POLITECNICO DI TORINO

Master's Degree Programme
in **Mechatronic Engineering**

Master's Degree Thesis

Thermal Management for Space Electronics and Topology-Optimized In-Chip Microfluidic Cooling



Politecnico
di Torino

Supervisor
Prof. Radu Bojoi

Candidate
MAHDY MOKDAD

Research affiliation
EPFL Power Lab
Center of MicroNanoTechnology (CMi)

Academic Year 2025/2026

Contents

I Literature review and comparative context 1

Introduction 1

Spacecraft Thermal-Control Architectures 1
- System overview and functional blocks . . . 1
- Heat extraction: passive and active approaches . . . 2
- Single-loop and dual-loop systems . . . 2
- Integration constraints . . . 2
- Examples of system-level performance . . . 2

Heat extraction 2
- Passive heat-transport technologies . . . 2
- Conventional heat pipes . . . 2
- Loop heat pipes . . . 3
- Oscillating heat pipes and additive manufacture . . . 3
- Performance limits . . . 3
- Active heat transport and cold plates . . . 3
- Mechanically pumped single-phase loops . . . 3
- Mechanically pumped two-phase loops . . . 4
- Cold-plate design . . . 4
- Liquid cooling and dielectric coolants . . . 4
- Single-phase and two-phase operation . . . 4
- Coolant selection . . . 5
- Single-phase experimental benchmarks . . . 5
- Microchannels and microfluidics (integrated within liquid cooling) . . . 5
- Why small channels work . . . 5
- Channel and fin geometry . . . 5
- Manifolds and maldistribution . . . 6
- Performance metrics and materials . . . 6
- Topology optimization of microchannels and heat sinks . . . 7
- Mathematical formulation . . . 7
- Two- and three-dimensional models . . . 8
- Competing objectives and Pareto selection . . . 8
- Manifold and transient optimization . . . 8
- Manufacturing constraints . . . 8
- Reported benchmarks . . . 8

Heat storage: phase-change materials 9
- Storage mechanism and sizing . . . 9
- Materials and spacecraft integration . . . 9

Heat rejection: space radiators 9
- Fixed and deployable radiators . . . 9
- Variable-emittance and fault-tolerant concepts . . . 10

Aerospace-relevant COP benchmark 10
- Implications for the present work . . . 10

Research Gaps and Future Directions 11
- What is already mature . . . 11
- The unresolved combination . . . 11
- Recommended research programme . . . 11

II Experimental case study 13

Chapter overview 13
- Main contributions of the work . . . 13

1. Device selection and heat-load definition **13**

2. Numerical design methodology **15**
2.1 Unit-cell screening . . . 15
2.2 Coupled thermo-fluid model . . . 15
2.3 Density-based topology optimization . . . 15
2.4 Adaptation of the reference topology-optimization methodology . . . 16
Material interpolation and penalization . . . 16
Iteration and continuation procedure . . . 18
Transfer of the method to the present microchannel cooler . . . 18

3. Manifold, cap, and PCB-level integration **20**

4. Silicon microfabrication at EPFL CMi **22**
4.1 Layout preparation and pattern polarity . . . 22
4.2 Surface cleaning, dehydration, and HMDS priming . . . 22
4.3 AZ 10XT-60 coating and soft bake . . . 23
4.4 Direct-write exposure with the MLA150 . . . 23
4.5 Development in AZ 400K . . . 23
4.6 Deep reactive-ion etching in the SPTS Rapier . . . 23
4.7 Resist stripping, cleaning, and inspection . . . 24
4.8 Fabrication fidelity and its influence on performance . . . 25

5. Experimental methodology **25**
5.1 Closed-loop architecture . . . 25
5.2 Instrumentation, electrical operation, and test sequence . . . 26

6. Experimental results **27**
6.1 Temperature rise and thermal resistance . . . 27
6.2 Thermo-hydraulic comparison . . . 28
6.3 Numerical validation . . . 29

7. Discussion **29**
7.1 Why the optimized geometry performed better . . . 29
7.2 Relevance to space electronics . . . 30
7.3 Limitations . . . 30

Conclusion **30**

Nomenclature and Acronyms

Acronyms

CAD	Computer-aided design
CMi	Center of MicroNanoTechnology, EPFL
COP	Coefficient of performance
DI	Deionized
DRIE	Deep reactive-ion etching
DSE	Deep silicon etching
EPFL	*École polytechnique fédérale de Lausanne*
ESC	Electrostatic chuck
GaN	Gallium nitride
HMDS	Hexamethyldisilazane
ICP	Inductively coupled plasma
IR	Infrared
ISS	International Space Station
LEO	Low Earth orbit
LHP	Loop heat pipe
MLA	Maskless aligner
MMA	Method of moving asymptotes
NI	National Instruments
OES	Optical emission spectroscopy
OHP	Oscillating heat pipe
ONB	Onset of nucleate boiling
PCB	Printed circuit board
PCM	Phase-change material
PR	Photoresist
RTD	Resistance temperature detector
SEM	Scanning electron microscopy
SiC	Silicon carbide
TIM	Thermal-interface material
TO	Topology optimization

Principal Symbols

Symbol	Description	Unit
A	Area	m^2
A_{die}	Active die area	m^2
c_p	Specific heat capacity at constant pressure	$\mathrm{J\,kg^{-1}K^{-1}}$
D_h	Hydraulic diameter	m
h	Convective heat-transfer coefficient	$\mathrm{W\,m^{-2}K^{-1}}$
J	Optimization objective function	–
k	Thermal conductivity	$\mathrm{W\,m^{-1}K^{-1}}$
K_{eff}	Effective thermal conductance	$\mathrm{W\,K^{-1}}$
L	Characteristic or transport length	m
m	Mass	kg
$\dot{m}$	Mass-flow rate	$\mathrm{kg\,s^{-1}}$
P_{diss}	Power dissipated by the device	W
P_{hyd}	Hydraulic pumping power	W
Q, $\dot{Q}$	Heat or heat-transfer rate, according to context	J, W
q''	Surface heat flux	$\mathrm{W\,m^{-2}}$
q'''	Volumetric heat generation	$\mathrm{W\,m^{-3}}$

Symbol	Description	Unit
R_{th}	Thermal resistance	$\mathrm{K\,W^{-1}}$
R''_{th}	Area-normalized thermal resistance	$\mathrm{m^2K\,W^{-1}}$
T	Temperature	K or °C
$\mathbf{u}$	Fluid-velocity vector	$\mathrm{m\,s^{-1}}$
$\dot{V}$	Volumetric-flow rate	$\mathrm{m^3\,s^{-1}}$
α	Inverse-permeability interpolation	$\mathrm{kg\,m^{-3}s^{-1}}$
β	Projection-slope parameter	–
γ	Topology density/design variable	–
Δp	Pressure drop	Pa
ΔT	Temperature difference	K
η	Pump efficiency	–
μ	Dynamic viscosity	Pa s
ρ	Density	$\mathrm{kg\,m^{-3}}$

Literature review and comparative context

Abstract

The thermal design of spacecraft is becoming more difficult as electronic power density rises while allowable mass remains tightly constrained. Phased-array radars, advanced communication payloads, and gallium-nitride (GaN) or silicon-carbide (SiC) power devices can create local heat fluxes above $200\,\mathrm{W\,cm^{-2}}$. Removing that heat without adding excessive mass, volume, or electrical demand is now a central spacecraft-design problem. This review considers the main technologies used to acquire heat at the electronic component, move or temporarily store it, and ultimately reject it to space. Particular attention is given to dual-loop systems, in which a relatively cool primary circuit serves the electronics and a separate, warmer secondary circuit serves the radiator. Such separation allows each loop to use a suitable fluid, pressure, and operating temperature. The discussion covers heat pipes, loop heat pipes, oscillating heat pipes, mechanically pumped loops, liquid cold plates, fixed and deployable radiators, variable-emittance surfaces, phase-change materials (PCMs), dielectric liquids, microchannel cooling, and topology-optimization methods. It also considers non-uniform and time-dependent heat loads, since the hottest part of a modern device may experience a load far above the chip average. The literature reveals a clear unresolved problem: topology-optimized microfluidic coolers have rarely been designed around dielectric fluids while also accounting for spacecraft-specific, heterogeneous, and transient heat loads. Addressing this combination is important if future high-power payloads are to remain within their temperature limits without an unacceptable thermal-control penalty.

Introduction

Spacecraft electronics are being asked to do much more than in earlier missions. On-board computers increasingly handle autonomous navigation and large data streams, while radar, communication, and electric-propulsion systems place additional demands on the electrical and thermal architecture [1]. Wide-bandgap devices such as GaN and SiC make high-frequency switching and compact power conversion possible, but the same compactness concentrates dissipation within millimetre-scale regions [2]. The load is not always steady. During an active radar scan, for example, a synthetic-aperture radar payload may briefly release two to five times its steady thermal load [3].

These devices still have to remain within qualified temperature limits. As a broad guide, commercial electronics are commonly kept below about 85°C, whereas military-grade parts may be rated to roughly 125°C. In orbit there is no surrounding air to carry heat away by convection. Heat must therefore be conducted from the component, transported through the spacecraft, and finally radiated to space.

The design freedom is limited by mass, envelope, and power. Even after the recent reduction in commercial launch prices, placing a kilogram in low Earth orbit still costs thousands of euros [4]. A kilogram assigned to pumps, pipes, or radiators cannot be used for an instrument, propellant, or structural margin. Deployable panels must also fit inside the launch fairing, while internal cold plates and lines compete with avionics and tanks for space. Any active device draws from the same solar-array and battery budget as the mission payload. Thermal design is therefore not simply a matter of maximizing heat transfer; temperature, mass, volume, pumping power, and reliability have to be balanced together.

Spacecraft Thermal-Control Architectures

A spacecraft thermal-control system performs three linked tasks: it acquires heat at the equipment, transports that energy to another location, and rejects it to the external environment [1]. The best arrangement depends on the total load, the location and timing of individual sources, the spacecraft geometry and orbit, and the intended mission duration.

System overview and functional blocks

Figure 1 summarizes a typical spacecraft thermal-control system in functional form: **heat extraction** at the equipment (passive or pumped interfaces), optional **heat storage** (e.g., PCM buffering for pulsed loads), and **heat rejection** through radiators to space. The remainder of this literature review is organized around these three functions.

Heat extraction (TIMs / spreaders / heat pipes / cold plates / microchannels) → **Transport bus** (single- or dual-loop) → **Heat storage** (PCM, optional) → **Heat rejection** (radiator)

Figure 1: Simplified functional diagram of a spacecraft thermal-control system (architecture overview).

Heat extraction: passive and active approaches

Heat can be extracted passively through thermal-interface materials, spreaders, or direct attachment to a heat-pipe evaporator. An active interface instead forces a liquid through a cold plate or microchannel structure adjacent to the heat source. Transport may likewise be passive (solid conduction, capillary two-phase devices, thermosyphons) or active (pump, accumulator, valves, controlled fluid circuit) [1].

Passive hardware is attractive because it has no moving parts, consumes no operating power, and has a long flight record. Eliminating pumps and control valves removes several single-point failures, but passive devices have limited control authority and are ultimately constrained by capillary, sonic, entrainment, and dry-out limits [2].

Single-loop and dual-loop systems

In a single-loop system, one fluid collects heat from the equipment and carries it directly to the radiator. This reduces component count, inventory, mass, and potential leak locations. It also gives a quick thermal response because the source and sink are directly coupled. The same coupling, however, means that a single coolant must satisfy two rather different sets of conditions. Near the electronics it should have low viscosity, high heat capacity, and acceptable electrical behaviour; at the radiator it must resist freezing, retain chemical stability, and operate at an appropriate vapour pressure. The common temperature and pressure must suit the cold plates, pump, pipework, and radiator at once [5, 6, 7, 8, 9].

A dual-loop arrangement inserts a heat exchanger between two circuits. The primary loop serves the cold plates or microchannels and can be kept comparatively cool. The secondary loop transfers heat from the interface exchanger to the radiator and can be run warmer, improving radiative output per unit area. This separation allows a dielectric, low-viscosity fluid and modest pressure on the equipment side, while the radiator side can use a fluid selected for heat capacity, freezing point, and pressure-drop tolerance. It also isolates contamination and pressure disturbances. The additional exchanger and plumbing carry a mass and temperature-drop penalty, but the architecture permits much more independent optimization of the source and rejection sides [1].

Integration constraints

Passive-system mass is concentrated in spreaders, heat pipes, radiator facesheets, and structural doublers. An active system adds the pump, accumulator, valves, fluid inventory, and control electronics. Radiator specific mass depends strongly on operating temperature and unobstructed view to space; a deployable system also needs hinges, actuators, latches, and flexible or articulated fluid connections. Internally, lines must be routed around tanks and avionics without losing good thermal contact. A conventional cold plate often adds 5–20 mm to an equipment chassis. A microchannel plate may reduce this to roughly 2–10 mm, although its smaller passages generally demand a larger pressure head [1].

Examples of system-level performance

Representative spacecraft thermal-control implementations span (i) passive acquisition and spreading (TIMs, straps, heat pipes/LHPs) feeding fixed radiators, and (ii) mechanically pumped loops (single- or two-phase) that distribute heat from multiple boxes to one or more radiator panels [1]. At the largest flight-proven scale, the ISS external active thermal-control system rejects tens of kilowatts through ammonia loops and deployable radiators [6].

For high-power payloads, system-level performance is often set by the combined constraints of radiator envelope, allowable bus temperature, pump power allocation, and internal routing/packaging. For example, radiator-panel integration choices (facesheet, tube bonding, embedded heat pipes, deployment mechanisms, and plumbing through hinges) directly trade radiator mass/area against temperature uniformity and allowable peak equipment temperature [1, 3].

Heat extraction

Passive heat-transport technologies

Passive two-phase devices use evaporation, condensation, and capillary action to achieve an effective conductance far above solid metal of similar size [1]. Heat pipes, loop heat pipes (LHPs), and oscillating heat pipes (OHPs) are the principal members of this family.

Evaporator (heat in) → **Vapour core** → **Condenser** (heat out)
Wick / capillary return closes the loop (liquid back to evaporator)

Figure 2: Simplified operating principle of a passive two-phase heat-transport device (heat pipe / LHP concept).

Conventional heat pipes

A conventional heat pipe is a sealed tube containing a working fluid and a wick. Heat entering the evaporator vaporizes liquid from the wick. The vapour travels through the open core to the condenser, releases its latent heat,

and condenses. Menisci in the wick then generate the capillary pressure needed to return liquid to the evaporator. Dry-out begins when that pressure can no longer overcome the liquid, vapour, gravitational, and wick losses.

The effective conductance is

$$K_{\text{eff}} = \frac{Q}{\Delta T}, \tag{1}$$

and the corresponding effective conductivity of a device with length L and area A is

$$k_{\text{eff}} = \frac{K_{\text{eff}} L}{A}. \tag{2}$$

Depending on geometry and operating point, the apparent conductivity can be one hundred to one thousand times that of copper. A representative spacecraft ammonia heat pipe with a sintered wick and an outer diameter near 10 mm may provide 10–50 $\mathrm{W\,K^{-1}}$ over 0.5–2 m, and carry roughly 50–500 W [1, 7, 8, 9].

Loop heat pipes

An LHP separates the liquid and vapour paths. Liquid passes from the condenser to a compensation chamber and through a porous evaporator wick; generated vapour returns by another line. Removing counter-current flow permits transport distances above 10 m with a small temperature drop. The compensation chamber stores fluid as temperature changes and helps prevent transient dry-out. Flexible lines can be routed around the spacecraft, and the evaporator can remain compact even when the radiator is remote. For an evaporator of similar size, the maximum transported load can be two to five times that of a conventional pipe, and stable operation is possible from roughly 5% to full design load.

Oscillating heat pipes and additive manufacture

An OHP consists of a long capillary passage folded into many turns and partially filled so that liquid slugs alternate with vapour plugs. Heating expands vapour near the evaporator; condensation lowers pressure at the cold end. The resulting self-excited motion repeatedly moves energy between the two regions. Because no wick is required, the geometry is simpler and can be embedded in a compact structure. Performance is most attractive at transport lengths below about 0.5 m, which makes OHPs candidates for small satellites. Additive manufacturing can form a monolithic metal body with internal meandering passages, removing brazed joints and allowing the channel layout to be tailored to the equipment.

Performance limits

Small, long conventional pipes may provide only about 5 $\mathrm{W\,K^{-1}}$, while larger and shorter units can approach 100 $\mathrm{W\,K^{-1}}$. Typical LHP values span about 20–200 $\mathrm{W\,K^{-1}}$, and compact OHPs about 2–20 $\mathrm{W\,K^{-1}}$. These figures refer to steady operation away from the extremes. Conductance falls at very low load, when the evaporator is incompletely wetted, and again close to the transport limit.

Several physical limits can intervene. At the capillary limit, liquid return is insufficient and the evaporator dries. The sonic limit appears when vapour flow chokes, particularly with light working fluids at high temperature. The entrainment limit is reached when fast vapour strips liquid from the wick. Finally, boiling within the wick may form bubbles that obstruct liquid supply. Depending on size and length, the capillary capacity of a conventional pipe is approximately 50–500 W; an LHP can reach 1 kW or more.

These totals can be misleading for modern devices because evaporator heat flux is often the tighter constraint. A 1 kW LHP attached over 100 $\mathrm{cm^2}$ is handling only 10 $\mathrm{W\,cm^{-2}}$, well below the 100–250 $\mathrm{W\,cm^{-2}}$ average that may occur in a high-power module. A conventional sintered wick with conductivity near 15 $\mathrm{W\,m^{-1}K^{-1}}$ and thickness 2 mm can add 0.4–2.0 $\mathrm{K\,cm^2\,W^{-1}}$. At 100 $\mathrm{W\,cm^{-2}}$, this alone corresponds to a 40–200°C rise. Once spreading and interface resistance are included, the reported junction-to-evaporator range of 0.5–3.0 $\mathrm{K\,cm^2\,W^{-1}}$ can consume most of the available margin. That is the main reason active, direct liquid cooling becomes necessary at high flux.

Active heat transport and cold plates

Pumps remove the capillary ceiling by supplying the pressure needed to circulate coolant. The price is parasitic power, moving parts, controls, and a more demanding fluid-management problem. Above the capacity of passive devices, however, a pumped system is often the only practical way to cool several distributed, high-power boxes.

Mechanically pumped single-phase loops

A single-phase circuit normally contains cold plates, a pump, a heat exchanger or radiator connection, and an accumulator for thermal expansion. The liquid remains below its phase-change condition, warms sensibly in the cold plate, and cools at the sink. Flow is predictable and can be regulated with pump-speed or valve control. The transported heat is set by mass flow, heat capacity, and the permitted inlet-to-outlet rise. At the largest flight-proven scale, the International Space Station uses two external ammonia circuits, each with a nominal rejection capability of about 35 kW through deployable radiator panels [3].

Table 1: Indicative comparison of passive heat-transport devices.

Technology	Conductance	Transport limit	Advantages	Limitations	Space-use examples
Conventional heat pipe	5–100 $\mathrm{W\,K^{-1}}$	50–500 W, usually capillary-limited	No operating power; simple; highly mature and reliable	Usually below 2 m; counter-current losses; possible orientation sensitivity	Communication satellites, Mars rovers, and ISS thermal hardware
Loop heat pipe	20–200 $\mathrm{W\,K^{-1}}$	1 kW or more	Long, flexible routing; broad load range; higher capacity	More complex evaporator; greater fluid inventory; evaporator resistance	Swift BAT, Mars Science Laboratory, and commercial satellites
Oscillating heat pipe	2–20 $\mathrm{W\,K^{-1}}$	50–200 W	Wickless; compact; compatible with additive manufacture	Shorter useful distance; filling-ratio sensitivity; limited heritage	CubeSat and small-satellite demonstrations

Mechanically pumped two-phase loops

If the coolant evaporates in the cold plate, much of the load is carried as latent rather than sensible heat. Vapour then condenses at the sink and a pump returns liquid to the evaporators. Parallel evaporators can serve many components at nearly the same saturation temperature. In the configuration studied by Middelhuis and co-workers, four parallel evaporators operated stably between 50 and 200 W each; temperature differences remained below 3°C, and natural pressure-drop balancing maintained the flow split [5]. The reported pump-power requirement was about 40% lower than for a comparable single-phase loop. The advantage is real, but so are the added issues of instability, phase separation, vapour-quality control, and dry-out.

Cold-plate design

The cold plate is the thermal bridge between an electronic package and its loop. Its resistance includes conduction through the base, spreading from local sources to the channels, and convection from the wall to the coolant. Modern plates commonly use passages with hydraulic diameters of 0.5–2.0 mm. Reducing the channel size increases area and thins the thermal boundary layer, but also raises pressure drop. Inlet and outlet manifolds must divide the flow evenly; otherwise nearby channels receive too much coolant and remote channels become hot.

Cold-plate performance is often reported as $(T_{\mathrm{surface}} - T_{\mathrm{in}})/Q$. Series connection is limited by accumulated pressure drop and coolant warming, whereas parallel branches can serve many boxes if orifices or control valves balance the flow. Passive solutions remain attractive below approximately 500 W and 10 $\mathrm{W\,cm^{-2}}$. In the 1–10 kW and 50–200 $\mathrm{W\,cm^{-2}}$ range, pumped microchannel plates are normally required.

Table 2: Passive and active transport technologies at system level (indicative ranges). Sources: [1, 6].

Technology	Capacity	Control and operating power	Mass and complexity	Typical use
Heat pipe	50–500 W	Passive; no operating power	0.5–2.0 kg per transported kW; low complexity and no moving parts	Spreading and moderate-power electronics
Loop heat pipe	0.1–1 kW	Passive with limited set-point control; no operating power	1.0–3.0 kg per transported kW; compensation chamber and flexible lines	Long-distance transport and precision platforms
Single-phase pumped loop	1–50 kW per loop	About ±2°C with active control; 50–300 W per loop	2.0–5.0 kg per transported kW; pump, accumulator, valves, and controls	Distributed high-power equipment and large spacecraft
Two-phase pumped loop	1–50 kW per loop	Nearly isothermal evaporators; 30–200 W per loop	1.0–3.0 kg per transported kW; phase and stability management	High-power systems where mass flow must be limited

Liquid cooling and dielectric coolants

PCMs can buffer a pulse, but the heat must first leave the semiconductor. GaN amplifiers, SiC converters, and radar modules may exceed 500 $\mathrm{W\,cm^{-2}}$ locally. Liquid flowing close to the die can provide the missing source-level conductance. For a spacecraft, the fluid must also remain electrically safe after a seal failure, stable for the mission duration, and compatible with every wetted material [13, 14, 41].

Single-phase and two-phase operation

In single-phase cooling, the liquid warms continuously:

$$\dot{Q} = \dot{m} c_p (T_{\mathrm{out}} - T_{\mathrm{in}}). \tag{3}$$

Water microchannels often provide coefficients of 5,000–50,000 $\mathrm{W\,m^{-2}K^{-1}}$, while dielectric liquids more commonly fall in the 1,000–10,000 $\mathrm{W\,m^{-2}K^{-1}}$ range unless channel geometry and velocity are particularly aggressive. Single-

phase flow is easier to control and avoids dry-out and vapour management. Its drawback is the mass flow or temperature rise needed to carry a large load.

Boiling uses a latent heat of roughly 100–300 $\mathrm{kJ\,kg^{-1}}$, which is about 50–150 times the sensible heat associated with a 10°C rise for many relevant fluids. Nucleate boiling can produce coefficients of 20,000–100,000 $\mathrm{W\,m^{-2}K^{-1}}$ and reduce mass flow by factors of 10–50. Those benefits come with critical-heat-flux limits, unstable pressure and vapour quality, surface dry-out, and the need to condense or separate vapour. This review therefore retains both approaches but emphasizes single-phase cooling where stability and predictable operation are more important than maximum flux.

Coolant selection

No fluid is best in every respect. A defensible selection considers conductivity, viscosity, heat capacity, density, freezing and boiling points, dielectric strength, chemical lifetime, flammability, toxicity, ozone-depletion potential, and global-warming potential [13]. To avoid burying the comparison in inline numbers, Table 3 summarizes representative properties for commonly used dielectric candidates.

Table 3: Representative thermophysical properties of dielectric coolants for electronics cooling. Liquid properties are reported at approximately 25 °C. Values are representative rather than specification values and should be verified against current manufacturer data before hardware selection [13, 41]

Coolant example	T_b (°C)	ρ ($\mathrm{kg\,m^{-3}}$)	c_p ($\mathrm{kJ\,kg^{-1}\,K^{-1}}$)	k ($\mathrm{W\,m^{-1}\,K^{-1}}$)	μ (mPa s)	Notes
HFE-7000	34	1,400	1.30	0.075	0.45	Very low viscosity and low boiling point. Suitable for low-temperature heat transfer, but offers a limited single-phase temperature margin.
HFE-7100	61	1,510	1.18	0.069	0.58	Commonly studied dielectric coolant. Its thermal conductivity and heat capacity are substantially lower than those of water.
HFE-7200	76	1,430	1.21	0.068	0.61	Higher boiling point than HFE-7100, providing a larger single-phase operating-temperature margin.
Novec 649 fluoroketone	49	1,600	1.10	0.059	0.64	Low-GWP fluoroketone with good dielectric properties. Its low boiling point is particularly suitable for two-phase cooling.
FC-72 perfluorocarbon	56	1,680	1.10	0.057	0.64	Established electronics coolant with high electrical resistivity, but high GWP and a long atmospheric lifetime.

Single-phase experimental benchmarks

Single-phase dielectric cooling is attractive for electrical safety, but its thermophysical limits (especially lower conductivity and heat capacity than water) can make it difficult to reach the highest heat-flux regimes relevant to orbital power electronics without incurring significant pressure drop and pumping power [13, 41]. Minichannel experiments with HFE-7000 covered heat fluxes of about 50–1300 $\mathrm{kW\,m^{-2}}$ at mass fluxes of 125–280 $\mathrm{kg\,m^{-2}s^{-1}}$ [14]. At the upper mass flux, the reported 1300 $\mathrm{kW\,m^{-2}}$, or 130 $\mathrm{W\,cm^{-2}}$, was removed with less than a 25°C wall-to-fluid difference [14].

Microchannels and microfluidics (integrated within liquid cooling)

Microchannels bring liquid within a fraction of a millimetre of the heat source and supply far more wetted area per unit volume than a conventional cold plate. They can therefore remove heat fluxes in the kilowatt-per-square-centimetre class, provided that pressure drop, fabrication, and flow distribution are controlled [15, 17, 18, 19].

Why small channels work

Surface area per volume varies inversely with characteristic size. A passage with a 100 μm hydraulic diameter offers roughly 40,000 $\mathrm{m^2\,m^{-3}}$, compared with 400 $\mathrm{m^2\,m^{-3}}$ for a 10 mm tube. For laminar flow, a useful scaling is $h \sim k/D_h$, so the smaller passage also raises the convective coefficient. With water, a 100 μm channel may provide around 6000 $\mathrm{W\,m^{-2}K^{-1}}$ in fully developed laminar flow, ten times the value estimated for a 1 mm channel. Near the inlet, the still-developing thermal layer can raise this coefficient by a further factor of two to five. Since laminar entrance length scales approximately as $L_e \sim 0.05 ReD_h$, a substantial portion of a short microchannel may remain in this favourable developing region.

Channel and fin geometry

Straight rectangular passages are easy to predict and manufacture. Widths of 50–500 μm, depths of 100–1000 μm, and hydraulic diameters of 80–600 μm are typical in the reviewed work. Their main thermal weakness is streamwise warming. For a heated area A,

$$\Delta T = \frac{q'' A}{\dot{m} c_p}. \tag{4}$$

At 500 $\mathrm{W\,cm^{-2}}$ over 1 $\mathrm{cm^2}$ and a water flow of 0.1 $\mathrm{L\,min^{-1}}$, the source estimates a 20°C inlet-to-outlet rise. More flow reduces the gradient but increases pump demand; dividing the heated length into shorter feed segments is

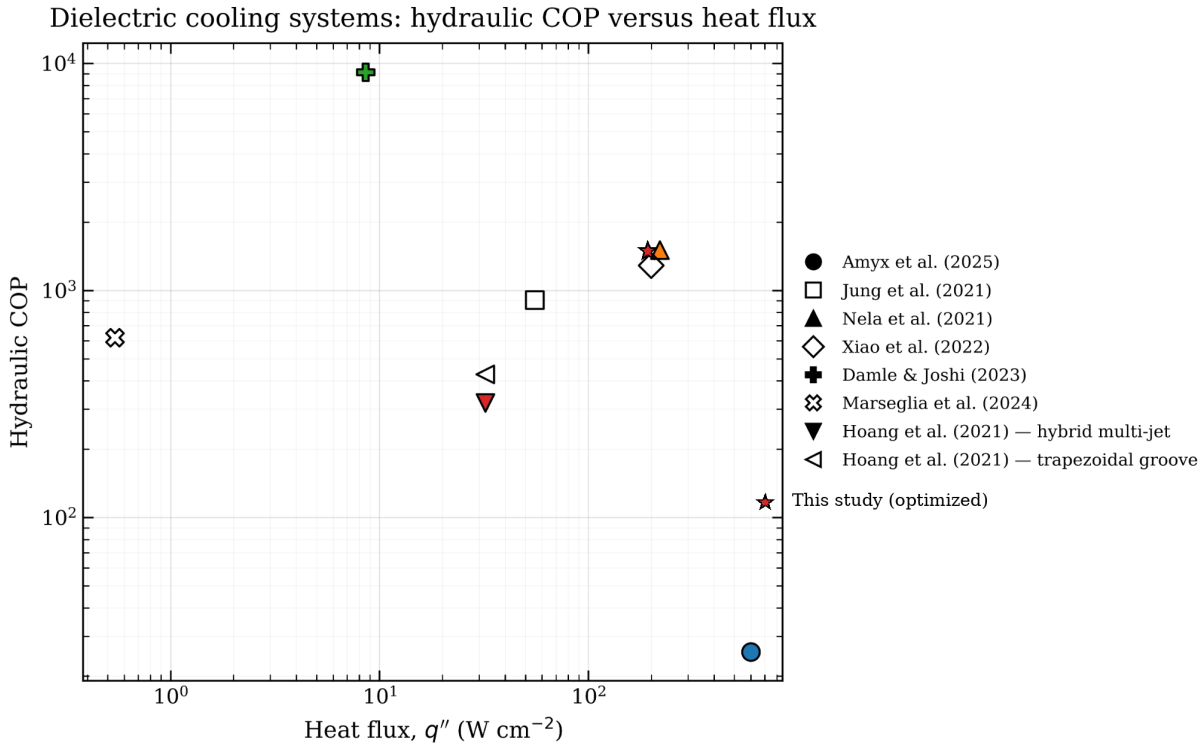

Each marker denotes one experimental study. Selected points correspond to single-phase or pre-ONB operation.

Figure 3: Hydraulic COP versus heat flux for selected dielectric-cooling studies (single-phase or pre-ONB points, as reported by the original sources). The point for this work must include all four measured flow-rate conditions and be computed from synchronized laboratory records. Sources: [56, 57, 58, 59, 60, 61, 62, 63].

another option.

Pin fins keep the boundary layer from becoming fully developed. Cylindrical pins of 50–200 μm diameter and 200–500 μm height can raise the heat-transfer coefficient by roughly 50–150% over straight channels [16]. Staggered arrays generate wakes and horseshoe vortices, producing reported water-side coefficients of 15,000–40,000 W m^{-2}K^{-1}. Their drag is lower than a long serpentine passage of comparable enhancement but still significant: reviewed designs with 10–20 rows reached approximately 50–150 kPa at 2 m s^{-1}.

Manifolds and maldistribution

A good channel array can fail thermally if the header feeds it unevenly. In a simple inlet and outlet manifold, pressure changes along the header make the nearest paths easier to enter. The review reports flow-rate differences of two to five between channels in poorly designed arrays. With 50–200 passages, channel losses may reach 50–200 kPa while the manifold itself loses only 5–10 kPa; small dimensional errors are then enough to disturb the split. A width or depth error of ±5–10 μm may cause a ±10–20% resistance change and a ±30–60% flow-rate difference. Streamlined or bifurcating manifolds, shortened channels, and local restrictions can improve uniformity.

Performance metrics and materials

Area-normalized resistance is

$$R''_{\mathrm{th}} = \frac{(T_{\mathrm{source}} - T_{\mathrm{inlet}})A}{Q}, \tag{5}$$

and typically lies around 0.05–0.30 cm^{2}K W^{-1} for single-phase microchannels, compared with about 0.3–0.8 cm^{2}K W^{-1} for conventional plates. Hydraulic power is

$$P_{\mathrm{pump}} = \frac{\Delta P\,\dot{V}}{\eta}, \tag{6}$$

and pressure drop may be estimated from

$$\Delta P = f\,\frac{L}{D_h}\,\frac{\rho U^2}{2}. \tag{7}$$

For laminar rectangular flow, $f = C/Re$, with C commonly between 50 and 65. Reported microchannel drops of 50–300 kPa over 10–20 mm correspond to approximately 0.5–5 W of pump power at 0.1–0.5 L min^{-1} and 60% efficiency.

Copper offers 385–400 $\mathrm{W\,m^{-1}K^{-1}}$ conductivity and is readily micro-machined, but its coefficient of thermal expansion, $16.5 \times 10^{-6}\,\mathrm{K^{-1}}$, is poorly matched to silicon at $2.6 \times 10^{-6}\,\mathrm{K^{-1}}$ and GaN at $5.6 \times 10^{-6}\,\mathrm{K^{-1}}$. Copper–tungsten composites lower the mismatch to roughly 6–9 $\times 10^{-6}\,\mathrm{K^{-1}}$ while retaining 180–250 $\mathrm{W\,m^{-1}K^{-1}}$ conductivity. They are more difficult and expensive to manufacture but can reduce cyclic stress by factors of two to three [17].

Table 4: Representative performance values for microchannel heat sinks and liquid cold plates. The listed thermal resistance is area-normalized and is reported as provided in the cited studies. Because the reference temperature, heated-area definition, and temperature-averaging method may differ between studies, these values should be verified before making a quantitative comparison. The general technology overview follows [15].

Technology	Heat flux ($\mathrm{W\,cm^{-2}}$)	Area-normalized thermal resistance ($\mathrm{cm^2\,K\,W^{-1}}$)	Pressure drop (kPa)	Coolant	Main feature
Copper straight microchannels [15]	450	0.18	85	Water	Straight channels with a hydraulic diameter of approximately 200 µm and a simple manifold.
Copper serpentine microchannels [17]	520	0.15	145	Water	Serpentine flow path providing heat-transfer enhancement through bend-induced secondary flow and mixing.
Copper micro-pin-fin array [16]	680	0.12	95	Water	Staggered 150 µm pin fins producing disrupted boundary layers and three-dimensional flow.
SiC microchannels [18]	1030	0.09	180	Water	100 µm channels integrated close to the GaN heat source to reduce conduction and spreading resistance.
Silicon micro-pin-fin heat sink [19]	1250	0.08	280	Water	40 µm silicon pins fabricated using deep reactive-ion etching for near-junction cooling.
CuW bifurcating microchannels [17]	580	0.16	95	Water	Tree-like flow distribution combined with improved coefficient-of-thermal-expansion matching.
Aluminium offset-fin cold plate [15]	380	0.22	65	PAO	Interrupted fins enhance boundary-layer redevelopment, while the aluminium substrate reduces component mass.
Additively manufactured copper cooler [19]	620	0.14	125	Water	Integrated three-dimensional manifold containing features of approximately 200 µm.

Notes: The area-normalized thermal resistance is generally expressed as

$$R''_{\mathrm{th}} = \frac{(T_{\mathrm{source}} - T_{\mathrm{in}})\,A_{\mathrm{h}}}{Q},$$

where T_{source} is the reported source temperature, T_{in} is the coolant inlet temperature, A_{h} is the heated area, and Q is the removed heat load. The exact temperature definition may differ among the cited studies. Heat flux, area-normalized thermal resistance, and pressure drop should therefore not be compared independently without also considering the heated area, flow rate, coolant properties, and allowable temperature rise.

The best laboratory examples exceed $1\,\mathrm{kW\,cm^{-2}}$ and fall below $0.1\,\mathrm{cm^2K\,W^{-1}}$. That makes microchannels promising for GaN, SiC, and phased-array modules, but a spacecraft implementation must still resolve pump power, radiation resistance, manufacturing cleanliness, leakage, and long-term compatibility with a dielectric fluid.

Topology optimization of microchannels and heat sinks

Most conventional heat sinks begin with a familiar shape—straight channels, a serpentine path, or a regular pin-fin array—and improve it by changing a few dimensions. That approach searches only the family of geometries selected at the outset. Topology optimization (TO) starts from a design region and asks where solid and fluid should be placed to meet a thermal-hydraulic objective. The resulting channels may branch, vary in width or depth, and place fins non-uniformly. Published studies report thermal-resistance improvements of about 15–30% and pressure-drop reductions of 20–50% relative to their chosen conventional baselines [20, 21, 22, 23, 24, 25].

Mathematical formulation

A common density-based formulation assigns a continuous design variable $\gamma(\mathbf{x})$ between 0 and 1 to every point. One end represents fluid and the other solid. Penalization discourages intermediate "grey" material so that the converged design can be interpreted as channels and walls. In a SIMP-type interpolation, properties may be written as

$$k(\gamma) = k_f + \gamma^p (k_s - k_f), \tag{8}$$

$$\alpha(\gamma) = \alpha_{\max} \frac{(1-\gamma)^q}{\gamma + \epsilon}, \tag{9}$$

where k_f and k_s are fluid and solid conductivities, α is a Darcy resistance used to suppress flow in solid regions, and p and q are usually chosen around 3–5 [20]. Temperature and velocity are obtained from coupled energy and flow equations, for example

$$\nabla\cdot[k(\gamma)\nabla T] + \rho_f c_p \mathbf{u}\cdot\nabla T = q''', \tag{10}$$

$$\rho_f(\mathbf{u}\cdot\nabla)\mathbf{u} = -\nabla p + \mu\nabla^2\mathbf{u} - \alpha(\gamma)\mathbf{u}, \tag{11}$$

$$\nabla\cdot\mathbf{u} = 0. \tag{12}$$

Adjoint sensitivities make it possible to update thousands or millions of design variables without solving a separate perturbation problem for each one.

Two- and three-dimensional models

A two-dimensional model treats channel depth as fixed and optimizes the plan view. It may be 10–100 times less expensive than a full three-dimensional calculation and is often reasonable when depth is much greater than width and velocity is mainly parallel to the heated plane. Such models have found branching networks, asymmetric routes, and unequal fin spacing with 15–25% lower resistance than straight-channel baselines [20]. They cannot capture out-of-plane vortices, vertical flow, depth changes, or multilayer passages. Those effects matter in pin fins, offset strips, tight bends, and manifold jets. Three-dimensional TO is therefore more expensive but can identify variable-depth channels and mixing structures that a planar model cannot represent.

Competing objectives and Pareto selection

The geometry that produces the lowest temperature is rarely the geometry that requires the least pumping. More wetted area and faster flow normally improve heat transfer while making the path harder to drive. A scalarized problem may use

$$J = w_1 R_{\mathrm{th}} + w_2 \Delta p. \tag{13}$$

Changing (w_1, w_2) generates a Pareto frontier. No point on that frontier can improve one objective without worsening the other. Instead of presenting a single supposedly universal optimum, the frontier allows a spacecraft designer to choose a point that respects the available pump power and temperature margin [28, 29].

Manifold and transient optimization

The inlet manifold is part of the thermal design, not merely plumbing. Gilmore and colleagues found that reshaping the transition from inlet to header, removing stagnant corners, and smoothing channel entrances reduced manifold pressure drop by about 17% relative to a rectangular form [27]. Concurrent optimization can determine the header and small channels together, producing a gradually changing cross-section that maintains velocity and distributes pressure more evenly [21].

Steady-state objectives are insufficient for duty-cycled payloads. A design with very little thermal mass may perform well at equilibrium but heat quickly during a short burst; an overly massive design may respond too slowly or carry unnecessary weight. Transient TO retains the storage term,

$$\rho c_p \frac{\partial T}{\partial t} + \nabla \cdot (k \nabla T) + \rho c_p \mathbf{u} \cdot \nabla T = q'''(t), \tag{14}$$

and can minimize either the maximum temperature or an integrated measure such as

$$J = \int_0^{t_f} \int_\Omega (T - T_{\mathrm{ref}})^2 \, \mathrm{d}\Omega \, \mathrm{d}t. \tag{15}$$

Gorman and co-workers used this kind of formulation for actively cooled microvascular materials, balancing steady resistance against transient capacitance [22].

Manufacturing constraints

An unconstrained optimizer may produce isolated islands, extremely thin fins, sealed cavities, or internal overhangs that cannot be inspected or fabricated. Length-scale filters and projection methods can impose minimum channel and wall sizes. For micro-milling with a $100\,\mu$m cutter, the source suggests design minima of roughly 150–200 μm. Additive-manufacturing constraints can similarly limit overhang angle, unsupported span, and minimum powder-clearance opening [23]. A manufacturable design may not achieve the unconstrained numerical optimum, but it is more valuable than a fragile geometry that cannot be built or qualified.

Reported benchmarks

Zhou and colleagues reported a topology-optimized manifold microchannel sink carrying up to $1570\,\mathrm{W\,cm^{-2}}$ at 11.9 kPa pressure drop [25]. At $687\,\mathrm{W\,cm^{-2}}$ and $0.4\,\mathrm{L\,min^{-1}}$, the area-normalized resistance was $0.067\,\mathrm{cm^2K\,W^{-1}}$, 35% below the $0.103\,\mathrm{cm^2K\,W^{-1}}$ baseline quoted in the study. The corresponding conventional drop was 21.6 kPa, so the optimized value was about 45% lower. Channel-to-channel inlet-flow discrepancy stayed below 6%. The paper reported a coefficient of performance of 78,273 at $687\,\mathrm{W\,cm^{-2}}$. Because different authors define COP differently, this number should be used only with the study's own definition and pumping-power boundary.

Across the wider literature, the gains depend heavily on baseline, objective, and constraints: the source review gives resistance reductions from 6.7% to 35% and pressure-drop reductions from 5.2% to 79.2%. Table 5 preserves the comparative values while making the cited baseline explicit.

Table 5: Selected topology-optimization results for microchannel heat sinks.

Study	Baseline	Optimized concept	Thermal result	Hydraulic result	Main contribution
Zhou et al. [25]	Rectangular manifold channels	Three-dimensional variable-depth manifold	35% reduction, 0.103 to 0.067 $\mathrm{cm^2K\,W^{-1}}$	45% reduction, 21.6 to 11.9 kPa	COP 78,273 and below 6% inlet-flow discrepancy
Yang et al. [24]	Inline pin fins	Non-uniform optimized micro-pin-fin field	6.7–7.9% reduction	5.2–12.4% reduction	11–12% better uniformity; performance-evaluation criterion 1.52
Gilmore et al. [27]	Rectangular manifold	Stagnation-suppressing manifold and nozzle	22.4% reduction in the reported comparison	Up to 79.2% with nozzle optimization	About 17% pressure reduction from manifold reshaping alone
Geng et al. [21]	Simple header	Concurrent multiscale manifold and channels	Not stated in source table	Not stated in source table	Flow uniformity within approximately ±5%
Duan et al. [28]	Uniform pin fins	Multi-objective non-uniform fin field	0.12 $\mathrm{cm^2K\,W^{-1}}$ at 85 kPa or 0.18 at 45 kPa	Pareto trade-off	Explicit thermal-hydraulic frontier
Xiong et al. [29]	Conventional multilayer cooler	Topology-optimized multilayer composite	Not stated in source table	Not stated in source table	48.86% lower thermal stress and 10,000-cycle assessment
Wang et al. [30]	Straight passages	TO translated into machinable fins	85–92% of predicted improvement retained	Not stated	Fabrication-aware geometry translation
Ozguc et al. [31]	Solid fins	Homogenized permeable structures	15–20% reduction	Not stated	Spatially varying porosity suited to additive manufacture

Heat storage: phase-change materials

Radiators are sized around a sustainable average load, yet many payloads produce short bursts far above that average. A phase-change material (PCM) can absorb the excess during the high-power interval and return it to the loop later, when generation has fallen below radiator capacity [10, 11, 12].

Storage mechanism and sizing

Ordinary sensible storage follows $Q = mc_p\Delta T$. During melting, a PCM stores approximately

$$Q = mL, \tag{16}$$

at a nearly constant melting temperature, where L is latent heat of fusion [10].

Materials and spacecraft integration

Organic PCMs (paraffins and fatty acids) cover melting temperatures from roughly 0 to 80°C with latent heats of 150–250 $\mathrm{kJ\,kg^{-1}}$, but low conductivity. Metallic PCMs and eutectics conduct better and offer higher volumetric storage at the cost of higher density and more demanding containment [10, 12].

Table 6: Representative PCM properties and performance reported in the literature. Sources: [10, 11, 12].

Material	Melting point	Latent heat	Conductivity	Density	Suitable use and enhancement	Reported observation
RT12 paraffin	12°C	160 $\mathrm{kJ\,kg^{-1}}$	0.2 $\mathrm{W\,m^{-1}K^{-1}}$	880 $\mathrm{kg\,m^{-3}}$	Low-temperature buffer; fins/foam/inserts	Example electronics tests in [11]
RT22 paraffin	22°C	190 $\mathrm{kJ\,kg^{-1}}$	0.2 $\mathrm{W\,m^{-1}K^{-1}}$	880 $\mathrm{kg\,m^{-3}}$	Moderate electronics temperature; lattice/foam	Example electronics tests in [11]
Gallium	29.8°C	80 $\mathrm{kJ\,kg^{-1}}$	29 (liq.)	6095 $\mathrm{kg\,m^{-3}}$	Fast transient storage; compliant containment	Freezing expansion and containment needs [12]

Heat rejection: space radiators

Whatever method is used inside the spacecraft, almost all waste heat is ultimately rejected radiatively. For a surface with area A, absolute temperature T, and hemispherical emittance ε, the ideal radiated power is

$$Q = \varepsilon\sigma AT^4, \tag{17}$$

where σ is the Stefan–Boltzmann constant [1].

Fixed and deployable radiators

Fixed radiators are typically aluminium honeycomb panels with high-emittance, low-solar-absorptance outer surfaces and bonded/brazed tubes. Deployable radiators trade added mechanisms and plumbing complexity for increased radiating area within the launch envelope [1, 46].

Variable-emittance and fault-tolerant concepts

Variable-emittance coatings and mechanically modulated surfaces can provide passive or semi-active turn-down without changing area, while fault-tolerant manifolds and parallel circuits improve resilience to leaks [50, 51, 52].

Table 7: Comparison of spacecraft radiator concepts (indicative ranges). Sources: [1, 45, 46, 47, 48, 50, 51, 52].

Concept	Rejection and emittance	Control method	Resilience and specific mass	Main application
Fixed panel	$\sim$300–400 W m^{-2} at 300 K; $\varepsilon \approx 0.8$–0.9	None	Line failures can disable area; $\sim$2–4 kg m^{-2}	Heritage missions
Deployable multi-panel	Similar areal rejection; larger area available	Deployment sequencing	Panel-level degradation tolerance; $\sim$3–6 kg m^{-2}	High-power, volume-limited
VO_2 variable emittance	ε modulation (reported $\sim$ 0.3–0.9)	Passive response	Coating ageing risk; $\sim$2–4 kg m^{-2}	Transient-load regulation
Hybrid parallel circuits	Similar areal rejection	Isolation / bypass	High retained capacity after a leak; $\sim$3–5 kg m^{-2}	High-reliability systems

Aerospace-relevant COP benchmark

The broader comparison expands the screening envelope beyond dielectric-only systems. It includes aerospace- and defense-enabling coolers as well as terrestrial high-heat-flux analogues. The purpose is to show what has been demonstrated at the device or cold-plate boundary, not to imply that every point is suitable for spacecraft use.

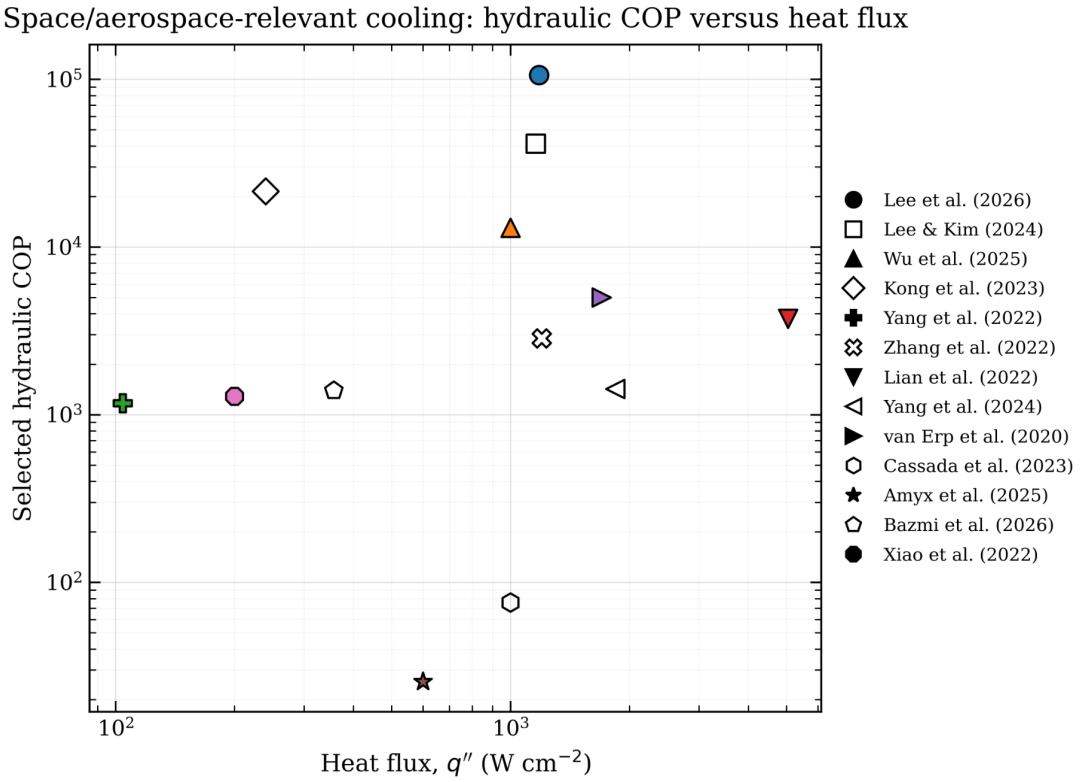

Tier B: aerospace/defense enabling systems; Tier C: terrestrial high-heat-flux analogs. No Tier-A space-test point is present.

Figure 4: Hydraulic COP versus heat flux for the selected space/aerospace-relevant and terrestrial high-heat-flux studies. Sources: [64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 56, 74, 59]. No plotted point is claimed as a complete flight-qualified thermal-control system.

The space/aerospace-relevant comparison in Figure 4 occupies a higher heat-flux band, extending from about 10^2 to several 10^3 W cm^{-2}, with COP values spanning more than three orders of magnitude. The strongest apparent points combine kilowatt-per-square-centimetre-class heat flux with COP above 10^4, but they should not be interpreted as flight-qualified solutions. Conversely, low-COP outliers can still be valuable if they demonstrate near-junction cooling, dielectric compatibility, compact packaging, or another constraint not encoded on the two axes.

Implications for the present work

The present topology-optimized cooler should be placed on these maps only after the same reporting boundary has been applied to every test point. At minimum, each plotted point should retain the source value of removed heat,

pressure drop, volumetric flow rate, heated area, coolant, operating regime, and uncertainty. A defensible company comparison should then report: (i) hydraulic COP using measured Δp and $\dot{V}$; (ii) a pump-electrical COP using the measured or assumed pump efficiency; and (iii) a system COP or power balance that includes the external loop when a spacecraft-level claim is made. This three-level reporting prevents a highly efficient cold plate from being mistaken for a complete high-efficiency thermal-control system.

Research Gaps and Future Directions

What is already mature

Conventional and loop heat pipes have decades of flight heritage. The review summarizes effective conductivities of 10,000–100,000 $\mathrm{W\,m^{-1}K^{-1}}$, transport lengths of 0.1–2 m for heat pipes and 5–10 m for LHPs, and heat-flux capability near 10–50 $\mathrm{W\,cm^{-2}}$ for ammonia pipes or 25–35 $\mathrm{W\,cm^{-2}}$ for selected LHP evaporators [7, 8, 9]. Variable-conductance devices can regulate by approximately ±2–5°C across 10–200 W. Ground tests above 100,000 hours and missions beyond 15 years support their role as the passive baseline, but their evaporator flux is below that required by the most concentrated electronics.

Pumped loops are also mature at system scale. The ISS external thermal-control system demonstrates tens of kilowatts of ammonia-loop rejection and more than two decades of service [6]. The source reports 98.5% availability, 35–70 kW rejection, and pump demand of approximately 0.8–1.2 kW per loop. Internal water–glycol or fluorocarbon buses commonly serve boxes of 50–500 W at only 0.5–2 $\mathrm{W\,cm^{-2}}$. Two-phase loops may reach 5–15 $\mathrm{W\,cm^{-2}}$ at lower flow but have less operational heritage because of stability and management complexity.

Radiators and PCM stores complete the mature system picture. Fixed panels commonly reject approximately 150–300 $\mathrm{W\,m^{-2}}$ near 300 K, with source values of 3–6 $\mathrm{kg\,m^{-2}}$; deployable panels reduce launch volume and can reach 200–400 $\mathrm{W\,m^{-2}}$ at 2.5–4.5 $\mathrm{kg\,m^{-2}}$ [45, 46, 47, 48]. Variable-emittance surfaces promise a turndown ratio near 3:1, and hybrid deployment/emittance concepts may eventually save 40–60% of radiator mass, although their readiness is lower. PCM flight and ground work reports peak reductions of 15–25°C and duty-cycle extensions of 30–50% for some applications [42, 43, 49, 53, 55].

The unresolved combination

The strongest recent TO results were obtained mainly with water, using approximately $\rho = 1000\,\mathrm{kg\,m^{-3}}$, $c_p = 4180\,\mathrm{J\,kg^{-1}K^{-1}}$, $k = 0.6\,\mathrm{W\,m^{-1}K^{-1}}$, and $\mu = 0.001\,\mathrm{Pa\,s}$. A few papers refer only to a generic liquid. These assumptions suit sealed terrestrial cold plates, but they are not neutral. A spacecraft system may require resistivity above $10^{10}\,\Omega\,\mathrm{cm}$ so that a leak does not short a power bus. Deionized water may begin around 10^5–$10^6\,\Omega\,\mathrm{cm}$ and fall to 10^3–$10^4\,\Omega\,\mathrm{cm}$ after contamination. It cannot provide the same fault tolerance.

Candidate HFE, fluoroketone, and perfluoropolyether fluids typically have density 1400–1900 $\mathrm{kg\,m^{-3}}$, heat capacity 1100–1300 $\mathrm{J\,kg^{-1}K^{-1}}$, conductivity 0.06–0.08 $\mathrm{W\,m^{-1}K^{-1}}$, and viscosity 0.0004–0.0012 Pa s [13, 41]. Relative to water, these differences can shift the optimum microchannel geometry and flow split; designs optimized under water assumptions should therefore be re-optimized with the intended dielectric properties.

This matters particularly in space because both available electrical power and allowable parasitic losses are constrained by the mission power budget, while mass and volume are constrained by launch and spacecraft integration. Quantified cost and power figures vary strongly by mission class and provider; any such values must be traceable to a specific published source before being used in design claims.

Recommended research programme

Four linked directions follow from the gap and from the dielectric-fluid and manufacturability challenges summarized above:

1. *Dielectric-fluid topology optimization.* Candidate fluids should be modelled with temperature-dependent properties and validated experimentally for the intended operating range [13, 41]. Geometry features that remain effective across a coolant family should be distinguished from coolant-specific features.
2. *Transient optimization coupled to PCM.* The objective should include the duty cycle, PCM phase change, charge/discharge dynamics, and radiator-side rejection constraints [22, 49].
3. *Thermal–hydraulic–structural optimization.* Thin features that improve heat transfer may be limited by launch loads and thermal cycling; coupled thermo-fluid-structural constraints and fatigue-relevant cycling assessments should therefore be included [29].
4. *Space-qualified fabrication.* Additive and microfabrication routes that enable integrated manifolds and variable-depth channels should be evaluated with fabrication-aware constraints and with qualification considerations (porosity, powder removal, roughness, outgassing, inspection) [23, 25, 33].

The proposed system emerging from the review uses three coordinated elements. First, a primary dielectric circuit and topology-optimized microchannels collect heat from GaN or SiC devices. Second, spatially variable fins and flow paths follow actual heat-flux maps with peak-to-average ratios of 5:1–10:1, targeting uniformity within

±5°C. Third, a PCM store and secondary radiator circuit shift energy from 10–30 minute peak periods into 30–90 minute recovery intervals. A primary temperature near 40–60°C would protect electronics while permitting a PCM melting range around 50–65°C; a radiator near 40–50°C could reject approximately 200–300 $\mathrm{W\,m^{-2}}$ with an optical-solar-reflector finish. These values are an initial design envelope, not a substitute for mission-specific modelling.

Experimental case study

Chapter overview

This chapter describes the development of a near-junction, single-phase microfluidic cooling platform for a compact gallium-nitride power device. The work was not limited to a numerical design exercise. It followed the complete route from a non-uniform power map, through topology optimization in COMSOL Multiphysics, to silicon microfabrication at EPFL's Center of MicroNanoTechnology, the design and 3D printing of the fluidic package, assembly on a dedicated printed circuit board, and finally thermal and hydraulic testing. A straight-channel device was retained as a reference so that the optimized geometry could be judged against a physically comparable baseline rather than against simulation alone.

The central design question was practical: could a flow field be shaped around the local heat-load distribution without paying an excessive pressure-drop penalty? The optimizer was therefore allowed to generate irregular solid features inside the cooling region, but intermediate material states were penalized so that the final geometry could be interpreted as either fluid or silicon and then fabricated. A projection continuation strategy sharpened the design as the solution evolved, with the projection slope parameter increased from $\beta = 1$ to $\beta = 32$.

The resulting device reduced the maximum thermal resistance by approximately 22% relative to the straight-channel reference. At comparable pumping power, the effective heat-transfer coefficient increased by approximately 17%, while the pressure-drop increase remained modest. Numerical predictions and measurements followed the same trends, with a maximum discrepancy of about 15% for the baseline and optimized samples.

Main contributions of the work

- Translation of the EPC23102 power map into a non-uniform cooling footprint, followed by a 256 µm unit-cell spacing study.
- Density-based topology optimization using the property-ramping parameters q_a and q_k, together with projection continuation in β, to suppress gray fluid–solid transition regions.
- Silicon fabrication at EPFL CMi, SEM inspection, and SOLIDWORKS/3D-printed packaging for PCB-level integration.
- Closed-loop thermal and hydraulic testing of straight and optimized samples, with comparison against COMSOL predictions.

1. Device selection and heat-load definition

The demonstrator was built around the EPC23102, a 100 V, 35 A integrated half-bridge. This component was selected because it combines a small footprint with a directly accessible top-side die surface. The exposed region is 2.60 mm by 4.47 mm, corresponding to an available cooling area of 11.622 mm^2, or 0.11622 cm^2. This small area makes the device thermally demanding, but it also makes it a useful platform for testing whether microfluidic cooling can remove heat close to the junction instead of carrying it through several packaging interfaces.

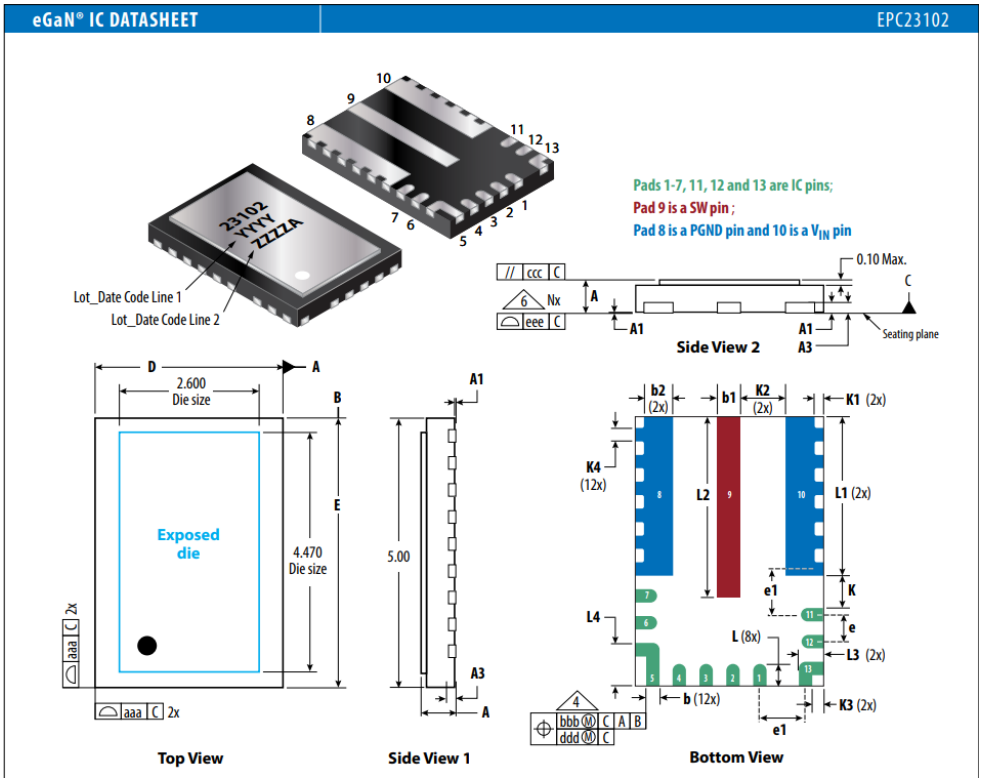

Figure 5: EPC23102 package and exposed-die dimensions used to define the available cooling footprint.

A nominal dissipated power of 35 W over the exposed area corresponds to a local heat flux of approximately 301 W/cm^2. The target quoted throughout the project as "about 300 W/cm^2" is therefore consistent with the device dimensions and the upper end of the experimental power range:

$$A_{\mathrm{die}} = 2.60\,\mathrm{mm} \times 4.47\,\mathrm{mm} = 0.11622\,\mathrm{cm}^2,$$
$$q'' = \frac{P_{\mathrm{diss}}}{A_{\mathrm{die}}} = \frac{35\,\mathrm{W}}{0.11622\,\mathrm{cm}^2} \approx 301\,\mathrm{W\,cm^{-2}}. \tag{18}$$

The surface heat load was not treated as perfectly uniform over the entire package. The power-conversion region carried the dominant load, whereas the control region dissipated roughly one order of magnitude less heat. This distinction affected both the optimization domain and the manifold concept. Rather than supplying every part of the die with an identical flow path, the design concentrated flow and heat-transfer area where the measured or estimated power map justified it.

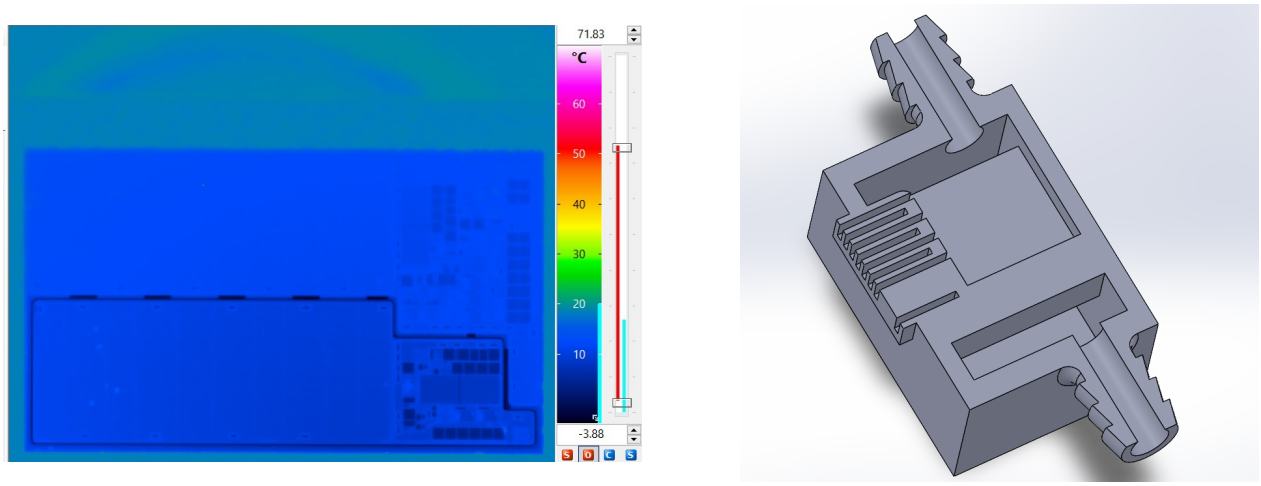

Figure 6: Power-map definition and the internal non-uniform routing developed to distribute coolant over the selected die region.

2. Numerical design methodology

2.1 Unit-cell screening

Before optimizing the full footprint, a smaller unit cell was used to establish a sensible manifold spacing. The inlet-to-outlet distance was varied and each case was evaluated using two competing indicators: pressure drop and thermal resistance. Shorter flow paths reduced thermal resistance but increased the hydraulic cost, whereas longer paths relieved the pressure requirement at the expense of surface temperature. The region where the two curves offered the most balanced response was selected rather than simply choosing the minimum of either quantity.

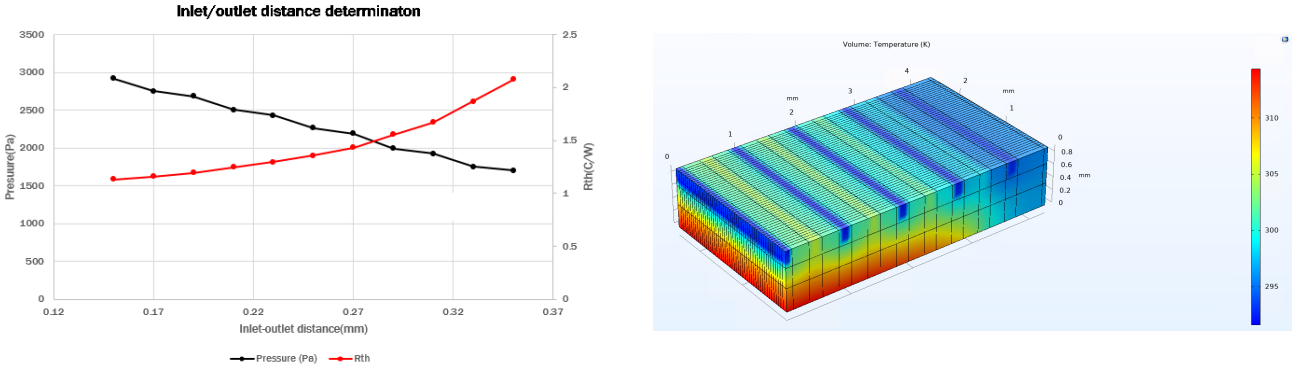

Figure 7: Unit-cell trade-off used to select a 256 µm inlet/outlet spacing and a representative COMSOL temperature field.

The selected inlet/outlet separation was 256 µm. This value was then carried into the manifold layout and the topology-optimization domain. The screening step mattered because it prevented the optimizer from solving an unrealistic local problem disconnected from the pressure losses and length scales of the intended package.

2.2 Coupled thermo-fluid model

The COMSOL model coupled laminar flow with heat transfer in the liquid and in the surrounding solid. The coolant was treated as a single-phase, incompressible dielectric liquid, and the silicon region conducted heat from the device surface toward the internal wetted structures. The numerical outputs required for design selection were the pressure difference between the fluid ports, the maximum device temperature, the spatially averaged temperature, and the corresponding thermal resistance. No boiling or latent-heat contribution was included.

$$R_{\mathrm{th,max}} = \frac{T_{\max} - T_{\mathrm{in}}}{P_{\mathrm{diss}}}, \qquad R_{\mathrm{th,avg}} = \frac{T_{\mathrm{avg}} - T_{\mathrm{in}}}{P_{\mathrm{diss}}}. \tag{19}$$

Boundary conditions were chosen to reproduce the available device footprint and the experimentally accessible mass-flow range.

2.3 Density-based topology optimization

The design region was described by a continuous density variable bounded between the two physical limits of fluid and solid. During the early iterations, intermediate values were permitted because they give the optimizer a differentiable design space. Left untreated, however, these gray regions would represent a fictitious mixture rather than a manufacturable channel or pillar. In the implemented property-ramping strategy, the two parameters were q_a and q_k. They were used to penalize intermediate fluid–solid states and drive the design toward physically distinct regions..

A smoothed projection was added on top of the raw density field. The projection slope β was initially set to 1, which allowed the material interface to move without making the optimization numerically brittle. It was then increased in stages up to $\beta = 32$. This continuation step progressively sharpened the interface and produced a nearly binary result. In practical terms, the irregular islands in the optimized map became discrete silicon features separated by continuous coolant passages.

Pressure drop and thermal resistance were considered together when comparing candidate geometries. Because the presentation does not provide the implemented objective function, its normalization, or weighting factors, no analytical objective equation is stated here.

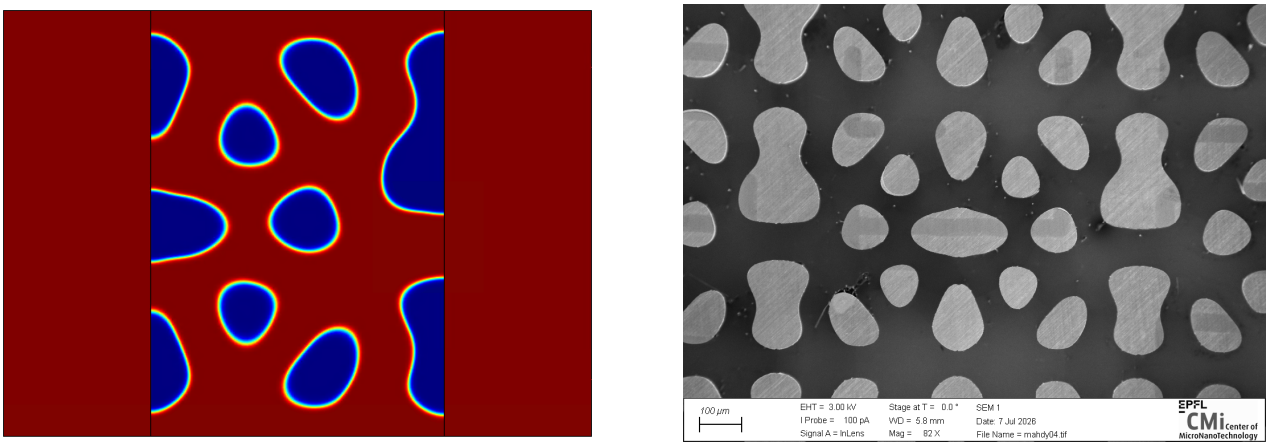

Figure 8: Final projected topology field and the corresponding fabricated distribution of irregular silicon features viewed from above.

The resulting shapes were not imposed as circles, fins, or a repeated lattice. Their size and orientation emerged from the local compromise between viscous dissipation, heat spreading, and hotspot suppression. This is the principal difference between the optimized sample and the straight-channel reference.

2.4 Adaptation of the reference topology-optimization methodology

The topology-optimization procedure used in the present work was developed from the density-based framework reported by Yan et al. [75] for liquid-cooled microchannel heat sinks. The reference work is important here not because its final geometry was copied, but because it provides a systematic method for allowing the solid–fluid distribution to evolve from an initially diffuse field into a nearly binary, manufacturable microchannel topology. The present study retained this numerical philosophy while replacing the square water-cooled benchmark geometry of Yan et al. with the actual EPC23102 cooling footprint, the experimentally relevant inlet/outlet arrangement, the non-uniform device power map, and the dielectric working fluid used in the laboratory campaign.

A central feature of the method is that the optimizer does not begin from predefined straight channels, pin fins, or a fixed branching pattern. Instead, every point inside the design domain is assigned a continuous design variable, denoted here by γ, with

$$0 \leq \gamma \leq 1, \tag{20}$$

where $\gamma = 1$ represents fluid and $\gamma = 0$ represents solid. Intermediate values are permitted during the optimization because they provide a continuous design space for gradient-based optimization. They do not, however, represent a physically manufacturable material. The numerical procedure must therefore progressively penalize these intermediate states and sharpen the fluid–solid interface as the solution develops.

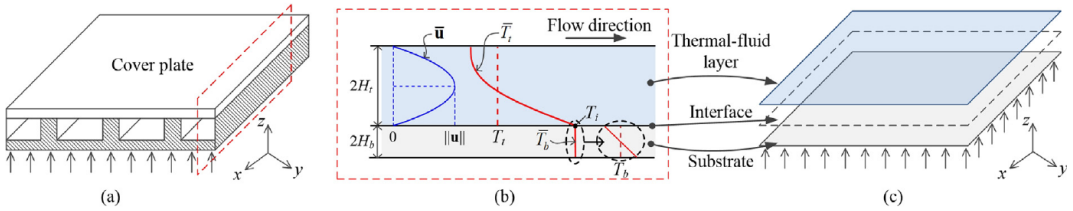

Figure 9: Two-layer microchannel heat-sink representation used by Yan et al. [75] as the basis for the reduced topology-optimization formulation. The thermal-fluid design layer is coupled to the solid substrate through an out-of-plane heat-transfer term. The present work adopts the same solid–fluid interpolation philosophy, while the actual geometry and boundary conditions are those of the EPC23102 cooler.

Material interpolation and penalization

Yan et al. used RAMP-type material interpolation to allow the same governing equations to represent both fluid and solid regions. The inverse permeability is interpolated according to

$$\alpha(\gamma) = \alpha_s + (\alpha_f - \alpha_s)\frac{(1+q_a)\gamma}{\gamma + q_a}, \tag{21}$$

where α_f is the inverse permeability assigned to the fluid and α_s is a much larger value used to suppress velocity in the solid. In the reference implementation, the permeability penalization parameter was kept fixed at

$$q_a = 0.1. \tag{22}$$

This interpolation gives a finite resistance in the fluid region and a very large resistance in the solid region, so the optimizer can gradually eliminate flow from locations that evolve toward solid material.

The effective thermal conductivity of the design layer is interpolated separately as

$$k_t(\gamma) = k_f + (k_s - k_f)\frac{1-\gamma}{1+q_k\gamma}, \tag{23}$$

where k_f and k_s are the thermal conductivities of the fluid and solid, respectively. The parameter q_k controls how strongly intermediate densities are penalized in the thermal problem. Unlike q_a, q_k is not kept constant. Yan et al. deliberately increase it during the optimization so that the first iterations remain numerically smooth and allow large changes in topology, while the later iterations increasingly discourage gray material and force the solution toward clearly separated coolant passages and silicon features.

The continuation sequence reported in the reference study is

$$q_k: \quad 1 \rightarrow 3 \rightarrow 10 \rightarrow 30, \tag{24}$$

with the value updated every 50 iterations or earlier if the current continuation stage has already converged. The maximum number of iterations in the reference implementation is 250, and convergence is identified when the maximum element-wise change in the design variable between two successive iterations becomes smaller than 0.01 [75]. This staged increase is particularly useful because applying a strong penalization from the first iteration can trap the solution in a poor local topology or cause numerical difficulties before a useful global flow path has formed.

For the present COMSOL model, the same continuation concept was combined with the projection strategy already described in Section 2.3. The projection slope was initially kept low, $\beta = 1$, so that the optimizer could move the material boundary smoothly. The projection was then strengthened progressively until $\beta = 32$. Thus, two mechanisms act in the same direction: increasing q_k penalizes thermally ambiguous intermediate material, while increasing β sharpens the projected design field. The solution from each stage is used as the initial condition for the next stage rather than restarting the optimization from a uniform field. This is essential for preserving the topology already developed while progressively improving its discreteness.

Table 8: Continuation parameters taken from the methodology of Yan et al. [75] and their role in the present implementation.

Parameter	Reference value / sequence	Status	Function in the optimization
γ	$0 \leq \gamma \leq 1$	Continuous design variable	Local fluid fraction; $\gamma = 1$ is fluid and $\gamma = 0$ is solid.
q_a	0.1	Fixed	Penalizes intermediate permeability and suppresses velocity as a region becomes solid.
q_k	$1 \rightarrow 3 \rightarrow 10 \rightarrow 30$	Increased by continuation	Progressively penalizes intermediate thermal conductivity states.
Continuation interval	50 iterations or convergence	Reference procedure	The converged field of one stage initializes the next stage.
Convergence criterion	$\max \lvert\gamma^{n+1} - \gamma^n\rvert < 0.01$	Reference procedure	Stops a continuation stage when the design field is no longer changing significantly.
Maximum iterations	250	Reference procedure	Upper bound used by Yan et al. for a complete optimization.
β	$1 \rightarrow 32$	Present COMSOL implementation	Gradually sharpens the projected density field and removes gray transition regions.
Density filter	$r_{\min} = L/100$	Reference procedure	Reduces checkerboard patterns and mesh-dependent features.
p-norm parameter	$p = 10$	Reference objective approximation	Smooth approximation of maximum substrate temperature in the reference formulation.

Iteration and continuation procedure

The practical optimization loop can be understood as a sequence of increasingly restrictive stages. At the beginning, the domain contains a large amount of fluid-like material and the interface is intentionally diffuse. Yan et al. initialized the designable square region with $\gamma = 0.9$ and prescribed the inlet and outlet regions as pure fluid, $\gamma = 1$. A density filter was applied before evaluating material properties in order to remove checkerboard patterns and reduce mesh dependence. In the present work, the same philosophy was used: the inlet and outlet passages were kept as non-design fluid regions, while only the internal cooling region was allowed to evolve.

At each optimization iteration, the current material field first defines the local inverse permeability and thermal conductivity. The laminar-flow problem is then solved, followed by the conjugate heat-transfer problem. The resulting pressure, velocity, and temperature fields provide the information required to update the design variables. Yan et al. used the Method of Moving Asymptotes (MMA) together with adjoint sensitivities. In the present implementation, this gradient-based procedure is carried out through the COMSOL Optimization interface, while the physical model is the project-specific laminar-flow and heat-transfer model described previously.

The reference paper minimizes the maximum substrate temperature, approximated by a smooth generalized p-norm with $p = 10$, while pressure drop can be imposed as a constraint when required. In the present work, temperature reduction and hydraulic penalty were evaluated together during design selection. The important point is that the material update is never based on temperature alone: every modification to the topology changes both the heat-transfer path and the flow resistance, and the next iteration is evaluated using the newly obtained thermo-hydraulic field.

A compact representation of the procedure used to generate the final topology is therefore

1. Initialize the design field with a predominantly fluid state and keep inlet/outlet regions fixed as fluid.
2. Apply the density filter and projection to obtain the physical design field used by the governing equations.
3. Solve the stationary laminar-flow problem and determine the pressure and velocity distribution.
4. Solve heat transfer in the fluid and silicon using the updated velocity field and the imposed non-uniform device heat load.
5. Evaluate the thermal objective and the hydraulic response, and compute the design sensitivity.
6. Update γ using the gradient-based optimizer.
7. Repeat until the design change is small for the current continuation stage.
8. Increase the thermal penalization q_k and strengthen the projection by increasing β; use the previous converged field as the initial condition for the next stage.
9. Continue until the final high-penalization/high-projection stage is converged, then convert the projected field into the binary geometry used for fabrication and three-dimensional validation.

The increase in penalization does not produce a perfectly monotonic objective history. Yan et al. showed that when q_k is increased, the objective and maximum temperature can exhibit a short transient or oscillation because the effective material properties have suddenly become more strongly penalized. The optimizer then re-adjusts the solid and fluid distribution and converges again. This behavior is visible in the reference optimization histories in Fig. 10; it is therefore not necessarily an indication of solver failure, provided that the solution stabilizes after each continuation step.

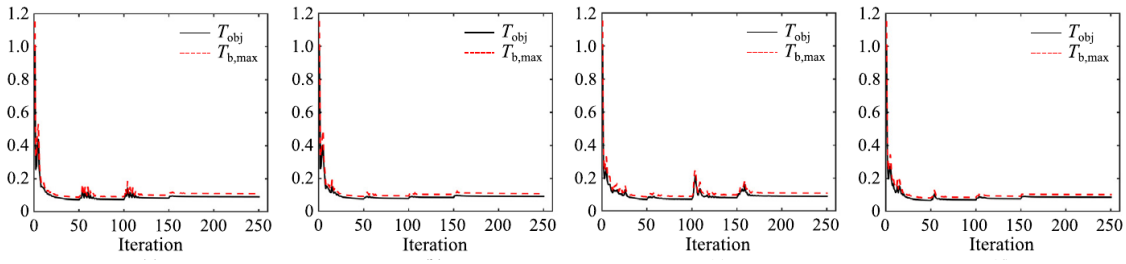

Figure 10: Optimization histories reported by Yan et al. [75]. Small disturbances in the objective occur when the conductivity penalization is increased, after which the design reconverges. This staged continuation is the methodological basis for progressively increasing the penalization and projection strength in the present COMSOL model.

Transfer of the method to the present microchannel cooler

Several aspects of the published benchmark were intentionally changed before applying the method to the present device. Yan et al. optimized a square silicon/water heat sink with a uniform heat flux and prescribed pressure-drop cases. The present problem instead uses the exposed EPC23102 footprint, a non-uniform power distribution, the experimentally selected 256 µm inlet/outlet spacing, and a single-phase dielectric coolant. Consequently, the optimized pattern in this thesis is not expected to reproduce the layouts in the reference paper. The value of

the reference study is the optimization strategy: continuous fluid–solid interpolation, filtering, continuation of penalization, and final projection.

The effect of changing the physical problem is significant. Yan et al. showed that increasing the available pressure drop, and therefore increasing the relative importance of advection, produces progressively more complex fin and channel networks. Their designs evolve from a small number of isolated streamlined fins at low pressure drop to dense branched networks at higher pressure drop. This trend is reproduced qualitatively in Fig. 11. It demonstrates why the final topology must be generated under the actual coolant properties and hydraulic boundary conditions rather than copied from a published optimized geometry.

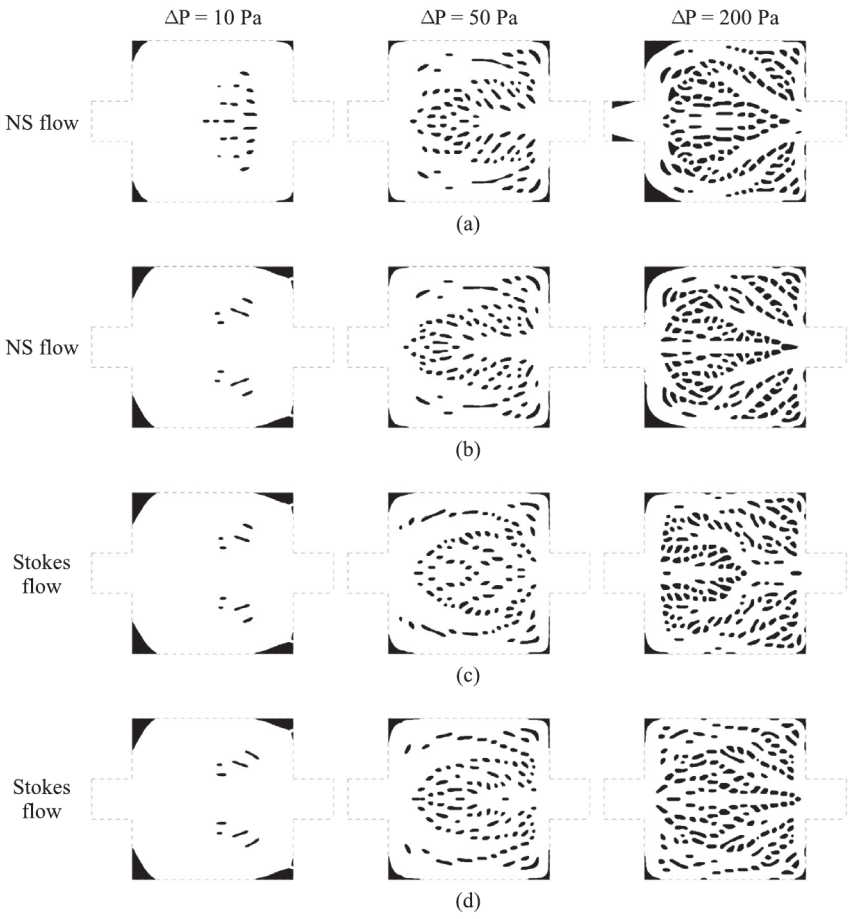

Figure 11: Examples from Yan et al. [75] showing the strong dependence of the optimized solid–fluid layout on the imposed pressure drop and flow model. Black regions are solid fins and white regions are coolant passages. The figure illustrates that topology is a result of the particular thermo-hydraulic operating point and cannot be transferred directly between different coolants or device geometries.

In the present design, the continuation process first produced broad, smooth regions that mainly established the preferred coolant paths. As q_k and the projection strength were increased, these diffuse regions separated into distinct silicon islands and continuous coolant passages. The final $\beta = 32$ solution was then interpreted as a binary geometry, cleaned where necessary to remove features below the fabrication length scale, and transferred into the three-dimensional device geometry. This projected field is the origin of the irregular silicon features shown in Fig. 8 and subsequently observed in the SEM images. Their shape is therefore not an imposed pin-fin pattern: it is the final discrete form of a continuous material-distribution problem solved through successive thermo-fluid analyses and increasingly strong penalization.

The reference study also provides an important justification for the validation strategy adopted here. Yan et al. used the reduced model to make topology optimization computationally practical, but subsequently rebuilt the optimized designs as full three-dimensional geometries and re-evaluated them with conjugate heat transfer in COMSOL Multiphysics. The same separation between optimization and validation is appropriate for the present work. The two-dimensional optimization identifies the material distribution efficiently, while the swept three-dimensional model and the fabricated samples are used to verify that the predicted thermal and hydraulic advantages persist after the real channel height, manifold geometry, and fabrication constraints are introduced.

In summary, the final optimized microchannel topology used in this thesis is best understood as an adaptation of the continuation-based density optimization of Yan et al., rather than as a direct reproduction of their heat sink. The parameters q_a and q_k control the physical interpolation of the solid–fluid mixture, while the projection parameter β controls the numerical sharpness of the interface. Starting from a weakly penalized, easily reconfigurable field

and increasing these parameters only after the preceding stage has stabilized allows the optimization to explore the design space first and impose manufacturability later. This staged procedure is what transforms the initial continuous density field into the final discrete topology used for the microfabricated cooler.

3. Manifold, cap, and PCB-level integration

The optimized silicon surface could not be tested as an isolated coupon. It required a package that brought coolant to the small active region, separated the inlet and outlet streams, attached reliably to the PCB, and left enough access for electrical connections and infrared imaging. The manifold and cap were therefore modeled in SOLIDWORKS after the flow-routing concept had been fixed.

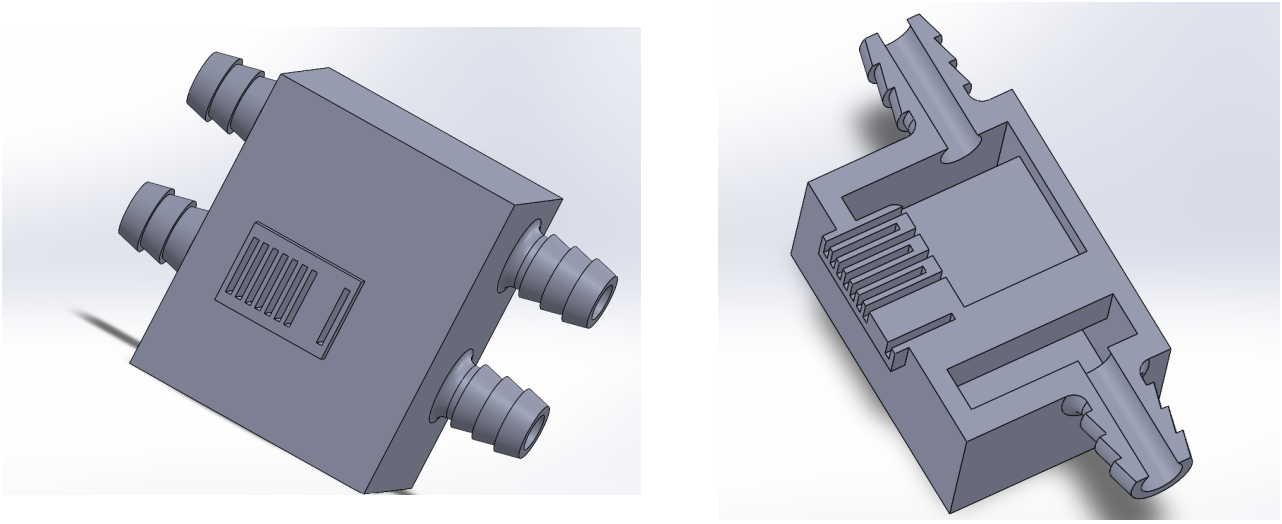

Figure 12: External and internal SOLIDWORKS views of the manifold/cap assembly designed for the PCB-mounted device.

The external model defined the hose interfaces and the footprint available above the board. Internally, the inlet and outlet passages were shaped to feed the non-uniform cooling region and to recover the fluid without short-circuiting between ports. The internal cavity also established the mechanical clearance around the active device and the microfabricated sample.

The parts were fabricated by 3D printing and used as the top-side fluidic package. This approach shortened the design cycle: changes to port orientation, cap thickness, or internal routing could be made in CAD and reprinted without modifying the silicon mask. The printed parts were then assembled over the device on the custom PCB. Black coating was applied to the relevant viewing surface to increase emissivity and reduce the sensitivity of the infrared measurement to the native reflectivity of the package.

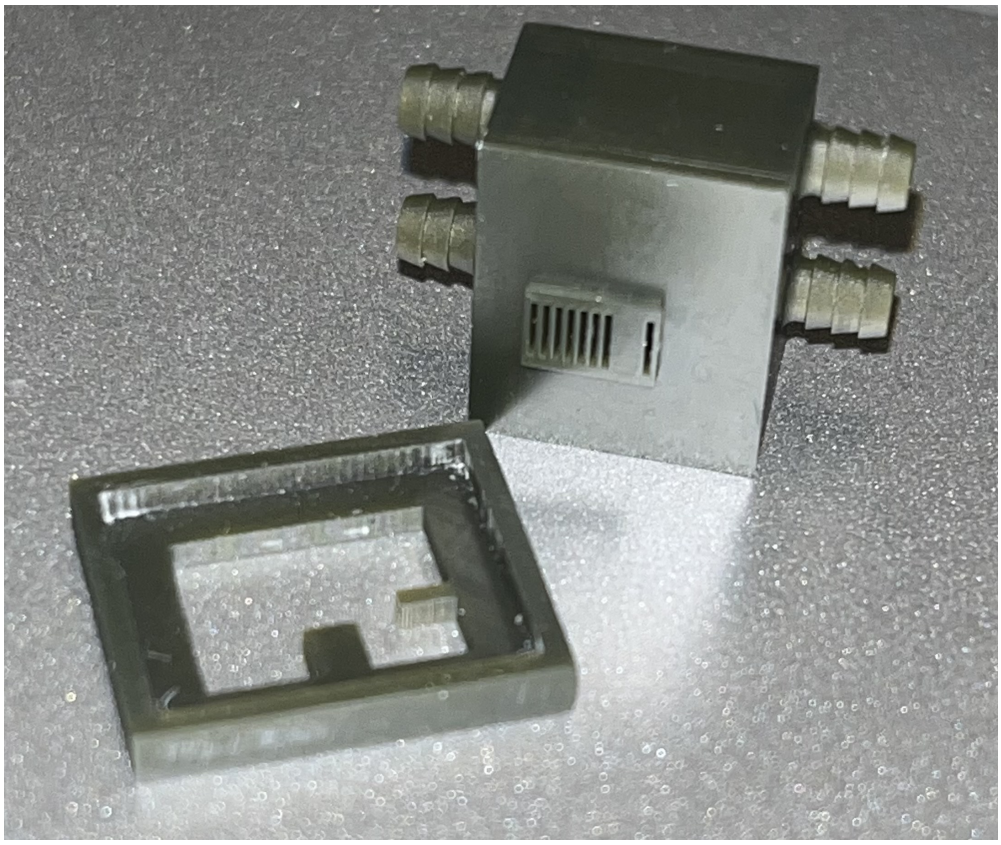

Figure 13: Fabricated manifold and alignment fixture. The separate fixture shown on the left was used during assembly to position the manifold accurately relative to the device surface. The manifold contains the main coolant inlet and outlet together with upstream and downstream pressure-measurement ports.

Figure 13 shows the fabricated manifold together with the separate alignment fixture used during assembly. The fixture positioned the manifold relative to the device surface so that the internal supply and return openings were correctly aligned with the microfabricated cooling region. This alignment was important for obtaining the intended flow distribution and for maintaining a uniform sealing perimeter around the device. Misalignment could partially obstruct the microchannel openings, create a coolant-bypass path, or reduce the available sealing area.

The four hose connections on the manifold perform two different functions. One connection on the inlet side carries coolant into the device, while the second connection on the same side acts as an upstream static-pressure tap. Similarly, one connection on the outlet side returns the coolant to the external loop, whereas the second provides the downstream pressure measurement. The two pressure taps were connected to the differential-pressure transducer, allowing the pressure loss across the packaged microchannel device to be measured directly. Because the sensing lines carry negligible flow during normal operation, they are more accurately described as pressure-measurement ports rather than as additional coolant inlets and outlets.

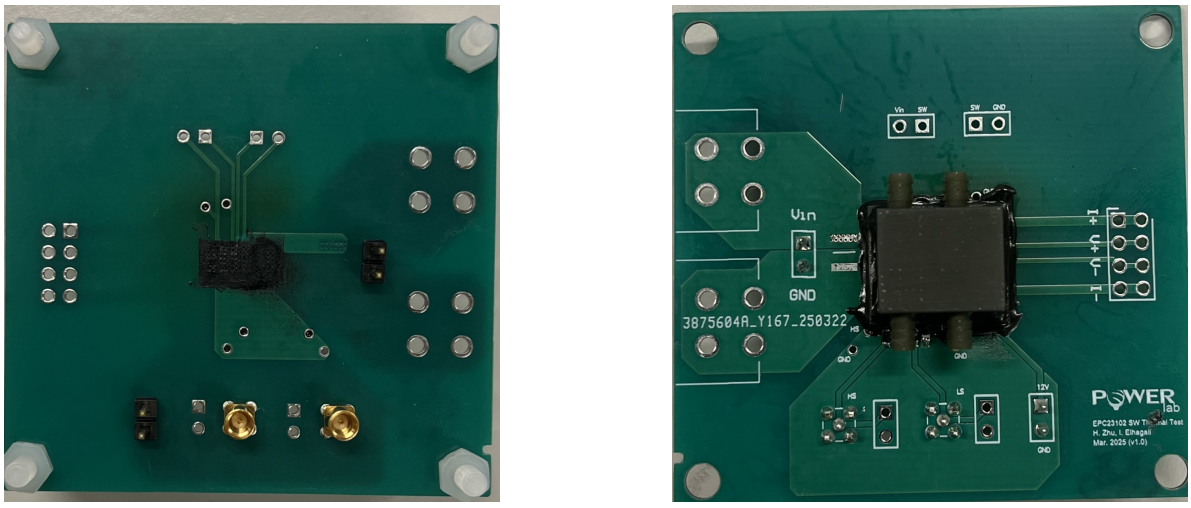

Figure 14: PCB-level test article before and after installation of the printed manifold package.

The package design was functional for the laboratory campaign, but it should be distinguished from a flight-

qualified assembly. The present study did not include vibration, launch-load, long-duration leak, thermal-vacuum, or radiation qualification. The package served to demonstrate the thermal concept and to create repeatable hydraulic boundary conditions around the device.

4. Silicon microfabrication at EPFL CMi

The numerical topology became useful only after it had been transferred into a three-dimensional silicon structure with open flow passages and mechanically continuous walls. This transfer was carried out at EPFL's Center of MicroNanoTechnology (CMi). The fabrication target was the silicon cooling surface associated with the EPC23102 test article. Both the straight-channel reference and the topology-optimized design were processed with the same main lithography, etching, stripping, and inspection sequence. Keeping the process route common to both samples was essential: otherwise, a difference in resist profile, etch depth, sidewall roughness, or surface cleanliness could be mistaken for a thermal benefit caused by the topology itself.

The complete sequence is summarized in Fig. 15. In brief, the surface was cleaned and dehydrated, primed with hexamethyldisilazane (HMDS), coated with AZ 10XT-60 positive photoresist, soft-baked and rehydrated, directly exposed with the MLA150, developed in AZ 400K, and then etched to a nominal channel depth of 300 µm in the SPTS Rapier deep-silicon etcher. The remaining photoresist was removed by wet stripping assisted by an ultrasonic bath, after which the samples were rinsed, dried, and inspected. The terminology is important: the hotplate operation before exposure was a *soft bake*, which removed coating solvent and stabilized the film; it was not the final hard bake sometimes used after development.

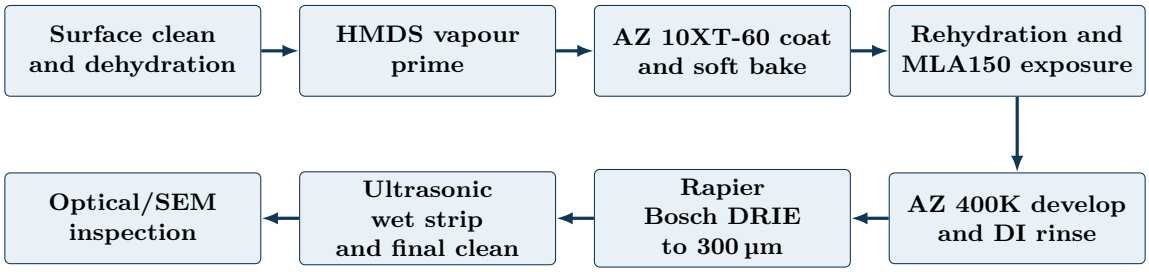

Figure 15: Fabrication route used to transfer the simulated fluid–solid topology to the silicon cooling surface.

4.1 Layout preparation and pattern polarity

The final projected COMSOL field was first converted from a continuous density map into a binary fabrication layout. Regions classified as fluid were assigned to the areas to be opened in the positive photoresist and subsequently etched, whereas regions classified as solid remained protected so that they would form silicon walls, fins, or islands. This polarity check was especially important for the optimized design: an inversion would exchange the intended flow passages and solid heat-spreading network.

The contour exported from the numerical model was cleaned before exposure. Isolated numerical pixels, extremely acute corners, and gaps below the practical lithographic length scale were removed or regularized without changing the main flow architecture. The resulting CAD pattern retained the inlet and outlet openings, the external sealing region, and the solid bridges needed to keep the remaining silicon mechanically connected. CMi notes that AZ 10XT-60 pattern transfer can require a critical-dimension bias of the order of 0.8 µm; consequently, the final mask should be understood as a fabrication-aware representation of the projected topology rather than an exact pixel-for-pixel copy of the raw density field [76].

Because the workpiece was chip-scale, safe and repeatable handling was also a process constraint. CMi's small-sample guidance specifies direct-write and coating precautions and a nominal minimum exposed area of 5 mm×5 mm; smaller or narrow samples must be processed with an approved carrier or pocket arrangement so that the coating remains uniform and the MLA150 pneumatic autofocus operates safely [80, 79]. The carrier arrangement used for a given run also had to provide adequate thermal contact during the long DRIE step.

4.2 Surface cleaning, dehydration, and HMDS priming

The starting silicon surface was cleaned to remove particles, organic residue, and handling contamination. The cleaning chemistry was selected to remain compatible with the device and any nearby metallization; aggressive wafer cleans cannot automatically be transferred to a finished power device. After cleaning, the sample was dried and dehydrated so that adsorbed water did not remain between the silicon and the resist.

HMDS vapour priming was then performed before spin coating. HMDS replaces hydrophilic surface hydroxyl termination with a more hydrophobic surface and therefore improves adhesion between the inorganic silicon surface and the organic photoresist. CMi strongly recommends vapour-phase HMDS for silicon and silicon-dioxide surfaces because poor adhesion can cause narrow mask features to lift during development or during the repeated plasma exposure of a Bosch etch [77]. In this work the priming and coating were performed on the Sawatec SM-200 manual line, which combines a programmable coater, an SSE VB20 HMDS vapour-prime module, and precision hotplates

[78].

4.3 AZ 10XT-60 coating and soft bake

AZ 10XT-60 was selected because it is a thick positive i-line photoresist developed for deep silicon etching and high-topography applications. CMi lists it as an appropriate mask for the Bosch process, with steep sidewalls, a nominal film-thickness range of approximately 5.5 µm to 20 µm, and compatibility with both the Sawatec SM-200 and MLA150 [76]. These properties made it more suitable than a thin general-purpose resist for the approximately 300 µm-deep channels, because the mask had to survive hundreds of alternating etch and passivation steps.

The resist was dispensed on the primed surface and spun at 6000 rpm. According to the CMi process curve, AZ 10XT-60 at this speed produces a nominal thickness of approximately 5.5 µm, rather than exactly 5 µm. For chip processing, the CMi Sawatec recipe uses a 40 s main spin and a short final acceleration to reduce edge-bead effects. The coated sample was then soft-baked on the hotplate at 105°C. CMi specifies the manual-bake duration as 90 s + 20 s/μm; for a 5.5 µm film this corresponds to approximately 200 s [77]. The bake removed solvent, stabilized the resist thickness, and reduced the risk of film flow or sticking during handling.

Thick positive resists require water to diffuse back into the film after the soft bake so that the photochemical reaction proceeds uniformly. The coated sample was therefore allowed to rehydrate before exposure. For a film near 5 µm, CMi specifies a minimum delay of 7 min at a cleanroom relative humidity of at least 40%; drier conditions require a longer delay. Respecting this interval improves dose stability and prevents incomplete development in the lower part of the resist [77].

4.4 Direct-write exposure with the MLA150

The binary layout derived from the simulated topology was imported into the Heidelberg Instruments MLA150. This maskless aligner converts standard CAD formats, including GDS, CIF, and DXF, and exposes the resist by scanning a focused 375 or 405 nm laser over the sample. Direct writing avoided the delay and cost of manufacturing a chrome mask and was particularly useful during development, because an updated topology or corrected dimensional bias could be exposed from a revised digital file [79].

Focus and dose were established for the actual resist thickness and layout. The MLA150 requires the user to set the exposure dose and defocus; the optimum depends on film thickness, surface reflectivity, feature density, and whether the pattern contains isolated channels, dense lines, or pillars. CMi gives 240 mJ cm^{-2} as a typical 405 nm starting dose for a 5.5 µm AZ 10XT film, while explicitly recommending a dose series for final dimensional matching [77]. During exposure, the regions corresponding to the future coolant passages received the laser dose. Because AZ 10XT-60 is a positive resist, these exposed regions became soluble in the developer, leaving the unexposed film as the silicon-etch mask.

4.5 Development in AZ 400K

After exposure, the pattern was developed using AZ 400K diluted with deionized water at a ratio of 1:3.5. AZ 400K is a KOH-based inorganic developer; CMi recommends it for the AZ 10XT family because it produces a comparatively sharp edge profile [77]. For manual processing, the recommended development time is approximately 45 s per micrometre of resist thickness. The nominal 5.5 µm film therefore required about 248 s, or slightly more than 4 min, with the endpoint confirmed by complete clearing of the exposed regions.

Development was followed by a 1 min DI-water rinse. A thorough rinse was necessary not only to stop development but also to remove alkaline residue from the front and back surfaces. CMi explicitly requires an additional backside DI rinse before subsequent etching, because developer contamination can damage or contaminate the etcher chuck [77]. The sample was then dried and inspected optically. At this checkpoint, the most important acceptance criteria were continuous resist coverage over the solid regions, fully opened channel areas, absence of pinholes, and no bridging across narrow passages.

4.6 Deep reactive-ion etching in the SPTS Rapier

The developed resist served as the sacrificial mask for deep reactive-ion etching in the SPTS Rapier DSE. CMi identifies the Rapier as a dual-source ICP tool for fluorine chemistry, bulk-silicon etching, and deep silicon etching by the Bosch process. The system uses electrostatic clamping and provides optical emission and white-light interferometric endpoint capabilities [81]. These functions are well suited to a deep microchannel process in which mask survival, thermal stability, and depth control are all critical.

The Bosch process alternates two plasma phases. During the etch phase, fluorine radicals remove silicon from the regions opened in the resist. During the passivation phase, a fluorocarbon polymer is deposited on the exposed surfaces. The next directional ion bombardment preferentially clears this polymer at the bottom of the features, allowing vertical etching to continue while the sidewalls remain protected. Repetition of the etch/passivation pair produces the high aspect ratio required here, although the cyclic mechanism also leaves the characteristic microscopic sidewall scalloping of a Bosch etch.

The nominal target depth was 300 µm. The fabrication run used approximately 350 Bosch cycles; this corresponds to an effective average depth increment of roughly 0.86 µm per cycle. This value is a useful run summary rather

than a universal etch rate. Local depth depends on feature width, open-area fraction, chamber condition, aspect-ratio-dependent etching, and pattern loading. The straight channels and the topology-optimized mask do not present identical local open areas, so depth and sidewall profile had to be checked after the run rather than inferred from cycle count alone.

Thermal management of the sample was especially important during this long etch. CMi explains that its ICP tools use helium beneath the wafer to conduct heat to the electrostatic chuck; the backside must therefore be clean and must clamp properly. When a chip is mounted on a carrier, good thermal contact between the chip and carrier is likewise required to prevent local overheating and resist failure [82]. Process interruptions or depth checks can be used when necessary to ensure that the mask remains intact and that the 300 µm target is approached without excessive notching, sidewall taper, or loss of the smallest silicon islands.

Table 9: Principal microfabrication parameters used for the silicon cooling structures. Public CMi reference values are distinguished from the reported Rapier run setting.

Operation	Parameter	Value or control criterion
Surface preparation	Adhesion promotion	Dehydration followed by HMDS vapour prime
Resist coating	Material / spin	AZ 10XT-60; 6000 rpm; 40 s chip recipe
Resist thickness	Nominal value	Approximately 5.5 µm
Soft bake	Temperature / time	105°C; approximately 200 s for 5.5 µm
Rehydration	Minimum delay	7 min at cleanroom RH $\geq$ 40%
Exposure	Tool / wavelength	MLA150; 405 nm (or calibrated 375 nm mode)
Exposure dose	CMi starting point	240 mJ cm^{-2} for 5.5 µm; final value by dose test
Development	Chemistry / time	AZ 400K:DI water = 1:3.5; about 45 s/μm
Post-development rinse	DI water	At least 1 min plus complete backside rinse
Silicon etch	Tool / process	SPTS Rapier; Bosch DRIE; approximately 350 cycles
Etch depth	Target	300 µm, verified by metrology
Final cleaning	Mask removal	Ultrasonic wet strip, rinse, dry, and inspection

4.7 Resist stripping, cleaning, and inspection

After DRIE, the remaining AZ 10XT-60 no longer served a function and had to be removed without damaging the high-aspect-ratio silicon structures. The sample was immersed in the CMi-approved stripping agent and treated in an ultrasonic bath to assist penetration and removal from the top surface and narrow features. It was then rinsed and dried carefully so that loosened resist or particles were not redeposited inside the channels. CMi provides a dedicated resist-stripping wet bench and ultrasonic-capable solvent facilities for this type of post-etch cleaning [81].

Ion bombardment during DRIE can cross-link the uppermost photoresist and form a hard crust that is less soluble than the original film. When this occurs, CMi recommends a short, high-power oxygen-plasma step to open the crust before wet removal [82]. In the present route, wet stripping with ultrasonic assistance was the final removal step; oxygen plasma should be reported as an additional operation only if it is confirmed in the individual run log. This distinction avoids claiming a treatment that was merely available as a process contingency.

The cleaned samples were inspected optically and by scanning electron microscopy (SEM). The top view verified that the coolant passages were open and that the optimized islands remained separated. Oblique views were used to assess etch depth, wall continuity, taper, and the qualitative magnitude of sidewall scalloping. Figure 16 shows the reference structure, while Fig. 17 shows the transfer of the non-periodic optimized field.

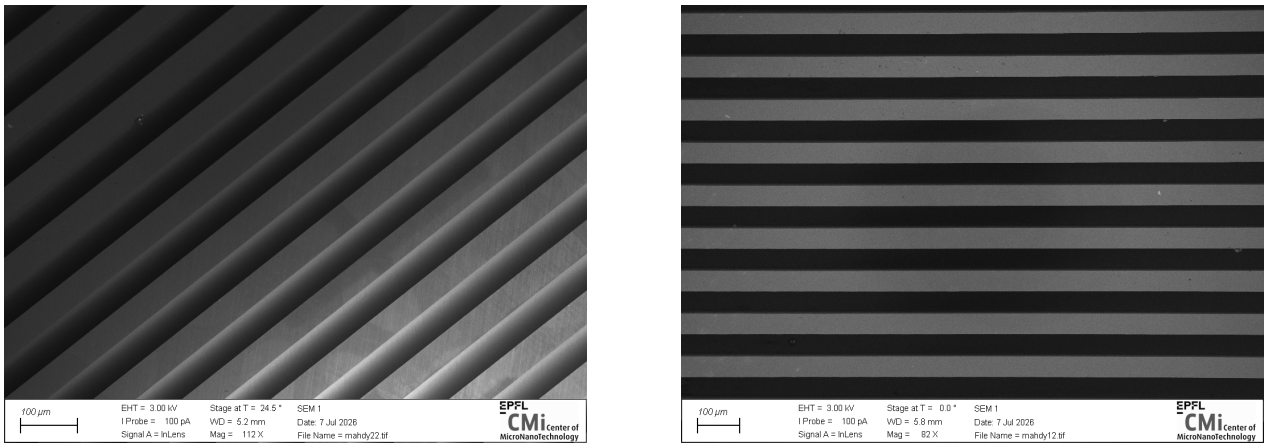

Figure 16: Oblique and top-view SEM images of the straight-channel reference after Rapier DRIE and resist stripping. The views are used to verify channel opening, wall continuity, and the achieved three-dimensional depth.

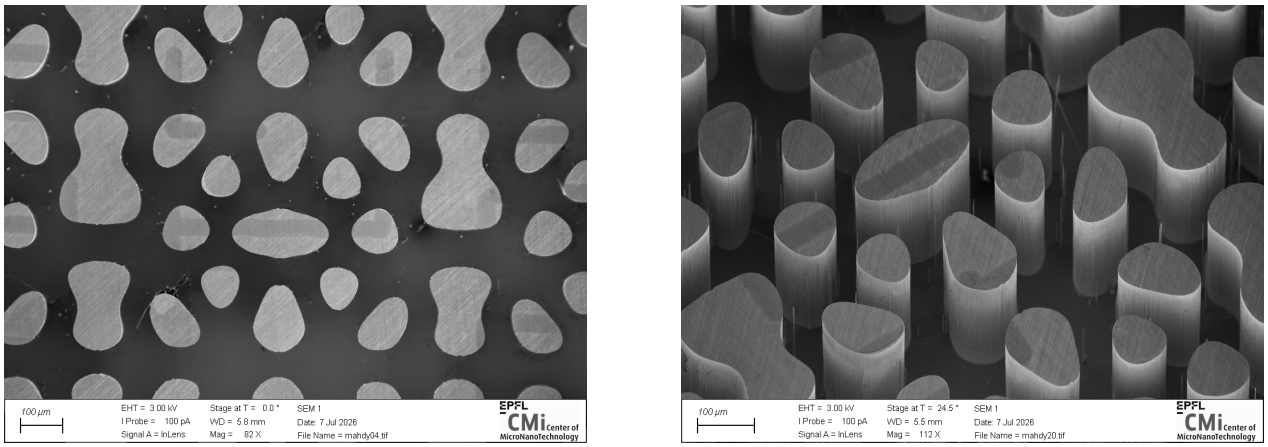

Figure 17: Top and oblique SEM views of the topology-optimized silicon structures fabricated at EPFL CMi. The principal flow passages and solid heat-spreading islands remain distinct after lithography, DRIE, and cleaning.

4.8 Fabrication fidelity and its influence on performance

The optimized layout contained a broader range of local radii, gaps, and solid feature sizes than the straight-channel design. Consequently, it was more sensitive to every transfer step. Resist edge rounding changed the initial opening; mask erosion and lateral plasma attack could enlarge the channel; aspect-ratio-dependent etching could make narrow gaps shallower than open regions; and residual polymer could increase roughness or partially obstruct a passage. These effects explain why a mathematically sharp COMSOL contour cannot be expected to appear identically in the final silicon.

The important fabrication outcome was that the main coolant network remained connected and that the larger solid islands responsible for heat spreading were retained. Nevertheless, the fabricated geometry—not only the CAD geometry—sets the experimental hydraulic diameter, wetted area, pressure drop, and conjugate heat-transfer path. For quantitative model validation, the final process record should therefore retain the measured etch depth, representative top and bottom feature widths, sidewall angle, and any visible scallop or residue. These dimensions can then be introduced into a sensitivity model or an as-fabricated COMSOL geometry. In this way, the CMi work is not a separate manufacturing appendix: it is the physical link between the optimized density field and the thermal and hydraulic measurements discussed in the following sections.

5. Experimental methodology

5.1 Closed-loop architecture

A single closed loop was assembled to circulate the dielectric coolant through the device while controlling flow rate and rejecting the extracted heat. Starting from the reservoir/pump unit, the liquid passed through a Coriolis mass-flow meter and a fan-cooled heat exchanger before reaching the device under test. Resistance temperature detectors measured the inlet and outlet liquid temperatures. Pressure taps immediately upstream and downstream of the device were connected to a differential-pressure transducer. The fluid then returned to the reservoir.

The working liquid was the single-phase dielectric coolant identified in the project presentation as Thermasolv IM7/HFE-7200. Its electrical compatibility permitted the coolant passages to be placed directly above the active power device, while the low-temperature properties motivated its consideration for future space-oriented systems.

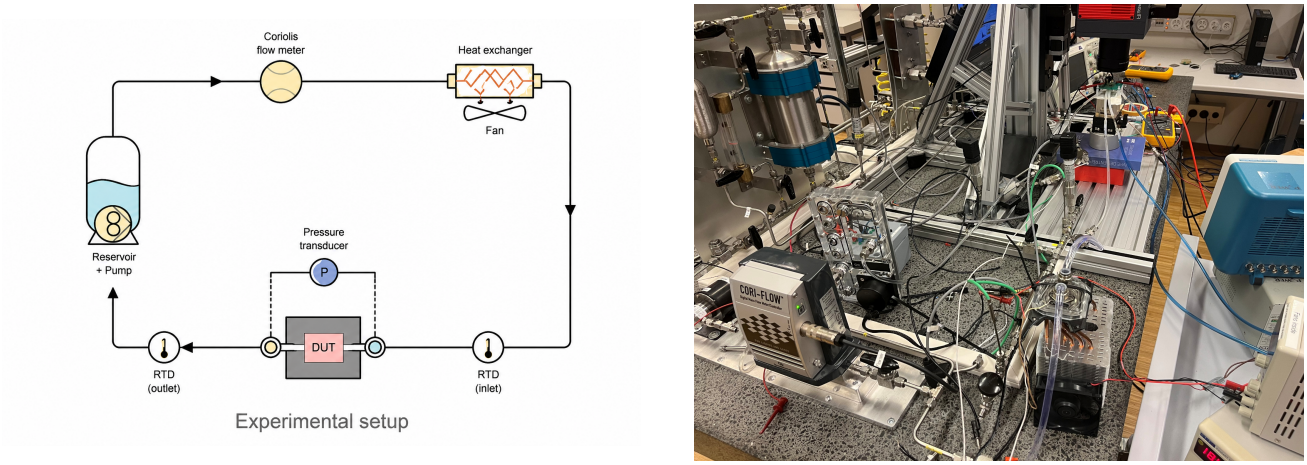

Figure 18: Hydraulic schematic and the corresponding laboratory test rig used for the single-phase experiments.

Placing the heat exchanger before the device reduced drift in the inlet condition and allowed each test point to be approached from a repeatable thermal state. The Coriolis meter supplied the mass-flow value used in the plots, while the differential transducer isolated the pressure loss of the device and package from most of the external tubing. This arrangement also allowed the hydraulic power associated with the test article to be estimated independently of the electrical power consumed by the circulation equipment.

$$P_{\mathrm{hyd}} = \Delta P\, \dot{V} = \Delta P \left(\frac{\dot{m}}{\rho} \right). \tag{25}$$

The experiments remained in the liquid-only regime. Consequently, heat removal was interpreted through sensible heating and convection rather than through evaporation. The hydraulic coefficient of performance, when required for comparison with other cooling technologies, is defined as $\mathrm{COP}_{\mathrm{hyd}} = Q_{\mathrm{removed}}/P_{\mathrm{hyd}}$; it should not be confused with a system-level coefficient that includes the fan, controls, and external heat exchanger.

5.2 Instrumentation, electrical operation, and test sequence

Table 1. Measurements used during thermal and hydraulic characterization.

Quantity	Instrument or method	Role in data reduction
Mass flow	Coriolis flow meter	Sets and records the coolant mass-flow rate.
Pressure loss	Differential-pressure transmitter	Measures ΔP across the packaged device.
Inlet/outlet temperature	RTDs and NI C-series temperature module	Defines the coolant thermal boundary conditions.
Device surface temperature	Infrared camera	Provides maximum and spatially averaged temperature after emissivity treatment.
Electrical loading	17-45 V input, switching control, 5 V auxiliary supply	Establishes the dissipated-power condition of the half-bridge.

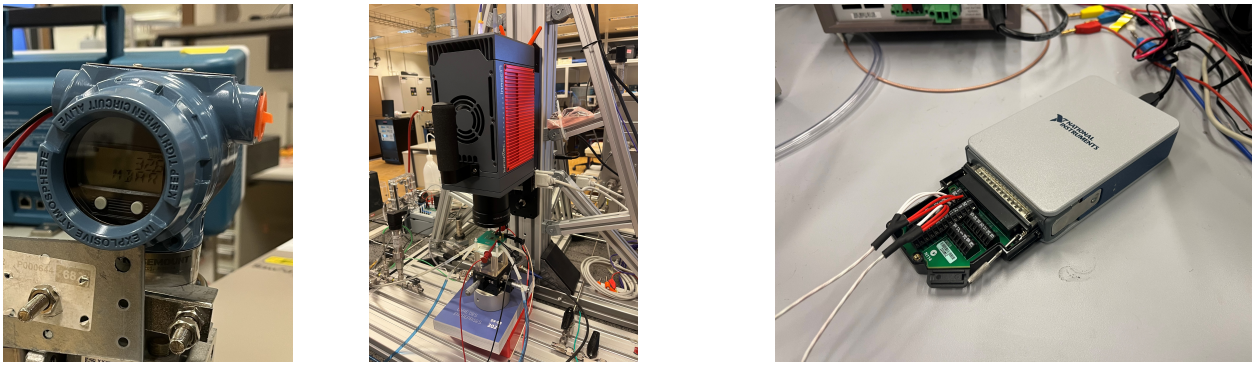

Figure 19: Differential-pressure transmitter, infrared camera, and NI temperature-acquisition hardware used in the campaign.

Each run was performed at one of four nominal mass-flow rates: 0.12, 0.17, 0.225, or 0.2825 kg/min. After the flow and inlet temperature stabilized, electrical power was increased in steps. Temperature, mass flow, and pressure difference were recorded after the device approached a steady response. The same sequence was repeated for the straight-channel reference and the optimized sample.

1. Prime the loop, remove visible gas pockets, and verify that the pressure signal returns to zero before flow is imposed.
2. Set the target mass-flow rate and allow the inlet temperature and pressure difference to stabilize.
3. Apply the electrical switching condition and increase dissipated power in controlled steps within the device limit.
4. Record ΔP, mass flow, inlet/outlet temperature, and the infrared temperature field at each steady operating point.
5. Return to the low-power condition before changing flow rate or replacing the sample.

The reported presentation did not include model numbers, calibration certificates, sampling duration, or a formal uncertainty propagation. These items should be added from the laboratory log before the chapter is submitted as a journal manuscript.

6. Experimental results

6.1 Temperature rise and thermal resistance

For both samples, the maximum and average temperature rises increased approximately linearly with dissipated power. Increasing the mass-flow rate reduced the slope of each curve, as expected for single-phase convection. The optimized structure shifted both the maximum-temperature and average-temperature curves downward relative to the straight-channel reference. This shift was visible at every tested flow rate, indicating that the benefit was not restricted to a single operating point.

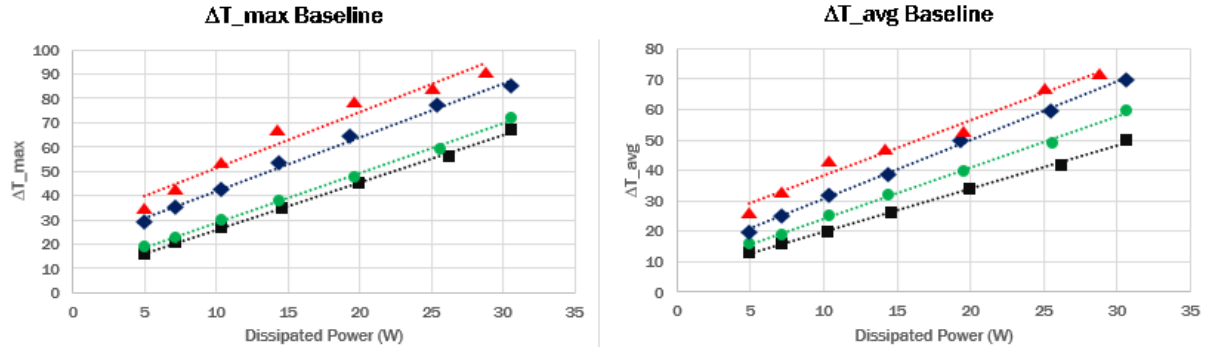

Figure 20: Maximum and average temperature rise of the straight-channel baseline as a function of dissipated power and mass-flow rate.

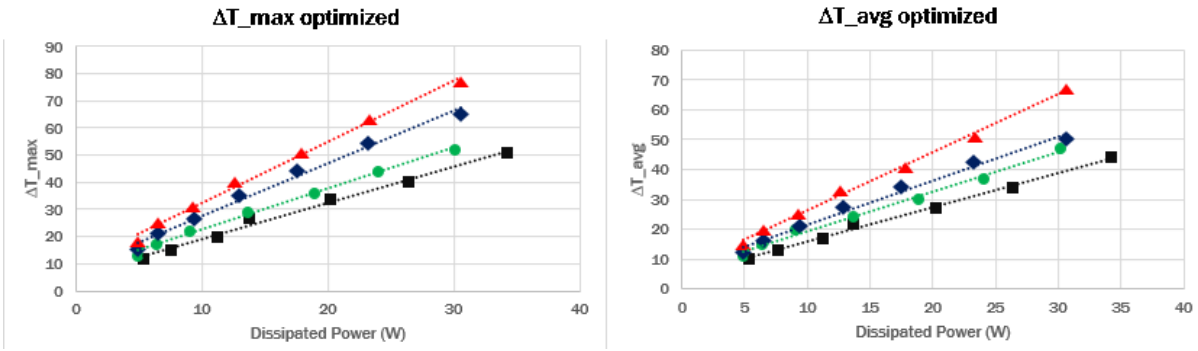

Figure 21: Maximum and average temperature rise of the topology-optimized device over the same test range.

The maximum-temperature metric is the stricter design criterion because the device is limited by local junction temperature rather than by the spatial average alone. Reducing the maximum slope therefore shows that the topology did more than enlarge the total wetted area: it redistributed coolant and solid conduction paths around the region that would otherwise become the dominant hotspot.

Thermal resistance was obtained from the slope of temperature rise against dissipated power at each flow rate. When plotted against inverse mass flow, the optimized sample remained below the baseline over the full range. The project summary reports a 22% reduction in maximum thermal resistance. Because the exposed die area is only 0.11622 cm^2, the upper electrical load of approximately 35 W corresponds to the intended 300 W/cm^2-class local heat flux.

6.2 Thermo-hydraulic comparison

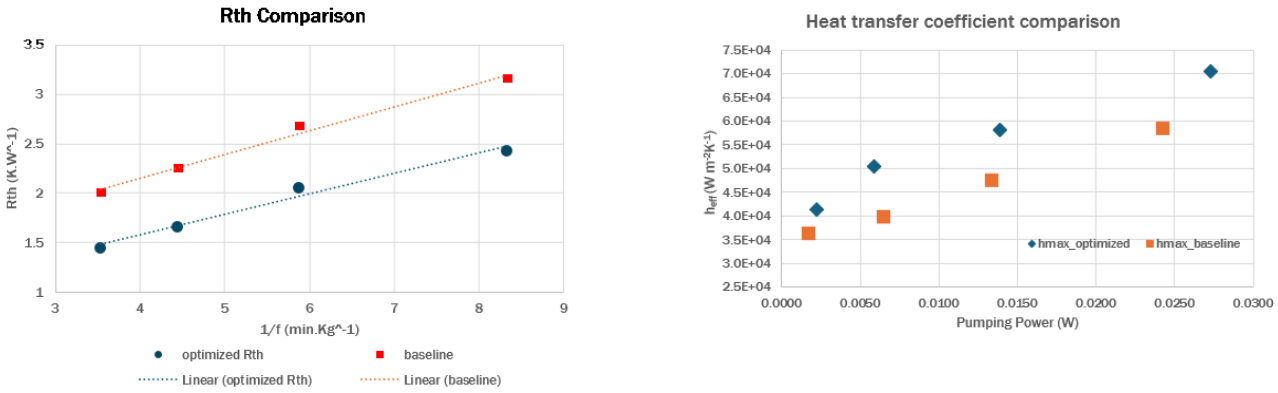

Figure 22: Maximum thermal-resistance trend and effective heat-transfer coefficient of the baseline and optimized samples.

The thermal-resistance comparison makes the design improvement direct: at a given flow condition, the optimized points lie consistently below the baseline points. The heat-transfer coefficient provides the complementary view. Across the tested pumping-power range, the optimized device produced a higher coefficient, with an improvement of approximately 17% at comparable hydraulic power.

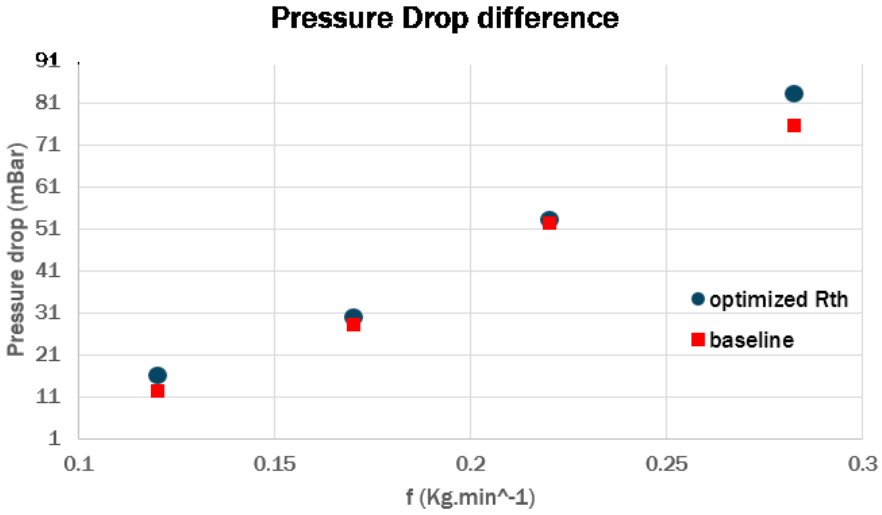

Figure 23: Measured device pressure drop for the baseline and optimized samples.

The optimized geometry did increase pressure drop, particularly at the upper end of the flow range. The increase was nevertheless small compared with the thermal gain. From the graph, the highest-flow point rose from roughly 76 mbar for the baseline to about 83 mbar for the optimized sample. At the lower flow points, the separation was only a few millibar. The topology therefore did not achieve its thermal improvement by accepting an order-of-magnitude hydraulic penalty.

Table 2. Principal measured outcomes of the topology-optimized cooler.

Performance indicator	Observed change	Interpretation
Maximum thermal resistance	$\approx 22\%$ lower	Lower hotspot temperature rise per watt.
Effective heat-transfer coefficient	17% higher	Improved heat extraction at comparable pumping power.
Pressure drop	Slight increase	Additional viscous loss remained modest across the test range.
Numerical/experimental difference	$\leq 15\%$	Model captured the principal thermal and hydraulic trends.

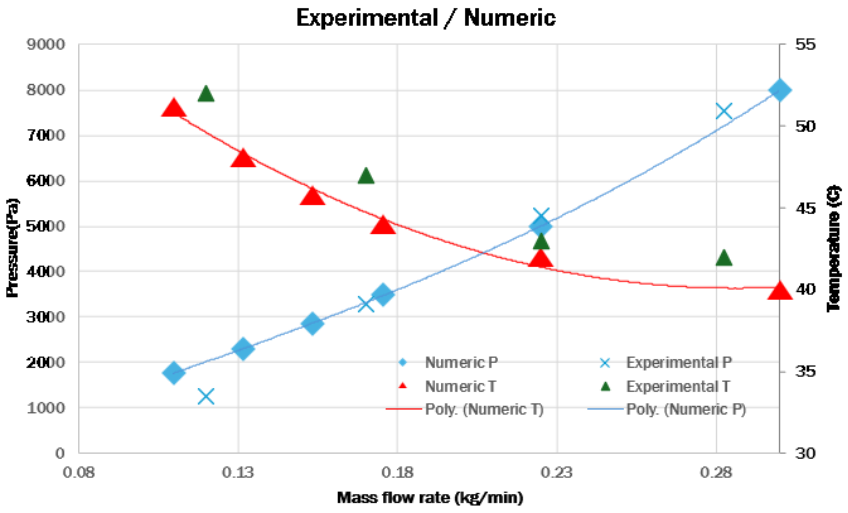

Figure 24: Experimental and numerical pressure/temperature response of the straight-channel reference.

6.3 Numerical validation

The numerical model was checked in two stages. First, the straight-channel sample was tested over several flow rates and its measured pressure and temperature were compared with the corresponding COMSOL predictions. This baseline comparison tested the material properties, thermal boundary conditions, and pressure-loss model without the additional geometric complexity of the optimized structure.

The same comparison was then repeated for the optimized geometry. The measured pressure points followed the predicted increase with flow, while the measured temperature followed the predicted decrease. The maximum reported disagreement was approximately 15% for both the baseline and optimized cases. This level of agreement is sufficient to support the design comparison, although it should not be presented as a complete validation of every local temperature or velocity in the device.

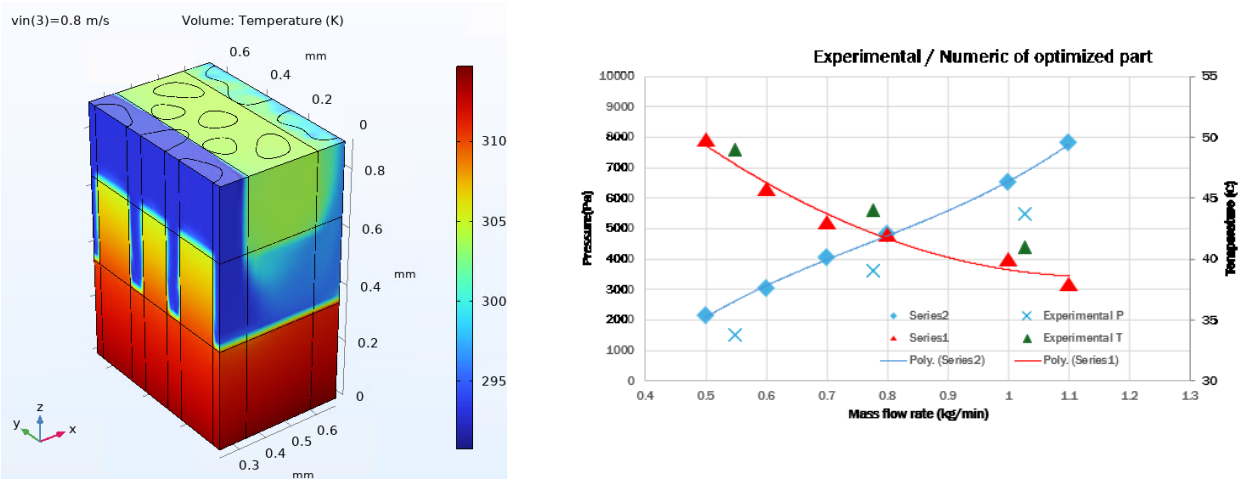

Figure 25: Optimized-device temperature field and comparison between simulated and measured pressure and temperature.

Several effects can explain the remaining discrepancy: uncertainty in the small internal dimensions after DRIE, the use of temperature-dependent coolant properties, heat loss from the PCB and package, contact resistance between the device and cooling structure, and the finite accuracy of the infrared emissivity correction. The simulation represented the designed geometry, whereas the experiment necessarily tested the fabricated and assembled geometry.

7. Discussion

7.1 Why the optimized geometry performed better

The optimized features altered the cooling problem in two linked ways. First, they introduced additional solid pathways that spread heat laterally away from the highest-load region. Second, their irregular spacing forced the liquid to interact with regions that would receive less flow in a simple straight-channel arrangement. The projection and material penalties were therefore not merely numerical devices; they were what allowed the final

map to become a set of manufacturable flow passages and heat-spreading features.

The result also illustrates why pressure drop cannot be excluded from microchannel optimization. A design focused only on temperature would tend to add material and constrict the fluid repeatedly around the hotspot. That choice might lower the simulated surface temperature but could make the required pumping power unacceptable. In this work, the modest rise in pressure drop, together with the 22% thermal-resistance reduction and the approximately 17% gain in heat-transfer coefficient, indicates that the combined objective found a useful compromise.

7.2 Relevance to space electronics

Near-junction cooling is attractive for space power electronics because it removes heat before the package, board, and external radiator stack impose additional thermal resistance. A dielectric coolant also reduces the electrical risk associated with routing liquid close to exposed conductors. The measured hydraulic demand was small relative to the heat removed, which is important in a power-limited spacecraft. The compact manifold concept may also be adapted to densely packed converters or computing modules where a conventional cold plate cannot reach the local hotspot efficiently.

These advantages do not by themselves establish space readiness. The experiments were performed in a laboratory loop under terrestrial gravity. A flight design would still need freeze management, radiation-compatible materials, long-duration seal qualification, vibration testing, contamination control, redundancy, and a system-level radiator/pump trade study. In microgravity, gas management becomes particularly important even in a nominally single-phase loop because trapped bubbles cannot be expected to rise out of the device.

7.3 Limitations

The present comparison is limited to one straight-channel baseline, one optimized architecture, and the flow range available in the test loop. The pressure and temperature agreement supports the model at the measured operating points, but it does not eliminate uncertainty in local heat flux, contact resistance, or internal velocity. The full manuscript should therefore include a measurement uncertainty table and should report the exact coolant property correlations, the fabricated dimensions, and the package sealing method.

Conclusion

The combined review and experimental programme demonstrate a credible route from high-heat-flux electronics requirements to a manufacturable near-junction cooling solution. The work translated the non-uniform heat load of a compact GaN half-bridge into a topology-optimized silicon flow field, fabricated both optimized and straight-channel devices, integrated them with a dedicated manifold and PCB, and validated their thermal and hydraulic response in a closed dielectric-fluid loop.

Relative to the straight-channel reference, the optimized cooler reduced maximum thermal resistance by approximately 22%, increased the effective heat-transfer coefficient by approximately 17% at comparable hydraulic power, and limited the highest-flow pressure-drop increase to approximately 9%. Numerical results remained within about 15% of measurement over the tested range. At approximately 301 $\mathrm{W\,cm^{-2}}$, the reconstructed hydraulic COP of about 1,280 places the device competitively within the selected dielectric-cooling literature while preserving the electrical-safety advantage of a dielectric coolant.

The result should be presented as a strong technology demonstrator rather than a flight-qualified product. Commercial or aerospace deployment now depends on closing the measurement uncertainty, confirming the exact COP from synchronized records, documenting fabrication tolerances, measuring full pump-electrical and system power, and completing package-life and environmental qualification. With those steps, the architecture offers a practical platform for compact converters, radar electronics, avionics processors, and other high-power assemblies in which a localized hotspot governs system performance.

References

[1] Y.-G. Lv, Y.-T. Wang, T. Meng, Q.-W. Wang, and W.-X. Chu, "Review on thermal management technologies for electronics in spacecraft environment," *Energy Storage and Saving*, vol. 3, no. 3, pp. 153–189, 2024. doi:10.1016/j.enss.2024.03.001.

[2] H. Huang and C. Bu, "Design and Verification of Thermal Control System of Communication Satellite," *Aerospace*, vol. 11, no. 10, p. 803, 2024. doi:10.3390/aerospace11100803.

[3] S. Li, S. Wu, T. Xie, Y. Han, and S. Lv, "Design, analysis and verification of thermal control system for satellite laser communication terminal," *Thermal Science and Engineering Progress*, vol. 58, p. 103239, 2025. doi:10.1016/j.tsep.2025.103239.

[4] G. Schunk, S. Babiak, and B. Evans, "Thermal Control System Architecture and Technology Challenges for a Lunar Surface Habitat." NASA Technical Reports Server, accessed May 31, 2026. `https://ntrs.nasa.gov/citations/20220003526`.

[5] M. Middelhuis, M. Muñoz Rojo, and W. Wits, "Analysis and experimental validation of a pumped two-phase loop for multi-component electronics cooling," *Heat and Mass Transfer*, vol. 60, pp. 1–23, 2023. doi:10.1007/s00231-023-03440-1.

[6] D. Cowan and T. A. Bond, "Coolant Leak from ISS External Active Thermal Control System (EATCS)—An Examination of Most Probable." NASA Technical Reports Server, accessed May 31, 2026. `https://ntrs.nasa.gov/citations/20220010221`.

[7] C. Wang et al., "Developments and Prospects in Temperature Control Technique of Loop Heat Pipe for Spacecraft," *Frontiers in Heat and Mass Transfer*, vol. 23, no. 4, pp. 1261–1280, 2025. doi:10.32604/fhmt.2025.066473.

[8] S. Zhong, R. Xie, Y. Li, and X. Sun, "Operational Characteristics of Loop Heat Pipe in Microgravity and Normal Gravity Environments," *Journal of Thermal Science*, vol. 33, no. 4, pp. 1394–1408, 2024. doi:10.1007/s11630-024-1955-7.

[9] M. Choi, *Swift BAT Loop Heat Pipe Thermal Performance After 19+ Years On Orbit*, 2024. doi:10.32865/JRIU5827.

[10] B. M. Diaconu, M. Cruceru, and L. Anghelescu, "Phase change materials in space systems. Fundamental applications, materials and special requirements—A review," *Acta Astronautica*, vol. 216, pp. 163–213, 2024. doi:10.1016/j.actaastro.2023.12.040.

[11] S. Wang, X. Hou, J. Yin, Y. Xing, and Z. Wang, "Comparative Study of the Thermal Enhancement for Spacecraft PCM Thermal Energy Storage Units," *Aerospace*, vol. 9, no. 11, p. 705, 2022. doi:10.3390/aerospace9110705.

[12] C. Hartsfield, T. Shelton, B. Palmer, and R. O'Hara, "All Metallic Phase Change Thermal Management Systems for Transient Spacecraft Loads," *Journal of Aerospace Engineering*, vol. 33, 2020. doi:10.1061/(ASCE)AS.1943-5525.0001150.

[13] S. Chakraborty, D. Shukla, and P. Kumar Panigrahi, "A review on coolant selection for thermal management of electronics and implementation of multiple-criteria decision-making approach," *Applied Thermal Engineering*, vol. 245, p. 122807, 2024. doi:10.1016/j.applthermaleng.2024.122807.

[14] O. A. Volodin, D. A. Shvetsov, V. S. Serdyukov, V. I. Zhukov, and A. N. Pavlenko, "Enhanced boiling and evaporation of dielectric fluids on modified surfaces for immersion cooling of electronic components—a review," *Applied Thermal Engineering*, vol. 277, p. 127088, 2025. doi:10.1016/j.applthermaleng.2025.127088.

[15] Z.-Q. Yu, M.-T. Li, and B.-Y. Cao, "A comprehensive review on microchannel heat sinks for electronics cooling," *International Journal of Extreme Manufacturing*, vol. 6, no. 2, p. 022005, 2024. doi:10.1088/2631-7990/ad12d4.

[16] I. Pinkus et al., "Ultra-low thermal resistance and pressure drop copper and copper-tungsten diamond-shaped pin fin cold plates for liquid cooling of electronics," *International Journal of Heat and Mass Transfer*, vol. 256, p. 128080, 2026. doi:10.1016/j.ijheatmasstransfer.2025.128080.

[17] C. Ao, B. Xu, and Z. Chen, "Microchannel cooling for high-performance microelectronic systems: Progress, challenges, and future directions," *International Communications in Heat and Mass Transfer*, vol. 172, p. 110494, 2026. doi:10.1016/j.icheatmasstransfer.2026.110494.

[18] N. Cassada et al., "SiC Microchannel Heat Sinks for High Heat Flux Dissipation of 1 kW/cm^2," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 13, 2023. doi:10.1109/TCPMT.2023.3281772.

[19] U. Ghani, M. A. Wazir, K. Akhtar, M. Wajib, and S. Shaukat, "Microchannel Heat Sinks—A Comprehensive Review," *Electronic Materials*, vol. 5, no. 4, pp. 249–292, 2024. doi:10.3390/electronicmat5040017.

[20] S. Sun, P. Liebersbach, and X. Qian, "3D Topology Optimization of Heat Sinks for Liquid Cooling," *Applied Thermal Engineering*, vol. 178, p. 115540, 2020. doi:10.1016/j.applthermaleng.2020.115540.

[21] D. Geng, C. Wei, Y. Liu, and M. Zhou, "Concurrent topology optimization of multi-scale cooling channels with inlets and outlets," *Structural and Multidisciplinary Optimization*, vol. 65, no. 8, p. 234, 2022. doi:10.1007/s00158-022-03336-4.

[22] J. Gorman, R. Pejman, S. R. Kumar, J. F. Patrick, and A. R. Najafi, "Transient topology optimization for efficient design of actively cooled microvascular materials," *Structural and Multidisciplinary Optimization*, vol. 67, no. 4, p. 60, 2024. doi:10.1007/s00158-024-03774-2.

[23] M.-É. Lamarche-Gagnon, M. Molavi-Zarandi, V. Raymond, and F. Ilinca, "Additively manufactured conformal cooling channels through topology optimization," *Structural and Multidisciplinary Optimization*, vol. 67, no. 8, p. 138, 2024. doi:10.1007/s00158-024-03846-3.

[24] H. Yang et al., "Topological optimization design of micro-pin fin heat sinks for high heat flux cooling application," *International Journal of Thermal Sciences*, vol. 218, p. 110102, 2025. doi:10.1016/j.ijthermalsci.2025.110102.

[25] J. Zhou, M. Lu, Q. Zhao, Q. Li, and X. Chen, "Topologically optimized manifold microchannel heat sink with extreme cooling performance for high heat flux cooling of electronics," *Applied Thermal Engineering*, vol. 241, p. 122426, 2024. doi:10.1016/j.applthermaleng.2024.122426.

[26] M. K. Sharma and B. Ramos-Alvarado, "Thermal management of 3-D heterogeneously integrated microelectronics: challenges and future research directions," *Communications Engineering*, vol. 5, p. 28, 2026. doi:10.1038/s44172-026-00590-y.

[27] N. Gilmore, V. Timchenko, and C. Menictas, "Manifold microchannel heat sink topology optimisation," *International Journal of Heat and Mass Transfer*, vol. 170, p. 121025, 2021. doi:10.1016/j.ijheatmasstransfer.2021.121025.

[28] Z. Duan, G. Xie, B. Yu, and P. Jin, "Multi-objective topology optimization and thermal performance of liquid-cooled microchannel heat sinks with pin fins," *Case Studies in Thermal Engineering*, vol. 49, p. 103178, 2023. doi:10.1016/j.csite.2023.103178.

[29] Y. Xiong, L. Gao, W. Sha, and J. Gao, "Multi-objective designs of microchannel heat sink for multilayer composite chip heat transfer systems using topology optimization," *Applied Thermal Engineering*, vol. 279, p. 127876, 2025. doi:10.1016/j.applthermaleng.2025.127876.

[30] Y. Wang, J. Wang, and X. Liu, "Topology Optimization Design of Micro-Channel Heat Sink by Considering the Coupling of Fluid-Solid and Heat Transfer," *Energies*, vol. 15, no. 23, p. 8827, 2022. doi:10.3390/en15238827.

[31] S. Ozguc, L. Pan, and J. A. Weibel, "Topology optimization of microchannel heat sinks using a homogenization approach," *International Journal of Heat and Mass Transfer*, vol. 169, p. 120896, 2021. doi:10.1016/j.ijheatmasstransfer.2020.120896.

[32] R. K. Padhy, K. Suresh, and A. Chandrasekhar, "TOMAS: topology optimization of multiscale fluid flow devices using variational auto-encoders and super-shapes," *Structural and Multidisciplinary Optimization*, vol. 67, no. 7, p. 119, 2024. doi:10.1007/s00158-024-03835-6.

[33] L. Gu, L. Chen, Q. Yang, X. Zheng, C. Guo, and W. Tao, "Topology optimization of convective heat transfer in microchannels under different working modes," *Applied Thermal Engineering*, vol. 269, p. 126049, 2025. doi:10.1016/j.applthermaleng.2025.126049.

[34] H. Han, Y. Han, Y. Lin, C. Wang, J. G. Korvink, and Y. Deng, "Topology optimization of microchannel heat sinks for laminar flows of thermal–fluid," *Applied Thermal Engineering*, vol. 270, p. 126153, 2025. doi:10.1016/j.applthermaleng.2025.126153.

[35] D. Lu et al., "Near-Junction Thermal Management of GaN-on-SiC MMIC Power Amplifier Through Substrate Embedded Microchannel," *IEEE Transactions on Electron Devices*, pp. 1–8, 2023. doi:10.1109/TED.2023.3338586.

[36] X. Cheng et al., "Thermal design and on-orbit validation of deployable synthetic aperture radar antennas for satellite constellations with emphasis on temperature uniformity," *Aerospace Science and Technology*, vol. 170, p. 111574, 2026. doi:10.1016/j.ast.2025.111574.

[37] D. Mitchao, T. Totani, and H. Sakamoto, "Thermal design exploration of flexible and lightweight antenna structures through design of experiments," *Applied Thermal Engineering*, vol. 264, p. 125323, 2025. doi:10.1016/j.applthermaleng.2024.125323.

[38] Y. Yan, C. Zhang, Y. Liu, Z. He, and J. Li, "Numerical study on hotspots adaptive cooling and thermal-hydraulic performance enhancement of fractal microchannel heat sink embedded with hydrogels," *International Journal of Thermal Sciences*, vol. 172, p. 107272, 2022. doi:10.1016/j.ijthermalsci.2021.107272.

[39] X. Li and Y. Xuan, "Embedded Cooling of High-Heat-Flux Hotspots Using Self-Adaptive Microchannel/Pin-Fin Hybrid Heat Sink," *IEEE Transactions on Components, Packaging and Manufacturing Technology*, vol. 12, no. 10, pp. 1627–1635, 2022. doi:10.1109/TCPMT.2022.3211009.

[40] J. Zhou, J. Chen, Q. Wang, X. Xie, P. Guan, and H. Zheng, "Numerical Study on Optimization of Manifold Microchannel Heat Sink," *Energies*, vol. 18, no. 22, p. 5883, 2025. doi:10.3390/en18225883.

[41] S. Aflatounian et al., "Comparative evaluation of dielectric liquids for single-phase immersion cooling of electronics," *International Journal of Heat and Mass Transfer*, vol. 265, p. 128765, 2026. doi:10.1016/j.ijheatmasstransfer.2026.128765.

[42] J. Khan and P. Singh, "Review on phase change materials for spacecraft avionics thermal management," *Journal of Energy Storage*, vol. 87, p. 111369, 2024. doi:10.1016/j.est.2024.111369.

[43] I. Garmendia, H. Vallejo, M. Seco, and E. Anglada, "Design and Fabrication of a Phase Change Material Heat Storage Device for the Thermal Control of Electronics Components of Space Applications," *Aerospace*, vol. 9, no. 3, p. 126, 2022. doi:10.3390/aerospace9030126.

[44] Y. Yan, Z. He, G. Wu, F. Xu, L. Li, and L. Zhang, "Automatically adaptive cooling of hotspots by a fractal microchannel heat sink embedded with thermo-responsive hydrogels," *International Journal of Energy Research*, vol. 46, no. 3, pp. 3132–3144, 2022. doi:10.1002/er.7370.

[45] S. Tachikawa, H. Nagano, A. Ohnishi, and Y. Nagasaka, "Advanced Passive Thermal Control Materials and Devices for Spacecraft: A Review," *International Journal of Thermophysics*, vol. 43, no. 6, p. 91, 2022. doi:10.1007/s10765-022-03010-3.

[46] "Advances in spacecraft thermal control," in *Advances in Heat Transfer*, vol. 56, Elsevier, pp. 1–50, 2023. doi:10.1016/bs.aiht.2023.04.001.

[47] Y. Akizuki et al., "Development and testing of the re-deployable radiator for deep space explorer," *Applied Thermal Engineering*, vol. 165, p. 114586, 2020. doi:10.1016/j.applthermaleng.2019.114586.

[48] R. B. Mulford et al., "Experimental demonstration of heat loss and turn-down ratio for a multipanel, actively deployed radiator," *Applied Thermal Engineering*, vol. 178, p. 115658, 2020. doi:10.1016/j.applthermaleng.2020.115658.

[49] Q. Meng, X. Chen, Y. Nian, W. Cheng, and Z. Zhao, "Experimental Study on the Transient Behaviors of Mechanically Pumped Two-Phase Loop with a Phase Change Energy Storage Device for Short Time and Large Heat Power Dissipation of Spacecraft," *Heat Transfer Engineering*, vol. 44, no. 21–22, pp. 2091–2106, 2023. doi:10.1080/01457632.2022.2164677.

[50] A. Butler and C. Argyropoulos, "Mechanically tunable radiative cooling for adaptive thermal control," *Applied Thermal Engineering*, vol. 211, p. 118527, 2022. doi:10.1016/j.applthermaleng.2022.118527.

[51] P. B. Chiranjeevi, K. Krishnaraj, K. G. Vinod, K. Srinivasan, and T. Sundararajan, "Numerical simulations, experimental investigation and optimization of hybrid space thermal radiators," *Applied Thermal Engineering*, vol. 234, p. 121275, 2023. doi:10.1016/j.applthermaleng.2023.121275.

[52] L. Wang, N. Boman, S. Taylor, and C. Stoops, "Reducing temperature swing and rectifying radiative heat transfer for passive dynamic space thermal control with variable-emittance coatings," *Applied Physics Letters*, vol. 127, no. 20, p. 201106, 2025. doi:10.1063/5.0303297.

[53] A. M. Elshaer, A. M. A. Soliman, M. Kassab, S. Mori, and A. A. Hawwash, "Thermal control of a small satellite in low earth orbit using phase change materials-based thermal energy storage panel," *The Egyptian Journal of Remote Sensing and Space Sciences*, vol. 26, no. 4, pp. 954–965, 2023. doi:10.1016/j.ejrs.2023.11.007.

[54] S. Miesner et al., *Thermal Orbital Spacecraft Analysis of an Additively Manufactured Deployable Radiator Oscillating Heat Pipes (AMDROHP) CubeSat*, 2024. doi:10.1115/IMECE2023-114220.

[55] M. G. Izenson, D. A. Knaus, and L. O'Neill, "Thermal Storage for High-Power Small Satellites."

[56] I. Amyx, C. Anderson, N. Cassada, and T. M. Bandhauer, "Thermal Management of Discretized Heaters Using CuW Microchannel Heat Sinks and FC3283," *IEEE Trans. Components, Packaging and Manufacturing Technology*, vol. 15, no. 3, pp. 470–478, 2025. doi:10.1109/TCPMT.2024.3477276.

[57] D. Jung et al., "Thermal Design and Management of Micro-Pin Fin Heat Sinks for Energy-Efficient Three-Dimensional Stacked Integrated Circuits," *Int. J. Heat Mass Transfer*, vol. 175, 121192, 2021. doi:10.1016/j.ijheatmasstransfer.2021.121192.

[58] L. Nela et al., "Impact of Embedded Liquid Cooling on the Electrical Characteristics of GaN-on-Si Power Transistors," *IEEE Electron Device Letters*, vol. 42, no. 11, pp. 1642–1645, 2021. doi:10.1109/LED.2021.3114056.

[59] R. Xiao, P. Zhang, L. Chen, Y. Hou et al., "Experimental Study on Cooling Performance of a Hybrid Microchannel and Jet Impingement Heat Sink," *Applied Sciences*, vol. 12, no. 24, 13033, 2022. doi:10.3390/app122413033.

[60] N. Damle and Y. Joshi, "Experimental and Analytical Study of the Hydrodynamic and Single and Two-Phase Convective Heat Transfer Performance of Flexible PDMS Microchannels with Micropillar Arrays," *Applied Thermal Engineering*, vol. 224, 119799, 2023. doi:10.1016/j.applthermaleng.2022.119799.

[61] G. Marseglia et al., "Experimental Investigation on the Effects of the Geometry of Microchannels Based Heat Sinks on the Flow Boiling of HFE-7100," *Applied Thermal Engineering*, vol. 236, 121479, 2024. doi:10.1016/j.applthermaleng.2023.121479.

[62] C. H. Hoang et al., "Hybrid Microchannel/Multi-Jet Two-Phase Heat Sink: A Benchmark and Geometry Optimization Study of Commercial Product," *Int. J. Heat Mass Transfer*, vol. 169, 120920, 2021. doi:10.1016/j.ijheatmasstransfer.2021.120920.

[63] C. H. Hoang et al., "Two-Phase Impingement Cooling Using a Trapezoidal Groove Microchannel Heat Sink and Dielectric Coolant HFE 7000," *Proc. IEEE iTherm*, 2021. doi:10.1109/ITherm51669.2021.9503283.

[64] Y. J. Lee, C. H. Hwang, H. Lee, I. Lee, and S. J. Kim, "Highly energy-efficient manifold microchannel for cooling electronics with a coefficient of performance over 100,000," *Energy Conversion and Management*, vol. 358, 121422, 2026. doi:10.1016/j.enconman.2026.121422.

[65] Y. J. Lee and S. J. Kim, "Experimental investigation on thermal-hydraulic performance of manifold microchannel with pin-fins for ultra-high heat flux cooling," *Int. J. Heat Mass Transfer*, vol. 224, 125336, 2024. doi:10.1016/j.ijheatmasstransfer.2024.125336.

[66] Z. Wu, W. Xiao, H. He, W. Wang, and B. Song, "Jet-enhanced manifold microchannels for cooling electronics up to a heat flux of 3,000 W cm^{-2}," *Nature Electronics*, vol. 8, pp. 810–817, 2025. doi:10.1038/s41928-025-01449-4.

[67] D. Kong, E. Jung, Y. Kim, and H. Lee, "An additively manufactured manifold-microchannel heat sink for high-heat-flux cooling," *Int. J. Mechanical Sciences*, vol. 248, 108228, 2023. doi:10.1016/j.ijmecsci.2023.108228.

[68] Y. Yang, J. Du, M. Li, W. Wang et al., "Embedded microfluidic cooling with compact double H type manifold microchannels for large-area high-power chips," *Int. J. Heat Mass Transfer*, vol. 197, 123340, 2022. doi:10.1016/j.ijheatmasstransfer.2022.123340.

[69] N. Zhang et al., "Experimental investigation of the embedded micro-channel manifold cooling for power chips," *Thermal Science*, vol. 26, no. 2B, pp. 1531–1543, 2022. doi:10.2298/TSCI210908328Z.

[70] T. Lian et al., "Thermal property evaluation of a 2.5D integration method with device level microchannel direct cooling for a high-power GaN HEMT device," *Microsystems & Nanoengineering*, vol. 8, 119, 2022. doi:10.1038/s41378-022-00462-3.

[71] S. Yang, J. Li, B. Cao, Z. Wu, K. Sheng et al., "Investigation of Z-type manifold microchannel cooling for ultra-high heat flux dissipation in power electronic devices," *Int. J. Heat Mass Transfer*, vol. 218, 124792, 2024. doi:10.1016/j.ijheatmasstransfer.2023.124792.

[72] R. van Erp, R. Soleimanzadeh, L. Nela, G. Kampitsis, and E. Matioli, "Co-designing electronics with microfluidics for more sustainable cooling," *Nature*, vol. 585, pp. 211–216, 2020. doi:10.1038/s41586-020-2666-1.

[73] N. Cassada et al., "SiC Microchannel Heat Sinks for High Heat Flux Dissipation of 1 kW/cm^2," *IEEE Trans. Components, Packaging and Manufacturing Technology*, vol. 13, no. 5, pp. 655–665, 2023. doi:10.1109/TCPMT.2023.3281772.

[74] B. Bazmi et al., "High-heat-flux liquid cooling of electronics with manifold microchannels: Thermal-hydraulic optimization and experimental validation," *Int. J. Heat Mass Transfer*, vol. 268, 128924, 2026. doi:10.1016/j.ijheatmasstransfer.2026.128924.

[75] S. Yan, F. Wang, J. Hong, and O. Sigmund, "Topology optimization of microchannel heat sinks using a two-layer model," *International Journal of Heat and Mass Transfer*, vol. 143, p. 118462, 2019. doi:10.1016/j.ijheatmasstransfer.2019.118462.

[76] EPFL Center of MicroNanoTechnology, "i-line photoresists available in CMi," equipment and process guidance. `https://www.epfl.ch/research/facilities/cmi/process/photolithography/iline-photoresist-selection/` (accessed 21 September 2026).

[77] EPFL Center of MicroNanoTechnology, "AZ 10XT Process steps," coating, soft-bake, rehydration, exposure, and development guidance. `https://www.epfl.ch/research/facilities/cmi/process/photolithography/iline-photoresist-selection/az-10xt/` (accessed 21 September 2026).

[78] EPFL Center of MicroNanoTechnology, "Photolithography equipment," Sawatec SM-200 manual coater, HMDS module, hotplates, and associated tools. `https://www.epfl.ch/research/facilities/cmi/equipment/photolithography/` (accessed 21 September 2026).

[79] EPFL Center of MicroNanoTechnology, "Heidelberg Instruments MLA150 – 1 & 2," equipment description and exposure guidance. `https://www.epfl.ch/research/facilities/cmi/equipment/photolithography/mla-150/` (accessed 21 September 2026).

[80] EPFL Center of MicroNanoTechnology, "Photolithography on chips," small-sample handling and carrier guidance. `https://www.epfl.ch/research/facilities/cmi/process/photolithography/photolithography-on-small-samples/` (accessed 21 September 2026).

[81] EPFL Center of MicroNanoTechnology, "Etching equipment," SPTS Rapier DSE, resist-strip, and solvent wet-bench capabilities. `https://www.epfl.ch/research/facilities/cmi/equipment/etching/` (accessed 21 September 2026).

[82] EPFL Center of MicroNanoTechnology, "Tips and Tricks," electrostatic-clamp cooling, chip mounting, and photoresist removal after dry etching. `https://www.epfl.ch/research/facilities/cmi/tips_and_tricks/` (accessed 21 September 2026).