



AQUASAR

Direct Re-Use of Waste Heat from Liquid-Cooled Supercomputers

Towards low power, high performance, zero-emission computing and datacenters

This project started in 2009.

Scope of project

The Information Technology (IT) industry has a key role in the global effort to reduce carbon-dioxide emissions as it consumes 2% of the world energy with a strong, demand-driven, upward trend. The IT industry already strives to digitally control and to optimize emission-intensive hardware and software processes. Examples include the replacement of traditional activities with digital information processing thus decoupling knowledge transfer from physically moving material or persons, resulting in an order of magnitude decrease in emissions. In the field of simulation based sciences, current computational approaches have an efficiency that is on average below 20% in terms of the hardware utilization. With computation becoming the third pillar of scientific enquiry along with experiment and theory in all scientific disciplines, it is essential to address the challenge of energy aware and low power high performance computing.

Summary of project

The goal of the present project is to demonstrate, for the first time, the employment of an emerging processor architecture, exhibiting a higher efficiency in terms of MFlops/W, for the solution of challenging scientific problems and to harness the heat that is inevitably generated as close as possible to the source with a very high temperature level. This allows, for example, the heat to be re-used for space heating or process heat, thus replacing combustion or other processes at the second user and offsetting the carbon dioxide emission by the generation of electricity. The goal is to demonstrate a high performance, low power consumption datacenter/computing operation approaching zero-net-emission.

As a first step, we will investigate and determine the exergy and energy efficiencies of a liquid cooled high performance system making a large fraction of the used energy available for space heating. The system uses ultra-high performance liquid coolers for CPU, interfaces, and DC-converters to collect 70–80% of the sensible heat at a >50°C temperature level and eliminates the need for a chiller for both free cooling in the summer and energy re-use in the winter.

The first objective is to characterize and co-optimize the electrical system-efficiency and the temperature level of the recovered coolant in order to minimize the effective emission of the system. The energy and exergy efficiency will be measured at different operation points with energy re-use at different temperature levels.

While the initial hardware will only capture 75% of the energy in liquid form, a small-scale demonstrator will allow us to demonstrate that up to 100% can be recovered, even at increased estimated temperature levels of 60–70°C. Higher output coolant temperature will be implemented through better CPU coolers that reduce the overall thermal resistance and focus the coolant flow and heat removal on the hottest regions of the chip. The target is to increase the coolant temperature another 5–8°C without increasing the transistor junction temperatures at the hottest regions of the chip. A two-phase microscale flow boiling cooler will also be investigated as a pumped system, whose potential is greatly reduced pumping power consumption (as much as 81% less) but requiring a higher level of system complexity and control. In addition we will further increase the temperature levels by means of introduction of a compressor in place of the

Main Investigator
Dimos Poulikakos, ETHZ

Project Partners
EPFL
ETHZ
IBM Research

pump to output higher grade waste heat (at the expense of higher energy consumption of the compressor). To reach the target heat transfer rates, the flow boiling process will use controlled bubble generation and throttles to optimize the flow distribution. These two microscale flow boiling cooling systems will be studied to minimize emission of the system and their performances will be compared to the liquid cooled system.

Goals and activities

In order to demonstrate the computational capabilities of an energy efficient computing system, we will implement particle based computational methods and will perform multiscale flow simulations as they pertain to the problems at the interface of nanotechnology and fluid mechanics. The software will be developed in synergy with a temperature and energy monitoring and control program that runs in parallel with the simulations and records the system efficiency and the emission as a function of CPU clock speed, supply voltage, coolant temperature and system load.

We will also optimize the overall efficiency of the scientific codes that run the system in MFlops/W compared to competing systems. The target is to exceed the current performance record of 488 MFlops/W. Furthermore we will use the system productively during operation parameter evaluations and to demonstrate record level values of MFlops/g CO₂. We will establish new frontiers in power aware computing and we will demonstrate that the solution of challenging scientific problems need not be adverse to the environment.

Expected results

- Operative supercomputer installation at ETH that runs in >75% direct energy-reuse mode at a >60°C temperature level to save 40% of datacenter energy and reduce emissions to a level below 15% compared to an air cooled datacenter (for both liquid cooled and microscale flow boiling cooled systems).
- Small scale prototypes of systems with 100% energy re-use and of systems that provide output temperature levels above 70°C demonstrated (for both liquid cooled and two-phase boiling cooled systems).
- State of the art multiscale computational methods implemented on manycore processor architectures. Applications on large scale effects of nanopatterned surfaces.
- Multiscale flow simulations, used in order to constantly analyze the performance in MFLOPS/g emitted CO₂ and full year record of data with application codes running.
- Software optimized with an efficiency >500 MFlop/W.

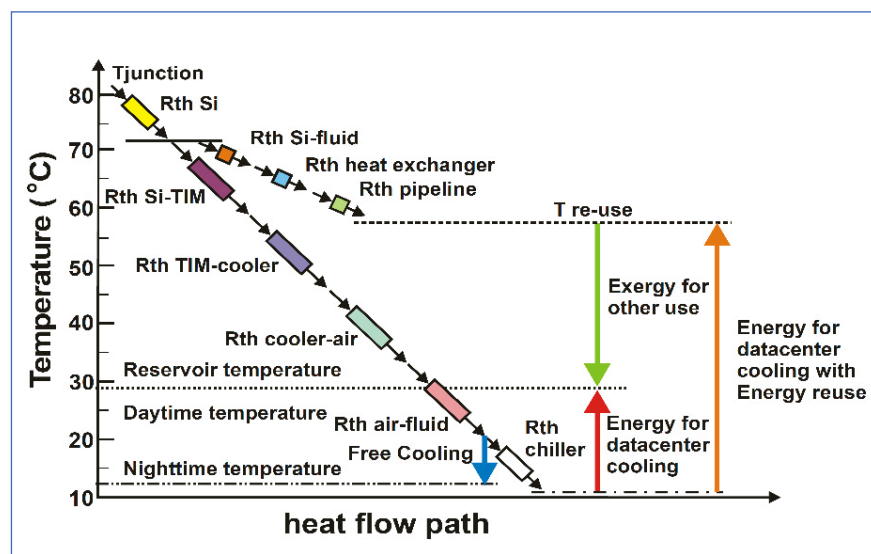


Figure 1: Thermal resistances from the transistor junction to the ambient for a normal air-cooled datacenter ($R_{th, Si} + R_{th, Si-TIM} + R_{th, TIM-cooler} + R_{th, cooler-air} + R_{th, air-fluid} + R_{th, chiller}$) and for an optimized liquid-cooled datacenter in which all resistances have been minimized and the number of components is smaller ($R_{th, Si} + R_{th, Si-fluid} + R_{th, heat\ exchanger} + R_{th, pipeline}$). Both cases start at the same transistor temperature $T_{junction}$. The small solid line shows that the liquid cooling solution captures the heat directly on the chip. The dotted line shows the reservoir or daytime temperature and the dash-double-dotted line shows nighttime temperature with free-cooling. The dashed line shows the energy re-use temperature and the dash-dotted line shows the internal chiller temperature. Vertical arrows show the available gradient for free-nighttime-cooling (blue), the available exergy for heating (green), and the required heat pump energy to either throw the heat away (red) or re-use it (orange). The vertical axis shows the temperature while the horizontal axis shows the heat flow path.