

Design of a Two-Phase Gravity-Driven Micro-Scale Thermosyphon Cooling System for High-Performance Computing Data Centers

André Seuret*, Arman Iranfar*, Marina Zapater*, John Thome[†], and David Atienza*

*Embedded Systems Laboratory (ESL), Swiss Federal Institute of Technology Lausanne (EPFL), Switzerland

[†]Heat and Mass Transfer Laboratory (LTCM), EPFL, Switzerland

{andre.seuret, arman.iranfar, marina.zapater, john.thome, david.atienza}@epfl.ch

ABSTRACT

Next-generation High-Performance Computing (HPC) systems need to provide outstanding performance with unprecedented energy efficiency while maintaining servers at safe thermal conditions. Air cooling presents important limitations when employed in HPC infrastructures. Instead, two-phase on-chip cooling combines small footprint area and large heat exchange surface of micro-channels together with extremely high heat transfer performance, and allows for waste heat recovery. When relying on gravity to drive the flow to the heat sink, the system is called a closed-loop two-phase thermosyphon. Previous research work either focused on the development of large-scale proof-of-concept thermosyphon demonstrators, or on the development of numerical models able to simulate their operation. In this work, we present a new ultra-compact micro-scale thermosyphon design for high heat flux components. We manufactured a working 8 cm height prototype tailored for Virtex 7 FPGAs with a heat spreader area of 45 mm × 45 mm, and we validate its performance via measurements. The results are compared to our simulator and accurately match the thermal performance of the thermosyphon, with error of less than 3.5%. Our prototype is able to work over the full range of power of the Virtex7, dissipating up to 60 W of power while keeping chip temperature below 60°C. The prototype will next be deployed in a 10 kW rack as part of an HPC prototype, with an expected Power Usage Effectiveness (PUE) below 1.05.

KEY WORDS: Thermosyphon, two-phase, liquid cooling, HPC data center, power usage effectiveness

Nomenclature

\dot{m}	Mass-flow rate [$\frac{kg}{s}$]
\dot{q}	Heat flux [$\frac{W}{m^2}$]
A	Surface area [m^2]
C	Channel
cp	Heat capacity [$\frac{J}{kg \cdot K}$]
F	Fin [m]
g	Gravity [$\frac{m}{s^2}$]
H	Enthalpy [$\frac{J}{kg}$], Height [m]
h	Heat transfer coefficient [$\frac{W}{m^2 \cdot K}$]
k	Pressure drop coefficient [$\frac{Pa \cdot m^2 \cdot s}{kg}$]
L	Length of evaporator or condenser [m]
N	Number of discretization or number of channels [-]
P	Pressure [Pa]
Q	Heat load [W]
S	Surface of the control Volume [m^2]

V	Volume [m^3]
v	Velocity [$\frac{m}{s}$]
W	Width of evaporator or condenser [m]
x	Vapor quality [-]
z	Z axis [-]

Greek symbols

β	Homogeneous void fraction [-]
λ	Heat conduction coefficient [$W/m/K$]
μ	Dynamic viscosity [$Pa \cdot s$]
ρ	Density [$\frac{kg}{m^3}$]
θ	Angle with the horizontal [rad]

Subscripts

c	Channel
end	At the end of the loop
F	Friction pressure drop
ftp	Footprint
g	Gas
h	Height [m]
i	Incrementation
l	Liquid
m	Module
$start$	At the beginning of the loop
$static$	Static pressure
tp	Two-phase
w	Width [m]
z	Along the z axis

INTRODUCTION

Energy efficiency is a major challenge in current and future High-Performance Computing (HPC) infrastructures. In the race towards exascale computing, rocketing rack densities require sustainable power consumption while keeping servers at safe thermal conditions. Cooling power has traditionally accounted for over 30% of the overall power consumption in air-cooled data centers [1]. Improvements in air-cooling subsystems, together with cooling-aware hardware and software techniques have allowed reductions on the average Power Usage Effectiveness (PUE) of data centers (i.e., the ratio of total facility power to IT power), improved from 2.5 in 2007 to 1.65 in 2013 [2], but has reached an stagnation in recent years. This is because, with increasing rack power densities, air-cooling becomes inefficient and insufficient due to its limited heat transfer capacity, impeding sufficient cooling provisioning. As a result, there has been a proliferation of pumped liquid-cooling techniques, both single- and two-phase.

Single-phase water-cooled systems are easier to manufacture and demonstrate efficient heat transfer characteristics [3]. However, several benefits of two-phase cooling systems make them more appealing than single-phase solutions. First, they have higher heat transfer coefficients and lower power consumption, which makes them suitable for high-density HPC systems. Second, two-phase on-chip cooling combines the small footprint area and large heat exchange surface of micro-channels together with the extremely high heat transfer performance of evaporating flow. Third, since the coolant can work at a higher temperature, it allows heat recovery.

The ultimate cooling solution is to deploy efficient passive two-phase cooling systems, which rely on gravity to drive the flow to the heat sink. Such a system is called a closed-loop two-phase thermosyphon. A thermosyphon enhances the thermal performance of current and future chip technologies and presents several advantages. First, its gravity-driven principle drastically reduces power consumption by eliminating pumping power. Second, it offers feasible opportunities for energy reuse of the waste heat. Finally, due to its passive nature, it is more reliable than traditional techniques as it does not utilize fans, pumps, and other mechanical components.

Previous research in this area either focused on the development of large-scale thermosyphon prototypes that served as proof-of-concept demonstrators [4], [5], or addressed the development of numerical models able to simulate the properties of mini-thermosyphons, in order to provide a design for 2U servers [6], [7]. To the best of our knowledge, no real micro-scale thermosyphon prototype has yet been built and validated in a commercial platform.

In this work, we design, manufacture, and validate a new ultra-compact micro-scale thermosyphon for high heat flux components. We design the thermosyphon using an in-house steady-state simulator, obtaining a prototype of 8 cm height, specifically tailored for Virtex 7 FPGA chips with a footprint of $45\text{ mm} \times 45\text{ mm}$ and a maximum Thermal Design Power (TDP) of 60 W. We manufacture a prototype and validate it against real measurements by using an in-house Thermal Test Vehicle (TTV) and an infrared microscope. The results obtained show the capability of our simulator to accurately match the real dynamics of the thermosyphon with less than 6.5% error. Our thermosyphon prototype is able to maintain the chip below 65°C for the whole power supply range and will be deployed in a 10 kW rack as part of a heterogeneous high-performance computing platform, with an expected PUE below 1.05.

In particular, our main contributions are as follows:

- We improve previous numerical simulation models by better convergence for the steady-state, increasing its flexibility to simulate every kind of design. We prove how our simulation framework is a suitable tool for designing a real thermosyphon, as it matches the real system dynamics with an error below 6.5%.
- We design and manufacture a real thermosyphon prototype of 8 cm height, specifically tailored for Virtex 7 FPGA chips with an area of $45\text{ mm} \times 45\text{ mm}$ and a TPD of 60 W, able to achieve a PUE of 1.05.
- By using an in-house TTV and an Infra-Red Microscope

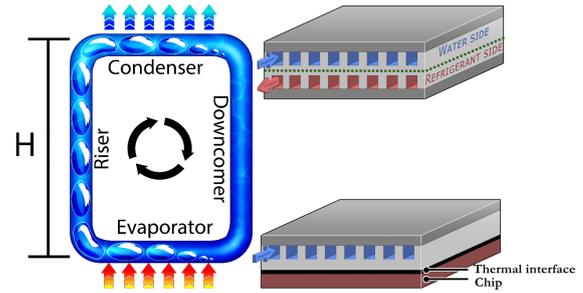


Fig. 1. Thermosyphon loop principle

we validate our prototype and prove that it is able to maintain the chip below 65°C for the whole power range.

THERMOSYPHON FUNDAMENTALS

Working principle of the thermosyphon

As shown in Fig. 1, a typical thermosyphon is composed of a closed loop with four main parts including downcomer, evaporator, riser, and condenser. The height difference (H) between the condenser and the evaporator is of great importance since gravity plays a major role in the loop. Firstly, the refrigerant in liquid state, which is in the downcomer, evaporates partially in the evaporator thanks to the heat received from a heat source (FPGA chip in our case). The two-phase mixture, composed of vapor and liquid, rises by buoyancy in the riser until the condenser. In the condenser, the heat exchange turns the vapor back into liquid, and it falls to the evaporator through the downcomer. Since the density of the two-phase mixture in the riser is lower than the liquid in the downcomer, a flow rate is created under the gravity field. In the literature [4], [5], an accumulator, which is mainly a downcomer with larger section, is added to these four parts. This accumulator reduces the height change of the liquid level and, as a consequence, stabilizes the system for a large operating domain. In this case, a secondary water loop is used to cool down the refrigerant in the condenser via a counter-flow heat exchanger. In addition, to deal with high heat fluxes, the exchange surfaces in the condenser and evaporator are maximized by using micro-channels with widths and heights of lower than 1 mm, and in order of millimeters, respectively. The height and the gravity field requirements for the thermosyphon pose a great challenge for design and manufacturing a micro-scale (less than 10 cm height and suitable for current chips) thermosyphons.

State-of-the-art on mini-thermosyphon design and manufacturing

The most recent thermosyphon built, which demonstrates comparable performance results to a pumped cooling loop system, has 15 cm height but a large footprint area of $1\text{ m} \times 1\text{ m}$ [7], [8]. This mini-thermosyphon validates the two flow regimes: i) gravitational dominant regime, and ii) frictional dominant regime. Moreover, it handles heat loads up to 158 W with a heat flux of $70\text{ W}/\text{cm}^2$ while keeping the temperature of the chip below 58°C . Moreover, the real data acquired allowed the validation of a steady-state numerical model, able to predict the chip temperature within 5% of error [6], while its dynamic simulation framework predicts temperature and pressure with

a maximum error of 0.4 K and 0.35 bar, respectively [4]. This setup and simulation framework allowed the design of a thermosyphon for a commercial 2U server [9]. However, this design was not manufactured and tested on a real system. The present work focuses on the design, manufacturing and testing of a micro-scale thermosyphon prototype for a commercial FPGA chip in an HPC system.

NUMERICAL SIMULATION MODEL

Transient thermosyphon numerical models have been proposed in previous work [4], [5], [9], [6]. In this paper, we present a simplified steady-state numerical model.

The thermosyphon is modeled through three modules: I) a pipe module for both downcomer and riser, II) a condenser module, and III) an evaporator module. In addition, a separate module is considered to account for the singular pressure drop in connections (i.e elbows or reduction/increase of cross-sectional flow area between different parts). In each module, the two-phase flow is modeled using a homogeneous plug flow model, with both phases in thermodynamic equilibrium. The two-phase flow is essentially reduced to a highly compressible single-phase flow that provides the possibility of keeping the same set of conservative equations for single- and two-phase flow. The overall fluid properties are thus derived from the vapor and the liquid properties. The latter is mixed thanks to the vapor quality and void fraction. A 1D steady-state flow is considered all over the loop. The set of conservative equations are derived from Navier-Stokes equations and simplified for this application. By considering a control volume V with a peripheral surface S , the conservation equations are as follows:

$$\frac{\partial(\rho v_z)}{\partial z} = 0 \quad (1)$$

$$\underbrace{\rho v_z \frac{\partial v_z}{\partial z}}_{\text{momentum}} = - \underbrace{\rho g \sin(\theta)}_{\text{static}} - \underbrace{\frac{\partial P}{\partial z}}_{\text{total}} - \underbrace{\left(\frac{\partial P}{\partial z}\right)_F}_{\text{friction}} \quad (2)$$

$$\underbrace{\rho v_z \frac{\partial}{\partial z}(H)}_{\text{convection}} + \underbrace{\frac{\partial}{\partial z} \left(\frac{1}{2} \rho v_z^3 \right)}_{\text{kinetic energy}} = \underbrace{v_z \left(\frac{\partial P}{\partial z} \right)_F}_{\text{viscous heating}} + \underbrace{\frac{\dot{q}S}{V}}_{\text{heat source from the perimeter}} - \underbrace{\rho g \sin(\theta) v_z}_{\text{potential energy}} \quad (3)$$

where z is flow direction, θ represents the angle between the horizontal and the flow direction, ρ is the density of the fluid, g is the gravity, P shows the pressure, H is the enthalpy, \dot{q} is the heat flux received, and v represents the velocity of the fluid. As shown later, the frictional pressure drop and heat flux terms are given by empirical correlations.

We apply discretization to all modules by dividing the flow into N finite volumes, numbered from $0 < i < N$. Thus, the total length is divided by N , which results in an infinitesimal length of Δz . In this work, we experimentally found N equal to 100 to be sufficient. For each cell, the three conservative equations (1), (2), and (3) are solved.

Because of the inter-dependency of the developed steady-state equations, a pressure velocity coupling algorithm such as SIMPLE [10] or SIMPLEC [11] is required. However, these algorithms add a convergence loop and, hence, complexity

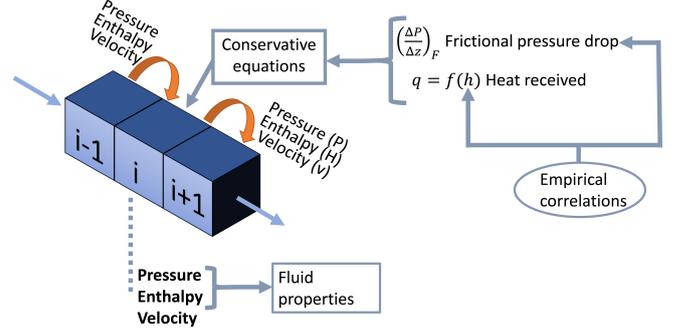


Fig. 2. Discretization of the 1D steady flow

to the model. We propose to use backward finite differences at every volume except for the pressure and enthalpy, for which the forward finite difference scheme is used in order to define the next finite volume. Since there could be error/delay between the considered volumes, they must be small enough to reduce the error effect. Although this inconsistency causes a spatial delay between the couples ρ/v , and H/P , it decreases the computational complexity of the model. In such way, pressure (P) and enthalpy (H) are known for the next volume, and we can entirely define the fluid state using CoolProp [12], a free thermophysical property database. The previous equations are thus discretized as follows:

$$v_{z,i} = \frac{\rho_{i-1} v_{z,i-1}}{\rho_i} \quad (4)$$

$$P_{i+1} = \left(-\rho_i g \sin(\theta) - \left(\frac{\partial P}{\partial z} \right)_{F,i} - \frac{v_{z,i} - v_{z,i-1}}{\Delta z} \rho_i v_{z,i} \right) \Delta z + P_i \quad (5)$$

$$H_{i+1} = \left(v_{z,i} \left(\frac{\partial P}{\partial z} \right)_{F,i} + \frac{\dot{q}S}{V} - \rho_i g \sin(\theta) v_{z,i} - \frac{1}{2} \frac{\rho_i v_{z,i}^3 - \rho_{i-1} v_{z,i-1}^3}{\Delta z} \right) \frac{\Delta z}{\rho_i v_{z,i}} + H_i \quad (6)$$

where the subscript i refers to the i^{th} volume, as shown in Fig. 2. The continuity (4), momentum (5), and energy equations (6) describe the 1D steady-flow and, thus, the flow in each thermosyphon part. The following subsections describe each thermosyphon module with their corresponding equations. The empirical correlations used in this work for different flow patterns are listed in Table 1. The viscosity is computed according to the model presented by Beattie *et al.* [13]:

$$\mu = \mu_l(1 - \beta)(1 + 2.5\beta) + \mu_g \beta \quad (7)$$

Pipe module

The pipe module is the simplest one since it only computes the conservative equations for a pipe, without any heat losses. The source term in the energy equation is thus equal to zero ($\frac{\dot{q}S}{V} = 0$). The frictional pressure drop is computed following the Friedel correlation for upward and downward flows [33]. This module takes as inputs the geometry of the pipe, the mass-flow rate, the pressure, and the enthalpy of the fluid at the inlet. The outputs are the velocity, pressure, and enthalpy at

Table 1. Correlations summary

Flow patterns	Heat transfer coefficient	Pressure drop
Single phase developing flow	Muzychka and Yovanovich [14]	Muzychka and Yovanovich [15]
Single phase laminar	Muzychka and Yovanovich (end)[14], Shah and London [16], Dharaiya and Kandlikar [17]	Shah and London[16]
Single phase turbulent	Gnielinski [18]	Petukhov [19]
Evaporating	Costa-Patry and Thome [20], updated three-zone model [21], [22], [23], annular flow model [24], [25], [26], [27]	Shah and London
Condensing	Koyama [28], Cavallini [29], [30], Kim and Mudawar [31]	Kim and Mudawar [32]

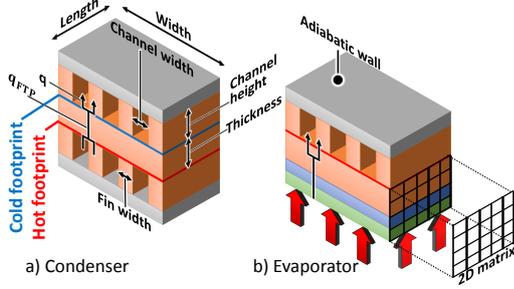


Fig. 3. Schematic view of the condenser and the evaporator

the end of the pipe. The pipe is divided into N finite volumes, and starting from the first one, the pressure and enthalpy are computed using equations (4), (5), (6).

Condenser module

In this subsection, we describe the model of a counter flow heat exchanger, as shown in Fig. 1 and 3. The inputs are the geometry of the channels in the hot and cold sides, the total dimension of the device, the mass-flow rate, and the pressure and enthalpy at the inlets of hot and cold sides. The outputs are the pressure and enthalpy at the outlets of the fluids. The model solves the conservative equations (4), (5), (6) for each finite volume in both hot and cold sides for one channel and assumes that it is the same across the channels. Starting from the inlet of the cold side, at each finite volume, the heat is conducted between both sides in 1D using the fin efficiency equations to provide the heat transfer coefficient to the footprint layer position. This layer is localized at the channel wall facing the opposite fluid side. The fin efficiency equations take into account the heat, in both hot and cold sides, which is spread into the fluid and the fins. The source term q in equation (6) is obtained from the footprint value.

$$\dot{q} = \frac{\dot{q}_{ftp} \cdot W \cdot \Delta z}{S} \quad (8)$$

$$m^2 = 2h \frac{\Delta z}{kF_w \Delta z} \quad (9)$$

$$\eta_{fin} = \frac{\tanh(mC_h)}{mC_h} \quad (10)$$

$$\frac{h_{ftp}}{h} = \frac{\dot{q}_{ftp}}{\dot{q}} = N_c (\eta_{fin} 2C_h + C_w) \quad (11)$$

$$\dot{q}_{ftp} = \frac{1}{\frac{1}{h_{ftp,hot}} + \frac{e}{\lambda} + \frac{1}{h_{ftp,hot}}} \quad (12)$$

$$\dot{q} = \frac{h_{ftp} W}{C_w \cdot L \cdot N_c} \quad (13)$$

where C_h , C_w , F_w , e are, respectively, the channel height, the channel width, the fin width and the thickness of the separating layer. L and W represent the whole condenser/evaporator

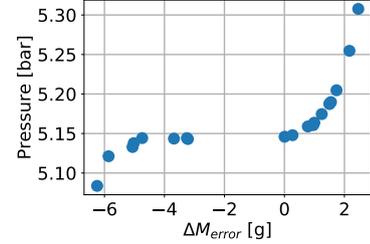


Fig. 4. Convergence of the pressure to match the initial mass

dimensions (see Fig. 3). The heat transfer coefficient h comes from the empirical correlations depending on the flow state as shown in Table 1. Besides, since the two-phase heat transfer coefficient (h) correlation is a function of the heat flux q , a convergence loop is performed in each step using the Nelder-Mead algorithm [34]. The initial value is obtained by:

$$\dot{q}_{hot} = \frac{Q_{tot}}{C_{w,hot} \cdot L \cdot N_{c,hot}} \quad (14)$$

$$\dot{q}_{cold} = \frac{Q_{tot}}{C_{w,cold} \cdot L \cdot N_{c,cold}} \quad (15)$$

where $Q_{tot} = (H_{hot,out} - H_{hot,in}) \cdot \dot{m}_{hot}$ is the heat load.

Since the computation starts from the inlet of the cold side, the finite volume in the hot side at this position is one of the outputs of the numerical model and, thus, the state values are unknown. Therefore, the pressure and enthalpy in the hot side have to be first initialized. For the pressure, we assume no pressure drop. Hence, the outlet pressure is equal to the inlet pressure. For the enthalpy, we assume that the outlet temperature is equal to $T_{hot} - x \cdot (T_{hot} - T_{cold})$. The value of x increases from zero to one until the first physical value is obtained, i.e. the hot fluid temperature is always higher than the cold one, for the whole channel length. Once all the finite volumes until the end of the heat exchanger are computed, the resulting inlets of the hot side are compared with the desired inputs. Thus, two errors exist: one over the pressure, and one over the enthalpy. The heat exchanger is now considered as a function whose inputs are outlet pressure and enthalpy for the hot side while the outputs are the norm of the errors. Finally, we use the Nelder-Mead algorithm to minimize this function and find the correct outlet values.

Evaporator module

The evaporator module is very similar to the condenser. Nonetheless, the conduction until the bottom of the chip is achieved through 2D steady-state conduction. In this 2D matrix, a constant heat flux is applied at the bottom boundary while the removal heat flux from the fluid is applied at the top.

To compute the heat received from the bottom of the package in the fluid finite volumes, the temperature at the bottom of the channels T_{ftp} (corresponding to the top boundary of the 2D conduction matrix) is set to 1°C higher than the inlet temperature of the fluid in the first step. Then, the heat flux q_{ftp} is computed in the same than for the condenser module. However, instead of equation (12), we use:

$$\dot{q}_{ftp} = h_{ftp} \cdot (T_{ftp} - T_{fluid}) \quad (16)$$

where h_{ftp} is the local heat transfer coefficient reported to the footprint location by fin efficiency equations (see Fig. 3). Since the correlation giving the heat transfer coefficient is a function of the heat flux, a convergence loop is performed at each fluid finite volume using Nelder-Mead algorithm. Therefore, the whole channel length is computed and the conduction matrix is solved using the removal heat flux calculated. Then, the new temperature T_{ftp} is obtained and we iterate the loop until the heat load received from the fluid matches the amount of heat applied at the bottom of the 2D conduction matrix. Note that many layers with different heat conduction can be set up in the conduction matrix to simulate the thermal interface. The maximum temperature at the bottom boundary matrix is considered as the temperature of the chip.

Thermosyphon loop

Having modeled all the thermosyphon parts independently, the challenge is how to link these modules together in a full simulation framework that finds the final values of the overall loop (i.e., the mass-flow rate, the pressure, and the enthalpy). The goal is to have a numerical model which takes the geometry, the filling ratio, the bottom chip heat flux, the mass-flow rate of the water, and the inlet temperature as inputs. The outputs are all characteristics of the thermosyphon, i.e., temperature, pressure and mass-flow rate of the refrigerant, chip temperature, enthalpy, etc.

The overall thermosyphon model runs each module one after the other by linking the outlets to the inlets. It starts right before the evaporator with initial values of pressure P_{start} , enthalpy H_{start} , and mass-flow rate \dot{m}_{start} . We assume that the saturated temperature of the refrigerant is equal to the outlet temperature of the water, which can be easily computed using the bottom heat flux. The initial enthalpy can be set such that the vapor quality is $x = -0.01$. The mass-flow rate is the most difficult parameter to approximate. As an initial guess, we assume the vapor quality at the outlet of the evaporator to be $0.4 \sim 0.5$ and then compute the mass-flow rate, as follows:

$$Q_{evap} = \dot{q}_{evap} \cdot S_{evap} \quad (17)$$

$$\dot{m} = \frac{Q_{evap}}{H_{out} - H_{in}} \quad (18)$$

The pressure is assumed constant in the whole loop. In the first step, the assumed pressure P_{start} does not change. The pressure and enthalpy at the end of the loop at the same initial location (before the evaporator), P_{end} , H_{end} should match the start values. If after one loop, the pressure is smaller than the initial assumption ($P_{end} < P_{start}$), the pressure drop is too large and the mass-flow rate should be reduced for the next loop.

Indeed, in the thermosyphon loop, the sum of all frictional elements, pressure drops, and static pressure differences should be equal to zero. Hence, the following equation holds:

$$\begin{aligned} \Delta P_{error} &= P_{end} - P_{start} \\ &= \sum_j^{module} (\Delta P_{F,j} + \Delta P_{static,j}) \\ &= \Delta P_{F,evaporator} + \Delta P_{static,evaporator} \\ &\quad + \Delta P_{F,riser} + \Delta P_{static,riser} \\ &\quad + \Delta P_{F,condenser} + \Delta P_{static,condenser} \\ &\quad + \Delta P_{F,downcomer} + \Delta P_{static,downcomer} \\ &\quad + \Delta P_{F,singulars} \end{aligned} \quad (19)$$

where ΔP_{error} is the error on the pressure and should be zero. To compute the next mass-flow rate, the total pressure drop of each module is considered as a singular pressure drop for which the pressure drop coefficient k should be computed as:

$$\begin{aligned} \Delta P_{F,m} &= k_m \frac{1}{2} \rho_m v_m^2 \\ &= k_m \frac{1}{2} \frac{\dot{m}^2}{\rho_m A_m^2} \end{aligned} \quad (20)$$

where A_m is the cross sectional area of the module, and ρ_m is the density at the beginning of the module of interest. By merging equations (19) and (20), one obtains:

$$\begin{aligned} \dot{m} &= \sqrt{\frac{-\sum \Delta P_{static,m} \cdot 2}{\sum \frac{-k_m}{\rho_m A_m^2}}} \\ &= \sqrt{\frac{-(\Delta P_{error} - \sum \Delta P_{F,m}) \cdot 2}{\sum \frac{-k_m}{\rho_m A_m^2}}} \end{aligned} \quad (21)$$

This method prevents occurrence of the non-stable points highlighted by Lamaison et al. [9]. The average value between the initial and the last enthalpy (i.e., $\frac{H_{start} + H_{end}}{2}$) is considered for the next loop. During this process, the initial pressure P_{start} does not change until the mass-flow rate and the enthalpy converge. Next, the total mass of the system is compared with the initial one. The latter is computed by the initial inputs of the system, as follows:

$$M_{initial} = \sum V_m \cdot FR \cdot \rho_{l,ST} \quad (22)$$

where FR is the filling ratio (i.e., input of the model), and $\rho_{l,ST}$ is the density of the liquid at standard temperature. The convergence loop is run until we obtain an error over the mass lower than 1% of the initial mass. Fig. 4 shows an example of this process. For each point, the mass-flow rate and the enthalpy have converged. A small change of pressure induces a large mass variation, demonstrating the sensitivity right below the initial mass. This is mainly due to the increase of pressure drop in two-phase, i.e., a small pressure favors the evaporation and, thus, the reduction of the total mass. However, this effect is enhanced by the fact that the mass-flow rate is also reduced in order to match the increased pressure drop. The convergence over the mass is achieved on the 7th iteration in the worst case. However, the convergence on the enthalpy and mass-flow rate takes around 30 iterations for the worst cases. The whole

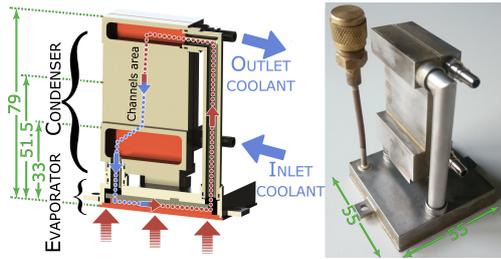


Fig. 5. Cut view of the prototype thermosyphon presenting the refrigerant loop (left) and the manufactured thermosyphon (right). The measurements are in millimeter

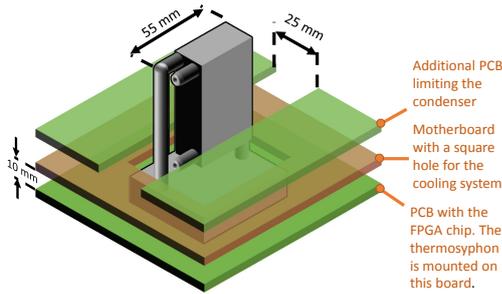


Fig. 6. Schematic view of the implementation of the thermosyphon

framework is coded in Python 3.4.

THERMOSYPHON DESIGN AND MANUFACTURING

Design Constraints

In this work, we consider a Virtex 7 FPGA with a footprint area of $45\text{ mm} \times 45\text{ mm}$ mounted on a $20\text{ cm} \times 20\text{ cm}$ PCB attached to server-size motherboard that hosts up to four PCBs. The motherboard contains a hole that provides space for the cooling system. The total power drawn by the FPGA is of 60 W, which we consider as TDP. Thus, using the models presented in the previous section, we ensure the capability of the design to handle 60 W.

Fig. 5 shows a schematic view of the design and the final manufactured prototype. The footprint dimension of the thermosyphon is considerably constrained by the PCB dimension, the size of the hole in the motherboard, and the available space between two consecutive motherboards, which limits the total height of the thermosyphon to 10 cm. The footprint of the channel area of the evaporator is $45\text{ mm} \times 45\text{ mm}$ to match the FPGA chip. Also, the channel area of the condenser is smaller and equal to $37\text{ mm} \times 34\text{ mm}$.

In addition, in order to make the PCB removable from the motherboard, quick connects for the coolant pipes were used. Moreover, as several motherboards can be stacked on top of the other, other PCBs constrain the dimension of the condenser, the riser, and the downcomer (see Fig. 6). For this reason, a vertical condenser (instead of horizontal) has been chosen, as shown in Fig. 5.

At the inlet of the evaporator, we must have a singular pressure drop to avoid dynamical perturbation in the evaporator channels. At the outlet of the condenser, a great volume with

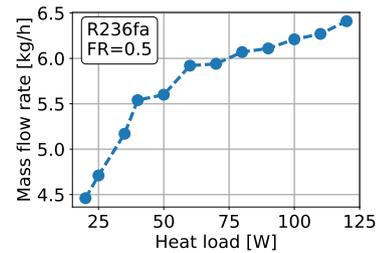


Fig. 7. Characteristic curve of the designed thermosyphon. Predicted mass-flow rate in function of Heat load.

a large cross section area is present which represents the accumulator.

Thermosyphon Characteristic Curve

The thermosyphon characteristic curve described by Bielski and Mikielwicz [35] presents two regimes. First, at low heat loads, the gravitational dominant regime (GDR) presents an increase in the mass-flow rate with the heat load. In this region, the vapor quality at the outlet of the evaporator increases and, thus, the density of the two-phase mixture decreases, which results in a larger density difference between the riser and the downcomer. In the second region, at higher heat loads, the impact of the pressure drop overcomes this mixture density, and the mass-flow rate decreases with the power until it becomes zero. This is called the frictional dominant regime (FDR). Thus, the thermosyphon should cover the whole power range (up to TDP) without reaching the dry-out. Fig. 7 shows the predicted characteristic curve of the designed thermosyphon with R236fa and, $FR = 0.5$, and coolant temperature of 20°C . Even until 120 W, the mass-flow rate increases with power and the thermosyphon stays in the GDR. The safety factor is, thus, at least 2.

Thermosyphon Manufacturing

Since the device fulfills the design requirements, it is manufactured using copper for the cold-plates (for evaporator and condenser), and stainless steel for the covers and the pipes. Brazing is used to assemble the copper cold-plates to the covers. Conventional brazing torch is challenging because of the small dimensions of the device, as brazing a spot may debraise neighbors through conduction. This is especially true for the condenser, which is a cold-plate with brazed covers on each side. For this reason, the device is assembled with brazing paste on the contact surfaces and, then, brazed at once in an oven at a controlled temperature of 750°C . In order to check whether any channels are blocked by the brazing and to check the overall brazing quality, a tomography scanner is used on both the evaporator and the condenser, as shown in Fig. 8. In this prototype, three channels over a total of 67 were blocked by the brazing, which is acceptable as it does not affect the performance of the prototype. In order to provide flexibility in the design, the condenser and the evaporator are not soldered together. Two screws (M3) maintain them attached and the air-tightness is obtained through four O-rings in perfluoroelastomer Kalrez at the extremities of the riser and the downcomer pipes. This material has been chosen to

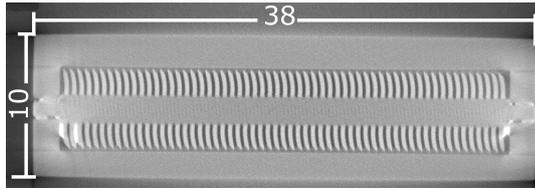


Fig. 8. X-ray scanning of the condenser. Water side on the top and refrigerant side on bottom. Three channels are blocked by brazing. Measurements in millimeter

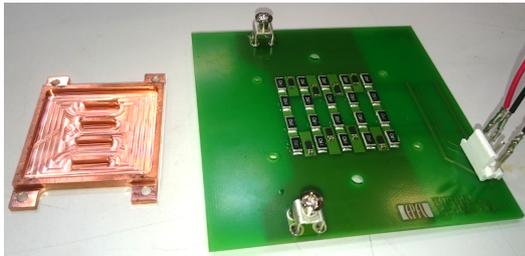


Fig. 9. In-house TTV composed of a 5x4 matrix of 5x4 of 1ω resistors, 8 thermal sensors, and a copper cover.

ensure the compatibility with a large range of refrigerants, also reducing any risk of wearout over time.

EXPERIMENTAL RESULTS AND VALIDATION

Experimental setup

In order to validate the manufactured thermosyphon and the numerical model we use a Thermal Test Vehicle (TTV) consisting of a heat-plate which generates homogeneous heat flux imitating a CPU/FPGA chip. We design a PCB composed of a matrix of 20 SMD resistors and eight temperature sensors, and collect the temperature data via I^2C bus using an Arduino board¹ every 100 ms. To provide an acceptably homogeneous heat flux, a copper plate is attached on top of the resistors. We fill the gap between the copper plate and the resistors with thermal paste² of high heat conductivity. The whole setup is shown in Fig. 9.

In order to calibrate the sensors, we use one offset point. In this context, inside a thermally isolated box, the ambient temperature given by the sensors was compared to two PT100 probes. Besides, in a thermal bath, four K-type thermocouples were calibrated with the same PT100 probes from 20 to 80 °C. These thermocouples were attached to the riser, downcomer, and at the inlet and outlet of the water, and measure the outside walls of the thermosyphon and not directly the internal fluid. Although the thickness of the walls are relatively small (~ 2 mm), the temperature readings are subject to conduction and the errors are thus difficult to estimate. To limit the losses by natural convection, an isolating material surrounds the device. A thermally controlled tank of five liters capacity equipped with a pump drives the water in the secondary loop. A filter of 90 μm is placed before the condenser to avoid any dirtiness inside the micro-channels. Finally, a flow-meter and a gate valve follow the thermosyphon. The whole facility is depicted in Fig. 10.

¹<https://www.arduino.cc/>

²<http://www.coolermaster.com/cooling/thermal-compound/mastergel-maker/>

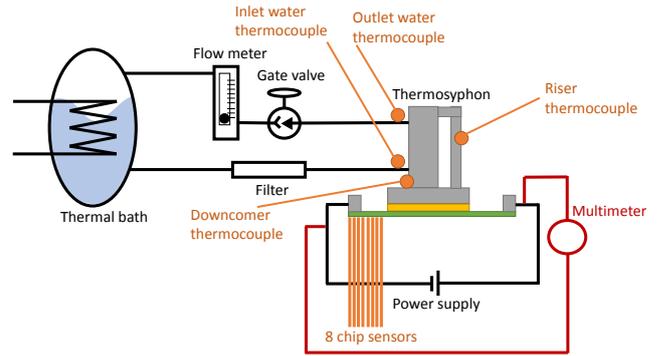


Fig. 10. Schema of the whole facility

Steady-state Results

The thermosyphon is charged with 8.04 g of R236fa, which represents a filling ratio of 50.4%. Fig. 11 shows the steady-state results for water mass-flow rate increasing from 3.59 to 6.88 kg/h with inlet temperature ranging from 20 to 30.7 °C for a heat load from 20 to 60 W. Fig. 11(a) illustrates the mean chip temperature, which increases linearly with the heat load. The variation of water mass-flow rate does not influence the chip temperature considerably. The slope increases slightly when the water mass-flow rate decreases. Moreover, when the mass-flow rate reduces by 48%, the temperature increases by 11% at 60 W. The change in the water temperature does not change the slope of the curve. However, it induces a larger rise in chip temperature, which is almost linear with the water temperature difference since a rise of 10.7 °C of the water induces an average increase of 11.35 °C of the chip. Although chip temperature depends highly in the specific packaging, the heat conductivity of our TTV is lower than in conventional chips and, therefore, represents a worst-case scenario.

The total thermal resistance, computed by $\frac{(T_{chip} - T_{water}) \cdot A_{chip}}{Q_{electric}}$ and plotted in Fig. 11(d), shows a decrease, reflecting an increase of the overall vapor quality. In Fig. 11(b) we plot the water temperature difference and show how inlet temperature does not affect its value. Nevertheless, the mass-flow rate reduction increases the difference between the inlet and outlet water temperatures. Not only is the curve shifted upward, but also the slope increases. Finally, Fig. 11(c) shows the heat loss by natural convection, measured as the difference between the input electrical power and the energy received from the water formulated as $q_{water} = \dot{m} \cdot c_{p_{water}} \cdot \Delta T_{water}$. The total loss increases with the heat load, which is coherent with the increase of the system temperature, and is below 12%.

Comparison with the numerical model

The predicted chip temperature values follow the trend that measurements with a tendency of under predicting the temperatures. The mean average percentage error (MAPE) is 3.46%. The increase of water temperature are predicted with a MAPE of 3.64%. Despite that the error over the thermal resistance is small since it is equal to 1.9%, the decreasing trend at low heat load is not represented. The explanation is that for these heat fluxes, the predicted chip temperatures show

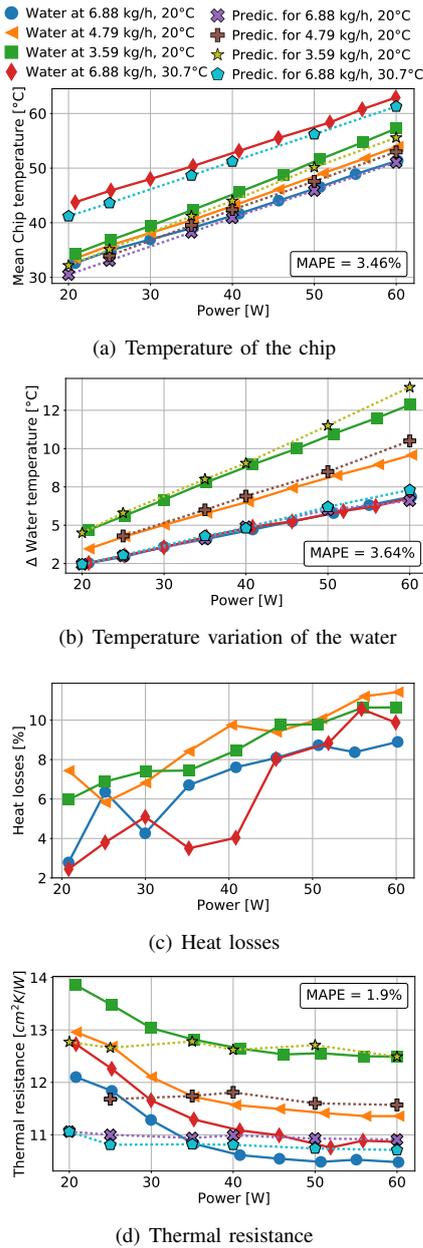


Fig. 11. Steady state results. Mean average percentage error (MAPE) is also shown

a larger error than for high heat fluxes and thus present a more horizontal curve.

Transient Measurements

In order to mimic the dynamics of real chips, additional dynamic measurements were performed with the thermosyphon charged with a FR=57%: from the same operating point (23.7 W), 3 heat load jumps were applied, 6.27 W, 14.56 W and 23.2 W (see Fig. 12). The initial state was set with inlet water at 18.9 °C, 6.88 kg/h, and a FR=0.57. For the first jump, the chip temperature needed 40 seconds to reach its steady temperature while for the second and the third jumps 60 and 90 seconds were necessary, respectively. Our designed thermosyphon properly follows all these jumps.

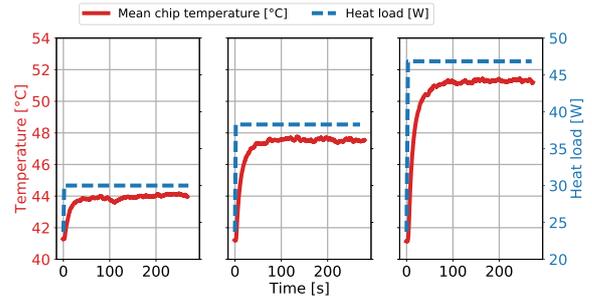


Fig. 12. Transient measurements

CONCLUSION

To enable energy-efficient HPC systems, there is a need to develop cooling solutions able to cool down high heat-flux current and future chips with the lowest possible energy expense. In this work, we presented a simplified numerical simulation model that we used to design of an ultra-compact micro-scale thermosyphon of 8 cm height, specifically tailored for Virtex 7 FPGAs chips with an area of 45 mm × 45 mm and a maximum TDP of 60 W. We manufactured a micro-scale thermosyphon prototype and validated it against real measurements for the whole power supply range allowed in the commercial FPGA chip, while always keeping chip temperature below 65 °C for the whole power supply range. Our results enable the usage of our thermosyphon solution in a 10 kW rack as part of a heterogeneous HPC prototype deployment in a data center belonging to a relevant industry in Switzerland, with an expected Power Usage Effectiveness (PUE) below 1.05.

Acknowledgements

This work has been partially supported by the YINS RTD project (No. 20NA21_150939), funded by Nano-Tera.ch with Swiss Confederation Financing and scientifically evaluated by SNSF, the EC H2020 MANGO FET-HPC project (Agreement No. 671668), and the ERC Consolidator Grant COM-PUSAPIEN (Agreement No. 725657)

References

- [1] T. J. Breen, E. J. Walsh, J. Punch, A. J. Shah, and C. E. Bash, "From chip to cooling tower data center modeling: Part i influence of server inlet temperature and temperature rise across cabinet," in *12th ITherm Conf.*, June 2010, pp. 1–10.
- [2] J. K. Matt Stansberry, "Uptime institute 2013 data center industry survey," Uptime Institute, Tech. Rep., 2013.
- [3] M. K. Tiwari, S. Zimmermann, C. S. Sharma, F. Alfieri, A. Renfer, T. Brunswiler, I. Meijer, B. Michel, and D. Poulikakos, "Waste heat recovery in supercomputers and 3d integrated liquid cooled electronics," in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITherm), 2012 13th IEEE Intersociety Conference on.* IEEE, 2012, pp. 545–551.
- [4] N. Lamaison, J. B. Marcinichen, C. L. Ong, and J. R. Thome, "Two-phase mini-thermosyphon electronics cooling, Part 3: Transient modeling and experimental validation," in *15th ITherm Conf.* IEEE, 2016, pp. 589–598.
- [5] N. Lamaison, C. L. Ong, J. B. Marcinichen, and J. R. Thome, "Two-phase mini-thermosyphon electronics cooling: Dynamic modeling, experimental validation and application to 2u servers," *Applied Thermal Engineering*, vol. 110, pp. 481–494, Jan. 2017.
- [6] J. B. Marcinichen, N. Lamaison, C. L. Ong, and J. R. Thome, "Two-phase mini-thermosyphon electronics cooling, Part 2: Model and steady-state validations," in *15th ITherm Conf.* IEEE, 2016, pp. 582–588.

- [7] C. L. Ong, N. Lamaison, J. B. Marcinichen, and J. R. Thome, "Two-phase mini-thermosyphon electronics cooling, part 1: Experimental investigation," in *Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM)*, 2016 15th IEEE Intersociety Conference on. Ieee, 2016, pp. 574–581.
- [8] C. L. Ong, R. L. Amalfi, J. B. Marcinichen, N. Lamaison, and J. R. Thome, "Two-phase mini-thermosyphon for cooling of datacenters: Experiments, modeling and simulations," in *ASME 2017 International Technical Conference and Exhibition on Packaging and Integration of Electronic and Photonic Microsystems collocated with the ASME 2017 Conference on Information Storage and Processing Systems*. American Society of Mechanical Engineers, 2017, pp. V001T02A003–V001T02A003.
- [9] N. Lamaison, J. B. Marcinichen, C. L. Ong, and J. R. Thome, "Two-phase mini-thermosyphon electronics cooling, Part 4: Application to 2u servers," in *15th ITherm Conf.* IEEE, 2016, pp. 599–609.
- [10] S. V. Patankar and D. B. Spalding, "A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows," *International journal of heat and mass transfer*, vol. 15, no. 10, pp. 1787–1806, 1972.
- [11] J. Van Doormaal and G. Raithby, "Enhancements of the simple method for predicting incompressible fluid flows," *Numerical heat transfer*, vol. 7, no. 2, pp. 147–163, 1984.
- [12] I. H. Bell, J. Wronski, S. Quoilin, and V. Lemort, "Pure and pseudo-pure fluid thermophysical property evaluation and the open-source thermo-physical property library coolprop," *Industrial & engineering chemistry research*, vol. 53, no. 6, pp. 2498–2508, 2014.
- [13] D. Beattie and P. Whalley, "A simple two-phase frictional pressure drop calculation method," *International Journal of Multiphase Flow*, vol. 8, no. 1, pp. 83–87, 1982.
- [14] Y. S. Muzychka and M. M. Yovanovich, "Laminar Forced Convection Heat Transfer in the Combined Entry Region of Non-Circular Ducts," *Journal of Heat Transfer*, vol. 126, no. 1, p. 54, 2004.
- [15] —, "Pressure Drop in Laminar Developing Flow in Noncircular Ducts: A Scaling and Modeling Approach," *Journal of Fluids Engineering*, vol. 131, no. 11, p. 111105, 2009.
- [16] R. K. Shah and A. L. London, "Laminar flow forced convection heat transfer and flow friction in straight and curved ducts—a summary of analytical solutions," Stanford Univ CA Dept of Mechanical Engineering, Tech. Rep., 1971.
- [17] V. Dharaiya and S. G. Kandlikar, "Numerical investigation of heat transfer in rectangular microchannels under h2 boundary condition during developing and fully developed laminar flow," *Journal of heat transfer*, vol. 134, no. 2, p. 020911, 2012.
- [18] V. Gnielinski, "New equations for heat and mass transfer in turbulent pipe and channel flow," *International Journal of Chemistry Engineering*, vol. 16, no. 2, pp. 359–368, 1976.
- [19] B. S. Petukhov, "Heat transfer and friction in turbulent pipe flow with variable physical properties," *Advances in heat transfer*, vol. 6, no. 503, p. i565, 1970.
- [20] E. Costa-Patry and J. R. Thome, "Flow pattern-based flow boiling heat transfer model for microchannels," *International Journal of Refrigeration*, vol. 36, no. 2, pp. 414–420, Mar. 2013.
- [21] M. Magnini and J. Thome, "An updated three-zone heat transfer model for slug flow boiling in microchannels," *International Journal of Multiphase Flow*, vol. 91, pp. 296–314, May 2017.
- [22] J. Thome, V. Dupont, and A. Jacobi, "Heat transfer model for evaporation in microchannels. Part I: presentation of the model," *International Journal of Heat and Mass Transfer*, vol. 47, no. 14–16, pp. 3375–3385, Jul. 2004.
- [23] V. Dupont, J. Thome, and A. Jacobi, "Heat transfer model for evaporation in microchannels. Part II: comparison with the database," *International Journal of Heat and Mass Transfer*, vol. 47, no. 14–16, pp. 3387–3401, Jul. 2004.
- [24] A. Cioncolini, J. R. Thome, and C. Lombardi, "Unified macro-to-microscale method to predict two-phase frictional pressure drops of annular flows," *International Journal of Multiphase Flow*, vol. 35, no. 12, pp. 1138–1148, Dec. 2009.
- [25] A. Cioncolini and J. R. Thome, "Algebraic turbulence modeling in adiabatic and evaporating annular two-phase flow," *International Journal of Heat and Fluid Flow*, vol. 32, no. 4, pp. 805–817, Aug. 2011.
- [26] —, "Entrained liquid fraction prediction in adiabatic and evaporating annular two-phase flow," *Nuclear Engineering and Design*, vol. 243, pp. 200–213, Feb. 2012.
- [27] —, "Void fraction prediction in annular two-phase flow," *International Journal of Multiphase Flow*, vol. 43, pp. 72–84, Jul. 2012.
- [28] S. Koyama, K. Kuwahara, K. Nakashita, and K. Yamamoto, "An experimental study on condensation of refrigerant R134a in a multi-port extruded tube," *International Journal of Refrigeration*, vol. 26, no. 4, pp. 425–432, Jun. 2003.
- [29] J. El Hajal, J. Thome, and A. Cavallini, "Condensation in horizontal tubes, part 1: two-phase flow pattern map," *International Journal of Heat and Mass Transfer*, vol. 46, no. 18, pp. 3349–3363, Aug. 2003.
- [30] J. Thome, J. El Hajal, and A. Cavallini, "Condensation in horizontal tubes, part 2: new heat transfer model based on flow regimes," *International Journal of Heat and Mass Transfer*, vol. 46, no. 18, pp. 3365–3387, Aug. 2003.
- [31] S.-M. Kim and I. Mudawar, "Universal approach to predicting heat transfer coefficient for condensing mini/micro-channel flow," *International Journal of Heat and Mass Transfer*, vol. 56, no. 1–2, pp. 238–250, Jan. 2013.
- [32] —, "Universal approach to predicting two-phase frictional pressure drop for adiabatic and condensing mini/micro-channel flows," *International Journal of Heat and Mass Transfer*, vol. 55, no. 11–12, pp. 3246–3261, 2012.
- [33] L. Friedel, "Improved friction pressure drop correlations for horizontal and vertical two phase pipe flow," *European Two Phase Flow Group Meeting, Ispra*, 1979.
- [34] V. Gnielinski, "A simplex method for function minimization," *Int. J. Chem. Engineering*, vol. 7, pp. 308–313, 1965.
- [35] H. Bielinski and J. Mikielwicz, "Natural circulation in single and two phase thermosyphon loop with conventional tubes and minichannels," *Heat Transfer - Mathematical Modelling, Numerical Methods and Information Technology*, 2011.
- [36] A. Cicchiti *et al.*, "Two-phase cooling experiments - pressure drop, heat transfer and burnout measurements," *Energ. Nucl.* 7, pp. 407–425, 1960.
- [37] V. V. Dharaiya and S. G. Kandlikar, "Numerical investigation of heat transfer in rectangular microchannels under H2 boundary condition during developing and fully developed Laminar flow," *Journal of Heat Transfer*, vol. 134, no. 2, p. 020911, 2012.
- [38] V. Gnielinski, "New equations for heat and mass transfer in turbulent pipe and channel flow," *Int. J. Chem. Engineering*, vol. 16, pp. 359–368–564, 1976.