BUILD FARM AGAIN

Kernel Recipes 2017
Willy Tarreau
HAProxy Technologies

WHAT'S NEW SINCE KR2016 ?

BUILD FARM PROTOTYPE

... as seen by Frank Tizzoni

Well, more like this in fact

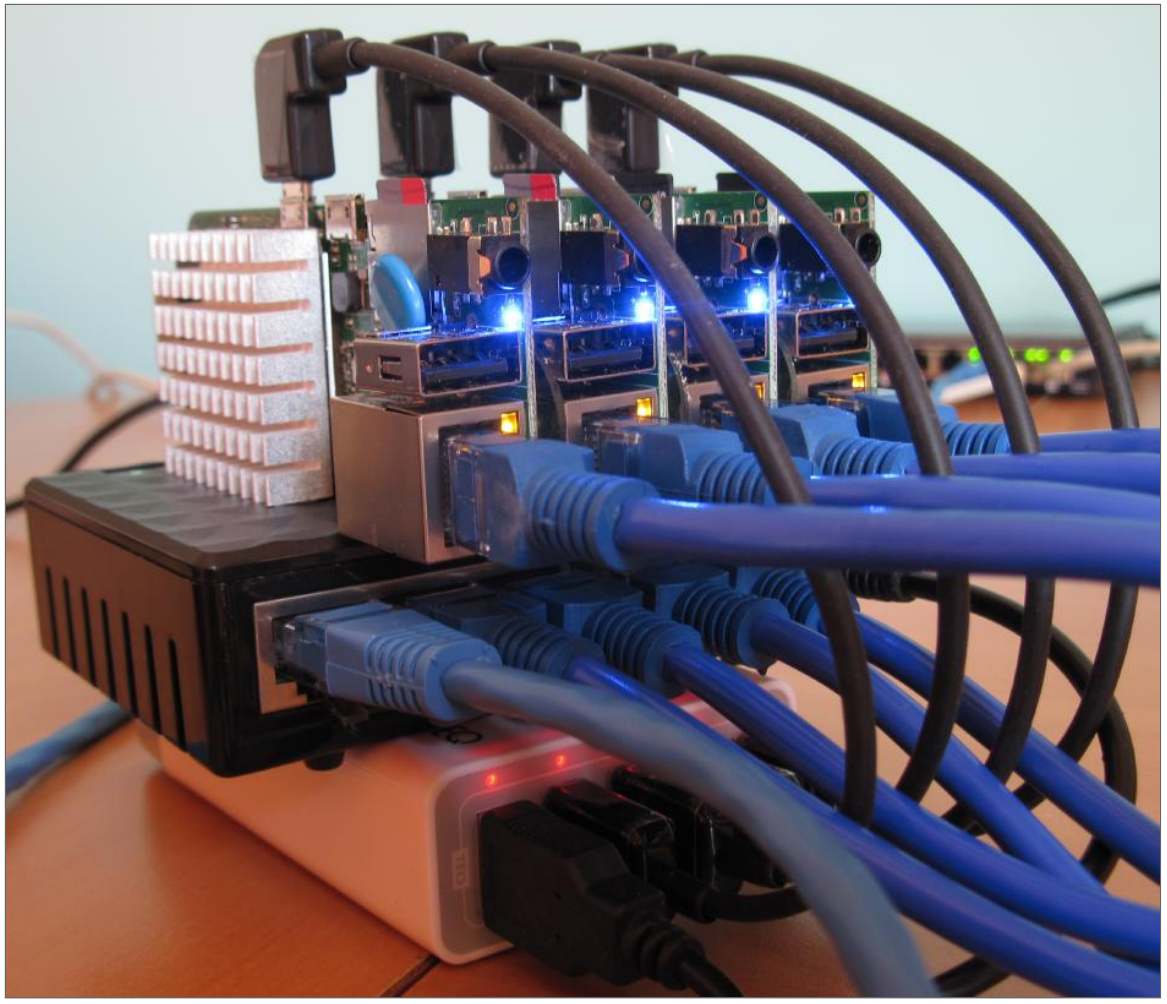

SHORTCOMINGS

- unreliable hardware (low-end HDMI sticks)
- power supply issues
- heating issues
- boots from Micro-SD (sensitive to wandering fingers)
- too many moving parts, no screws

STRONG POINTS

- RK3288's still unbeaten performance (quad-A17 @2GHz)
- even benefit small projects (haproxy: ~80 files)

BOARD ISSUES NOW FIXED

- Finally got 10 MiQi boards (thanks HapTech)
- These ones are even faster (dual channel DDR3-1600)

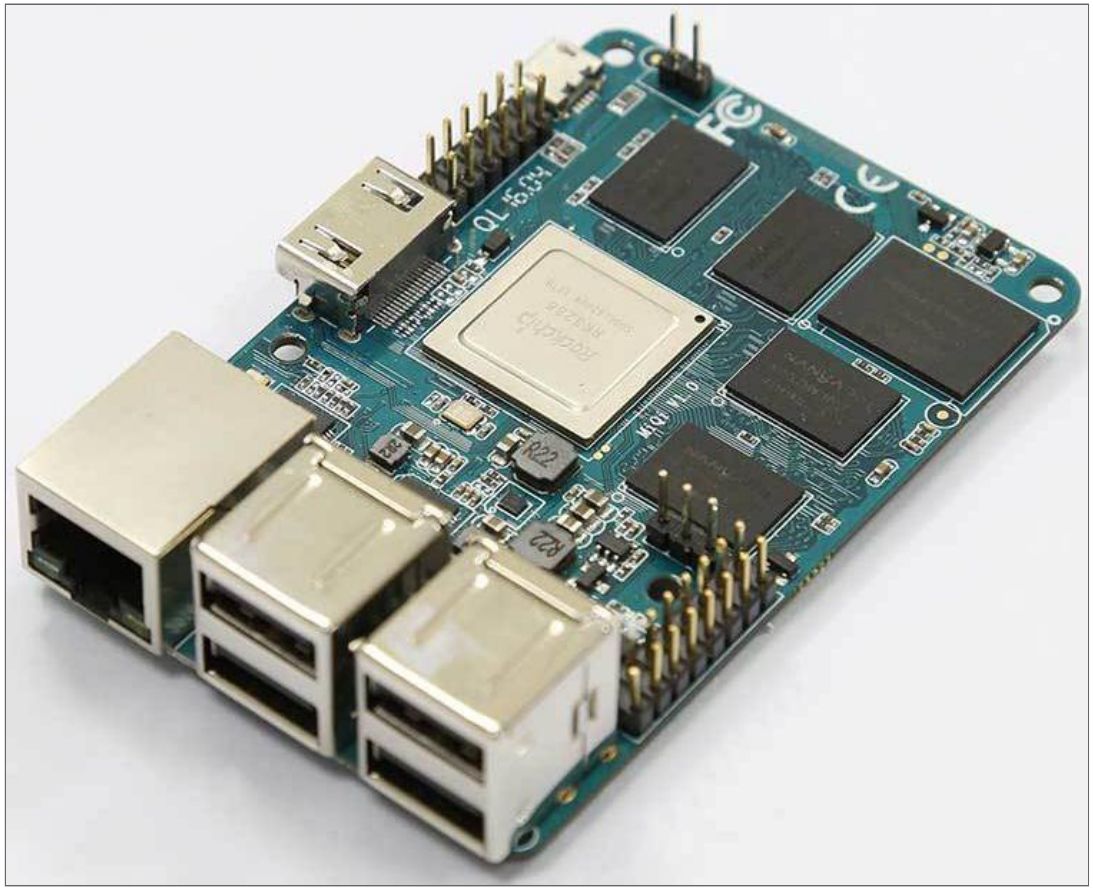

REMAINING/NEW ISSUES

- Find a better way to tie them together
- Still hard to attach a heatsink (no hole)
- Power losses in micro-USB (~0.3V @2A, peaks to 4A)
- Use less switch ports

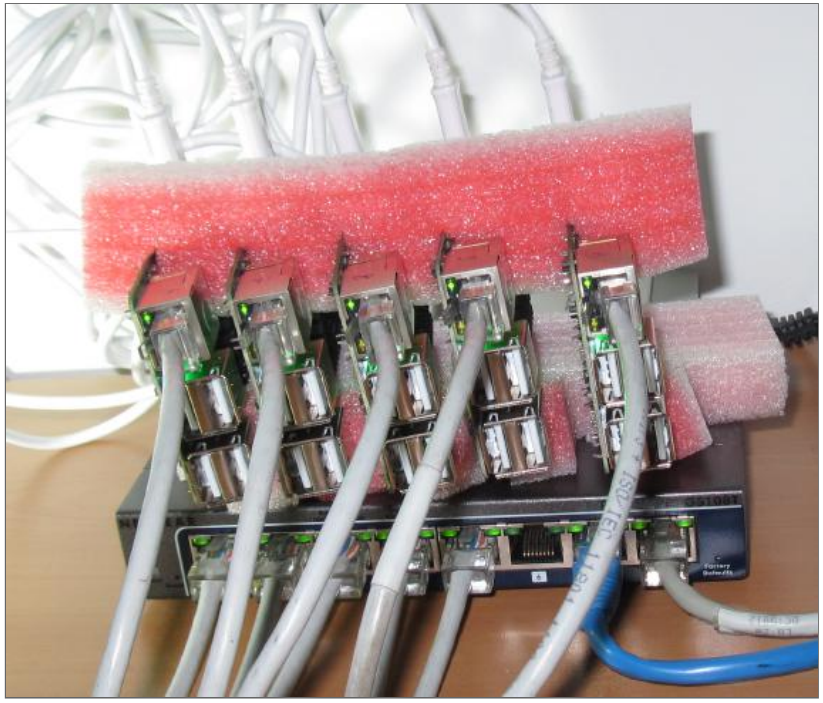

TYING THE BOARDS TOGETHER

Used M3 22mm spacers... and killed a LAN led!

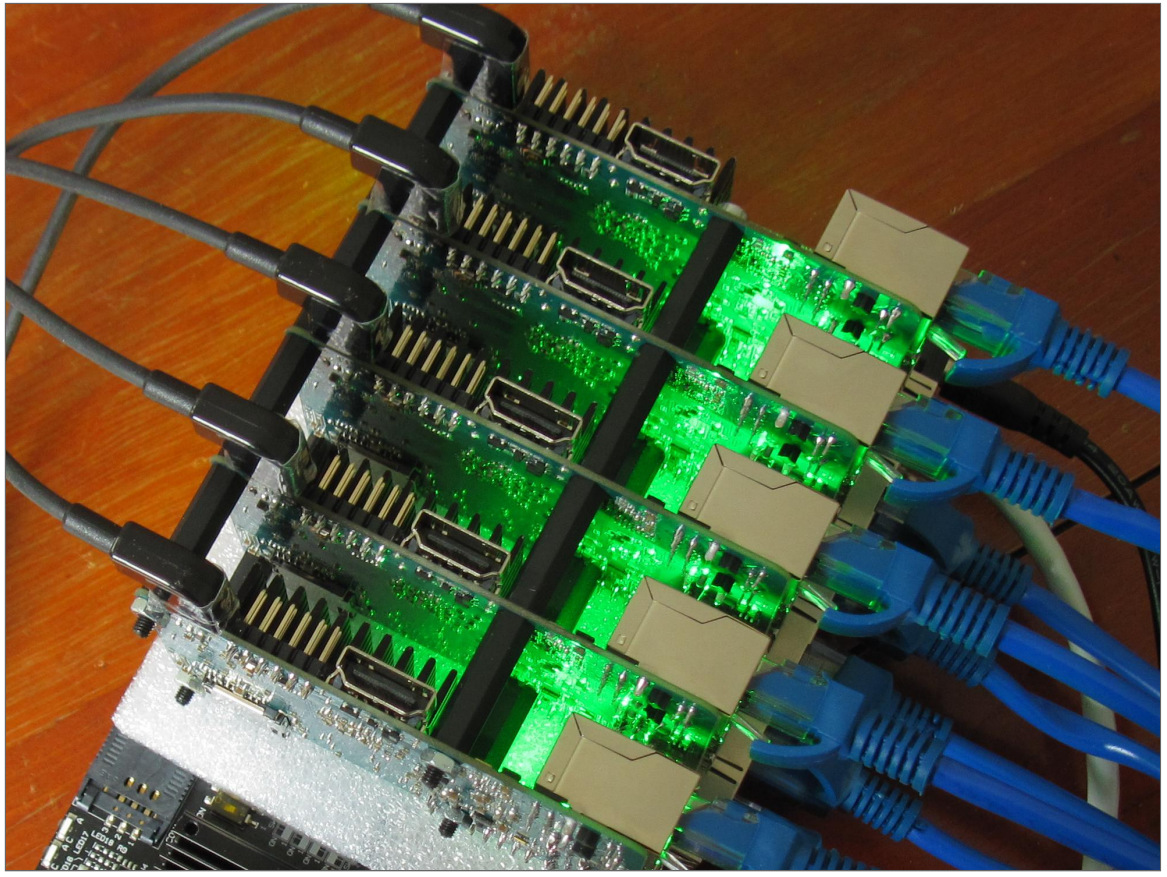

FIXING THE HEATSINK

Used 15mm wide 3M thermal tape (good but not perfect)

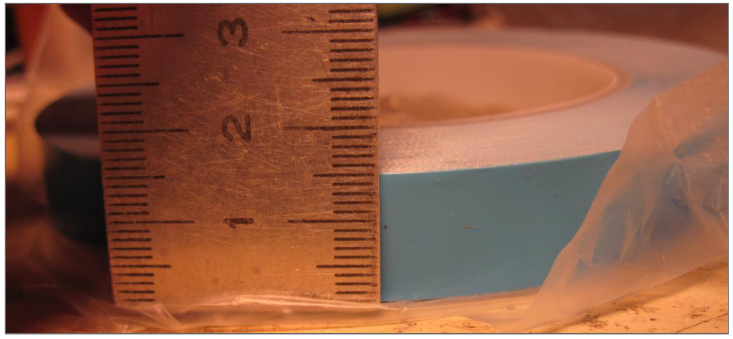

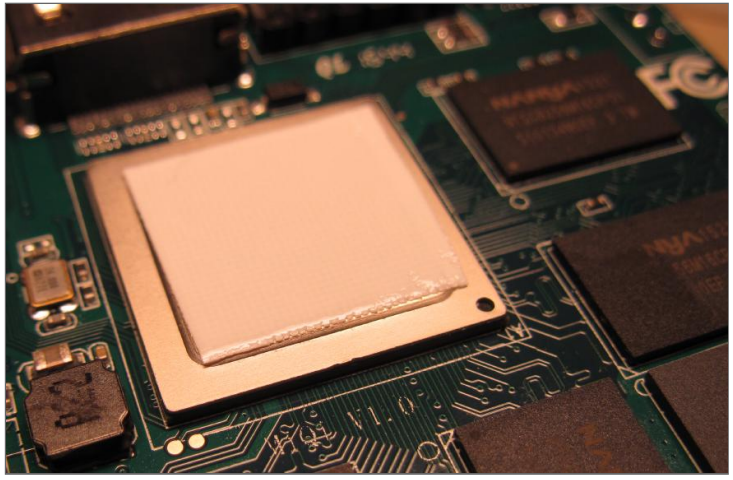

POWER LOSSES

Simply soldered thick wires

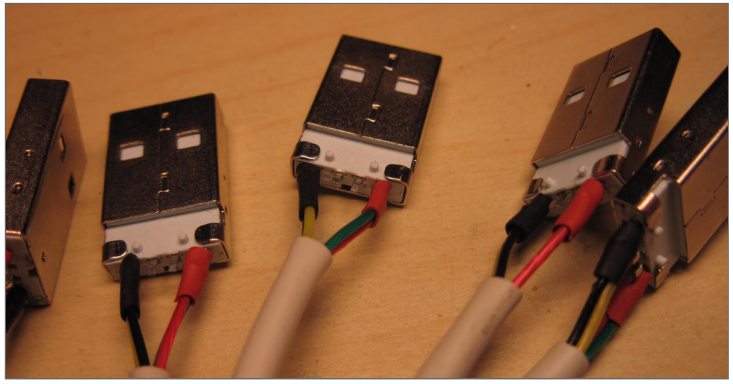

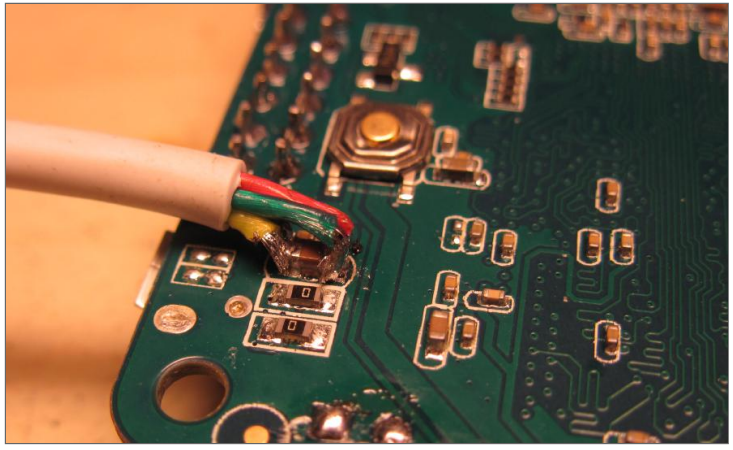

SWITCH PORTS

- I had a seldom used Clearfog-A1 board with 7 GigE ports
- Not powerful enough to serve as a build node (dual-A9)
- But interesting enough to open new perspectives :-)

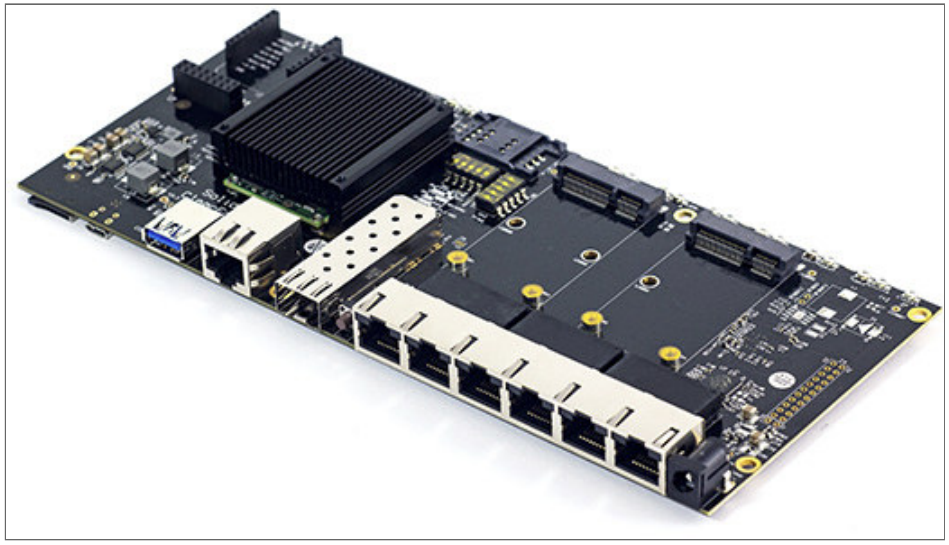

DISTCC

- Need a huge amount of parallelism to saturate 6 CPUs / 24 cores (60-100 processes)
- upgraded to latest -git to get more flexibility (eg: #local cpp jobs)
- enabled LZO source compression to reduce upload time
- had a hardcoded limit of 50 parallel jobs in the lock file lookup!
- used to disable remotes for -Wa<*anything*>

DISTCC USAGE

```
DISTCC_SKIP_LOCAL_RETRY=0 \
DISTCC_FALLBACK=0 \
DISTCC_BACKOFF_PERIOD=0 \
DISTCC_PAUSE_TIME_MSEC=50 \
TMPDIR=/dev/shm \
DISTCC_HOSTS="--localslots_cpp=100 localhost/6 192.168.0.212/100,lzo"\
PATH=/f/tc/flx2/x86_64-gcc47_glibc218-linux-gnu/bin:$PATH \
make \
  -j 60 \
  CROSS_COMPILE=/tc/x86_64-gcc47_glibc218-linux-gnu/bin/x86_64-gcc47_glibc218-linux-gnu- \
  CC=/tc/distcc-wrappers/x86_64-gcc47_glibc218-linux-gnu-gcc
```

OBSERVATIONS

- very uneven load : 0 to 4 tasks per board
- lots of remaining idle due to upload + fork/exec time
- local machine often is highly loaded
- death of a machine causes some long timeouts
- adding/removing machines requires fiddling with environment vars

IMPROVEMENTS

- place haproxy in front (on the Clearfog) with leastconn algorithm
- limit #of concurrent connections per CPU to #cores in haproxy
- double this limit on build node to preload pre-processed code
- use huge TCP buffers on the build node to buffer total input and output (pipelining)
- use remaining local CPUs as well

GAINS

- about 15-20% faster
- less idle
- less contention

COMPLETE X86_64 BUILD OF 3.10.106

(make -jXX allmodconfig)

- local (core i7-6700k - 4C8T @4.4 GHz) : 13 minutes
- local + 6 MiQi : 4m45, up to 120 files per second!

HAPROXY ON THE LOCAL MACHINE

- not as much interesting (contention with gcc)
- higher latencies again
- not convenient for me as haproxy developer :-)

SEEKING SMALL MACHINES TO HOST HAPROXY

Found EdgeRouter-X for ~$50

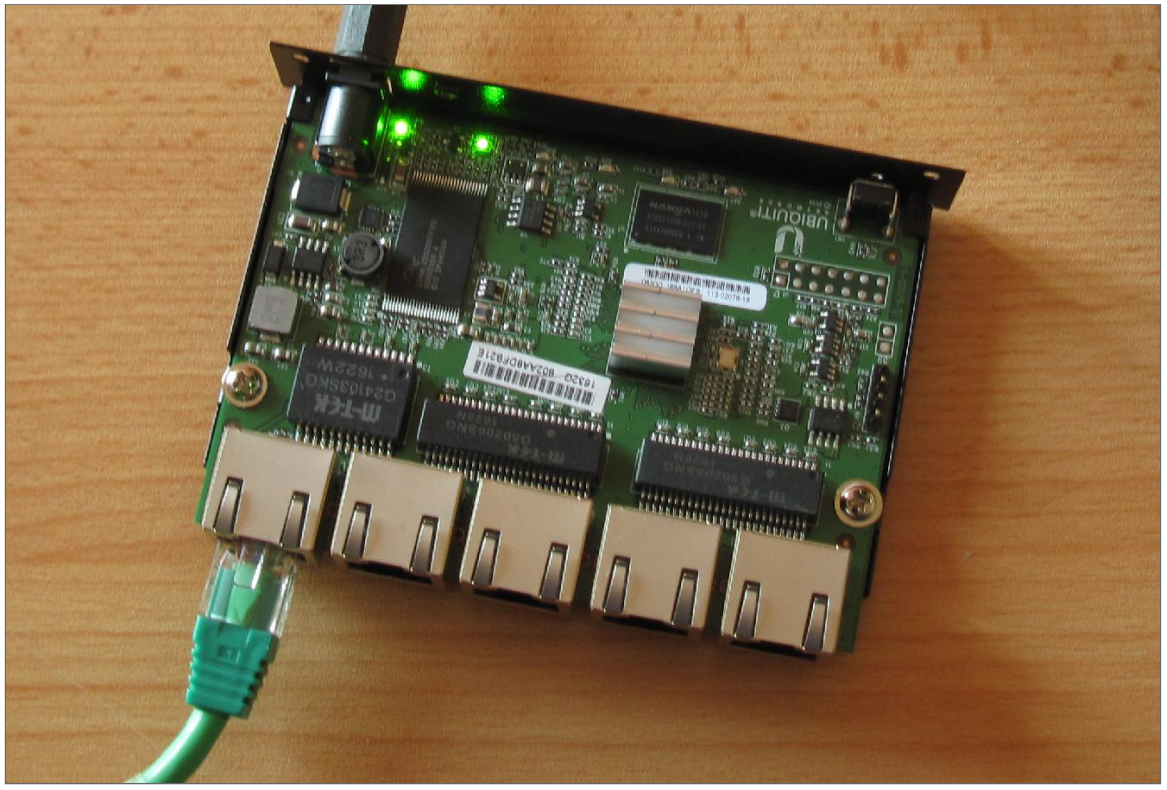

EDGEROUTER-X

- 5 GigEthernet ports
- dual-core 880 MHz MIPS (MT7621N)
- 128 MB RAM
- 128 MB flash
- running "Linux" (well, Vyatta + out-of-tree kernel in fact).
- open SSH with root access and R/W file system
- Consumes less than 4W.
- haproxy cross-compiled, sustaining 500 Mbps of traffic
- *need to upgrade the firmware though*

OFFICE FARM

OFFICE FARM (PLUGGED)

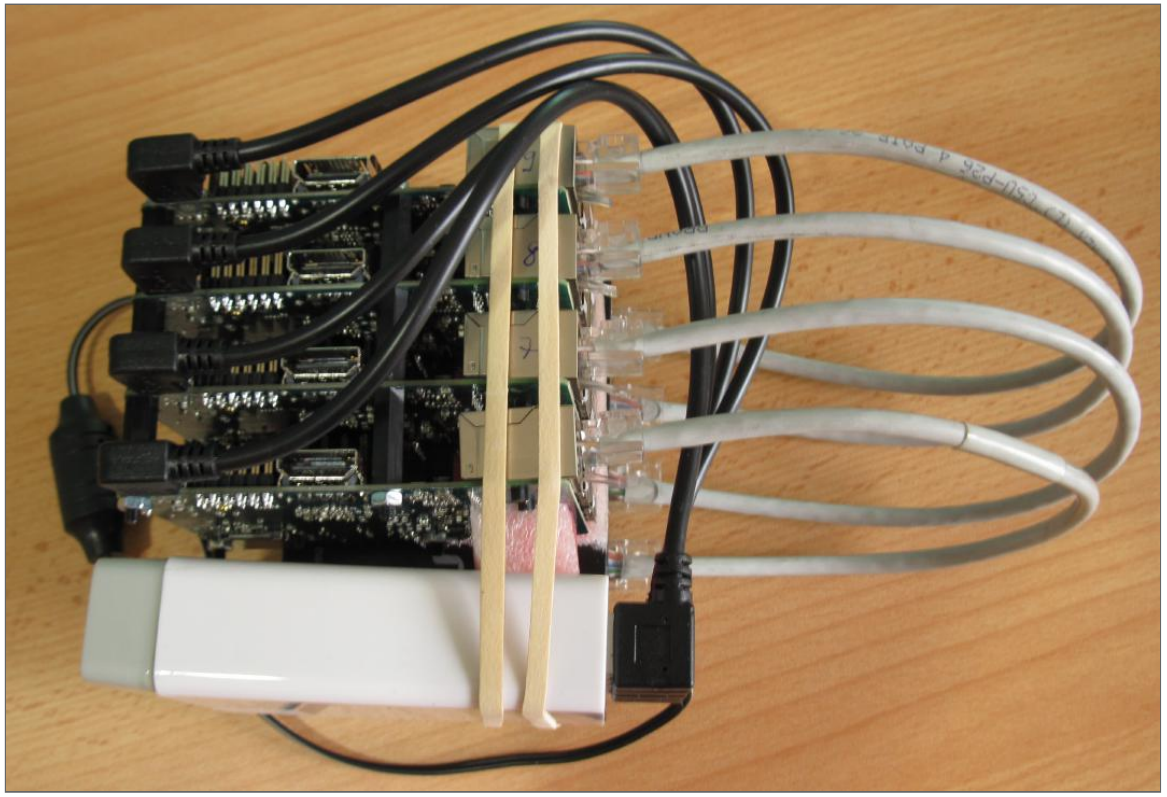

NEW ISSUES

- power supply sometimes cut-off at ~45W instead of 50W.
- some instabilities at full load (voltage drops by 0.1V at 2A)
- small PSU patch to get some margin (75W now)

PSU PATCH

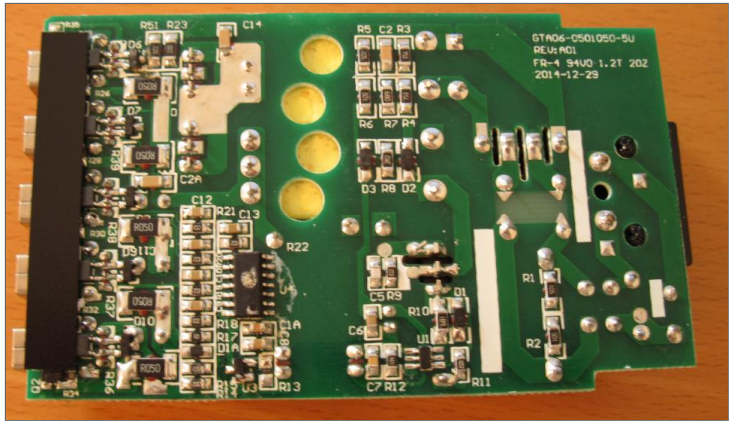
GTA06-0501050-5U
REV:A01
FR-4 SHUO 1.2T 202
2014-12-29

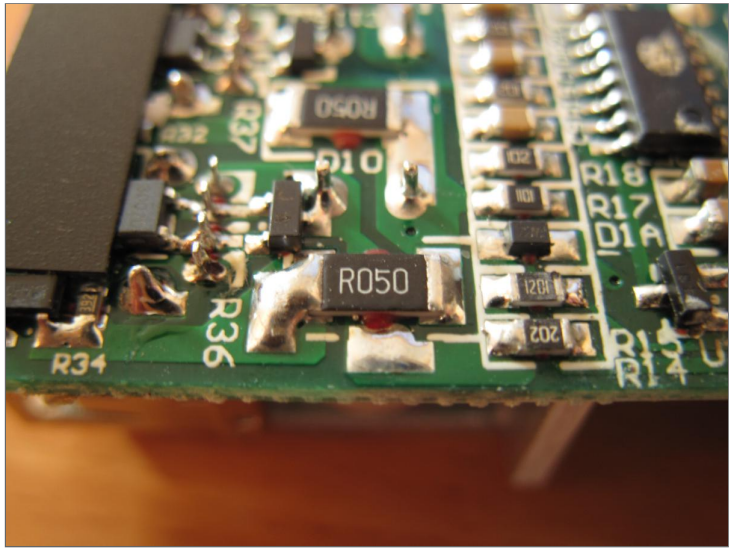
R050
R10
R18
R17
D1A
R050
R36
R34

MONITORING

Created a new "led-activity" LED trigger to monitor the CPU usage level of each board and ensure they're still alive.
Submitted a few times late, now accepted for 4.15.

BENEFITS

- office: 4 boards compile haproxy in 3s instead of 11s
- up to 200 full builds a day, saves less than half an hour
- real gain: high incentive to build and run after any simple change
- on release days, builds ~15 ALOHA kernels during lunch hour.
- home: lots of builds during 3.10 backport sessions, saves hours on week-end.

FEEDBACK / LESSONS

- feedback sent to MiQi's vendor (mqmaker) about overall board design, very welcomed
- patches sent to distcc
- never believe USB power supplies' advertised power, count on about half.
- lots of work remains local, need a good build controller
- not all cores always in use, could be nice to use DVFS based on #cores in use
- no need for complicated / optimal setups, the savings are excellent already

ALTERNATIVES

- RK3399 ? (2xA72 + 4xA53) => bought a cheap one (H96 Max), porting attempts in progress
- UP Board (Atom x8350 4*1.68 or 2*1.92 GHz). Not as good but still very good.
- Odroid-MC1 (4*A15+4*A7) => should be even faster and cheaper, impatient to test. May require distcc patching.

ALTERNATIVES

eg: Odroid-MC1

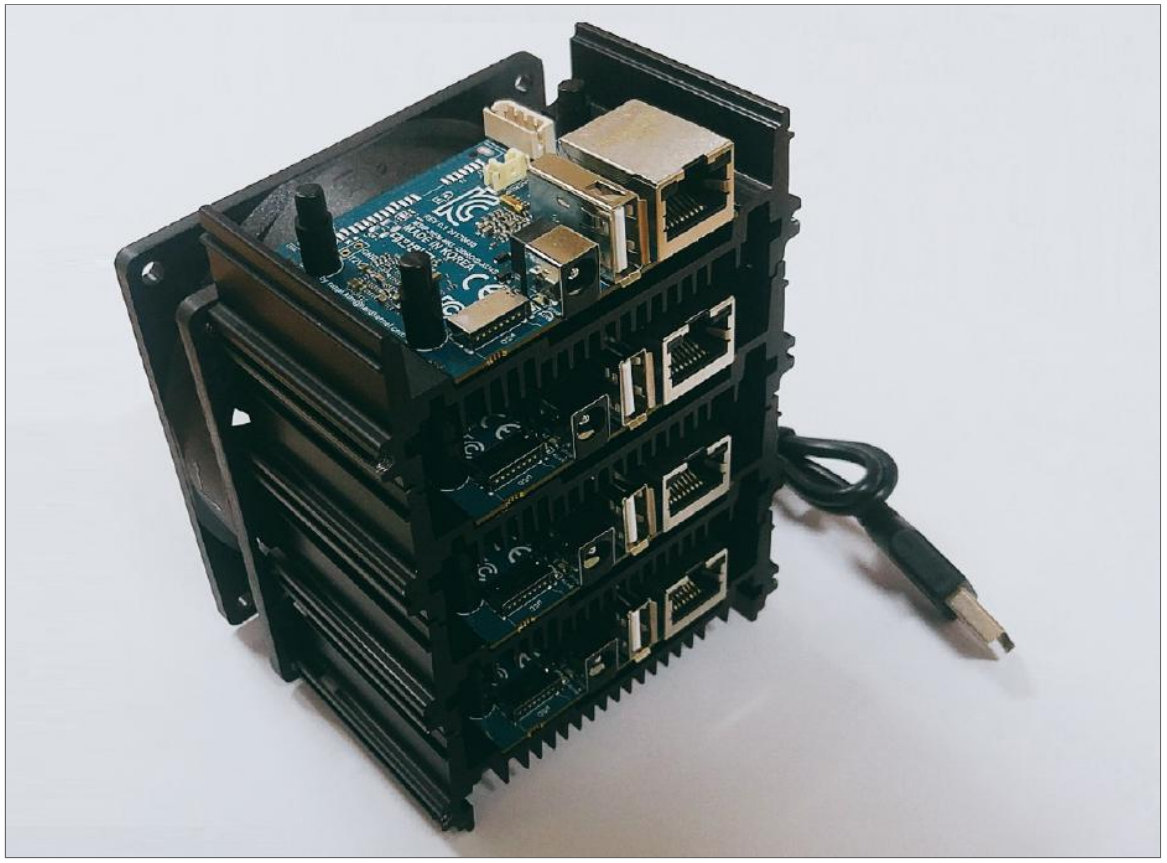

OTHER OPTIONS

- distcc's pump mode => too complicated to set up during a few tests, not convinced about the result
- icecc instead of distcc => didn't find it convenient in this context

LINKS - CONCLUSION

- distcc: https://github.com/distcc/distcc
- Odroid-MC1: http://www.hardkernel.com/
- mqmaker forum post : https://forum.mqmaker.com/t/miqi-based-build-farm-finally-up-and-running/605/24
- wiki : http://wiki.ant-computing.com/Choosing_a_processor_for_a_build_farm

That's all!