Temperature Control When Testing the RD53 Chip in Lab: Technical Report

Mario d'Andrea (University of California, Davis)
Supervisor: Andromachi Tsirou
August 17th, 2018

1. Introduction

The phase II readout chip (RoC) for Si-detectors[1] generates a significant amount of heat during operation, typically around 2W but hypothetically up to 8W. This level of heat generation on such a small area demands some form of heat dissipation. The final implementation within the CMS detector will deal with this via the overall cooling system of the Si- Pixel/Tracker detector, a pressurized CO_2 based evaporation cooling system that will allow temperatures around -25^0 C. However, for the purposes of R&D we want a more flexible and manageable system for use in the lab.

Towards this end, it is adequate to attach a standard cpu heat sink to the thermal contact surface of the chip. With a fan and aluminum heat sink block, this should be able to dissipate the 8W at an operating temperature of under 40 C. However, this approach leaves the chip temperature floating as an uncontrolled variable. In addition to simply removing the heat we would like to be able to directly control the temperature. This allows for an investigation into the chip's temperature dependant properties.

Design Goals

The temperature control system should be able to:

- Dissipate up to 8W of chip power
- Operate over a temperature range large enough to reveal temperature dependant effects in the chip's operation[2] (the chip should not go above 30^0 C)
- Operate at chip temperatures at or below room temperature with adequate protection when around dewpoint.
- Most important point - Remain *stable* while within specified operating zones. That is, it should not deviate sizeably (< .5^{0}C) from the target temperature or undergo thermal runaway.

[1] The CMS Collaboration, CMS Technical Design Report for the Pixel Detector Upgrade, CERN-LHCC-2012-016. Chapters 5,6,9.
[2] It is unclear (at least to the author) what range this is

- Come to equilibrium at the requested temperature within a reasonable amount of time (<5 mins)
- Control the temperature of multiple chips (goal: 4) simultaneously and independently.

2. Implementation: Thermoelectric Cooling Overview

Seeing the requirements, a Peltier element chiller is a solid first approach. There are numerous commercial Peltier control systems that allow for controlling the power of a Peltier element using the achieved temperature at the "cold" side of it as controlling variable of the control process. I worked using the thermoelectric cooling (TEC) control system from Meerstetter engineering[3]. The system controls the power entering a Peltier element to regulate the object temperature (cold side of the Peltier), pumping heat from the chip to the heat sink according to the electrical current provided to the Peltier element. An overview of this setup is given in Fig. 1.

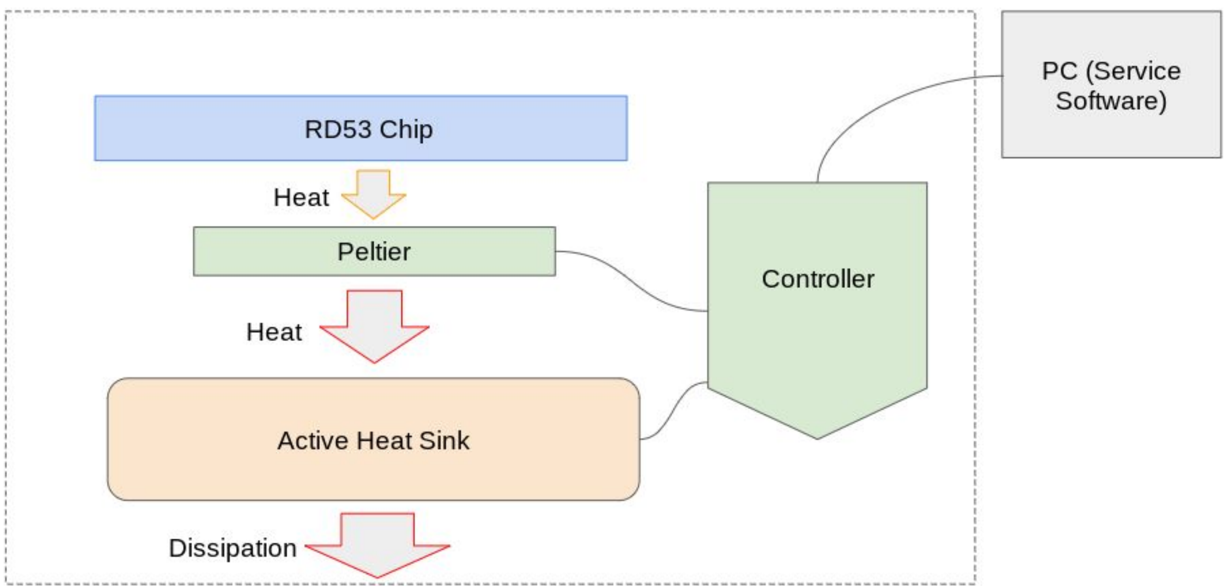

Figure 1. *The basic parts of the thermoelectric cooling system.*

The key players in the operation are the Peltier element, the heat sink, and the TEC controller. The TEC controller monitors the object temperature and constantly compares it to a target object temperature, r(t). The target temperature r(t) is under your control. The difference between the target object temperature and the measured object temperature $T_o(t)$ forms an error e(t) = $T_o(t)$ -

[3] See their website for detailed product information https://www.meerstetter.ch/

r(t). The TEC controller's mission is to drive this error function e(t) to zero via a PID control system. In this case, the control variable which the system uses to minimize |e(t)| is the current supplied to the Peltier. The process is made possible by, and is critically limited by, the ability of the heat sink to dissipate the heat delivered to it by the Peltier element.

Peltiers

Peltier elements are commonly used if we need temperature control where a cooling system is difficult to install/function. A Peltier element utilizes the thermoelectric effect to convert an electrical current into a temperature difference. By reversing the direction of current, the element can be made to drive heat in either direction.

In applications, the behavior of the device is specified by its spatial dimensions and operating parameters, I_{max}, V_{max}, ΔT_{max}, and Q_{max}.

Note that I_{max} and V_{max} are not actual limits on how much power can be supplied. Rather they are limits on how much power you would practically ever want to provide. Beyond this point, the cooling capabilities begin to deteriorate. This is because the device must transport not only the object heat, but also the resistive losses I^2 R associated with the power supplied to the device. Beyond the point (I_{max}, V_{max}), the diminishing rate of returns makes it futile to add more power. The TEC controller prevents the system from entering this inefficient region.

The most important parameters are I_{max}, V_{max} and ΔT_{max}, available from the devices' datasheets. These are needed to calibrate the TEC controllers properly.

Peltiers should never be driven without a heat sink. They are unable to dissipate heat effectively on their own, and will quickly overheat.

Temperature Sensors

General Information

In order to function, the controller must monitor the object temperature (and optionally the heat sink temperature too). The object temperature sensor is ideally placed as near as possible to the critical part of the object being cooled.

I have been using thermistors and resistance thermometers as temperature sensors. Both are just resistors whose resistance is dependant on temperature, more so than a standard resistor. There are two specific types of temperature sensors I have been using: NTC's and Pt's. An NTC (negative-temperature coefficient) is a thermistor whose resistance decreases with increasing temperature. A Pt (platinum) resistance thermometer is a piece of pure platinum, which also happens to have a positive temperature coefficient.

Application-Specific Information

Placement of temperature sensors has proven to be one of the more difficult aspects of designing the cooling system. The chips are delicate and small, and offer no immediate place to attach an external temperature sensor. The optimal object temperature measuring method would use the on-board NTC, but this comes with its own drawbacks. Each NTC is a little bit different. This can be corrected by adding an offset[4] and/or making precise measurements of resistance vs. temperature[5] and entering the results to calibrate the controller.

The current system uses a pt1000 as an object temperature sensor. It does not directly measure the chip temperature. Instead it sticks on the underside of the mounting plate, in effect monitoring the Peltier's cold side temperature. This was chosen because all pt1000 sensors are identical (every piece of platinum is the same).

PID Tuning

The TEC controller uses a PID (proportional-integral-derivative) algorithm to control the Peltier. There are a number of free parameters which significantly impact the system's performance[6]. In addition to the three PID parameters[7], there is a "course temperature ramp" parameter. Specified in [C/s], the controller will attempt to ramp the temperature at this rate when transitioning from one temperature to another.

The "proximity width" parameter specifies where the ramp stops and the PID begins.

---

[4] Possible in the Object Temperature and Sink Temperature tabs. Make two measurements of room temperature and compare to form an offset.

[5] The standard approximation method for an NTC involves measuring the resistance and temperature at three values distributed across the intended range of use. See Steinhart-Hart Equation. These values are entered into the Advanced -> Temperature Conversion tab of the service software.

[6] All these parameters can be changed in the Temperature Control tab

[7] These are proportional, integral, and derivative parameters: https://en.wikipedia.org/wiki/PID_controller

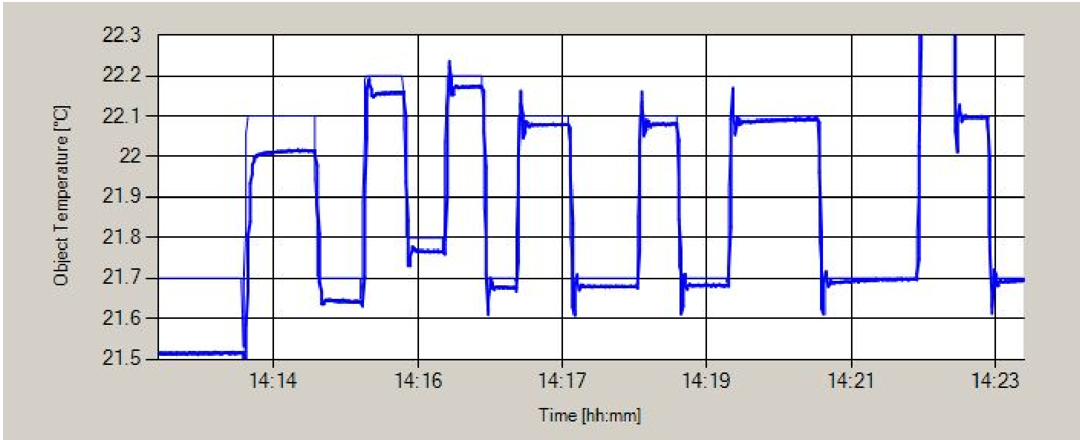

Figure: A large course temperature ramp creates some initial wobble as the system overshoots.

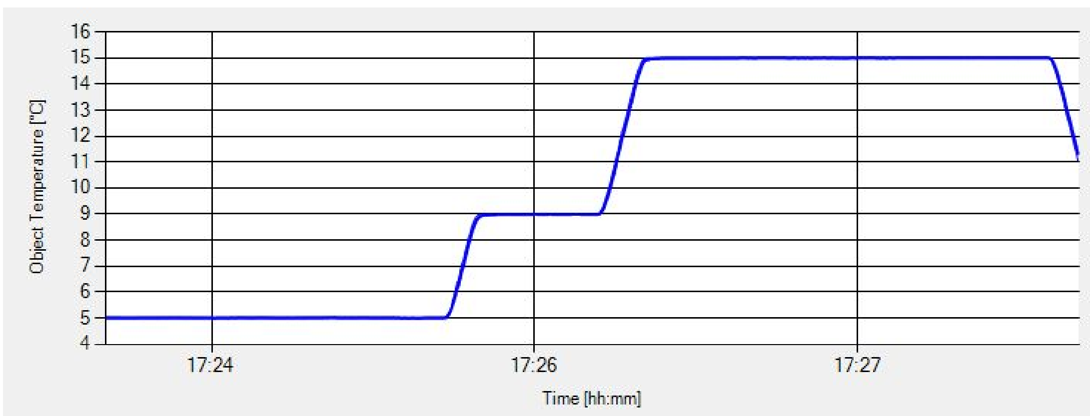

Figure: A well-tuned configuration smoothly transitions between temperatures and remains stable.

For the current system using the Adaptive™ ET-063-10-13-RS Peltier, I am using the following values:

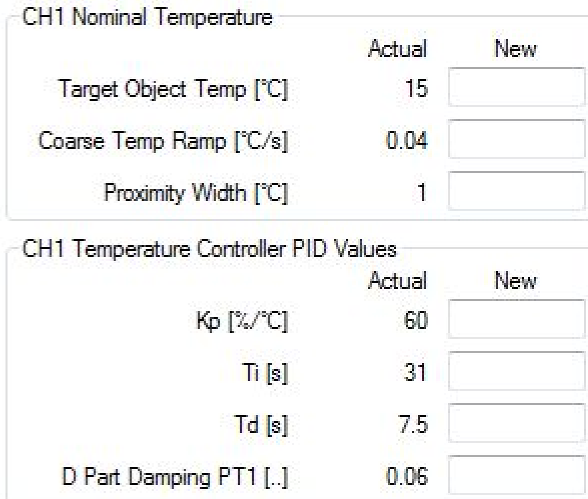

These are PID tuning parameters I have used with the Adaptive ET-063-10-13-RS Peltier and the Pt1000 temperature sensor. Found in the Temperature Control tab.

The service software supports an autotuning system. I have found the manual tuning to be more reliable; the autotuning system sometimes suggests values of Kp that are too large, causing instability.

Lab Set-Up

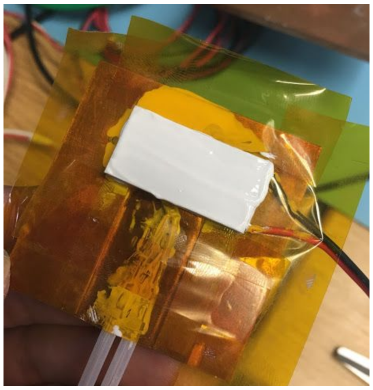

A Peltier element and temperature sensor are attached to the mounting plate.

The mounting plate is then secured to the RD53 chip with 3mm screws.

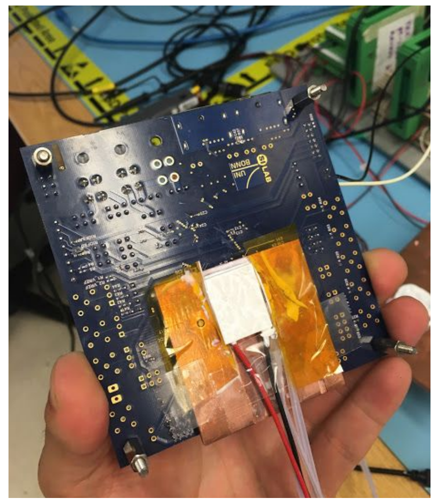

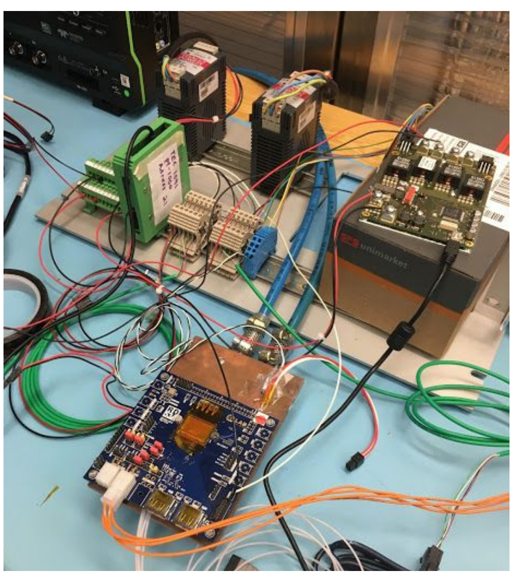

The assembly is secured to the heat sink. The Peltier current and object temperature are connected to the controller.

3. Results and Limitations

The system as it stands meets the design goals in section 1. Testing shows that it can draw up to 8W from the chip while keeping the chip anywhere between 15 and 30 C. The temperature stabilizes within a few minutes, and remains stable indefinitely[8]. The service software and controllers support an addressing mechanism[9], allowing multiple chips to be

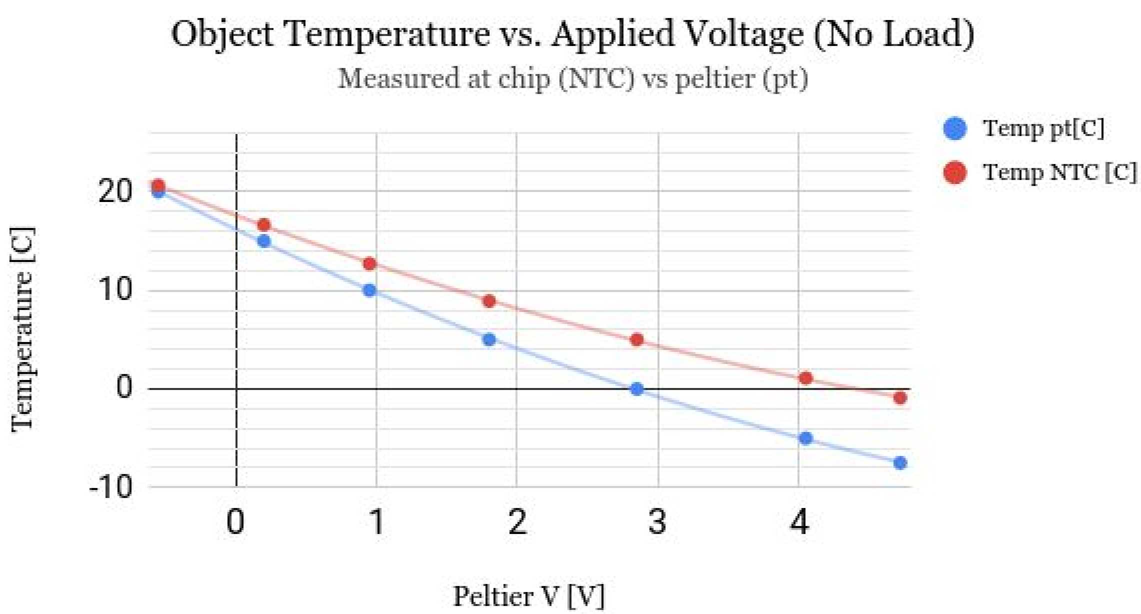

controlled and monitored independently and simultaneously.

A major limitation of the current system is the object temperature sensor. As discussed, the onboard NTC is not being used as the object temperature. Therefore, you cannot expect the chip to be held at the specified target temperature. The plots illustrate the difference between the target temperature and the chip temperature.

When there is no load applied to the chip, the disparity stays under 10° C while the chip itself can be brought down to 0° C. As the power supplied to the chip increases, so does the difference between the object temperature and chip temperature. The lowest achievable chip temperature also becomes much more limited. This is illustrated in the next two plots, where the measurements are repeated with chip power of 2W and 7.4W.

[8] Successfully kept the object temperature at 10 C for 20 hours, while never deviating more than 0.1 C.
[9] Device address is set in Operations tab and the device to be addressed is set in the Maintenance tab. More information in the user manual.

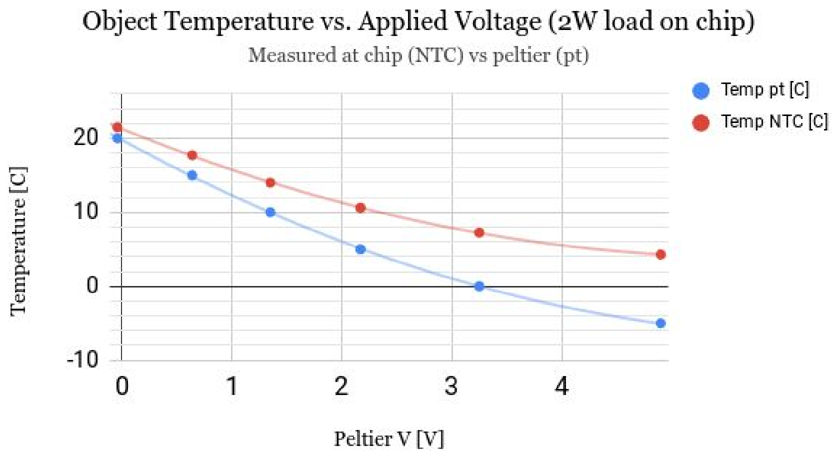

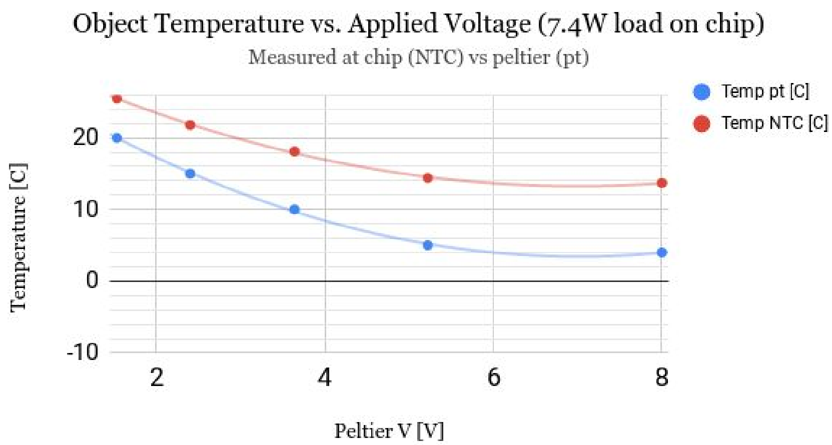

Details on the system's operating points can be found on the operating points spreadsheet:

- [v1 Operating Points Spreadsheet]

Another limitation is the amount of hardware. There is currently one liquid-cooled heatsink (large enough for only one card) and one custom card mounting piece. There are also currently two single-channel TEC controllers, each capable of driving a single Peltier.

4. Future Development

There are two problems worth attacking: (1) object temperature (2) scale.

(1) There are a variety of tradeoffs associated with using the onboard NTC versus a Pt to monitor the object temperature. Some of there are:
- The pt measurement will always be an inaccurate way of setting the chip temperature due to thermal resistance between the chip and Peltier.
- The pt needs a place to be attached, while the NTC is already in place.
- The TEC controllers' hardware must be modified before it can accept the NTC as an object temperature sensor (see chap 7.2 of the TEC Family-User Manual and the TEC-1091 Datasheet for more information).
- The Pt is a generic part, while the NTC's characteristics need to be entered and its performance tested for validity.

Ultimately either option is viable. I lean towards using a pt for the object temperature while passively monitoring the chip temperature via the onboard NTC.

(2) There has been talk that a good goal would be to handle up to four chips at a time. For this, a larger heat sink is needed, along with more custom card-mounting parts. Other parts such as Peltiers, Pts and controllers can just be ordered online.

The simplest solution is to order two more TEC 1091 single-channel controllers, as these are proven to work in this application. However, there are also already two TEC 1123 dual-channel controllers sitting in the lab, with hardware ready for Pt1000 object temp sensing. In the next week I hope to set up and test out one of these controllers. If they work the way I expect them to, the four Peltiers can be controlled by the two dual-channel controllers.

References

[1] The CMS Collaboration, *CMS Technical Design Report for the Pixel Detector Upgrade*, CERN-LHCC-2012-016

[2] CMS Collaboration. "*Technical Proposal for the Phase-II Upgrade of the CMS Detector"*, CERN-LHCC-2015-010 ; LHCC-P-008 ; CMS-TDR-15-02. CMS Collaboration ; CERN. Geneva. The LHC experiments Committee ; LHCC. 2015.