Ethernet Cable Diagnostics using Netlink Ethtool API

Andrew Lunn

andrew@lunn.ch

LPC 2019

THE Sun
NEWS WEBSITE OF THE YEAR

THAT'S JUST NUTS!

Pesky squirrels gnawing through superfast BT broadband cables in woodland because they taste of nuts

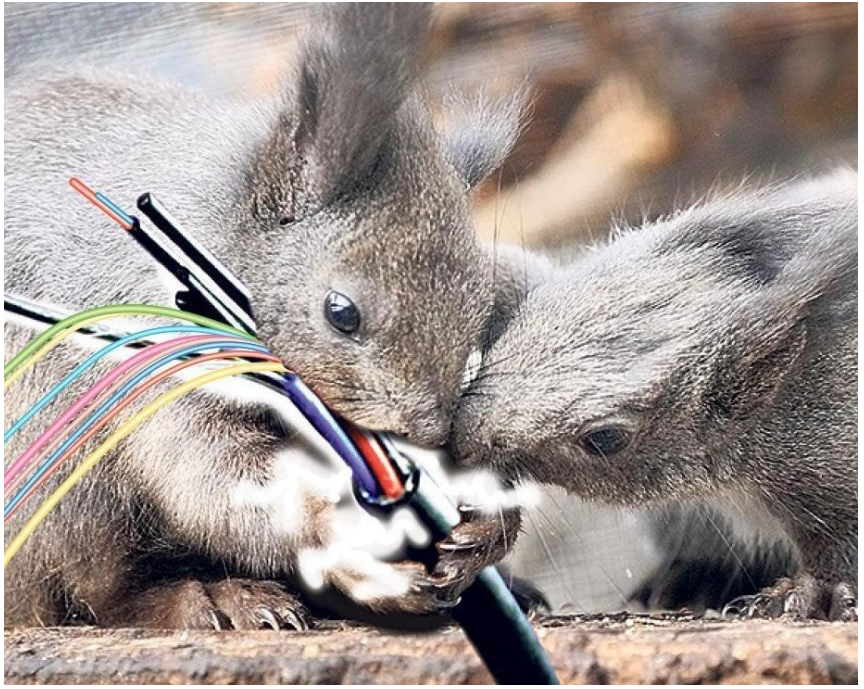

Structured Wiring: Where is the Fault?

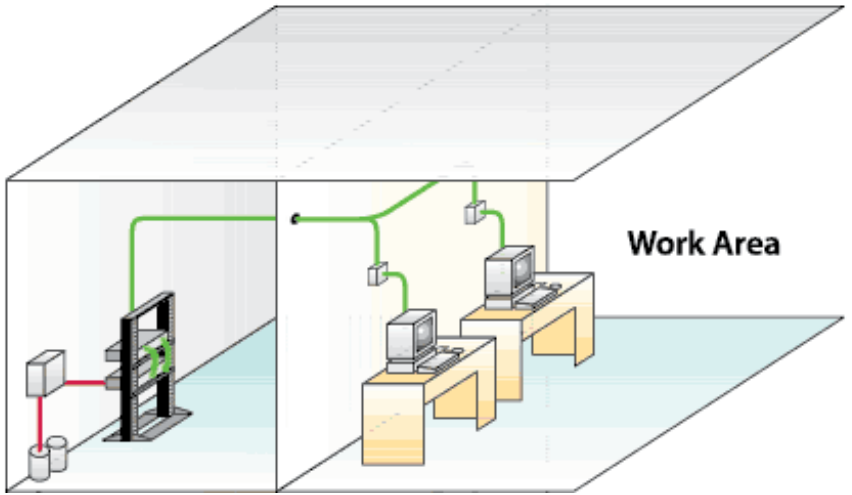

- PC to wall patch cable?
- In the wall?
- Closet patch cable?

Predict the Fault Before it Happens

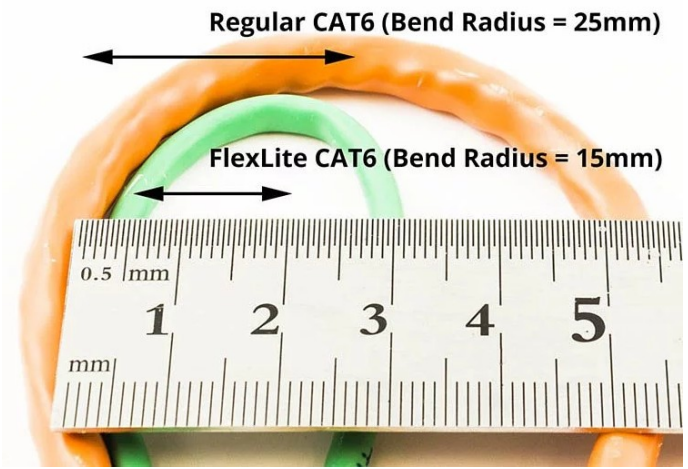

- Twisting, bending, vibration, water ingress, etc, can degrade and eventually break a cable or its connectors
- The diagnostics don’t need to be binary good/ bad.

Ethernet Cables

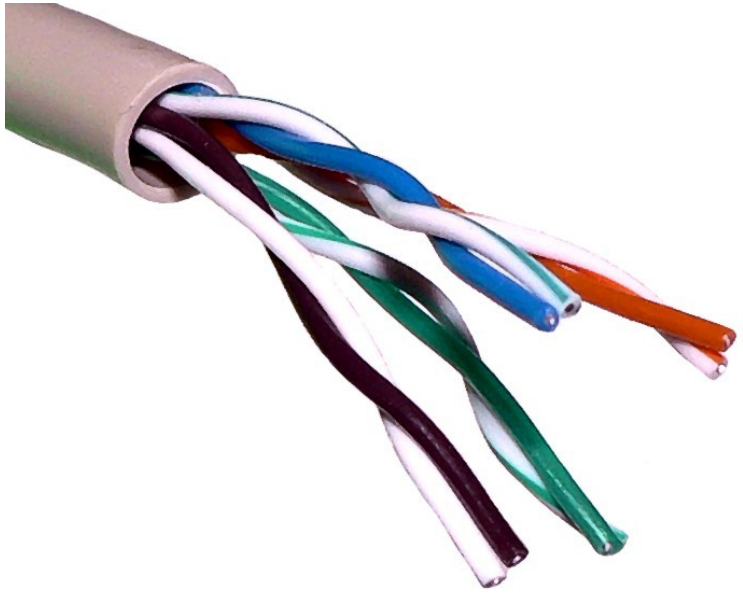

Wikimedia

- Two (10/100Mbps) or four (1Gbps-10Gbps), twisted pairs.
- Pair and Twist gives good immunity to cross talk and emissions/induction

Differential Pair

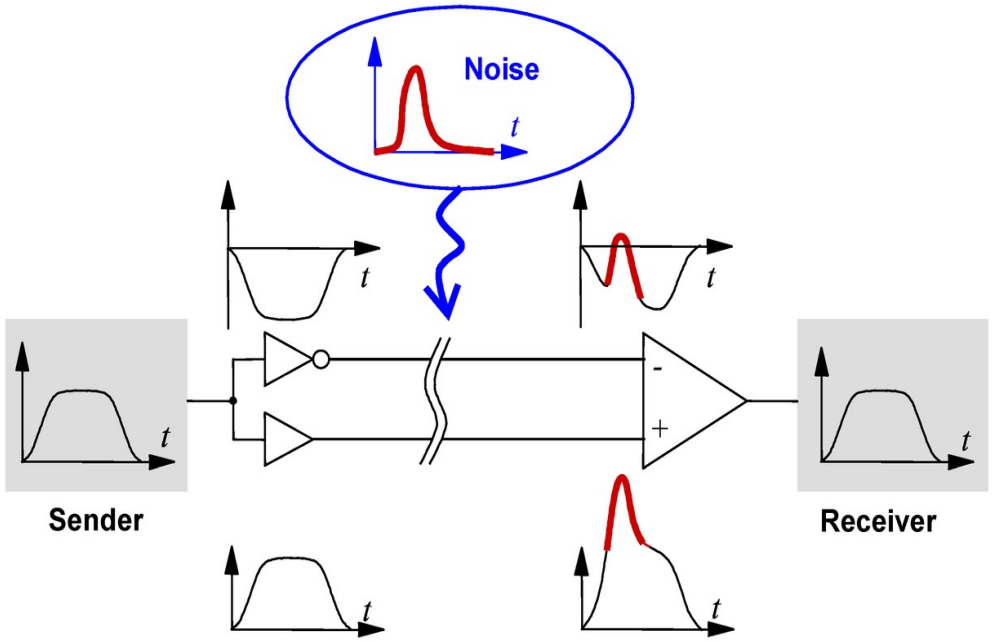

Wikimedia

- Transmit opposites on the pair
- Cable acts like a radio antenna
- Both cables pick up about the same noise signal
- Receiver subtracts out the noise

1687: Sir Isaac Newton and the Speed of Sound

https://mountainmystery.files.wordpress.com

- Newton measured the speed of sound by timing the echo of a clap, over a known distance.
- Maybe the beginning of Time-Domain Reflectometry, TDR.

Cables and TDR

- The same principle works for an electrical pulse sent down a cable.
- Moves down the cable at about $280*10^6$m/s.
- The pulse is reflected by changes in impedance in the cable.
- Measure the time of the reflect, derive the distance.

Open Ended Pair

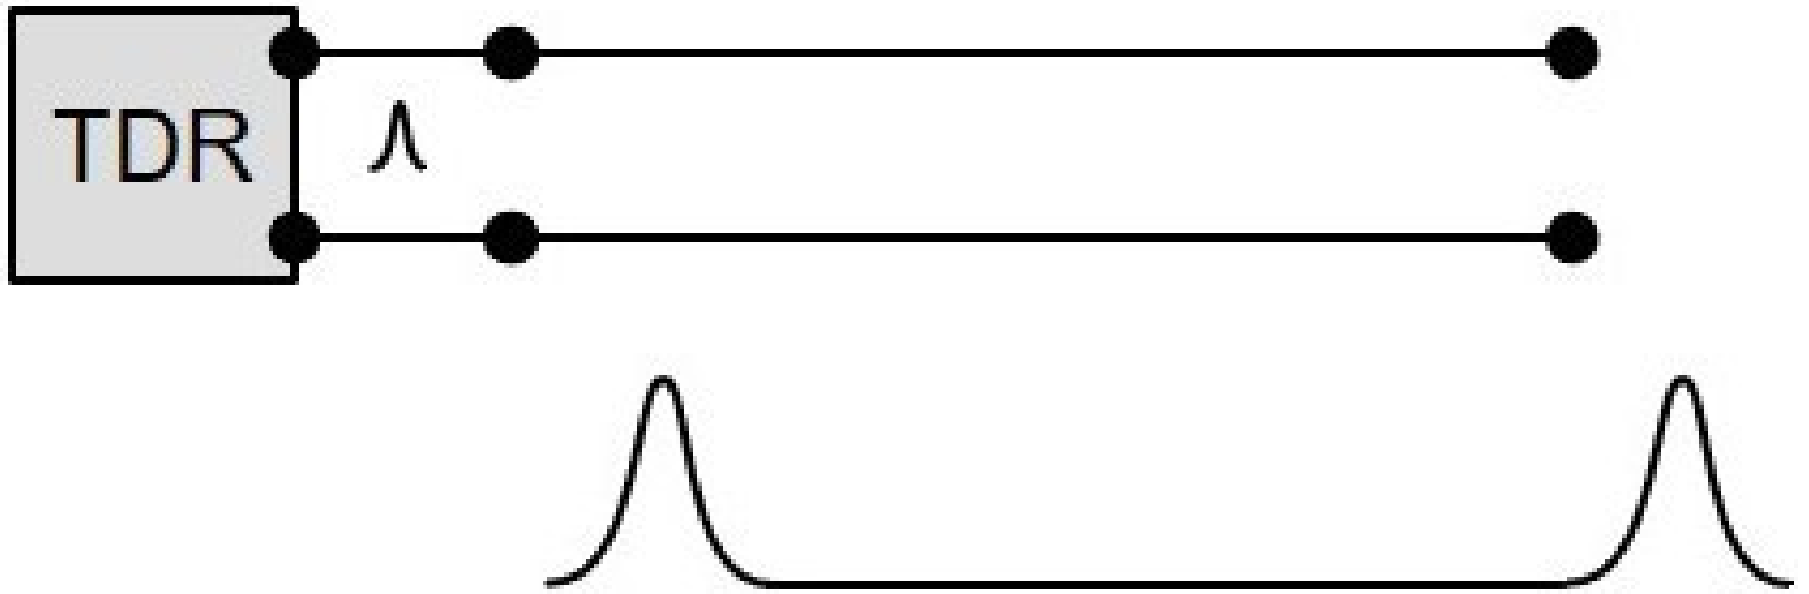

Shorted Pair

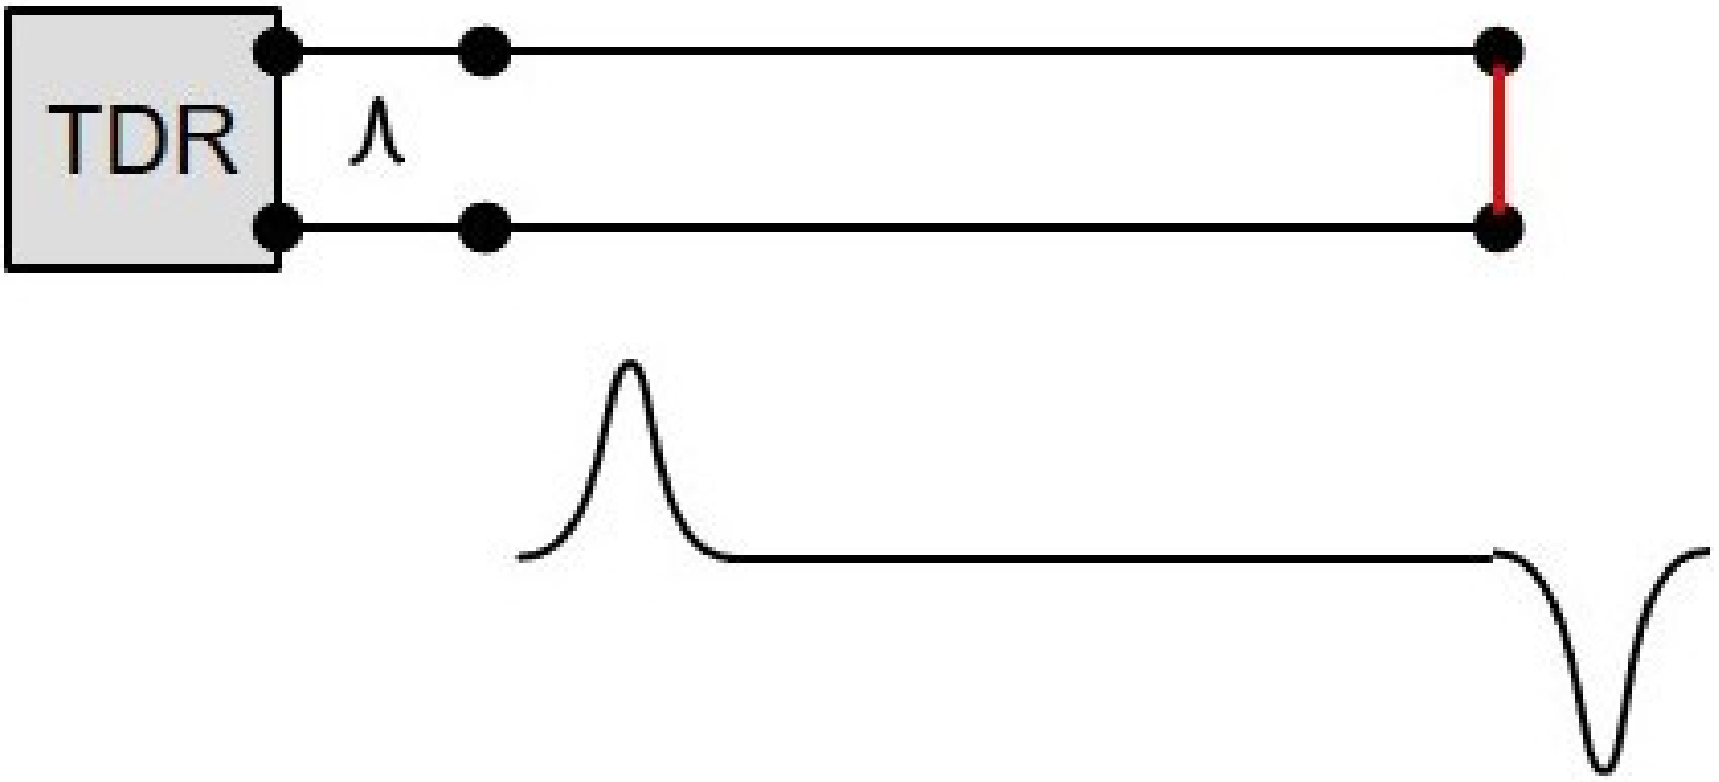

Break in one cable

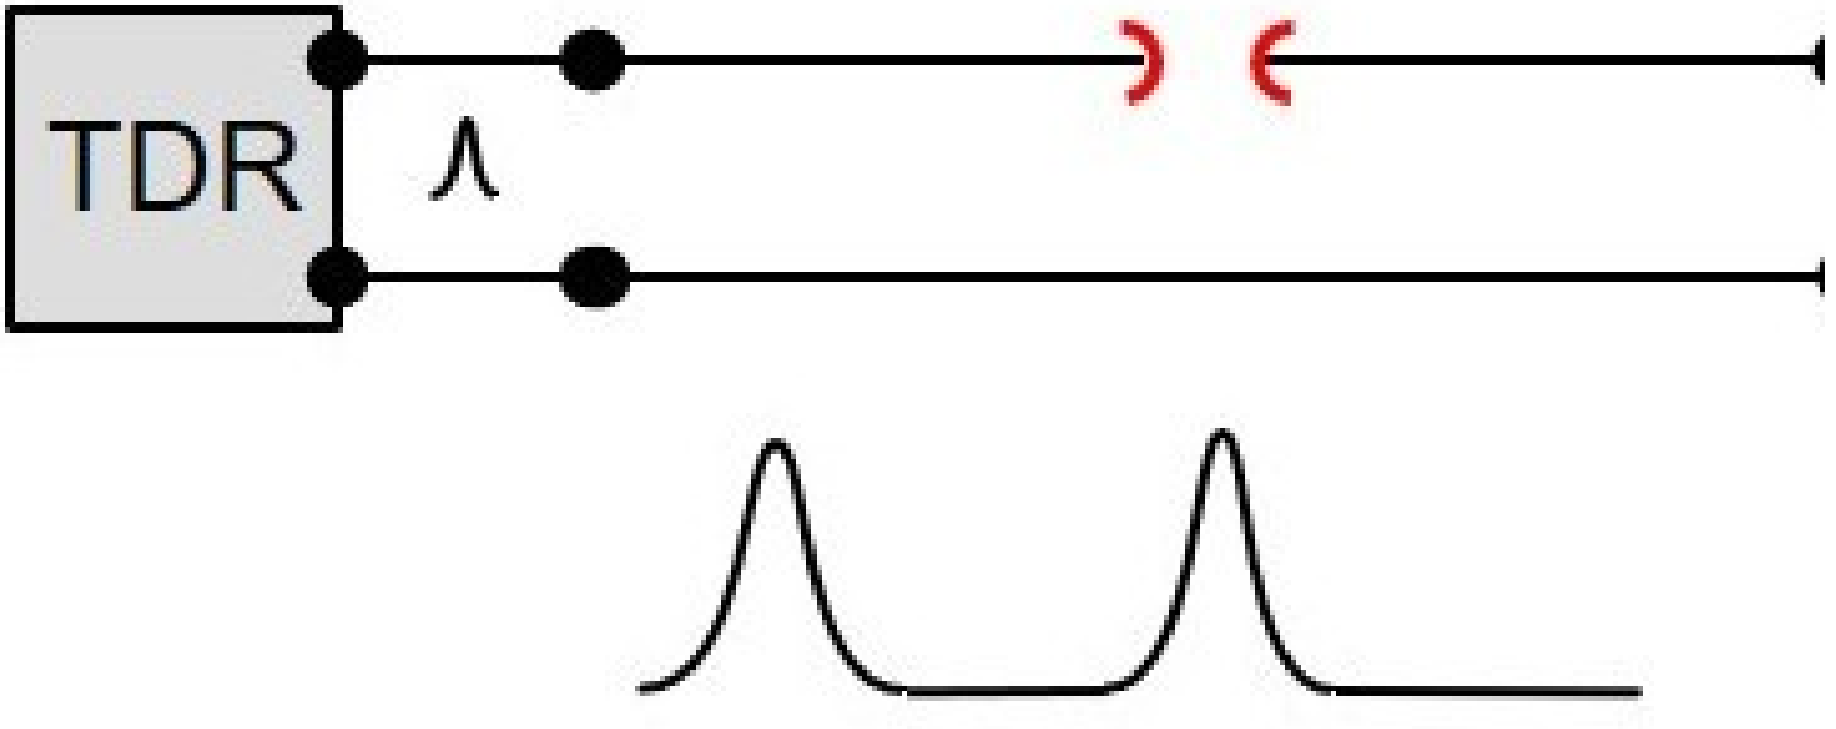

Everything O.K.

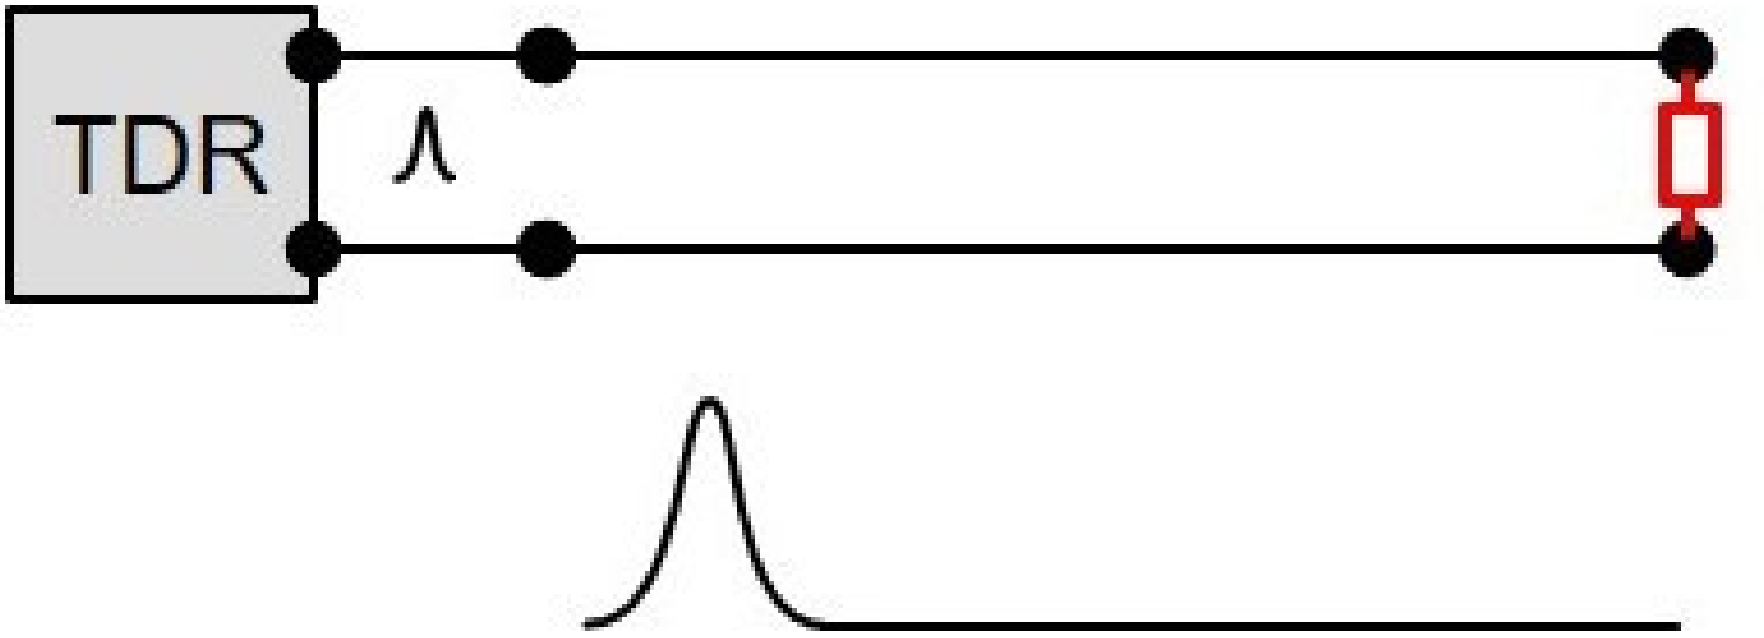

Problems Plumbing this into Linux

- No two PHYs are the same
  - Different fault codes
  - First fault vs 4 faults
  - Raw TDR data
  - 1, 2 or 4 pairs
  - Pulse configuration
- Does not fit well into an ethtool IOCTL call

Challenges

- Slow, a few seconds to run diagnostics
  - ethtool holds the RTNL while calling driver
  - Not acceptable to hold RTNL this long
  - Possible race conditions if we don’t hold RTNL
  - ethtool IOCTL has no async mechanism

Netlink Ethtool

- Michal Kubecek <mkubecek@suse.cz>
- This is first part of netlink based alternative userspace interface for ethtool. The aim is to address some long known issues with the ioctl interface, mainly lack of extensibility, raciness, limited error reporting and absence of notifications. The goal is to allow userspace ethtool utility to provide all features it currently does but without using the ioctl interface.

NL ethtool provides the missing Pieces

- Netlink ethtool
  - Flexible selection of result attributes
  - RTNL held during get/set
  - Spontaneous notification messages to user space
    - Trigger diagnostics using set
    - Report results later, spontaneously
- Still some potential for race conditions

PHY State Machine

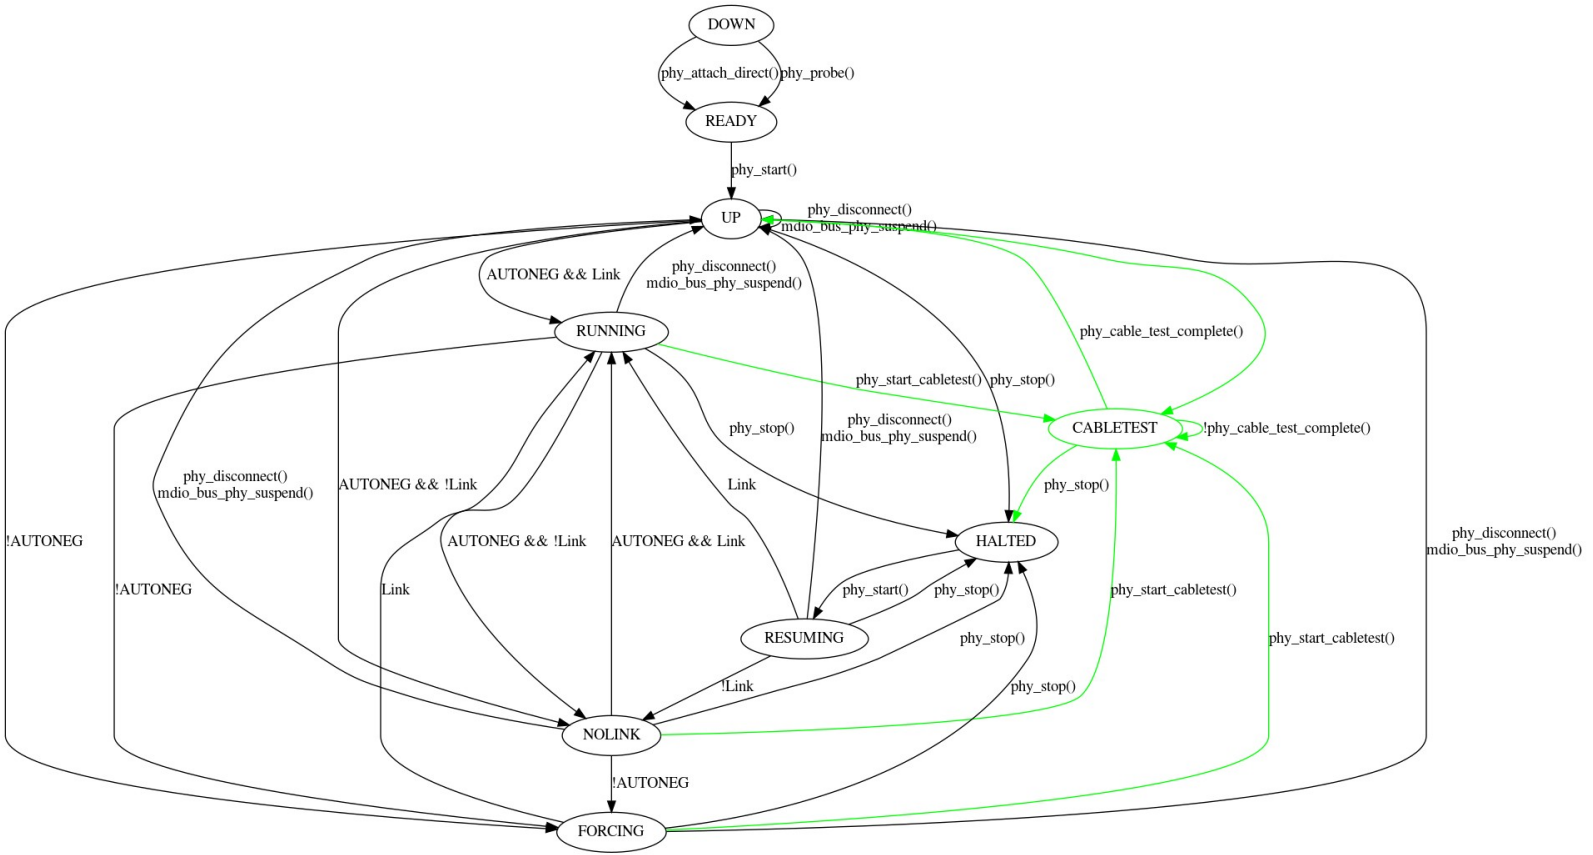

Flow

- ETHNL_CMD_ACT_CABLE_TEST
  - Start Cable test
  - Allocate SKB for results
  - Change to PHY_CABLETEST state
  - Return via ETHNL O.K
- Poll PHY/wait for interrupt until test complete
  - PHY uses helpers to fill SKB with results
  - Core multicasts ETHTOOL_A_EVENT_CABLE_TEST with results
  - Change to PHY_UP state to restart auto-neg, etc.

Driver API

```
struct phy_driver {
...
        /* Start a cable test */
        int (*cable_test_start)(struct phy_device *dev, int options);

        /* Once per second, or on interrupt, request the status of the
         * test.
         */
        int (*cable_test_get_status)(struct phy_device *dev, bool *finished);
}
```

```
int phy_cable_test_result(struct phy_device *phydev, u8 pair, u16 result);
int phy_cable_test_fault_length(struct phy_device *phydev, u8 pair,
                                u16 cm);
int phy_cable_test_amplitude(struct phy_device *phydev, int distance, u8 pair,
                             int mV);
int phy_cable_test_pulse(struct phy_device *phydev, int mV);
```

Ethtool(1) support

```
./ethtool --cable-test lan6
Cable test for device lan6.
Pair: 0, result: OK
Pair: 1, result: OK
Pair: 2, result: OK
Pair: 3, result: OK
```

```
./ethtool --cable-test lan2
Cable test for device lan2.
Pair: 0, result: Open Circuit
Pair: 1, result: Open Circuit
Pair: 2, result: Open Circuit
Pair: 3, result: Open Circuit
Pair: 0, fault length: 14.40m
Pair: 1, fault length: 15.20m
Pair: 2, fault length: 14.40m
Pair: 3, fault length: 15.20m
```

```
./ethtool --cable-test lan5
Cable test for device lan5.
Pair: 0, result: OK
Pair: 1, result: OK
Pair: 2, result: Short within Pair
Pair: 3, result: Short within Pair
Pair: 2, fault length: 1.60m
Pair: 3, fault length: 0.80m
```

Ethtool(1) – Raw TDR data

```
./ethtool --cable-test lan2 amplitude-graph
Cable test for device lan2.
Cable test Pulse: 1000mV
  Distance     Pair 0     Pair 1     Pair 2
         0        109         85         39
         1        -15         46         -7
         2          7          0          0
         3          7          7          0
         4         15          7          7
         5         23          0          0
         6         15          0          0
         7         23          0          0
         8          7          0          0
         9          0          0          0
        10          0          0          0
        11          0          0          0
        12         -7          0          0
        13         -7          0         31
        14        359        140        273
        15        523        523        609
        16        601        570        632
        17        640        617        648
        18        562        625        179
        19        109        265         78
        20         85         93         70
        21         54         54         39
        22         46         46         31
        23         39         31         23
        24         15         23         15
        25          7         15         15
        26          0         15          7
        27          0          7          7
```

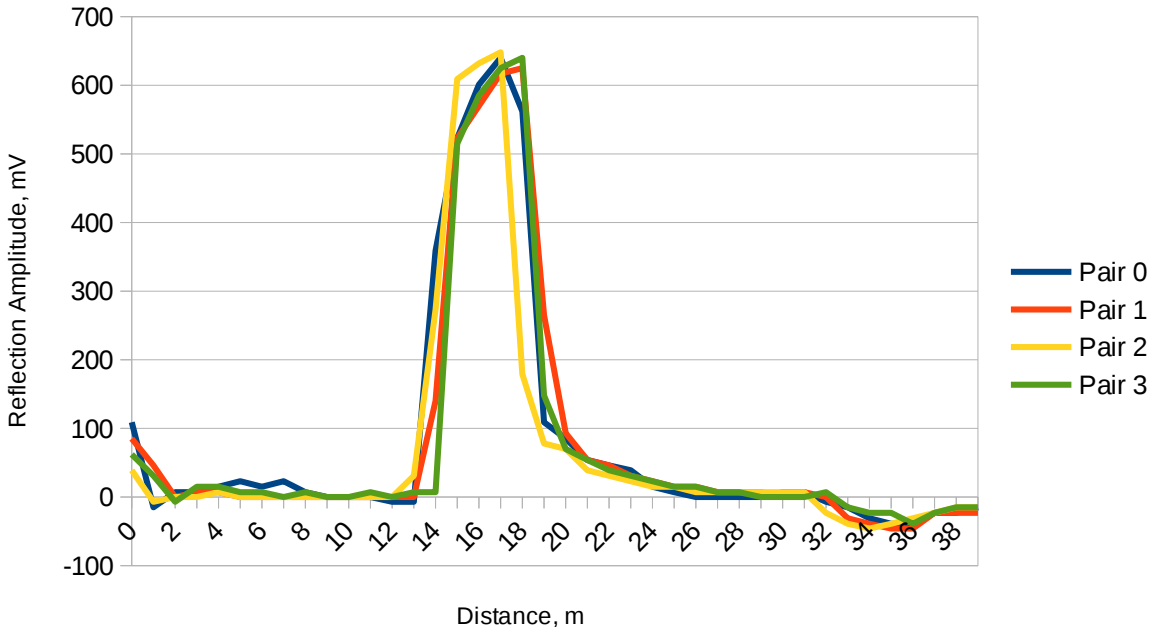

Status

- Waiting for Ethtool Netlink to be merged
- Code posted once for review
- Marvell 1G PHY driver supported
- Planned: 10G Marvell & Aquantia PHY drivers
- Microchip will likely add support
- Broadcom will likely add support