Correlator numbers/use cases

Paul Boven

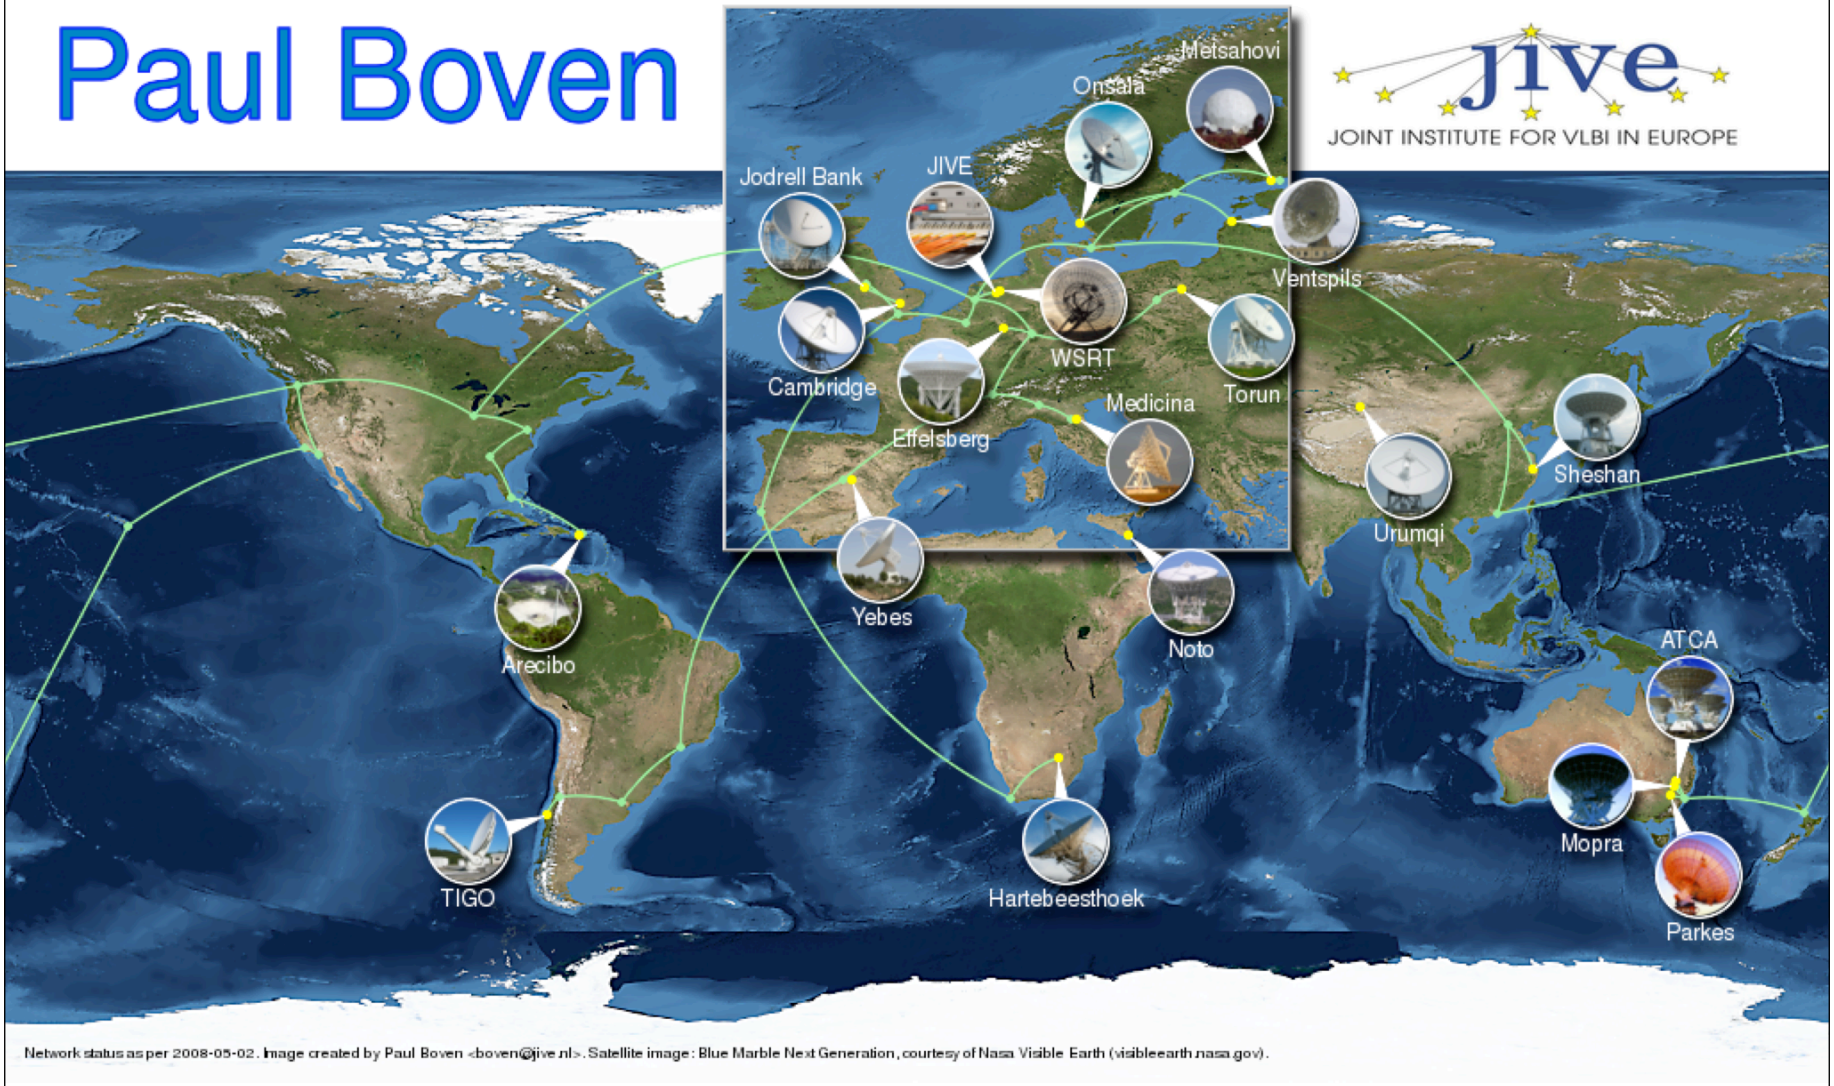

JIVE
JOINT INSTITUTE FOR VLBI IN EUROPE

Uniboard Correlator Overview

- Data arrives in subbands
- Transpose: full L-R mesh
- Route per-subband to correct input chip
- Input chip: buffer, delay tracking, PFB, fringe rotation
- 2x DDR3 (2x 52.6Gb/s)
- 2x 8GB -> 3.2 s
- Correlator chip: cross- and auto correlations, integration, all baselines per chip
- QDR2+ memory: 64.8 Gb/s, 4MB

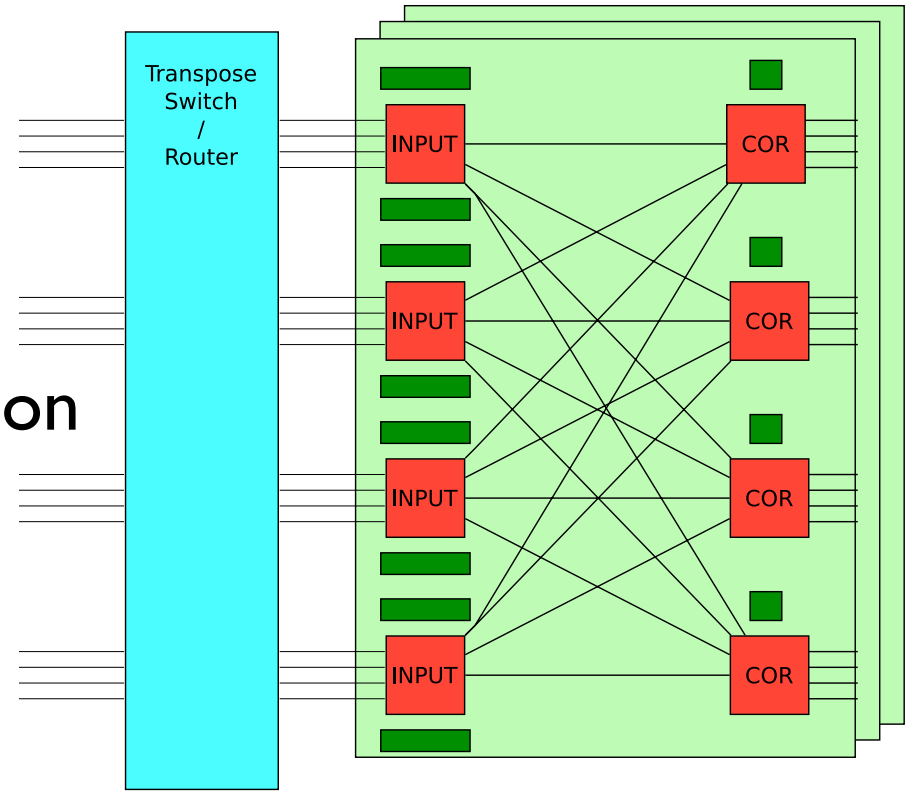

Bandwidths and bottlenecks

- Data arrives in subbands
- Input chip forms frequency channels (2kHz - 1MHz)
- Frequency channels can be sent to 4 cor. chips
- $BW_{SUB} \leq 4x\ BW_{COR}$ (without backplane)
- Correlator chip has ≤ 40 Gb/s input BW
- This contains 32 stations, 2 polarisations: 64 streams
  - 625Mb/s per stream, complex data
- For 64MHz per correlator chip: max. 4 bits
- For 32MHz per correlator chip: max. 9 bits
- Multipliers are 18x18 (or 18x25), only useful >9 bits!
- Run the on-board interconnect at double rate?
- Goal is to get most GMACs out, not be IO limited

Bitgrowth in FFTs

- FFT wordsize grows by N_{Stages} + 1
  - Noise-power diminishes with smaller bandwidth
  - But CW carriers (RFI) will keep the same amplitude
- RFI is local to telescope (we hope)
  - Use bitgrowth at top for RFI immunity
  - Use bitgrowth at bottom for increased precision
  - Ignore bitgrowth for more BW, sensitivity
- Example: 3 bit subbands of 32MHz, 4kHz bins:
  - bitgrowth is 15 bits, FFT output is 18 bits
- At lower frequencies (e.g. L-band) there is more RFI, but less spectrum to sample
- At higher frequencies, go for highest bandwidth
- Bitgrowth means TP growth: $TP_{cor} > TP_{inp}$

Input chip limitations

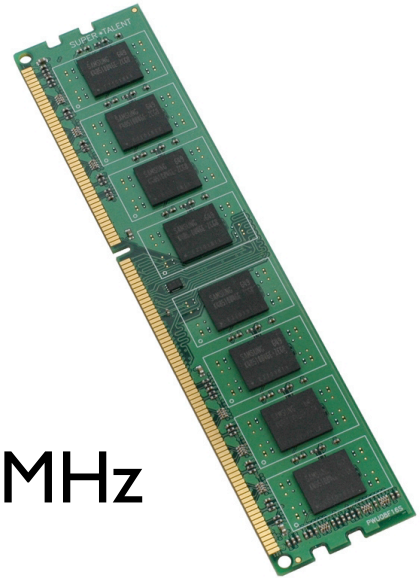

- Must store data in memory
- Network jitter compensation
- Delay tracking
- Best bandwidth so far: Virtex-5 DDR3 400MHz
  - 64 bit dimms - 50Gb/s
- Need at <u>least</u> twice that, so 2 memory interfaces
- Still only enough bandwidth to go in & out once
- Combine jitter buffer and delay
  - Data hits memory bottleneck only once
- 1x 32bit DDR3 interface: 1760 slices on V5
  - 5% of a V5 SX240T
- Reduce input data rate? (also allows for bitgrowth)
- Investigate performance on actual usage pattern!

Care and feeding of multipliers

- Use built-in DSP resources (Multiplier/accumulator)
- Make sure these are always busy, always have data
- Use built-in accumulator for integration

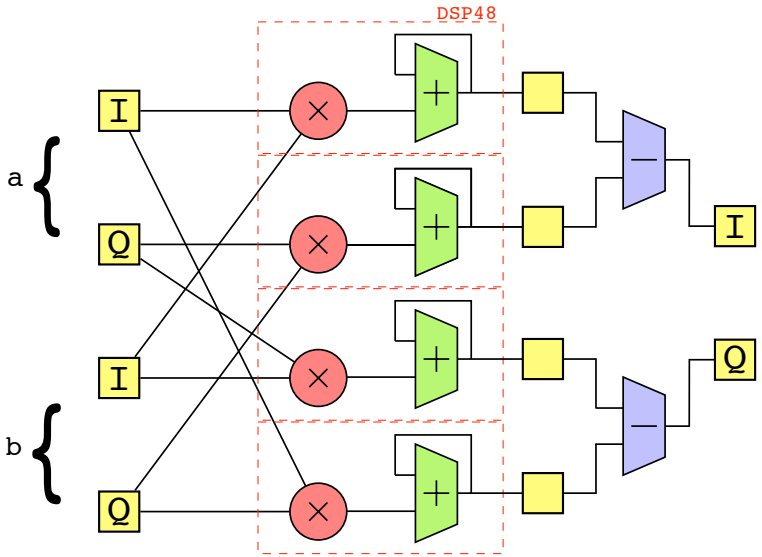

- Input chip PFB creates time-ordered spectra
- Re-order to per-subband
- Second transpose needed
- QDR lacks bandwidth
  - Two QDR interfaces?

- Nbr of integrations depends on size of 2nd transpose
- 64 streams, 8 bit (compl) in 4MB: 65k datapoints
  - 32MHz, 64 bins (500kHz): 1024 integrations,
  - 8192 bins (4kHz): 8 integrations.

Correlator chip performance

- Process a single band, all stations per correlator unit
- 32 stations, all stokes => 2112 products
- Complex multiplication, 8192 MACs needed
  - Virtex 5 SX240T: 1056 DSP48 (R_M = 7.7x)
  - Stratix 4 SGX230: 1288 18x18 (R_M = 6.4x)
  - Virtex 6 SX475T: 2016 DSP48E (R_M = 4.1x)
- Re-using multipliers means re-issuing data

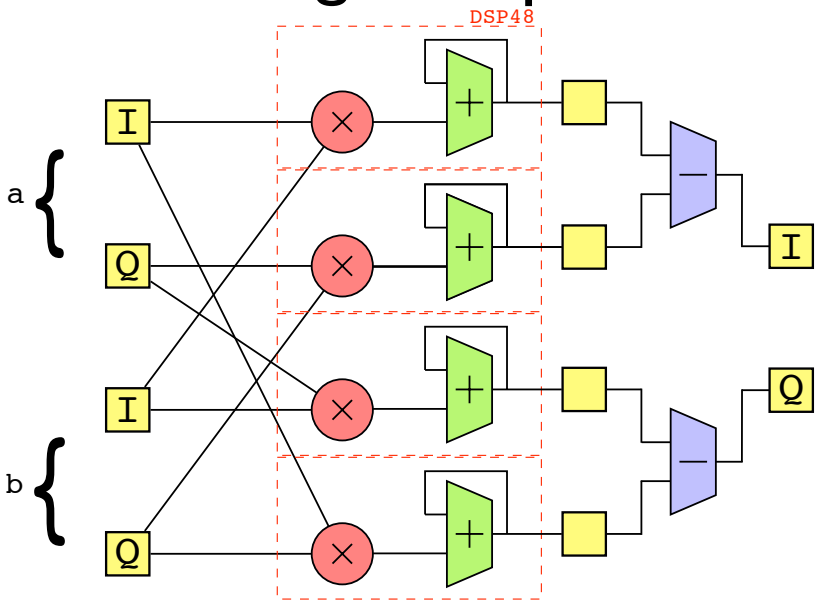

- $BW_{cor} \times R_M = F_M$

- 400MHz design goal?
  - V5: BW_{cor} < 50MHz
  - S4: BW_{cor} < 60MHz
  - V6: BW_{cor} < 95MHz

A 31 station correlator?

- Virtex 6 SX475T: 2016 DSP48E, R_{M31}=3.8
- 1064x dual-ported BRAM, 36kbit
  - 18 bits complex data: each BRAM is1024 points
- 32 telescopes x 2 pol = 64 datastreams
  - 16 x 1024 = 16k points - probably too optimistic
- 2nd acc. on-chip (in LUTs) to limit output BW

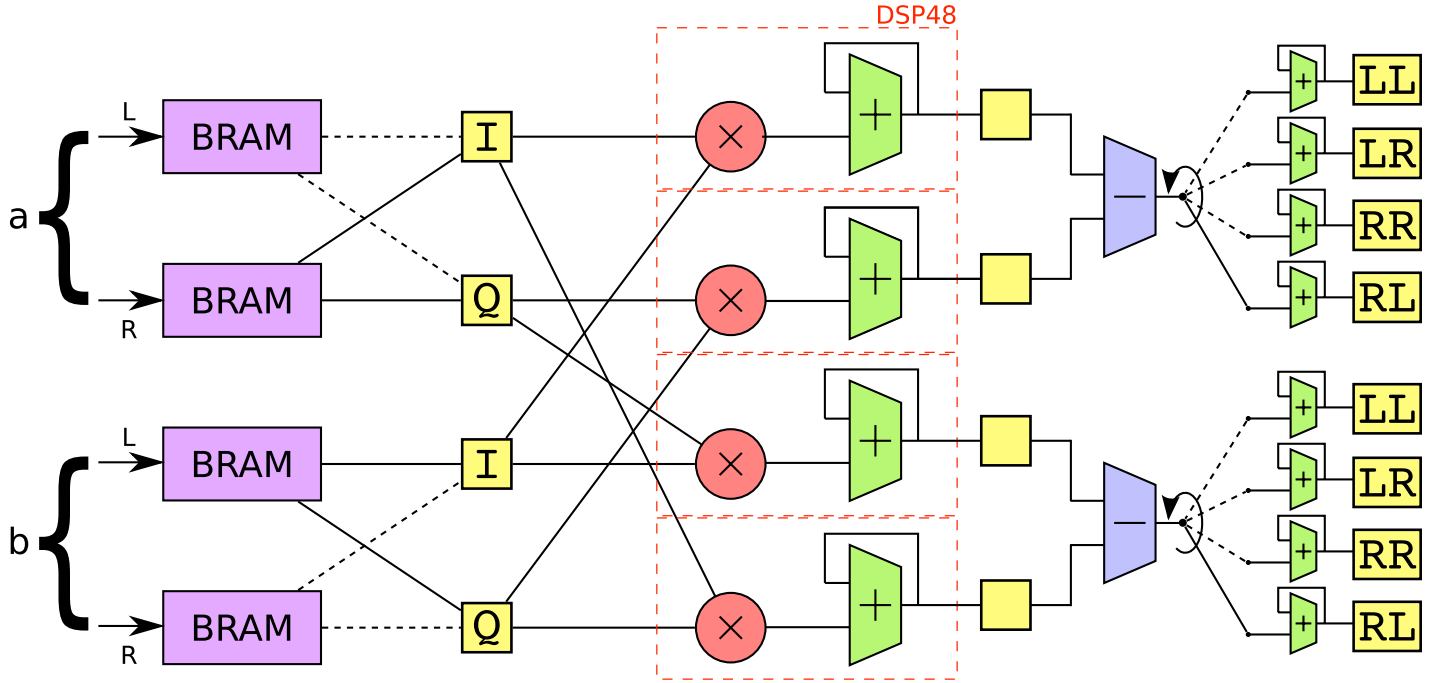

A 31 station correlator?

- 31 stations would need:
  - 1922 DSP48 (out of 2016)
- Store data close to multipliers (after duplication):
  - Single BRAM must store both L and R
  - 961 BRAM (out of 1064)
- Store data only once:
  - 64 BRAM (or multiple), much more routing

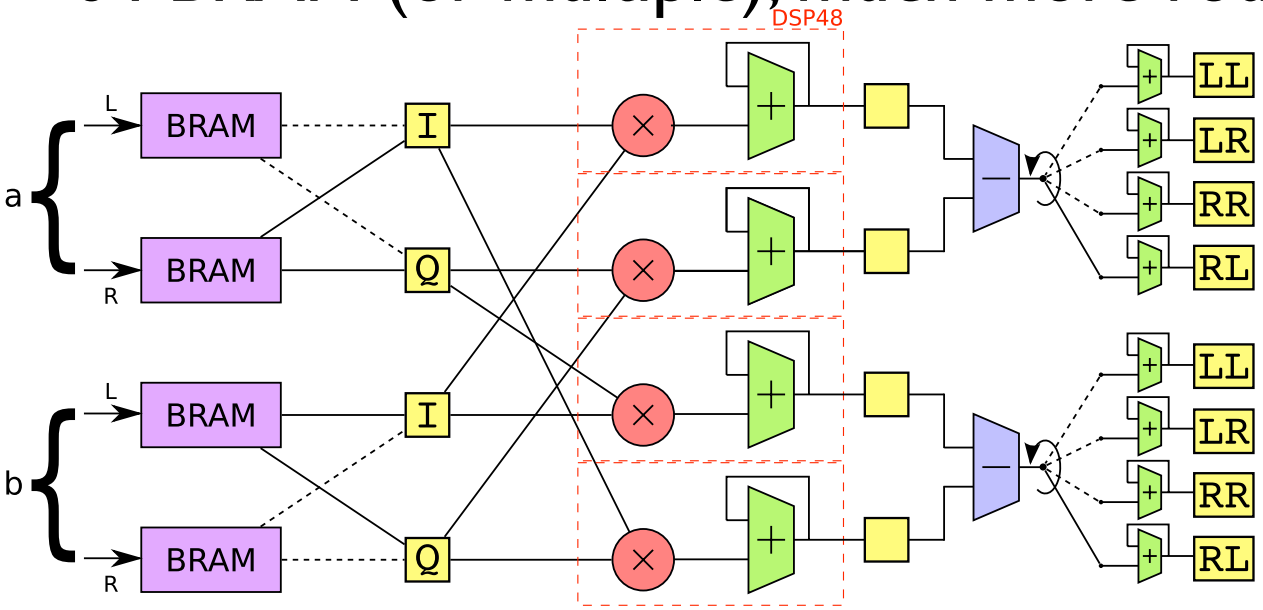

A closer look at DSP48E

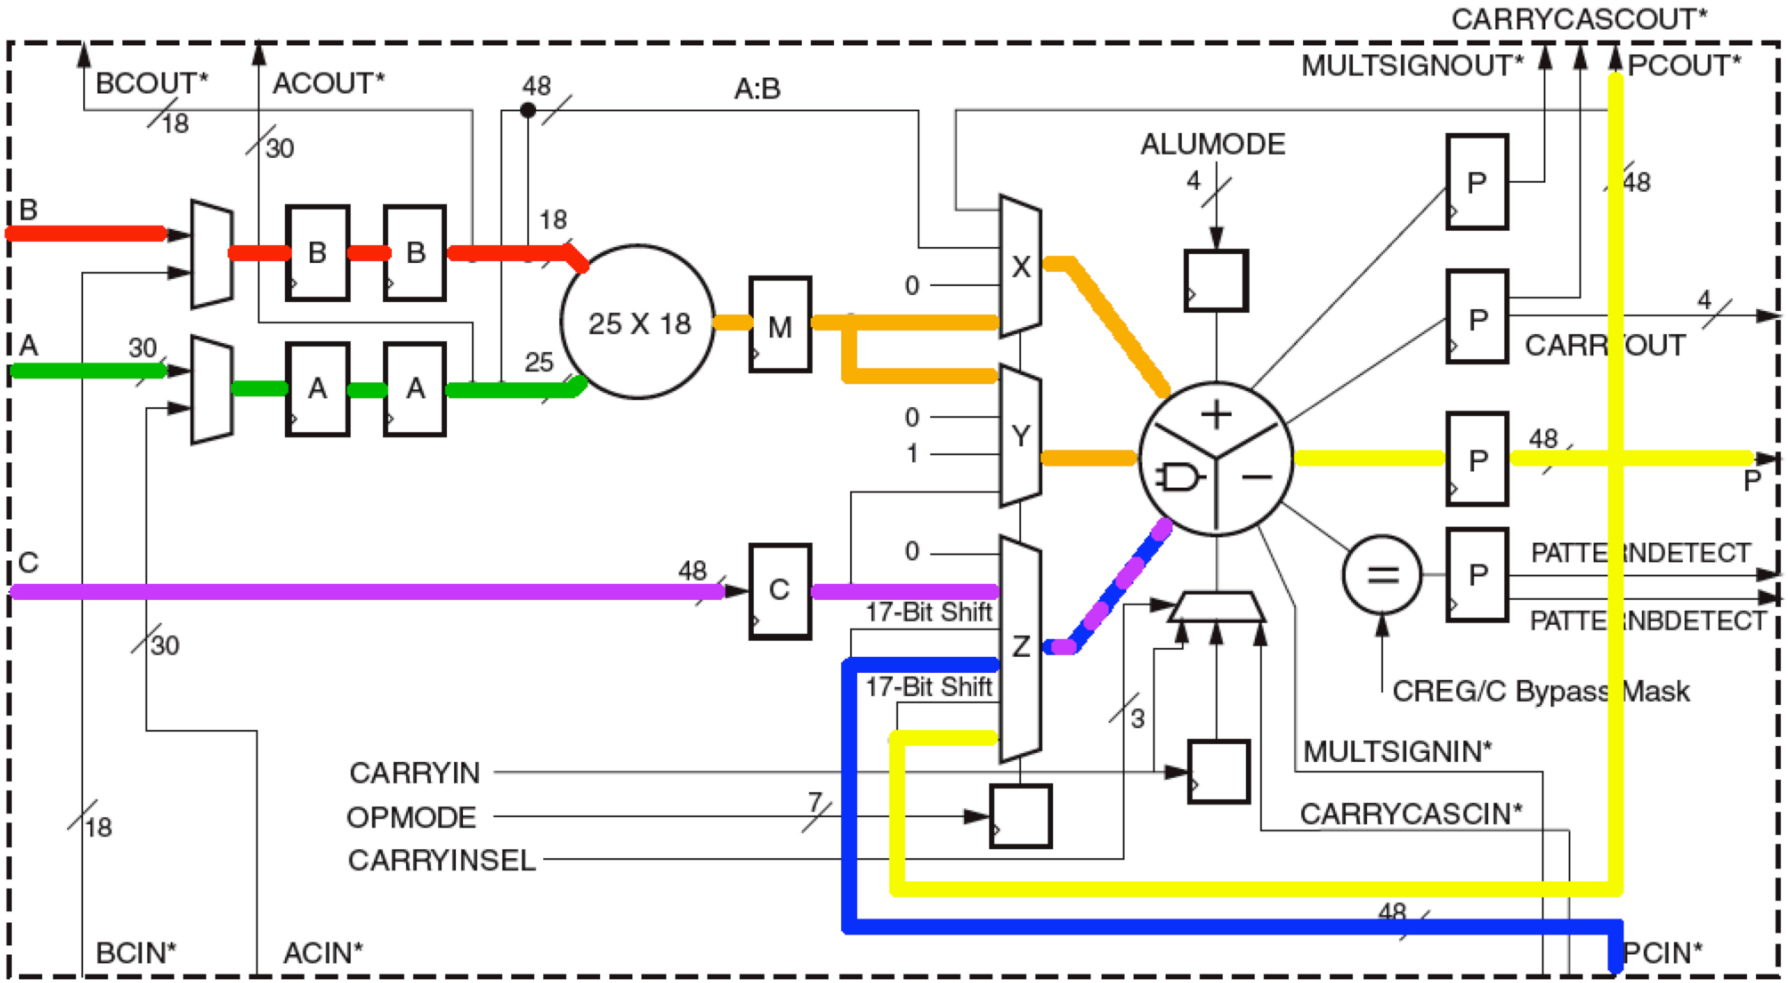

*These signals are dedicated routing paths internal to the DSP48E column. They are not accessible via fabric routing resources.

UG193_c1_01_032806

Multiple accumulators

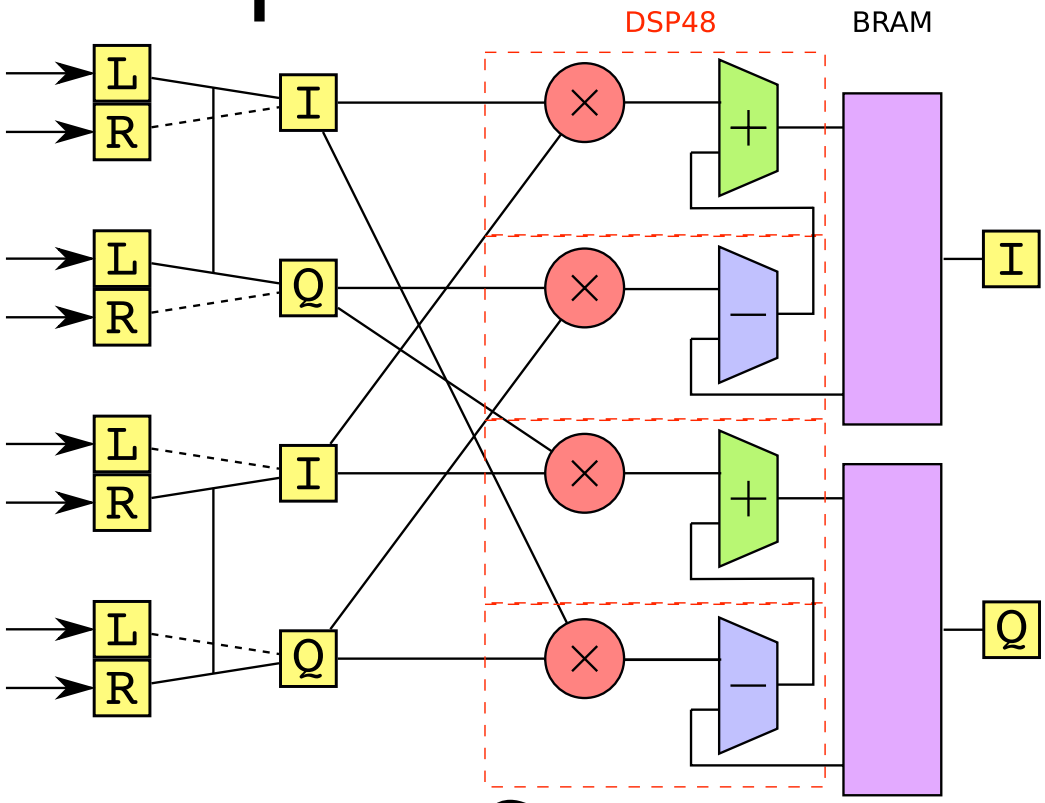

Pros:

- No 2nd transpose
- 36 bit accumulators
- No CLB math
- Matches resources

Cons:

- 1024 accumulators
- 256 bins, 4 stokes
- Uses almost all BRAMs
- Pipeline for speed

Summary (so far)

- Only Virtex 6 can do 64MHz per chip
- But does not exist yet
- I have no Xilinx (or Altera) shares
- Either 2 or 0 QDR chips for correlator
- Correlator chip is bandwidth limited - 80Gbs please?
- On Virtex 6, 31 stations is much easier than 32
- 2nd transpose or multiple accumulators
- On-chip routing resources might be final bottleneck
- 32 (31) stations, 4 Stokes is worst (best) case
- These are all just upper limits
- A lot of design paramters to review

- Hence : The Correlator Construction Kit

The numbers game

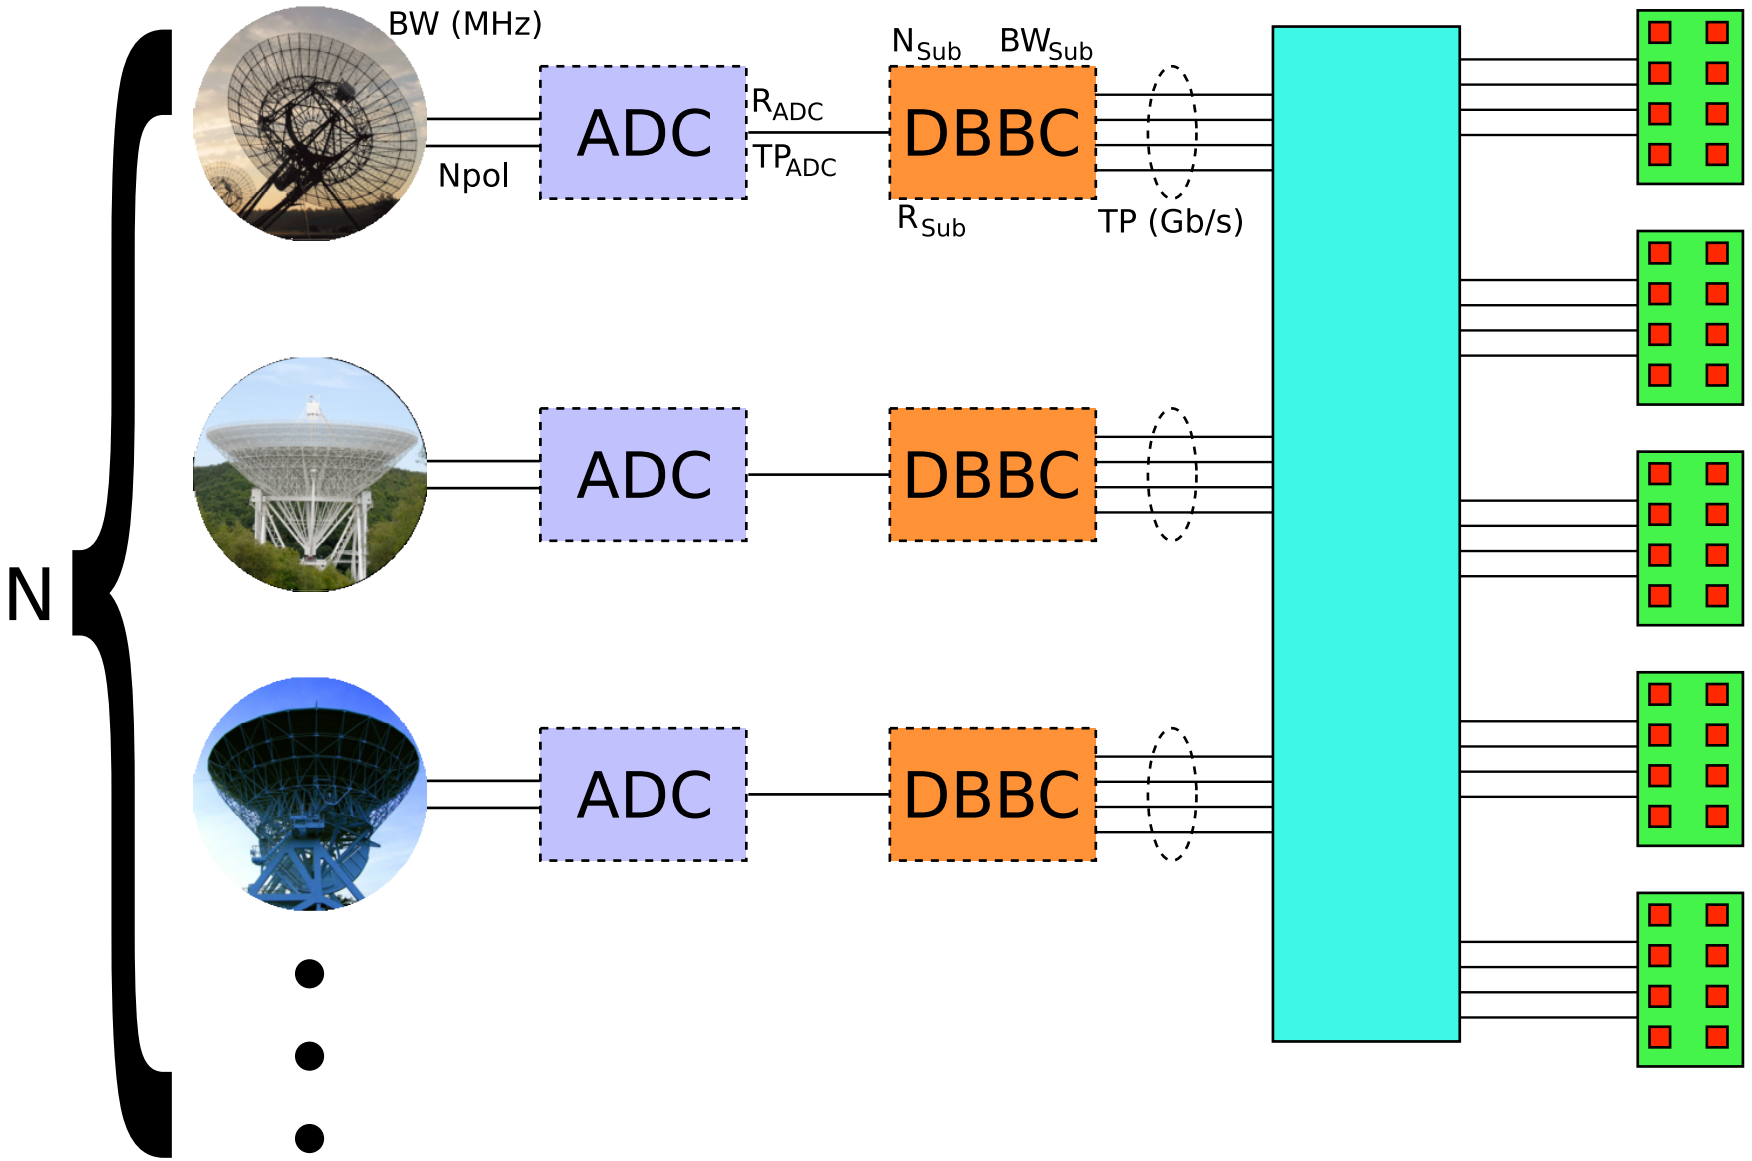

The numbers game (II)

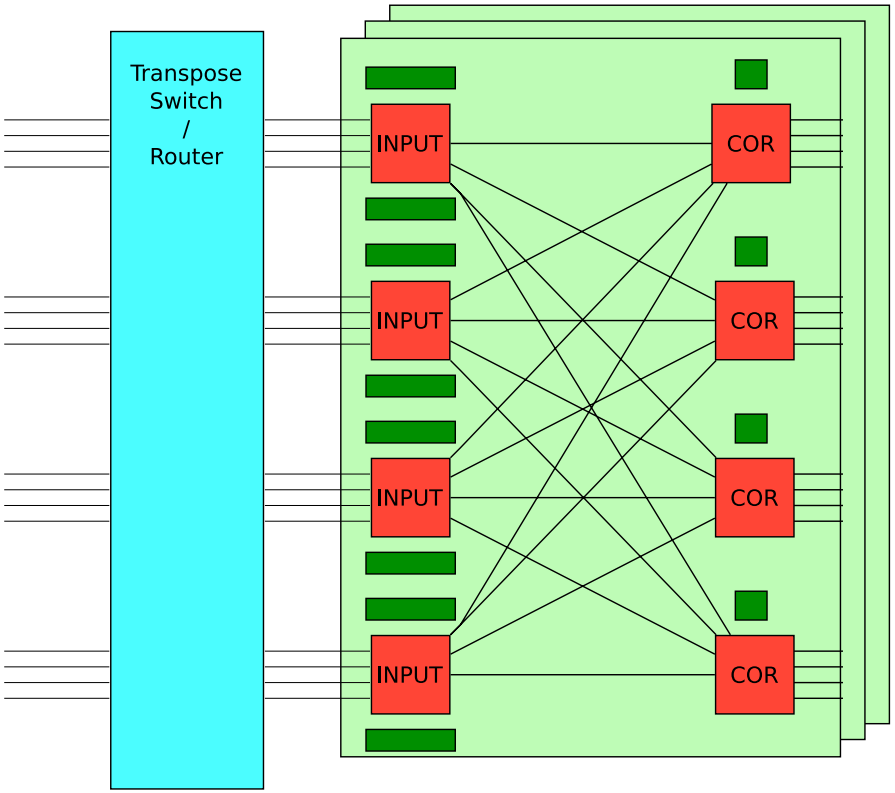

- TP_{INP} (Gb/s)
- TP_{COR} (Gb/s)
- BW_{COR} (MHz)
- N_{bin}
- BW_{bin} (MHz)
- R_{bin} (bits)

Some use cases

- Current EVN correlator
  - N = 16, BW = 128MHz, T_{int} = 0.25s
- Next-gen correlator
  - N = 32, BW = 512MHz, T_{int} = 0.1s
- EVN2015
  - N = 32, BW = 8GHz - 16GHz, T_{int} << 0.01s
- Spectral line research
  - N = 32, BW = 128MHz (1000m/s), BW_{bin} < 10m/s
  - Full BW at lower spectral resolution
- Space
  - BW_{sub} = 128MHz (VSOP2), N_{bin} = 1E6, delay track

Current EVN correlator

Input side

N	16		Number of stations
Nb	120		N*(N-1)/2
BW	128	MHz	Input bandwidth
Npol	2		Number of polarizations
Radc	2	bits	Resolution of ADC
TP	1.024	Gb/s	Data rate per station, BW * 2 * Npol * Radc
			This would be the network TP without on-site PFB
ATP	16.384	Gb/s	All stations, N * TP

Subbands

Nsub	8		For one polarization
BWsub	16	MHz	BW / Nsub
Rsub	2	bits	Quantization after PFB
TPsub	2048	Mb/s	Subband data rate, N * 2 * BWsub * Rsub * Npol
			TP is the input rate for 1 correlator FPGA (<40Gb/s

Frequency Resolution

Npoints	16		Frequency points per subband
dF	1000	kHz	BWsub / Npoints
Fobs	C band (4-8 GHz)		Observing frequency/band
ΔV	0.0	km/s	Doppler resolution
Vmax / Vmin	/691	km/s	Doppler range, Fobs - BW/2 to Fobs + BW/2

Output

Cpol	X pol, cross + auto		
Nc	544		Number of correlation products: Cpol * (Nb + N)
Nbin	139264		Number of bins, 2 * Nc * Npoints * Nsub
Ro	24	bits	Output resolution
TPo	6.7	Mb/s	Nbin * Ro * 2 / Tint (times 2 because complex)
TPo	3	GB/hr	

FPGA Resources

Fg	25	MHz	Frequency goal
Nm	195		Number of Multipliers or MACs.
G	270	GMAC	Estimated single FPGA performance
	DBBC		One of these at each telescope
TPpfb	512	Mb/s	PFB input (per polarization): 2 * BW * Radc
Fadc	256	MHz	ADC clock rate, 2 * BW
Fadc/Fg	1		ADC samples per FPGA clock (>1 requires DDR/striping) BW * 2 * log2 (Nsub) * Npol / 1000
TPtel	1024	Mb/s	Network throughput per telescope, 2 * BW * Rsub * Npol
	Input chip (PFB)		One for each subband
TPfft	2048	Mb/s	Input rate for FFT, 2* BW / Nsub * Rsub * Npol * N
Gfft	4	GMAC	BW / Nsub * 2 * log2(Npoints) * Npol * N / 1000
	Output chip (correlator)		One for each subband
Gcor	35	GMAC	4 * BW / Nsub * Nc / 1000
Mcor	0	MB	Storage for integration, Nc * Ro * Npoints * 2 / 8 / 1024^2