Experience powering Virtex-7 FPGAs

Andrew W. Rose, Gregory M. Iles

Imperial College, London

MP7

- 1.5Tb/s optical signal processor
- Xilinx Virtex-7 FPGA:
  - XC7VX485T or XC7VX690T
- Advanced boot-loader & diagnostics (full system test at start-up)
- On-board firmware repository
- 2×144Mbit 550MHz QDR RAM (optional)
- Been in hand for over a year
  - Continuous testing over that period
  - Exceptionally well understood

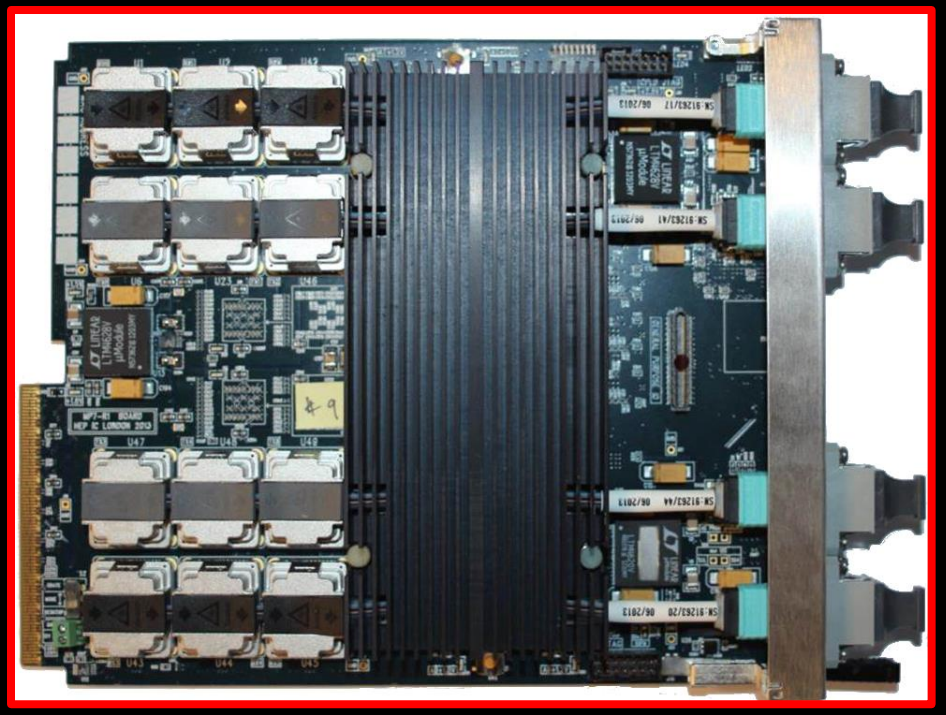

Imperial MP7 processor board

See also TWEPP 2012: https://indico.cern.ch/contributionDisplay.py?contribId=86&confId=170595&sessionId=51
https://indico.cern.ch/contributionDisplay.py?contribId=97&confId=170595&sessionId=53

FPGA architecture

- MP7 uses XC7VX485T or XC7VX690T
- Pin-compatible FFG1927 package

- This is how the XC7VX690T looks in the documentation:

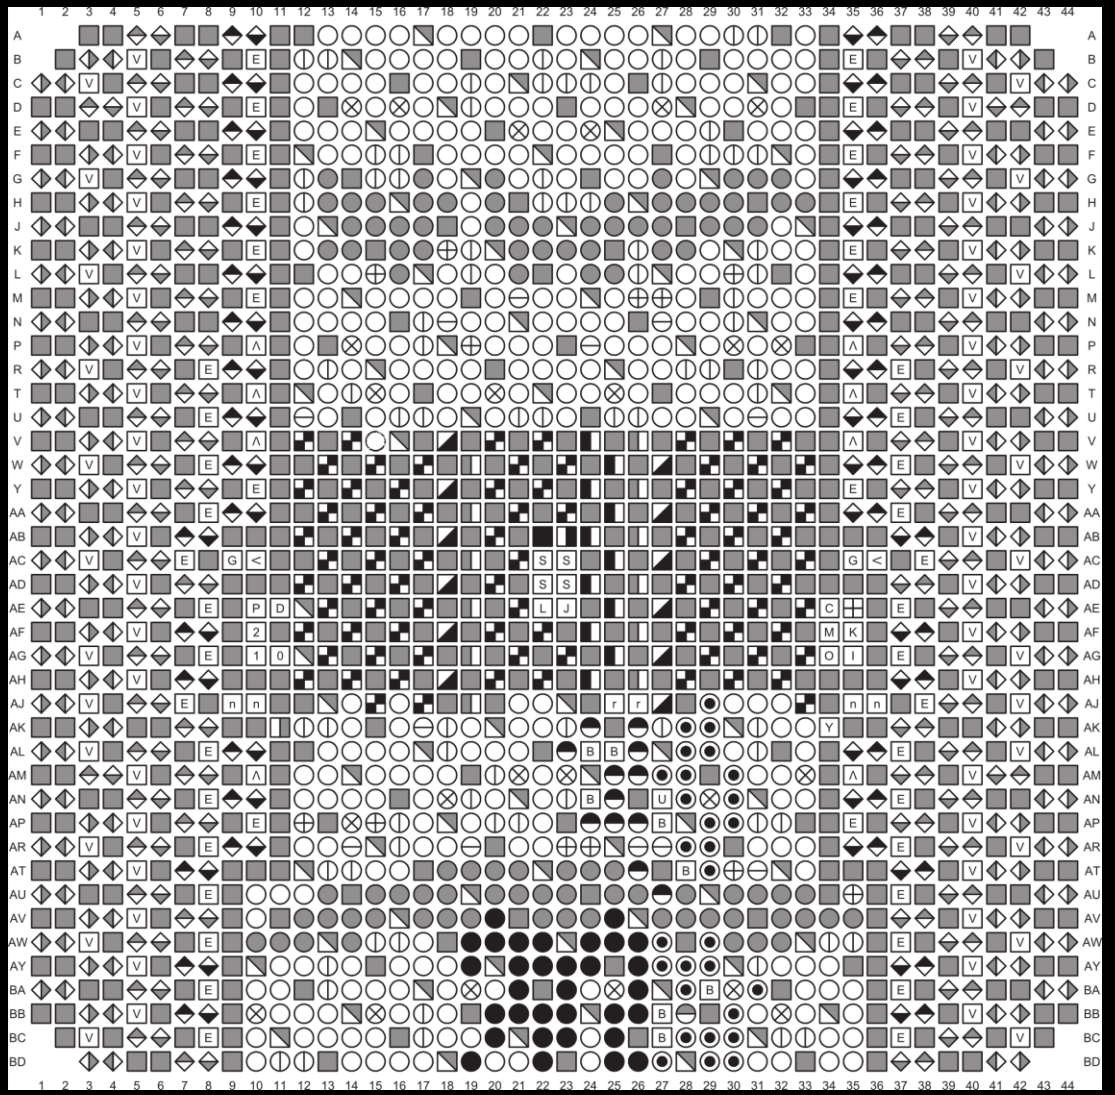

FPGA architecture

- MP7 uses XC7VX485T or XC7VX690T
- Pin-compatible FFG1927 package
- And if we colour in the power pins:
- Half of the pins on the chip are dedicated to powering the thing

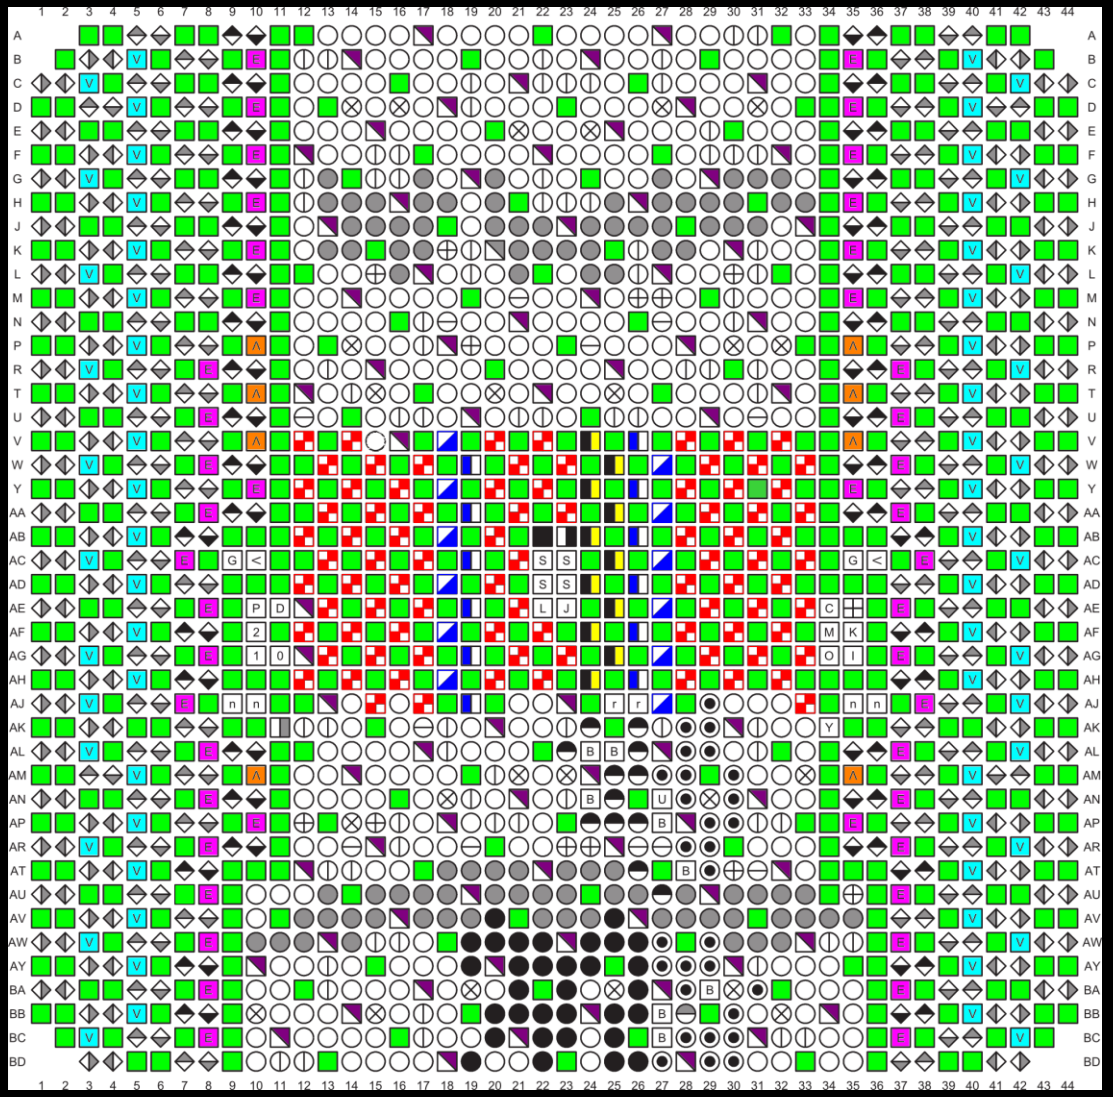

MP7 board – Top half

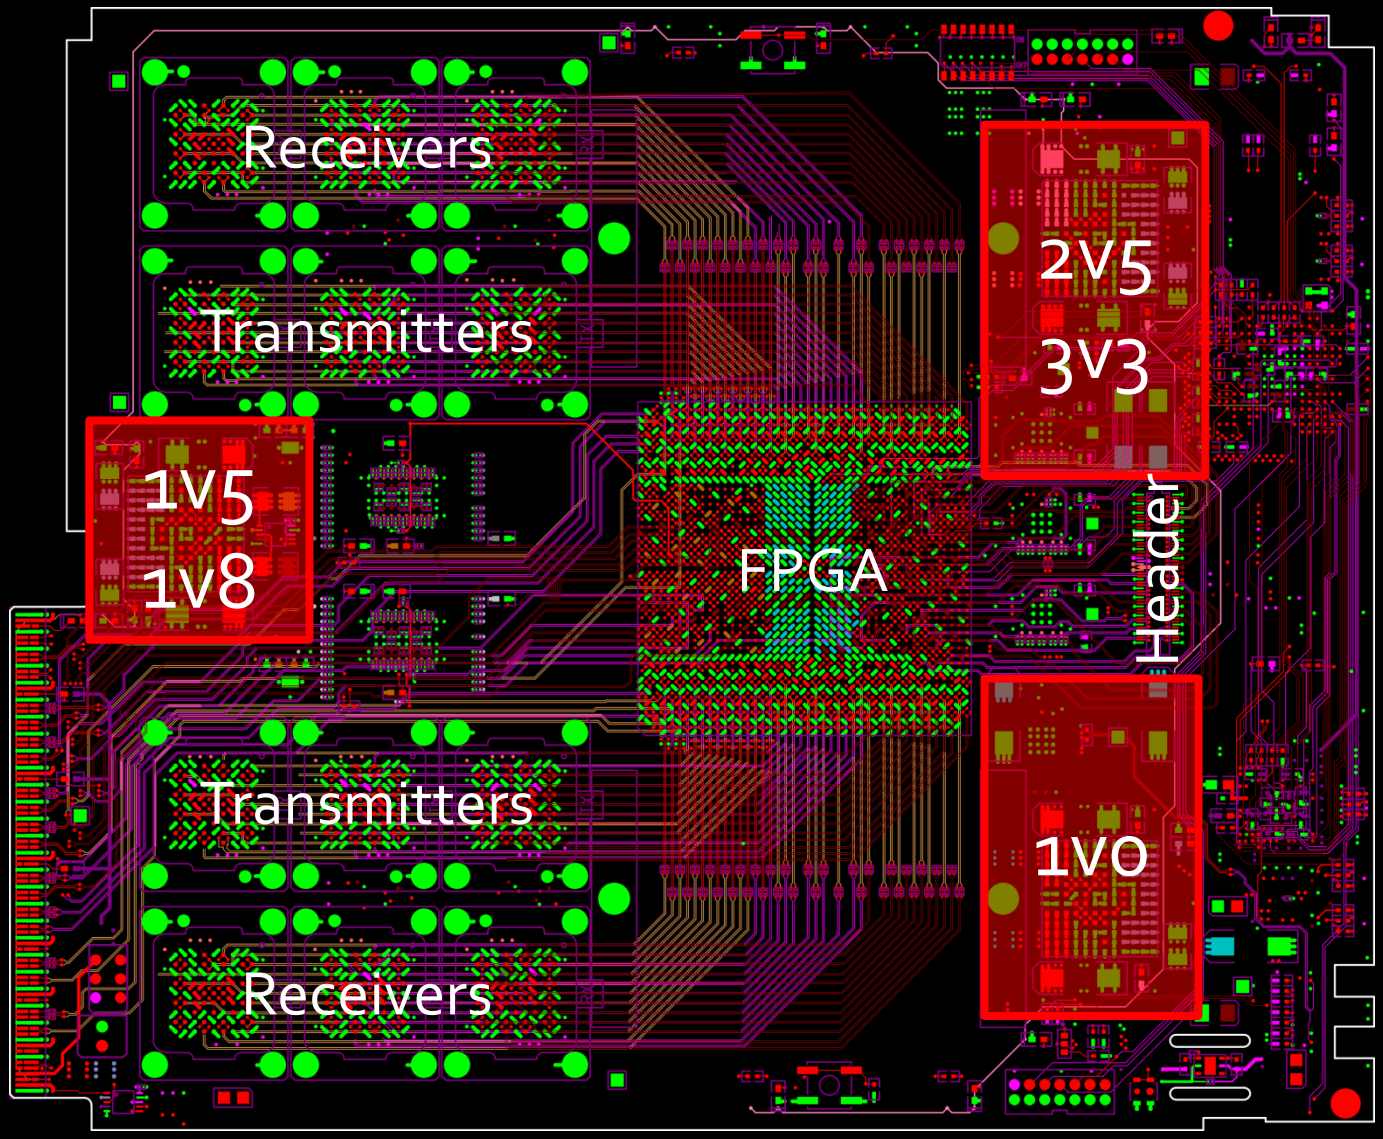

MP7 board – Bottom half

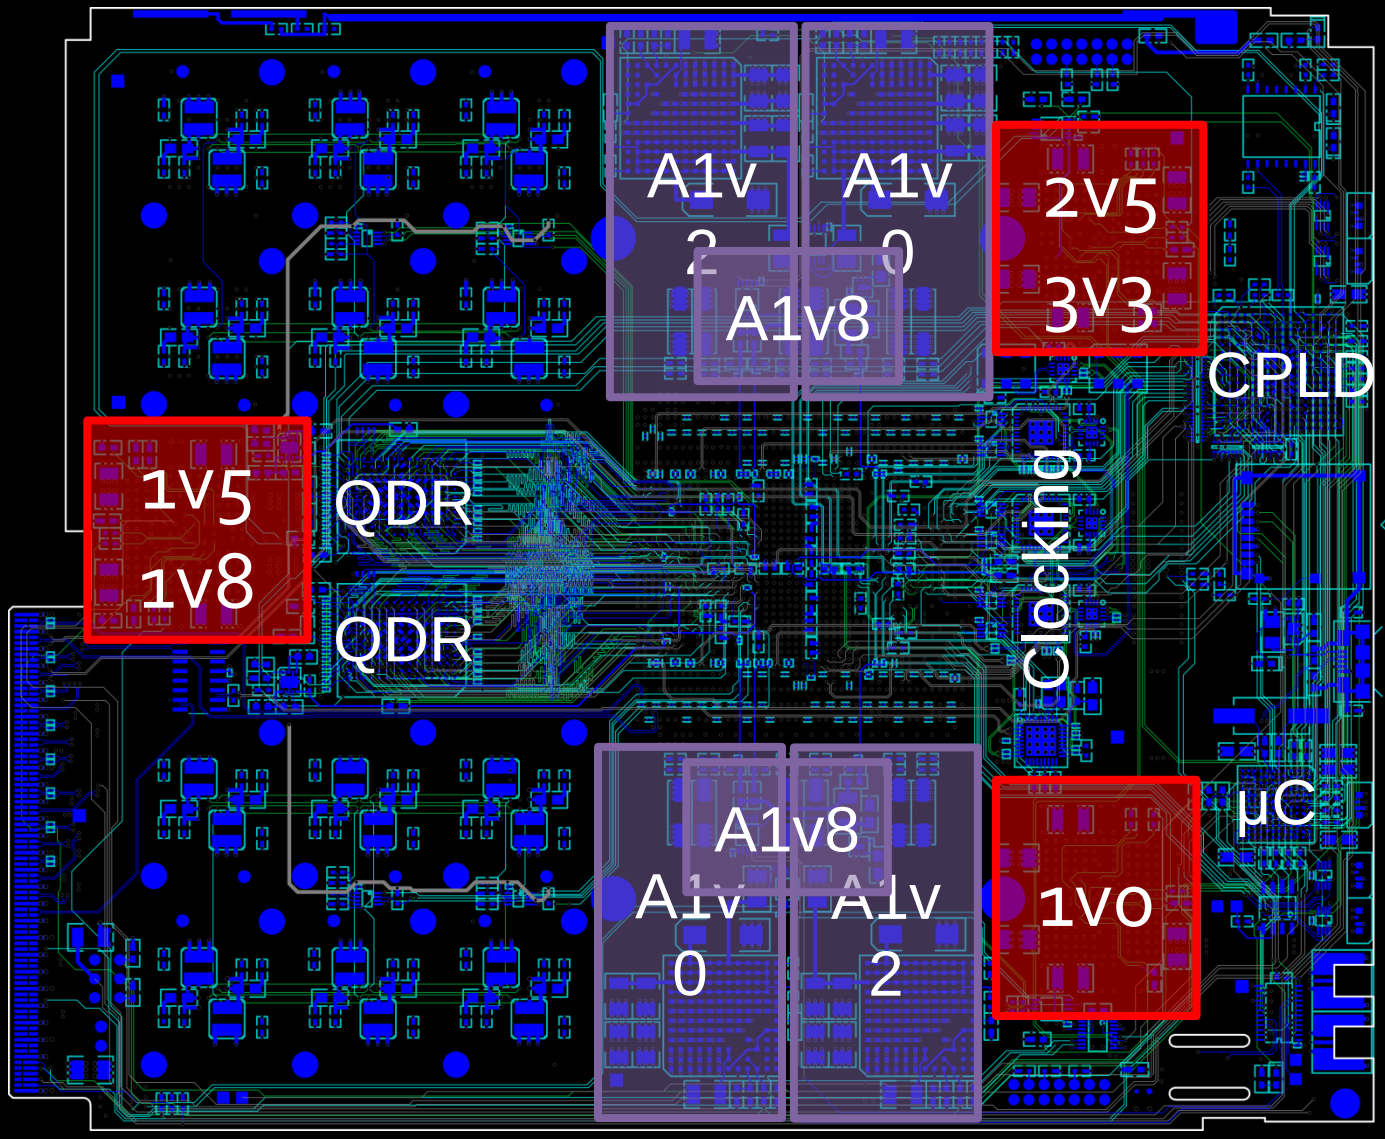

MP7 Ro power architecture

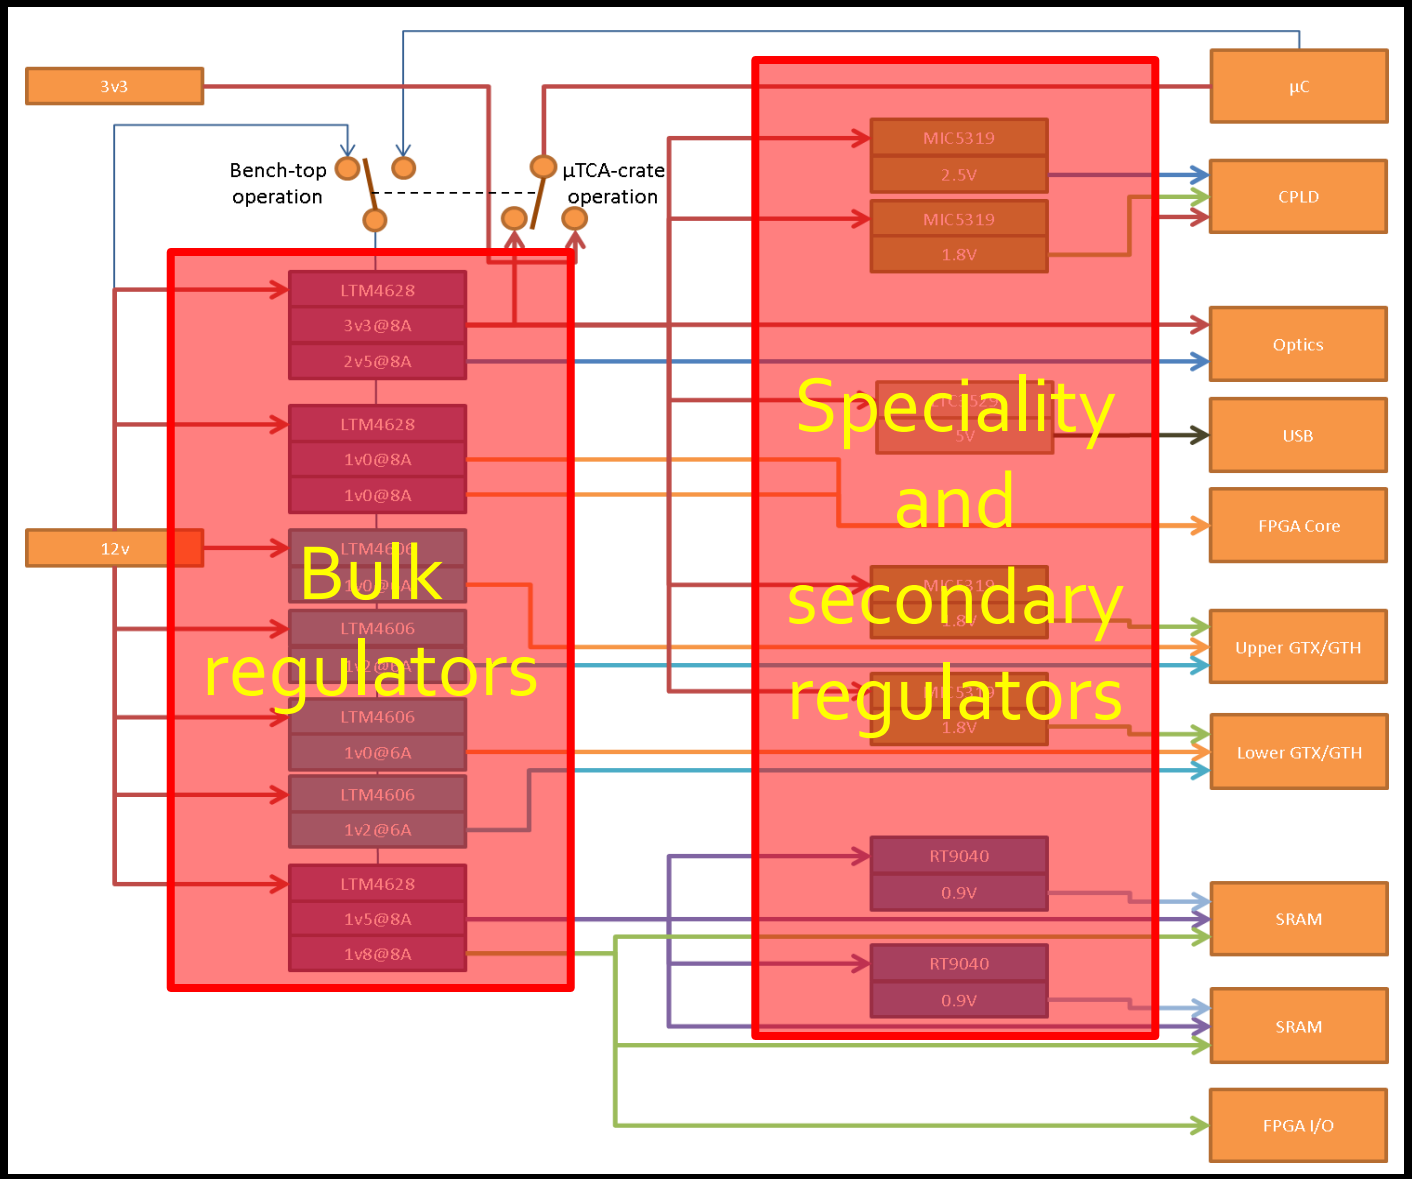

MGTs

- MP7 is an optical stream processor
- 72 links @ >10Gb/s, fully bidirectional
- 7-series MGTs require:
  - 1v0
  - 1v2
  - 1v8
- but there are very tight constraints!

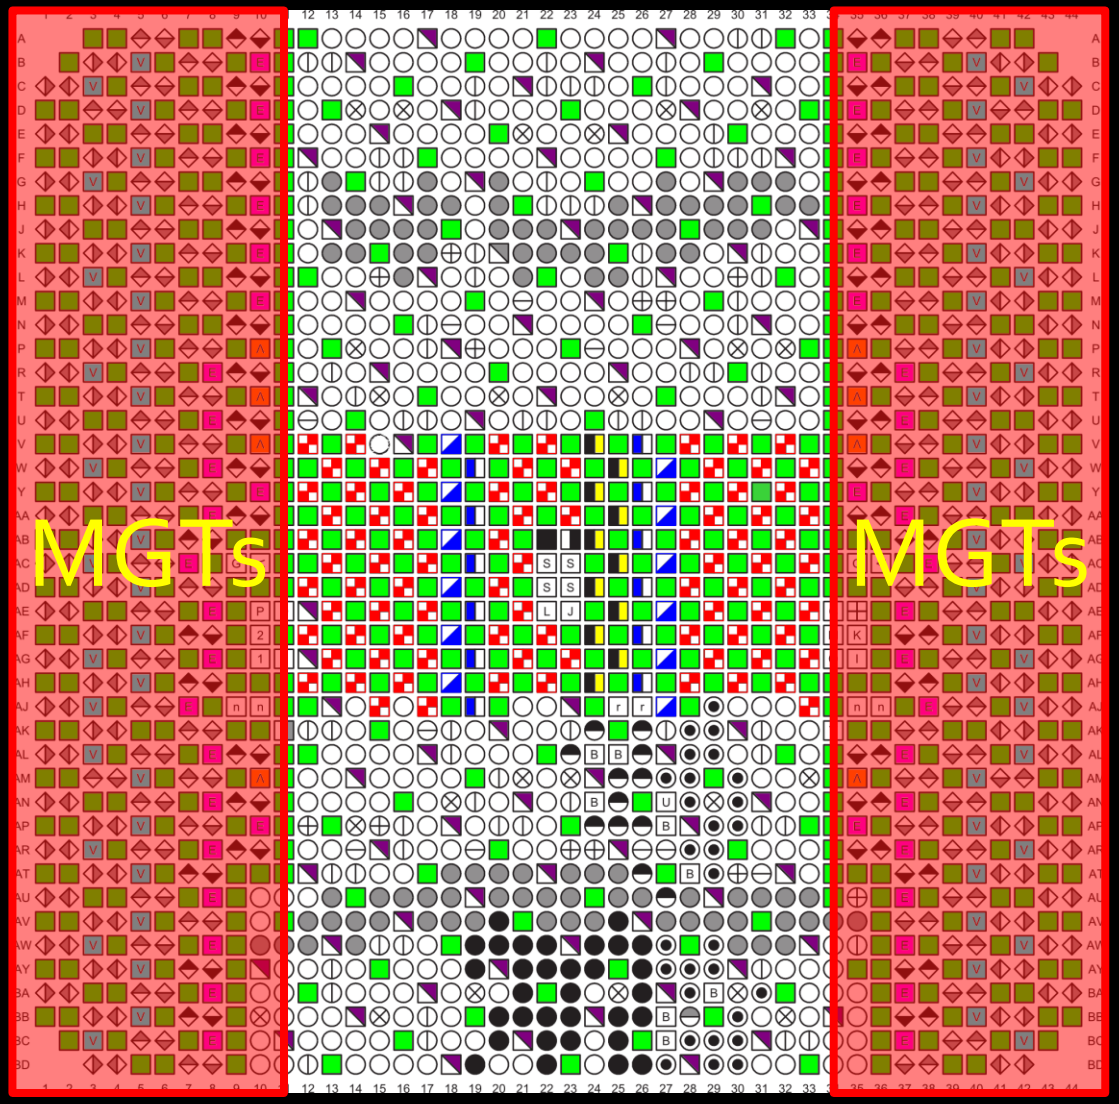

MGTs

- The noise limits and tolerances on the MGT power supplies are tight!
- There are also constraints on voltage drop across the chip
  - Ohmic losses in the planes required that the regulators are as close as possible to the MGTs
  - There are two banks of MGTs
  - Each needed powering independently
- Space constraints also mean that the 10Gb/s signals run underneath the regulators
- Noise was a concern…

MGTs: Revision 0

- Favoured LTM4606 6A ultralow-EMI switch-mode regulators by Linear, since:
  - These are designed for transceiver applications
  - Are sufficiently low that the fitted on the bottom of the board
  - We had used these successfully on the Mini-T5 card

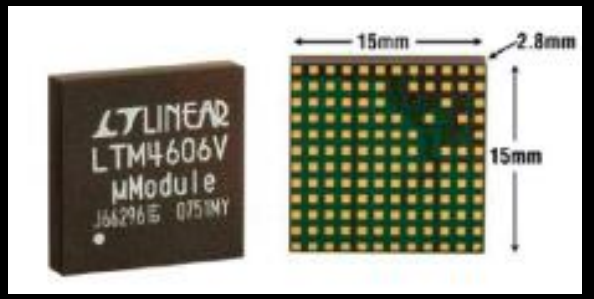

- Included a regulator on a test board to test for interference
- Also, taped a Samtec kapton cable to top of regulator and tested 10Gb/s signal integrity that way
- Good news – regulator had no effect on a 10G signal passing under it

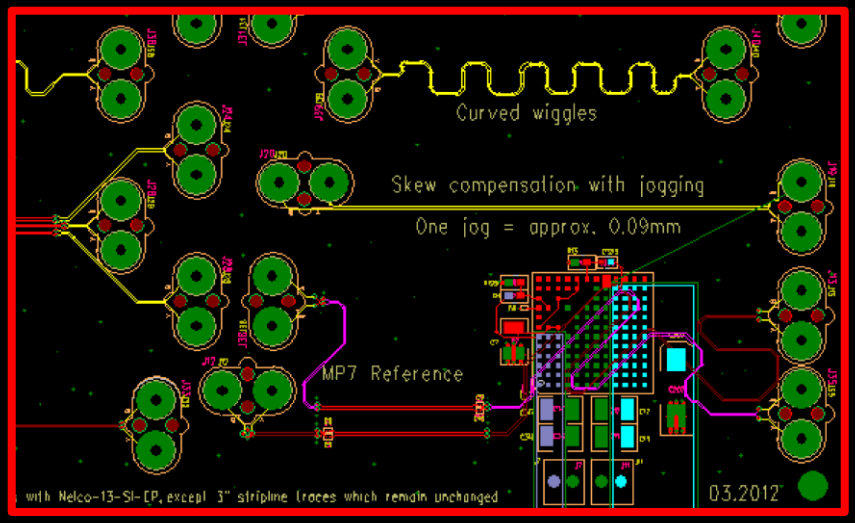

MGTs: Revision 0

- Card was designed before 7-series engineering silicon was available
- LTM4606 should have had ~35% headroom based on Xilinx's power estimator
- When card assembled and tested, power consumption 30% to 220% higher than Xilinx had predicted
- A lot of discussion with Xilinx engineers – found a lot of "features" that Xilinx weren't aware of. The price you pay for living at the cutting edge.

DFE: RXLPMEN = 0		24 Chans, Quads: 113-118 Theory XPE 14.2 Current (A)	24 Chans, Quads: 113-118 Theory & Errata Current (A)	24 Chans, Quads: 113-118 Measured (ES Parts) Card 2, 1mOhm Resistor Current (A)	24 Chans, Quads: 113-118 Measured (ES Parts) Card 2, 1mOhm Resistor Measured/Predicted Current	24 Chans, Quads: 113-118 Measured (Production) Card 4, 1mOhm Resistor Current (A)	24 Chans, Quads: 113-118 Measured (Production) Card 4, 1mOhm Resistor Measured/Predicted Current
MGTAVcc (V)	1.00	5.37	7.73	6.59	1.23	6.74	1.16
MGTAVtt (V)	1.20	1.94	2.73	6.17	3.18	5.45	2.81
Power (W)		7.70	11.00	13.99	1.82	12.78	1.66

Low Power: RXLPMEN = 1		24 Chans, Quads: 113-118 Theory XPE 14.2 Current (A)	24 Chans, Quads: 113-118 Theory & Errata Current (A)	24 Chans, Quads: 113-118 Measured (ES Parts) Card 2, 1mOhm Resistor Current (A)	24 Chans, Quads: 113-118 Measured (ES Parts) Card 2, 1mOhm Resistor Measured/Predicted Current	24 Chans, Quads: 113-118 Measured (Production) Card 4, 1mOhm Resistor Current (A)	24 Chans, Quads: 113-118 Measured (Production) Card 4, 1mOhm Resistor Measured/Predicted Current
MGTAVcc (V)	1.00	4.07	6.11	5.46	1.34	5.29	1.30
MGTAVtt (V)	1.20	2.11	2.94	6.54	3.10	5.16	2.45
Power (W)		6.60	9.63	13.31	2.02	11.48	1.74

Greg Iles made a huge number of measurements

MGTs: Revision 0

- Very impressed with LTM4606
  - Rated 6A nominal
  - 8A peak
- But performed excellently even when run flat-out at 25% above its nominal rating

MGTs: Revision 1

- We were very concerned about the possibility of power supply changes introducing noise
- Several test cards were made to test alternative power supply designs on an Ro card
- Noise was measured both electrically and by its effect on the error-rate of the 10Gb/s optical links
- LTM4601 switch-mode regulator by Linear won the day:
  - Same size as the LTM4606 (although different footprint)
  - Similar external components
  - Simple replacement

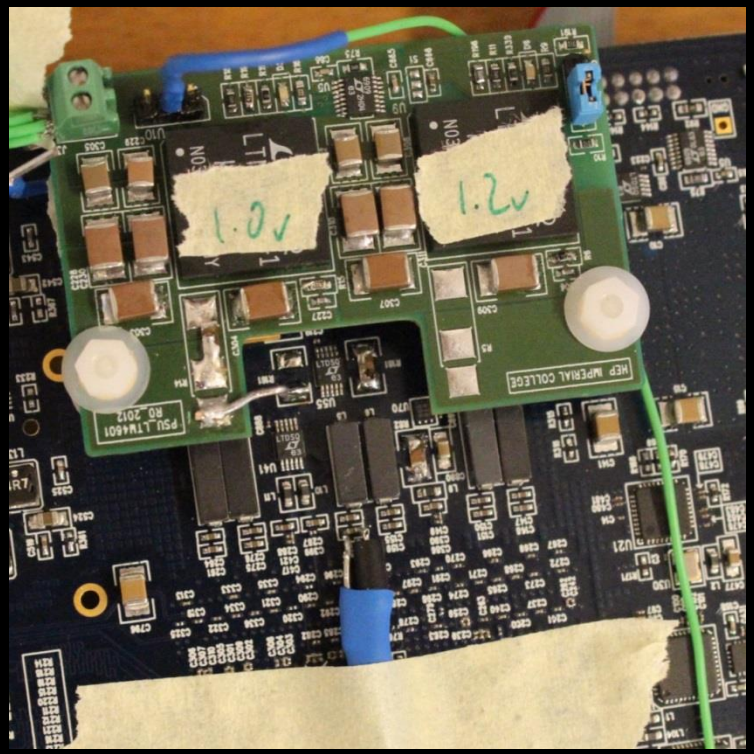

MGTs: Revision 1

- 485, RPBS7, 10G, QPLL, 24 links (one side),
- Plots show bathtub from 6 different links: 1 from each quad

Original PSU

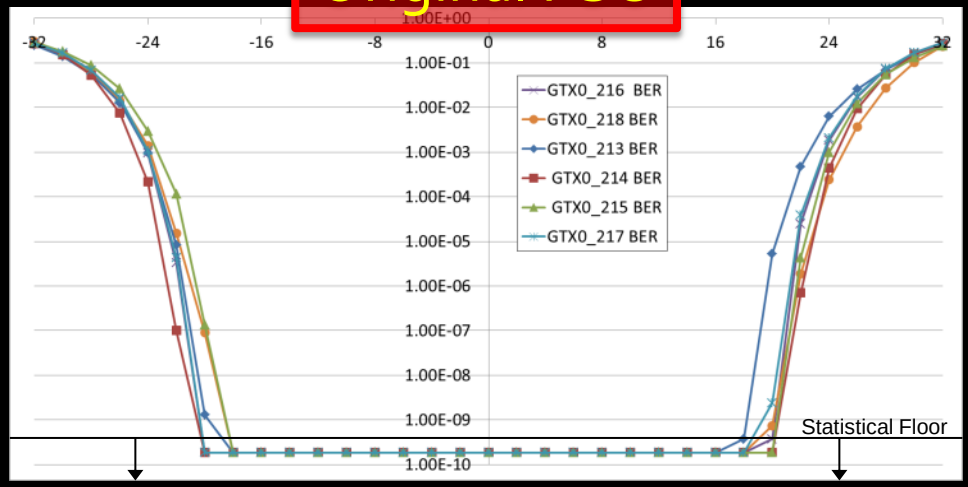

New PSU

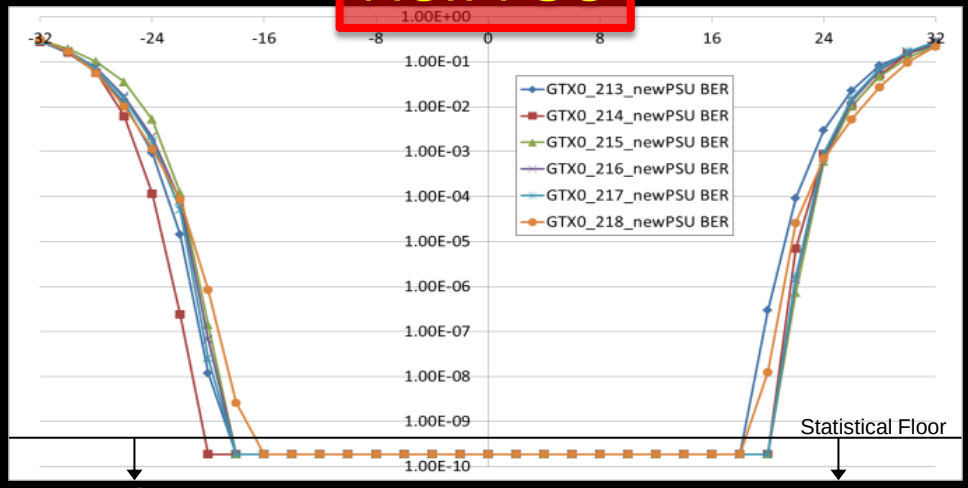

Magnus Loutit

Core power

- Large BGAs puncture the board with vast forest of vias
  - Increase the effective resistivity of the power planes
- Core power pins are at centre of BGA
- Use fills in three layers to get power into core to ensure sufficient current
  - Prefer layers with 1oz copper over ½oz copper layers
- Use remote sense, even though the distance is only a few cm

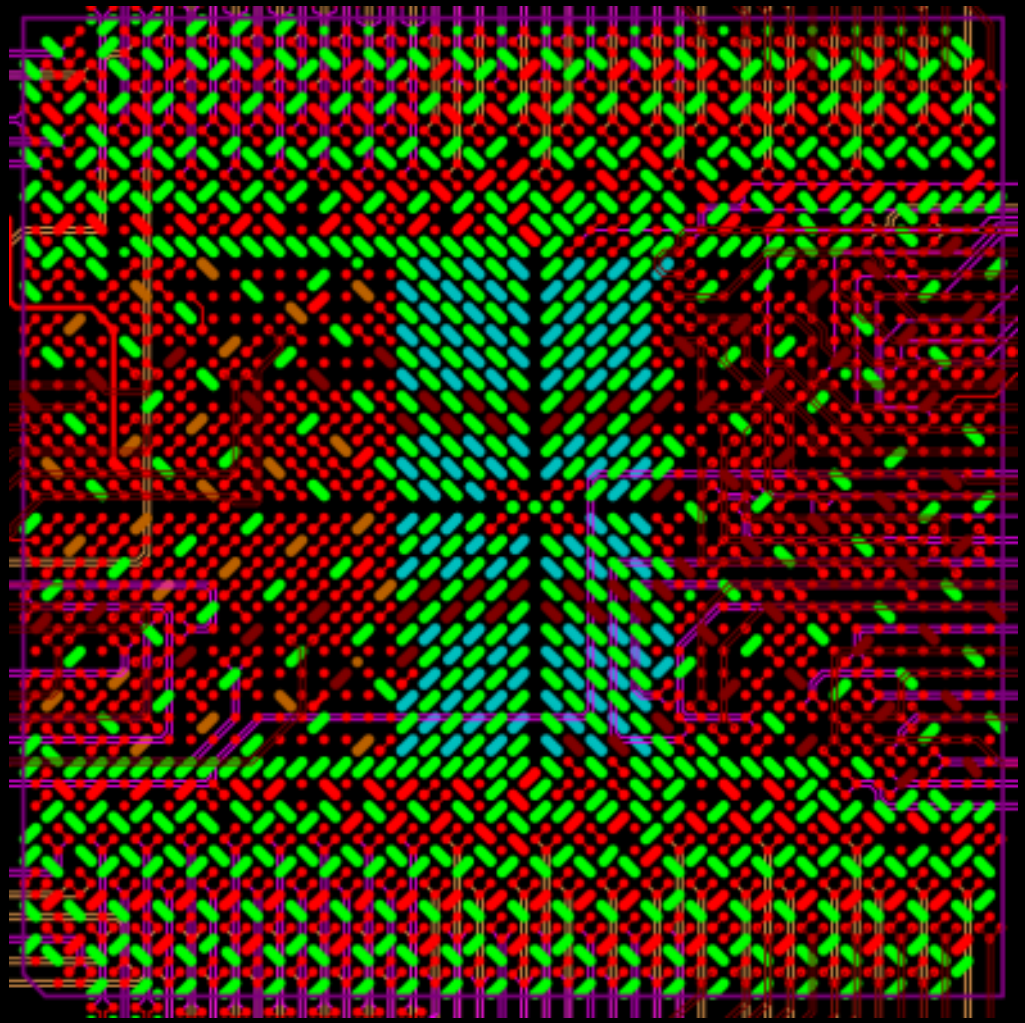

Power supply monitoring

- MP7 makes extensive use of LTC2990 four channel voltage/current/power and temperature monitor
  - Sub-millivolt resolution
  - 1% current resolution
  - 1% temperature resolution
  - 10-Lead MSOP Package – size is important in the space-constrained µTCA environment

- MP7 measures:
  - Both incoming supplies
  - ALL bulk supplies on the board
  - Subset of secondary/speciality supplies

- Sensors are distributed around the board allowing an approximate temperature profile of the board

Power consumption sensitivity

- Power consumed by the Virtex-7 can be extraordinarily sensitive to configuration flags
- Unused and non-optimally configured resources can contribute massively to power consumption. Not always immediately obvious.
- Need to consider the entire design (all configuration flags) before making statements about power consumption
- No substitute for hands-on experience

With great power come great heat dissipation…

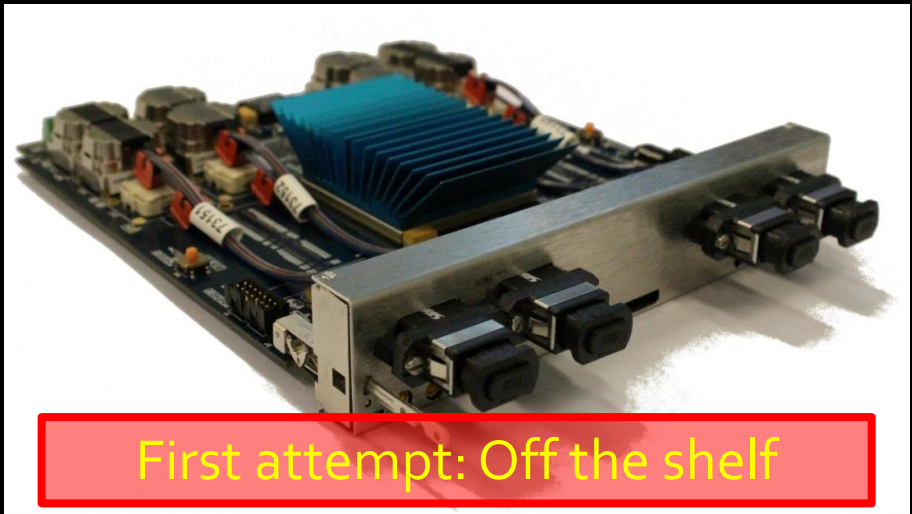

- With early revision of board, 48-link design hit thermal cut-out when on the bench (no fans). Reached 60°C in a crate.
- Designed a heatsink ourselves (old-school educated guesswork, no simulations) and prototyped in-house. Dissipated 40% more power than the off-the shelf part. Temperature didn't exceed 45°C in 48-link design.
- Production heatsinks manufactured externally and anodized (increase radiative transfer by further 25%)

Conclusions

- The compact nature and excellent performance of Linear's switch-mode modules make them an excellent match for 7-series FPGAs and the µTCA environment, where space is constrained
- Xilinx power estimator should be considered exactly that… AN ESTIMATOR
  - Don't rely on it being correct
  - It is no substitute for hands-on experience
- 7-series FPGAs have a large number of configuration flags and it is not always immediately obvious how these will affect power consumption