Fire and Ice: How Temperature Affects GPU Performance

Danny Price
Harvard-Smithsonian Center for Astrophysics
dprice@cfa.harvard.edu

CfA: Danny Price, Ben Barsdell, Lincoln Greenhill
NVIDIA: Mike Clark, Ron Babich

GTC 2014 | MARCH 24-27 | SAN JOSE, CA

Introduction

- Watts, leakage & performance
- GPUs in radio astronomy
- Experimental setup
- Results & conclusions

The Green 500

Listed below are the November 2013 The Green500's energy-efficient supercomputers ranked from 1 to 10.

Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)
1	**4,503.17**	GSIC Center, Tokyo Institute of Technology	TSUBAME-KFC - LX 1U-4GPU/104Re-1G Cluster, Intel Xeon E5-2620v2 6C 2.100GHz, Infiniband FDR, NVIDIA K20x	27.78
2	**3,631.86**	Cambridge University	Wilkes - Dell T620 Cluster, Intel Xeon E5-2630v2 6C 2.600GHz, Infiniband FDR, NVIDIA K20	52.62
3	**3,517.84**	Center for Computational Sciences, University of Tsukuba	HA-PACS TCA - Cray 3623G4-SM Cluster, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband QDR, NVIDIA K20x	78.77
4	**3,185.91**	Swiss National Supercomputing Centre (CSCS)	Piz Daint - Cray XC30, Xeon E5-2670 8C 2.600GHz, Aries interconnect , NVIDIA K20x Level 3 measurement data available	1,753.66
5	**3,130.95**	ROMEO HPC Center - Champagne-Ardenne	romeo - Bull R421-E3 Cluster, Intel Xeon E5-2650v2 8C 2.600GHz, Infiniband FDR, NVIDIA K20x	81.41
6	**3,068.71**	GSIC Center, Tokyo Institute of Technology	TSUBAME 2.5 - Cluster Platform SL390s G7, Xeon X5670 6C 2.930GHz, Infiniband QDR, NVIDIA K20x	922.54
7	**2,702.16**	University of Arizona	iDataPlex DX360M4, Intel Xeon E5-2650v2 8C 2.600GHz, Infiniband FDR14, NVIDIA K20x	53.62
8	**2,629.10**	Max-Planck-Gesellschaft MPI/IPP	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	269.94
9	**2,629.10**	Financial Institution	iDataPlex DX360M4, Intel Xeon E5-2680v2 10C 2.800GHz, Infiniband, NVIDIA K20x	55.62
10	**2,358.69**	CSIRO	CSIRO GPU Cluster - Nitro G16 3GPU, Xeon E5-2650 8C 2.000GHz, Infiniband FDR, Nvidia K20m	71.01

Source: http://www.green500.org/lists/green201311

The Green 500

#1

TSUBAME-KFC

Tokyo Institute of Technology
Oil-cooled

#2

Wilkes Cluster

University of Cambridge
Air-cooled

Source: http://www.green500.org/lists/green201311

The Square Kilometer Array

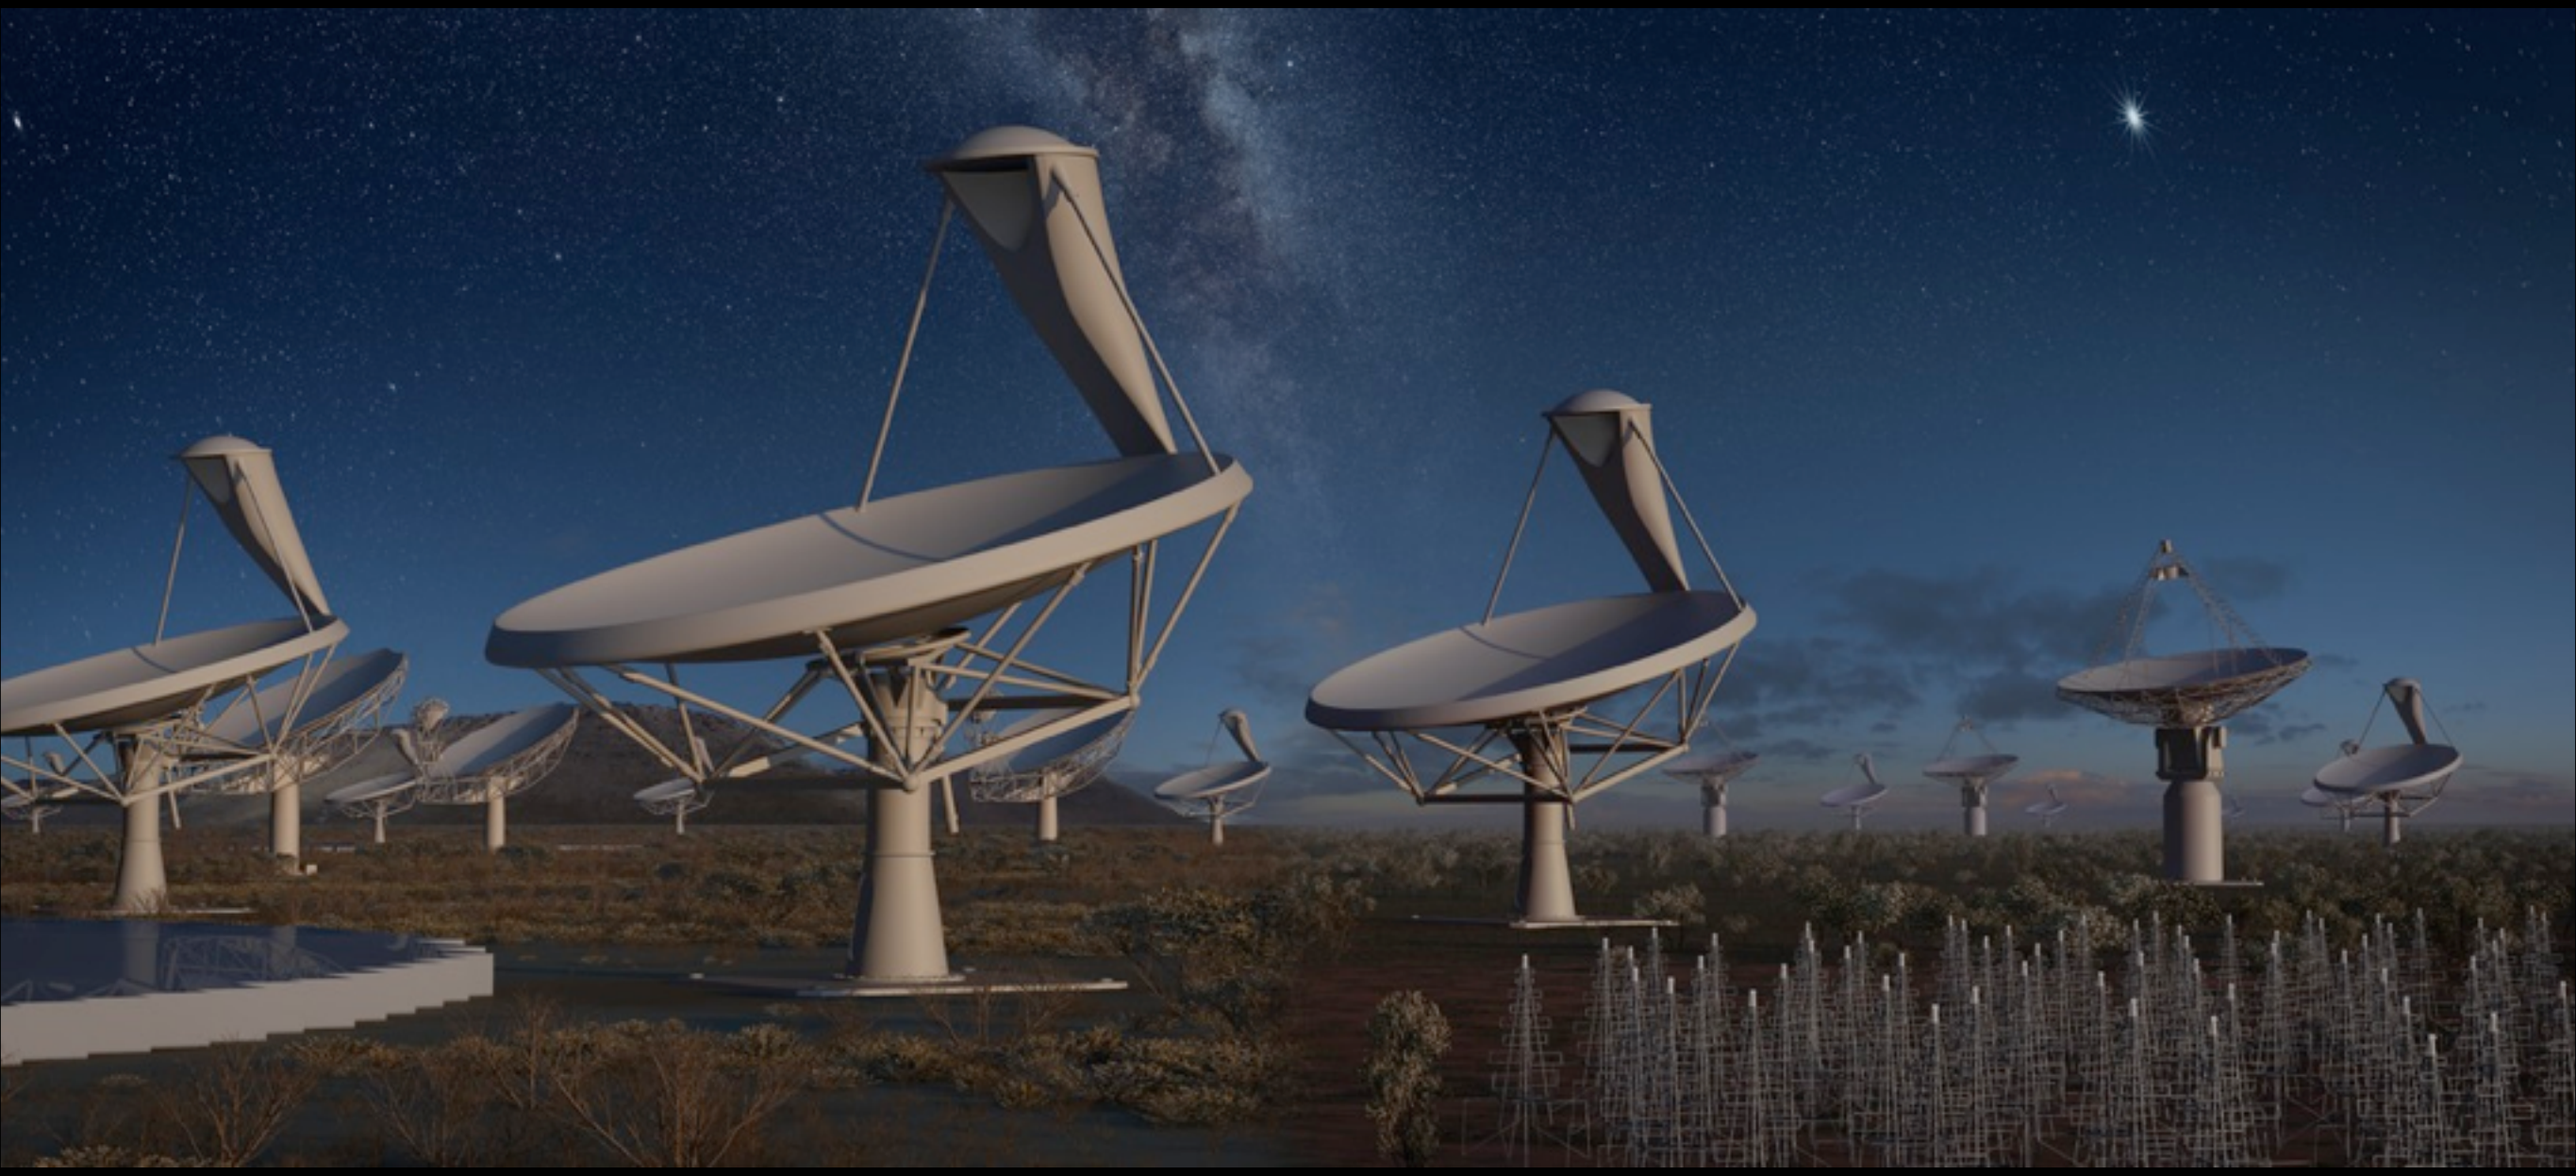

Image: www.skatelescope.org

Radio Interferometry - lots of computations

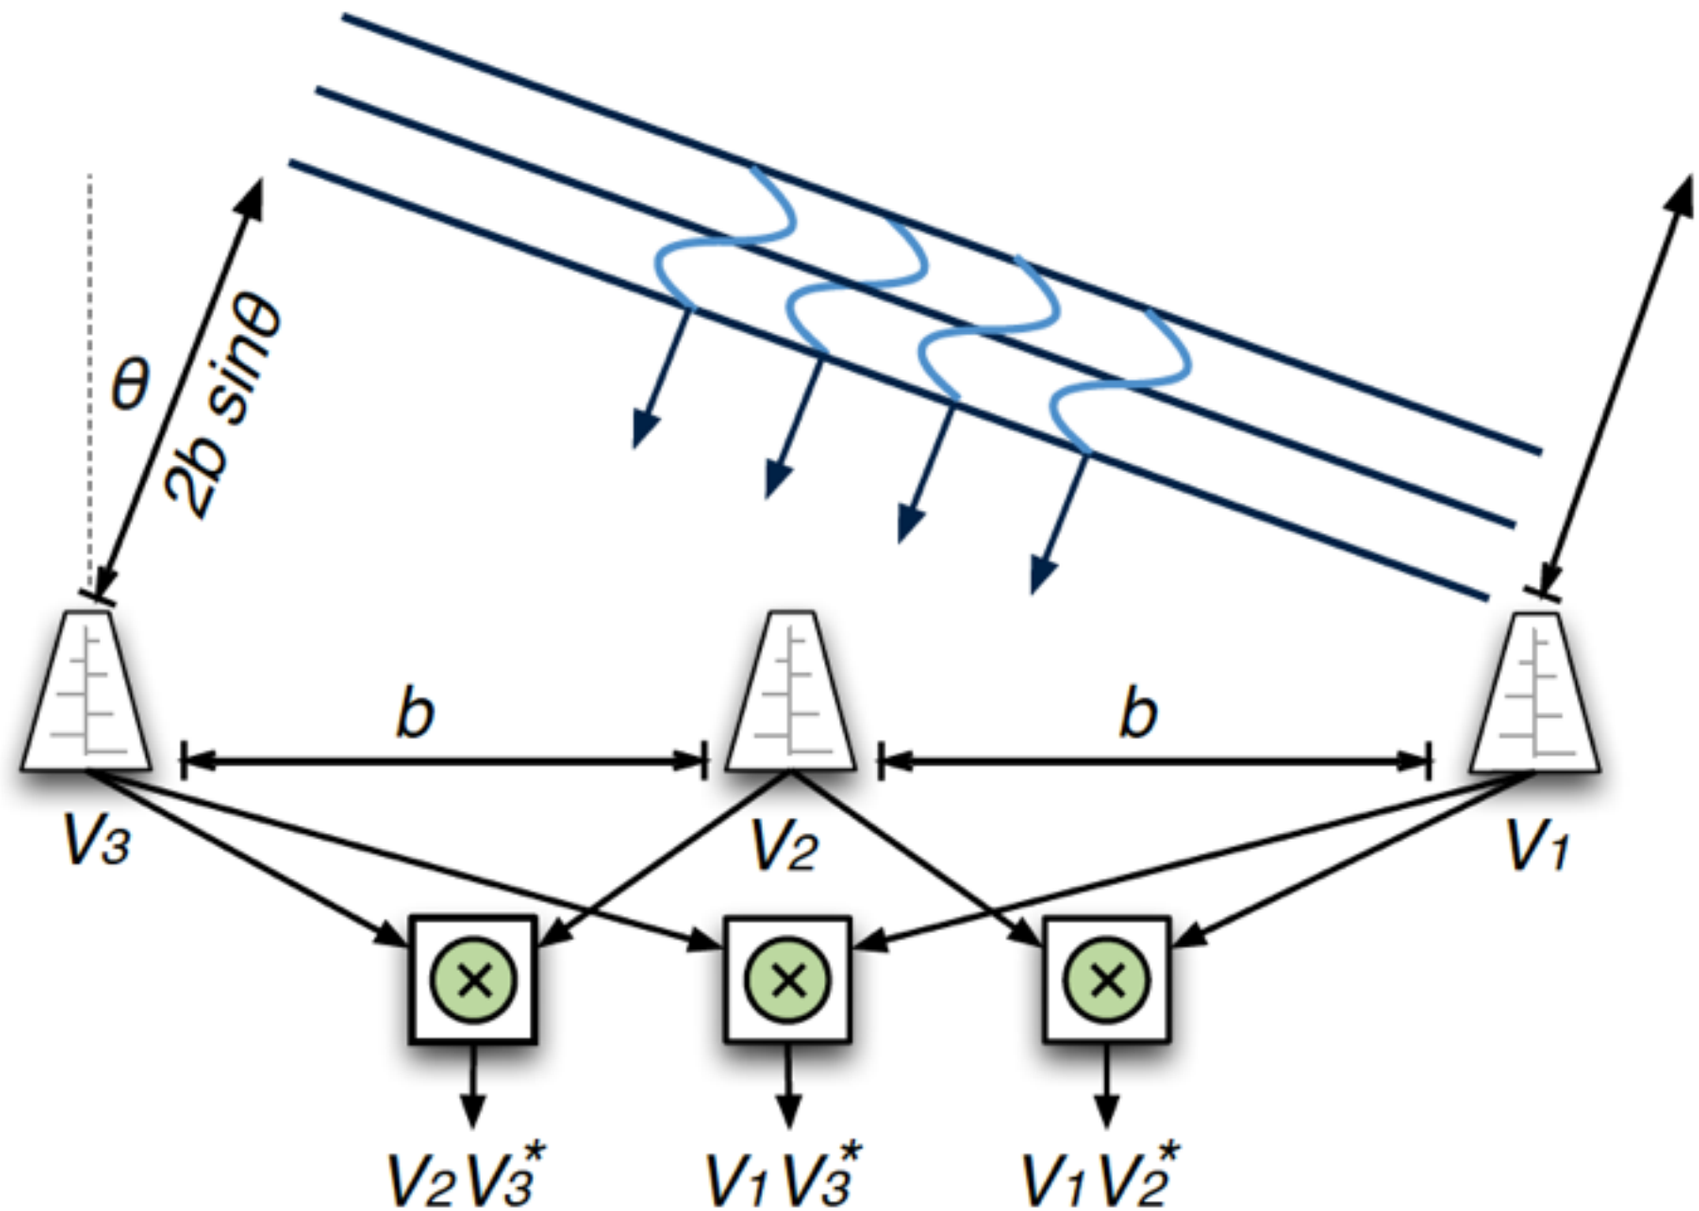

Radio Interferometry - lots of computations

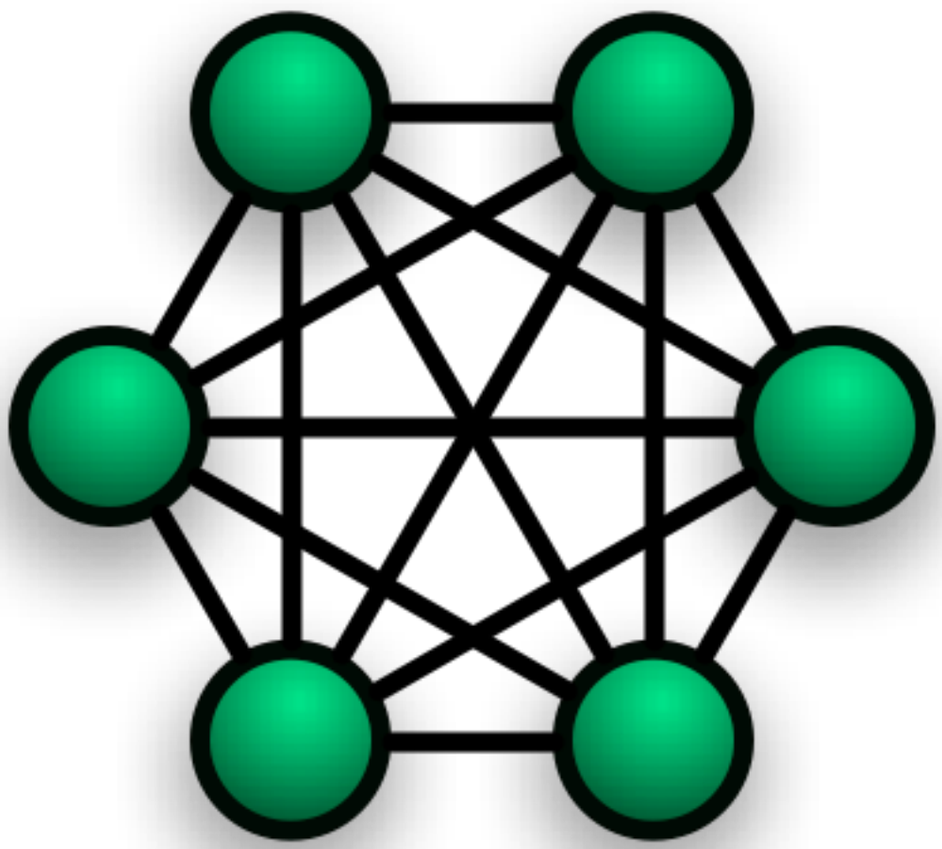

- More antennas = lots more baselines

$$N_B = \frac{N_A(N_A - 1)}{2}$$

for example, $N_A = 512 \Rightarrow N_B = 130816$

- Computations done on a frequency channel basis
- Trivially parallelizable: over frequency and baseline

xGPU - a GPU correlator for radio astronomy

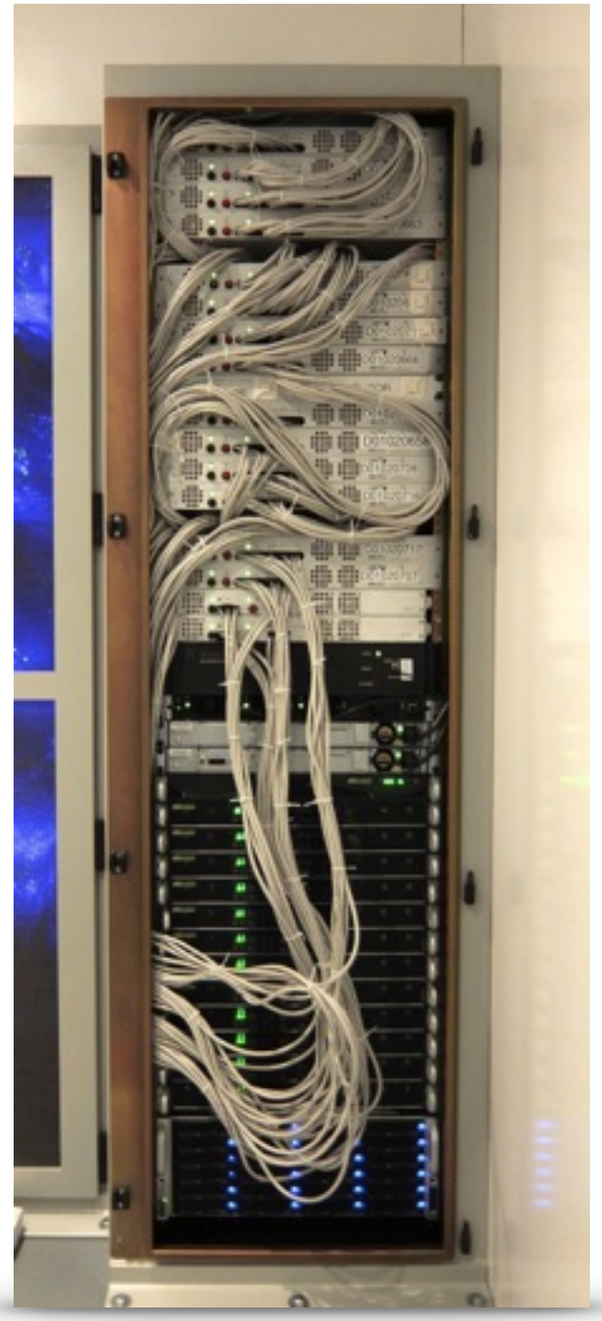

The LEDA-512 correlator

xGPU is a GPU-based cross correlator for radio astronomy.

It powers LEDA (the world's largest *Nant* radio telescope), PAPER, and the MWA (among others).

M. A. Clark, P. C. La Plante, and L. J. Greenhill, "Accelerating Radio Astronomy Cross-Correlation with Graphics Processing units", *[arXiv:1107.4264 [astro-ph]].*

Authors:
Michael Clark (NVIDIA)
Paul La Plante (Loyola University Maryland)
Lincoln Greenhill (Harvard-Smithsonian Center for Astrophysics)
David MacMahon (University of California, Berkeley)
Ben Barsdell (Harvard-Smithsonian Center for Astrophysics)

Watts, leakage & performance

$$P = P_{dynamic} + P_{static}$$

Leakage power (bad)

Watts, leakage & performance

Switching frequency

$$P_{dynamic} = CfV^2$$

Transistor capacitance

Voltage

Watts, leakage & performance

- Leakage mechanisms:
  - Reverse-biases junction leakage
  - Gate-induced drain leakage
  - Gate direct-tunneling leakage
  - Subthreshold (weak inversion) leakage

Watts, leakage & performance

- Leakage mechanisms:
  - Reverse-biases junction leakage
  - Gate-induced drain leakage
  - Gate direct-tunneling leakage
  - Subthreshold (weak inversion) leakage

How temperature affects performance

$$I_{subthreshold} = A_s \frac{W}{L} \left(\frac{kT}{q} \right)^2 e^{\frac{q(V_S - V_{th})}{nkT}}$$

How temperature affects performance

$$I_{subthreshold} = A_s \frac{W}{L} \left(\frac{kT}{q} \right)^2 e^{\frac{q(V_S - V_{th})}{nkT}}$$

Thermal voltage

8.62×10^{-5} eV/K

26 mV at room temperature

How temperature affects performance

Device specifics:

A: Technology specific constant

L, W: device channel length & width

$$I_{subthreshold} = A_s \frac{W}{L} \left(\frac{kT}{q}\right)^2 e^{\frac{q(V_S - V_{th})}{nkT}}$$

Thermal voltage

8.62×10^{-5} eV/K

26 mV at room temperature

How temperature affects performance

Device specifics:

A: Technology specific constant

L, W: device channel length & width

Voltage terms:

Vs: Gate to source voltage

Vth: Switching threshold voltage

n: transistor subthreshold swing coeff

$$I_{subthreshold} = A_s \frac{W}{L} \left(\frac{kT}{q}\right)^2 e^{\frac{q(V_S - V_{th})}{nkT}}$$

Thermal voltage

8.62×10^{-5} eV/K

26 mV at room temperature

Results Part 0:

Green GPU: A water-cooled test rig

Green GPU: A water-cooled test rig

- GreenGPU specifications:
  - Gigabyte GA-Z68MX motherboard
  - Intel Core i7-2600
  - 16GB RAM
  - **NVIDIA Tesla K20m**
  - EK-FCTK20 water block

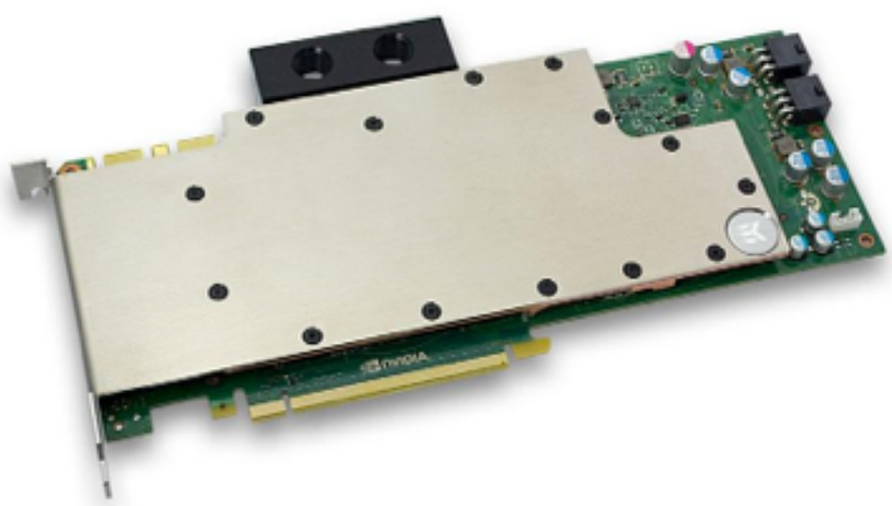

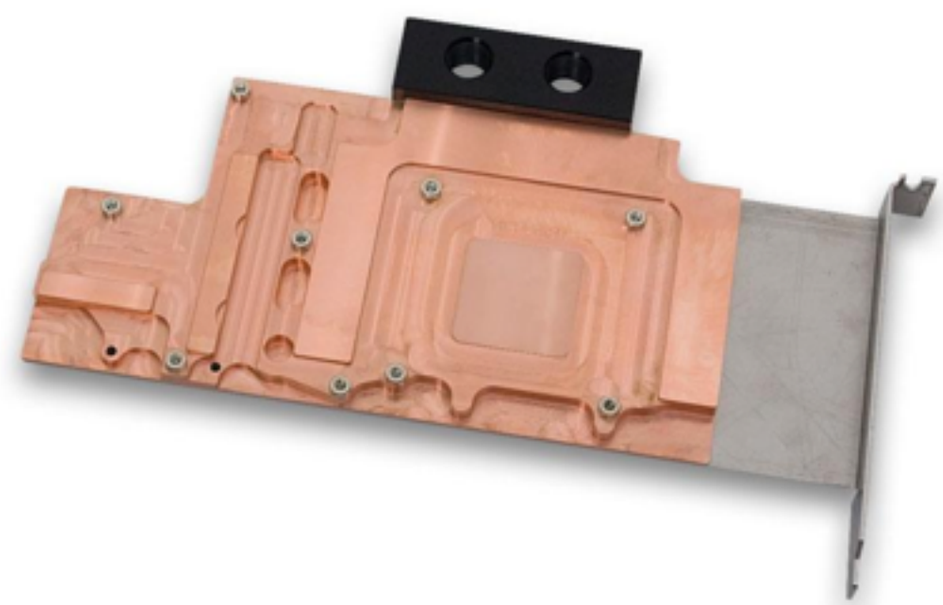

Warranty void: K20m from a HP server

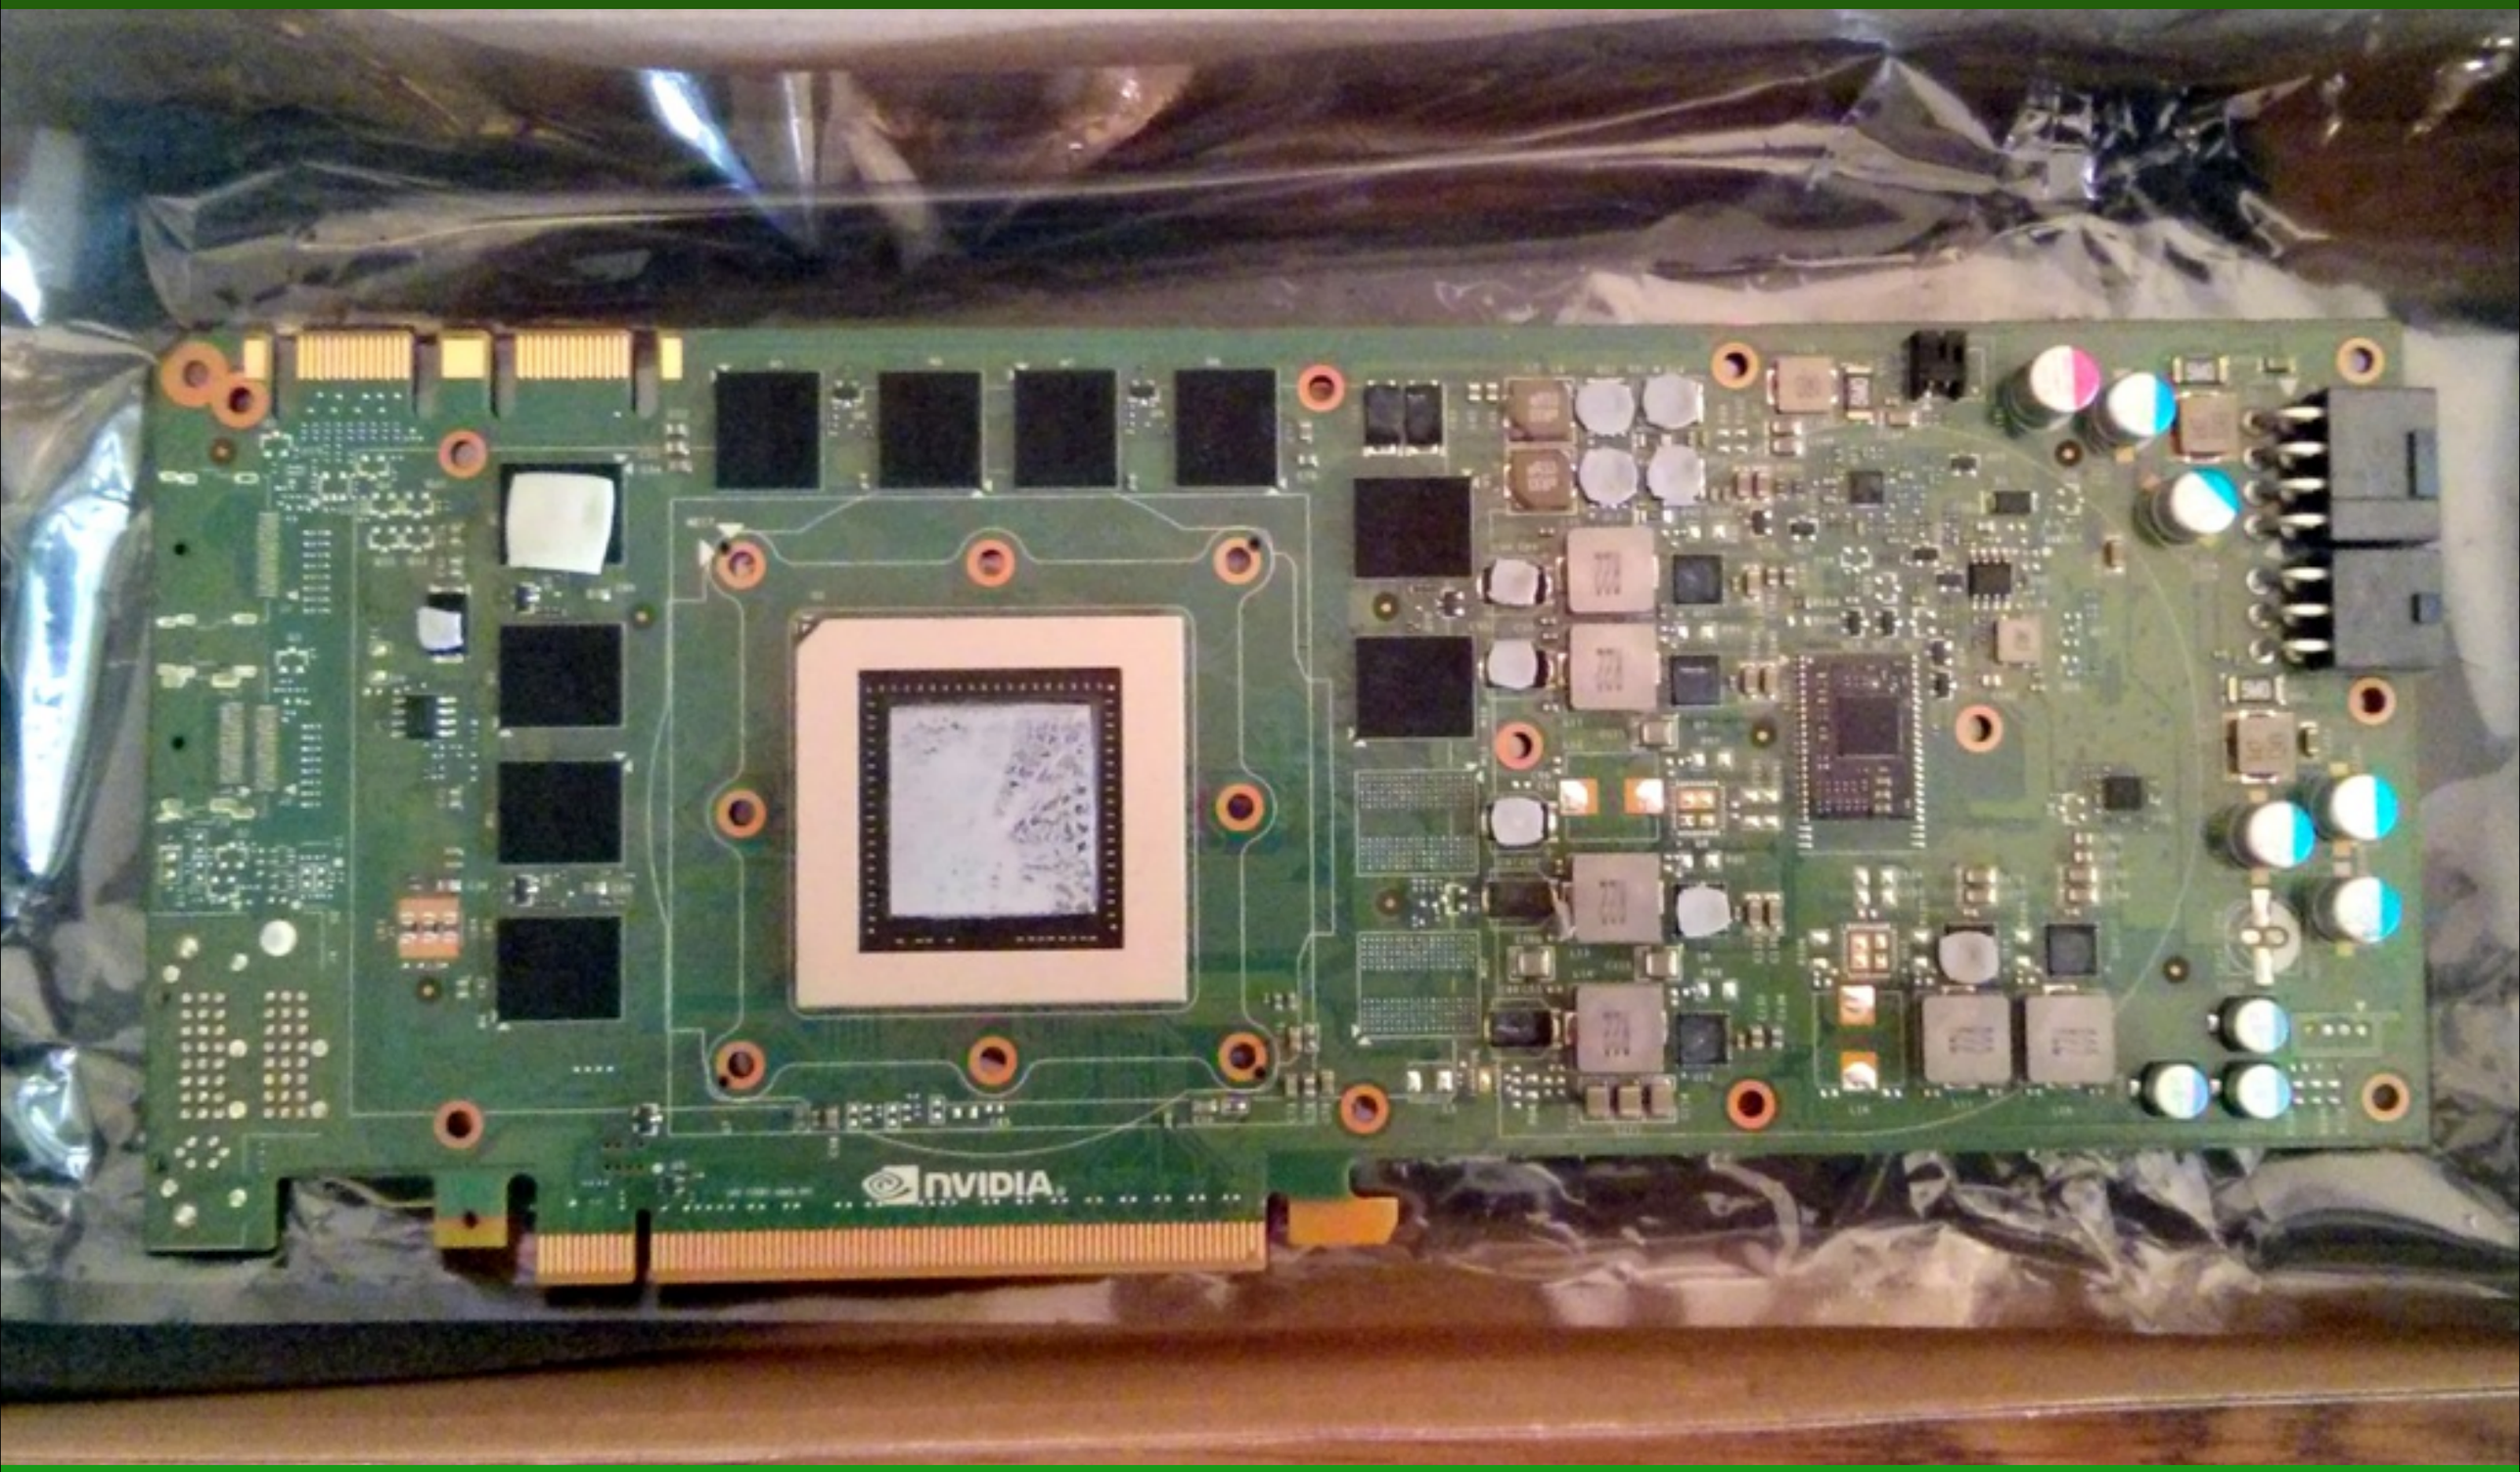

Warranty void: K20m from a HP server

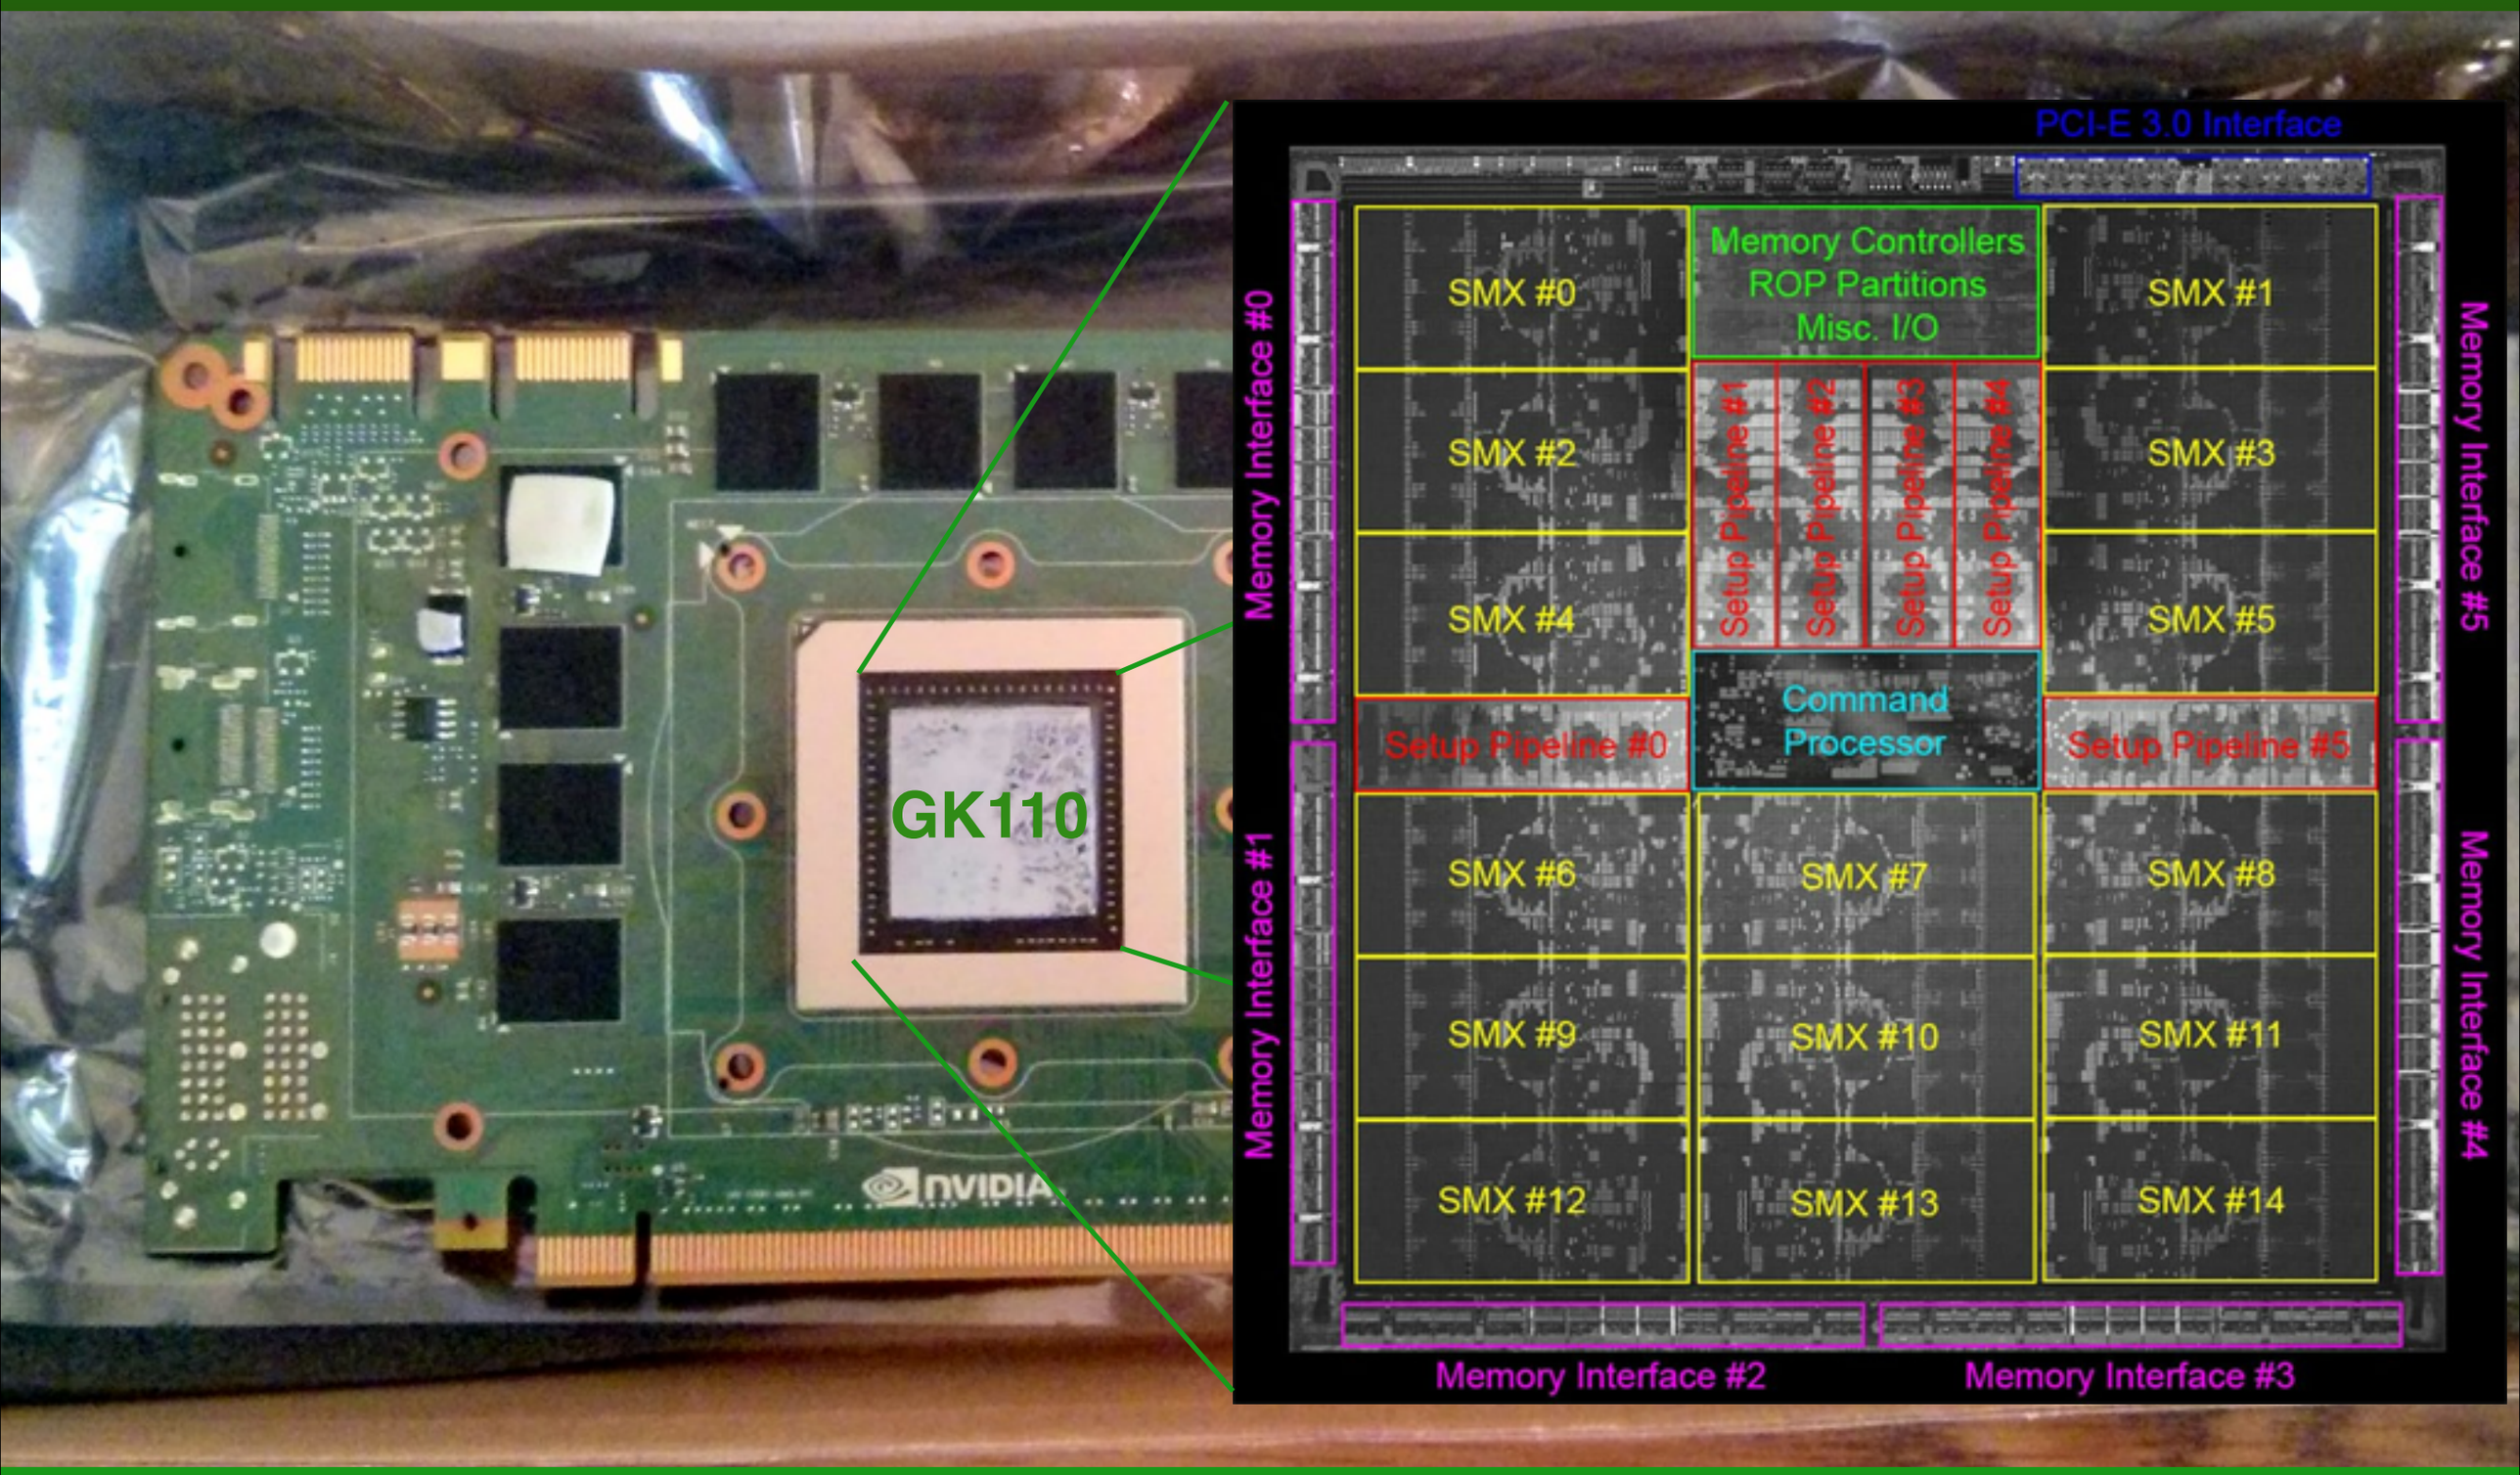

GK110 die image: news.softpedia.com

Green GPU

Green GPU

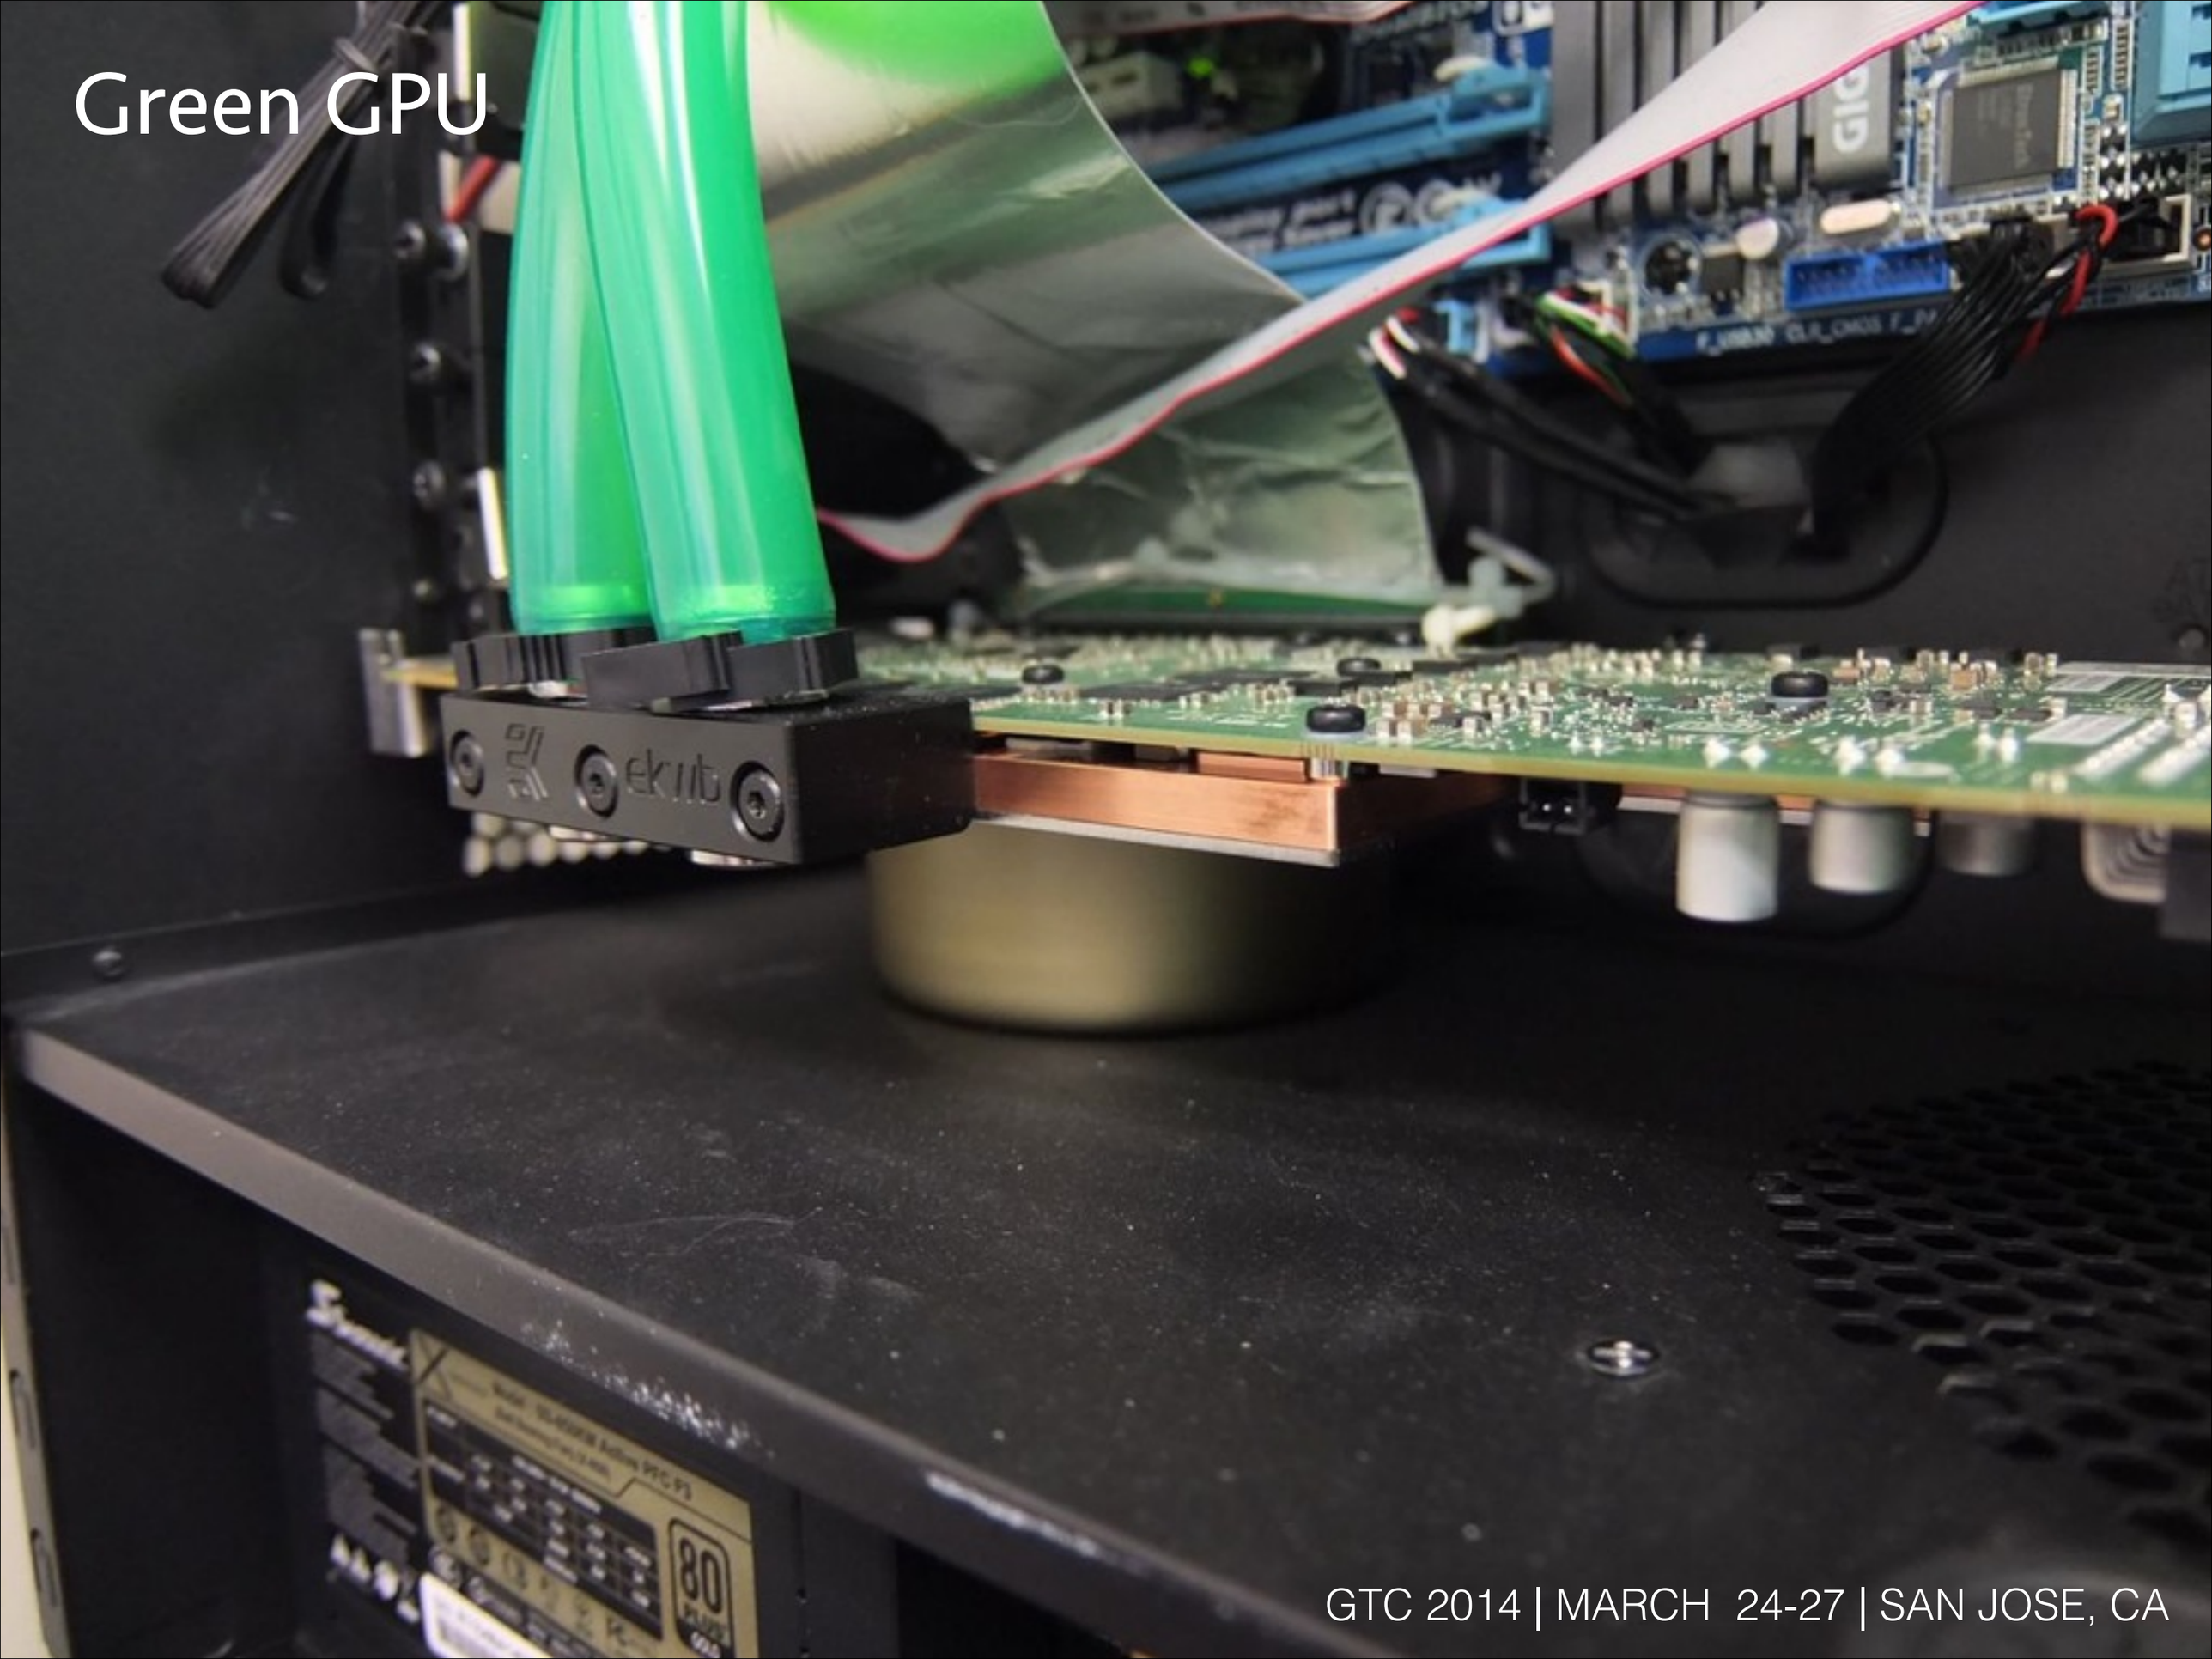

Green GPU

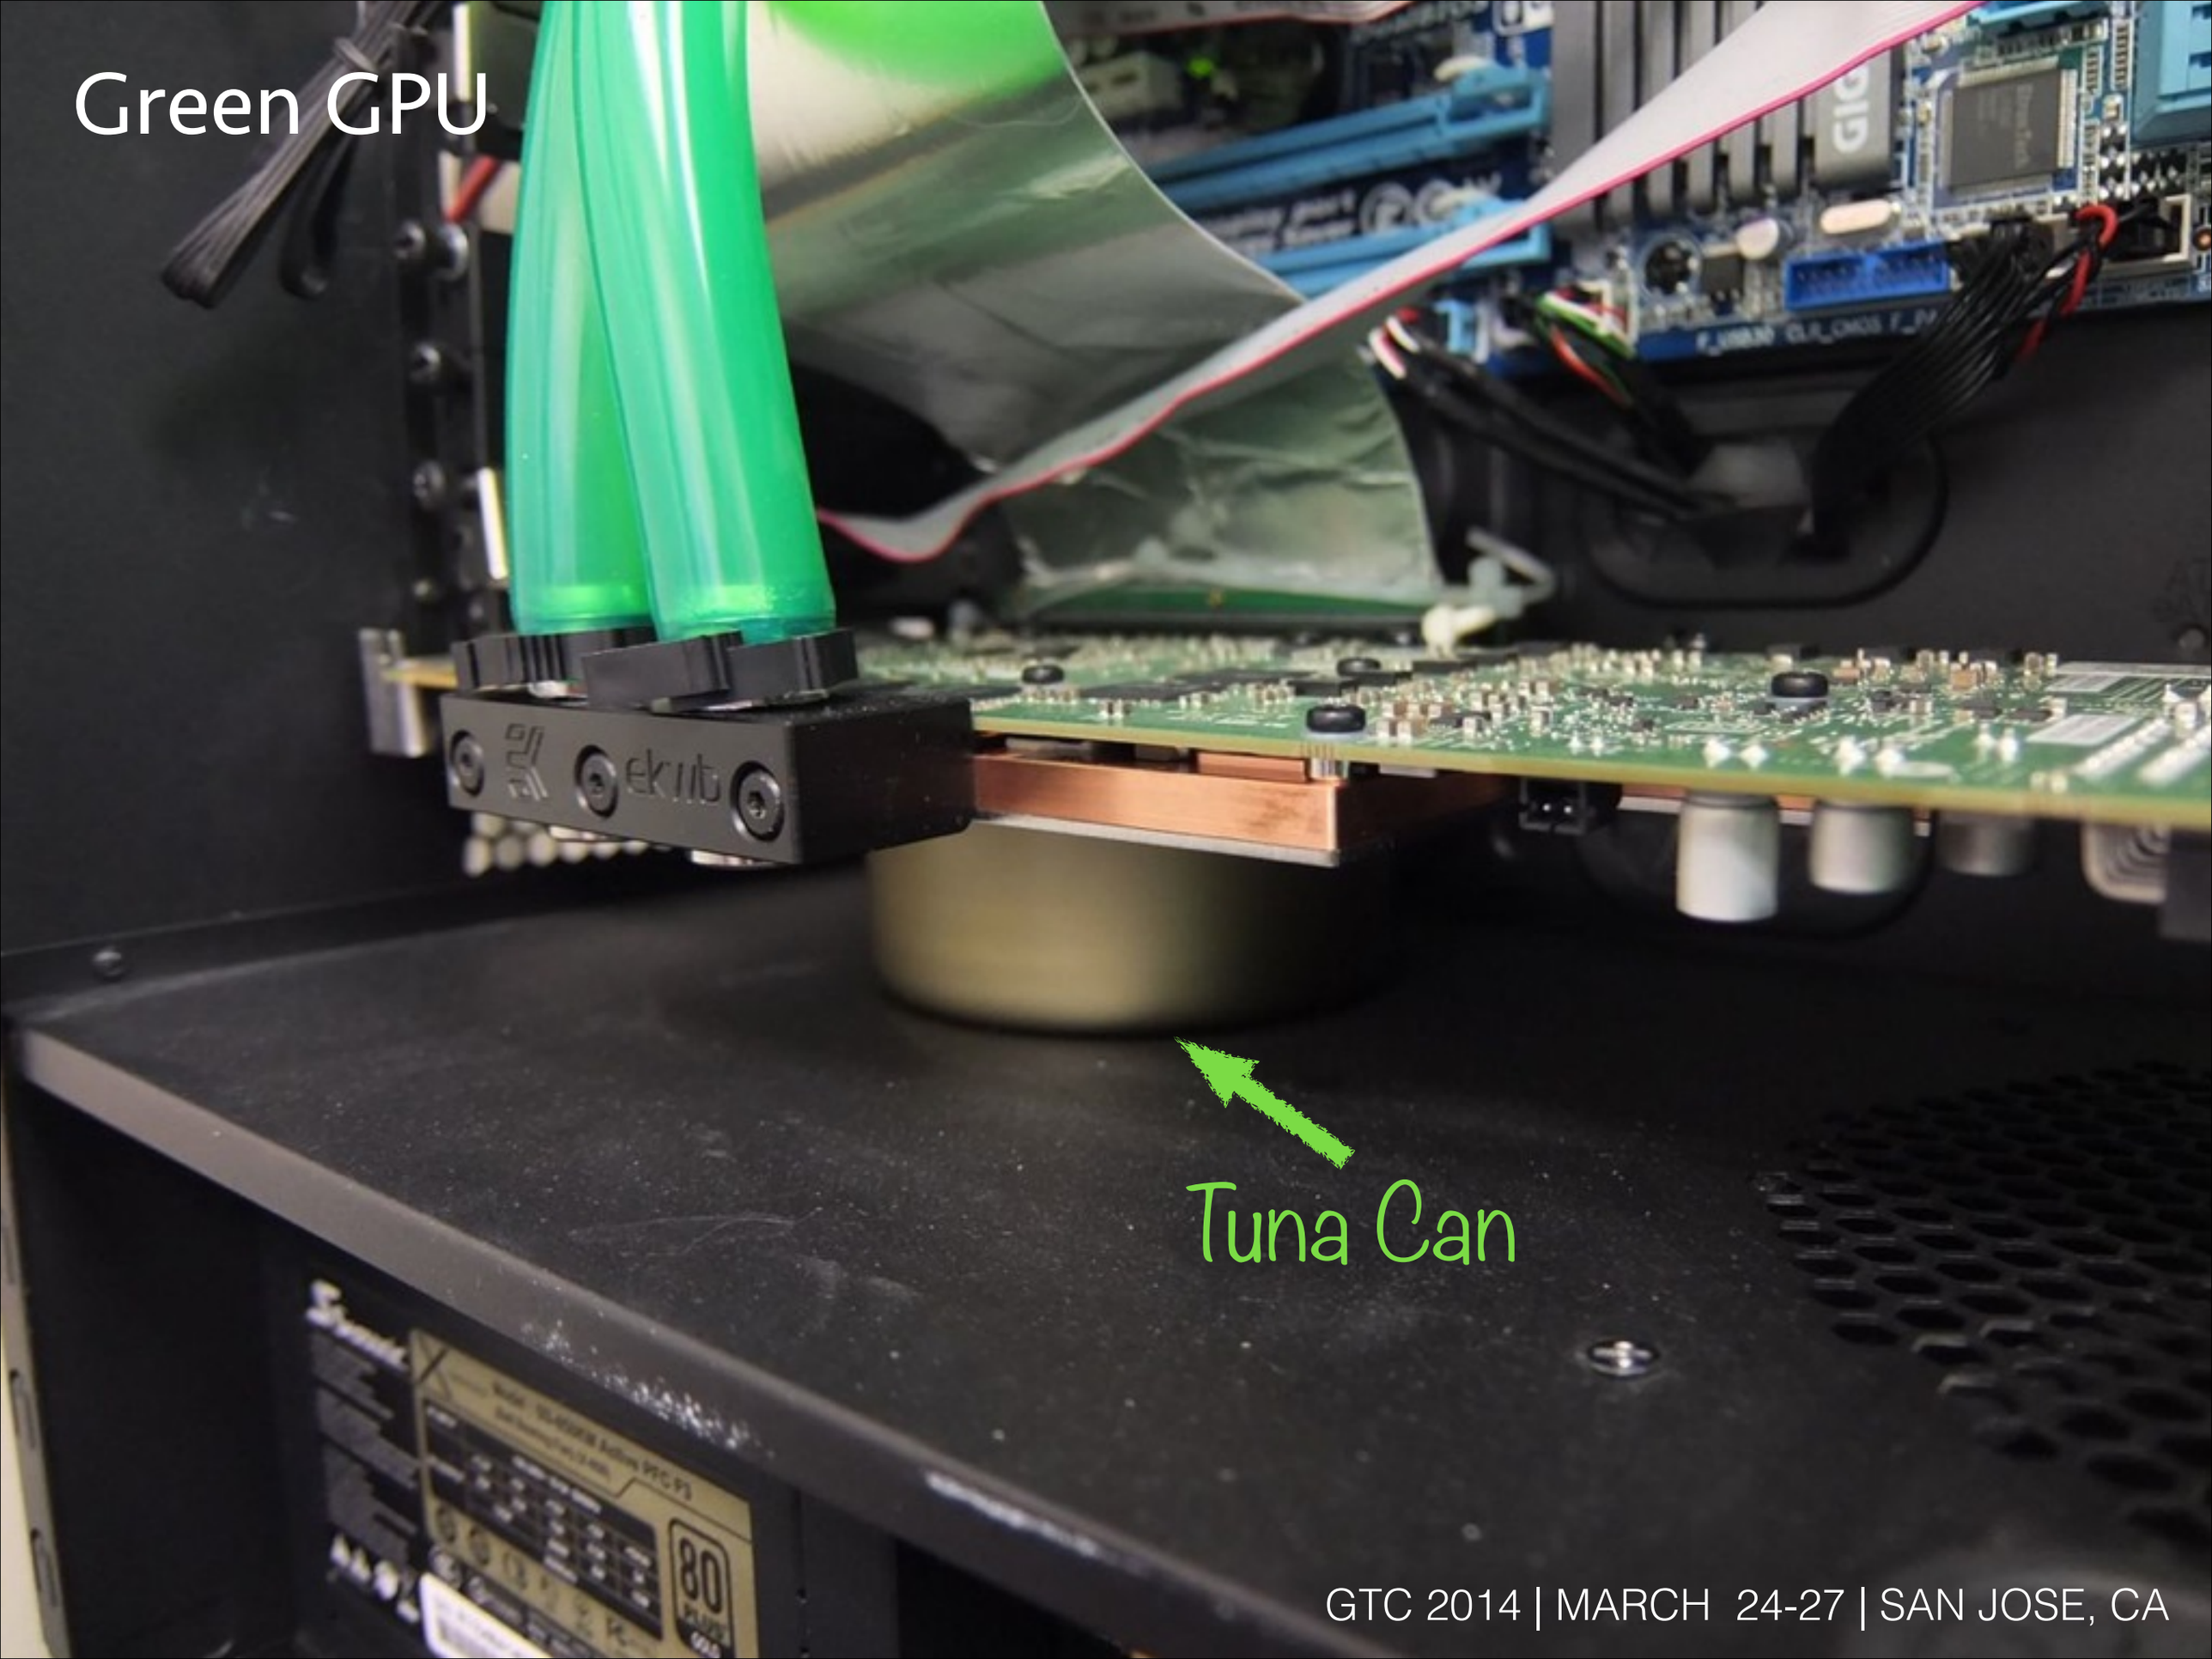

The experiment

- Run xGPU in benchmark mode
- Run SMI tool in loop
- Monitor GPU die temp and power

NVIDIA System Management Interface

```
dprice@greengpu: ~/xGPU/src
dprice@greengpu: ~/xGPU/src    dprice@greengpu: ~/xGPU/src
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K20m          Off  | 0000:01:00.0     Off |                    0 |
| N/A   31C    P0   151W / 225W |      340MB /  4799MB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                                          GPU Memory |
|  GPU       PID  Process name                                                Usage      |
|=============================================================================|
|    0      9159  ./cuda_correlator                                            325MB  |
+-----------------------------------------------------------------------------+
Thu Feb 13 15:23:00 2014
+------------------------------------------------------+
| NVIDIA-SMI 5.319.37   Driver Version: 319.37         |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K20m          Off  | 0000:01:00.0     Off |                    0 |
| N/A   31C    P0   151W / 225W |      340MB /  4799MB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                                          GPU Memory |
|  GPU       PID  Process name                                                Usage      |
|=============================================================================|
|    0      9159  ./cuda_correlator                                            325MB  |
+-----------------------------------------------------------------------------+
Thu Feb 13 15:23:05 2014
+------------------------------------------------------+
| NVIDIA-SMI 5.319.37   Driver Version: 319.37         |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K20m          Off  | 0000:01:00.0     Off |                    0 |
| N/A   31C    P0   151W / 225W |      340MB /  4799MB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                                          GPU Memory |
|  GPU       PID  Process name                                                Usage      |
|=============================================================================|
|    0      9159  ./cuda_correlator                                            325MB  |
+-----------------------------------------------------------------------------+
```

NVIDIA System Management Interface

```
dprice@greengpu: ~/xGPU/src
dprice@greengpu: ~/xGPU/src    dprice@greengpu: ~/xGPU/src
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K20m          Off  | 0000:01:00.0     Off |                    0 |
| N/A   31C    P0   151W / 225W |      340MB /  4799MB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                                          |
|  GPU       PID  Process name                                    GPU Memory |
|                                                                 Usage      |
|=============================================================================|
|    0      9159  ./cuda_correlator                                    325MB  |
+-----------------------------------------------------------------------------+
Thu Feb 13 15:23:00 2014
+------------------------------------------------------+
| NVIDIA-SMI 5.319.37   Driver Version: 319.37         |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K20m          Off  | 0000:01:00.0     Off |                    0 |
| N/A   31C    P0   151W / 225W |      340MB /  4799MB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                                          |
|  GPU       PID  Process name                                    GPU Memory |
|                                                                 Usage      |
|=============================================================================|
|    0      9159  ./cuda_correlator                                    325MB  |
+-----------------------------------------------------------------------------+
Thu Feb 13 15:23:05 2014
+------------------------------------------------------+
| NVIDIA-SMI 5.319.37   Driver Version: 319.37         |
|-------------------------------+----------------------+----------------------+
| GPU  Name        Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp  Perf  Pwr:Usage/Cap|         Memory-Usage | GPU-Util  Compute M. |
|===============================+======================+======================|
|   0  Tesla K20m          Off  | 0000:01:00.0     Off |                    0 |
| N/A   31C    P0   151W / 225W |      340MB /  4799MB |     99%      Default |
+-------------------------------+----------------------+----------------------+

+-----------------------------------------------------------------------------+
| Compute processes:                                                          |
|  GPU       PID  Process name                                    GPU Memory |
|                                                                 Usage      |
|=============================================================================|
|    0      9159  ./cuda_correlator                                    325MB  |
+-----------------------------------------------------------------------------+
```

NVIDIA SMI — GPU Boost on K20m

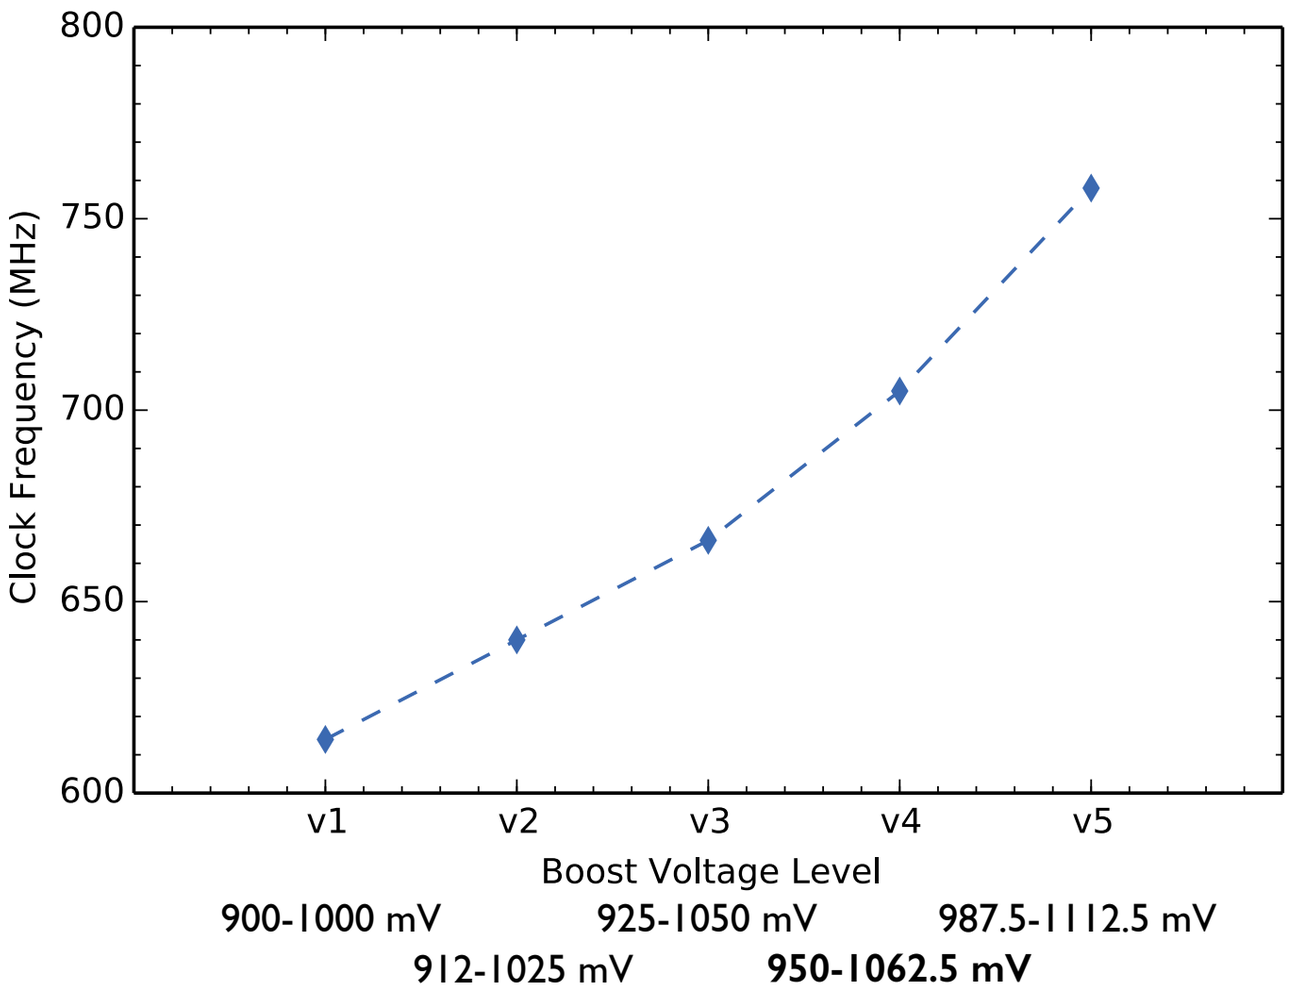

- Can control *"boost clock"* through nvidia-smi:

```
nvidia-smi –q –d SUPPORTED_CLOCKS
nvidia-smi –ac <MEM clock, Graphics clock>
```

- Changing boost clock changes voltage levels (more on this later)
- ***Undervolting*** uses less power, but becomes unstable at higher temperatures
- ***Overvolting*** uses more power, but is more stable at higher temperatures. There is a big "overclocker" community that exploits this.
- ***V4*** is the K20m default. Note max power (TDP) is *225W*.

How to *really* void your warranty

11-23-2012, 06:23 PM #1

CrazyNutz

Xtreme Addict

Join Date: Dec 2007
Location: Texas
Posts: 1,785

KGB - Kepler BIOS Editor/Unlocker

KGB (Kepler Golden BIOS) Is an Nvidia Kepler BIOS Editor/Unlocker program I have written. Enjoy ☺

KGB supports: GTX690, GTX680, GTX670, GTX660Ti, GTX660OEM and GTX660.

Instructions: KGB is a console application, you must use it from a command prompt in windows. Start menu then type cmd, then cd to the folder containing kgb.
You will need to extract your bios from your card with this version of GPU-Z http://www.techpowerup.com/forums/at...2&d=1344772659
When you are done modding your bios you will need to flash it back to your card with nvflash.

Usage Example 1: kgb.exe your_bios.rom
This will display info from the bios.

Usage Example 2: kgb.exe your_bios.rom unlock
This will **unlock** your BIOS. Make sure you have a backup of your original BIOS file!!!!!!!!!!!!!!!!!!!

http://www.xtremesystems.org/forums/showthread.php?284014-KGB-Kepler-BIOS-Editor-Unlocker

Don’t try this at home

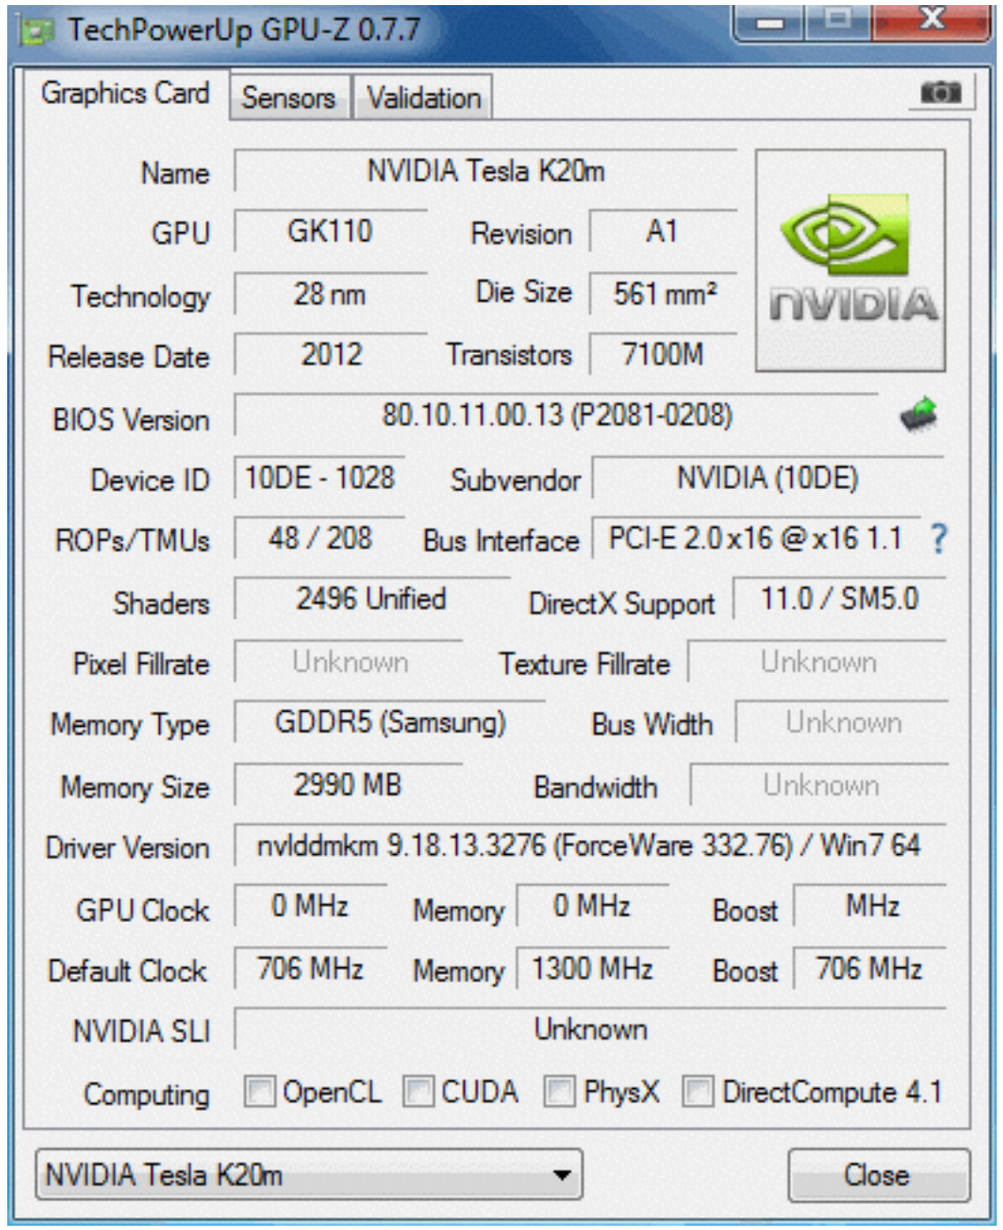

GPU-Z
(www.techpowerup.com/gpuz/)

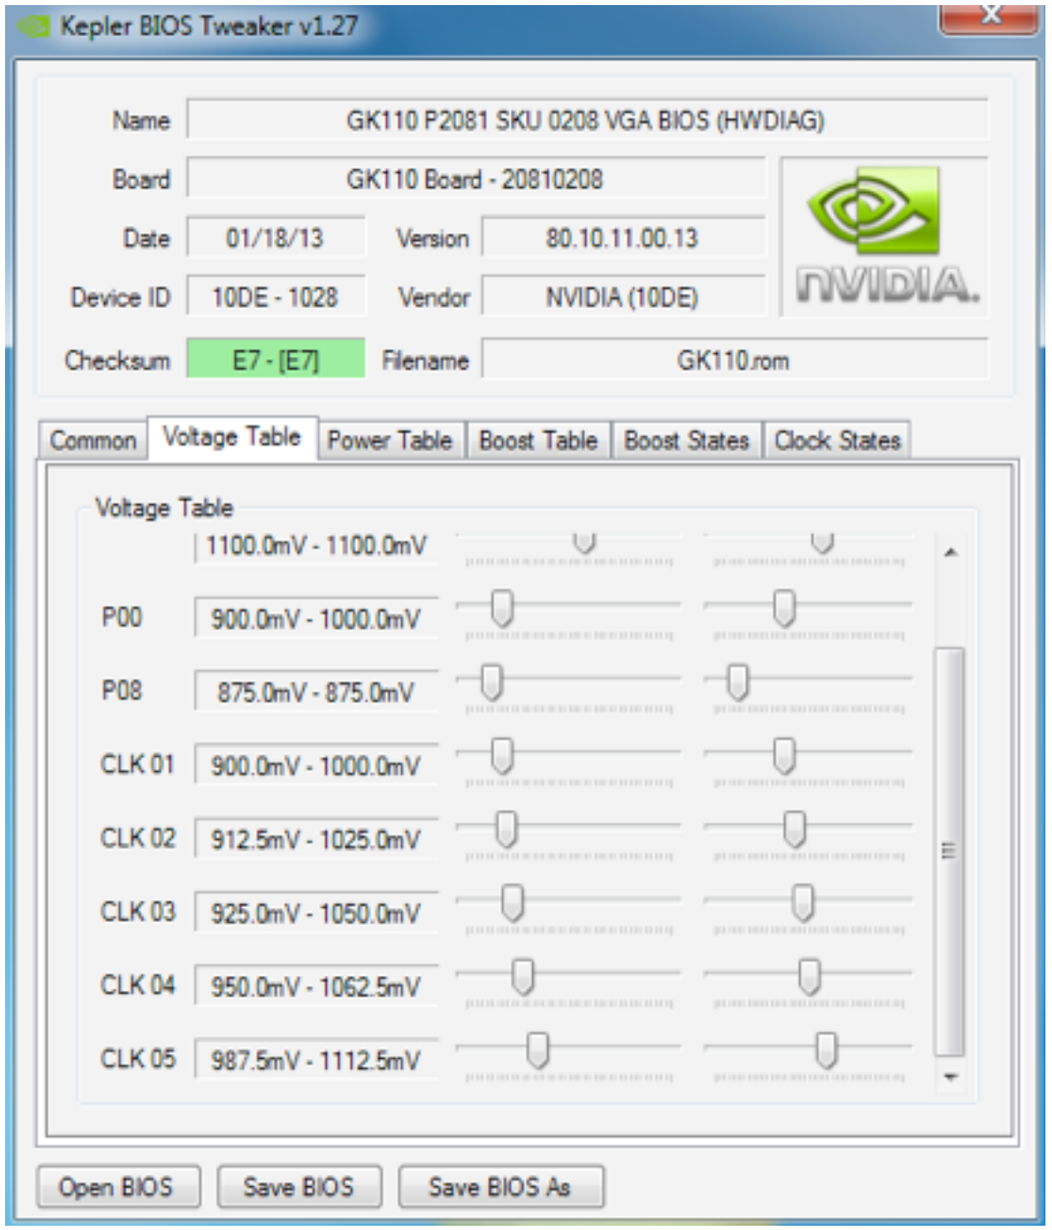

Kepler Bios Tweaker v1.27
(overclock.net forums)

Overclocking and xGPU performace

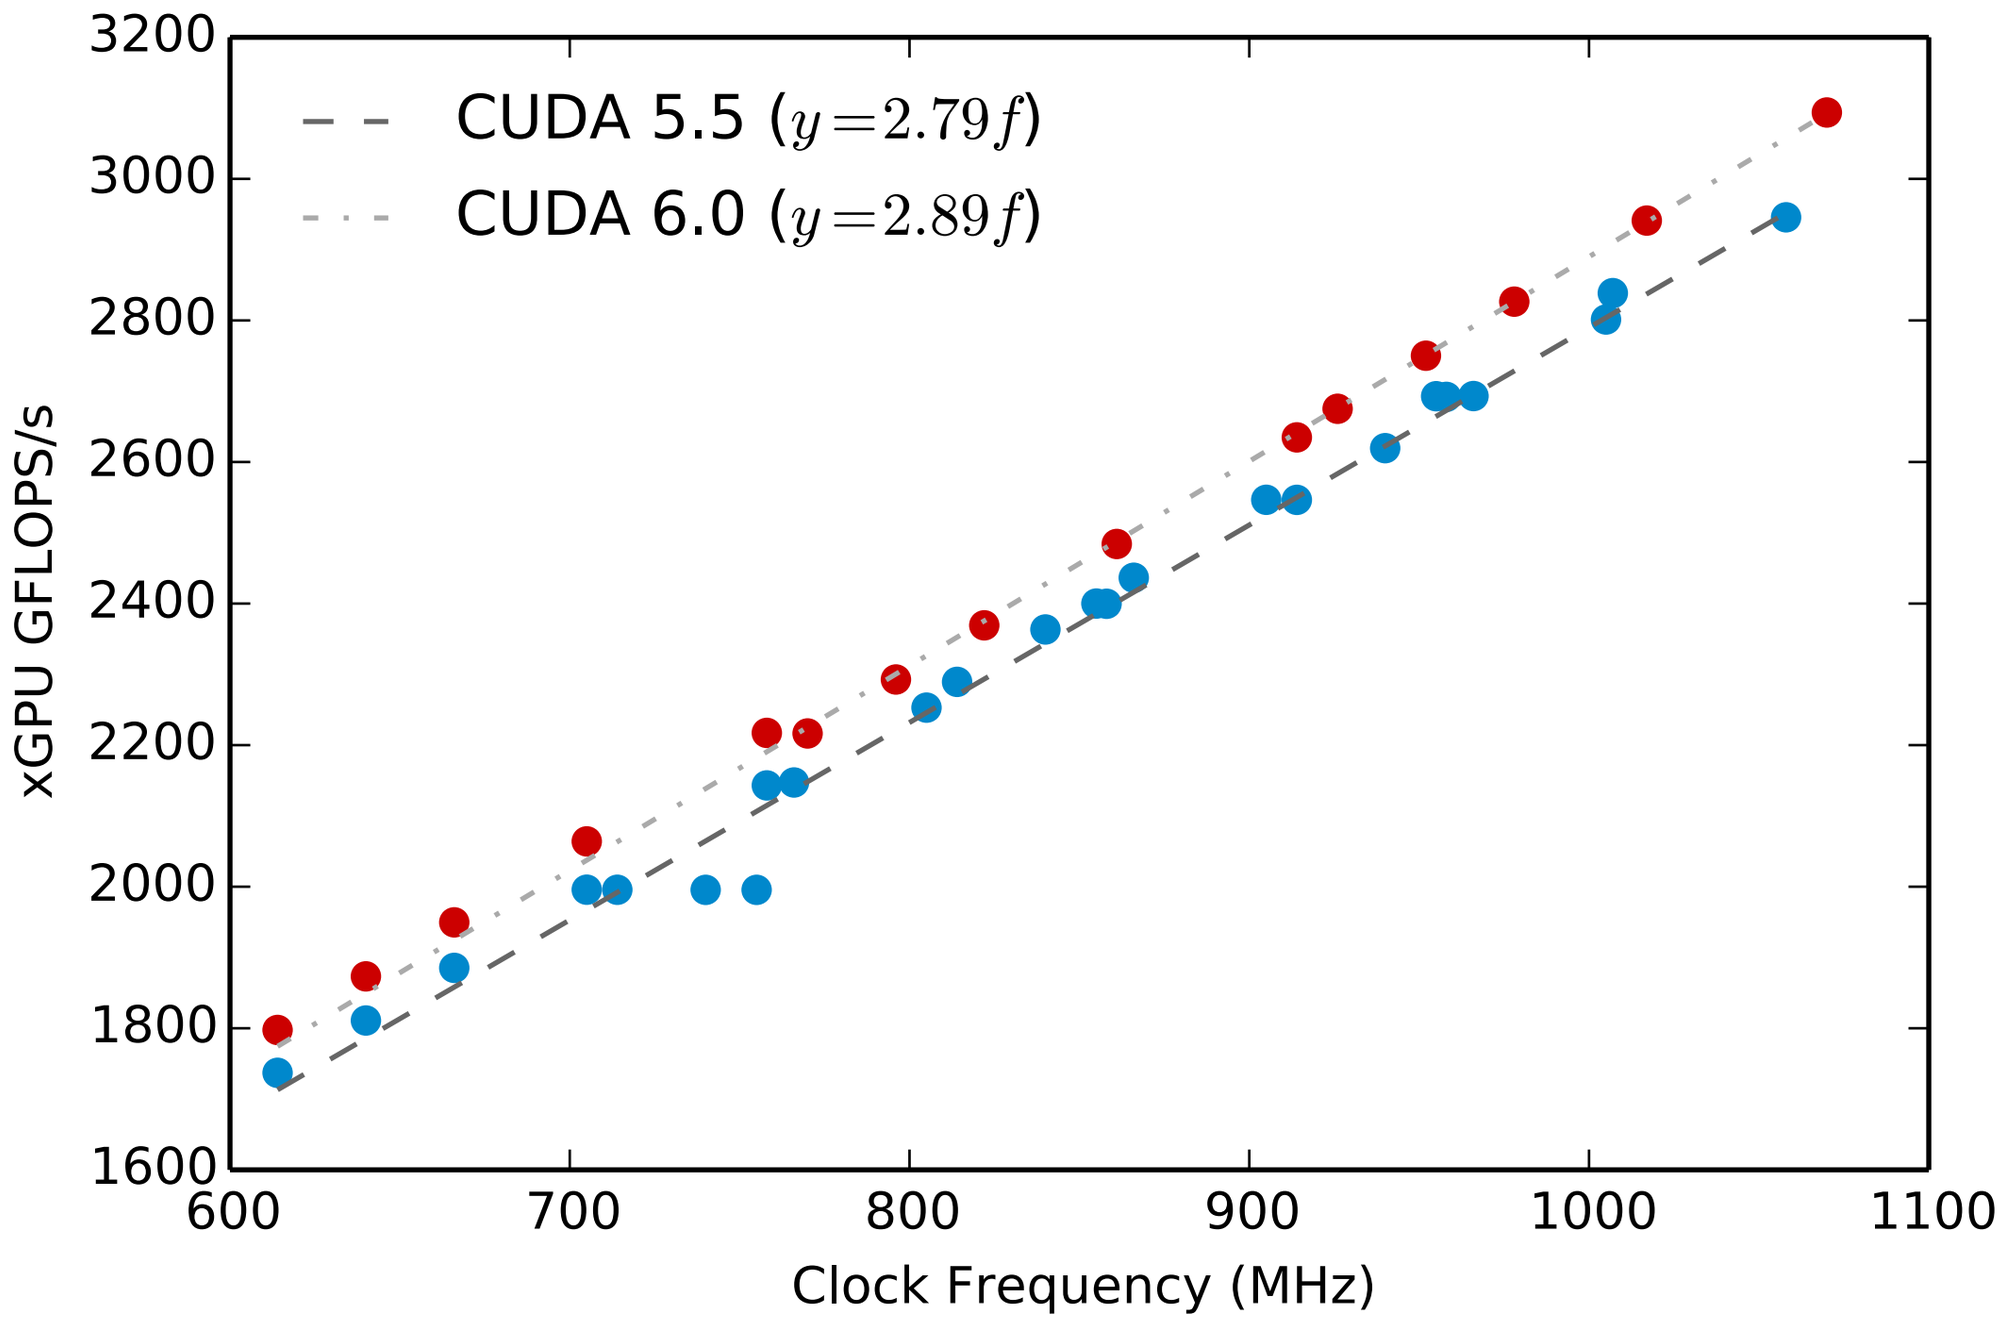

Drifting temperature, drifting power

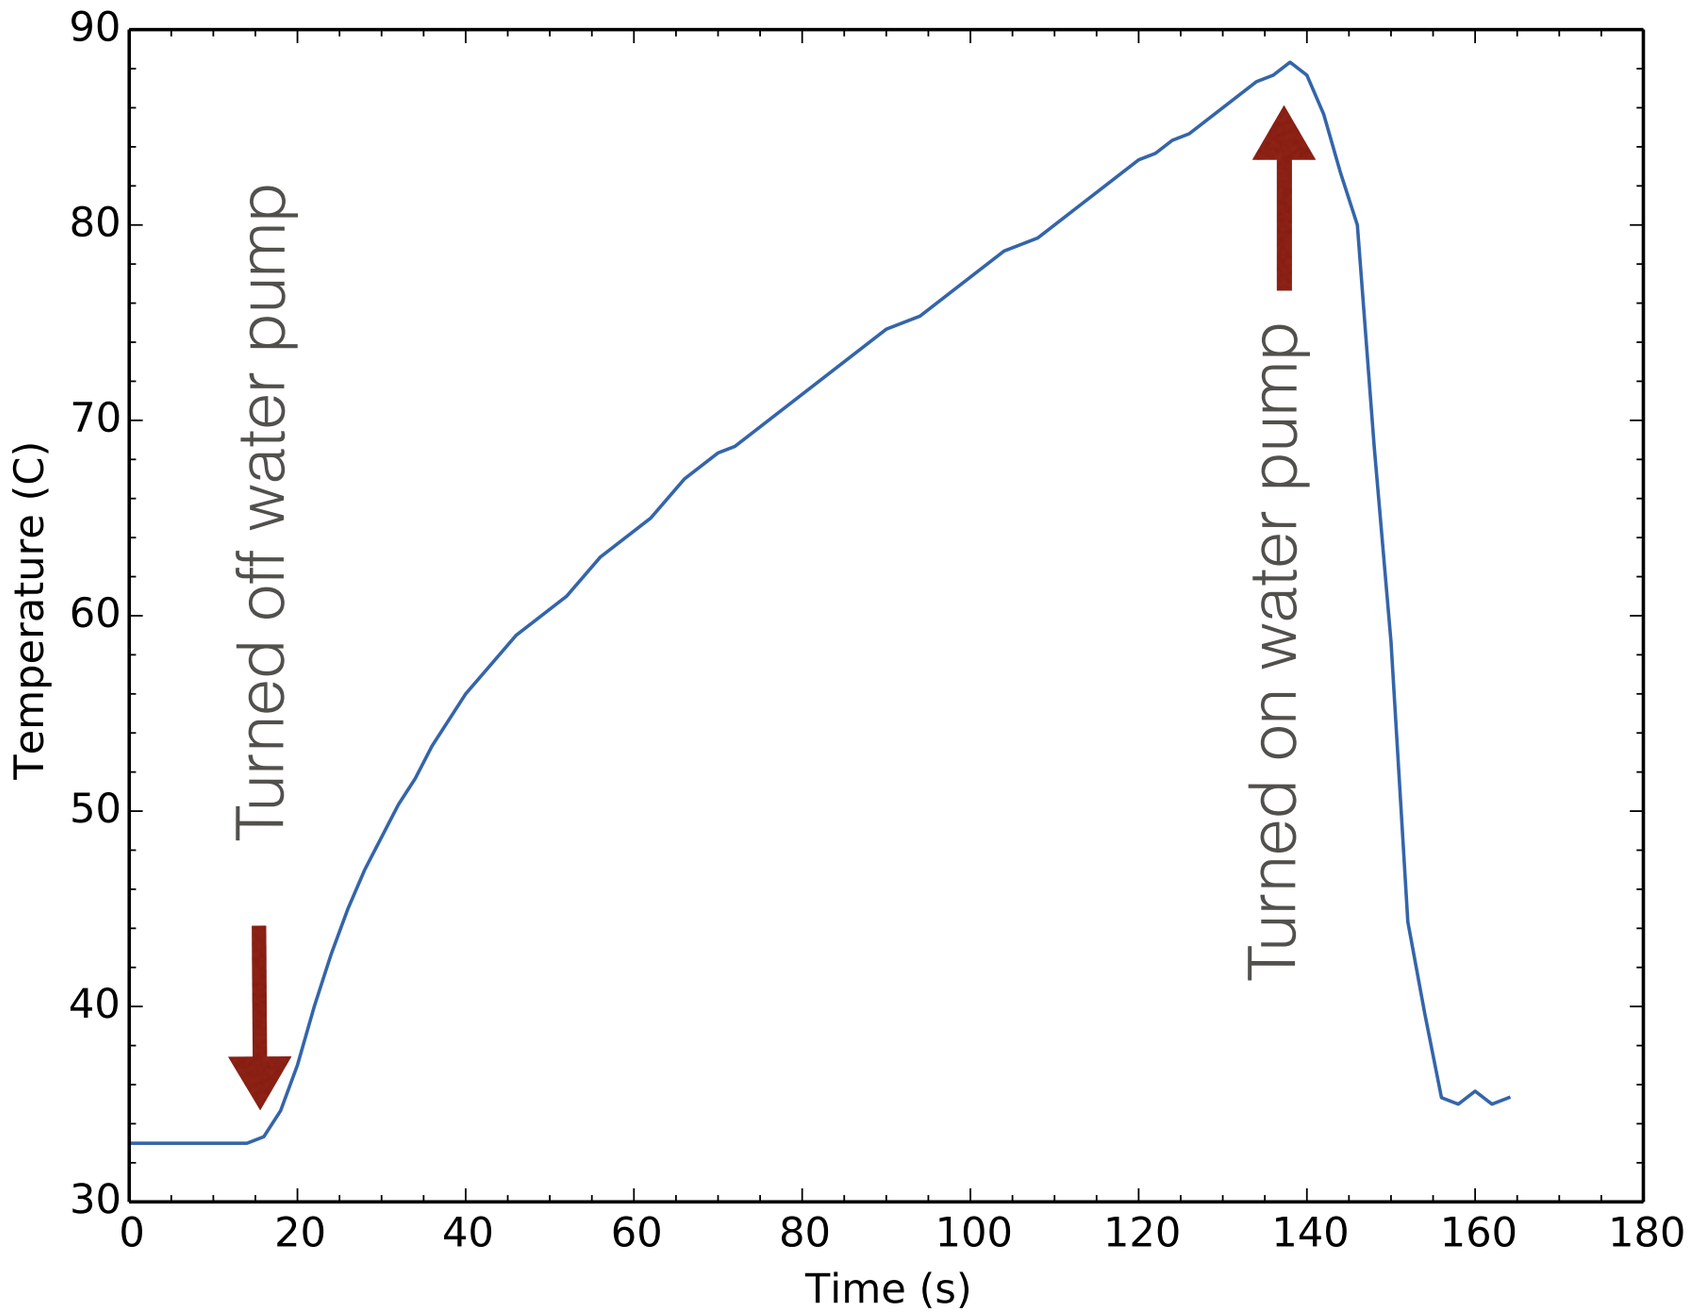

Results: clock & voltage control

Measured power vs temperature

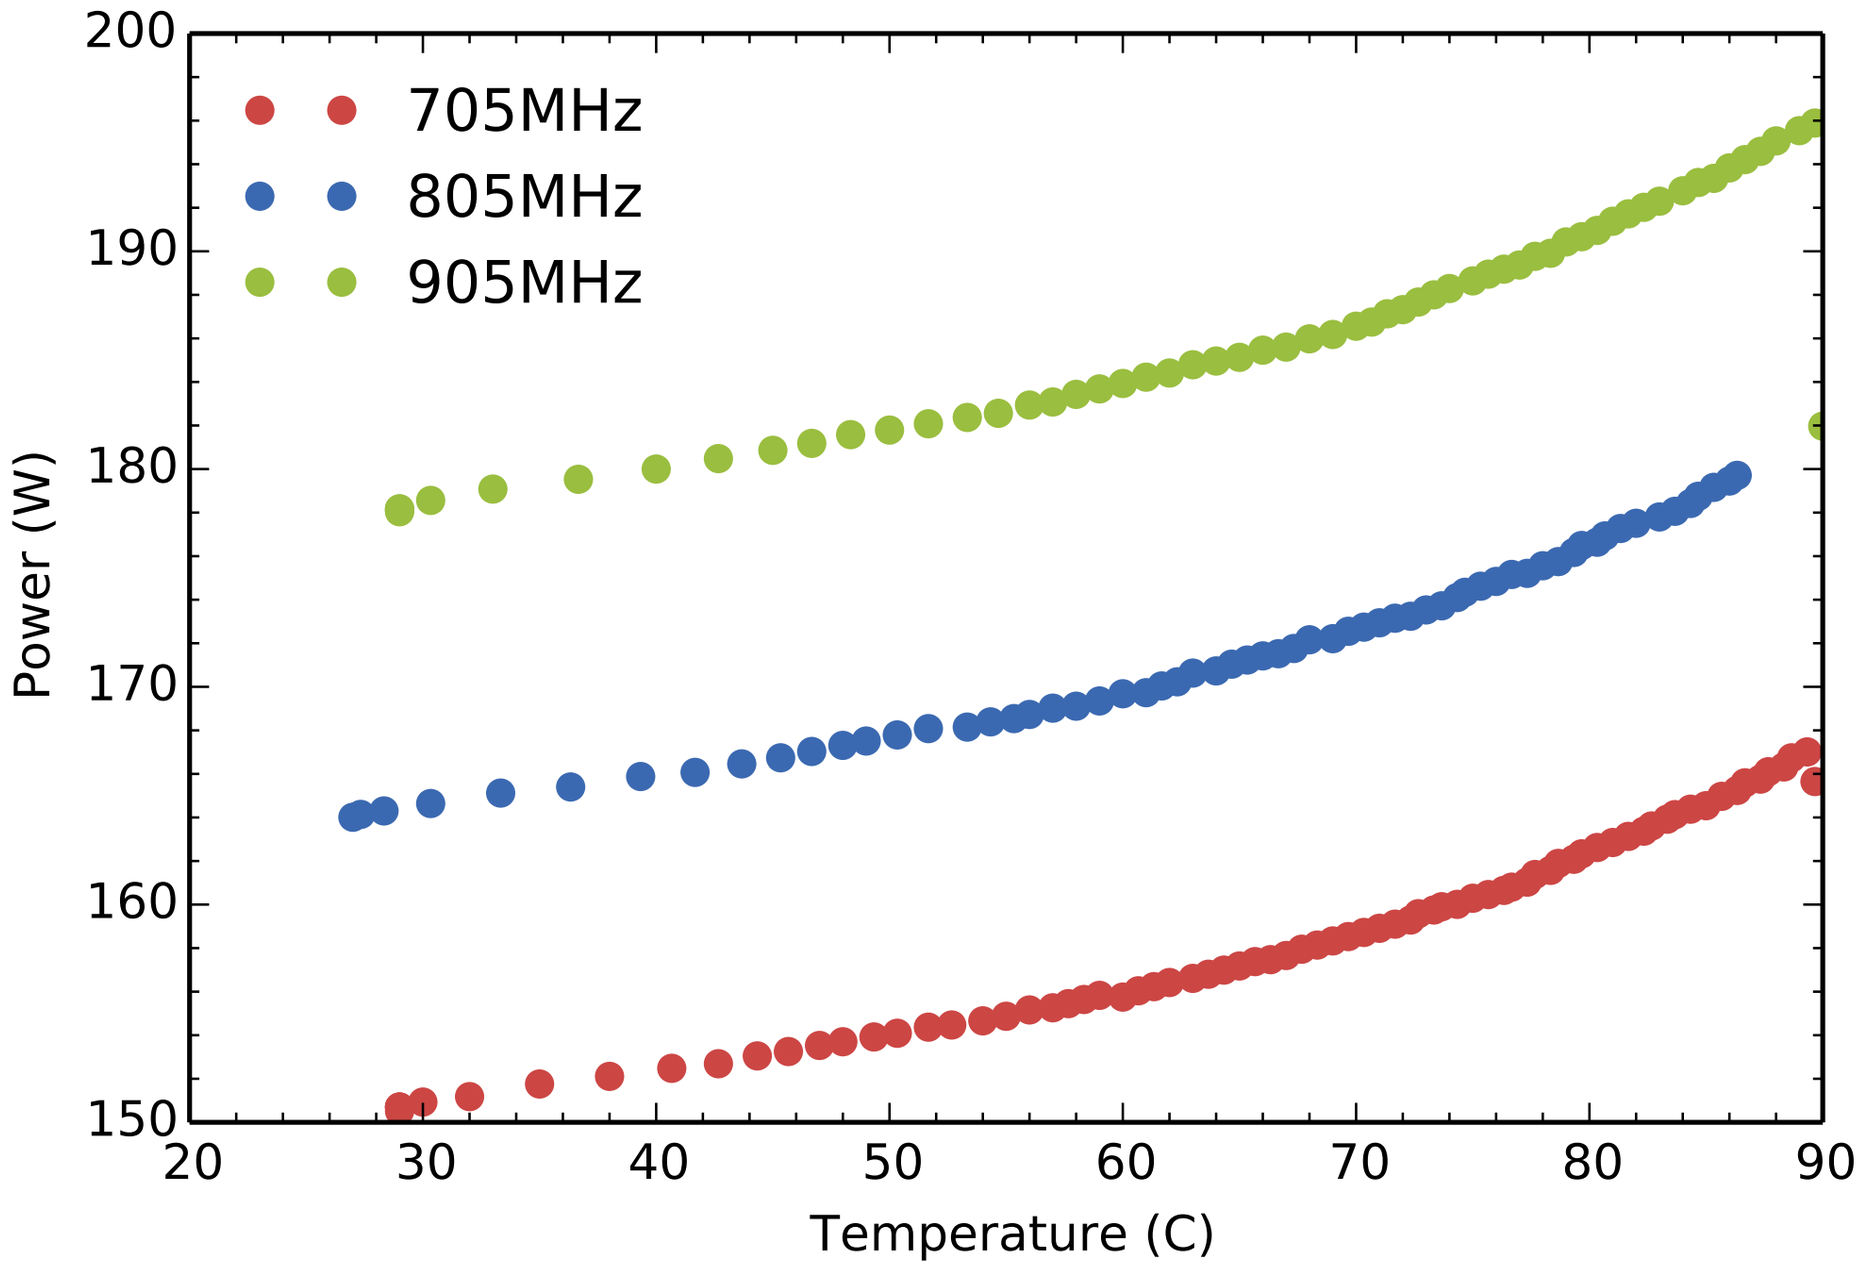

V = V4

Measured power vs temperature

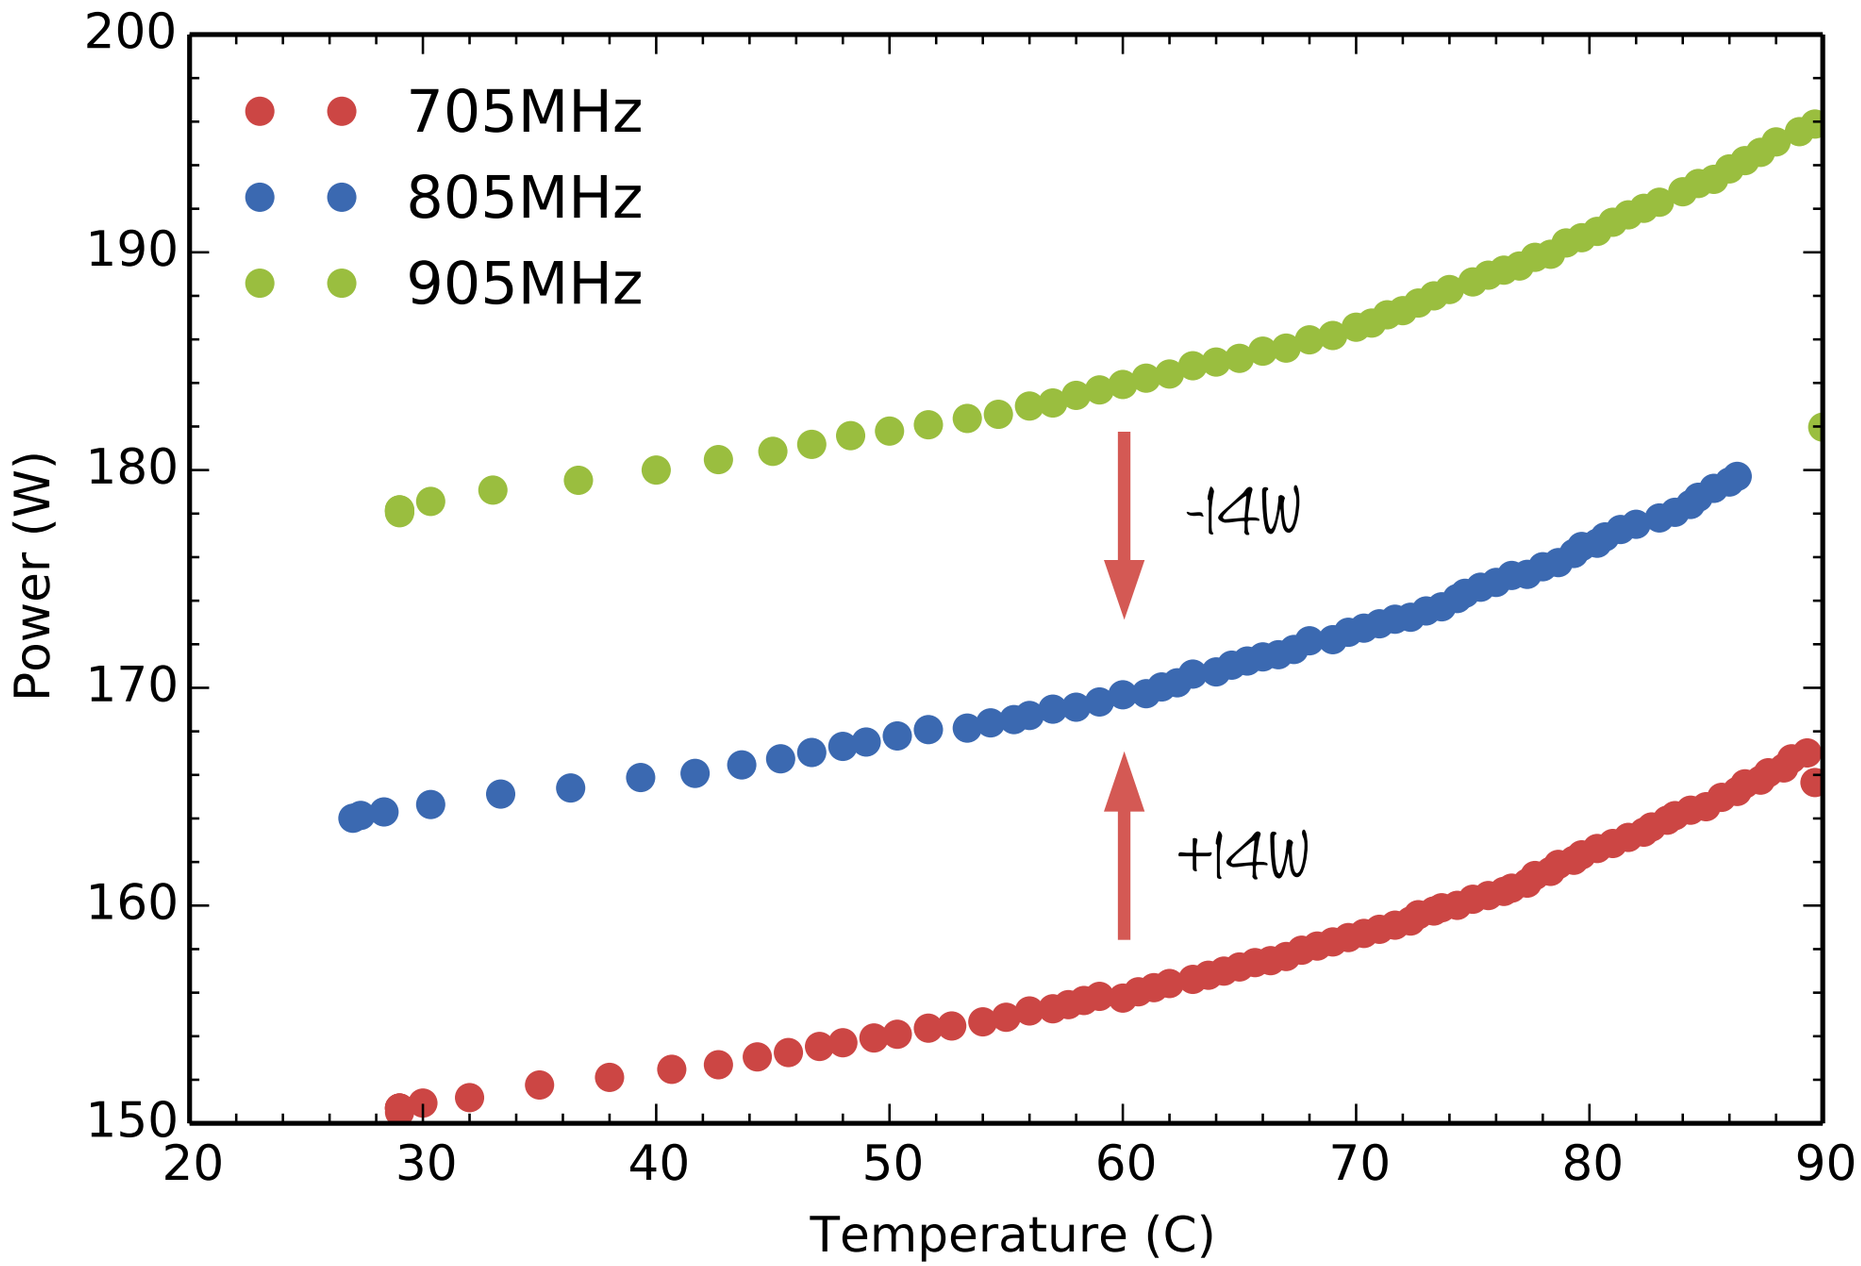

V = V4

Measured power vs temperature

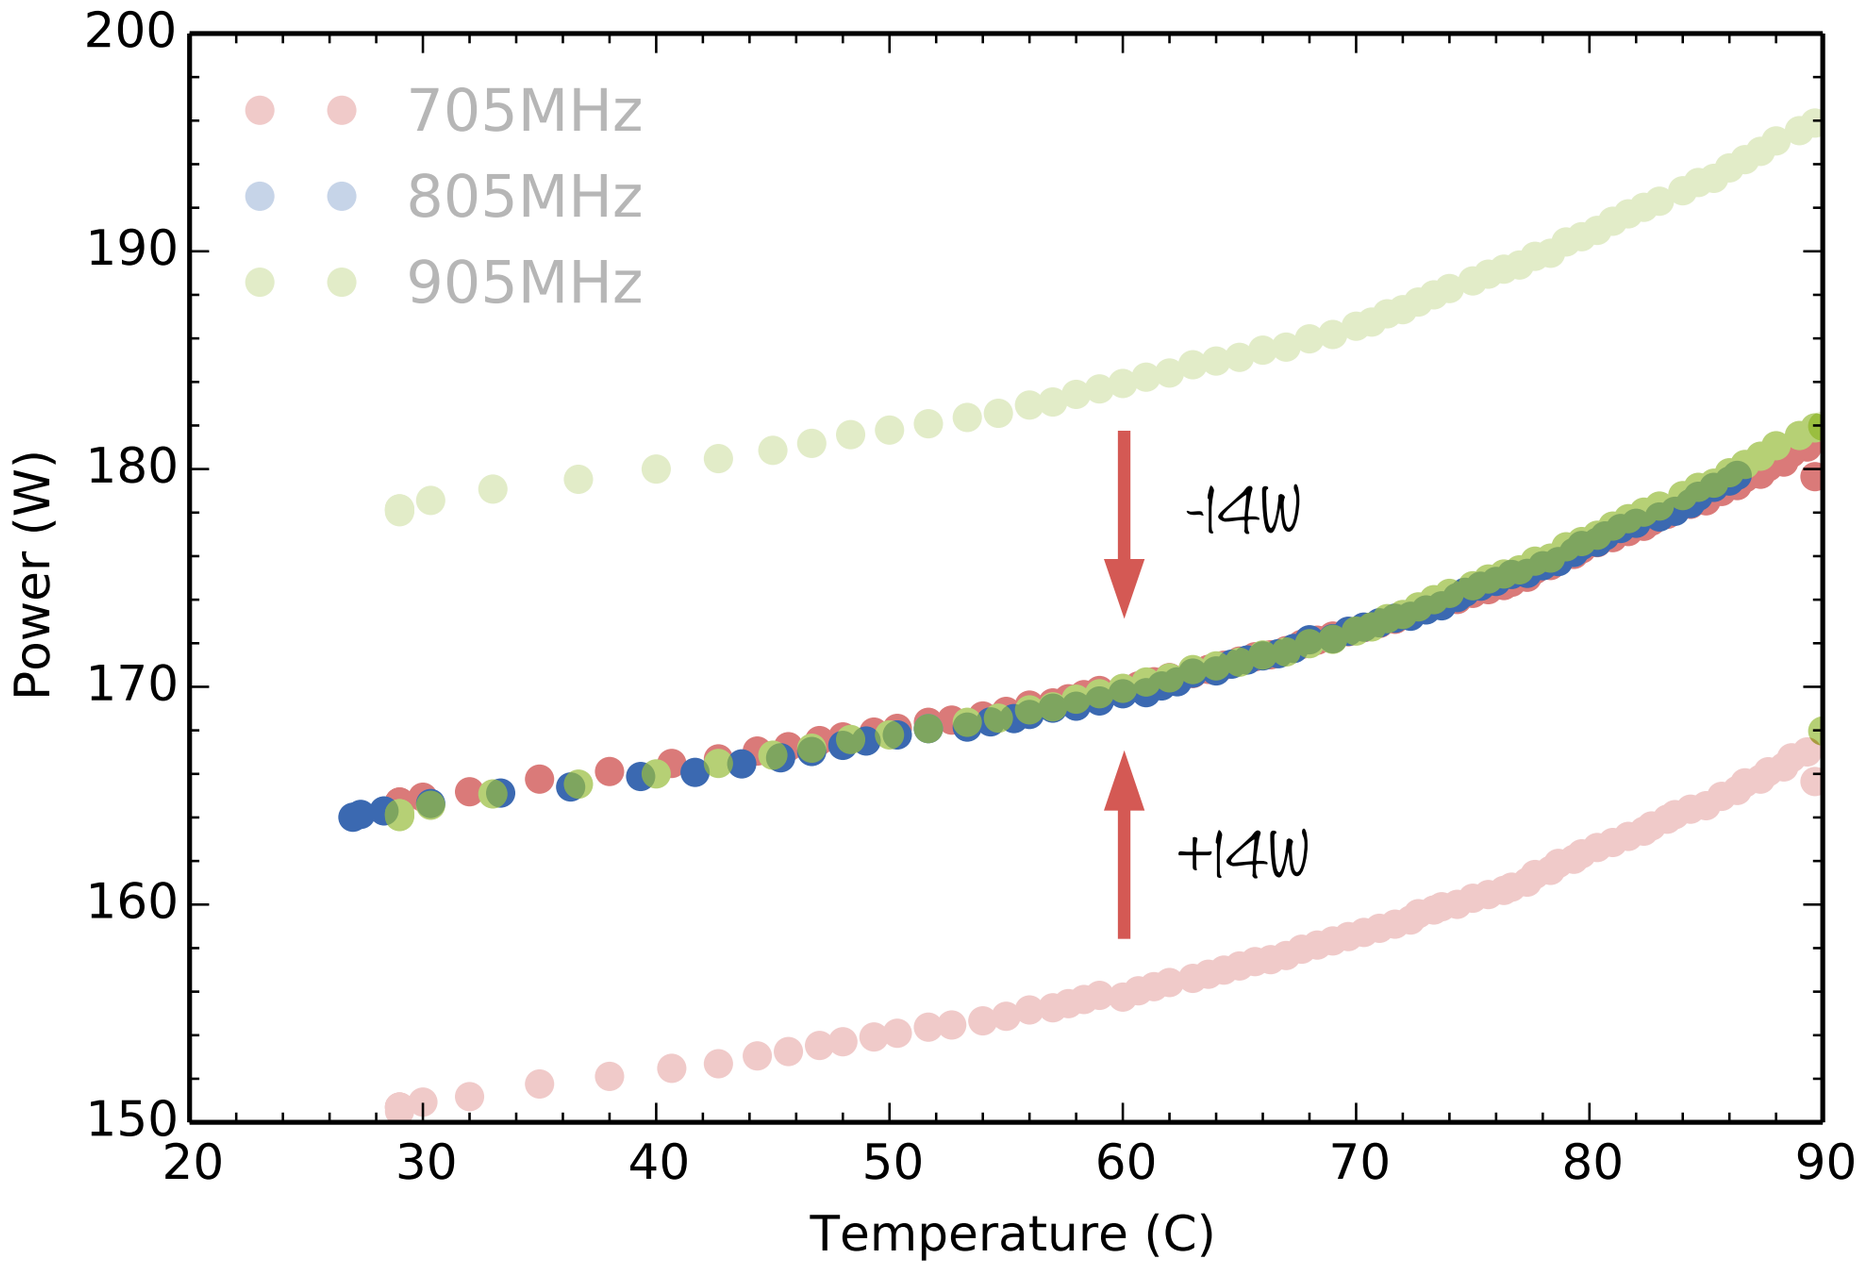

V = V4

Measured power vs temperature

$$P = P_{dynamic} + P_{static}$$

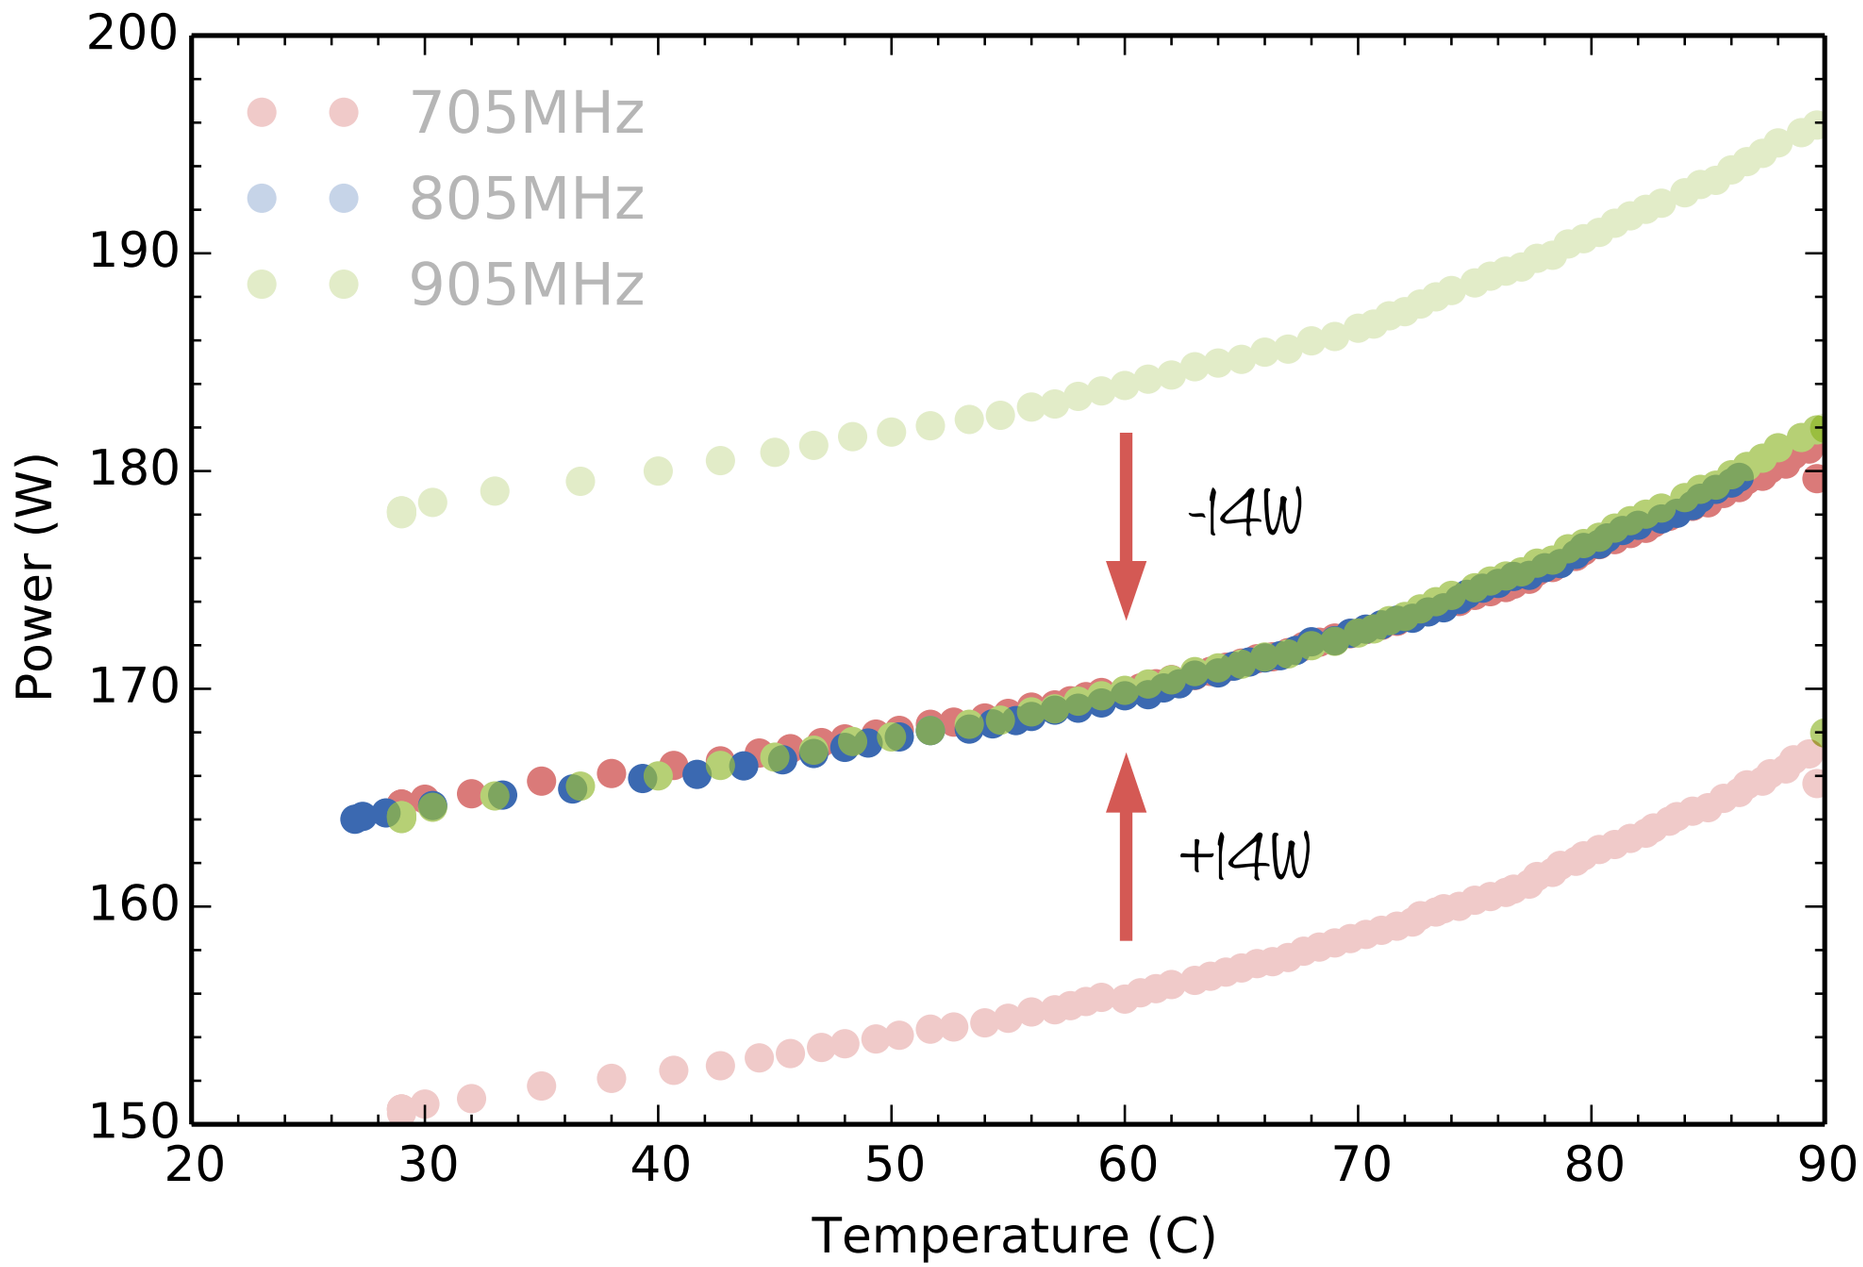

V = V4

How temperature affects performance

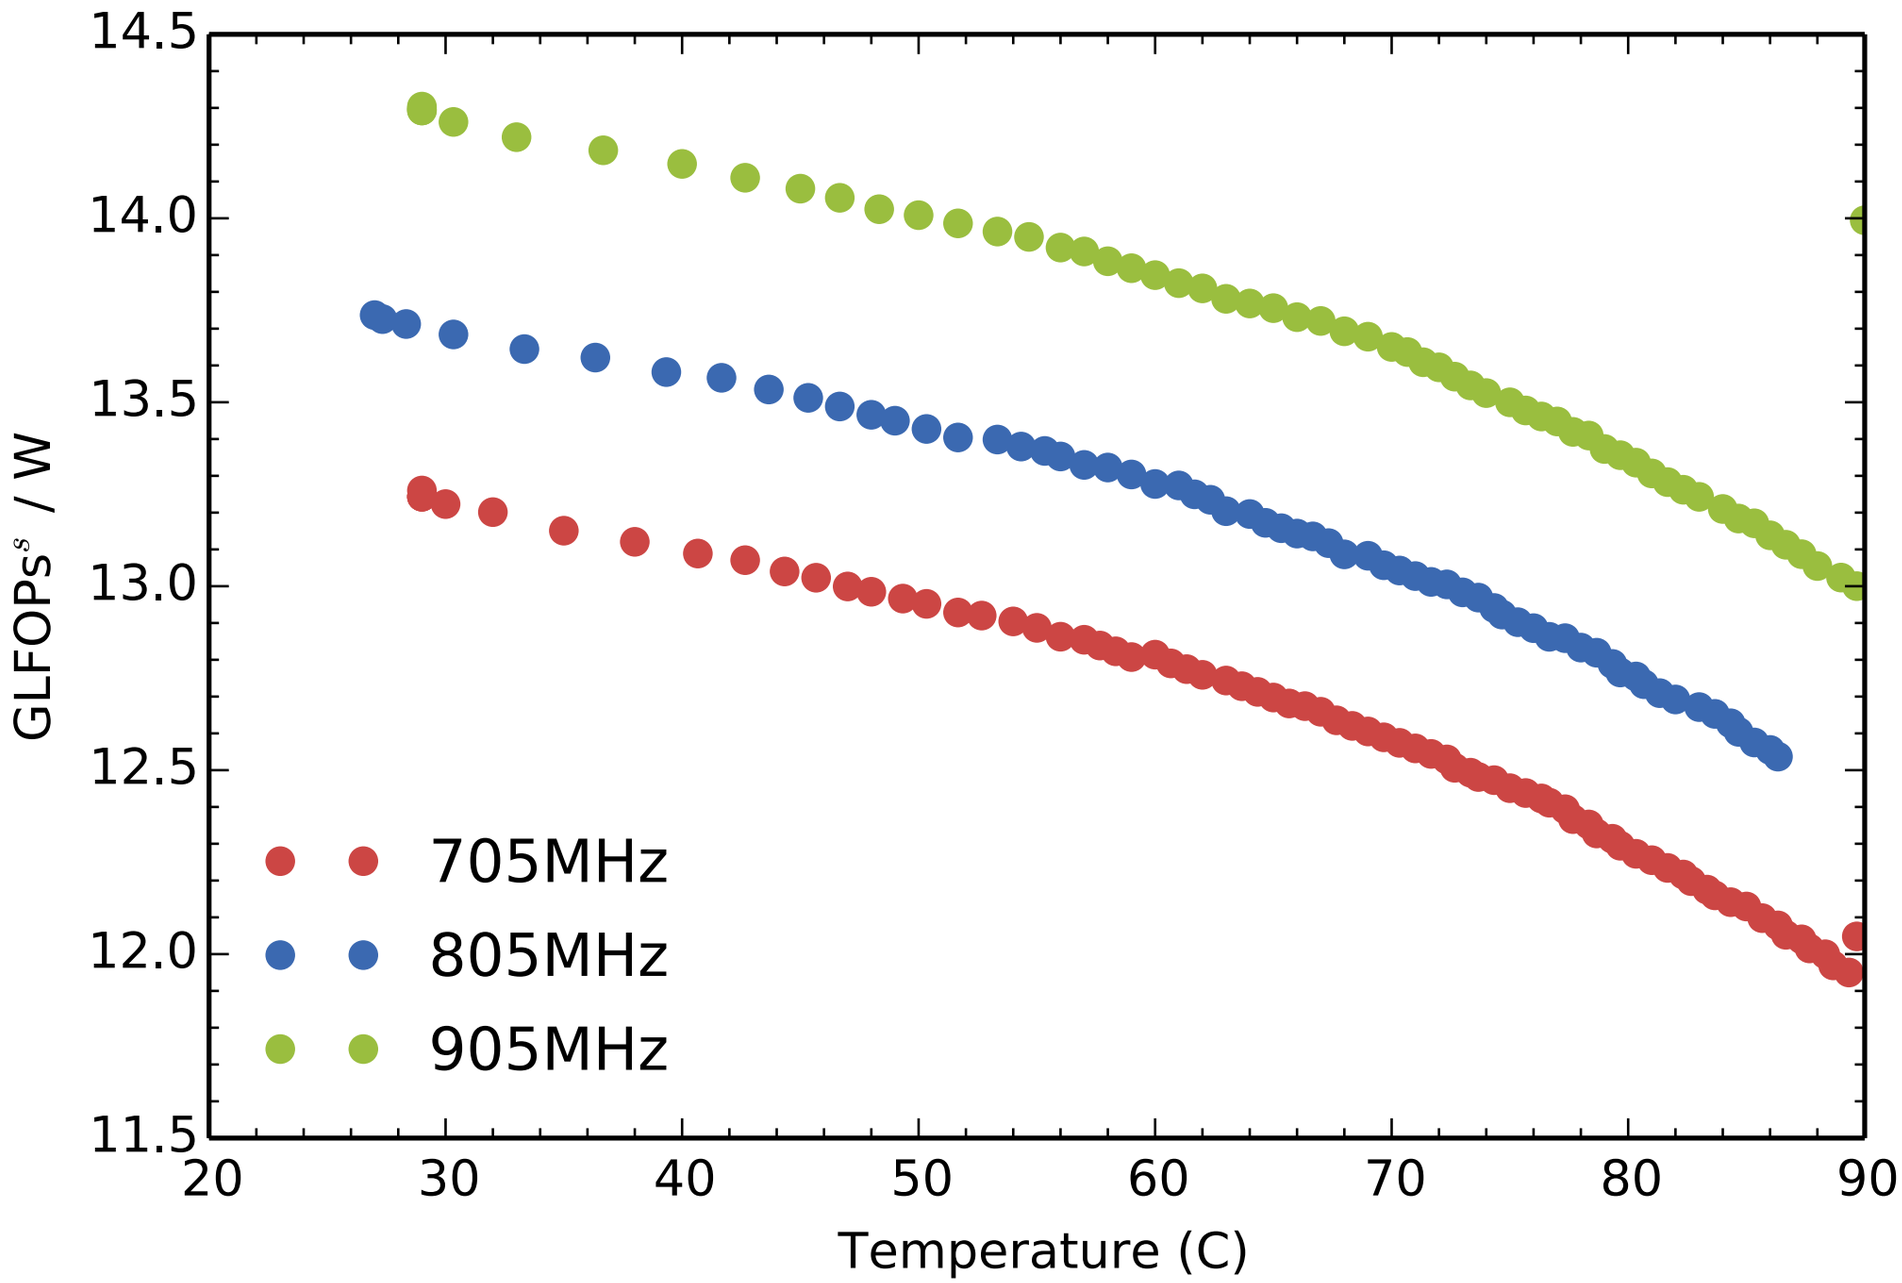

V = V4

How voltage affects performance

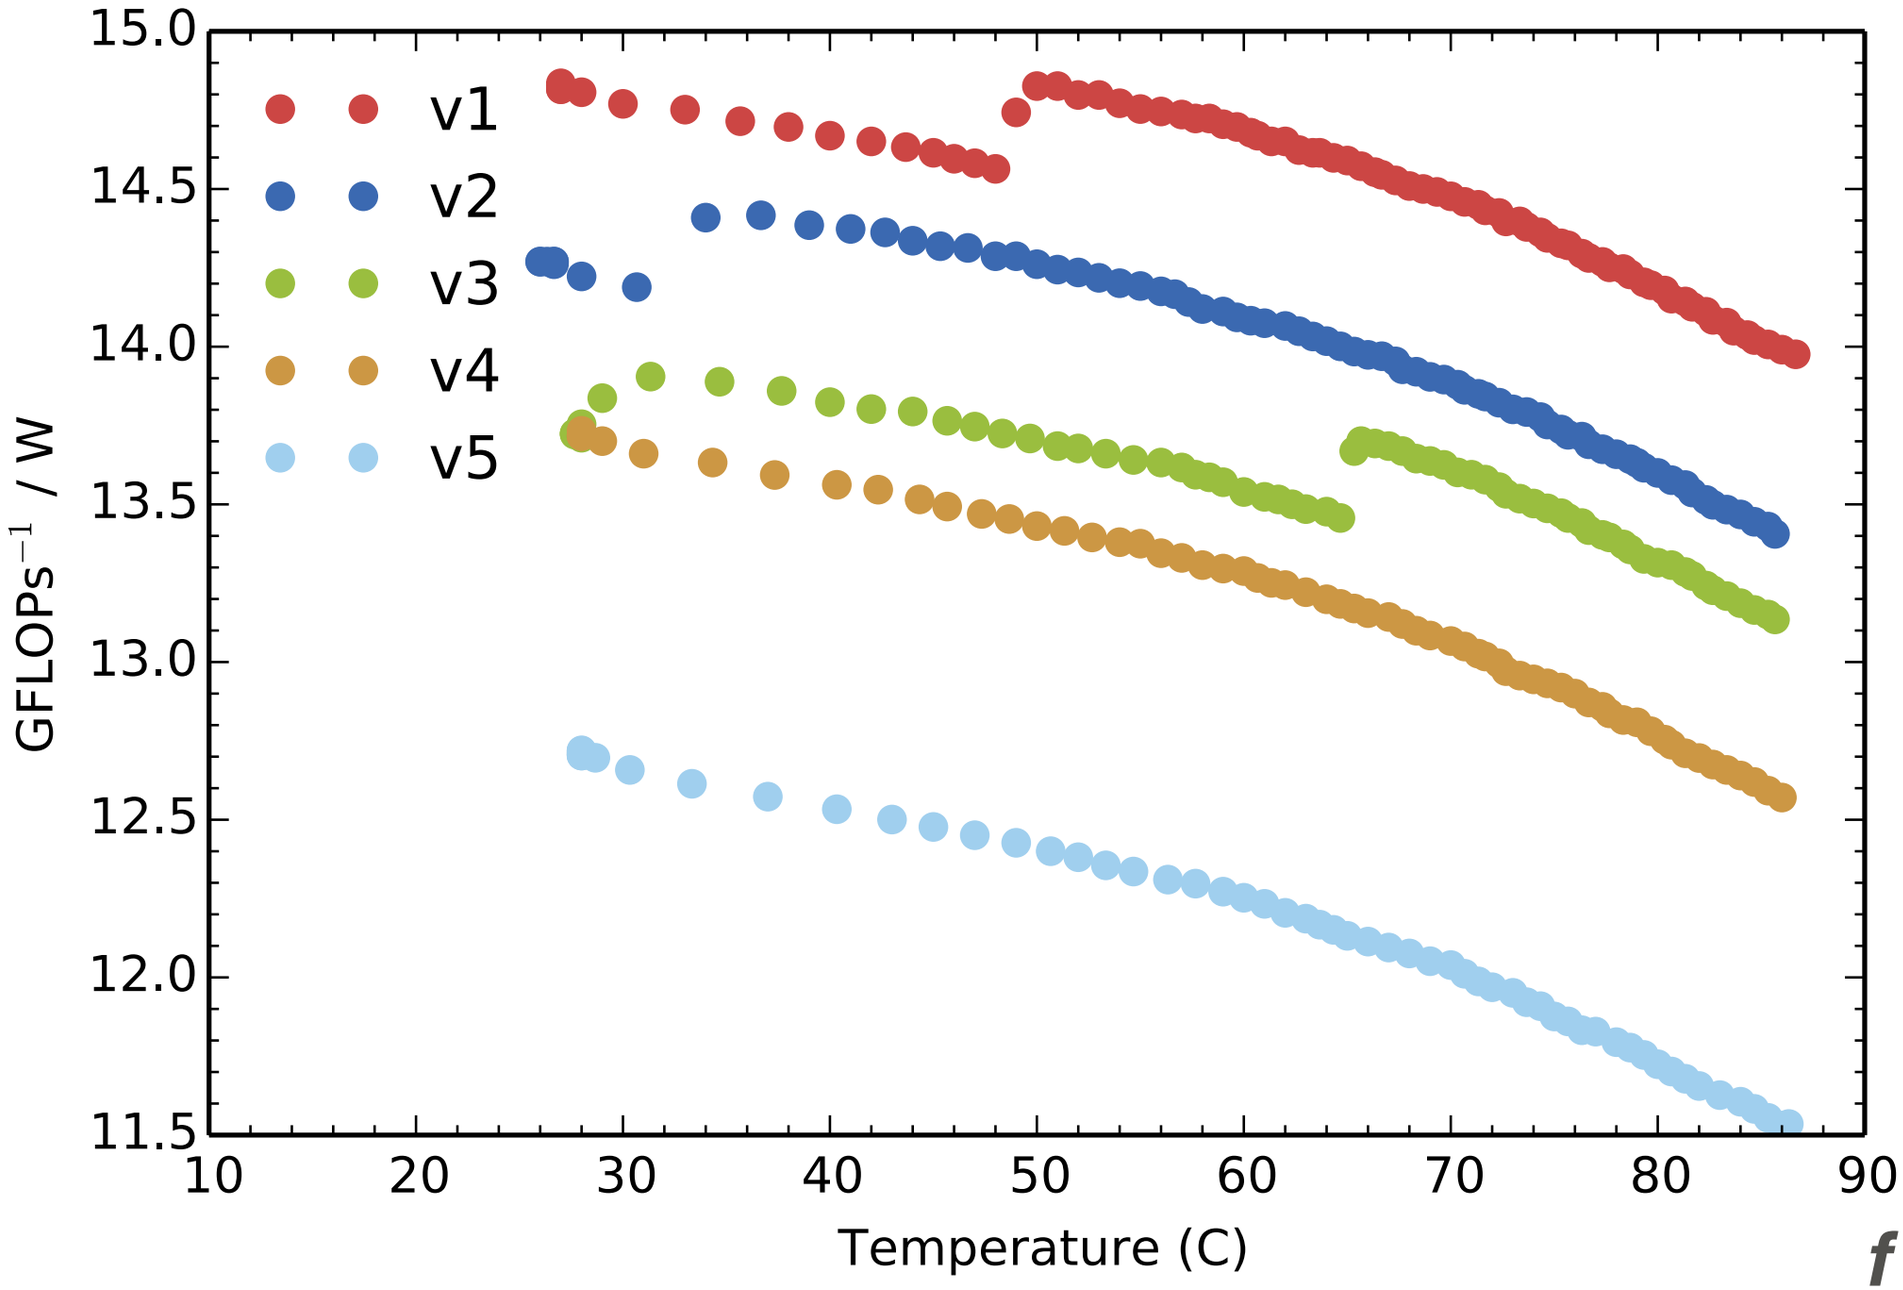

f = 800 MHz

Undervolted, different clock frequencies

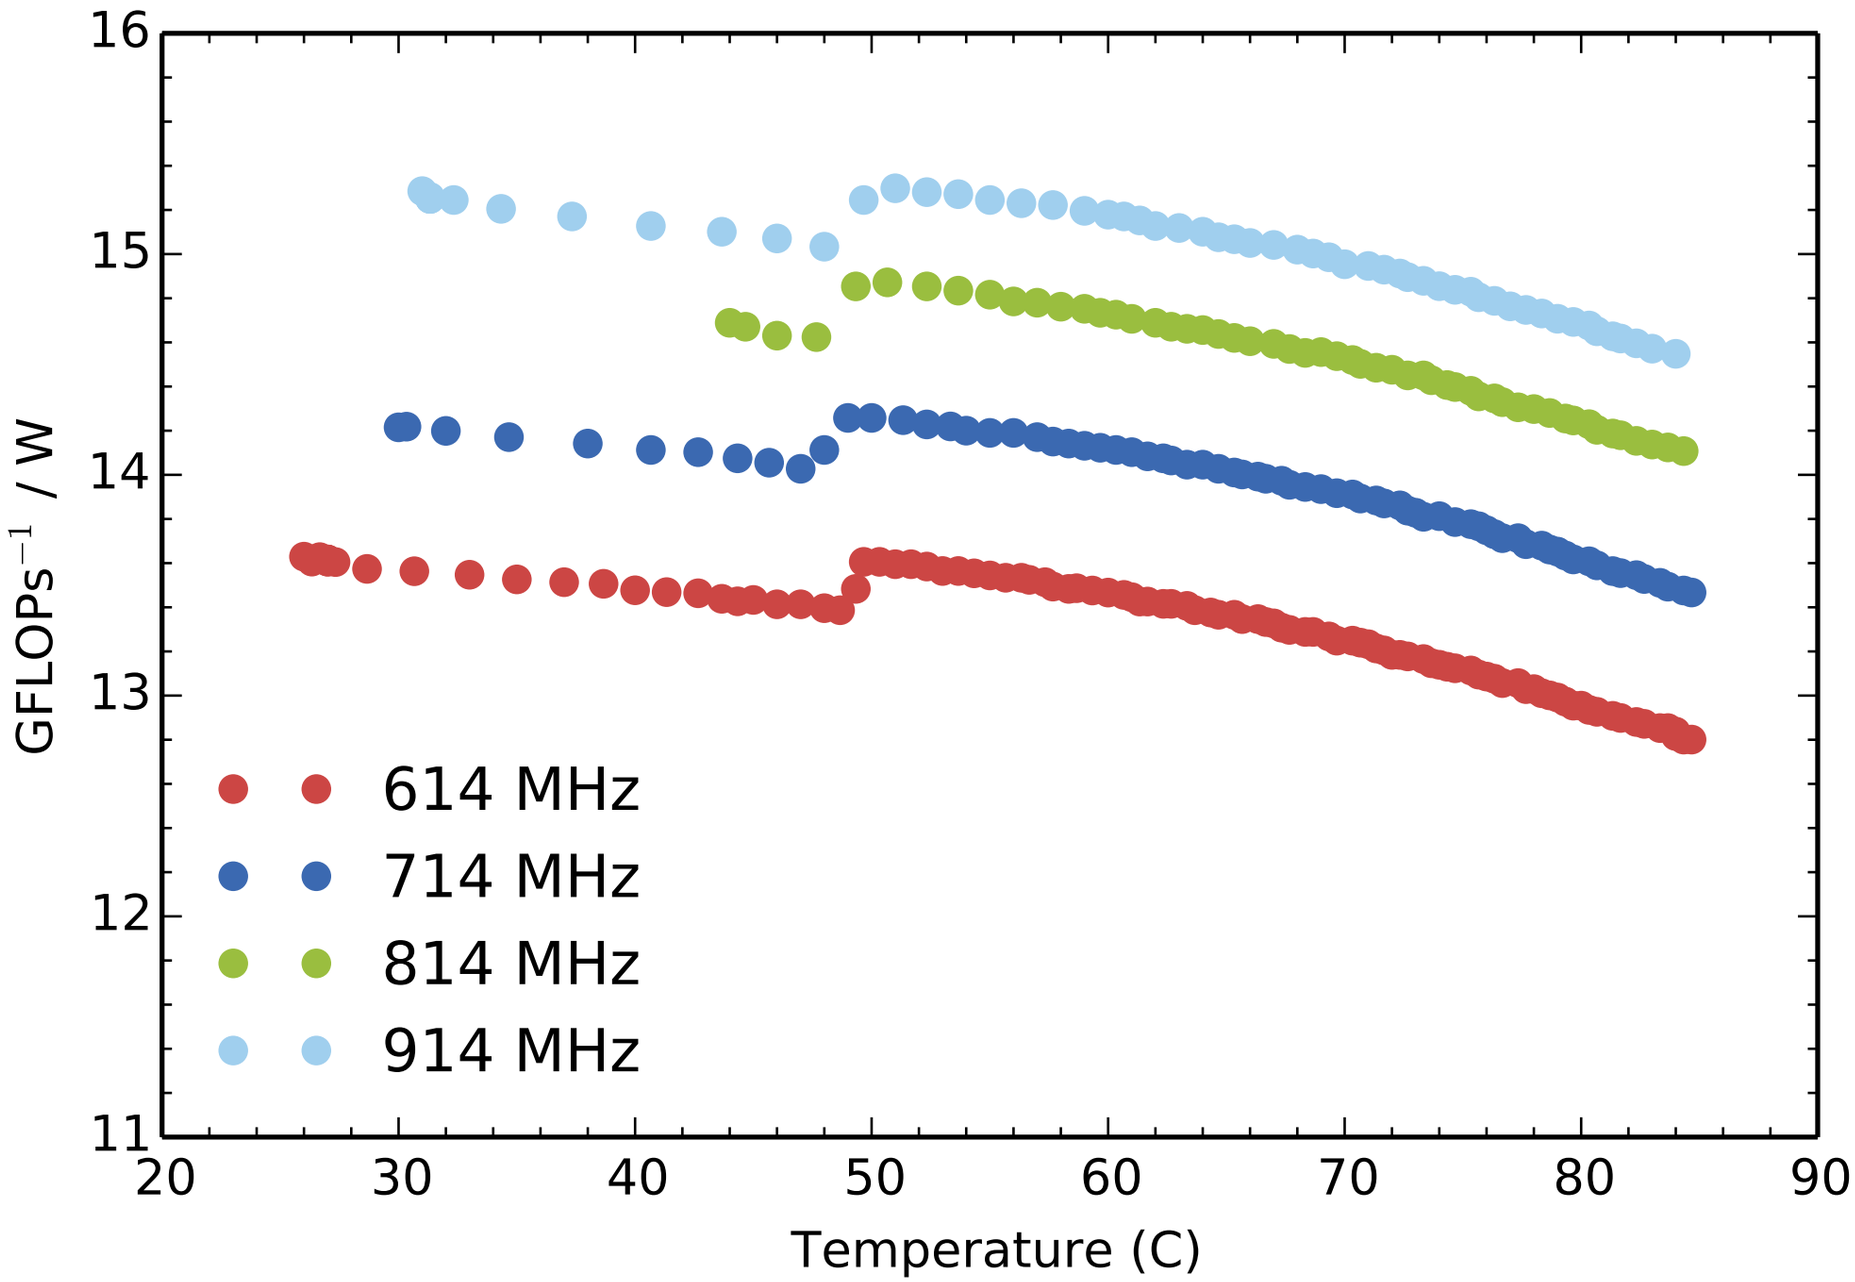

V = V1

Heavily undervolted

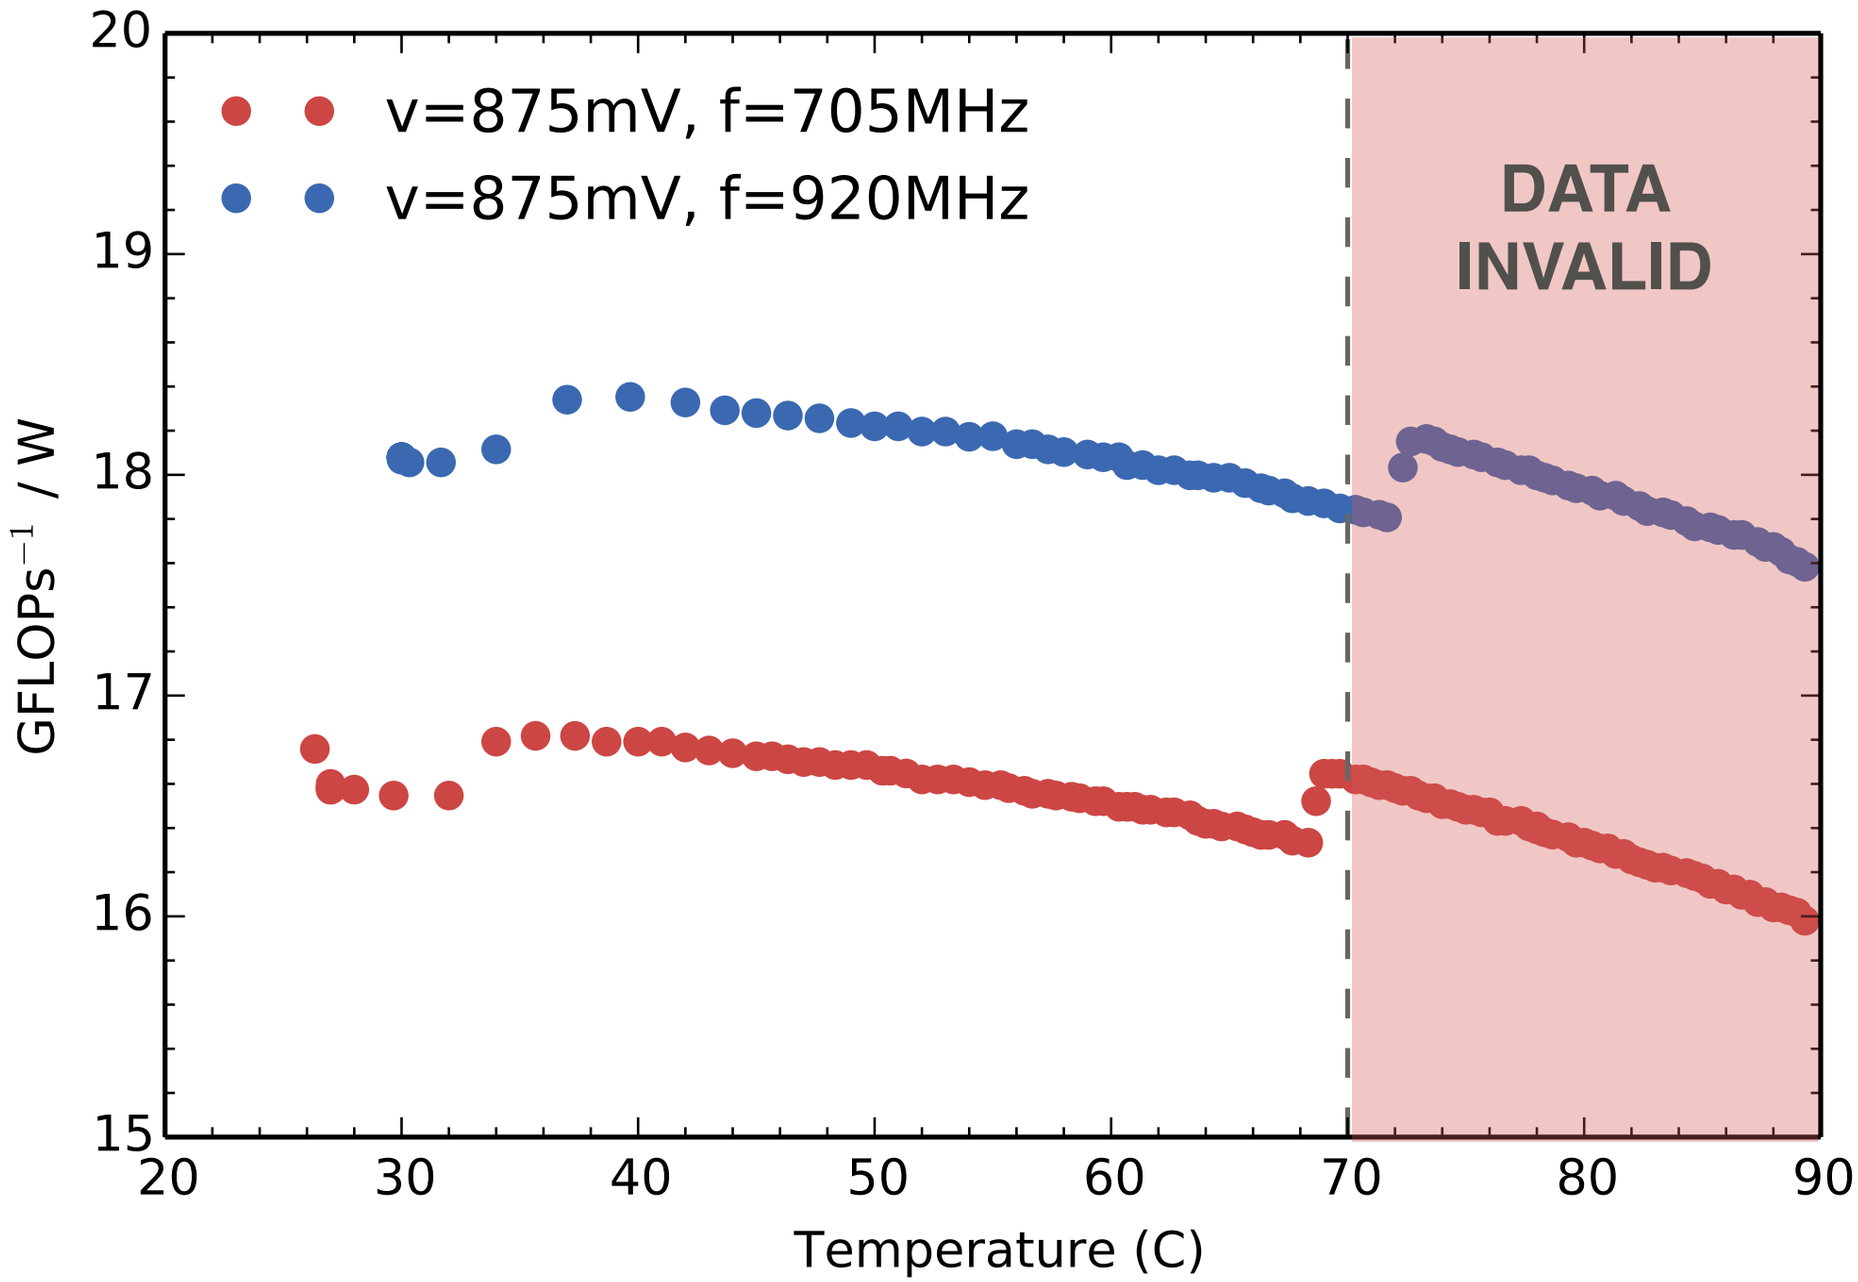

Green GPU and the *Test Equity 1000* Chamber of Doom

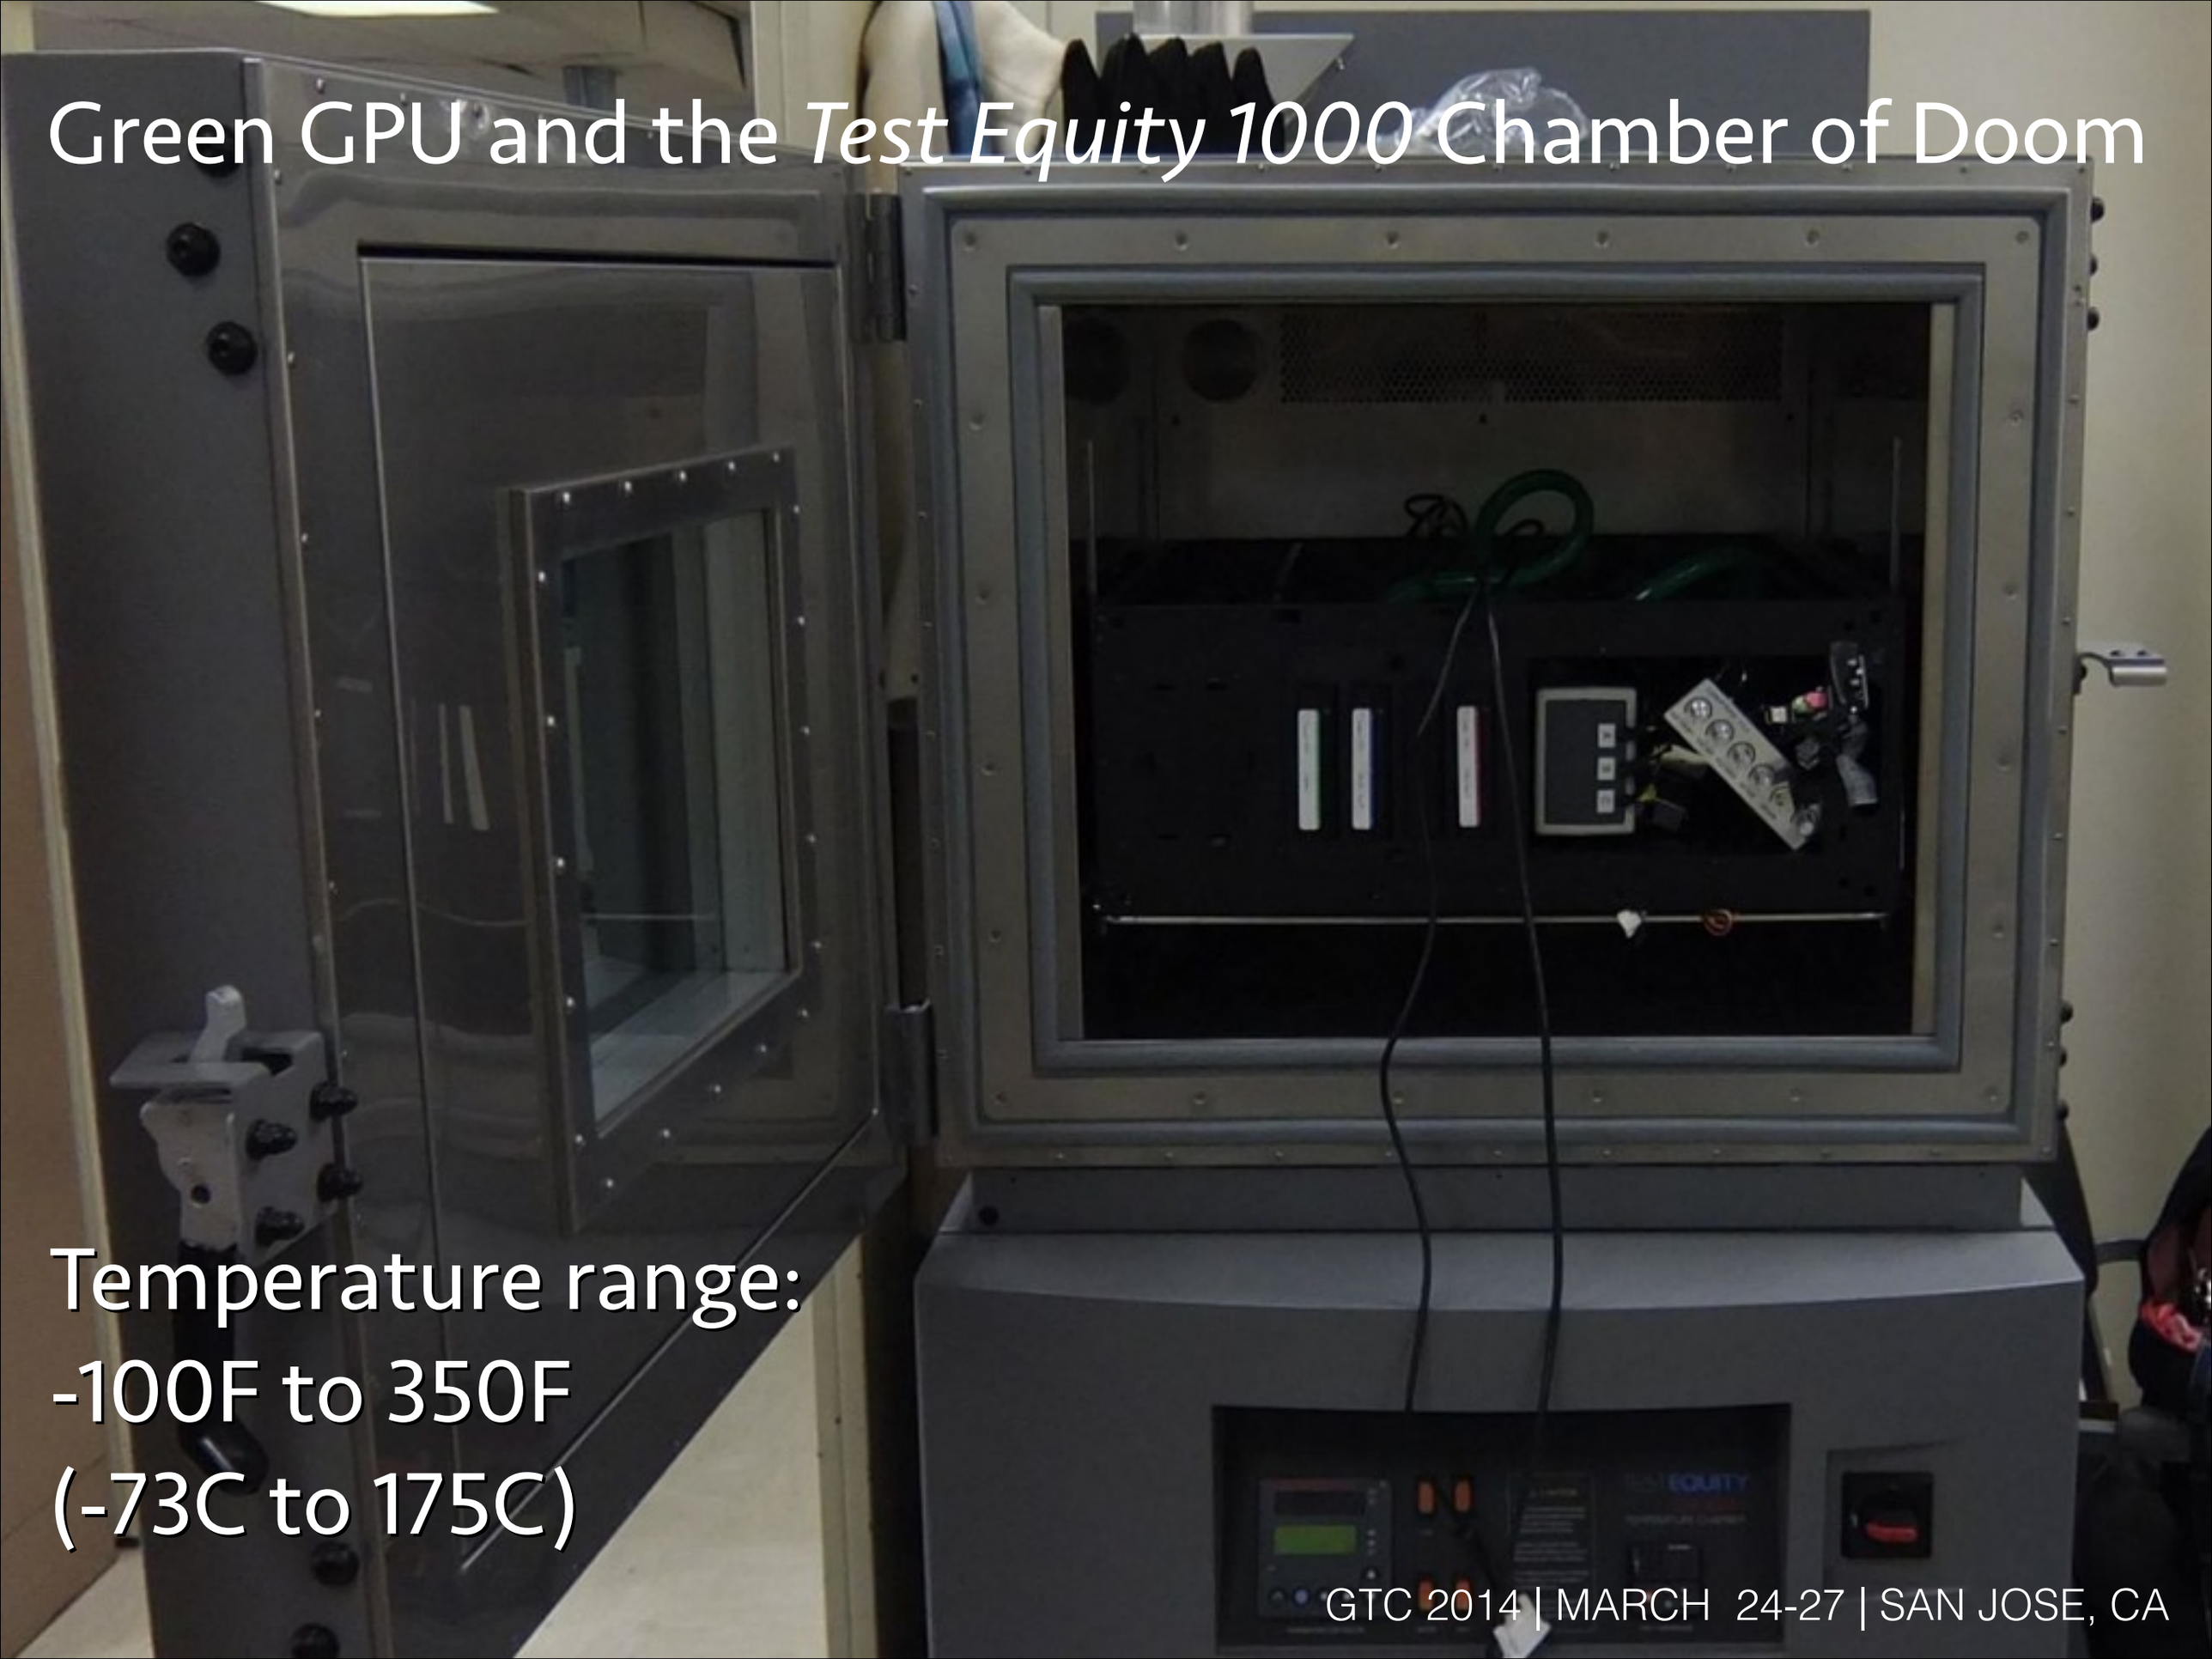

Temperature range:
-100F to 350F
(-73C to 175C)

Green GPU and the *Test Equity 1000* Chamber of Doom

or

Thermal shock: how to *really really* void your warranty

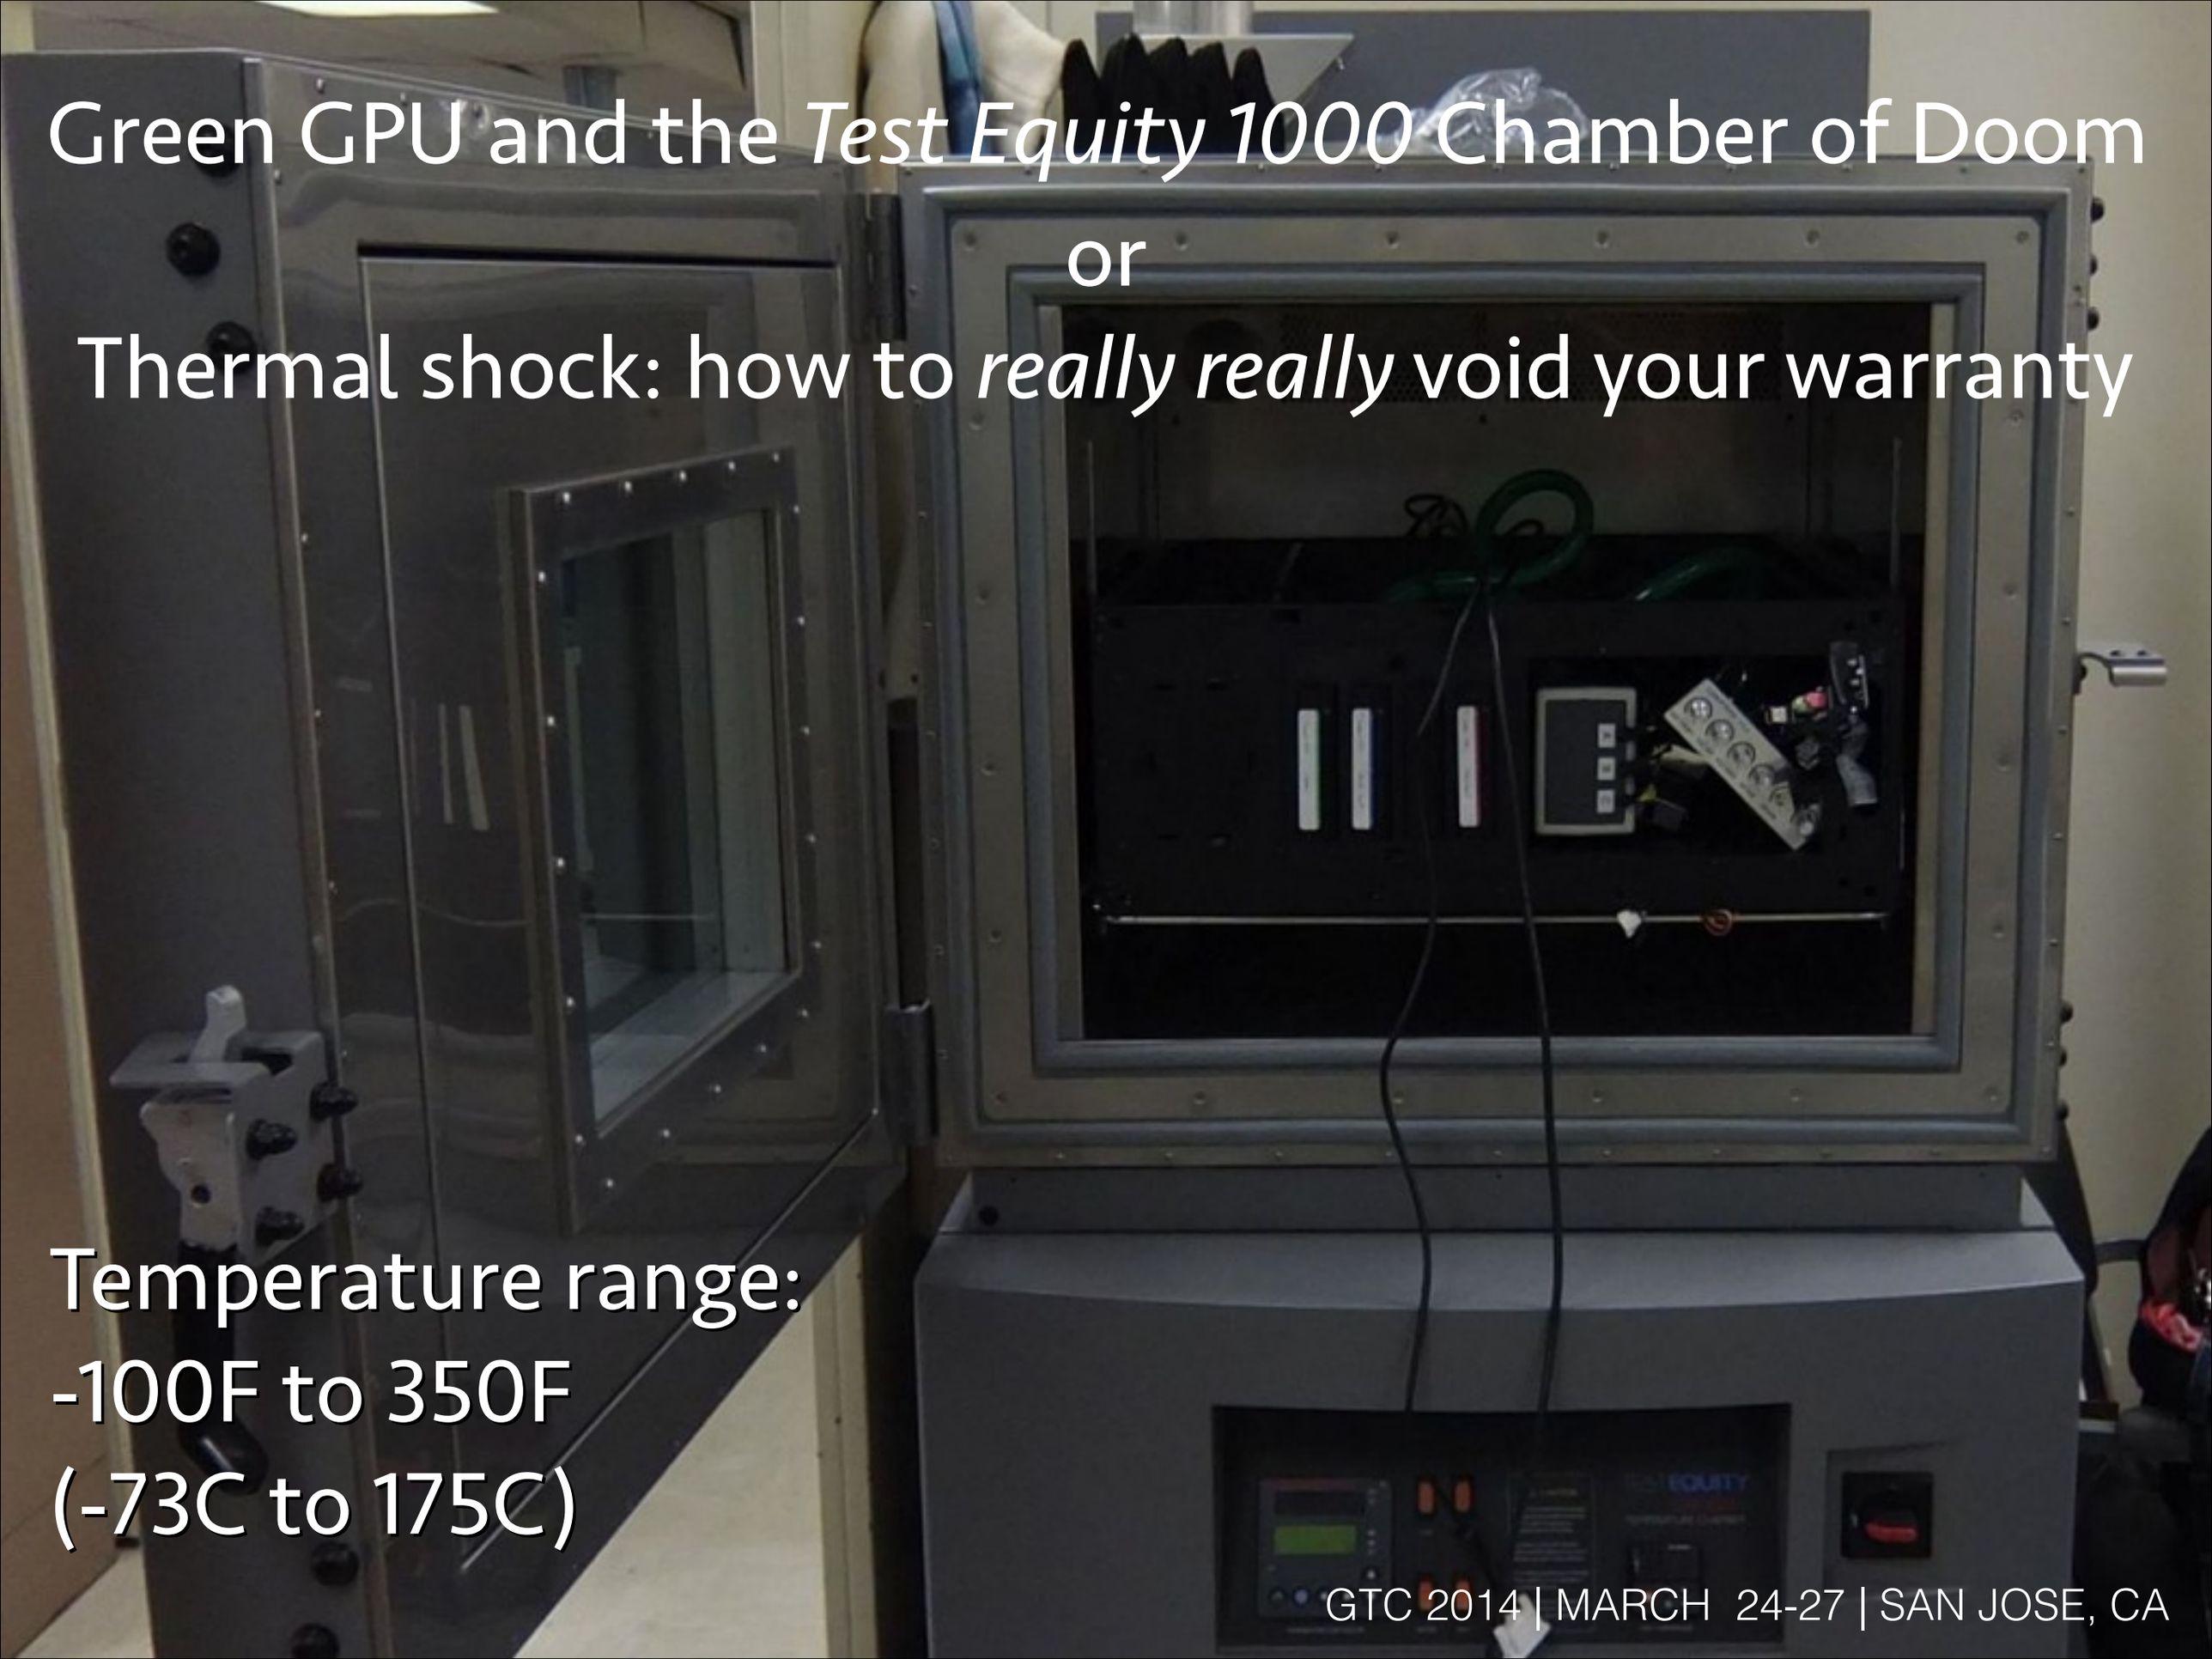

Temperature range:
-100F to 350F
(-73C to 175C)

Green GPU and the *Test Equity 1000* Chamber of Doom

or

Thermal shock: how to *really really* void your warranty

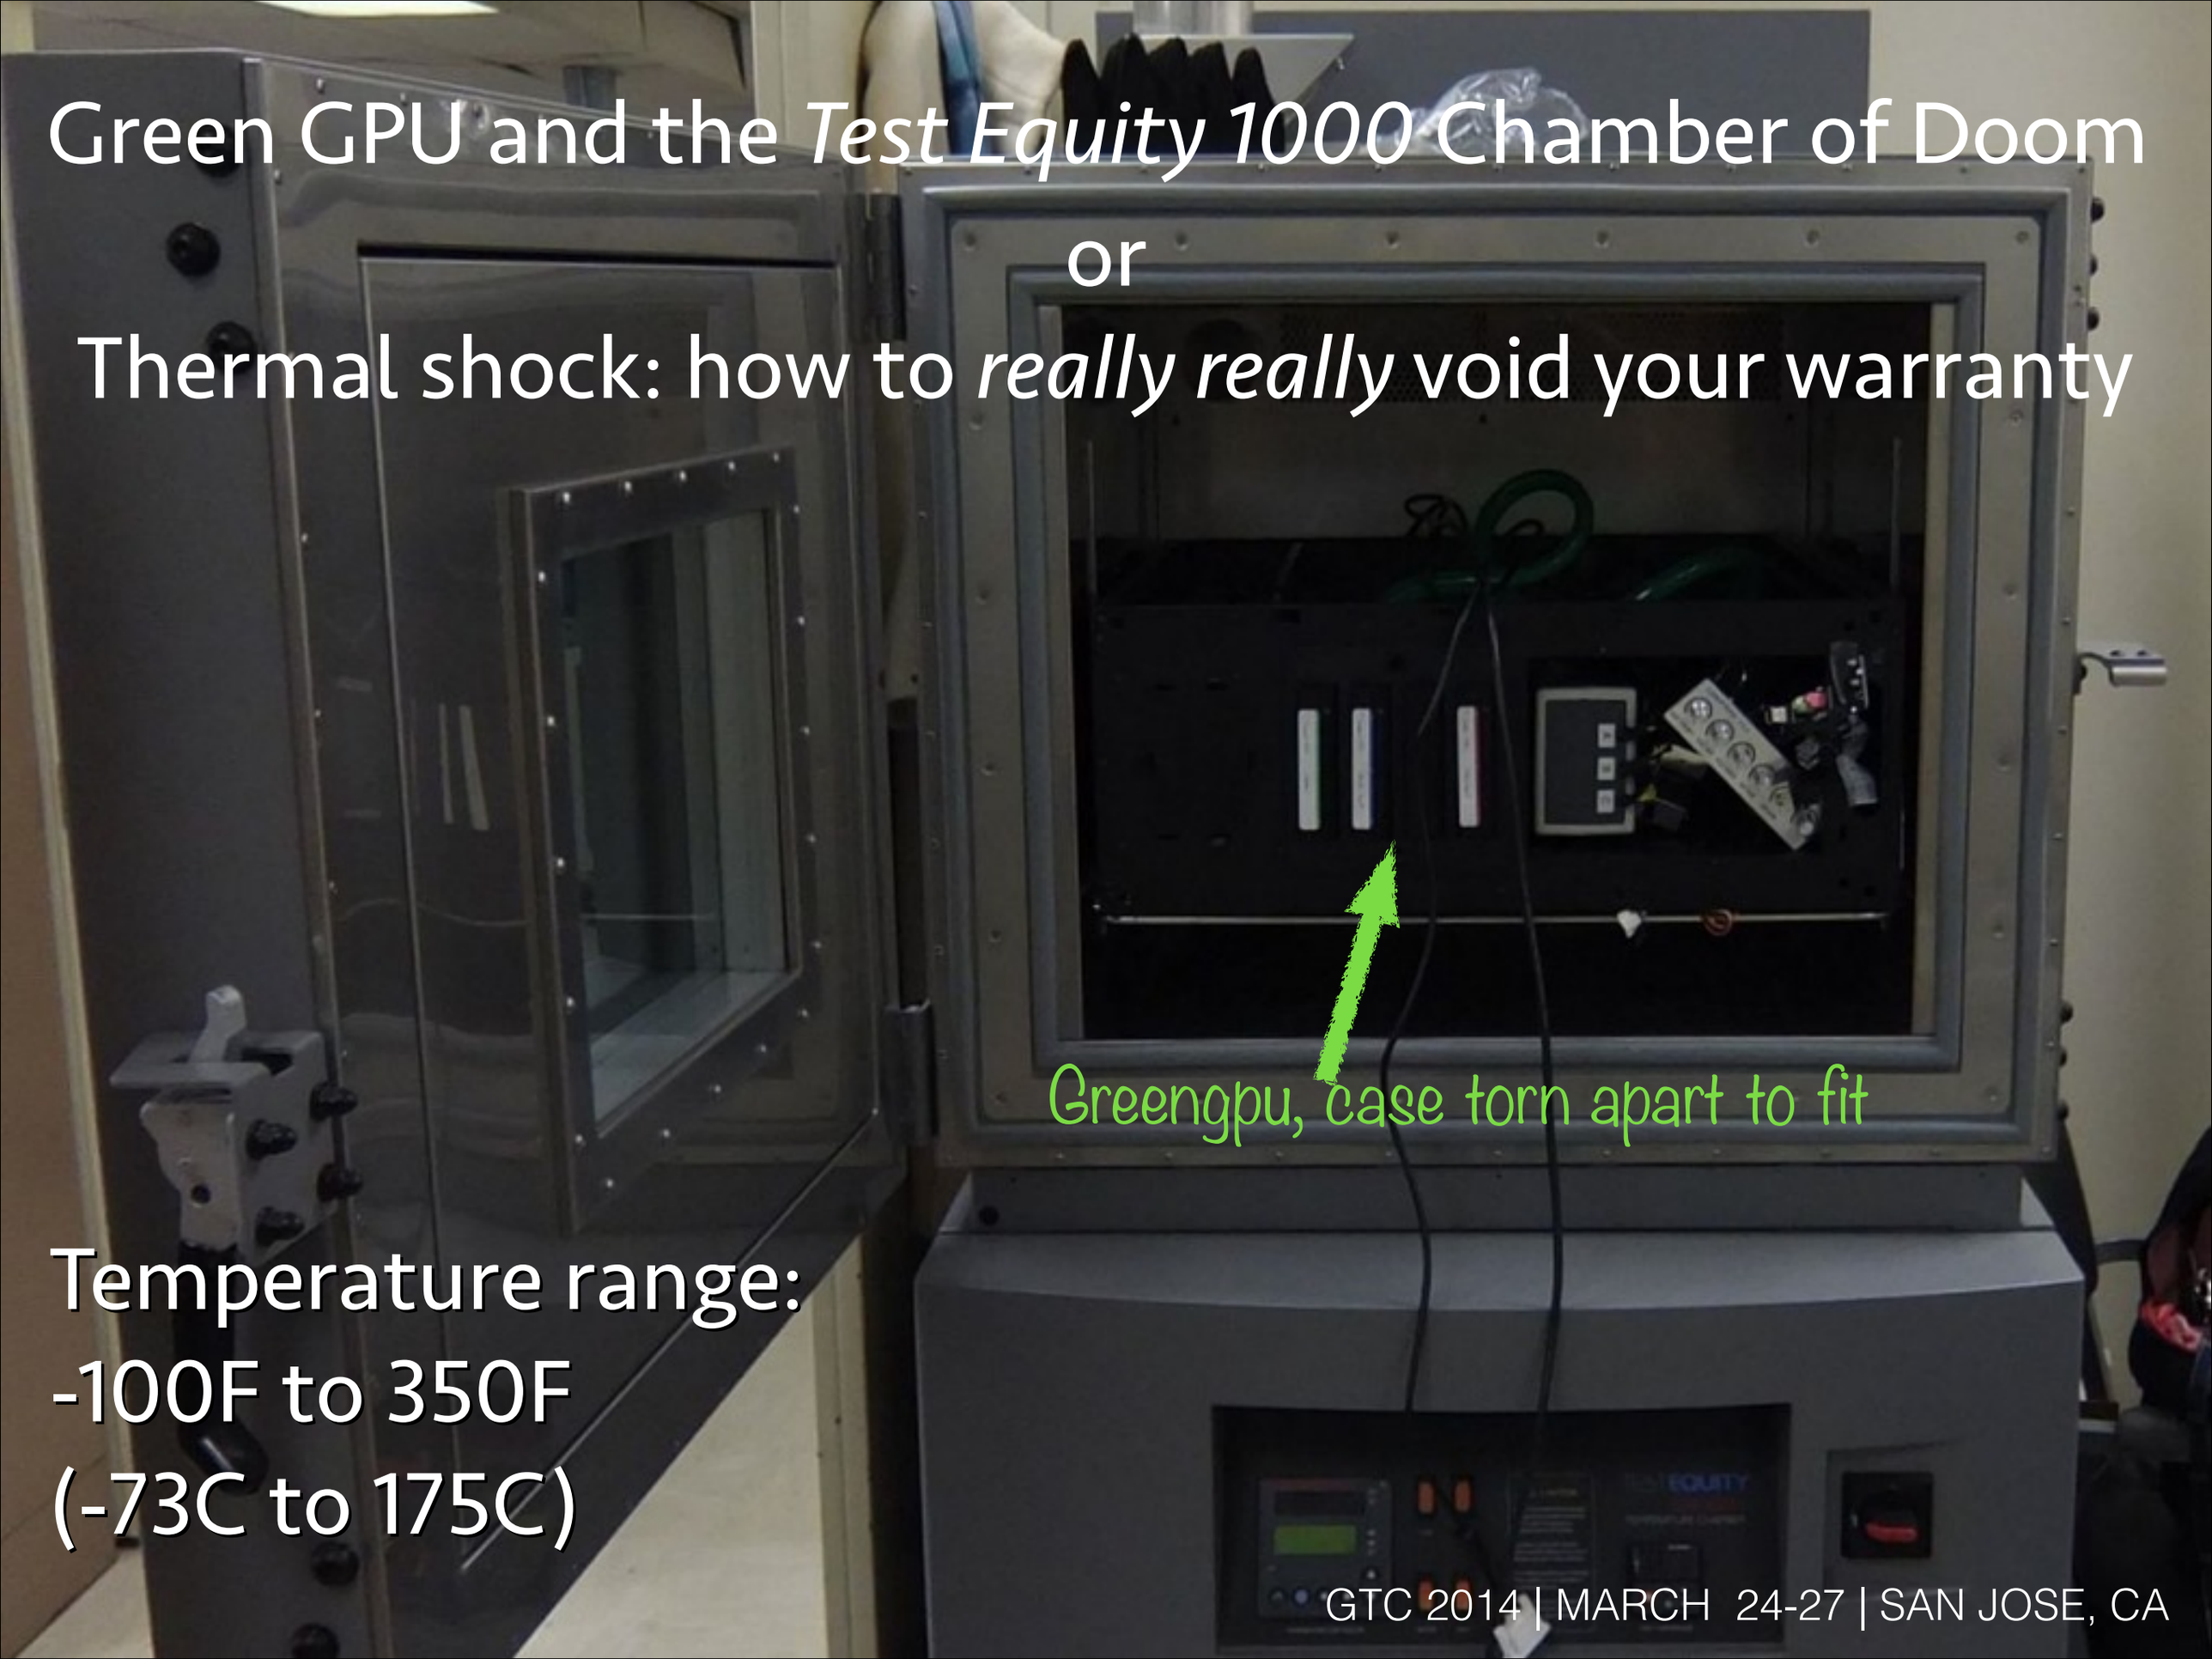

Temperature range:
-100F to 350F
(-73C to 175C)

Overclocked & Undervolted

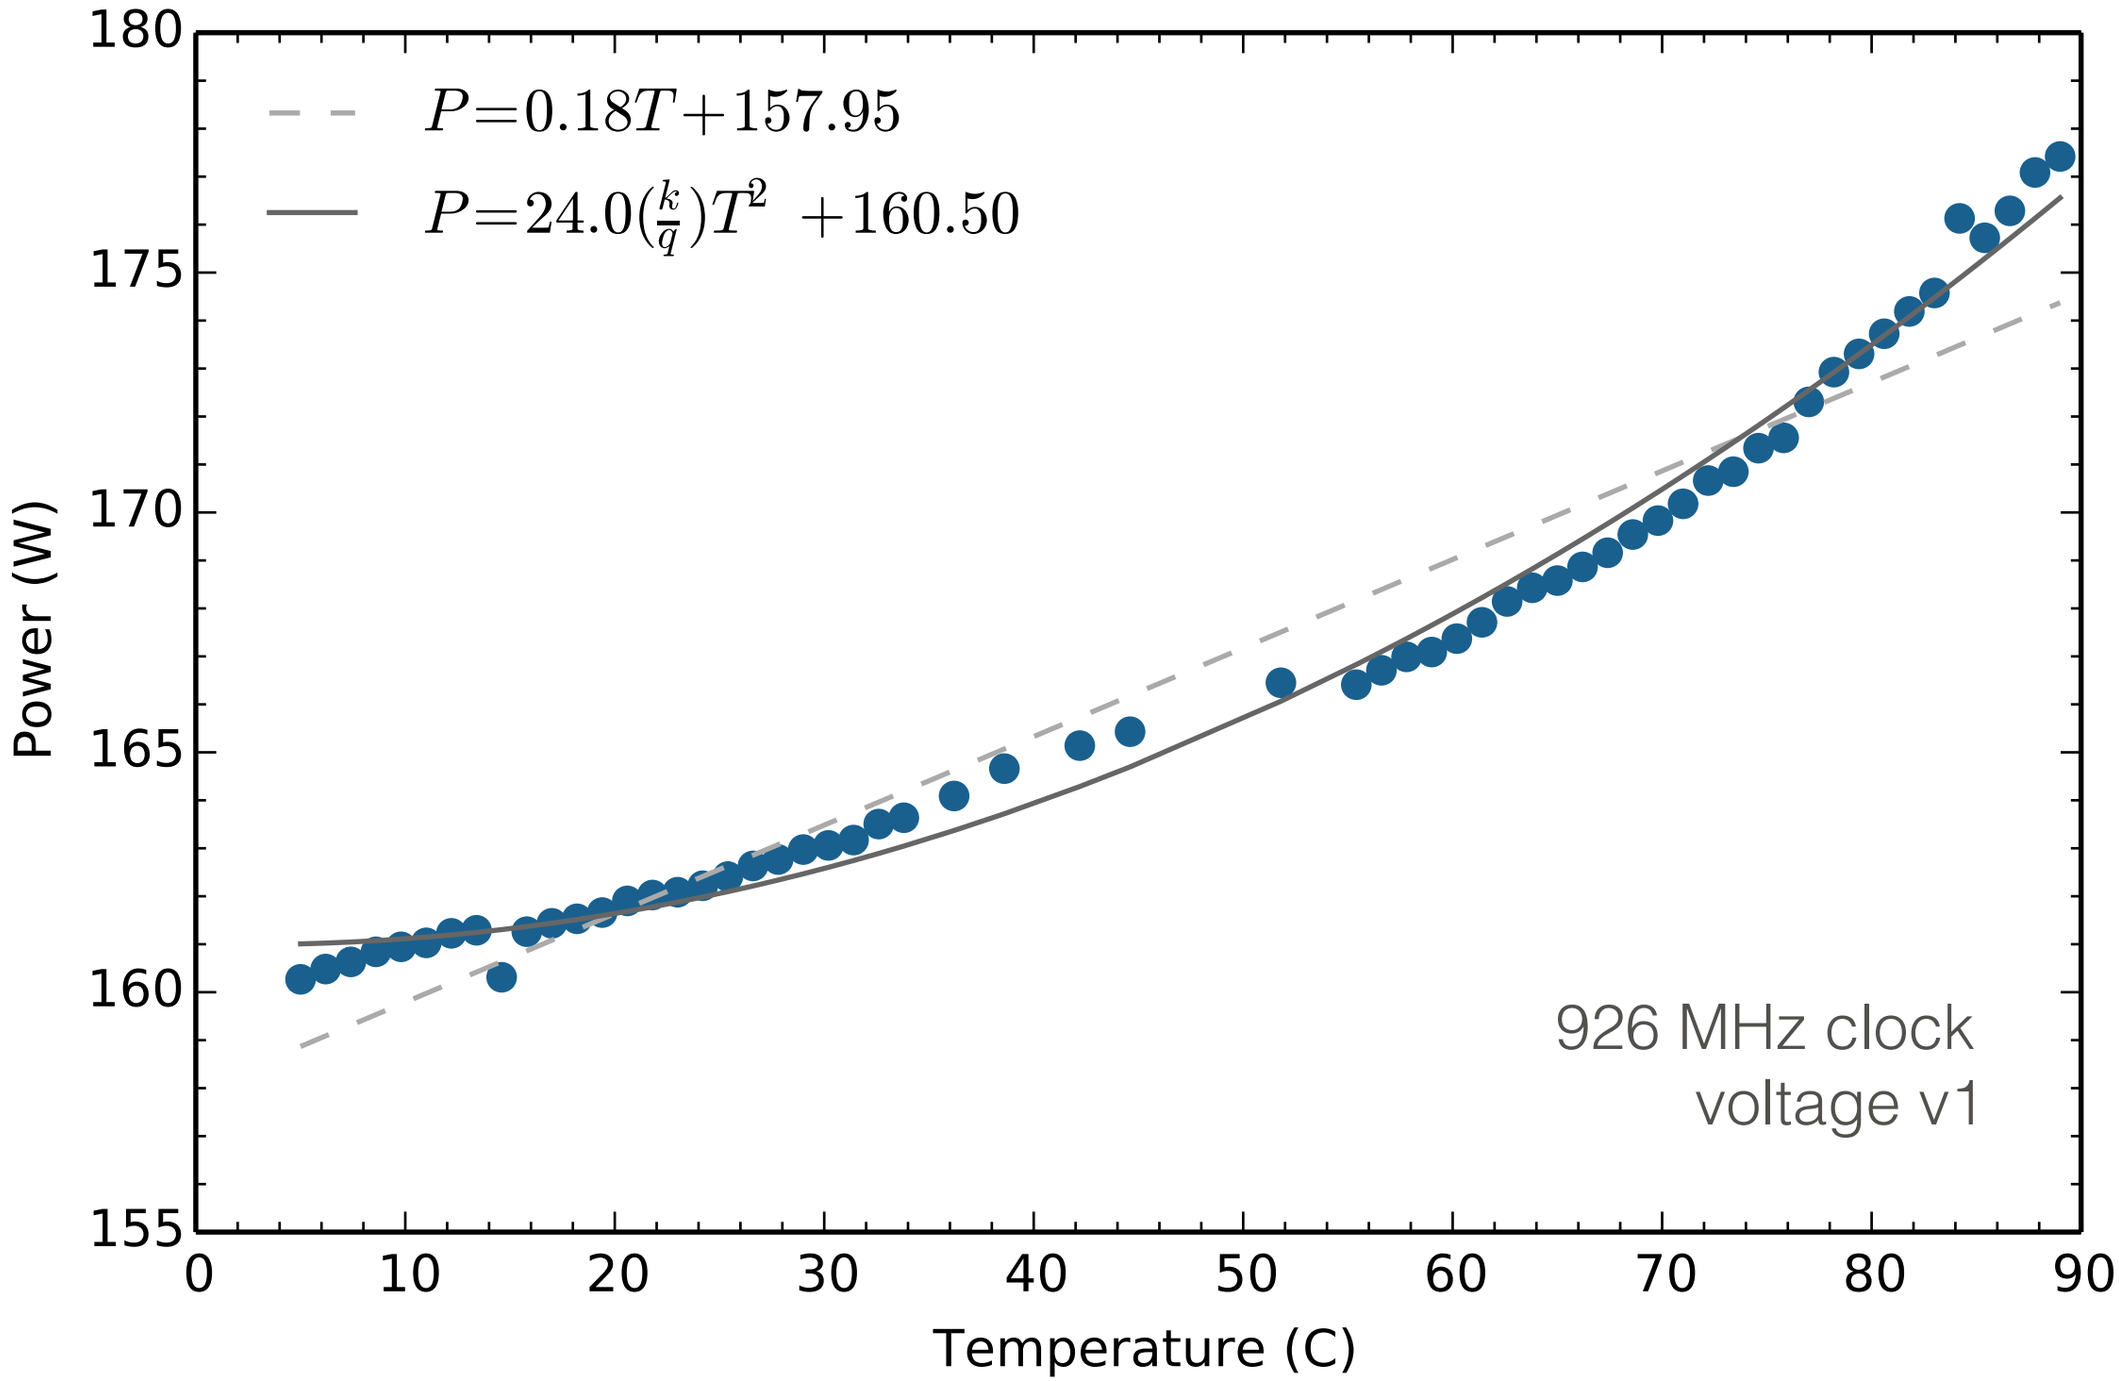

Overclocked & Undervolted

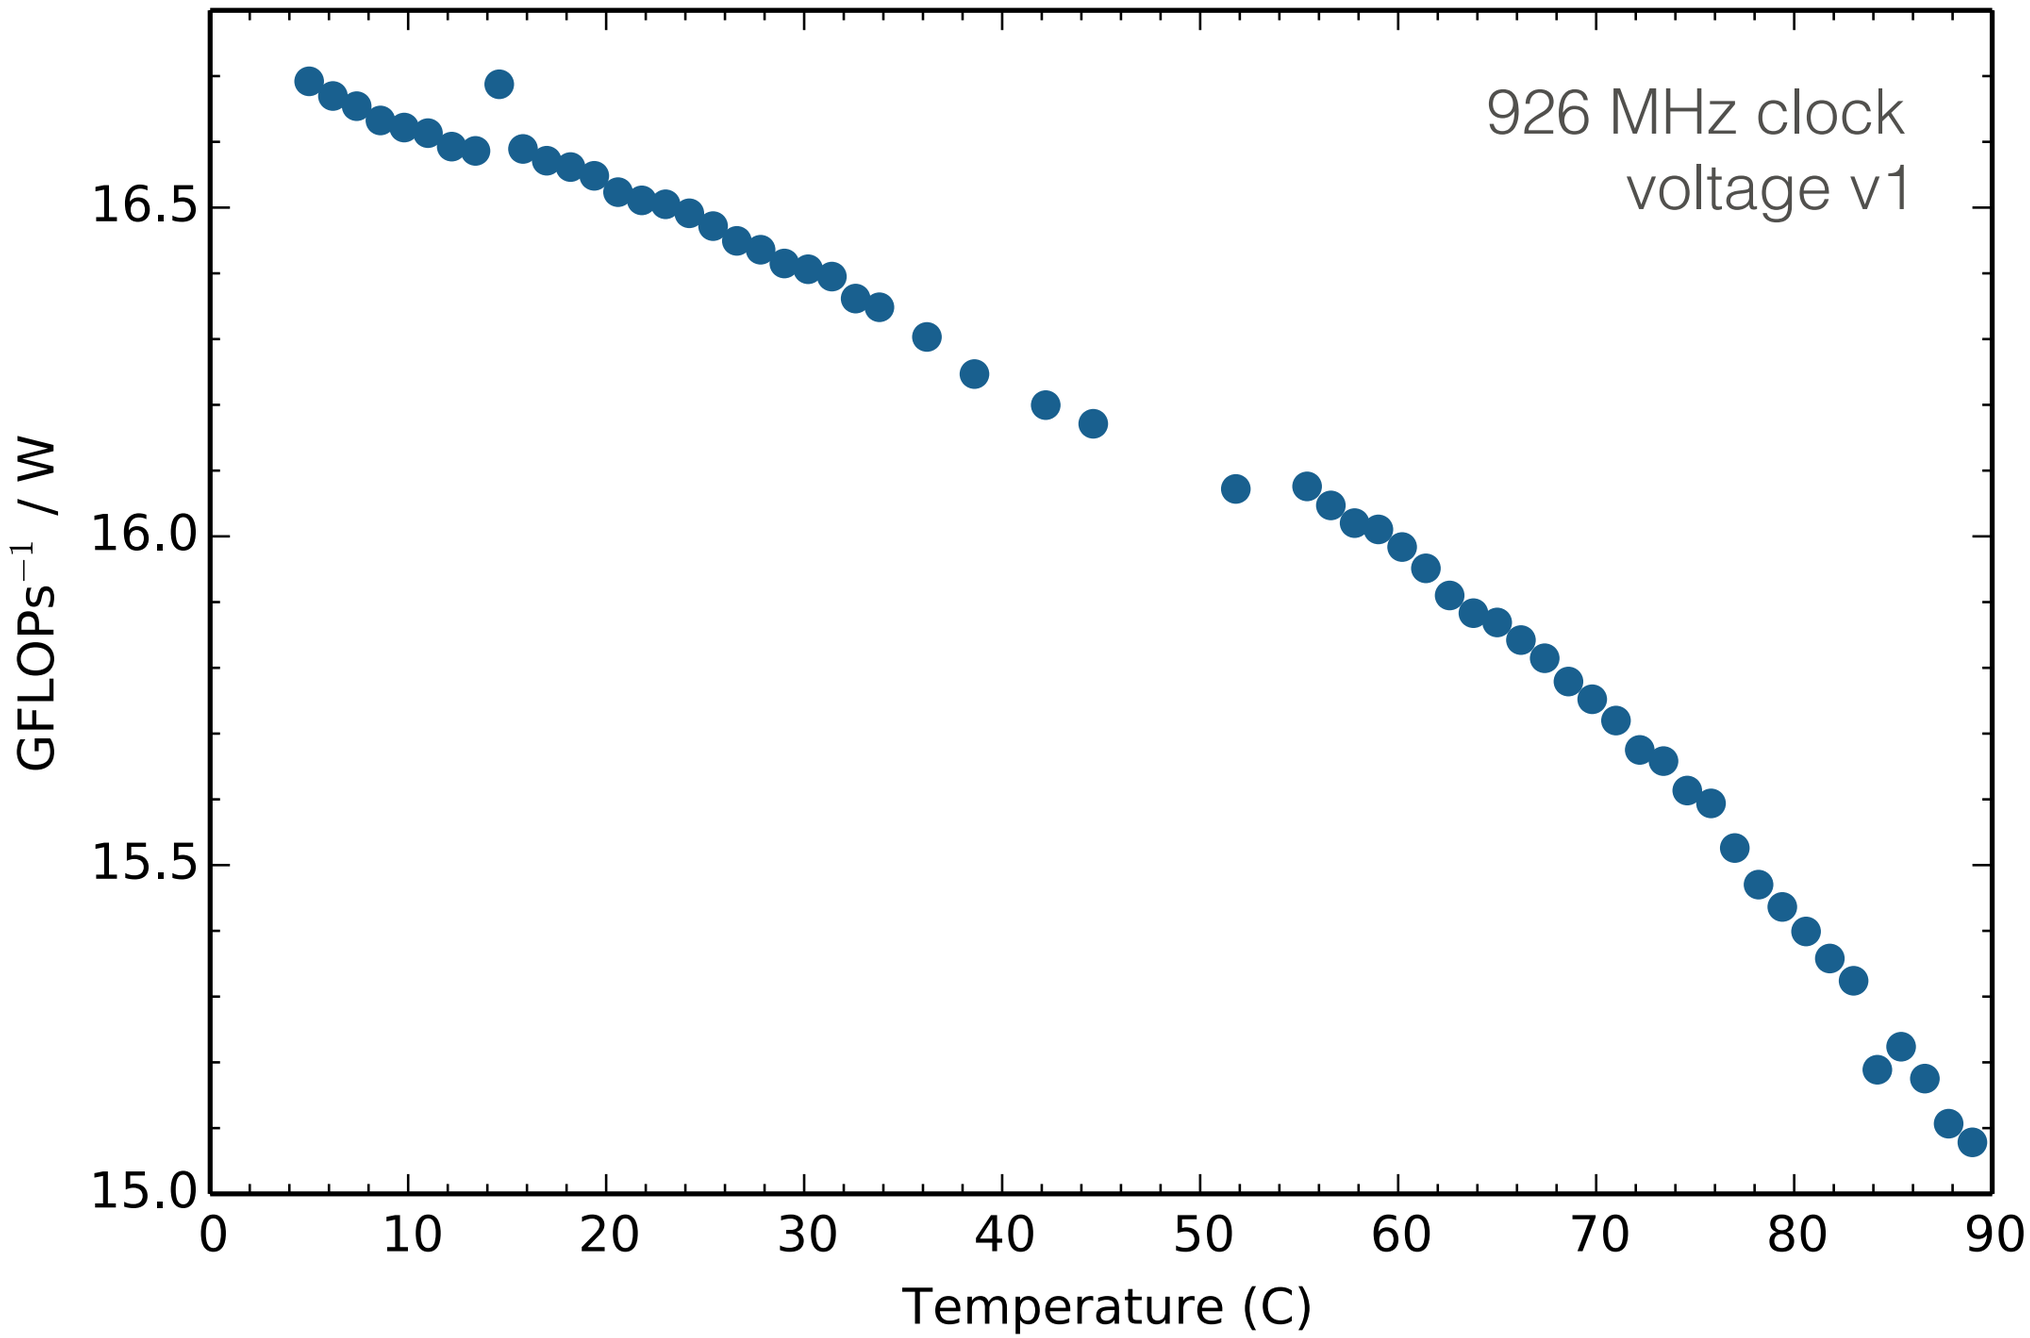

Part III: Comparisons & Conclusions

Kepler vs. Fermi

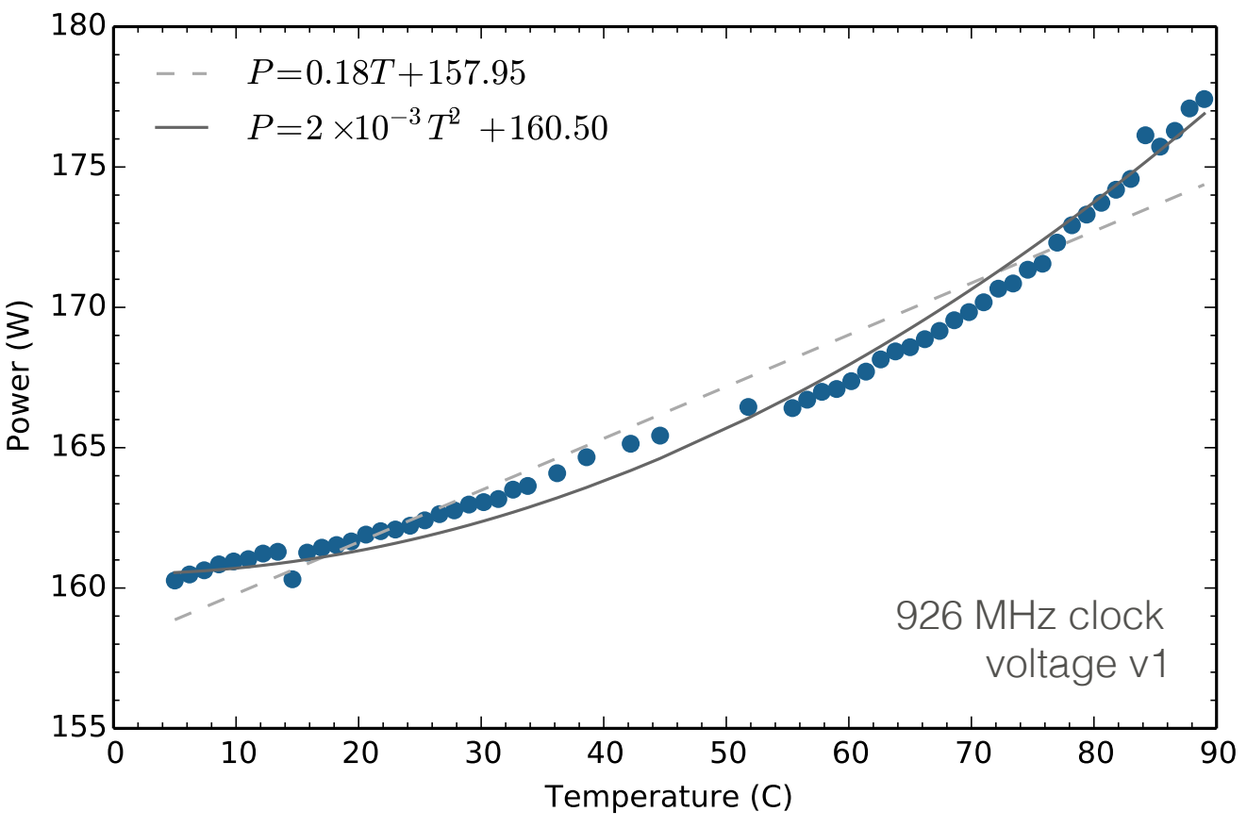

Tesla K20m

GK110, 28nm

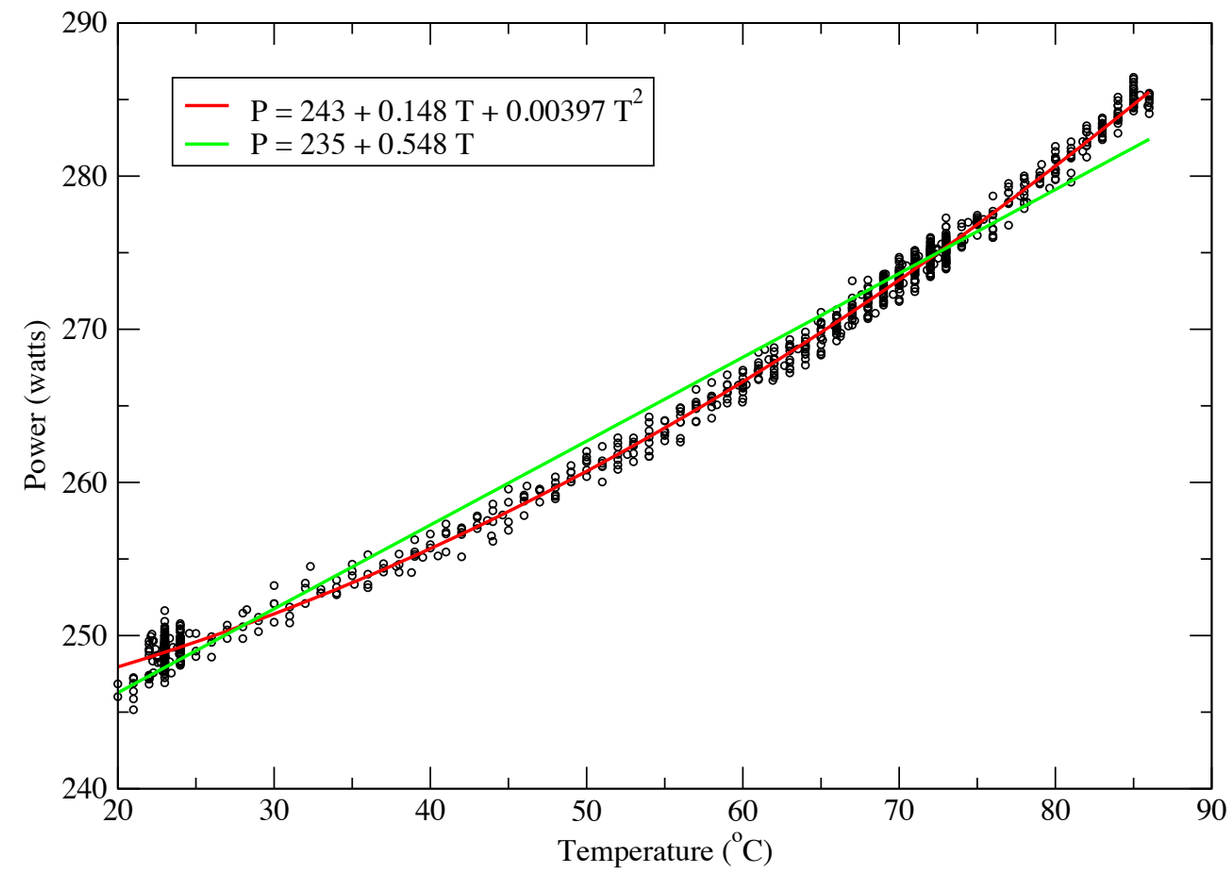

Geforce GTX 580

GF110, 40nm

Maxwell results (preliminary)

GTX 750 Ti

- 640 CUDA Cores
- 1020 MHz base Clock (MHz)
- 1085 MHz boost Clock (MHz)
- 60W TDP

xGPU

- 1062 GFLOP/s performance
- 76% efficiency of code
- expect >80% efficiency with code optimization
- already >16 GFLOPS/W
- can’t control clock / voltage (yet…)

Conclusions

- 12.96 —> 18.34 GFLOPs/W, i.e. **41% increase** by overclock + undervolt on K20m for xGPU.
- Maxwell is >16 GFLOPs/W for xGPU out of the box!

Conclusions

- xGPU: a useful CUDA benchmark and “burn-in” tool
- Lower temperature = lower power consumption
- Overclocking w/o overvolting = good
- Overclocking + unvervolting = better
- Tesla K20m is pretty hard to brick

Acknowledgements

- Robert Kimberk (CfA lab tech)
- Overclocking community
- NVIDIA

CfA: Danny Price, Ben Barsdell, Lincoln Greenhill
NVIDIA: Mike Clark, Ron Babich

dprice@cfa.harvard.edu