures, and the coolant energy balance over the control volume. After convergence, the station output is stored before the coolant state is advanced to the next channel station.

The pressure loss is evaluated station-wise using the same local coolant state and geometric quantities as the coolant-side heat-transfer model. The pressure-loss model receives the local hydraulic diameter, flow area, surface roughness, coolant properties, and mass flow rate, and returns the pressure drop over the current control volume. The coolant pressure is then updated together with the coolant temperature before the solver proceeds to the next station. This coupling is one-way within each axial step: the local pressure affects coolant properties and saturation conditions, while the computed pressure drop determines the inlet pressure for the following station.

4 Verification and benchmark cases

This chapter assesses the implemented model before it is applied to the P25 and P26 cooling-channel cases. The verification and benchmark cases are chosen to test the main parts of the modelling chain: the thermochemical gas-state calculation, the coolant-side forced-convection implementation, and the coupled one-dimensional cooling-channel solver.

The thermochemical calculations are compared with established equilibrium and rocket-performance tools, including *NASA Chemical Equilibrium with Applications* (CEA) [18, 19] and *Rocket Propulsion Analysis* (RPA) [20, 21]. The coolant-side heat-transfer implementation is checked against single-phase forced-convection measurements, and the full cooling-channel model is benchmarked against a published one-dimensional RL10A-3-3A cooling-jacket model. Together, these cases provide a practical validation basis for the range of calculations performed later in the thesis.

The nucleate-boiling model described in Section 2.4.4 is not assessed separately in this chapter. The available validation cases focus on equilibrium chemistry, single-phase forced convection, and full-channel benchmark behaviour. A relevant future extension would be to compare the boiling implementation against the subcooled-boiling wall heat-transfer model of Steiner et al. [22].

4.1 Chemical equilibrium solver verification

This section verifies the Cantera-based gas-state calculation used by the thermal model. The purpose is to check that the implemented equilibrium workflow gives combustion temperatures, thermodynamic properties, and rocket-performance quantities that are consistent with established reference tools for comparable assumptions and input conditions.

The comparisons use the San Diego combustion mechanism as the Cantera source file, consistent with the gas-property implementation described in Section 3. Since the selected mechanism defines the available species, thermodynamic data, reaction set, and transport data, the comparisons in this section verify the complete gas-state setup used in the thesis rather than Cantera in isolation.

Three verification cases are considered. The first compares constant-enthalpy, constant-pressure equilibrium calculations for methane–air combustion against CEA for lean, stoichiometric, and rich mixtures. This provides a simple check of equilibrium temperature and major thermodynamic properties at atmospheric pressure. The second case compares chamber, throat, and exit states for a rocket-like hydrogen–oxygen expansion against CEA. The third case compares an ethanol–nitrous-oxide rocket-performance calculation against RPA, which provides convenient reference data for the propellant combination used in the Propulse NTNU engine cases.

The second case checks whether the implementation can reproduce ideal rocket-performance states for a hydrogen–oxygen engine. Liquid hydrogen and liquid oxygen are evaluated at 3000 psi, which is representative of Space Shuttle Main Engine operating pressure. The chamber state is calculated from constant-enthalpy, constant-pressure equilibrium, and the throat and nozzle states are obtained from an isentropic expansion procedure. Selected thermodynamic properties are compared with CEA, together with derived rocket-performance quantities such as characteristic velocity and specific impulse.

The third case considers ethanol and nitrous oxide at moderate pressure, a propellant combination relevant to student-developed liquid rocket engines. The CEA input data used in this work do not include liquid nitrous oxide in the same convenient form as RPA. The Cantera results are therefore compared with RPA for this case. RPA has also been compared against CEA for common propellant mixtures in previous work [20, 21].

The resulting comparisons are summarized in Tables 4.1–4.3.

Table 4.1: Constant enthalpy-pressure equilibrium results for methane-air at 1 atm for lean, stoichiometric, and rich mixtures.

Parameter	$\phi = 0.8$			$\phi = 1.0$			$\phi = 1.2$		
	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)
P [bar]	1.0132	1.0132	0.00%	1.0132	1.0132	0.00%	1.0132	1.0132	0.00%
T [K]	1996.6	2003.2	-0.33%	2225.2	2231.6	-0.29%	2136.6	2141.3	-0.22%
ρ [kg/m ³]	0.1705	0.1700	0.34%	0.1508	0.1504	0.28%	0.1512	0.1509	0.20%
H [kJ/kg]	-208.41	-204.32	1.96%	-257.08	-252.94	1.61%	-304.68	-300.59	1.34%
U [kJ/kg]	-802.56	-800.41	0.27%	-929.04	-926.79	0.24%	-974.65	-971.86	0.29%
G [kJ/kg]	-19222	-19283	-0.32%	-22136	-22199	-0.28%	-21850	-21900	-0.23%
S [kJ/kg-K]	9.5231	9.5245	-0.01%	9.8325	9.8345	-0.02%	10.0840	10.0870	-0.03%
M [1/n]	27.941	27.937	0.01%	27.533	27.535	-0.01%	26.515	26.522	-0.03%
Mass fractions	CEA	Cantera	Error (abs.)	CEA	Cantera	Error (abs.)	CEA	Cantera	Error (abs.)
Ar	1.234%	1.235%	0.000%	1.221%	1.221%	0.000%	1.208%	1.208%	0.000%
CO	0.051%	0.053%	-0.002%	0.913%	0.877%	0.036%	4.771%	4.730%	0.042%
CO ₂	12.132%	12.083%	0.049%	13.652%	13.663%	-0.011%	10.402%	10.422%	-0.020%
H	0.000%	0.000%	0.000%	0.001%	0.001%	0.000%	0.002%	0.002%	0.000%
H ₂	0.002%	0.002%	0.000%	0.026%	0.025%	0.001%	0.203%	0.201%	0.002%
H ₂ O	9.887%	9.884%	0.003%	11.963%	11.966%	-0.003%	12.758%	12.778%	-0.020%
NO	0.326%	0.000%	0.326%	0.202%	0.000%	0.202%	0.013%	0.000%	0.013%
N ₂	72.017%	72.179%	-0.162%	71.283%	71.388%	-0.105%	70.597%	70.614%	-0.017%
O	0.007%	0.008%	-0.001%	0.012%	0.014%	-0.002%	0.001%	0.001%	0.000%
OH	0.111%	0.114%	-0.003%	0.197%	0.209%	-0.012%	0.040%	0.042%	-0.002%
O ₂	4.232%	4.442%	-0.210%	0.529%	0.636%	-0.107%	0.003%	0.004%	-0.001%

Table 4.2: Chemical equilibrium results using constant enthalpy-pressure for chamber and constant entropy-pressure for throat and nozzle exit, for SSME parameters.

Parameter	Chamber			Throat			Nozzle exit		
	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)
P [bar]	206.84	206.84	0.00%	118.85	117.43	1.19%	0.2152	0.1680	21.93%
T [K]	3598.8	3603.4	-0.13%	3381.7	3381.2	0.01%	1233.8	1172.6	4.96%
ρ [kg/m ³]	9.411	9.404	0.08%	5.8080	5.7431	1.12%	0.0296	0.0243	17.84%
H [kJ/kg]	-986.31	-986.28	0.00%	-2161.66	-2187.20	-1.18%	-10544.80	-10722.00	-1.68%
U [kJ/kg]	-3184.12	-3185.80	-0.05%	-4208.00	-4231.80	-0.57%	-11271.80	-11413.00	-1.25%
G [kJ/kg]	-62769	-62871	-0.16%	-60216	-60255	-0.06%	-31727	-30859	-2.74%
S [kJ/kg-K]	17.1677	17.1740	-0.04%	17.1677	17.1740	-0.04%	17.1677	17.1740	-0.04%
M [1/n]	13.614	13.621	-0.05%	13.740	13.749	-0.07%	14.111	14.112	-0.01%
γ_s [-]	1.1475	1.169	-1.87%	1.1487	1.175	-2.29%	1.2539	1.259	-0.41%
Sonic velocity [m/s]	1588.1	1604.0	-1.00%	1533.2	1551.0	-1.16%	954.8	932.5	2.34%
Mach [-]	0.000	0.000	0.00%	1.000	1.003	-0.30%	4.579	5.064	-10.59%
Mass fractions	CEA	Cantera	Error (abs.)	CEA	Cantera	Error (abs.)	CEA	Cantera	Error (abs.)
H	0.188%	0.191%	-0.003%	0.149%	0.151%	-0.002%	0.000%	0.000%	0.000%
HO ₂	0.008%	0.009%	-0.001%	0.004%	0.004%	0.000%	0.000%	0.000%	0.000%
H ₂	3.663%	3.652%	0.011%	3.594%	3.582%	0.012%	3.486%	3.485%	0.001%
H ₂ O	90.821%	91.003%	-0.182%	92.443%	92.629%	-0.186%	96.514%	96.515%	-0.001%
H ₂ O ₂	0.004%	0.004%	0.000%	0.002%	0.000%	0.002%	0.000%	0.000%	0.000%
O	0.238%	0.248%	-0.010%	0.143%	0.147%	-0.004%	0.000%	0.000%	0.000%
OH	4.571%	4.364%	0.207%	3.347%	3.157%	0.191%	0.000%	0.000%	0.000%
O ₂	0.505%	0.530%	-0.025%	0.318%	0.329%	-0.011%	0.000%	0.000%	0.000%
Transport property (Frozen)	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)
Viscosity [Pa-s]	1.083E-04	1.04E-04	3.71%	1.03E-04	9.95E-05	3.82%	4.60E-05	4.14E-05	10.06%
Cp [kJ/kg-K]	3.796	3.801	-0.14%	3.760	3.759	0.05%	2.910	2.866	1.48%
Conductivity [W/m-K]	0.5778	0.6336	-9.66%	0.5436	0.5982	-10.05%	0.1904	0.1999	-5.02%
Prandtl number [-]	0.7117	0.6254	12.13%	0.7159	0.6254	12.64%	0.7029	0.5931	15.62%
Rocket parameter	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)	CEA	Cantera	Error (rel.)
C* [m/s]	N/A			2322.8	2308.3	0.62%	2322.8	2308.3	0.62%
C_f [-]				0.6601	0.674	-2.10%	1.8823	2.046	-8.68%
Isp [m/s]				1533.2	1555.7	-1.46%	4372.3	4722.2	-8.00%

The results show good agreement between CEA and the present Cantera implementation for the thermodynamic properties at the chamber and throat, with deviations typically within 0–2%.

Specific heat is not compared directly because mixture-averaged values are not formed identically in all tools. Larger discrepancies appear in derived rocket-performance quantities, such as thrust coefficient and specific impulse. These quantities depend on the full quasi-one-dimensional nozzle procedure and on implementation details beyond the local equilibrium calculation. The discrepancy is therefore interpreted primarily as a limitation of the present performance-postprocessing routine rather than as evidence of a large thermochemical-property error.

For detailed performance prediction, errors on the order of 10% in thrust coefficient or specific impulse would not be acceptable. For the heat-transfer analysis in this thesis, the consequence is less severe. The gas-side heat-transfer correlations depend on the thermodynamic state, but the dominant sensitivity is often through transport properties such as viscosity, thermal conductivity, and Prandtl number. The verification therefore supports the use of the Cantera implementation for thermal analysis, while not establishing the code as a high-accuracy rocket-performance tool.

Table 4.3: Chemical equilibrium and performance parameters for Case 3.

Parameter	Chamber			Throat			Nozzle exit		
	RPA	Cantera	Error (rel.)	RPA	Cantera	Error (rel.)	RPA	Cantera	Error (rel.)
P [bar]	20.00	20.00	0.00%	11.39	11.20	1.65%	1.0130	0.9696	4.29%
T [K]	2935.8	3007.9	-2.46%	2792.7	2796.3	-0.13%	1801.4	1851.7	-2.79%
ρ [kg/m ³]	2.079	2.022	2.72%	1.2820	1.2287	4.16%	0.1742	0.1622	6.89%
H [kJ/kg]	-291.22	-106.81	63.32%	-812.26	-657.49	19.05%	-2580.66	-2496.00	3.28%
U [kJ/kg]	-1253.46	-1095.90	12.57%	-1700.56	-1569.40	7.71%	-3162.18	-3093.70	2.17%
S [kJ/kg-K]	10.3296	10.3930	-0.61%	10.3296	10.3930	-0.61%	10.3296	10.3930	-0.61%
M [1/n]	25.367	25.284	0.33%	25.550	25.497	0.21%	25.757	25.756	0.00%
γ_s [-]	1.159	1.217	-4.98%	1.1784	1.222	-3.70%	1.2379	1.242	-0.33%
Sonic velocity [m/s]	1056.2	1097.0	-3.86%	1020.8	1055.0	-3.35%	848.4	861.8	-1.58%
Mach [-]	0.000	0.000	0.00%	1.000	1.000	0.00%	2.522	2.577	-2.18%
Mass fractions	RPA	Cantera	Error (abs.)	RPA	Cantera	Error (abs.)	RPA	Cantera	Error (abs.)
CO	11.595%	11.762%	-0.167%	10.926%	11.050%	-0.124%	9.107%	9.230%	-0.123%
CO ₂	16.172%	15.909%	0.263%	17.224%	17.029%	0.195%	20.083%	19.888%	0.195%
H	0.024%	0.030%	-0.006%	0.015%	0.019%	-0.004%	0.000%	0.000%	0.000%
H ₂	0.337%	0.337%	0.000%	0.325%	0.323%	0.002%	0.406%	0.397%	0.009%
H ₂ O	19.426%	19.227%	0.199%	19.811%	19.697%	0.114%	19.483%	19.561%	-0.078%
N ₂	50.708%	50.920%	-0.212%	50.817%	50.920%	-0.103%	50.918%	50.920%	-0.002%
NO	0.450%	0.000%	0.450%	0.217%	0.000%	0.217%	0.000%	0.000%	0.000%
O	0.066%	0.109%	-0.043%	0.025%	0.044%	-0.019%	0.000%	0.000%	0.000%
O ₂	0.338%	0.553%	-0.215%	0.138%	0.243%	-0.105%	0.000%	0.000%	0.000%
OH	0.882%	1.151%	-0.269%	0.499%	0.675%	-0.176%	0.000%	0.004%	-0.004%
Transport property (Frozen)	RPA	Cantera	Error (rel.)	RPA	Cantera	Error (rel.)	RPA	Cantera	Error (rel.)
Viscosity [Pa-s]	9.339E-05	8.77E-05	6.08%	8.89E-05	8.34E-05	6.19%	6.67E-05	6.27E-05	6.01%
Cp [kJ/kg-K]	1.758	1.762	-0.22%	1.745	1.750	-0.26%	1.647	1.654	-0.44%
Conductivity [W/m-K]	0.2429	0.2442	-0.54%	0.2275	0.2264	0.48%	0.1561	0.1599	-2.43%
Prandtl number [-]	0.676	0.638	5.64%	0.682	0.645	5.45%	0.704	0.649	7.84%
Rocket parameter	RPA	Cantera	Error (rel.)	RPA	Cantera	Error (rel.)	RPA	Cantera	Error (rel.)
C* [m/s]	N/A						1528.2	1526.0	0.14%
Cf [-]							1.4002	1.455	-3.94%
Isp [m/s]							2411.6	2220.9	7.91%

The agreement between the present implementation and RPA in the ethanol–nitrous-oxide case is better than the agreement observed for some of the derived performance quantities in the hydrogen–oxygen CEA comparison. The most likely reason is that the pressure and nozzle area ratio are more moderate, giving less extreme thermodynamic conditions and reducing sensitivity to the details of the nozzle-performance calculation.

Other works have shown closer agreement between CEA and Cantera for ideal rocket calculations, indicating that the remaining performance discrepancies are associated with the present implementation choices rather than with an inherent inconsistency between the tools [23]. Based on the comparisons in this section, the Cantera-based implementation is considered suitable for providing gas states and gas transport properties for the heat-transfer analyses in the remainder of the thesis.


4.2 Single-phase forced-convection channel validation

The implementation of the solver is verified using a simple single-phase forced-convection problem. A rectangular channel is heated from one side with a constant heat flux. The fluid is water at no more than 2 bar and $20 - 60^\circ\text{C}$. The Nusselt number is calculated along the length of the channel and compared to reference work. [24] The reference work uses both empirical data as well as numerical results with multiple heat transfer coefficient models, including one that was developed specifically for this flow configuration. Two different Reynolds numbers are used for verification, both being in the fully turbulent regime.


The solver uses Dirichlet boundary conditions for the gas-side, where the temperature is prescribed. To solve the thermal system in [24], it is necessary to change to a Neumann boundary. This is done by constructing the gas-side heat transfer coefficient to achieve the desired heat flux.

$$q'' = h_{\text{gas}}(T_{\text{ref}} - T_i) \iff h_{\text{gas}} = \frac{q''}{T_{\text{ref}} - T_i} \quad (38)$$


The reference temperature is arbitrarily set to 1000 K. The heat flux is calculated as 93% of the electrical power. [24] The mass flow rate is calculated from the user-provided Reynolds number and then used directly in the solver.




(a) Present work, $\text{Re} = 4536$



(b) Reference, $\text{Re} = 4536$ (Fig. 17a [24])



(c) Present work $\text{Re} = 22185$



(d) Reference, $\text{Re} = 22185$ (Fig. 17b [24])

Figure 4.1: Comparison between present work and reference at two different Reynolds numbers.

For each Reynolds number, the total pressure loss and temperature rise are independent of the selected liquid-side heat-transfer coefficient, confirming energy conservation on an integral scale. The Nusselt-number distributions agree well with the reference data, indicating that the single-phase forced-convection implementation is correct for this case. However, this verification does not establish accuracy for rocket-like flows, where the result is strongly model-dependent.

4.3 Pratt & Whitney RL10A-3-3A cooling-jacket benchmark

The Pratt & Whitney RL10 is one of the most extensively used cryogenic upper-stage rocket engines. Developed in the late 1950s and first flown in 1963 on the Centaur upper stage, it was the first operational liquid-hydrogen/liquid-oxygen rocket engine. Owing to its long operational history and the availability of published performance and thermal data, the RL10 is frequently used as a reference case in studies of rocket engine heat transfer and regenerative cooling systems [6, 7, 25].




Figure 4.2: An RL10 engine stands in a vacuum chamber at Aerojet Rocketdyne's test stand in West Palm Beach, Florida. Photo by Dimitri Gerondidakis; public domain.

The RL10 employs liquid hydrogen as both fuel and coolant. Prior to injection, the hydrogen is routed through cooling channels along the chamber and nozzle walls, removing a large fraction of the heat load imposed by the hot combustion gases while simultaneously preheating the fuel. The resulting coupling between wall heat transfer, coolant pressure loss, and fuel temperature rise makes the RL10 a useful reference case for evaluating a one-dimensional regenerative-cooling model.

In the present work, the RL10A-3-3A is used as a system-level benchmark to assess the coupled thermodynamic and heat-transfer modeling of the cooling-channel solver. Unlike the preceding verification cases, which isolate individual sub-models, this case evaluates the combined behavior of gas-side convection, wall conduction, and coolant-side convection under representative engine conditions. The comparison is performed against previously published numerical models, including the Pratt & Whitney internal simulation, the RTE model, and the New System Model [6, 7, 25, 26]. Reference distributions of wall temperature, heat flux, coolant temperature, and coolant pressure are taken from the compiled results reported by Binder et al. [7]. Available experimental coolant inlet and outlet data provide a limited validation of the overall pressure and energy balance. The objective is therefore not exact agreement with a single reference solution, but rather to reproduce the expected magnitude and axial trends of the thermal response and to obtain consistent outlet conditions.

The RL10 cooling jacket employs a tubular-wall construction with a one-and-one-half pass configuration. As shown schematically in Fig. 4.3, the coolant enters through the inlet manifold and is distributed into a set of long tubes forming the upstream chamber section. At the turnaround manifold, the flow is redirected into a second set of shorter tubes, increasing the number of cooling passages in the nozzle section before exiting through the collector manifold.




Figure 4.3: Schematic of the RL10 tubular cooling configuration, showing the inlet manifold, turnaround manifold, collector manifold, and the division between long and short cooling tubes. Reproduced from NASA-CR-120145 [26]; public use permitted.

Figure 4.4 shows the inlet manifold region of the cooling jacket. Upstream of the manifold, the thrust chamber wall is formed by 180 full-length tubes, while downstream an additional set of tubes is introduced, resulting in 360 cooling passages in the nozzle section. This one-and-one-half pass configuration allows the passage count to increase with nozzle expansion while retaining a compact tubular-wall design.

The cooling channels also exhibit significant axial variation in cross-section. As shown in Fig. 4.5, the tubes are approximately round in the upstream chamber section, while downstream of the inlet manifold they become narrower and more flattened. This variation motivates representing the channel geometry as axial profiles rather than assuming a constant cross-section.






Figure 4.4: Overview of the inlet manifold region of the RL10 thrust chamber, showing the transition from the upstream section (left) with 180 full-length tubes to the downstream section (right) where an additional 180 tubes are introduced. Image: Mike Jetzer/heroicrelics.org [27].



(a) Cooling tubes just aft of the throat (left side of Fig. 4.4), where the wall is formed by 180 approximately round tubes.



(b) Cooling tubes downstream of the inlet manifold (right side of Fig. 4.4), where 180 additional tubes are introduced and the cross-sections become flatter.

Figure 4.5: Cut-away views of the RL10 tubular cooling jacket, illustrating the axial variation in cooling-tube geometry. Images: Mike Jetzer/heroicrelics.org [28, 29].

4.3.1 Geometry reconstruction

The hot-wall geometry of the RL10A-3-3A thrust chamber is reconstructed from published data in Binder et al. [7] and Tomsik [6]. Key geometric parameters, including chamber and throat radii and characteristic contour dimensions, are taken from tabulated values where available, as summarized in Table 4.4. These parameters define the overall scale and consistency of the reconstructed geometry.

Table 4.4: RL10 hot-wall geometry parameters used for the benchmark.

Parameter	Symbol	Value	Source
Chamber radius	R_c	5.13 in	[6, 7]
Throat radius	R_t	2.47 in	[6, 7]
Nozzle exit radius	R_e	$R_t\sqrt{61.5} \approx 19.3$ in	[6]
Contraction radius	R_2	$4R_c$	[6]
Throat entry radius	R_1	$2R_t$	[6]
Throat exit radius	R_n	$0.2R_t$	[6]
Contraction angle	b	21°	[6]
Nozzle entry angle	θ_n	32.5°	[6]
Nozzle exit angle	θ_e	Not reported	—
Manifold split area ratio	A_{split}/A_t	14.3	[6]
Chamber and nozzle length	L	1.476 m	[7]
Nozzle length fraction		75% bell	[6]

The axial variation of the chamber contour and cooling-channel geometry is reconstructed from graphical data reported by Tomsik [6]. The extracted data are interpolated onto the computational grid using monotone piecewise-cubic interpolation, yielding a smooth and physically consistent radius distribution along the chamber and nozzle. Although key geometric parameters are available in tabulated form, the full analytic reconstruction of the chamber contour is not uniquely defined. In particular, the standard thrust-chamber parametrization (cf. Fig. 1.2) requires the nozzle exit angle θ_e , which is not reported for the RL10A-3-3A (Table 4.4). The adopted reconstruction therefore preserves the dominant geometric features governing the gas-side heat-transfer behaviour, while remaining consistent with the available geometric constraints. The resulting chamber contour is shown in Fig. 4.6.




Figure 4.6: Axial distribution of chamber radius used in the reconstructed RL10 hot-wall geometry.

The one-and-one-half-pass cooling configuration is reconstructed by splitting the channel path at the nozzle location corresponding to $A/A_t \approx 14.3$. Upstream of this location, the thrust chamber wall consists of 180 full-length tubes, while downstream the effective channel count is increased to account for the combined short- and long-tube region. The internal tube cross-section is approximated as an ellipse when evaluating coolant flow area, wetted perimeter, and hydraulic diameter.

The resulting axial variation in per-channel flow area and hydraulic diameter is shown in Fig. 4.7 and Fig. 4.8, respectively.




Figure 4.7: Axial variation of per-channel areas in the reconstructed RL10 cooling jacket.




Figure 4.8: Axial variation of hydraulic diameter in the reconstructed RL10 cooling jacket.

The reconstructed channel geometry exhibits a pronounced decrease in hydraulic diameter and a reduction in per-channel flow area at approximately 300 mm downstream of the throat, corresponding to the location where the cooling configuration transitions from 180 full-length tubes to the combined 360-tube region. This increase in channel count reduces the dimensions of the individual flow passages, which is clearly reflected in the distributions shown in Fig. 4.7 and Fig. 4.8. It is noted that the per-channel heat transfer area decreases significantly at the manifold split, as the number of cooling passages is doubled. However, the corresponding reduction in per-channel flow area is less pronounced. As a result, the hydraulic characteristics do not scale directly with tube count, and the local flow area is partially preserved despite the increase in the number of channels.

As a consistency check, the total coolant-channel volume of the reconstructed geometry was compared against the reference value reported by Binder et al. [7]. The reconstructed volume is within 5% of the reference value, indicating that the overall geometric representation is consistent with the published data despite the use of digitized geometry input.

4.3.2 Modelling decisions

The RL10 case is used as a system-level consistency check against the cooling-jacket model of Binder et al. [7]. The purpose of this comparison is not to assess which heat-transfer correlation is most accurate, but to verify that the present solver can reproduce the behaviour of a published regenerative-cooling model when supplied with comparable geometry, operating conditions, and empirical submodels. For this reason, the gas- and liquid-side heat-transfer formulations used by Binder et al. are adopted for the RL10 benchmark. This reduces the number of modelling differences and makes discrepancies more directly attributable to the present implementation, geometry reconstruction, and remaining modelling assumptions.

This makes the RL10 benchmark different from the Propulse engine simulations presented later in the thesis. For P25 and P26, the objective is to apply the modelling approach developed in this work, using the default heat-transfer models described in Section 2.4. For the RL10 case, the objective is instead to reproduce a published cooling-jacket calculation as closely as possible. The gas- and liquid-side heat-transfer models are therefore treated as case-specific benchmark settings, chosen to match the formulation of Binder et al. [7].

The heat flux is expressed by Binder [7] in terms of an enthalpy difference,

$$q_g'' = \frac{h_{gr}}{c_{p,gr}} (H_{aw} - H_{hw}), \quad (39)$$

rather than a temperature difference as in the default Bartz formulation. Since the one-dimensional solver in the present work is formulated in terms of a temperature-driven heat transfer coefficient, derived from $q'' = h_g(T_{aw} - T_w)$, the enthalpy-driven heat flux is converted to an equivalent coefficient,

$$h_g = \frac{h_{gr}}{c_{p,gr}} \frac{H_{aw} - H_{hw}}{T_{aw} - T_{hw}}. \quad (40)$$

before being passed to the solver. This conversion preserves the heat flux while adapting the formulation to the solver interface.

In this model, the gas-side wall enthalpy is first evaluated from the local gas static pressure and wall temperature,

$$H_{hw} = H(P_g, T_{hw}). \quad (41)$$

The gas properties used in the heat transfer coefficient are evaluated at an Eckert-type reference enthalpy state,

$$H_{gr} = \frac{1}{2} (H_{hw} + H_g) + 0.18 (H_{g0} - H_g), \quad (42)$$

where H_g is the local static gas enthalpy and H_{g0} is the chamber stagnation enthalpy. The reference-state properties are then evaluated from P_g and H_{gr} .

The adiabatic-wall enthalpy is given by

$$H_{aw} = H_g + \text{Pr}_{gr}^{1/3} (H_{g0} - H_g), \quad (43)$$

where Pr_{gr} is evaluated at the reference enthalpy state. The corresponding Bartz-type heat transfer coefficient is given by

$$h_{gr} = 0.026 \frac{k_{gr}}{D_h} \left(\frac{c_{p,gr}}{k_{gr} \mu_{gr}} \right)^{0.4} \left(\frac{\dot{m}_g D_h}{A_{\text{chmb}}} \right)^{0.8} \left(\frac{T_g}{T_{gr}} \right)^{0.8} \phi_B. \quad (44)$$

Here, the subscript gr denotes the reference state, and ϕ_B is an empirical correction factor. Binder et al. used $\phi_B = 1.08$ to improve agreement with RL10 test data [7]. In the present study, $\phi_B = 1.0$

is used, corresponding to the uncorrected form of the correlation. This avoids introducing an additional empirical adjustment beyond the heat-transfer formulation itself.

For the RL10 benchmark, the coolant-side heat transfer coefficient is evaluated using the Colburn-type correlation reported by Binder et al. [7]. This choice is made to maintain consistency with the reference model, rather than to reflect the default modelling approach used elsewhere in this work.

The coolant-side heat flux is expressed in terms of the local temperature difference between the coolant-side wall and the bulk coolant,

$$q_c'' = h_c (T_{cw} - T_c), \quad (45)$$

where T_{cw} is the coolant-side wall temperature and T_c is the local bulk coolant temperature.

Following Binder et al. [7], the heat transfer coefficient is calculated as

$$h_c = 0.023 \frac{k_f}{D_h} \left(\frac{c_{p,f} \mu_f}{k_f} \right)^{0.4} \left(\frac{\dot{m}_c D_h}{A_c \mu_f} \right)^{0.8} \phi_{\text{curv}}, \quad (46)$$

where the coolant properties are evaluated at the film temperature

$$T_f = \frac{1}{2} (T_c + T_{cw}), \quad (47)$$

and the term $\dot{m}_c D_{h,c} / (A_c \mu_f)$ corresponds to the Reynolds number expressed in terms of mass flow rate rather than velocity.

Binder et al. [7] also include an empirical curvature correction factor, ϕ_{curv} , to account for the influence of channel curvature on the coolant-side boundary layer. In the present benchmark, this factor is set to unity, $\phi_{\text{curv}} = 1$, in order to avoid introducing an additional empirical correction in a case where the overall heat transfer is already over-predicted relative to the reference model.

The coolant-side model therefore follows the formulation used by Binder et al. [7], with the exception that the curvature correction is omitted to avoid introducing an additional empirical tuning parameter.

4.3.3 Operating conditions

The operating conditions listed in Table 4.5 define the steady-state conditions used in the RL10 benchmark. The primary thermodynamic and flow parameters are taken from Binder et al. [7], while injector and coolant inlet temperatures are taken from Tomsik [6]. The coolant mass flow rate reported corresponds to the total flow through the cooling jacket.

Table 4.5: RL10A-3-3A operating conditions used for the benchmark.

Parameter	Value	Source
Chamber pressure	482.0 psia	[7]
Mixture ratio (O/F)	5.26	[7]
Combustion-gas mass flow rate	37.36 lbm/s	[7]
Total coolant mass flow rate	5.973 lbm/s	[7]
Coolant inlet pressure	1025 psia	[7]
Coolant inlet temperature	61.0 deg R	[6]
Cooling-jacket pressure drop	242.1 psi	[7]
Cooling-jacket temperature rise	344.4 deg R	[7]
Oxidizer injector temperature	182.0 deg R	[6]
Fuel injector temperature	358.0 deg R	[6]

4.3.4 Results

The results obtained with the present solver are compared with the RL10 cooling-jacket model of Binder et al. [7]. The comparison includes axial distributions of wall temperature, heat flux, coolant temperature, and coolant pressure, as well as integral quantities such as total coolant temperature rise and pressure drop.

The reconstructed RL10 model represents a system-level approximation based on digitized geometry and published operating conditions. The comparison should therefore be interpreted as a consistency check against the modelling approach of Binder et al. [7], rather than a direct validation against the physical engine. Differences are expected due to variations in geometric representation and in the specific heat-transfer and friction correlations employed.

Figure 4.9 shows the predicted wall temperature distribution along the chamber and nozzle, with the axial coordinate defined relative to the throat. All models predict a peak near the throat, corresponding to the location of maximum heat flux and flow acceleration. The present model reproduces the overall shape of the temperature distribution and captures the throat region accurately. The predicted throat gas-side wall temperature deviates by only 2.2% from the reference value (Table 4.6). Lower temperatures are predicted in the nozzle region, which is likely due to differences in the coolant-side heat transfer model and the simplified representation of the cooling-channel geometry.




Figure 4.9: Predicted wall temperature distribution along the RL10 combustion chamber and nozzle. Negative distances correspond to the combustion chamber, while positive distances correspond to the nozzle region downstream of the throat. Reference data were digitized by Hovik [30] from the RL10 data reported in [7].

The heat flux distribution is shown in Figure 4.10. The heat flux increases through the chamber and reaches a maximum at the throat due to the combined effects of high gas temperature, velocity, and thermal gradients. Downstream, the heat flux decreases rapidly as the flow expands. The overall shape of the heat flux distribution is well reproduced. The predicted peak heat flux deviates by 6.8% from the reference value (Table 4.6), indicating that the gas-side heat transfer model provides a consistent estimate of the thermal loading.




Figure 4.10: Heat flux along the RL10 chamber and nozzle predicted by the present model and by previously reported system-level models. Reference data were digitized by Hovik [30] from the RL10 data reported in [7].

The coolant temperature distribution is shown in Figure 4.11. As the coolant flows through the channels, it absorbs heat from the wall and its temperature increases monotonically along the channel length. The predicted coolant temperature rise agrees well with the reference model, with a deviation of 4.7% (Table 4.6). This agreement is significant, as the coolant temperature rise represents an integral measure of the total heat load absorbed by the cooling system. The result therefore indicates that the solver captures the overall heat transfer from the gas to the coolant with good accuracy.




Figure 4.11: Coolant temperature along the RL10 cooling channels. Reference data were digitized by Hovik [30] from the RL10 data reported in [7].

Figure 4.12 shows the predicted coolant pressure along the cooling channels. The pressure increases slightly near the inlet due to flow acceleration and inlet effects, and remains relatively constant downstream with moderate pressure losses.




Figure 4.12: Coolant pressure along the RL10 cooling channels. Reference data were digitized by Hovik [30] from the RL10 data reported in [7].

Table 4.6: Summary of key RL10 benchmark results compared with the New System Model of Binder et al. [7].

Quantity	Present	Reference [7]	Error
Peak heat flux	30.1 MW/m ²	32.3 MW/m ²	-6.8%
Throat gas-side wall temperature	921.8 K	942.4 K	-2.2%
Coolant temperature rise	200.3 K	191.3 K	4.7%
Coolant pressure drop	9.36 bar	16.69 bar	-43.9%
Total heat transfer rate	9.20 MW	8.14 MW	13.0%

A significant discrepancy is observed in the predicted coolant pressure drop. Using the surface roughness value reported by Binder et al. [7], corresponding to $\varepsilon \approx 1.2 \mu\text{m}$, the present model predicts a pressure drop of 9.36 bar compared with 16.69 bar in the reference model, corresponding to an underprediction of 43.9% (Table 4.6).

A sensitivity study showed that the predicted pressure drop is highly dependent on the assumed roughness, increasing from 6.28 bar for a hydraulically smooth channel to 16.20 bar for $\varepsilon = 10 \mu\text{m}$, which approaches the reference value. However, even with the roughness specified by Binder, a substantial discrepancy remains.

This suggests that the difference cannot be attributed solely to surface roughness, but is also influenced by the simplified channel geometry and the modelling of high-speed hydrogen flow near the throat. In this region, the coolant velocity exceeds 400 m/s and the fluid undergoes strong temperature rise and density reduction, leading to large property variations and potentially compressibility effects not fully captured by standard friction correlations.

Importantly, variations in the assumed roughness were found to have negligible impact on the predicted heat flux, wall temperature, and coolant temperature rise. The pressure-drop discrepancy should therefore be interpreted primarily as a hydraulic modelling uncertainty rather than as a deficiency in the thermal predictions.

4.3.5 Conclusion

Taken together, the results demonstrate that the present solver reproduces the main thermal characteristics of the RL10 cooling system. The predicted heat flux distribution and throat wall temperature show good agreement with the reference model, and the coolant temperature rise is predicted within 5%.

Larger discrepancies are observed in the coolant pressure drop, reflecting the sensitivity of hydraulic predictions to friction modelling and geometric detail. These differences are expected for a one-dimensional model and primarily affect the hydraulic response rather than the thermal predictions.

The RL10 benchmark should therefore be interpreted as a consistency check of the coupled thermodynamic and heat-transfer models. The agreement in key thermal quantities indicates that the solver provides physically reasonable predictions for a realistic rocket engine configuration, supporting its application to the Propulse engine analyses in the following chapters.

5 Development of a surface roughness model

In the one-dimensional coolant-flow model, surface roughness enters the pressure-loss calculation through the Darcy friction factor. For a channel control volume, the frictional pressure loss is calculated from Eq. (8), where the friction factor depends on Reynolds number and relative roughness, ε/D_h . The roughness model developed in this chapter therefore addresses a specific modelling input: the effective sand-grain roughness ε used in the channel pressure-loss calculation.

For rough-wall turbulent flow, the hydraulic effect of roughness depends on the size of the roughness elements relative to the near-wall viscous layer. In the hydraulically smooth limit, roughness elements are effectively submerged within the viscous sublayer and the friction factor is mainly Reynolds-number dependent. In the fully rough limit, the friction factor becomes largely independent of Reynolds number and is controlled primarily by relative roughness. The LPBF cooling channels considered here are treated as conventional internal-flow passages operating between these limits, so an estimate of effective relative roughness is required for pressure-loss prediction.

This chapter presents the development of a surface roughness model tailored to additively manufactured cooling channels. The approach consists of two main steps. First, surface roughness is characterized experimentally as a function of local overhang angle using dedicated test specimens, and a parameterized representation suitable for one-dimensional modelling is derived. Second, the resulting model is compared with flow-test data obtained from printed channel specimens to assess whether it provides pressure-loss predictions of the correct magnitude and trend for channels manufactured under the same process conditions. All roughness coupons and channel-flow specimens used in this chapter were manufactured with the same process parameters as P26, but not P25. The resulting roughness model should therefore be interpreted within that manufacturing domain.




Figure 5.1: Workflow used to develop and assess the roughness model. Roughness coupons were used to establish a relation between measured surface texture and model roughness input, while printed channels were tested hydraulically to assess whether the resulting model provides reasonable pressure-loss predictions.

5.1 A brief introduction to surface roughness

The roughness parameters used in this work are briefly defined below. Profile heights are measured relative to the filtered mean profile, while areal heights are measured relative to the filtered mean plane. The parameters R_a and R_z are profile-based quantities, whereas S_a and S_z are areal quantities. Continuous definitions are provided for completeness, followed by the discrete forms used in practice.

The arithmetic mean profile height is defined as

$$R_a = \frac{1}{l_r} \int_0^{l_r} dx |z(x)|, \quad (48a)$$

$$R_a \approx \frac{1}{N} \sum_{i=1}^N |z_i|, \quad (48b)$$

where l_r is the evaluation length, $z(x)$ is the filtered profile height relative to the mean profile, and z_i are the sampled profile heights.

The corresponding peak-to-valley profile measure is defined as

$$R_z = Z_p + Z_v, \quad (49a)$$

$$R_z \approx \max_i z_i + \left| \min_i z_i \right|, \quad (49b)$$

where the continuous peak and valley terms are

$$Z_p = \max_{x \in [0, l_r]} z(x), \quad (50a)$$

$$Z_v = \left| \min_{x \in [0, l_r]} z(x) \right|. \quad (50b)$$

In the discrete case, the same peak-to-valley definition is applied to the sampled profile data. These definitions follow the profile-texture framework of ISO 21920-2:2021 [31].

The areal analogue of R_a is defined as

$$S_a = \frac{1}{A} \iint_A dx dy |z(x, y)|, \quad (51a)$$

$$S_a \approx \frac{1}{N_x N_y} \sum_{i=1}^{N_x} \sum_{j=1}^{N_y} |z_{ij}|, \quad (51b)$$

where A is the evaluation area, $z(x, y)$ is the filtered surface height relative to the mean plane, and z_{ij} are the sampled surface heights.

The corresponding areal peak-to-valley measure is defined as

$$S_z = S_p + S_v, \quad (52a)$$

$$S_z = \max_{i,j} z_{ij} + \left| \min_{i,j} z_{ij} \right|, \quad (52b)$$

where the continuous peak and pit terms are


$$S_p = \max_{(x,y) \in A} z(x, y), \quad (53a)$$

$$S_v = \left| \min_{(x,y) \in A} z(x, y) \right|. \quad (53b)$$


In the discrete case, the same peak-to-valley definition is applied to the sampled surface grid. These definitions follow the areal-texture framework of ISO 25178-2:2021 [32].

5.2 Design of experiment


To isolate the effect of overhang angle on surface roughness, a set of dedicated coupons was designed. Each coupon corresponds to a single overhang angle and incorporates three wall-thickness regions, allowing thickness effects to be assessed within a single specimen. Four coupons were included in the final dataset, corresponding to overhang angles of 0° , 15° , 36° , and 45° . An additional coupon at 50° was manufactured but excluded from the analysis, as this angle lies outside the recommended process window of the printer and may therefore not be representative of typical build conditions.



(a) Overview of the roughness-coupon specimens. The specimens include the overhang angles used in the final dataset, with an additional 50° specimen shown for reference.



(b) Coupon showing the three wall-thickness levels included for each overhang angle.



(c) Close-up of a measured coupon surface, showing the characteristic LPBF surface texture.

Figure 5.2: Roughness-coupon specimens used for surface roughness characterization. The coupons were designed to separate the effects of overhang angle and wall thickness while preserving representative as-printed LPBF surfaces. Photos by Philip Draget, Propulse NTNU.

Each coupon contains both upskin and downskin surfaces. For each surface, the roughness parameters R_a , R_z , S_a , and S_z were evaluated across all three thickness regions. The complete dataset therefore consists of $4 \times 3 \times 2 \times 4 = 96$ individual roughness values.

In addition to the roughness coupons, a cooling-channel test specimen was manufactured for hydraulic assessment of the model. The specimen corresponds to a single cooling channel extracted from the P25 geometry. The channel geometry is therefore not repeated here; a detailed description of the P25 cooling system is provided in Chapter 6.

Table 5.1: Design of the roughness-coupon specimens. Each specimen corresponds to one overhang angle and contains three wall-thickness regions.

Quantity	Values
Overhang angle, θ [deg]	0, 15, 36, 45
Wall thicknesses per specimen [mm]	0.8, 1.2, 2.4
Number of specimens	4

Unlike the original P25 chamber, the channel specimen was manufactured using the updated process parameters described above. This provides a direct comparison between the previously modelled P25 cooling-channel geometry and an equivalent printed channel for which experimental pressure-drop data are available.



Figure 5.3: Cooling-channel test specimen manufactured for hydraulic assessment of the model: a single channel extracted from the P25 cooling-channel geometry and reprinted with the updated process parameters.

5.3 Surface roughness characterization

Surface roughness measurements were performed using a Keyence VHX-X1 optical microscope equipped with VHX-E20 and VHX-E100 lenses at a nominal magnification of $50\times$. Profile roughness parameters were evaluated from traces measured in the vertical (Z) direction, while areal parameters were obtained from the corresponding surface topography data. Prior to evaluation, the raw measurements were filtered to remove form and longer-wavelength waviness components. For the profile data, $\lambda_s = 0$ and $\lambda_c = 8$ were used, while the areal data were evaluated using an S -filter of 25 and an L -filter of 0.8. These settings follow the ISO framework for profile and areal surface-texture evaluation [33, 34].

Inspection of the measurement data, found in Appendix A, indicate that the specimen thickness does not correlate well with the surface roughness. The data were therefore grouped solely by overhang angle for subsequent analysis.

Measured roughness data as a function of overhang angle are presented in Fig. 5.4, with upward-pointing triangles representing upskin surfaces and downward-pointing triangles representing downskin surfaces. A clear distinction is observed between the two surface types: upskin roughness varies only weakly with overhang angle, so linear fits were used for all upskin roughness parameters, while downskin roughness increases more strongly at larger overhang angles and was represented using second-order polynomial fits. The fitted interpolation functions used in the roughness model are shown explicitly in the figure.




Figure 5.4: Surface roughness data for 0° , 15° , 36° , 45° overhang angles, with measurements from all three specimen thicknesses collapsed into common upskin and downskin datasets. The fitted curves shown in the figure are obtained from these thickness-combined datasets. 0° indicates a vertical surface, normal to the print bed. Line measurements use $s = 0$, $c = 8$ and surface measurements use S -filter = 25, L -filter = 0.8 (refer to ISO 21920-3 [33] and ISO 25178-3 [34] for details).

The fitted trends are not equally well behaved for all roughness parameters. For the downskin surfaces, both R_a and S_a exhibit a non-monotonic dependence on overhang angle, with a local decrease at low angles before increasing at higher angles. This behaviour is not expected physically, since downskin roughness in LPBF is generally observed to increase with increasing overhang angle.

By contrast, the R_z downskin fit is nearly monotonic over the relevant range, showing only a negligible deviation from monotonic behaviour at very low angles. In addition, R_z provides the highest coefficient of determination for both the upskin and downskin fits. This makes R_z the most consistent and physically plausible descriptor for the subsequent roughness model.

The quantities R_a and S_a are average-amplitude measures and therefore provide a compact description of the overall roughness level. In contrast, R_z and S_z are extreme-height measures and are more sensitive to isolated peaks, partially fused particles, and deep valleys. This distinction is particularly relevant for LPBF surfaces, which are typically anisotropic and three-dimensional, with local asperities that may not be represented adequately by a single profile trace. Areal parameters are therefore, in principle, better suited to describe the full topography of LPBF surfaces, whereas profile parameters remain useful because they are widely established in engineering practice and are easier to relate to conventional roughness literature.

5.4 Formulation of the roughness model

The model distinguishes between upskin and downskin surfaces and maps from measured geometrical roughness to an equivalent sand-grain roughness suitable for use in one-dimensional flow calculations.

For turbulent flow correlations, surface roughness is typically expressed in terms of an equivalent sand-grain roughness. Adams [35] proposed a conversion in which profilometer definitions are applied to a hypothetical surface consisting of a monolayer of uniform spheres. This leads to linear relations between measured roughness parameters and sand-grain roughness, including

$$k_s = 0.978R_{zd}, \quad (54)$$

and

$$k_s = 5.863R_a, \quad (55)$$

where R_{zd} is a segmented peak-to-valley roughness parameter.

The parameter R_{zd} is defined as the arithmetic mean of peak-to-valley heights evaluated over five sampling lengths [35]. The surface profile is divided into segments of equal length, and the peak-to-valley height is computed for each segment before averaging. In contrast, the present measurements provide the standard R_z parameter, defined as the peak-to-valley height over a single evaluation length. The use of R_z instead of R_{zd} therefore corresponds to using a single-segment estimate rather than an averaged value, which may increase sensitivity to local surface features.

While the Adams conversion is formulated for profile-based parameters, the same correlations are applied here to areal parameters as a practical approximation. The relation in (55) is used for both R_a and S_a , while (54) is applied to R_z and S_z . This introduces two simplifications: R_z is used as a proxy for R_{zd} , and S_z is treated as an equivalent extreme-height descriptor despite being defined over an area. These approximations should be regarded as first-order engineering assumptions.

Using these conversions, angle-dependent sand-grain roughness values are obtained.




Figure 5.5: Angle-dependent equivalent sand-grain roughness computed from the measured roughness data using the Adams conversion.

Peak-to-valley roughness metrics are preferred for hydraulic roughness modelling. Adams [35] showed that segmented peak-to-valley measures (R_{zd}) provide more accurate estimates of sand-grain roughness than mean-amplitude measures such as R_a . This is consistent with rough-wall turbulent-flow physics, where the additional pressure loss is dominated primarily by form drag associated with the largest surface asperities rather than by the average surface height [36].

Between the two peak-to-valley candidates, R_z is selected as the primary descriptor going forward. Although S_z is conceptually attractive because it is computed over an area, the measured downskin

S_z values exhibit behaviour that is inconsistent with the expected monotonic increase in roughness with overhang angle. Furthermore, since S_z is defined by the single highest peak and deepest pit within the sampled area, it is particularly sensitive to rare outliers and to the chosen evaluation area, which can lead to unstable estimates when the surface contains occasional large defects. The R_z values, by contrast, provide a more robust and physically plausible trend across the relevant angle range in the present dataset while still representing an extreme-asperity metric. For these reasons, the angle-dependent sand-grain roughness model in the remainder of this thesis is based on R_z .

For $\theta \in [0^\circ, 45^\circ]$, the fitted sand-grain roughness based on R_z is

$$k_s(\theta) = \begin{cases} 0.624 \theta_{\text{deg}} + 12.684, & \text{upskin,} \\ 0.0504 \theta_{\text{deg}}^2 - 0.219 \theta_{\text{deg}} + 17.383, & \text{downskin,} \end{cases} \quad (56)$$

Given a set of wall segments i with wetted areas A_i and angles θ_i , the local sand-grain roughness for each segment is evaluated as

$$k_{s,i} = \begin{cases} k_s^{\text{up}}(\theta_i), & i \in \text{upskin} \\ k_s^{\text{down}}(\theta_i), & i \in \text{downskin} \\ k_s^{\text{vert}}, & i \in \text{vertical walls / ribs} \end{cases} \quad (57)$$

For one-dimensional analysis, an effective sand-grain roughness $\hat{k}_s$ is defined to account for the multiple walls of the cooling channel:

$$\hat{k}_s = \left(\frac{\sum_i A_i k_{s,i}^p}{\sum_i A_i} \right)^{1/p} \quad (58)$$

where A_i is the area of wall i . The choice of p may vary, with typical values being $p = 1$ (arithmetic mean), $p = 2$ (root-mean-square), and $p \rightarrow \infty$ (max). The effective relative roughness $\hat{k}_s/D_h$ is then computed and can be used directly in the pressure-loss correlations of the thermal system solver.

For channels with small hydraulic diameters and relatively large surface roughness, the roughness may also reduce the effective flow area available to the fluid. In such cases, the roughness elements occupy a non-negligible fraction of the cross-sectional area, and the nominal geometry may overestimate the available flow region. To account for this effect, a constricted-area formulation is considered. In this approach, the channel dimensions are reduced by a roughness-dependent length scale prior to evaluating the flow area, wetted perimeter, and hydraulic diameter. This results in an effective hydraulic diameter that is smaller than the nominal value, thereby increasing the predicted pressure drop.

This type of correction has been proposed for flows with high relative roughness and small hydraulic diameters, where conventional friction-factor correlations may underestimate the influence of surface roughness. Kandlikar et al. [36] showed that roughness affects pressure drop not only through increased wall shear stress, but also through a reduction in the effective flow area due to the presence of surface asperities. For sufficiently large roughness elements, the flow may not fully reattach to the wall, and the effective hydraulic diameter is reduced accordingly.

The constricted-area formulation adopted here follows this interpretation and is intended to capture the first-order geometric effects of roughness on the flow. The approach does not resolve the detailed rough-wall flow field, but provides a practical correction for channels with high relative roughness.

It should be noted that the conversion model proposed by Adams [35] is based on an idealized surface consisting of a monolayer of spheres of uniform diameter. Real LPBF surfaces deviate significantly from this assumption, exhibiting irregular, anisotropic features and partially fused particles of varying size. The resulting sand-grain roughness values should therefore be interpreted as approximate, first-order estimates.

5.5 Comparison with channel flow measurements

Flow tests were conducted using water at a controlled temperature of 4°C. The test specimen was connected to a constant-pressure supply, and the mass flow rate was determined by collecting the discharged fluid in a container of known volume (20 L) and measuring the corresponding filling time. The mass flow rate was then calculated from the collected volume, fluid density, and measurement time. The pressure drop across the test section was measured using differential pressure instrumentation under steady flow conditions.

Table 5.2 compares the predicted and measured pressure drops for the P25 cooling-channel geometry using different modelling approaches. The comparison is performed at the closest achievable operating conditions for each configuration. The experimental data were obtained from in-house flow testing conducted at Propulse NTNU.

Table 5.2: Comparison of predicted and measured pressure drops for the P25 cooling-channel geometry under different modelling approaches.

Process	Model	$\dot{m}$ [g/s]	Δp_{model} [bar]	Δp_{exp} [bar]	Error [%]
Original	Constant roughness	18.1	0.911	1.80	-49.4
	Constant roughness + constricted area	18.1	1.04	1.80	-42.2
Updated	Roughness model	20.5	1.34	1.92	-30.2
	Roughness model + constricted area	20.5	1.71	1.92	-10.9

The results show that a constant roughness assumption leads to a significant underprediction of the pressure drop. Introducing an angle-dependent roughness model based on the measured surface data improves the prediction, but still underpredicts the observed pressure losses.

A substantial improvement is obtained when the effect of surface roughness is also incorporated as a reduction in the effective flow cross-sectional area. This constricted-area formulation accounts for the finite size of surface asperities, which reduce the available flow area and thereby increase the local flow velocity and associated pressure losses. This behaviour is consistent with the findings of Kandlikar et al. [36], who demonstrated that surface roughness influences pressure drop through both increased wall shear stress and geometric flow constriction. For channels with high relative roughness, the latter effect may become dominant, and conventional friction-factor correlations based solely on equivalent sand-grain roughness may underestimate the pressure losses.

This behaviour can be understood by considering the distribution of roughness and flow conditions along the channel. The largest contributions to the pressure drop occur in the nozzle throat region, where the local flow velocities are highest and where downskin surfaces introduce significant roughness. In these regions, the angle-dependent roughness model predicts values ranging from approximately 25 μm near the inlet to about 57 μm in the most severe overhang regions. The corresponding reduction in channel dimensions therefore lies in the range of roughly $2 \times 25 \mu\text{m}$ to $2 \times 57 \mu\text{m}$, which represents a substantial fraction of the local hydraulic dimensions. Since pressure losses scale strongly with velocity, the effective constriction in these sections governs the overall pressure drop. A constant roughness value of 50 μm therefore acts as a reasonable average for the dominant flow regions, leading to similar pressure-drop predictions.

In the present analysis, the effective roughness is evaluated using the downskin roughness values only, which corresponds to the limiting case $p \rightarrow \infty$ in Eq. 58. In this limit, the effective roughness is governed by the maximum local roughness, implying that the roughest surfaces dominate the pressure-loss behaviour. This assumption is consistent with the observed flow physics, as the dominant pressure losses originate in the high-velocity, high-roughness regions of the channel.

These findings suggest that, for additively manufactured cooling channels with high relative roughness, accurate prediction of pressure losses requires accounting for both frictional and geometric effects of the rough surface. The angle-dependent roughness model remains important, however, as it provides a physically grounded description of how roughness varies with local geometry and manufacturing conditions, which is required for predictive modelling of more complex channel configurations.

6 Case study: Propulse NTNU P25

The developed thermal model is applied to the Propulse NTNU P25 combustion chamber. Propulse NTNU is a student rocket team developing liquid-propellant rocket engines and flight vehicles, where short development cycles and limited instrumentation require extensive use of engineering models during design and testing. The P25 engine is designed to deliver 8 kN sustained thrust at sea level for the *Heimdall* flight vehicle, using liquid nitrous oxide and an ethanol–water fuel blend. The fuel is routed through the regenerative cooling channels before injection and therefore also acts as the chamber coolant. The P25 engine represents a relevant case study for the present work because it combines regeneratively cooled additively manufactured hardware with available hot-fire and flight-test data.

The P25 combustion chamber is manufactured from Inconel 625 using laser powder-bed fusion. This manufacturing route enables the chamber wall, cooling channels, manifold features, and external structure to be produced as a single integrated component. For a student-developed engine, this eliminates all welded or brazed joints, simplifies assembly, and allows rapid iteration of the cooling-channel geometry. It also makes the design suitable for compact regeneratively cooled chambers where conventional machining would be difficult or impractical.



Figure 6.1: The P25 engine assembly, consisting of the combustion chamber, injector bulkhead, and injector plate (not visible). Photo by Propulse NTNU.

The chamber was used without post-build heat treatment. Consequently, residual stresses from the additive manufacturing process may remain in the part, and the as-built material state may differ from stress-relieved or solution-annealed Inconel 625. This is relevant when interpreting thermal-margin predictions and observed damage, since thermal gradients during firing can combine with residual manufacturing stresses and local geometric stress concentrations. In the present work, the material is nevertheless modelled using temperature-dependent bulk properties, and the effect of residual stress, anisotropy, and local microstructure is not explicitly included.

Only limited post-processing was performed on the chamber. Sealing surfaces and mechanical interfaces were machined, while the cooling channels and most internal surfaces remained in their additively manufactured condition. The resulting channel roughness and possible local constrictions are therefore important modelling considerations. In this case study, the P25 data are used to assess the sensitivity of the predicted coolant state, wall temperature, and heat flux to different coolant-side modelling assumptions, including surface roughness, reduced effective channel area, and nucleate boiling.

The available experimental data consist primarily of global coolant-side measurements, such as inlet and outlet coolant conditions and pressure loss. These data are used to compare model variants and assess the overall coolant energy balance and hydraulic behaviour.



Figure 6.2: Static ground test of the P25 engine. The thermal analysis in this section is based on flight-test operating data, while the image is included to show the P25 test-firing configuration. Photo by Propulse NTNU.

The P25 engine was subsequently used in the *Heimdall* vehicle, shown in Figure 6.3. This makes the chamber particularly useful as a case study: the hardware progressed from design and ground testing to flight use within a rapid development cycle, while still relying on a limited set of thermal measurements.



Figure 6.3: Launch of the Heimdall flight vehicle from Tarva, Norway in October 2025. Photo by Propulse NTNU.

6.1 Geometry

The P25 thrust chamber was designed for a nominal sea-level thrust of 8 kN at an absolute chamber pressure of 21.013 bar. The nozzle shape is a parabolic bell curve. The injector sits on top of the chamber assembly, allowing the chamber to utilize the full height of the print volume. The resulting total chamber height is approximately 295 mm. The parameters of the parametrized P25 hot-wall contour are listed in Table 6.1.

Table 6.1: P25 hot-wall geometry parameters.

Parameter	Symbol	Value
Chamber radius	R_c	64.52 mm
Throat radius	R_t	30.77 mm
Nozzle exit radius	R_e	56.81 mm
Combustor length	L_c	212.02 mm
Nozzle length	L_e	78.98 mm
Contraction angle	b	30.0°
Nozzle entry angle	θ_n	21.0°
Nozzle exit angle	θ_e	14.5°
Throat entry radius	R_1	$1.5R_t$
Throat exit radius	R_n	$0.6R_t$
Contraction ratio	A_c/A_t	4.40
Expansion ratio	A_e/A_t	3.41
Characteristic length	L^*	780 mm

Figure 6.4 shows the wall-radius distribution along the chamber, which governs the local area ratio and thereby the axial variation in gas pressure, temperature, velocity, and heat flux used in the thermal model.




Figure 6.4: P25 hot-wall geometry

The P25 cooling system consists of 80 axial cooling channels separated by ribs of constant width. The number of channels is determined from a minimum allowable circumferential sector length, such that the sum of channel width and rib width satisfies manufacturing constraints at the smallest radius (the throat). This ensures that the channel layout remains feasible along the entire chamber contour. A constant rib width of 0.70 mm is used throughout the chamber. For a given channel height, this leads to smaller flow areas in regions of reduced circumference, particularly near the throat. The resulting increase in coolant velocity increases the Reynolds number and enhances

convective heat transfer in the region where the heat flux is highest.

Figure 6.5 shows the total coolant-channel flow area used in the model. Variations in the total flow area directly affect the local coolant velocity and pressure loss, and therefore play a central role in determining both the coolant temperature rise and the wall-temperature distribution. The coolant flow is directed opposite to the hot-gas flow, which increases the temperature difference in the most thermally loaded region near the throat.




Figure 6.5: P25 cooling system total flow area

Finally, the hot-wall thickness is set to 0.70 mm. This value is chosen as a lower bound based on manufacturability and structural integrity, in order to reduce thermal resistance and limit thermally induced stresses. For the P25 chamber, thermally induced stresses associated with steep temperature gradients are expected to dominate over pressure-induced stresses, and minimizing the wall thickness is therefore favourable from a thermal-margin perspective.

6.2 Operating conditions

The thermal analysis of the P25 chamber is based on flight-test data from the *Heimdall* vehicle. Although both static ground-test and flight-test data are available, only the flight condition is used in the present case study. This data set is closest to the nominal P25 operating point, whereas the available static ground test ended at a comparatively lean mixture ratio and is therefore less representative of the intended thermal design condition. The data set does not provide local wall-temperature or heat-flux measurements.

Table 6.2: P25 operating and simulation parameters used in the thermal analysis.

Parameter	Value
Chamber pressure	21.0 bar absolute
Oxidizer-to-fuel ratio	3.0
Coolant composition	90 wt% ethanol, 10 wt% water
Coolant mass flow rate	1.0 kg/s
Coolant inlet pressure	27.0 bar
Coolant inlet temperature	−20.0°C
Oxidizer inlet temperature	−25.0°C
Measured coolant outlet temperature	126.0°C
Measured cooling-channel pressure loss	2.0 bar
Assumed TBC thickness	40 μm
TBC thermal conductivity	1.2 W/(m K)

The pressure operating point is extracted from the flight data at the representative nominal steady-state point indicated in Figure 6.6. At this point, the chamber pressure is close to the nominal design value and the cooling-system pressure loss has settled near 2 bar. The supply pressure is 27.0 bar and the chamber pressure is 20.0 bar. The plotted chamber pressure is shown as gauge pressure; the value used in the combustion calculation is converted to absolute pressure.




Figure 6.6: Pressure sensor data from the flight-test of *Heimdall*, powered by the P25 engine.

The coolant temperature rise is extracted from the flight-test thermocouple data shown in Figure 6.7. The channel-inlet temperature remains approximately constant at -20°C during the burn, while the channel-outlet/fuel-injector temperature stabilizes near 126°C . The representative coolant temperature rise used for comparison with the thermal model is therefore $\Delta T_c = 146\text{ K}$. The oxidizer supply temperature is approximately -25°C and is used as the oxidizer inlet condition in the Cantera equilibrium calculation along with the fuel injector temperature.




Figure 6.7: Thermocouple data from the *Heimdall* flight test, showing the coolant inlet and outlet temperatures used to estimate the coolant temperature rise. The dashed vertical line indicates the representative nominal steady-state point used for the thermal analysis.

The thermal-barrier-coating parameters are included in Table 6.2 because they are used directly in the thermal-resistance calculation. Their implementation and interpretation are discussed in the following section.

6.3 Modelling decisions

The P25 simulations use the one-dimensional thermal model described in Section 2. The hot-gas state is calculated using Cantera from the measured chamber pressure, mixture ratio, and reactant inlet temperatures listed in Table 6.2. The fuel is represented as a 90 wt% ethanol and 10 wt% water mixture in the combustion calculation, while the oxidizer is represented by nitrous oxide at the measured inlet temperature. The resulting combustion-gas state is expanded along the chamber contour and used as the gas-side boundary condition for the thermal model.

The coolant properties are evaluated along the channel as the pressure and temperature change, allowing the model to account for the variation in density, viscosity, thermal conductivity, and heat capacity. The measured coolant temperature rise and pressure loss are therefore used as integral checks on the predicted coolant state and hydraulic behaviour.

The additively manufactured cooling channels are not expected to behave as hydraulically smooth channels. The cooling-channel walls were left in their as-built condition, and the local channel dimensions may deviate from the nominal CAD geometry due to manufacturing tolerances, powder adhesion, and partially fused surface features. To account for this, the P25 case is evaluated using both a baseline forced-convection model and a roughness-corrected forced-convection model using the roughness model developed in Section 5. The roughness model increases the coolant-side friction factor and heat-transfer coefficient through the assumed equivalent surface roughness. In addition, a reduced effective channel area is considered to represent possible constrictions in the as-built channels. This affects the coolant velocity, pressure loss, and convective heat transfer.

A TBC is included between the hot gas and the chamber wall. For P25, the coating is represented by an assumed thickness of 40 μm and a thermal conductivity of 1.2 W/(m K), as listed in Table 6.2. The basis for these coating assumptions and the difficulty of predicting the in-situ coating thickness are discussed in Section 7.5. The coating is therefore not modelled as a separate structural layer with temperature-dependent properties, but only through its thermal resistance. This captures the first-order effect on wall heat flux, while neglecting local thickness variation, cracking, spallation, and interface resistance.

The final model variant includes nucleate boiling on the coolant side. This is motivated by the single-phase solutions, in which the predicted coolant-side wall temperature remains above the local coolant saturation temperature over substantial parts of the chamber. Under these conditions, the coolant-side heat transfer is unlikely to be represented fully by single-phase forced convection. The nucleate-boiling model therefore provides a physically motivated sensitivity case for estimating the effect of enhanced coolant-side heat transfer on the predicted wall temperature and heat flux. The boiling model is still a one-dimensional engineering model and does not resolve two-phase flow structure, vapour transport, dry-out, or local flow maldistribution. The nucleate-boiling result should therefore be interpreted as a sensitivity study rather than a complete two-phase cooling prediction.

The three model variants considered for P25 are therefore: a baseline forced-convection model, a forced-convection model with roughness and reduced effective channel area, and a nucleate-boiling model applied to the same roughness-affected channel geometry. These variants are compared using the measured coolant pressure loss and outlet temperature from the flight data. Since no local wall-temperature or heat-flux measurements are available, agreement with these global coolant-side quantities does not uniquely validate the predicted wall-temperature distribution. Instead, the comparison is used to assess whether the model variants produce physically plausible coolant states and thermal margins for the P25 chamber.

The P25 calculations are performed using 300 axial stations along the cooling channel. This is a relatively fine discretisation for the present one-dimensional model. For comparison, the RL10 cooling-jacket model of Binder et al. [7] used 20 axial stations for an engine with a total length of approximately 1.4 m, while Hovik [37] found that the peak heat flux in an RL10 regenerative-cooling calculation was already reproduced with reasonable accuracy using 10 or more axial nodes. Given the considerably shorter length of the P25 chamber, the 300-station grid is considered sufficient to resolve the axial trends in coolant pressure, coolant temperature, wall temperature, and heat flux. A separate grid-convergence study is therefore not performed for the P25 case.

6.4 Results

The results are shown as steady-state axial distributions at the selected flight operating point. The axial coordinate follows the hot-gas flow direction, while the coolant flows in the opposite direction through the regenerative cooling channels. Consequently, the coolant enters near the nozzle end and exits near the injector end. The coolant pressure and temperature distributions can be compared directly with the measured inlet and outlet values from the flight data, whereas the heat flux and wall-temperature distributions are model predictions without local experimental validation.

The same axial discretization, chamber contour, channel geometry, and gas-side heat transfer model are used for all P25 cases, so the differences between the curves reflect changes in the coolant-side modelling assumptions rather than changes in the numerical grid.

The predicted coolant pressure and temperature distributions are shown in Figure 6.8. The pressure distribution is primarily governed by the channel geometry, mass flow rate, and friction model, and is therefore only weakly affected by whether nucleate boiling is included in the heat-transfer calculation. In contrast, the coolant temperature is more sensitive to the heat-transfer model because it depends directly on the heat absorbed from the chamber wall. The nucleate-boiling case predicts a higher coolant temperature rise, indicating increased heat transfer from the wall to the coolant.




Figure 6.8: Comparison of coolant pressure and temperature distributions for the baseline, roughness, and boiling models.

The roughness-corrected cases produce a pressure loss closer to the measured flight value, while the baseline case underpredicts the hydraulic resistance of the cooling channels. This supports the assumption that the as-built additively manufactured channels cannot be represented accurately as hydraulically smooth passages. Significant pressure drop occurs in the region around the throat, where the velocity is the highest, as expected.

Figure 6.9 shows the predicted wall heat flux for the different coolant-side modelling assumptions. The heat flux peaks in the throat region, where the gas-side heat-transfer coefficient and thermal loading are highest.




Figure 6.9: Predicted heat flux along the chamber wall for the different models.

The increase in heat flux in the boiling case does not imply a higher wall temperature. Rather, the additional coolant-side heat transfer reduces the thermal resistance between the wall and the coolant, allowing more heat to be removed while maintaining a lower wall temperature.

The corresponding gas-side wall-temperature distributions are shown in Figure 6.10. The baseline forced-convection model predicts wall temperatures above the operational temperature limit over a significant part of the chamber. Including roughness and reduced effective channel area improves the cooling performance by increasing coolant velocity and convective heat transfer, but the predicted wall temperature still remains above the limit in the most highly loaded region.





Figure 6.10: Gas-side wall temperature distribution for the different models.

The nucleate-boiling model gives a substantially lower wall-temperature prediction and keeps the wall temperature closer to the intended thermal margin. This indicates that boiling-enhanced heat transfer may be important for interpreting the P25 thermal behaviour. The nucleate-boiling model gives a substantially lower wall-temperature prediction and keeps the wall temperature closer to the intended thermal margin. This indicates that boiling-enhanced heat transfer may be important for interpreting the P25 thermal behaviour and motivates treating the coolant-side heat-transfer closure as a dominant modelling sensitivity.

A crack was observed in the P25 chamber wall after the final ground-test campaign, as shown in Figure 6.11. The crack is included here as an experimental observation relevant to the thermal-margin assessment, but the present steady-state model is not used to predict its formation. Possible contributing mechanisms, including shutdown transients, residual stresses from additive manufacturing, and local film-cooling effects, are discussed in Section 6.5. The chamber was subsequently repaired using high-temperature adhesive, hydrostatically tested beyond working pressure, and later flown on the *Heimdall* vehicle without observed cooling-system failure.



(a) Microscope close-up of the crack area



(b) Crack location after decommissioning

Figure 6.11: P25 chamber wall crack observed after the final ground-test campaign. (a) Close-up image of the crack area before repair. (b) Crack area after decommissioning. The white material around the crack area is residue from the high-temperature adhesive used during repair. The crack morphology is consistent with a high-temperature wall failure mode. Photos by Propulse NTNU.

6.5 Discussion

6.5.1 Pressure model

The P25 flight-test pressure data show that the present hydraulic model underpredicts the measured pressure loss for all model variants. The baseline model, which uses the nominal channel geometry and a constant sand-grain roughness of $\varepsilon = 20 \mu\text{m}$, gives the lowest pressure loss. Introducing the constricted-area formulation increases the predicted pressure loss and therefore moves the result in the correct direction, but the discrepancy remains large. Including the roughness model together with the constricted-area formulation further increases the predicted loss, but still does not reproduce the measured flight value. The model changes therefore improve the pressure-loss prediction directionally and show that effective area reduction and roughness are important for interpreting the P25 hydraulic behaviour.

The axial pressure trend is nevertheless physically reasonable. The largest calculated pressure gradient occurs near the throat, where the coolant mass flux is highest. In the straighter and larger-area sections of the cooling channel, the predicted pressure gradient is comparatively small. This indicates that the model captures the dominant geometric trend of the cooling-channel pressure loss, even though the total pressure drop is underpredicted.

The underprediction is also consistent with the separate water-flow comparison presented previously. In that case, the same modelling changes increased the predicted pressure loss: the baseline

model underpredicted the measured pressure loss by almost 50%, and the constricted-area formulation reduced the underprediction to approximately 40%, and the combination of roughness and constricted area reduced it to approximately 10%. However, the water-flow specimen was manufactured with improved process parameters and should not be interpreted as a direct validation specimen for the P25 flight chamber. The result is therefore used only as supporting evidence that the current modelling assumptions tend to underpredict hydraulic losses unless effective area reduction and roughness effects are included.

A likely contributor to the remaining discrepancy is a mismatch between the modelled and measured pressure boundaries. The model represents the pressure loss through the cooling channels, whereas the flight-test pressure measurement may include additional losses upstream of the channels. The inlet pressure sensor is located before the inlet tube and distribution manifold, so the measured pressure drop can include losses in the feed hardware, the manifold, the contraction into each channel, and the square-edged channel entrances. The inlet tube is relatively wide compared with the coolant flow rate and is therefore expected to contribute less than the manifold and entrance regions. The manifold, however, has a varying flow area as coolant is distributed to the individual channels, and each channel inlet introduces a local contraction. These losses are not included in the present one-dimensional channel model and may explain part of the systematic underprediction.

The weak dependence of the pressure distribution on the heat-transfer model is also physically reasonable for this case. The P25 coolant remains liquid, and the density variation along the channel is modest compared with the dominant effects of geometry, mass flux, hydraulic diameter, and friction factor. The nucleate-boiling model modifies the coolant-side heat-transfer coefficient, but the hydraulic model does not introduce void fraction, mixture density, acceleration pressure loss, or two-phase friction. As a result, changing from forced convection to nucleate boiling can strongly affect wall temperature and coolant temperature rise while having only a small effect on the predicted pressure loss.

There is also significant uncertainty in the effective hydraulic roughness. Roughness measurements from additively manufactured surfaces are not directly equivalent to the equivalent sand-grain roughness used in conventional pipe-flow correlations. In addition, roughness measurements from coupons or channel specimens manufactured with different process parameters may not represent the hydraulic roughness of the actual P25 chamber. Partially fused powder, local waviness, local overhang effects, and small geometric constrictions can increase the effective hydraulic resistance without being fully captured by the present roughness model.

The one-dimensional model further assumes equal flow distribution between all cooling channels. In the real chamber, additive-manufacturing tolerances, local constrictions, and the manifold geometry may cause channel-to-channel maldistribution. If some channels have higher resistance or smaller effective flow area than assumed, the total system pressure loss can increase relative to the idealized equal-flow model. This effect cannot be captured by the present single-channel-equivalent formulation.

Overall, the present pressure model should be described as a channel-only hydraulic model that underpredicts the measured P25 system pressure loss. The discrepancy is reduced when effective area reduction and roughness are included, but the remaining error indicates that additional hydraulic effects are missing. The most likely contributors are manifold losses, channel entrance losses, uncertainty in the effective roughness, geometric deviations from the nominal channel dimensions, and possible channel-to-channel maldistribution. A useful diagnostic extension would be to determine what equivalent roughness or effective area reduction would be required for the channel-only model to match the measured flight pressure loss. If the required value is physically unreasonable, this would support the interpretation that omitted system-level losses dominate the remaining discrepancy.

6.5.2 Thermal model

The thermal-model results are used to assess physical consistency and sensitivity to coolant-side heat-transfer closure. The main distinction between the forced-convection and nucleate-boiling cases is the treatment of the coolant-side wall condition. In the forced-convection cases, the predicted coolant-side wall temperature exceeds the local saturation temperature of the coolant over a substantial part of the chamber. This indicates that a purely single-phase forced-convection model is unlikely to be appropriate. The Gnielinski correlation used in the forced-convection model is intended for single-phase internal flow and does not account for vapour formation, bubble dynamics, or the heat-transfer enhancement associated with nucleate boiling. The nucleate-boiling model is therefore better motivated physically when the wall temperature is above the local saturation condition.

The wall-superheat distribution in Fig. 6.12 provides a useful check on whether the nucleate-boiling correction is physically motivated. Along the wall, the coolant-side wall temperature is above the local saturation temperature, so $T_{cw} - T_{\text{sat}}$ remains positive. The largest values occur in the throat region, where the gas-side heat load is highest and the coolant channel area is smallest. Rocket-engine cooling literature gives typical wall-superheat values of approximately 10–50 K for nucleate boiling [38]. The predicted P25 values are therefore of the same order, which supports using a boiling heat-transfer correction in the region where the thermal loading is most severe.




Figure 6.12: Wall superheat ($T_{cw} - T_{\text{sat}}$) for the P25 cooling channel using the nucleate-boiling model. The saturation temperature is evaluated at the local coolant pressure.

A high wall superheat also introduces an additional uncertainty. The present boiling model treats boiling as a nucleate-boiling heat-transfer enhancement and does not include a critical heat-flux criterion, dryout, transition boiling, or film boiling. If the local wall heat flux approaches or exceeds the critical heat flux, the nucleate-boiling model may overpredict the coolant-side heat-transfer coefficient and underpredict the wall temperature. The nucleate-boiling result therefore improves the coolant-side heat-transfer representation in regions where the wall exceeds the saturation temperature, while the absence of a critical-heat-flux criterion remains an important uncertainty.

The heat-flux distributions are similar in shape for the different thermal models, with the highest heat flux occurring near the throat. The nucleate-boiling model gives a peak heat flux approximately 10% higher than the forced-convection models, while still predicting a substantially lower gas-side wall temperature. This is physically consistent: nucleate boiling increases the coolant-side heat-transfer coefficient, which reduces the wall-to-coolant temperature difference required to transfer heat. The result is that a modest increase in heat flux can occur together with a large reduction in wall temperature. The heat-flux comparison should therefore not be interpreted only in terms of magnitude; it also explains why the wall-temperature prediction is highly sensitive to the coolant-side heat-transfer model.

The gas-side wall-temperature results show the strongest model dependence. The forced-convection

models predict significantly higher wall temperatures than the nucleate-boiling model, particularly near the throat. Only the nucleate-boiling model predicts a maximum wall temperature below the design temperature limit. This is an important result because reduced-order thermal models are often used as a basis for deciding whether a chamber design should proceed to manufacturing and testing. For P25, the design assessment would therefore depend strongly on whether single-phase forced convection or nucleate boiling is considered the more appropriate coolant-side model.

The forced-convection cases also predict a secondary increase in gas-side wall temperature near the nozzle end, while this feature is not present in the nucleate-boiling result. This does not necessarily indicate a numerical error. A high wall temperature can occur even when the local gas-side heat load is moderate if the coolant-side thermal resistance is large. The physical mechanisms in the two models are different. In the forced-convection model, the coolant-side heat-transfer coefficient is driven primarily by Reynolds number, hydraulic diameter, and coolant thermal properties. If the local channel area increases and the coolant velocity decreases, the forced-convection heat-transfer coefficient can become small even though the local heat flux is moderate. In the nucleate-boiling model, part of the heat transfer is instead driven by wall superheat relative to the local saturation temperature. The nucleate-boiling contribution can therefore remain significant in lower-velocity regions, and the suppression factor becomes less restrictive as the liquid Reynolds number decreases. This provides a plausible explanation for why the forced-convection models develop a secondary nozzle-end temperature rise while the nucleate-boiling model does not.

It is also important that agreement with the measured coolant outlet temperature cannot uniquely validate the predicted wall-temperature distribution. The coolant outlet temperature depends primarily on the integrated heat input to the coolant. Different axial heat-flux distributions can produce the same total coolant temperature rise while giving different local wall temperatures. The wall-temperature prediction is therefore more sensitive to the local heat-transfer model than the coolant outlet temperature alone can reveal.

Overall, the thermal comparison indicates that the coolant-side heat-transfer model is one of the dominant uncertainties in the P25 analysis. The forced-convection models are inconsistent with their own predicted wall state because they predict coolant-side wall temperatures above the local saturation condition while using single-phase correlations. The nucleate-boiling model is more physically appropriate for this regime and gives wall temperatures consistent with the design limit, but it also remains incomplete because critical heat flux, dryout, and post-CHF heat transfer are not modelled. The thermal results therefore identify boiling heat transfer as a likely controlling effect in the P25 wall-temperature prediction, while also showing that the available coolant measurements constrain the integrated heat input more strongly than the local wall-temperature field.

6.5.3 Observed chamber-wall crack

Following ground testing, a crack was observed in the P25 chamber wall, as shown in Figure 6.11. The crack is not interpreted here as a direct prediction made by the one-dimensional thermal model. Instead, it is used as an indication that thermal and mechanical margins must be considered together, rather than assessing the chamber only from the peak steady-state wall temperature.

The most plausible explanation for the crack is a combination of reduced material margin in the as-built LPBF IN625 chamber and an off-nominal shutdown transient. LPBF parts can retain significant thermally induced residual stresses after manufacture, and stress-relief heat treatment is commonly used to reduce residual stress and associated distortion or cracking susceptibility in IN625 components [39, 40]. Since the P25 chamber was used without post-build stress relief or solution heat treatment, residual manufacturing stresses may have been superposed on the firing-induced thermal stresses. This does not prove that residual stress initiated the crack, but it reduces the available margin during rapid thermal transients.

The reported ox-rich shutdown condition is also important because it is outside the steady-state flight operating point analysed in this chapter. During shutdown, the propellant mixture ratio, wall heat flux, film-cooling effectiveness, and coolant-side boundary conditions may change rapidly while the wall temperature field is still transient. Large thermal gradients and thermal cycling

are established drivers of crack formation in regeneratively cooled rocket combustion-chamber liners [41, 42]. The observed crack is therefore consistent with a transient thermal-stress failure mode, but the present model cannot determine the crack-initiation location, initiation time, or stress intensity.

A further local contribution may have come from the film-cooling arrangement. The P25 injector introduced film coolant through discrete orifices rather than through a continuous annular slit. This can create local jet impingement and circumferentially non-uniform wall temperatures near the injector end. Such local gradients are not captured by the one-dimensional model, which represents the chamber using axially averaged thermal loads and coolant-side conditions. The crack therefore also highlights the importance of local injector and film-cooling effects when assessing damage mechanisms near the chamber head end.

The subsequent repair and successful flight test provide useful operational context. The absence of observed crack growth during flight indicates that the chamber retained functional integrity during the nominal flight segment after repair. However, this should not be interpreted as proof that the original thermal margin was large. It instead suggests that the repaired chamber was able to tolerate the nominal operating condition, while the original crack formation remains consistent with limited margin during a more severe off-nominal shutdown transient.

6.6 Conclusion

The P25 case provides a useful assessment of the one-dimensional cooling-channel model using the available engine data. The model captures several physically reasonable trends: the pressure gradient is largest near the throat where the coolant mass flux is highest, the heat flux peaks in the throat region, and the predicted wall temperature is strongly dependent on the coolant-side heat-transfer model. The coolant temperature rise also provides a useful global energy-balance check on the calculation. These results show that the model is useful for comparing modelling assumptions and identifying dominant sensitivities in the P25 cooling system.

The pressure-loss prediction remains the clearest discrepancy in the P25 analysis. All model variants underpredict the measured flight-test pressure loss. The baseline model, using nominal channel geometry and a constant sand-grain roughness of $\varepsilon = 20 \mu\text{m}$, gives the lowest pressure loss. The constricted-area formulation and the spatial roughness model both increase the predicted pressure loss in the correct direction, but they do not fully reproduce the measured value. This should be interpreted as a channel-only hydraulic model underpredicting the measured system pressure loss, rather than as proof that the cooling-channel friction formulation alone is invalid. The remaining discrepancy is likely caused by a combination of missing manifold and entrance losses, uncertainty in the effective hydraulic roughness, deviations from the nominal channel geometry, and possible channel-to-channel maldistribution. The separate water-flow comparison supports the importance of roughness and effective area reduction, but it should not be treated as direct validation of the P25 flight chamber because the tested specimen was manufactured with different process parameters.

The thermal results are also model-dependent. The forced-convection models predict coolant-side wall temperatures above the local coolant saturation condition over the entire chamber wall, which is outside the intended application range of single-phase Nusselt-number correlations. The nucleate-boiling model is therefore physically motivated and gives lower wall temperatures, including a maximum gas-side wall temperature below the design limit. However, this result should not be interpreted as a validated thermal margin. The boiling model does not include critical heat flux, dryout, transition boiling, film boiling, or a resolved two-phase pressure-drop model. The thermal margin is therefore bounded by significant uncertainty: forced convection is likely too conservative in regions with wall boiling, while the nucleate-boiling model may be non-conservative if the wall heat flux approaches a post-critical-heat-flux regime.

The observed crack in the P25 chamber should be interpreted in this context. The one-dimensional model does not predict crack initiation and does not include structural stress, residual stress, local injector effects, or shutdown transients. Nevertheless, the crack supports the broader conclusion

that thermal and mechanical margins cannot be assessed from steady-state wall temperature alone. The most plausible explanation is a combination of residual stress in the as-built LPBF chamber and an off-nominal oxygen-rich shutdown transient. The lack of stress-relief heat treatment likely reduced the available margin, while the discrete film-cooling injection pattern may have introduced local temperature gradients not captured by the one-dimensional model. The subsequent repair and successful flight test show that the chamber retained functional integrity during the nominal flight segment, but they do not prove that the original design had large thermal or structural margin under transient off-nominal conditions.

Overall, the P25 study shows that roughness and effective channel area cannot be ignored in additively manufactured cooling channels, boiling can strongly affect the predicted wall temperature, and global inlet–outlet measurements constrain the internal thermal state only weakly. Stronger validation would require additional instrumentation, such as local wall-temperature measurements, heat-flux estimates, or better-separated pressure measurements across the cooling channels alone. These findings directly motivated the P26 design work, including improved post-processing and stress relief, revised film-cooling implementation, and more careful treatment of channel geometry and roughness in the modelling workflow.

7 Design study: Propulse NTNU P26

The P26 engine is developed for the *Fossekall* flight vehicle, a student rocket designed and built by Propulse NTNU. The vehicle is planned for flight in July 2026, which places the engine development on a short design and test cycle with limited time, manufacturing access, instrumentation, and test opportunities. Since the engine has not flown at the time of writing, no flight data are available for P26. This chapter is therefore treated as a design study rather than a validation case: the purpose is to apply the modelling framework to the revised chamber design and evaluate the expected thermal and hydraulic consequences before flight testing.



Figure 7.1: The P26 combustion chamber after printing, stress relieving and machining. Photo by Philip Draget, Propulse NTNU.

P26 is a direct design iteration of P25. The P25 flight demonstrated that the regenerative cooling concept was functional, but the test campaign also exposed limitations relevant to the chamber design. In particular, the liner crack observed after ground testing motivated a redesign aimed at improving tolerance to thermal gradients, residual stress, and off-nominal transients. The main P26 changes are stress relieving after manufacture, a revised injector-side film-cooling arrangement, and bifurcated cooling channels in the upper part of the chamber.

The P26 chamber was stress relieved after printing to reduce residual stresses introduced during LPBF. Martucci et al. [39] showed that stress-relief heat treatment of LPBF IN625 can substantially reduce residual stress, with treatments in the range 750–800 °C providing a reduction in residual stress while limiting undesirable microstructural changes compared with higher-temperature treatment. For P26, this is important because residual manufacturing stresses can be superposed on firing-induced thermal stresses during startup and shutdown transients. The chamber was removed from the build plate before stress relieving due to equipment size limitations. This is less ideal than stress relieving the part while still attached to the build plate, since part removal can release residual stress before the heat treatment and may cause distortion. However, the stress-relief step still represents a significant improvement over the P25 manufacturing route, where no corresponding post-build stress relief was performed.

The cooling-channel layout was revised by introducing channel bifurcation in the upper part of the chamber. As the chamber radius increases upstream of the throat, a constant number of channels leads to wider passages and a larger unsupported hot-wall span between ribs. In P26, each channel is split into two narrower branches where the channel width would otherwise become too large. This reduces the unsupported wall span and increases the structural stiffness of the hot wall, while keeping the total coolant flow area comparable with the P25 design. The intention is therefore to improve the mechanical behaviour of the wall without relying on a large change in coolant-side performance.

The role of the model in this chapter is different from its role in the P25 case. For P25, measured flight data were available and the model could be compared against observed coolant pressure and temperature trends. For P26, the model is instead used during design iteration to assess whether the revised chamber architecture gives reasonable predicted pressure loss, coolant temperature rise, heat flux, and wall temperature. The following sections describe the P26 hot-wall geometry, operating conditions, cooling-channel design, and predicted thermal and hydraulic behaviour.

7.1 Hot-wall geometry

The P26 thrust chamber was designed for a nominal sea-level thrust of 8 kN at a chamber pressure of 21.013 bar absolute, corresponding to 20 bar gauge at sea level. The hot-wall contour is a parabolic bell nozzle. This gives a shorter nozzle than a comparable conical contour while retaining a relatively efficient expansion and remaining within the available LPBF build height. In contrast to P25, the injector is recessed into the chamber assembly, which limits the available axial length for the combustor and nozzle to approximately 275 mm instead of the approximately 295 mm available in the P25 design. The principal hot-wall geometry parameters are listed in Table 7.1, and the corresponding radius distribution is shown in Figure 7.2.

Table 7.1: Principal P26 hot-wall geometry parameters.

Parameter	Symbol	Value
Chamber radius	R_c	67.0 mm
Throat radius	R_t	30.0 mm
Nozzle exit radius	R_e	55.0 mm
Combustor length	L_c	200.0 mm
Nozzle length	L_e	74.6 mm
Contraction angle	b	30.0°
Nozzle entry angle	θ_n	20.8°
Nozzle exit angle	θ_e	14.0°
Contraction radius	R_2	115.63 mm
Throat entry radius	R_1	$1.5R_t$
Throat exit radius	R_n	$0.382R_t$
Contraction ratio	A_c/A_t	4.99
Expansion ratio	A_e/A_t	3.36
Characteristic length	L^*	794 mm




Figure 7.2: P26 hot-wall radius distribution used in the one-dimensional thermal model.

7.2 Bifurcated cooling channels

The P25 analysis showed that the predicted wall temperature is highly sensitive to the coolant-side heat-transfer model and to the local cooling-channel geometry. In particular, wide channels increase the unsupported span of the hot wall between adjacent ribs and reduce the channel aspect ratio. This is undesirable both thermally and structurally: a larger wall span is more compliant under pressure and thermal loading, while a lower aspect ratio generally gives a less favourable distribution of coolant-side heat transfer for the same wall thickness and flow area.




Figure 7.3: Effect of cooling-channel aspect ratio on hot-wall structural response. Increasing aspect ratio reduces the unsupported hot-wall span and improves the thermal and mechanical behaviour under cyclic loading. Reproduced from NASA-TM-105679 [43]; public use permitted.

The benefit of high-aspect-ratio passages has been demonstrated experimentally in rocket combustion chambers. Carlile and Quentmeyer [43] tested low- and high-aspect-ratio cooling passages and found that the high-aspect-ratio configuration gave substantially lower hot-gas-side wall temperature for the same coolant pressure drop. The high-aspect-ratio chamber also showed improved cyclic durability, with no observed fatigue damage after 440 thermal cycles [43]. These results do not define the P26 geometry directly, but they motivate the same design principle: reducing the channel width relative to its height can improve both cooling effectiveness and hot-wall structural response.

A practical difficulty is that the local chamber circumference changes strongly along the engine. If the number of channels is kept constant, the channels become wider in the large-radius chamber

section. If the number of channels is increased everywhere, the hydraulic penalty and manufacturing complexity increase throughout the full cooling circuit. Bifurcation provides an intermediate solution. In this context, bifurcation means that one regenerative cooling passage divides into two narrower passages over a short axial distance [4]. The purpose is not primarily to increase the total coolant flow area, but to redistribute approximately the same flow area into narrower channels where the chamber circumference would otherwise force the passages to become too wide.




Figure 7.4: Historical examples of non-bifurcated and bifurcated cooling-channel concepts used in liquid rocket combustion chambers. Reproduced from NASA-SP-8087 [4]; public use permitted.

The concept is not new. NASA-SP-8087 documents bifurcated coolant-tube geometries and discusses bifurcation joints as a practical solution for adapting the cooling-passage layout to changes in chamber circumference [4]. Although the manufacturing route was different from that of the additively manufactured P26 chamber, the underlying geometric problem is the same. A single wide passage can be replaced by two narrower passages in order to maintain a more favourable local channel aspect ratio and reduce the unsupported hot-wall span.

Wadel later compared continuous, bifurcated, and stepped high-aspect-ratio channel layouts analytically [44]. The study found that high-aspect-ratio cooling could reduce hot-gas-side wall temperature with only a moderate pressure-drop penalty, and that the bifurcated configuration gave the best overall compromise between wall-temperature reduction and pressure loss. After fabrication constraints were included, the bifurcated design retained an approximately 20% reduction in

hot-gas-side wall temperature with a pressure-drop increase of about 9% [44]. This is especially relevant for P26, since the design objective is not simply to maximise the number of channels, but to obtain the thermal and structural benefit of narrower passages only where they are needed.

A recent additively manufactured example with closer geometric relevance is the GRCop-42 LOX/LCH₄ combustion-chamber demonstrator presented by Waugh et al. [45]. The demonstrator uses an additively manufactured GRCop-42 liner inside an IN718 structural jacket, and includes 88 coolant channels with bifurcations in the nozzle contraction and expansion regions. Trial builds of the throat and bifurcation regions were also produced to assess the manufacturability of the narrow channel features. This example is particularly relevant to P26 because the bifurcation is used for the same general purpose: adapting the local channel count to a changing hot-wall radius while avoiding excessively wide coolant passages in regions where a fixed channel count would give an unfavourable channel aspect ratio.

For P26, bifurcation is therefore introduced as a deliberate design measure in the upper part of the chamber. The channel count is increased only where the chamber radius becomes large enough that the primary channels would otherwise become too wide. The following section describes how this principle was applied to the final P26 cooling-channel layout.

7.3 P26 cooling-channel design

The P26 cooling-channel layout was constrained by both manufacturability and cooling performance. Narrow channels are desirable because they reduce the unsupported hot-wall span between ribs and help maintain a favourable channel aspect ratio. However, reducing the channel width also increases the risk of incomplete powder removal, local constriction, roughness sensitivity, and excessive pressure loss. The minimum channel width was therefore treated as a manufacturing and hydraulic design constraint, rather than as a purely thermal optimum.

Earlier designs used minimum channel widths of approximately 1.2 mm or larger, but these limits were based on older manufacturing assumptions. With the updated process parameters used for P26, $w_{\min} = 1.0$ mm was accepted as the minimum design channel width. This value should not be interpreted as an optimized hydraulic width, but as the smallest tested width that could be printed, cleared, and flow-tested without anomalous behaviour.

The primary channel count was then set by the throat region, where the available circumference is smallest. Both the hot-wall thickness and the rib width were set to 0.7 mm, which was treated as the practical lower limit for the selected manufacturing process. Using the throat diameter, hot-wall thickness, rib width, and accepted minimum channel width, the largest allowable number of primary channels in the non-bifurcated throat region is

$$N = \left\lfloor \frac{2\pi(R_t + t_w)}{w_{\min} + b_{\min}} \right\rfloor = 113. \quad (59)$$

The final design therefore uses 113 primary channels. This gives a throat-region channel width consistent with the accepted 1.0 mm minimum, while leaving the bifurcations to control the channel width further downstream where the chamber circumference increases.

As the chamber radius increases upstream of the throat, the circumferential arc length available to each channel also increases. If the 113 primary channels were continued unchanged through the full chamber, the channels would therefore become progressively wider in the large-radius chamber section. This would reduce the local aspect ratio and increase the unsupported hot-wall span between ribs. The bifurcation criterion was defined from the width required to replace one parent channel by two minimum-width daughter channels separated by a new rib. With a minimum channel width of $w_{\min} = 1.0$ mm and a rib width of $b_{\min} = 0.7$ mm, this gives a bifurcation trigger width of 2.7 mm. For the selected primary channel count, this width is reached at a local hot-wall radius of 60.45 mm. This occurs only in the chamber-side part of the cooling passage. The nozzle exit radius is 55 mm, so the same criterion is not reached in the nozzle section. The final P26 design therefore bifurcates only upstream of the throat, where the chamber radius becomes large

enough that the primary channels would otherwise exceed the chosen width limit.

The channel-height distribution was generated from the P25 per-channel flow area distribution. In the 113-channel section of P26, each primary channel was therefore assigned approximately the same local flow area as one P25 channel. Since P25 used 80 channels, the total P26 flow area in the non-bifurcated section becomes approximately $113/80$ times the P25 total flow area. A strict constant-total-area redesign would instead have required scaling the P25 per-channel area by $80/113$ before applying it to P26, but this scaling was not applied in the final manufactured geometry.

In the bifurcated section, each of the 113 primary channels is split into two daughter channels. To avoid a discontinuous doubling of the total flow area, each daughter channel was assigned approximately half of the corresponding primary-channel area. The 226-channel section therefore retains approximately the same total flow area as the 113-channel section at the same axial location, while redistributing the flow into narrower channels. The resulting area distribution should therefore be interpreted as the final modelled P26 geometry, not as a constant-area transformation of P25. This distinction is important because the larger total flow area reduces coolant velocity and pressure loss, but may also reduce the local forced-convection heat-transfer coefficient.




Figure 7.5: Total cooling-channel flow area for the two segments in the P26 engine.

The splitter is introduced with a 30° slope to reduce the likelihood of a local stagnation region and to remain compatible with the overhang constraints of the additive-manufacturing process. The resulting splitter and the channel turning region are shown in Figure 7.6.




Figure 7.6: Internal cooling channel geometry of the P26 combustion chamber. The main view shows the full channel distribution, while detailed views highlight the channel turning region and the bifurcation feature. Note that the metal-volume is hidden; only the fluid-volume is visible. Photos by Jan Vik, Propulse NTNU.

7.4 Operating conditions

The thermal analysis was carried out for a coolant consisting of 90 wt% ethanol and 10 wt% water. The regenerative cooling-channel mass flow rate was set to $\dot{m}_c = 1.1 \text{ kg/s}$, with a channel inlet pressure of 28 bar and an inlet temperature of -20°C . The gas-side calculation uses a global mixture ratio of $\text{O/F} = 3.0$. These values define the baseline operating point used in the P26 design study. The analysis does not attempt to validate the local thermal field, since no flight or hot-fire thermal data are available for P26 at the time of writing. Instead, the conditions are used to evaluate the expected coolant pressure loss, coolant temperature rise, heat flux, and wall-temperature trends for the revised chamber geometry.

Table 7.2: P26 operating and simulation parameters used in the thermal analysis.

Parameter	Value
Chamber pressure	21.013 bar absolute
Oxidizer-to-fuel ratio	3.0
Coolant composition	90 wt% ethanol, 10 wt% water
Coolant mass flow rate	1.1 kg/s
Coolant inlet pressure	28 bar
Coolant inlet temperature	-20°C
Oxidizer inlet temperature	-25°C
Assumed TBC thickness	10 μm
TBC conductivity	1.2 W/(mK)

The thermal-barrier-coating thickness used for P26 should be interpreted as an assumed lower-bound modelling value rather than as a measured coating thickness. The coating is expected to form in situ during operation, and its final thickness and uniformity are therefore difficult to prescribe before testing. Previous internal deposition modelling of similar in-situ silica-based coatings indicated plausible layer thicknesses ranging from the order of 10 μm to several hundred micrometres, with strong dependence on the assumed deposition conditions and injector-driven near-wall flow field [46]. This is also consistent with experimental observations from Copenhagen Suborbitals, where TEOS-derived silica deposition was observed but not as a uniform, well-controlled coating [2]. Since no post-test thickness measurement is available for P26 at the time of writing, a thickness of 10 μm is used as a conservative baseline value. This avoids relying on an optimistic level of thermal protection while still representing the additional thermal resistance expected from an in-situ coating.

7.5 Modelling decisions

The P26 simulations use the one-dimensional thermal model described in Section 2. Since no flight data are available, the gas-side boundary condition is calculated from the design operating point in Table 7.2. The hot-gas state is obtained with Cantera using nitrous oxide and a 90 wt% ethanol, 10 wt% water fuel mixture, and is then expanded along the chamber contour for use in the thermal model.

The coolant properties are evaluated locally as the pressure and temperature change along the channel. The predicted coolant outlet temperature and pressure loss are therefore treated as design outputs rather than validation quantities. They are used to assess whether the revised chamber gives reasonable hydraulic and thermal behaviour under the assumed operating conditions.

The P26 hydraulic geometry is modelled using the overhang-angle-dependent roughness model from Section 5 together with the reduced effective-area formulation. This is the intended application of the roughness model, which was developed to represent the increased roughness and partial constriction associated with as-built LPBF cooling channels. Unlike the P25 comparison, where a constant-roughness model was retained as a reference case, the P26 design study uses the angle-dependent roughness and constricted-area formulation as the default hydraulic representation.

A thermal-barrier-coating resistance is included between the hot gas and the chamber wall. The coating thickness is set to 10 μm with a thermal conductivity of 1.2 W/(mK), as listed in Table 7.2. This is an assumed lower-bound modelling value, not a measured P26 coating thickness. Local coating-thickness variation, adhesion, cracking, spallation, and interface resistance are not resolved.

Two thermal model variants are evaluated on the same roughness- and area-corrected channel geometry: forced convection and nucleate boiling. The nucleate-boiling case is included because the P25 analysis indicated that single-phase forced convection can predict coolant-side wall temperatures above the local saturation condition. Since P26 uses the same coolant mixture and operates in the same general thermal design space, nucleate boiling is a physically relevant sensitivity case. It is not a complete two-phase pressure-drop or dry-out model, and does not resolve vapour transport,

void fraction, flow-regime transitions, critical heat flux, or channel-to-channel maldistribution.

The P26 calculations are performed using 300 axial stations along the cooling channel, matching the discretisation used for P25. This is considered sufficient to resolve the axial trends in coolant pressure, coolant temperature, wall temperature, and heat flux over the relatively short P26 chamber. A separate grid-convergence study is therefore not performed.

7.6 Results

The P26 simulations give a thermal and hydraulic response similar to that predicted for P25, which is the intended outcome of the redesign. The cooling system was not redesigned to produce a fundamentally different operating point, but to retain comparable cooling behaviour while improving the local wall geometry and manufacturability of the chamber. The results therefore primarily indicate whether the revised P26 geometry remains in the same hydraulic and thermal design space as P25.

The only experimental comparison available for P26 at the time of writing is the water flow test of the printed cooling-channel test geometry. As shown in Table 7.3, the angle-dependent roughness model with the constricted-area formulation predicts the measured pressure drop to within 3.9%. This supports the use of the roughness and effective-area formulation as the default hydraulic representation of the printed P26 channels.

Table 7.3: Comparison of predicted and measured pressure drop for the P26 cooling-channel water flow-test geometry. The Reynolds number in the throat region is $1.4 \cdot 10^4$.

Model	$\dot{m}$ [g/s]	Δp_{model} [bar]	Δp_{exp} [bar]	Error [%]
Angle-dependent roughness + constricted area	20.9	1.99	1.91	3.9

The predicted coolant pressure and temperature for the P26 design point are shown in Figure 7.7. The forced-convection and nucleate-boiling cases give nearly identical pressure distributions because both use the same roughness- and area-corrected channel geometry. The heat-transfer formulation only affects the pressure loss indirectly through changes in coolant properties, and these effects are small compared with the geometric and frictional contributions. The coolant temperature rise is also similar between the two cases, indicating that the total heat absorbed by the coolant is not strongly changed by the boiling model.




Figure 7.7: Predicted coolant pressure and temperature for the forced-convection and nucleate-boiling models.

The predicted heat flux is shown in Figure 7.8. Both cases give the expected axial trend, with the maximum heat flux located in the throat region. The nucleate-boiling model gives a higher peak heat flux than the forced-convection model because the enhanced coolant-side heat transfer reduces the wall-side thermal resistance and allows more heat to be transferred into the coolant. This is consistent with the P25 results: the boiling model mainly changes the local wall-temperature distribution, not the global coolant energy balance.




Figure 7.8: Predicted heat flux along the chamber wall for the forced-convection and nucleate-boiling models.

The gas-side wall-temperature distribution is shown in Figure 7.9. As in the P25 case, the forced-convection model predicts the highest wall temperatures because the coolant-side heat transfer is limited to single-phase convection. The nucleate-boiling model reduces the wall temperature in the high-heat-flux region, but still predicts a peak temperature slightly above the design limit. This exceedance identifies the throat region as the critical thermal-margin location for future hot-fire evaluation, especially given the sensitivity to coating thickness, surface condition, and boiling behaviour.




Figure 7.9: Gas-side wall-temperature distribution for the forced-convection and nucleate-boiling models.

The forced-convection model also predicts a secondary wall-temperature increase near the nozzle exit, similar to the behaviour observed for P25. Since the heat flux in this region is much lower than at the throat, this feature is interpreted as a coolant-side resistance effect rather than as evidence of a high gas-side heat load in the nozzle.

Overall, the P26 results support the design objective of retaining a P25-like thermal and hydraulic response while introducing a more favourable cooling-channel layout. The water-flow comparison gives confidence in the hydraulic treatment of the printed channels, while the hot-fire predictions show that the thermal margin remains controlled by the coolant-side heat-transfer model and the assumed thermal-barrier-coating resistance.

7.7 Discussion

The main engineering result from the P26 study is that the bifurcated channel layout retains a P25-like thermal and hydraulic response while improving the local channel geometry in the large-radius chamber section. The bifurcation redistributes the available flow area into narrower passages, giving a higher local aspect ratio and a smaller unsupported hot-wall span without forcing the entire chamber to use excessively narrow channels. The water-flow comparison provides the strongest quantitative support for the hydraulic representation, while the thermal analysis identifies the throat region, coating thickness, and coolant-side boiling behaviour as the main quantities to assess during hot-fire testing.

7.8 Conclusion

The P26 case demonstrates the use of the modelling framework during design evaluation of a revised regenerative-cooling chamber. The redesigned cooling system uses 113 primary channels and one bifurcation in the chamber-side region to maintain acceptable channel widths and aspect ratios as the chamber radius increases. This gives a more favourable local channel layout than a non-bifurcated continuation of the same primary-channel count, while retaining a thermal and hydraulic response similar to the P25 cooling concept.

The water-flow comparison provides the strongest quantitative support for the P26 hydraulic model. The angle-dependent roughness model with the constricted-area formulation predicts the measured pressure drop to within 3.9%, supporting its use for the manufactured P26 channels. The hot-fire thermal simulations identify the throat as the critical thermal-margin location and show that the predicted wall temperature is mainly controlled by the assumed thermal-barrier-coating resistance and coolant-side boiling behaviour. Further confidence in the final thermal margin would require hot-fire pressure and temperature data, improved knowledge of the coating thickness, and local wall-temperature measurements.

8 Discussion

The purpose of this thesis has been to develop, assess, and apply a reduced-order model for regenerative cooling analysis of liquid rocket engine thrust chambers. The model has been used both as a numerical framework for studying heat transfer and pressure loss, and as an engineering tool for evaluating real cooling-channel geometries used in Propulse NTNU engines. The following discussion considers the model in that context. First, the intended scope of the model and the validation basis are summarized. The thermal and hydraulic performance are then discussed, followed by the implications of the surface-roughness model and the lessons from the P25 and P26 case studies. Finally, the main limitations and suggested improvements are considered.

8.1 Model purpose and validation basis

The model developed in this work connects gas-state evaluation, gas-side heat transfer, wall thermal resistance, coolant heating, coolant pressure loss, material properties, and additive-manufacturing surface roughness in a single regenerative-cooling calculation. This allows the thermal and hydraulic behaviour of a thrust chamber to be evaluated along the coolant flow path and makes the coupling between geometry, heat transfer, pressure loss, and material selection directly visible during design iteration. The model is therefore best assessed by whether it captures the coupled trends that govern regenerative-cooling design: heat flux, wall temperature, coolant temperature, coolant pressure, and pressure loss.

This modelling strategy is consistent with the role of reduced-order and correlation-based methods in rocket thrust-chamber design, where empirical thermal and hydraulic models are commonly used before more detailed analysis is performed [4, 6, 8]. A fully resolved treatment of the gas flow, wall conduction, coolant flow, local rib geometry, channel curvature, boiling structure, and boundary-layer development would require substantially more geometry preparation, meshing, turbulence modelling, and case-specific boundary conditions. Such methods are valuable for detailed verification, while the present model prioritizes computational speed, modularity, and interpretability for rapid comparison of cooling-channel concepts, as described in Sections 2 and 3.

The confidence in the model comes from several levels of comparison rather than from a single validation case. The gas-state model was compared with established equilibrium calculations, the coolant-side heat-transfer implementation was compared with forced-convection reference data, and the full cooling-channel solver was benchmarked against the RL10 cooling-jacket model of Binder et al. [7]. The model was then applied to the P25 and P26 engines in Sections 6 and 7, where available hydraulic and hot-fire data were used to assess whether the predicted trends were physically reasonable. These comparisons do not all provide the same level of validation. Literature benchmarks and controlled forced-convection cases test isolated submodels, the RL10 case tests the integrated cooling solver against a known engine model, while P25 and P26 test the usefulness of the model for real additively manufactured student engines with limited measurement data.

This validation hierarchy is important. The results show that the model can reproduce relevant first-order behaviour, but they do not imply that every local wall temperature or local heat flux is known accurately. In particular, the P25 and P26 comparisons are limited by the available measurements. Coolant inlet and outlet conditions, or total pressure drop, are valuable quantities, but they are integral results. They do not uniquely determine the axial heat-flux distribution, the local gas-side wall temperature, or the coolant-side wall temperature. Many different internal heat-flux distributions can produce similar total coolant temperature rises. The model is therefore most useful as a tool for engineering judgement, design comparison, and identification of the measurements or modelling assumptions that most strongly affect the predicted thermal margin.

8.2 Performance of the reduced-order thermal model

The thermal results obtained with the model are broadly consistent with the expected behaviour of regeneratively cooled thrust chambers. The peak heat flux and peak wall temperatures occur

near the throat region, where the gas-side heat-transfer coefficient and thermal loading are largest. The coolant temperature increases along the channel as heat is absorbed from the wall, and the coolant pressure decreases due to frictional losses. These qualitative trends are important, since a reduced-order model that fails to capture the correct spatial behaviour would not be useful for cooling-channel design.

The RL10 benchmark in Section 4.3 provides the strongest integrated comparison against an external cooling-jacket model. The present solver reproduces the general behaviour of the reference case, including the location of the main thermal loading and the overall coolant heating trend. However, the gas-side wall temperature is not reproduced with the same accuracy as the integral coolant quantities. As summarized in Table 4.6 and shown in Fig. 4.9, the present model predicts a lower peak gas-side wall temperature than the reference model, with the largest differences occurring in the nozzle region. The RL10 case therefore supports the integrated modelling approach for overall thermal and hydraulic behaviour, while also showing that local wall-temperature predictions are particularly sensitive to geometry reconstruction, assumed channel dimensions, and heat-transfer closure [7, 25].

The P25 and P26 cases test the model against the types of data currently available from Propulse NTNU chambers, rather than against local wall-temperature or heat-flux measurements. For P25, the available pressure and temperature measurements constrain the global cooling-system response, including coolant pressure loss and overall thermal behaviour, but they do not provide local gas-side wall temperatures, coolant-side wall temperatures, or axial heat-flux measurements. For P26, the water-flow pressure-drop comparison in Table 7.3 provides a controlled assessment of the hydraulic model for the manufactured channel geometry, while the hot-fire thermal predictions remain dependent on assumptions regarding coating thickness, boiling behaviour, gas-side heat transfer, and local wall temperature. The two cases therefore support the use of the model for design interpretation and uncertainty identification, while also showing that future validation would benefit most from local wall-temperature or heat-flux measurements.

A key conclusion from these comparisons is that the largest validation gap is not necessarily the reduced dimensionality of the model, but the lack of local experimental data. Inlet and outlet pressure or temperature measurements are useful for constraining the integrated cooling-system response, but they cannot validate the axial distributions of wall temperature and heat flux predicted by the model. Historical rocket-engine cooling studies have used embedded or brazed thermocouples to measure wall or rib-base temperatures [41, 43, 44]. Similar instrumentation, preferably at several axial locations, would provide a much stronger basis for assessing the predicted thermal field in future Propulse NTNU chambers. Additive manufacturing may also make such measurements more accessible by enabling embedded sensor features or dedicated instrumented test articles.

8.3 Hydraulic modelling and surface roughness

The surface-roughness model is one of the most practically useful outcomes of this work. The flow resistance of additively manufactured cooling channels cannot generally be represented well by smooth-wall friction correlations. The internal surfaces produced by laser powder-bed fusion are rough, anisotropic, and dependent on local build orientation. For narrow cooling channels, where the hydraulic diameter is small, even moderate absolute roughness values can become large relative roughness values. This makes surface quality a first-order hydraulic design parameter rather than a secondary manufacturing detail.

The measurements presented in Section 5 show a clear distinction between upskin and downskin surfaces. Upskin surfaces showed comparatively weak dependence on overhang angle, while downskin surfaces showed a much stronger angle dependence. This is physically reasonable, since downskin regions are more sensitive to partially supported melt tracks, powder adhesion, and local overhang effects. A roughness model based only on a single constant roughness value would therefore miss an important feature of the manufactured geometry.

The use of an angle-dependent equivalent sand-grain roughness gives a simple way to include this manufacturing effect in a one-dimensional pressure-loss model. The model is easy to apply: the

local wall orientation is used to estimate a local roughness, which is then used in a rough-wall friction-factor relation. This is a useful engineering compromise. It does not resolve the true three-dimensional roughness topology, but it converts measured LPBF surface data into a form that can be used directly in channel pressure-loss calculations. The approach is consistent with the idea of relating measured roughness to equivalent sand-grain roughness, while still recognizing that this conversion is approximate [35, 36].

The choice of roughness parameter is also important. Peak-to-valley roughness measures are more relevant to hydraulic roughness than average-amplitude measures, because turbulent momentum loss is strongly influenced by protruding asperities and form drag rather than only by the average height of the surface. In the present dataset, the R_z -based fit was also the most physically consistent for the downskin surfaces. This supports the use of R_z as the primary descriptor in the current model. At the same time, the use of R_z as a proxy for equivalent sand-grain roughness remains approximate. Surface-texture parameters and filtering choices are standardized, but the conversion from measured surface texture to hydraulic roughness is not unique [31–34]. The fitted model should therefore be interpreted as a first-order engineering model rather than a universal LPBF roughness law.

The hydraulic results also suggest a broader design implication: cooling-channel performance may be manufacturing-limited as much as geometry-limited. In an idealized design process, pressure loss is controlled by channel hydraulic diameter, length, curvature, and mass flow rate. In an additively manufactured chamber, the achievable surface quality, local overhang angle, and process parameters can shift the actual pressure loss significantly. This means that channel layout, print orientation, process development, and hydraulic performance are coupled design problems.

Surface roughness should not be interpreted only as a negative effect. Increased roughness increases pressure loss, which is usually undesirable because it raises the required feed pressure and reduces pressure margin. However, roughness may also increase the coolant-side heat-transfer coefficient by enhancing turbulence and near-wall mixing. This effect is represented only indirectly in the present model through the friction-factor dependence of the forced-convection correlation, and was not independently validated. The optimum surface quality therefore depends on whether the design is limited primarily by pressure margin or thermal margin. A smoother channel may reduce pressure drop but may also reduce coolant-side heat transfer, while a rougher channel may improve cooling at the cost of hydraulic performance [10, 24].

Several limitations remain. The roughness dataset is limited to a small number of coupons and overhang angles, and the measured surfaces may not fully represent enclosed internal cooling channels. LPBF surface quality can vary with machine, process parameters, powder condition, part orientation, shielding-gas conditions, and post-processing route. The filter settings used for surface-texture evaluation also influence the resulting roughness values. Future work should therefore extend the dataset with more overhang angles, more repeated samples, and more direct measurements of internal channels. Areal roughness parameters such as S_z should also be investigated further, since areal measurements may better represent the three-dimensional nature of LPBF surfaces. Nevertheless, the present roughness model represents a useful step from smooth-wall assumptions toward a manufacturing-aware hydraulic model.

8.4 Engineering lessons from P25 and P26

The P25 and P26 case studies show how the model can be used to reason about regenerative-cooling design choices for additively manufactured geometries. P25 provides the strongest comparison with firing and flight-related data, while P26 demonstrates how the same framework can be used during design evaluation of a newer cooling-channel concept. Together, the cases show how cooling-channel geometry, manufacturing effects, heat transfer, and pressure loss interact in additively manufactured thrust chambers.

For P25, the model provides a way to interpret the observed coolant pressure and temperature behaviour in terms of channel geometry, heat transfer, and pressure loss. The pressure-loss comparison in Section 6 shows that smooth-wall assumptions are insufficient for the printed cooling

channels, and that roughness and effective-flow-area effects must be included to obtain realistic hydraulic predictions. The thermal analysis identifies the expected heat-flux and wall-temperature distributions and shows that the predicted thermal margin is strongly affected by the coolant-side heat-transfer closure. The observed chamber-wall crack discussed in Section 6 further shows that cooling performance cannot be assessed from wall temperature alone. The structural response also depends on pressure loading, thermal gradients, material strength, heat treatment, residual stress from printing, and cyclic loading, all of which are central to the final assessment of chamber reliability [39–42].

The P26 case demonstrates the use of the model during cooling-channel design. The bifurcated channel concept, shown in Fig. 7.6a, improves the local channel aspect ratio and reduces unsupported wall span in the chamber region. This is beneficial because wide channels can increase wall deformation and thermal stress, while bifurcation allows a larger chamber circumference to be cooled without making the individual channels excessively wide. The P26 water-flow comparison in Table 7.3 also shows that the hydraulic behaviour of the manufactured channel geometry can be predicted reasonably well when roughness and effective-flow-area effects are included.

Bifurcation also introduces additional design considerations. The splitter and recombination regions may produce local three-dimensional flow, additional pressure losses, and maldistribution between daughter channels. These effects are represented only indirectly in a one-dimensional channel description through prescribed local geometry and flow area. Future bifurcated designs should therefore be supported by additional hydraulic testing or local three-dimensional analysis, especially if the flow split is expected to be sensitive to manufacturing deviations.

The P25 and P26 designs also raise a more general design question: what quantity should the cooling system be optimized for? A cooling-channel design can be judged by maximum gas-side wall temperature, coolant-side wall temperature, pressure drop, coolant outlet temperature, thermal stress, manufacturability, or robustness to uncertain operating conditions. These objectives are not always aligned. Increasing coolant velocity may improve heat transfer but increase pressure loss. Increasing roughness may improve heat transfer but reduce pressure margin. Adding a thermal barrier coating may reduce metal temperature but also changes the heat available to the coolant. Without an explicit objective function, it is difficult to say whether a design is optimal, even when all predicted temperatures remain within acceptable limits. A future version of the tool should therefore include simple stress and margin estimates so that cooling-channel geometries can be compared using more meaningful engineering comparison metrics.

8.5 Main model limitations

The main limitations of the present model follow from its intended role as a reduced-order design tool. The most important thermal simplification is the representation of the chamber wall as a wall-normal resistance network at each axial station. This neglects axial conduction and does not resolve the two-dimensional temperature field in the wall cross-section. For many regenerative cooling problems, the strongest temperature gradient is through the wall thickness, so this approximation is reasonable as a first-order model. However, local rib geometry, channel curvature, wall-thickness variation, and bifurcation regions can produce cross-sectional temperature variations that are not captured.

It is therefore not obvious that the main limitation of the model is the quasi-one-dimensional fluid treatment. A more likely limitation is the simplified representation of the metal temperature field. Solving the heat equation in the chamber-wall cross-section would be a natural extension while keeping the gas and coolant flow quasi-one-dimensional. A quasi-two-dimensional wall-conduction solver coupled to the present axial marching formulation would improve predictions of rib-base temperature, local thermal gradients, and stress-driving temperature differences without requiring a full three-dimensional conjugate heat-transfer model.

The gas-side heat-transfer model is also empirical. The Bartz correlation is widely used for rocket-engine heat-transfer estimates, but it remains a correlation and cannot fully represent all effects of acceleration, boundary-layer development, wall roughness, finite-rate chemistry, film cooling, or

three-dimensional flow [4, 5]. Other accepted gas-side correlations may give different results for the same geometry and operating condition, which would directly affect the predicted wall-temperature margin. This uncertainty should be treated as part of the model uncertainty, especially when the predicted wall temperature is close to an allowable limit.

The gas-state calculation is similarly simplified. The gas-state model provides the thermodynamic and transport properties required for heat-transfer calculations, but it is not a detailed combustion model. Discrepancies between equilibrium tools can arise from species sets, thermodynamic databases, nozzle-equilibrium assumptions, transport-property models, or implementation details. Similar differences between Cantera-based calculations and CEA have been observed in related work [37]. For the present model, the main consequence is that uncertainty in gas temperature, viscosity, thermal conductivity, and Prandtl number propagates directly into the Bartz heat-transfer coefficient and the adiabatic wall temperature.

The coolant-side model also contains significant uncertainties. The baseline forced-convection correlations assume internal turbulent flow and are based on simplified channel geometries. The actual cooling channels are rough, curved, non-circular, and may include developing flow near inlets or bifurcations. The present model assumes fully developed flow locally, which is not always valid near channel entrances, contractions, expansions, or manifold regions. Entrance effects and minor losses should therefore be included in future pressure-loss modelling.

The nucleate-boiling model is the least mature heat-transfer closure in the present solver. It is implemented as a correction to the coolant-side heat-transfer coefficient and has not been verified for the ethanol–water mixture, pressure range, channel geometry, or heat fluxes considered here. A more complete treatment should include onset of nucleate boiling, critical heat flux, departure from nucleate boiling, and possible transition to film boiling or dry-out. The present implementation should therefore be used only to indicate possible heat-transfer enhancement once the wall exceeds the local saturation temperature.

The thermophysical-property model for the ethanol–water coolant is another source of uncertainty, especially for viscosity and surface tension over the relevant temperature and pressure range. Coolant properties are evaluated using CoolProp [47], but the CoolProp mixture-viscosity treatment is approximate. The CoolProp issue tracker notes that transport properties for mixtures were not originally implemented in the same way as for pure fluids, and the implemented fallback is based on a simple mixture rule rather than a dedicated transport-property model [48]. This is important because viscosity directly affects Reynolds number, Prandtl number, friction factor, pressure loss, and coolant-side heat-transfer coefficients. Uncertainty in mixture viscosity therefore propagates through both the thermal and hydraulic predictions. The parachor-based surface-tension estimate used for the boiling model was also implemented as a practical approximation and was not validated against experimental mixture data [15]. Future work should therefore compare the implemented ethanol–water properties against measured mixture-property data before using the boiling model or detailed pressure-loss predictions for final design.

The pressure-loss model should also be extended. In the present work, dynamic-pressure changes were included where needed for the RL10 case, but they should be incorporated into the general solver formulation. Minor losses from entrances, exits, contractions, expansions, bends, manifolds, and bifurcations are also not fully represented. These losses may be important in compact additively manufactured cooling systems, where the manifold and transition regions can contribute significantly to the total pressure drop. Including these effects would make the hydraulic model more generally applicable.

Gas-side radiation and film cooling are not included in the present thermal model. Film cooling may strongly reduce gas-side heat transfer in engines where a fraction of the fuel is injected along the wall, and a future film-cooling model should distinguish between liquid-film evaporation near the injector and gaseous film cooling farther downstream. Existing NASA film-cooling models and recent student-engine work provide useful starting points for this extension [4, 12]. Radiation from combustion products such as H₂O, CO₂, and CO may also contribute to the wall heat flux, depending on chamber pressure, gas composition, optical path length, wall temperature, and the convective heat-flux level. Both effects should therefore be assessed with simplified engineering models before being neglected in future engine designs.

Thermal barrier coating properties are another source of uncertainty. For P25 and P26, the effective coating thickness and uniformity are not known precisely, even though the coating has a direct influence on the predicted metal temperature and heat absorbed by the coolant. For detailed design, the coating should be applied using a controlled process with measured thickness and material properties rather than treated as an uncertain in-situ effect.

Finally, the solver does not currently include structural analysis. Wall temperature alone is not a complete safety metric. The relevant failure mechanisms depend on thermal stress, pressure stress, creep, fatigue, material strength at temperature, heat-treatment condition, and residual stresses from LPBF manufacturing [39, 40, 42]. The material strength before and after heat treatment, and the residual stress state before and after stress relief, should be quantified for future chambers. Even a first-order stress estimate would make the tool more useful, because it would allow geometries to be compared based on mechanical margin rather than only on temperature and pressure drop.

8.6 Future improvements

A useful near-term improvement would be to extend the wall model from a wall-normal resistance network to a quasi-two-dimensional wall-conduction solver. The gas and coolant streams could remain quasi-one-dimensional, while the chamber-wall cross-section is solved with temperature-dependent conductivity. This would improve estimates of rib-base temperatures, local thermal gradients, and stress-driving temperature differences.

The hydraulic model should also be extended. Dynamic-pressure changes, entrance effects, minor losses, manifolds, contractions, expansions, and bifurcation losses should be included in the general solver rather than treated only in selected cases. These effects are especially relevant for additively manufactured chambers, where compact manifolds and complex channel transitions are common. The roughness model should be refined using more repeated samples, direct internal-channel characterization, and additional flow tests so that the relation between measured roughness descriptors and hydraulic performance can be assessed more robustly.

Several alternative heat-transfer submodels should be implemented for sensitivity studies. On the gas side, the Bartz correlation should be supplemented with other empirical models to quantify the sensitivity of predicted wall temperature to gas-side heat-transfer closure. On the coolant side, alternative forced-convection and boiling correlations should be available. The modular solver architecture is useful for this purpose, since gas models, coolant-property models, heat-transfer correlations, and pressure-loss models can be exchanged without rewriting the full program.

Film cooling and radiation should also be added in future work. Film cooling is particularly relevant for engines where fuel is deliberately injected near the wall to reduce the thermal load, and existing NASA film-cooling models provide a useful starting point [4]. Radiation should be assessed at least with simplified correlations for combustion products dominated by H₂O, CO₂, and CO.

A practical software improvement would be to add a lightweight Python interface or case-template layer to the C++ solver. This would make the tool easier to use for future students and design teams while preserving the performance and modularity of the underlying implementation.

The experimental basis should also be improved. Future engines or dedicated test articles should include local pressure and temperature measurements along the cooling path, with the greatest benefit coming from direct metal-temperature measurements at multiple axial locations. Additive manufacturing may make it possible to design chambers with embedded instrumentation features or accessible thermocouple locations. Controlled thermal barrier coating deposition and post-test coating characterization would also reduce one of the major uncertainties in the present P25 and P26 analyses.

9 Conclusion

This thesis has developed and assessed a reduced-order model for regenerative cooling analysis of liquid rocket thrust chambers. The objective was not to replace detailed CFD or conjugate heat-transfer analysis, but to establish a fast and transparent engineering tool that connects chamber geometry, cooling-channel design, heat transfer, pressure loss, material behaviour, and additive-manufacturing effects. The model treats the gas and coolant flow using a quasi-one-dimensional formulation, while the wall heat transfer is represented through a local thermal-resistance network. This gives a computationally efficient framework that can be used during early design iteration, where many channel concepts and operating conditions must be compared before more detailed analysis is justified.

The model was implemented as a modular C++ solver with separate components for geometry, gas-state evaluation, coolant properties, heat-transfer correlations, pressure-loss modelling, and material properties. This structure proved useful because the dominant uncertainties in regenerative cooling are not confined to one part of the model. Gas-side heat transfer, coolant-side forced convection, boiling, pressure loss, surface roughness, and coating properties can each significantly affect the predicted wall temperature or pressure margin. The modular implementation therefore makes the solver suitable not only for the cases studied here, but also for future sensitivity studies and model extensions.

The verification and benchmark cases showed that the implemented model captures the main behaviour expected for regeneratively cooled thrust chambers. The chemical-equilibrium calculations and forced-convection comparisons verified important submodels, while the RL10 benchmark demonstrated that the integrated solver can reproduce the main thermal trends of a published cooling-jacket model. The RL10 comparison also showed that local wall-temperature and pressure-drop predictions remain sensitive to geometry reconstruction, roughness assumptions, and empirical correlation choices. This supports the use of the model as an engineering design tool, but not as a source of exact local wall-temperature validation.

A separate contribution of the thesis was the development of a surface-roughness model for additively manufactured cooling channels. The roughness measurements showed that LPBF surface quality depends strongly on local build orientation and overhang angle. This is important because internal channel roughness directly affects hydraulic performance, and because regeneratively cooled chambers often contain long, narrow channels where pressure loss is a major design constraint. The comparison with channel flow measurements showed that including both area constriction and roughness gives significantly more realistic pressure-drop predictions than treating the channels as smooth ideal geometries. The roughness model is still limited by the available measurements and should not be treated as universally transferable, but it provides a practical method for including manufacturing effects in early cooling-channel design.

The P25 case demonstrated both the usefulness and the limitations of the model when applied to a real student-developed engine. The pressure-loss comparison indicated that the hydraulic model must include roughness and channel constriction to represent the manufactured chamber. The thermal analysis showed that the predicted wall temperature is strongly dependent on the assumed coolant-side heat-transfer regime. In particular, the forced-convection-only models predicted coolant-side wall temperatures above the local saturation condition, making nucleate boiling physically relevant for the operating point considered. However, the boiling model remains incomplete because CHF, dryout, and post-CHF behaviour are not included. The observed crack in the P25 chamber also showed that steady-state wall temperature alone is not a sufficient safety metric. Thermal stress, pressure stress, residual stress from LPBF manufacturing, material condition, heat treatment, and transient operation must also be considered.

The P26 case showed how the model can be used as a design-study tool for a new cooling-channel concept. The bifurcated channel geometry improves manufacturability by reducing unsupported wall spans and limiting extreme aspect ratios, but it also introduces local three-dimensional flow effects, splitter losses, possible maldistribution, and manufacturing risk that are not fully represented in the one-dimensional formulation. The water-flow comparison gave confidence that the hydraulic model can represent the manufactured channel geometry reasonably well under con-

trolled conditions. The thermal results should instead be interpreted as design estimates, since no local hot-fire wall-temperature measurements are available and the results depend on assumptions regarding coating thickness, boiling behaviour, and gas-side heat transfer.

The main conclusion is that the reduced-order approach is valuable precisely because it exposes the coupled tradeoffs in regenerative cooling design. Lower pressure loss, higher coolant-side heat transfer, manufacturable channel geometry, acceptable wall temperature, and structural margin cannot be optimized independently. Surface roughness is a clear example: it increases pressure loss, but may also increase coolant-side heat transfer. Similarly, bifurcated channels improve manufacturability and geometric feasibility, but introduce hydraulic and modelling uncertainties. A useful design tool must therefore make these tradeoffs visible rather than hide them behind a single predicted wall temperature.

The most important future improvement is to extend the wall model from a purely wall-normal resistance network to a quasi-two-dimensional conduction model. The gas and coolant streams could remain quasi-one-dimensional, while the metal cross-section is solved with temperature-dependent conductivity. This would improve predictions of rib-base temperatures, local thermal gradients, and stress-driving temperature differences without requiring a full three-dimensional conjugate heat-transfer model. Other high-priority improvements include a more complete hydraulic loss model, additional heat-transfer correlations for sensitivity studies, implementation of film cooling and radiation models, improved boiling and CHF treatment, controlled coating characterization, and structural margin estimation.

Overall, the thesis demonstrates that a relatively simple and transparent model can provide useful engineering insight for regeneratively cooled student rocket engines. The model captures the dominant axial thermal and hydraulic trends, highlights the importance of LPBF surface roughness, supports comparison of cooling-channel concepts, and identifies where the largest uncertainties remain. Its main value is therefore not that it eliminates the need for testing or higher-fidelity simulation, but that it provides a practical framework for making better early design decisions and for deciding which uncertainties should be addressed first.

Bibliography

- [1] Ibrahim Chehimi, Jay Johal and Sam Owen. *Critical Design Review. Internal Report*. Team Name: UCL Rocket, Academic Year: 2024–2025. Apr. 2025.
- [2] Thomas Pedersen. *A Few μm of SiO_2 , Please!* Published by Copenhagen Suborbitals. Photo by Carsten Olsen. Copenhagen Suborbitals. June 2018. URL: <https://copenhagensuborbitals.com/a-few-um-of-sio2-please/>.
- [3] Emmanuel De Leon et al. ‘A Review of Modeling, Simulation, and Process Qualification of Additively Manufactured Metal Components via the Laser Powder Bed Fusion Method’. In: *Journal of Manufacturing and Materials Processing* 9.1 (2025), p. 22. DOI: 10.3390/jmmp9010022.
- [4] National Aeronautics and Space Administration. *Liquid Rocket Engine Fluid-Cooled Combustion Chambers*. NASA Special Publication NASA-SP-8087. National Aeronautics and Space Administration, Apr. 1972.
- [5] D. R. Bartz. *Turbulent Boundary-Layer Heat Transfer from Rapidly Accelerating Flow of Rocket Combustion Gases and of Heated Air*. NASA Contractor Report NASA-CR-62615. Available from NASA Technical Reports Server, document ID 19650013685. Jet Propulsion Laboratory, California Institute of Technology, Dec. 1963.
- [6] Thomas M. Tomsik. *A Hydrogen-Oxygen Rocket Engine Coolant Passage Design Program (RECOP) for Fluid-Cooled Thrust Chambers and Nozzles*. Tech. rep. NASA N95-70893. Cleveland, Ohio: NASA Lewis Research Center, 1994.
- [7] Michael Binder, Thomas Tomsik and Joseph P. Veres. *RL10A-3-3A Rocket Engine Modeling Project*. NASA Technical Memorandum NASA-TM-107318. NASA report number E-10414. National Aeronautics and Space Administration, Lewis Research Center, Jan. 1997.
- [8] Tae Jun Jeon and Tae Seon Park. ‘Thermal recycling analysis in regenerative cooling channels based on liquid rocket engine cycles’. In: *Applied Thermal Engineering* 256 (2024), p. 124095. ISSN: 1359-4311. DOI: <https://doi.org/10.1016/j.applthermaleng.2024.124095>. URL: <https://www.sciencedirect.com/science/article/pii/S1359431124017630>.
- [9] Hermann Schlichting and Klaus Gersten. *Boundary-Layer Theory*. 9th ed. Berlin, Heidelberg: Springer, 2017. ISBN: 978-3-662-52917-1. DOI: 10.1007/978-3-662-52919-5.
- [10] Y. A. ÇENGEL. *Heat and mass transfer. Fundamentals applications*. McGraw-Hill Education, 2020.
- [11] J. C. Chen. ‘Correlation for boiling heat transfer to saturated fluids in convective flow’. In: *Industrial & Engineering Chemistry Process Design and Development* 5.3 (1966), pp. 322–329.
- [12] Antoni Barredo Juan and Alberto López Platero. ‘Design of the regenerative cooling system for a 4kN LOX/Ethanol student-built liquid rocket engine’. In: *Aerospace Europe Conference 2023 – 10th EUCASS – 9th CEAS*. 2023. DOI: 10.13009/EUCASS2023-035.
- [13] Weiwei Chen and Xiande Fang. ‘A note on the Chen correlation of saturated flow boiling heat transfer’. In: *International Journal of Refrigeration* 48 (2014), pp. 100–104.
- [14] H. K. Forster and N. Zuber. ‘Dynamics of Vapor Bubbles and Boiling Heat Transfer’. In: *AIChE Journal* 1.4 (1955), pp. 531–535.
- [15] J. A. Hugill and A. J. Van Welsenes. ‘Surface Tension: A Simple Correlation for Natural Gas + Condensate Systems’. In: *Fluid Phase Equilibria* 29 (1986), pp. 383–390. DOI: 10.1016/0378-3812(86)85038-5.
- [16] David G. Goodwin et al. *Cantera: An Object-oriented Software Toolkit for Chemical Kinetics, Thermodynamics, and Transport Processes*. <https://www.cantera.org>. Version 3.0.0. 2023. DOI: 10.5281/zenodo.8137090.
- [17] University of California San Diego Combustion Research Group. *Chemical-Kinetic Mechanisms for Combustion Applications: San Diego Mechanism*. <https://web.eng.ucsd.edu/mae/groups/combustion/mechanism.html>. Accessed 2026-05-14. 2024.

-
- [18] Sanford Gordon and Bonnie J. McBride. *Computer Program for Calculation of Complex Chemical Equilibrium Compositions and Applications. Part 1: Analysis*. NASA Reference Publication 1311. Cleveland, Ohio: NASA Lewis Research Center, 1994. URL: <https://ntrs.nasa.gov/citations/19950013764>.
- [19] Bonnie J. McBride and Sanford Gordon. *Computer Program for Calculation of Complex Chemical Equilibrium Compositions and Applications. Part 2: Users Manual and Program Description*. NASA Reference Publication 1311. Cleveland, Ohio: NASA Lewis Research Center, 1996. URL: <https://ntrs.nasa.gov/citations/19960044559>.
- [20] Alexander Ponomarenko. *RPA: Tool for Liquid Propellant Rocket Engine Analysis*. <http://software.lpre.de>. May 2010.
- [21] Alexander Ponomarenko. ‘RPA – Tool for Rocket Propulsion Analysis’. In: *Proceedings of the Space Propulsion Conference*. May 19–22, 2014. Cologne, Germany, May 2014. URL: <http://www.propulsion-analysis.com>.
- [22] Helfried Steiner, Alexander Kobor and Ludwig Gebhard. ‘A Wall Heat Transfer Model for Subcooled Boiling Flow’. In: *International Journal of Heat and Mass Transfer* 48 (2005), pp. 4161–4173. DOI: 10.1016/j.ijheatmasstransfer.2005.03.032.
- [23] Kyle Niemeyer. *Rocket Propulsion Notes*. Supplementary notes for ME 499/505: Propulsion, Oregon State University. Licensed under CC BY-SA 4.0. 2024. URL: <https://kyleniemeyer.github.io/rocket-propulsion-notes/>.
- [24] Daeseong Jo et al. ‘Experimental Investigation of Convective Heat Transfer in a Narrow Rectangular Channel for Upward and Downward Flows’. In: *Nuclear Engineering and Technology* 46.2 (Apr. 2014), pp. 195–206. DOI: 10.5516/NET.02.2013.057.
- [25] Francesco Di Matteo, Marco De Rosa and Marcello Onofri. ‘Transient Simulation of the RL-10A-3-3A Rocket Engine’. In: *Space Propulsion Conference*. Bordeaux, France, 2012.
- [26] A. Adams. *Design Study of RL10 Derivatives. Volume 2: Engine Design Characteristics*. Tech. rep. NASA-CR-120145. Report No. FR-6011-VOL-2. West Palm Beach, FL, USA: Pratt & Whitney Aircraft, 1973.
- [27] Mike Jetzer. *RL10 Thrust Chamber Section Gallery: dsc85180.jpg*. <https://www.heroicrelics.org>. Image of RL10 thrust chamber section at U.S. Space & Rocket Center. Accessed: 2026.
- [28] Mike Jetzer. *RL10 Engine (Old Museum) Gallery: dsc01696.jpg*. <https://www.heroicrelics.org>. Image of regenerative cooling tubes near the hydrogen inlet manifold. Accessed: 2026.
- [29] Mike Jetzer. *RL10 Engine (Old Museum) Gallery: dsc06821.jpg*. <https://www.heroicrelics.org>. Image of regenerative cooling tubes just aft of the throat. Accessed: 2026.
- [30] Ask Hovik. *pyskyfire – A Simulation Framework for Liquid Propellant Rocket Engines*. 2025. DOI: 10.5281/zenodo.15791527. URL: <https://github.com/ask-hovik/pyskyfire>.
- [31] International Organization for Standardization. *ISO 21920-2:2021 Surface texture: Profile — Part 2: Terms, definitions and surface texture parameters*. 2021.
- [32] International Organization for Standardization. *ISO 25178-2:2021 Surface texture: Areal — Part 2: Terms, definitions and surface texture parameters*. 2021.
- [33] International Organization for Standardization. *ISO 21920-3:2021 Surface texture: Profile — Part 3: Specification operators*. 2021.
- [34] International Organization for Standardization. *ISO 25178-3:2012 Surface texture: Areal — Part 3: Specification operators*. 2012.
- [35] Thomas Adams, Christopher Grant and Heather Watson. ‘A Simple Algorithm to Relate Measured Surface Roughness to Equivalent Sand-grain Roughness’. In: *International Journal of Mechanical Engineering and Mechatronics* 1.1 (2012), Article ID 008. ISSN: 2929-2724. DOI: 10.11159/ijmem.2012.008.
- [36] Satish G. Kandlikar et al. ‘Characterization of Surface Roughness Effects on Pressure Drop in Single-Phase Flow in Minichannels’. In: *Physics of Fluids* 17.10 (2005), p. 100606. DOI: 10.1063/1.1896985.
- [37] Ask Haugerud Hovik. ‘A Simulation Framework for Liquid Propellant Rocket Engines’. Master’s thesis. Norwegian University of Science and Technology, July 2025.
-

-
- [38] George P. Sutton and Oscar Biblarz. *Rocket Propulsion Elements*. 9th ed. Hoboken, NJ: John Wiley & Sons, 2017. ISBN: 978-1118753651.
- [39] Alessandra Martucci et al. ‘Effects of Stress-Relieving Temperature on Residual Stresses, Microstructure and Mechanical Behaviour of Inconel 625 Processed by PBF-LB/M’. In: *Metals* 13.4 (2023), p. 796. DOI: 10.3390/met13040796.
- [40] Giulio Marchese et al. ‘Heat-Treated Inconel 625 by Laser Powder Bed Fusion: Microstructure, Tensile Properties, and Residual Stress Evolution’. In: *Journal of Materials Engineering and Performance* 33 (2024), pp. 6825–6834. DOI: 10.1007/s11665-024-09235-7.
- [41] R. J. Quentmeyer. ‘Rocket Combustion Chamber Life-Enhancing Design Concepts’. In: *26th Joint Propulsion Conference*. AIAA 90-2116. Orlando, Florida, July 1990. DOI: 10.2514/6.1990-2116.
- [42] R. W. Miller. *Cyclic Fatigue Analysis of Rocket Thrust Chambers. Volume 1: OFHC Copper Chamber Low Cycle Fatigue*. NASA Contractor Report NASA-CR-134641-VOL-1. National Aeronautics and Space Administration, June 1974.
- [43] Julie A. Carlile and Richard J. Quentmeyer. *An Experimental Investigation of High-Aspect-Ratio Cooling Passages*. NASA Technical Memorandum NASA-TM-105679. Also published as AIAA-92-3154. National Aeronautics and Space Administration, Lewis Research Center, July 1992.
- [44] Mary F. Wadel. *Comparison of High Aspect Ratio Cooling Channel Designs for a Rocket Combustion Chamber*. NASA Technical Memorandum NASA-TM-107473. Also published as AIAA-97-2913. National Aeronautics and Space Administration, Lewis Research Center, June 1997.
- [45] Iain Waugh et al. ‘Green-Laser Additive Manufacture of a GRCop-42 LOX/LCH₄ Combustion Chamber with Compliant Firewall’. In: *Space Propulsion 2024*. SP2024.274. Airborne Engineering Ltd., Fraunhofer ILT, and Manufacturing Technology Centre. Glasgow, United Kingdom, May 2024.
- [46] William Dugan. *Modeling In-Situ SiO₂ Particle Deposition for Thermal Barrier Coatings in Rocket-Like Flows*. Course report. Technische Universiteit Delft, Aug. 2025.
- [47] Ian H. Bell et al. ‘Pure and Pseudo-pure Fluid Thermophysical Property Evaluation and the Open-Source Thermophysical Property Library CoolProp’. In: *Industrial & Engineering Chemistry Research* 53.6 (2014), pp. 2498–2508. DOI: 10.1021/ie4033999. URL: <https://doi.org/10.1021/ie4033999>.
- [48] CoolProp Contributors. *Transport Properties for Mixtures #307*. GitHub issue, CoolProp/-CoolProp. Accessed 5 June 2025. 2014. URL: <https://github.com/CoolProp/CoolProp/issues/307>.

Appendix

A Surface-roughness data




Figure A.1: Surface roughness data for different overhang angles, separated by upskin/downskin and wall-thickness.

B Inconel 625 properties




Figure B.1: Inconel 625 (annealed) material properties. Curve fits are made with data from ANSYS Granta EduPack 2023.

C 347 Stainless steel properties




Figure C.1: Stainless steel 347 material properties. Curve fits are made with data from Sandmeyer Steel.

D Ethanol–water mixture properties




Figure D.1: Thermophysical properties as a function of temperature for a binary mixture of 90 wt% ethanol and 10 wt% water at $p = 2.5$ MPa. The data was generated with the HEOS-backend in CoolProp.




Figure D.2: Surface tension of a binary mixture of 90 wt% ethanol and 10 wt% water as a function of saturation pressure. The parachor-model prediction is compared with mass-fraction-weighted and mole-fraction-weighted mixture averages.

The parachor model predicts a lower surface tension than the linear mixture estimates because surface tension is not generally additive with composition. Ethanol–water mixtures exhibit non-ideal interfacial behaviour, and the lower-surface-tension ethanol component may influence the liquid–vapour interface more strongly than implied by the bulk mass or mole fractions. The parachor model accounts for mixture behaviour through the mixture parachor and the liquid–vapour density difference, whereas the weighted averages are included only as simple reference estimates.

E Reconstructed RL10 geometry with source data




Figure E.1: Reconstructed RL10 cooling-jacket geometry used as input to the benchmark model. The chamber-wall radius, channel width, and channel height were obtained by fitting smooth curves through the available digitized source data. The plotted markers show the underlying digitized points, while the continuous lines show the fitted geometry used in the present calculations. The source data are taken from the reconstructed RL10 dataset reported by Hovik [30, 37].

F Cantera input file for cryogenics at normal boiling point


Listing 1: Custom Cantera input file used for cryogenic reactants.

```
1 description: |-
2   Thermodynamic data for common cryogenic fluids modeled as constant heat
3   capacity liquids. Data is taken from McBride. Only to be used at NBP as
4   reactants.
5
6 phases:
7   - name: liquid_oxygen
8     thermo: ideal-gas
9     elements: [O]
10    species: [O2(L)]
11
12   - name: liquid_nitrous
13     thermo: ideal-gas
14     elements: [N, O]
15     species: [N2O(L)]
16
17   - name: liquid_hydrogen
18     thermo: ideal-gas
19     elements: [H]
20     species: [H2(L)]
21
22   - name: liquid_methane
23     thermo: ideal-gas
24     elements: [C, H]
25     species: [CH4(L)]
26
27 species:
28   - name: O2(L)
29     composition: {O: 2}
30     thermo:
31       model: constant-cp
32       T0: 90.170
33       h0: -12979.0 J/mol
34       s0: 0.0
35       cp0: 0.0
36
37   - name: N2O(L)
38     composition: {O: 1, N: 2}
39     thermo:
40       model: constant-cp
41       T0: 184.67
42       h0: 61079.0 J/mol
43       s0: 0.0
44       cp0: 0.0
45
46   - name: H2(L)
47     composition: {H: 2}
48     thermo:
49       model: constant-cp
50       T0: 20.270
51       h0: -9012.0 J/mol
52       s0: 0.0
53       cp0: 0.0
54
55   - name: CH4(L)
56     composition: {C: 1, H: 4}
57     thermo:
58       model: constant-cp
59       T0: 111.66
60       h0: -89233.0 J/mol
61       s0: 0.0
62       cp0: 0.0
```

G SUNDIALS KINSOL residual function

Listing 2: Residual function used in the local nonlinear thermal solve.

```
1 SUNErrCode Solver::residual(N_Vector y, N_Vector f, void* data_ptr)
2 {
3     SolverData* data = static_cast<SolverData*>(data_ptr);
4
5     // *****
6     // GET CURRENT SOLUTION
7     // *****
8     const real_t Ti = NV_Ith_S(y, 0); // [K] Inner-most wall temperature
9     const real_t Thwg = NV_Ith_S(y, 1); // [K] Gas-side wall temperature
10    const real_t Thwl = NV_Ith_S(y, 2); // [K] Liquid-side wall temperature
11    const real_t Tl = NV_Ith_S(y, 3); // [K] Liquid bulk temperature
12
13    #ifdef SPECTRA_DEBUG
14    if (data->debug)
15    {
16        SPECTRA_CORE_TRACE(
17            "Ti {0:6.1f} | Twg {1:6.1f} | Twl {2:6.1f} | Tl {3:6.1f}",
18            Ti, Thwg, Thwl, Tl
19        );
20    }
21    #endif
22
23    // *****
24    // CALCULATE CURRENT STATE
25    // *****
26    const real_t h_gas = data->model.htc_gas(Ti);
27    const real_t h_liquid = data->model.htc_liquid(Thwl, Tl);
28    const real_t r_wall = data->model.resistance_wall(Thwg, Thwl);
29    const real_t r_TBC = data->model.resistance_TBC(Ti, Thwg);
30    const real_t heatFlux = (data->model.referenceTemperature() - Tl) / (
31        1.0 / h_gas + r_wall + r_TBC + 1.0 / h_liquid);
32
33    // *****
34    // STORE CURRENT STATE
35    // *****
36    data->output = {
37        .tempTBC = Ti,
38        .tempGasWall = Thwg,
39        .tempLiquidWall = Thwl,
40        .tempCoolant = Tl,
41        .gasHTC = h_gas,
42        .liquidHTC = h_liquid,
43        .heatFlux = heatFlux
44    };
45
46    // *****
47    // RESIDUALS
48    // *****
49    // Combustion gas - TBC convection
50    NV_Ith_S(f, 0) = h_gas * (data->model.referenceTemperature() - Ti)
51        - heatFlux;
52
53    // Wall conduction
54    NV_Ith_S(f, 1) = (Thwg - Thwl) / r_wall - heatFlux;
55
56    // Wall - coolant convection
57    NV_Ith_S(f, 2) = h_liquid * (Thwl - Tl) - heatFlux;
58
59    // Coolant energy balance (temperature change)
60    NV_Ith_S(f, 3) = data->model.getMassFlow()
61        * data->coolantSpecificHeat
62        * (Tl - data->prevCoolantTemp)
63        - heatFlux * data->model.contactArea();
64
65    #ifdef SPECTRA_DEBUG
66    if (data->debug)
67    {
68        SPECTRA_CORE_TRACE(
69            "r0 {0:8.1e} | r1 {1:8.1e} | r2 {2:8.1e} | r3 {3:8.1e}",
70            NV_Ith_S(f, 0), NV_Ith_S(f, 1), NV_Ith_S(f, 2), NV_Ith_S(f, 3)
71        );
72    }
73    #endif
74
75    return KIN_SUCCESS;
76 }
```



Deklarasjon om KI-hjelpemidler

Har det i utarbeidingen av denne rapporten/avhandlingen blitt anvendt KI-baserte hjelpemidler?

Nei



Ja

Hvis nei, signer nederst. Hvis ja: spesifiser type av verktøy og bruksområde under.

Tekst



Stavekontroll. Er deler av teksten kontrollert av: Grammarly, Ginger, Grammarbot, LanguageTool, ProWritingAid, Sapling, Trinkai.ai eller lignende verktøy?

Tekstgenerering. Er deler av teksten generert av: ChatGPT, GrammarlyGO, Copy.AI, WordAi, WriteSonic, Jasper, Simplified, Rytr eller lignende verktøy?

Skriveassistanse. Er en eller flere av ideene eller fremgangsmåtene i oppgaven foreslått av: ChatGPT, Google Bard, Bing chat, YouChat eller lignende verktøy?

Hvis ja til anvendelse av et tekstverktøy - spesifiser bruken her:

Stavekontroll, rettskriving, tips til setningsoppbygning, støtte til figurtekster.

Koder og algoritmer



Programmeringsassistanse. Er deler av koden/algoritmene som i) fremtrer direkte i rapporten eller ii) har blitt anvendt for produksjon av resultater slik som figurer, tabeller eller tallverdier blitt generert av: GitHub Copilot, CodeGPT, Google Codey/Studio Bot, Replit Ghostwriter, Amazon CodeWhisperer, GPT Engineer, ChatGPT, Google Bard eller lignende verktøy?

Hvis ja til anvendelse av et programmeringsverktøy - spesifiser bruken her:

Github copilot og ChatGPT har blitt brukt til å generere plots, grafer or tabeller (ikke innhold). Claude har blitt brukt til å lage Tikz diagrammer.

Bilder og figurer



Bildegenerering. Er ett eller flere av bildene/figurene i rapporten blitt generert av: Midjourney, Jasper, WriteSonic, Stability AI, Dall-E eller lignende verktøy?

Hvis ja til anvendelse av et bildeverktøy - spesifiser bruken her:

ChatGPT har blitt brukt til å fjerne bakgrunn på fotografier for å gjøre dem mer passende for oppgaven.

Andre KI-verktøy. Har andre typer av verktøy blitt anvendt? Hvis ja spesifiser bruken her:



Jeg er kjent med NTNUs regelverk for bruk av kunstig intelligens. Jeg har redegjort for all anvendelse av kunstig intelligens enten i) direkte i rapporten eller ii) i dette skjemaet.

William Dugan,
Underskrift/Dato/Sted