

---

# Building a Moderately Complex Modem with Spare Parts

Dan CaJacob

Director of Space Systems  
HawkEye 360, Inc.



GNURadio Conference 2017

Thursday 14 September 2017

# Abstract

We present a QPSK modem with concatenated coding and IP encapsulation implemented entirely in GnuRadio. The modem is made with core GnuRadio components to the maximum extent possible, then relying on a few existing OOT modules and finally, after exhausting all other options, the author deigns to write a few custom blocks (in Python). Scary topics like constellation rotation/inversion correction and latency reduction are addressed in passing. Fun is had by all. An ill-advised live demo may be demanded. Hecklers point out that this is painfully close to a CCSDS recommended standard.

# Overview

- Goal
- A Moderately Complex Modem?
- Spare Parts?
- Quick Success
- Modem
- A Few Custom Parts
- Five Stages of Avoiding Writing Code
- Defeating Latency
- Live Demo
- Future Work
- Acknowledgements




# Goal

- Build a full-duplex modem with concatenated FEC codes corresponding closely to CCSDS recommended standard
- Support IP / UDP encapsulation
- Minimize latency
- Bonus: Add TCP acceleration via some PEP magic



# A Moderately Complex Modem?

- Higher order modulation (QPSK)
- Concatenated Coding
  - $\frac{1}{2}$  Rate Convolutional
  - RS + Scrambling and Interleaving
- Basically the CCSDS recommendation, without their insane mishmash of higher layers
- Phase Ambiguity Issues
- IP Encapsulation
  - Custom, toy protocol, but you could easily replace with HDLC, etc.
- Latency Reduction
- Do as much as possible in GRC



# Spare Parts?

- Basically, I bend over backwards to not write a custom block (and fail)
- Because I'm lazy
- And because GNURadio largely makes this possible
- Also, there are some awesome OOT solutions out there
- And others that I completely missed
  - Sorry, gr-mapper



# Quick Success

- Not too hard to build a modem with concatenated FEC codes
  - Thanks to awesome OOTs like André Løfaldli's gr-ccsds
  - Experimented early with the old convolutional encoder block, quickly moved on to newer FEC API blocks
  - GR makes it easy to add Layers incrementally
- Use stock TUN block
  - Network interface to user-land code
- Latency was heinous (like 20 seconds)
  - Adding fill frames before modulator
- Works, but not really a good solution



# Modem

Stock

Existing  
OOT

Hier Block

Custom  
OOT

TUN

IP  
Encapsulator

RS Encoder,  
Interleaver,  
Scrambler

Convolutional  
Encoder

Constellation  
Modulator

Idle  
Frame  
Insert

USRP  
Sink

Channel

USRP  
Source

AGC

FLL BE  
Filter

Clock Sync

Equalizer

Phase Sync

Constellation  
Decoder

Convolutional  
Decoder

RS Decoder,  
De-Interleaver,  
De-Scrambler

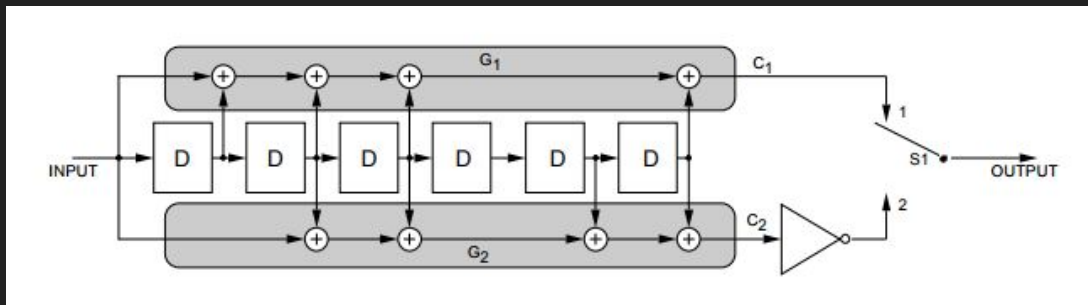
IP  
Decapsulator

TUN



# Modem: Convolutional Code

- CCSDS recommends  $R=1/2$ ,  $K=7$ 
  - Each output bit is dependent on  $2(K-1) = 12$  previous bits
  - So, we should ignore the first 12 bits sometimes
  - As a compromise for simplicity, ignore the first 16 bits
  - ASM is 32 bits, 64 bits after conv. code, so 52 of those are deterministically encoded
    - We just use 48 since it obeys byte boundaries
  - I'm actually not using the inversion yet, but Daniel Estevez made this easier to do, recently



# Modem: Reed Solomon Code

- CCSDS recommends
  - RS (255, 223)
  - Interleaving depth = 5
  - Scrambling
- I use André Løfaldli's gr-cssds OOT
  - Uses libfec (Phil Karn)

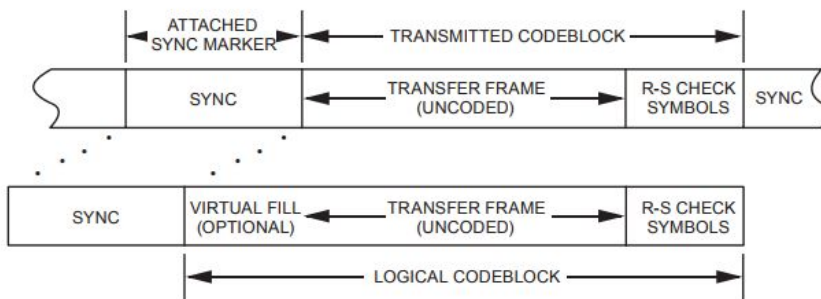


Figure 4-1: Reed-Solomon Codeblock Partitioning

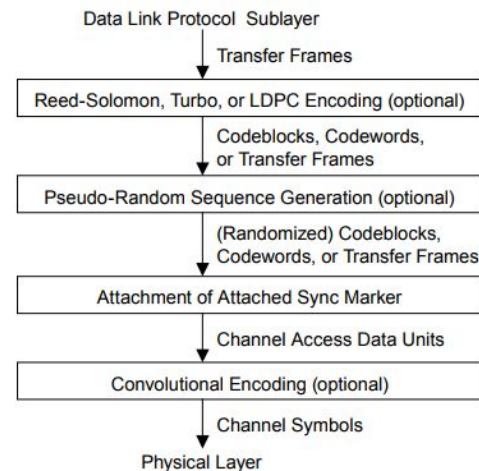
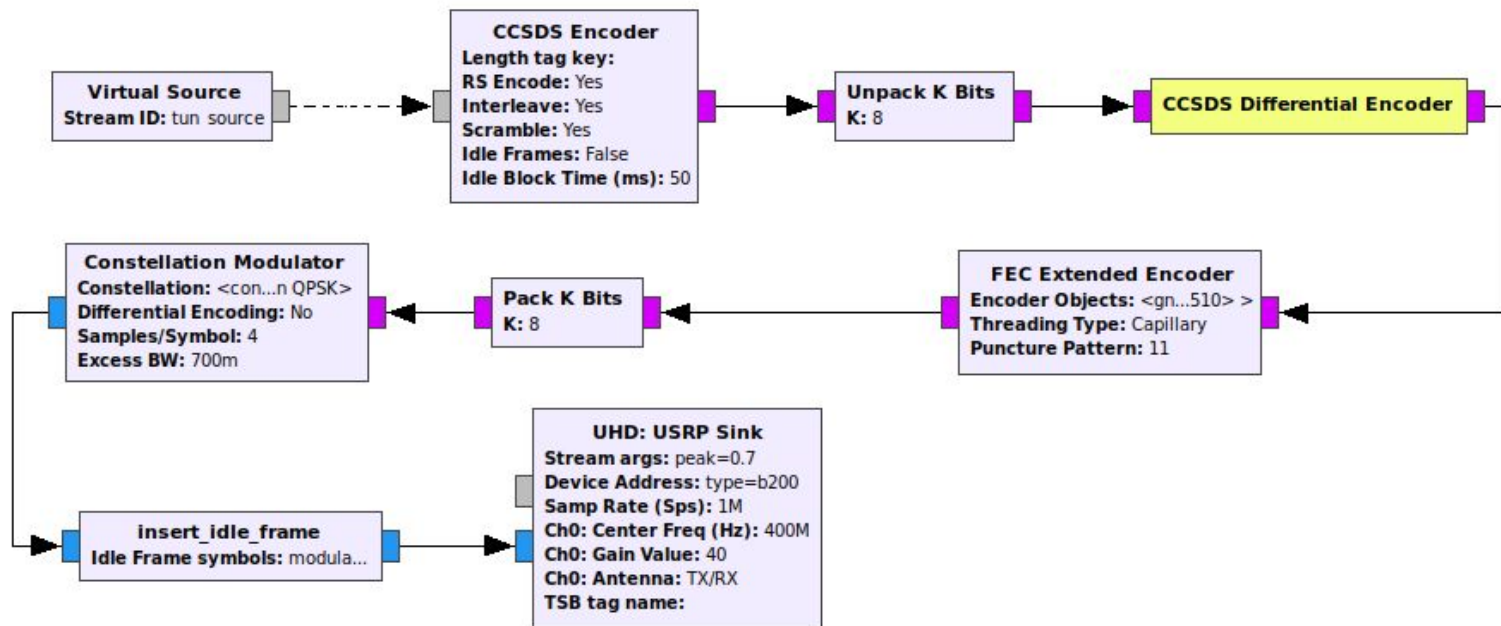
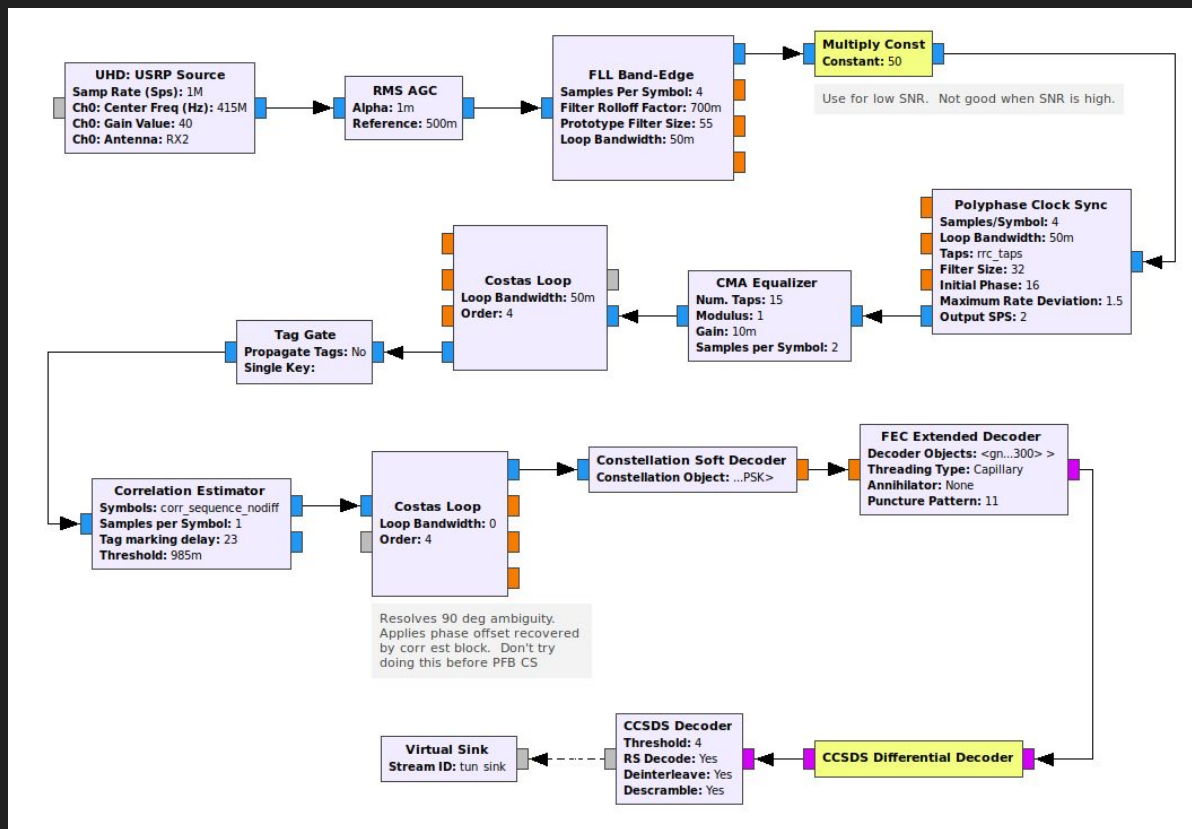


Figure 2-2: Internal Organization of the Sublayer at the Sending End

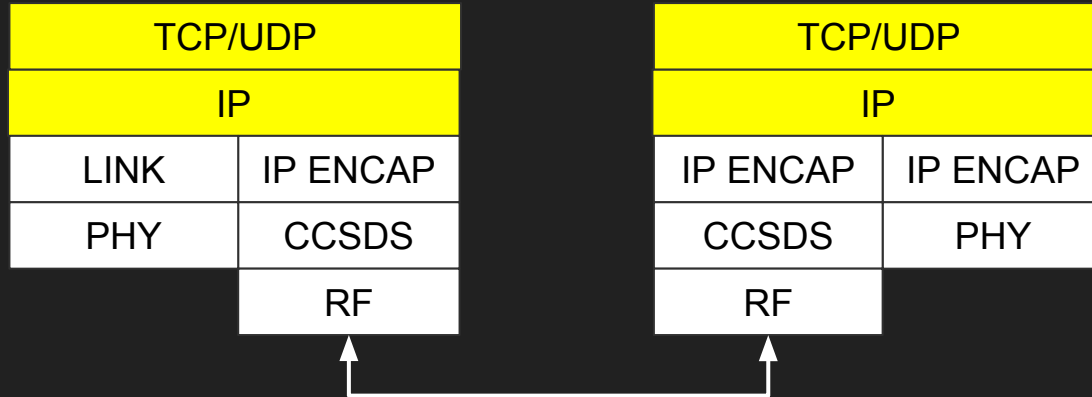

# Modem: Transmitter



# Modem: Receiver



# Modem: IP Encap/Decapsulation Interface



# A Few Custom Parts

- Tricky problems force me to write some code
  - Phase lock ambiguity
    - Spent a lot of time spinning my wheels on this
    - Eventually settled on a compromise solution that uses spare parts
  - Fill frame insert
    - Stupidly simple
    - Can probably do this with spare parts
  - Post Modulator, Convolutional Encoder Fill frame stitching
    - Not really custom code, but still hard to get right
  - IP Encapsulator / Decapsulator
    - Also stupidly simple



# Five Stages of ~~Grief~~ Avoiding Writing Code

## 1. Denial

- Surely this exists...

## 2. Anger

- Why the hell doesn't this exist?

## 3. Bargaining

- Can you help me make this exist?

## 4. Depression

- I'll never make this work.

## 5. Acceptance

- OK, I'll write some code.
- OK, I'll do the convolutional code by hand (Thanks for your help, Darek)



## 6. My own addition

- Can somebody help me write this code?



# Latency Reduction

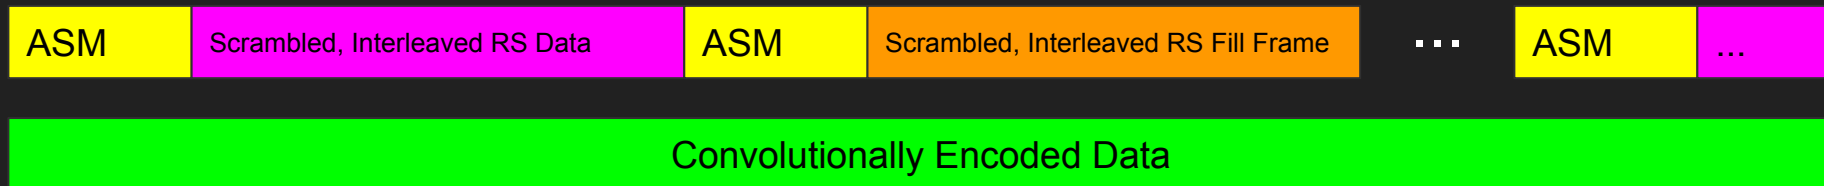
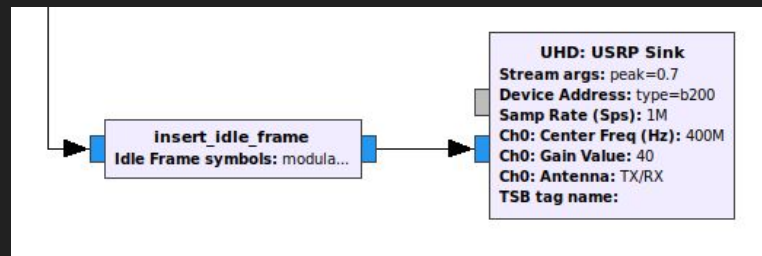
- Add fill frames / symbols at the last possible moment
  - Really... for reals
  - It makes phase synchronization a lot harder (because of the convolutional code)
- But go ahead and try *everything else* first
  - GR buffer manipulation (nope)
  - USB buffer buggery (kinda nope)
  - Bargaining (nope nope)
- Smaller gains
  - Tweak convolutional encoder / viterbi decoder work length (meh)
  - Increase the data rate (at least until the Os come)
- Gr-perf-monitorx keeps you honest






# Add Fill Frames Last

- Fill frames need to include all channel coding (convolutional coding!) and modulation
- If work is called and no items pending, send a fill frame
- Fill frame needs to be match up with previously sent data and data to come after (unpredictable)
  - Convolutional encoder and other blocks have memory, so this is hard!
  - But there's no better way to defeat latency




# Building the Fill Frame

- Fill frames are just a data block of 0s
- Several support flowgraphs
  - Idle\_frame\_encode
    - Encode fill sequence with RS, interleaving, scrambling, convolutional code
  - Sub graph in main flowgraph
    - Use Modulate Vector block to generate a modulated IQ vector
- Zero pad to flush, slice intelligently to combat memory in certain blocks
- $[0]*223*5 \rightarrow ((-0.6960...0.9399...j), \dots (-0.9150...+0.7737...j))$



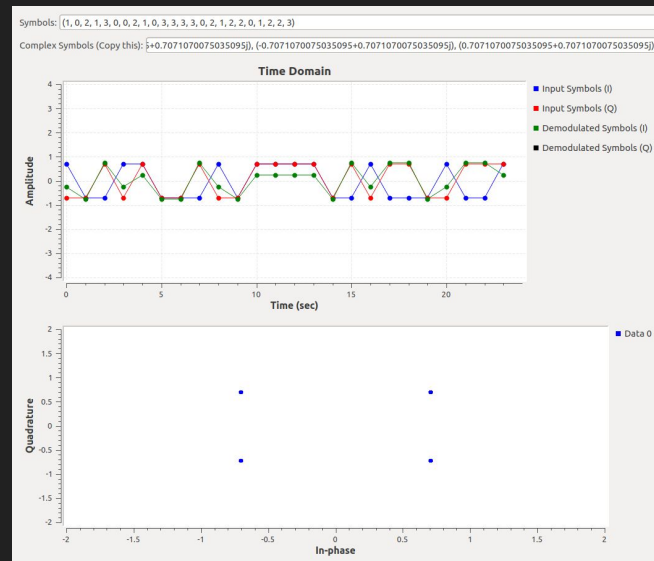
# Building the Fill Frame

- Hand encode ASM
  - After much prodding and help from colleagues
  - Crucial to getting to bit-accurate generation
  - Generator blocks have some memory, so need to trim intelligently



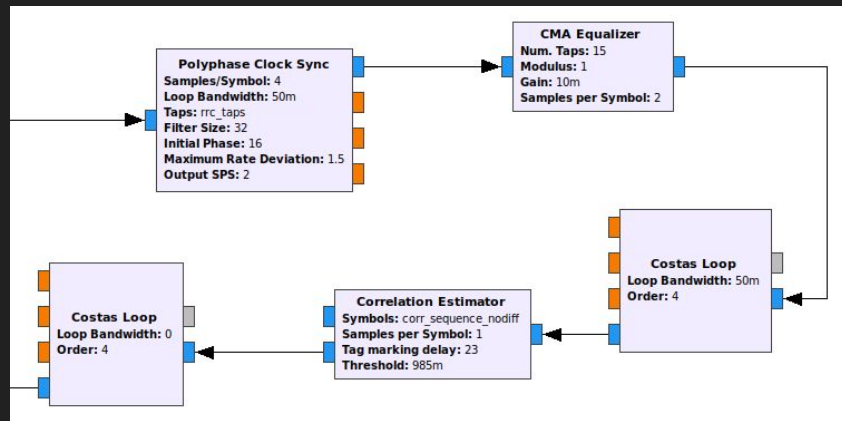


# Building the Correlation Sequence

- Several support flowgraphs
  - Asm\_test
    - Encode ASM with convolutional code
  - Asm\_post\_costas
    - Map encoded ASM to complex symbols
- Zero pad to flush, slice intelligently to combat memory in certain blocks
- 0x1ACFFC1D →  
 $((0.7071...-0.7071...j), \dots (0.7071...+0.7071...j))$



# Phase Synchronization

- Harder than BPSK
- Costas loop will lock, but perhaps not to the right rotation, order
  - Simulations can fool you (unless you simulate real world effects)
  - Over the air, you might be right 50% or 25% of the time, depending on coding
- Correlation estimator and 2nd Costas Loop clean up the ambiguity
- There are better ways to do this



# Phase Synchronization

- There are better ways to do this

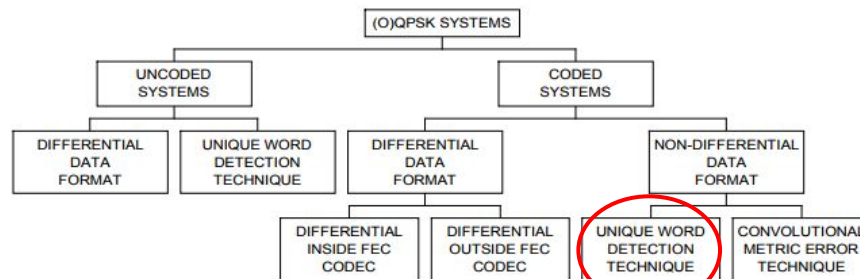


Figure 2.4.11-1: List of Phase-Ambiguity Resolution Techniques

TABLE 2.4.11-1: RELATIONSHIPS BETWEEN THE TRANSMITTED AND RECEIVED DATA

| CARRIER<br>PHASE ERROR<br>(DEGREES) | RECEIVED DATA |        |
|-------------------------------------|---------------|--------|
|                                     | $I_R$         | $Q_R$  |
| 0                                   | $I_T$         | $Q_T$  |
| 90                                  | $-Q_T$        | $I_T$  |
| 180                                 | $-I_T$        | $-Q_T$ |
| 270                                 | $Q_T$         | $-I_T$ |

TABLE 2.4.11-2: SUMMARY OF THE SALIENT FEATURES OF THE PREFERRED TECHNIQUES

| AVAILABLE<br>TECHNIQUES   | BIT ERROR RATE<br>(BER) DEGRADATION                                      | ADVANTAGES &<br>DISADVANTAGE  |
|---|--|---|
| UNIQUE WORD DETECTION   | NONE   | - INCREASE EARTH STATION COMPLEXITY   |
| DIFFERENTIAL DATA FORMATTING<br>WITHOUT<br>FORWARD-ERROR-CORRECTION<br>(FEC)          | INCREASES BY APPROXIMATELY<br>A FACTOR OF TWO                            | - SIMPLE TO IMPLEMENT<br>- CAN CAUSE DEGRADATION IN THE<br>DETECTION OF THE TRANSMITTED<br>SYNC MARKERS |
| DIFFERENTIAL DATA FORMATTING<br>INSIDE THE<br>FEC ENCODER AND DECODER PAIR<br>(CODEC) | ABOUT<br>3 dB FOR CONVOLUTIONAL<br>CODE WITH $R = \frac{1}{2}$ , $K = 7$ | - PROVIDES QUICK PHASE AMBIGUITY<br>RESOLUTION<br>- REQUIRES OVERPOWERED LINK                           |
| DIFFERENTIAL DATA FORMATTING<br>OUTSIDE<br>THE FEC CODEC                              | SMALL  | - REQUIRES DIFFERENTIAL DECODERS<br>AT THE STATION  |

# IP Encapsulation / Decapsulation

- Toy implementation
  - Dead simple Python blocks
  - 1 IP / UDP packet per CCSDS frame
  - Just add a 2 byte header to define length of encapsulated packet
  - Enforce MTU size limit (2 header bytes + 1113 data bytes)
- Better solutions
  - HDLC encode to delimit packets
- Worse solutions
  - Actually follow CCSDS, don't do that



Header

IP Packet

Free (0 - n)

ASM

Scrambled, Interleaved RS Code blocks

...

ASM

...

# Live Demo: Loopback Simulation


- Note unchanging fully encoded ASM symbols at start of frames
  - Accurate sticking of fully pre-encoded fill frames allows > 100x latency reduction
  - This was the principal challenge and success of this work



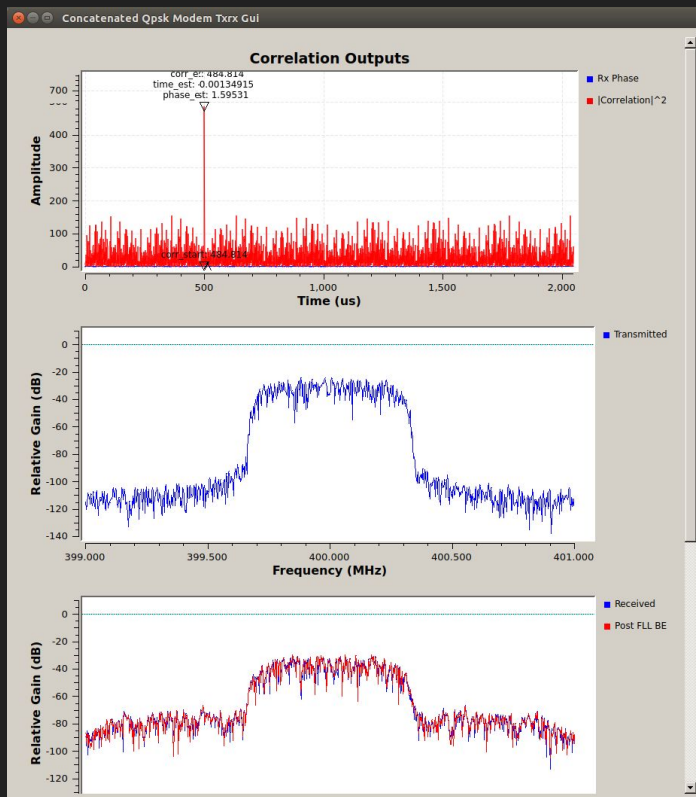


# Live Demo: OTA

- Connect
- Ping
  - Highlight low latency
- Disconnect channel and demonstrate self-healing
- SSH / Mosh
  - Hollywood time!
- File Transfer?
- Demo PEP acceleration on a simulated high latency, high BER channel



# Live Demo



```
dan@dan-XPS-15-9550: ~/target_gnuradio/src/gr-ccsds/examples
dan@dan-XPS-15-9550: ... x dan@dan-XPS-15-9550: ... x root@dan-XPS-15-9550: ... x dan@dan-XPS-15-9550: ... x
64 bytes from 10.11.10.1: icmp_seq=804 ttl=64 time=186 ms
64 bytes from 10.11.10.1: icmp_seq=805 ttl=64 time=187 ms
64 bytes from 10.11.10.1: icmp_seq=806 ttl=64 time=191 ms
64 bytes from 10.11.10.1: icmp_seq=807 ttl=64 time=175 ms
64 bytes from 10.11.10.1: icmp_seq=808 ttl=64 time=180 ms
64 bytes from 10.11.10.1: icmp_seq=809 ttl=64 time=184 ms
64 bytes from 10.11.10.1: icmp_seq=810 ttl=64 time=188 ms
64 bytes from 10.11.10.1: icmp_seq=811 ttl=64 time=192 ms
64 bytes from 10.11.10.1: icmp_seq=812 ttl=64 time=176 ms
64 bytes from 10.11.10.1: icmp_seq=813 ttl=64 time=181 ms
64 bytes from 10.11.10.1: icmp_seq=814 ttl=64 time=185 ms
64 bytes from 10.11.10.1: icmp_seq=815 ttl=64 time=189 ms
64 bytes from 10.11.10.1: icmp_seq=816 ttl=64 time=193 ms
64 bytes from 10.11.10.1: icmp_seq=817 ttl=64 time=177 ms
64 bytes from 10.11.10.1: icmp_seq=818 ttl=64 time=182 ms
64 bytes from 10.11.10.1: icmp_seq=819 ttl=64 time=186 ms
64 bytes from 10.11.10.1: icmp_seq=820 ttl=64 time=190 ms
64 bytes from 10.11.10.1: icmp_seq=821 ttl=64 time=174 ms
64 bytes from 10.11.10.1: icmp_seq=822 ttl=64 time=178 ms
64 bytes from 10.11.10.1: icmp_seq=824 ttl=64 time=187 ms
64 bytes from 10.11.10.1: icmp_seq=825 ttl=64 time=190 ms
64 bytes from 10.11.10.1: icmp_seq=826 ttl=64 time=174 ms
64 bytes from 10.11.10.1: icmp_seq=827 ttl=64 time=179 ms
64 bytes from 10.11.10.1: icmp_seq=828 ttl=64 time=184 ms
64 bytes from 10.11.10.1: icmp_seq=829 ttl=64 time=188 ms
64 bytes from 10.11.10.1: icmp_seq=830 ttl=64 time=192 ms
64 bytes from 10.11.10.1: icmp_seq=831 ttl=64 time=176 ms
64 bytes from 10.11.10.1: icmp_seq=832 ttl=64 time=180 ms
64 bytes from 10.11.10.1: icmp_seq=833 ttl=64 time=185 ms
64 bytes from 10.11.10.1: icmp_seq=834 ttl=64 time=189 ms
64 bytes from 10.11.10.1: icmp_seq=835 ttl=64 time=190 ms
64 bytes from 10.11.10.1: icmp_seq=836 ttl=64 time=174 ms
64 bytes from 10.11.10.1: icmp_seq=837 ttl=64 time=179 ms
64 bytes from 10.11.10.1: icmp_seq=838 ttl=64 time=183 ms
64 bytes from 10.11.10.1: icmp_seq=839 ttl=64 time=188 ms
64 bytes from 10.11.10.1: icmp_seq=840 ttl=64 time=191 ms
64 bytes from 10.11.10.1: icmp_seq=841 ttl=64 time=175 ms
64 bytes from 10.11.10.1: icmp_seq=842 ttl=64 time=174 ms
64 bytes from 10.11.10.1: icmp_seq=843 ttl=64 time=179 ms
64 bytes from 10.11.10.1: icmp_seq=844 ttl=64 time=183 ms
64 bytes from 10.11.10.1: icmp_seq=845 ttl=64 time=187 ms
64 bytes from 10.11.10.1: icmp_seq=846 ttl=64 time=191 ms
64 bytes from 10.11.10.1: icmp_seq=847 ttl=64 time=175 ms
64 bytes from 10.11.10.1: icmp_seq=848 ttl=64 time=180 ms
64 bytes from 10.11.10.1: icmp_seq=849 ttl=64 time=183 ms
64 bytes from 10.11.10.1: icmp_seq=850 ttl=64 time=188 ms
64 bytes from 10.11.10.1: icmp_seq=851 ttl=64 time=192 ms
64 bytes from 10.11.10.1: icmp_seq=852 ttl=64 time=176 ms
64 bytes from 10.11.10.1: icmp_seq=853 ttl=64 time=181 ms
64 bytes from 10.11.10.1: icmp_seq=854 ttl=64 time=185 ms
64 bytes from 10.11.10.1: icmp_seq=855 ttl=64 time=189 ms
64 bytes from 10.11.10.1: icmp_seq=856 ttl=64 time=193 ms
64 bytes from 10.11.10.1: icmp_seq=857 ttl=64 time=177 ms
64 bytes from 10.11.10.1: icmp_seq=858 ttl=64 time=182 ms
64 bytes from 10.11.10.1: icmp_seq=859 ttl=64 time=186 ms
64 bytes from 10.11.10.1: icmp_seq=860 ttl=64 time=190 ms
64 bytes from 10.11.10.1: icmp_seq=861 ttl=64 time=173 ms
```

# Future Work

- Replace toy IP encapsulation protocol with HDLC / MPoFR
- Figure out how to synchronize OQPSK
  - Technically, CCSDS recommends OQPSK, not QPSK
- Replace hacky phase synchronizer with more robust and more efficient implementation
  - Maybe try gr-mapper?
- Fix bug with long-running file transfers
  - ASM seen in data packet, confusing PLL?
- Script the correlation and idle frame generation
- Measure BER vs.  $E_s/N_0$
- Bonus: Get all the OOT stuff in tree
  - RS  $\rightarrow$  FEC API, etc.
- Get an intern to do it all for me



# Acknowledgements

- Darek Kawamoto
- Nick McCarthy
- Andy Walls & the IRC gang
- André Løfaldli
  - gr-ccsds
- Daniel Estevez
  - Read his blog: <http://desteveez.net/>
- HawkEye 360



# Questions?

- HE360 is hiring!



*"We'd now like to open the floor to shorter speeches disguised as questions."*