NRC · CNRC

Expanded Very Large Array

Overview of New Connectivity Scheme for the EVLA Correlator

B. Carlson

Outline

- Review/motivation for old scheme; some problems.
- Overview/motivation for new scheme.
- “Modes”
- Related topics:
  - phasing, auto-correlations, dump control, expansion, VLBA correlator.
- Summary.

Review/motivation for old scheme

- Originally devised as method to handle ~40+ stations.
  - New Mexico Array.
  - VLBA; lower bandwidth, more antennas.
- Could tradeoff bandwidth for number of antennas, without rewiring the internals of the correlator.
- Could easily upgrade the size of the correlator to handle more antennas, without touching existing wiring.

Review/motivation for old scheme

- Consisted of 16 “sub-band correlators”.
- Each sub-band correlator correlated 1 sub-band pair from all antennas from all BB (a.k.a. I/F) pairs.
- Can spread lags across sub-band corrs with Station Board cross-bar switch.
- No cross-bar switch required between Station Boards and Baseline Boards…just “Fanout Boards” and cable.

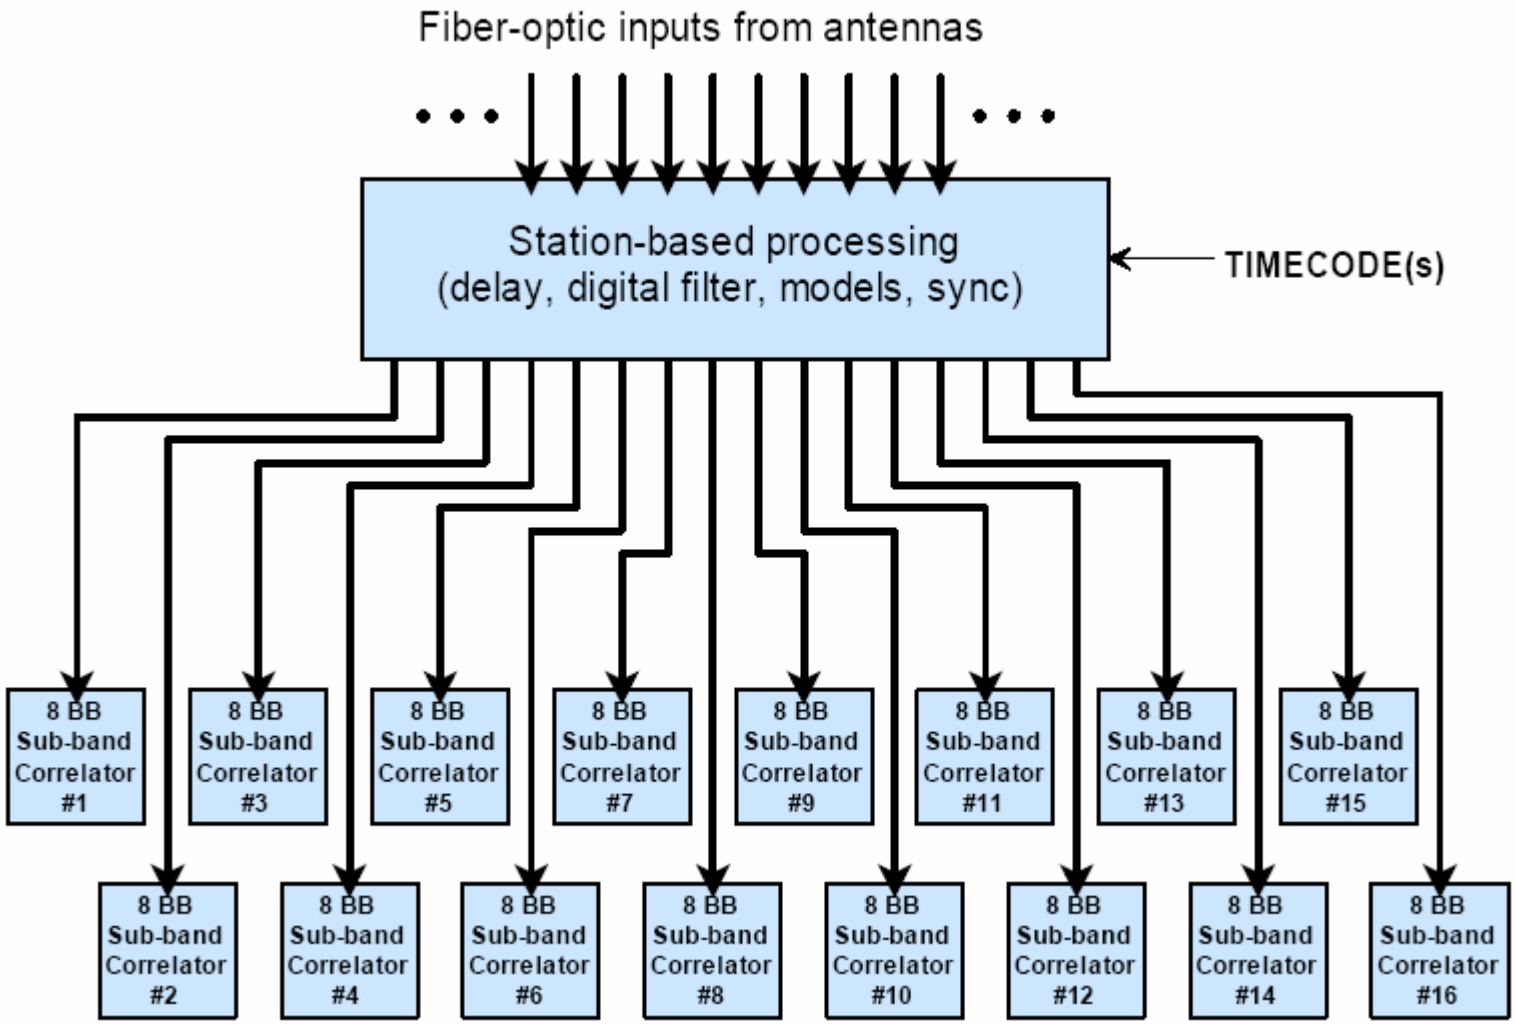
Fiber-optic inputs from antennas
Station-based processing
(delay, digital filter, models, sync)
TIMECODE(s)
8 BB Sub-band Correlator #1
8 BB Sub-band Correlator #3
8 BB Sub-band Correlator #5
8 BB Sub-band Correlator #7
8 BB Sub-band Correlator #9
8 BB Sub-band Correlator #11
8 BB Sub-band Correlator #13
8 BB Sub-band Correlator #15
8 BB Sub-band Correlator #2
8 BB Sub-band Correlator #4
8 BB Sub-band Correlator #6
8 BB Sub-band Correlator #8
8 BB Sub-band Correlator #10
8 BB Sub-band Correlator #12
8 BB Sub-band Correlator #14
8 BB Sub-band Correlator #16

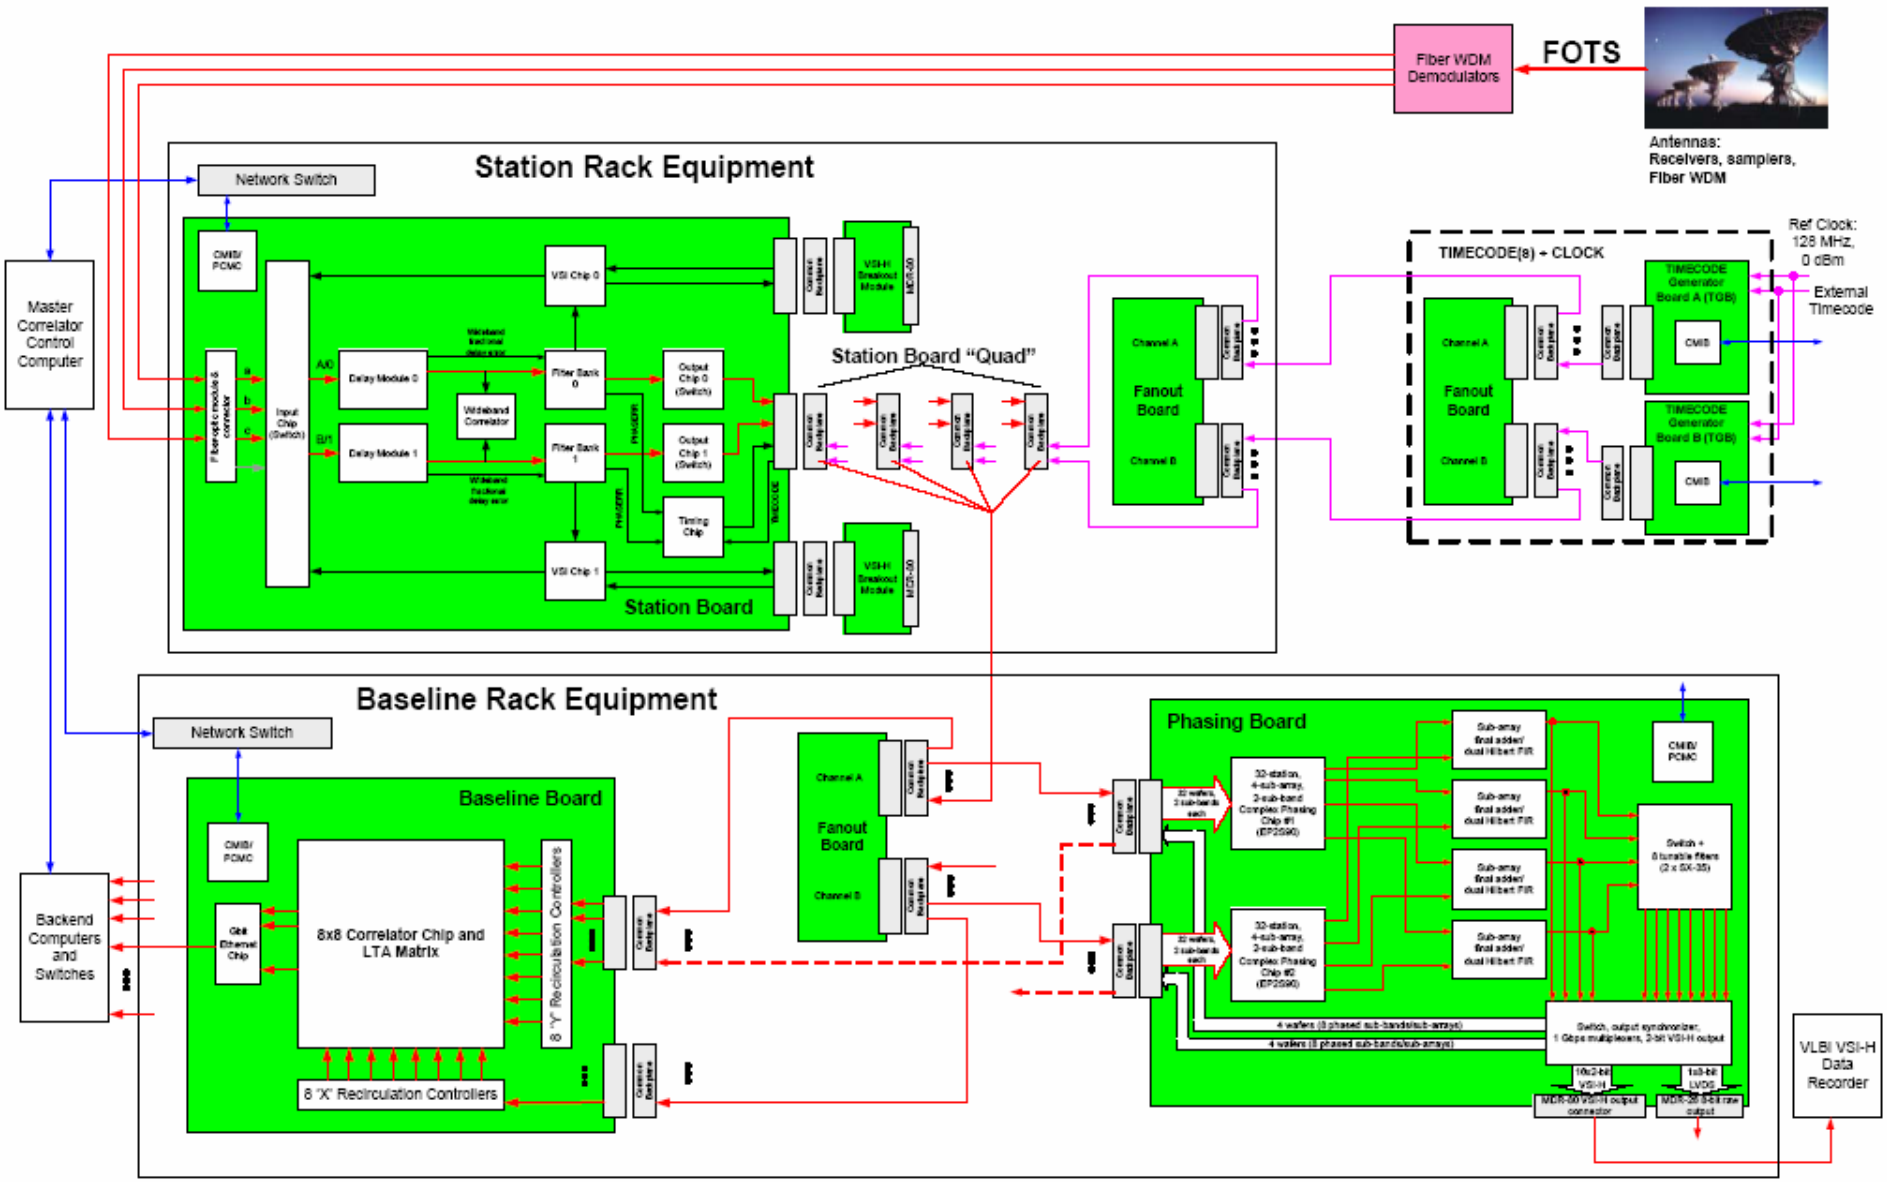

Figure 8-2 Correlator module connectivity diagram.

Review/motivation for old scheme

- Baseline Board to CBE switches can be distributed; 1 switch takes input from 1 Baseline Board in each sub-band correlator.
- 8 Station Racks, 16 boards in each.
- 16 Baseline Racks, 10 boards in each (for 32 station, 15 boards for 40 station). Each board does a triangle or parallelogram in the correlation matrix.
- VLBA correlator uses unused correlator capacity…but requires it to be at the VLA site.

Review/motivation for old scheme

- Some nagging problems:
  - Configurations for sub-bands in one BB pair affect what can be done for the same sub-band(s) in the other BB pair(s). User understanding and software—complex to understand.
  - The Baseline Boards that do the triangle part of the matrix are ½ unused…wasted hardware.
  - Due to Fanout Board drivers, if one Baseline Board is removed from the rack, several other boards' inputs are messed up.
  - Large volume of cable from 16 Fanout Boards to Baseline Boards in each rack…possible, but still voluminous (10x16=160 "short" cables in each baseline rack).
  - Phased-array output fixed by wiring…limited flexibility…no entirely satisfactory place to put the "Phasing Board".
  - "R2" recirculation expensive…limits spectral-line processing…makes software/understanding a bit painful.

Overview/motivation for new scheme

- Expansion beyond 32 stations largely unlikely.
  - New Mexico Array is dead.
  - VLBA real-time into correlator unlikely possibility in the near future.
  - Sense that VLBA correlator at the VLA site is undesirable.

- Budget very tight…no contingency left…desirable to free-up some contingency to prevent having to de-scope.

- Baseline Board has 8 x 4-pair inputs; connectivity and protocol allows it to be 16 x 2-pair inputs, 32 x 1-pair inputs, or 64 x single stream inputs.

Overview/motivation for new scheme

- Baseline Board array is much more efficient if X and Y inputs are the same since ½ board does "similar thing" or mirror image of the other ½ of the board (e.g. lead and lag lags).

- But, still want to keep expansion ability…hmmm.

- A thought for a new connectivity scheme…but some problems to overcome.

Overview/motivation for new scheme

- Requirements:
  - Can't re-spin the corr chip…too expensive and time consuming. Must be able to do all required pol'n products at widest bandwidth with no reduction in spectral-line capacity/flexibility.
  - Minimal to no h/w and s/w impact…must minimize schedule impact since it is so late in the project, and schedule has already slipped substantially.
  - Science improvement must be positive…no loss in flexibility to meet EVLA <=32 station requirements.
  - Contingency savings must be substantial…can't be *more* expensive.

Overview/motivation for new scheme

- Requirements cont'd:
  - Simplification of system…less hardware…easier to understand is always desirable.
  - Deal with that nagging black sheep called "phasing" (although this evolved later…).
  - Desirable for expansion still to be possible.
  - Desirable for non-real-time VLBA correlation…but not at VLA site.

Overview/motivation for new scheme

- Result:
  - "Quadrant" correlator…each quadrant correlates one BB pair.
  - Distributed X-bar switch between Station Boards and Baseline Boards to overcome fixed quadrant limitations…restores flexibility of original design. Possible with flexibility of Meritec HM "wafer" cabling…no monolithic backplanes or backplane wiring.
  - "R2" recirculation *required* and comes for "free"… ½ of recirculation done with external RAM, R2 recirculation done with internal FPGA RAM.
    - Just barely fits in affordable Recirculation FPGA (EP2S30…)…although lots of spare logic left over. Required dump time of 200 usec to LTA proven in real h/w.
    - Saves huge cost for additional external RAM. Freed-up contingency more than pays for the cost.

Overview/motivation for new scheme

- Result cont'd:
  - Reduces the number of baseline racks from 16 to 8. Each rack is fully utilized with 16 boards.
  - Reduces the number of Baseline Boards from 160 to 128…most substantial cost saving.
  - Reduces the number of cables from ~3000 to ~1000…higher system reliability, vastly reduced rack assembly complexity.

Overview/motivation for new scheme

- Result cont'd:
  - More seamless and logical tradeoff of bandwidth for #spectral channels. Sub-bands in one BB don't affect other BBs. Each sub-band pair configuration is now independent of other sub-band pairs.
- New observational capabilities:
  - Potentially more pulsars, mixed normal/pulsar observing.
  - Recirculation available on all streams, more flexibility: each sub-band can have different bandwidth and recirc factor.
  - Phase entire bandwidth all the time…several ways to use phased data.
- Still allows for expansion > 32 antennas…partial or full bandwidth.
- Separate standalone VLBA correlator makes more sense.

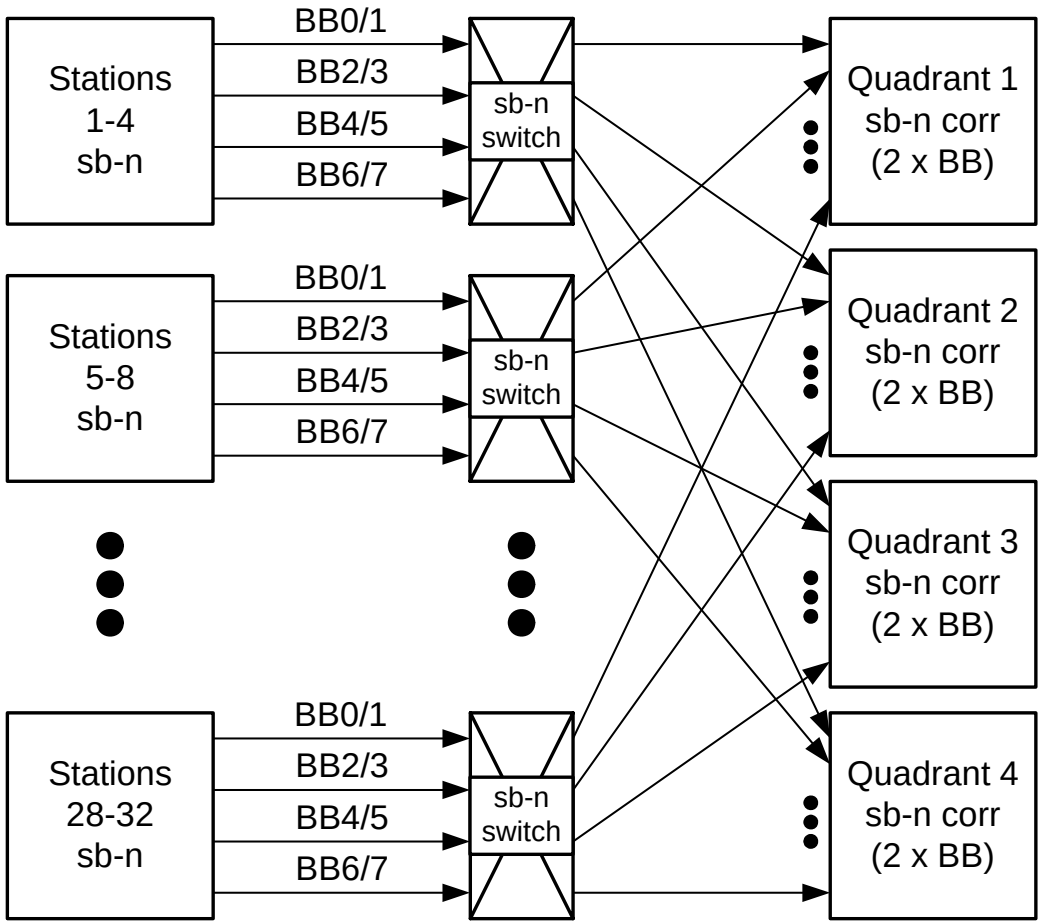

EVLA Correlator new connectivity: simplified diagram (sub-band-n)

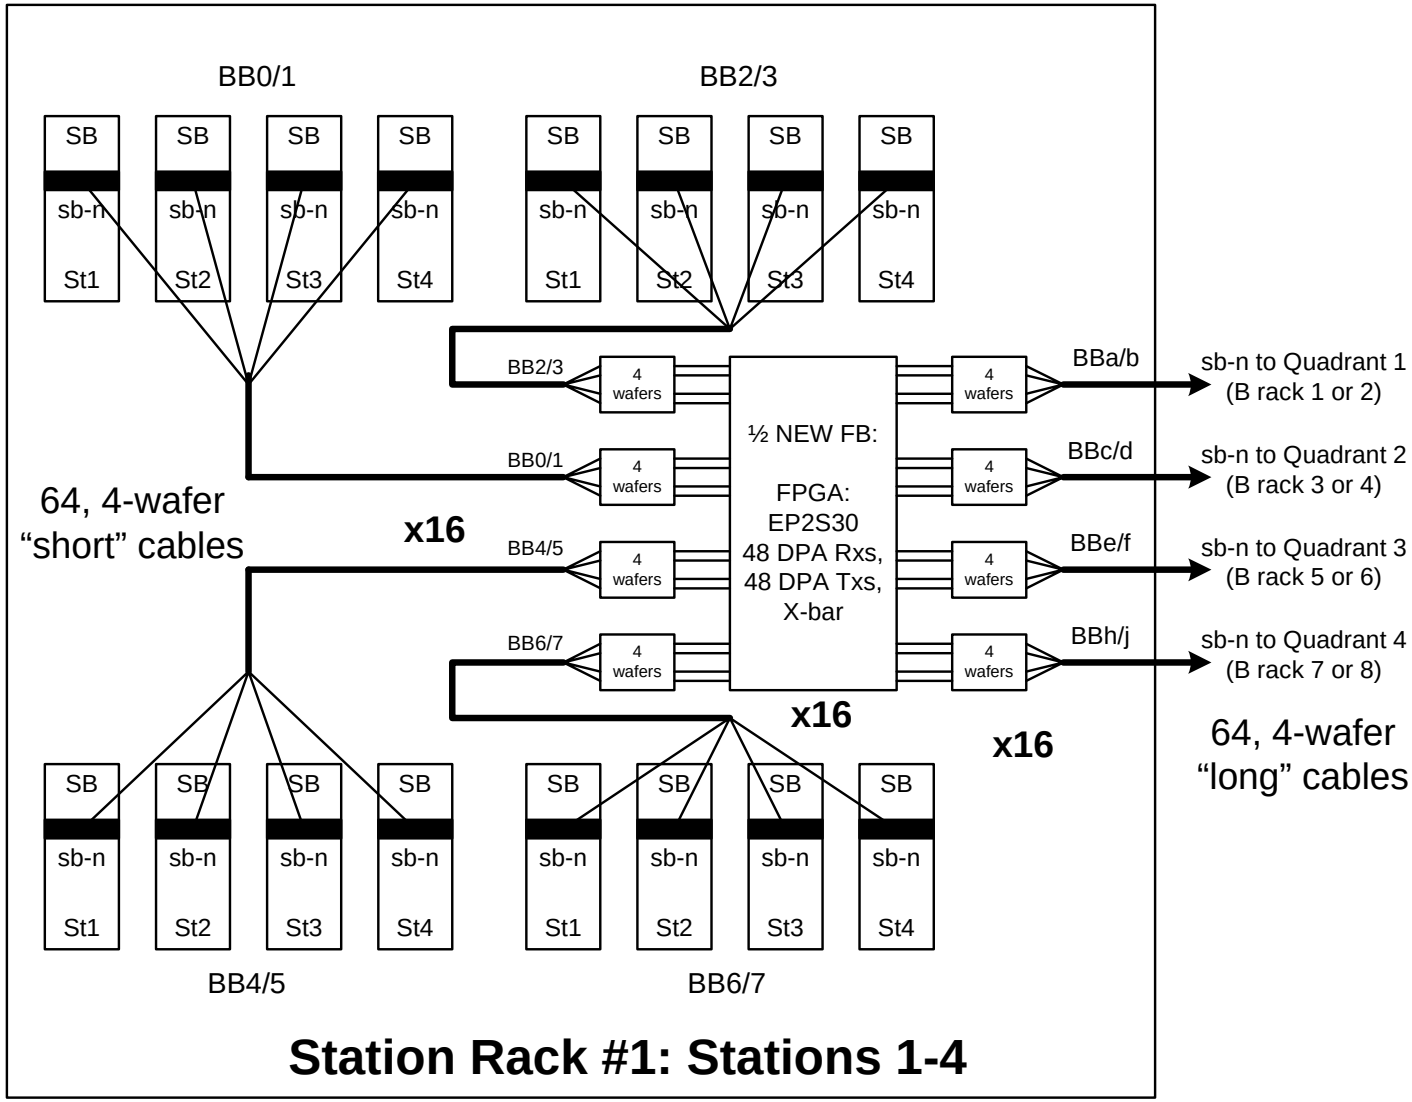
BB0/1
BB2/3
SB
sb-n
St1
St2
St3
St4
BB2/3
BB0/1
BB4/5
BB6/7
4 wafers
½ NEW FB:
FPGA:
EP2S30
48 DPA Rxs,
48 DPA Txs,
X-bar
BBa/b
BBc/d
BBe/f
BBh/j
sb-n to Quadrant 1
(B rack 1 or 2)
sb-n to Quadrant 2
(B rack 3 or 4)
sb-n to Quadrant 3
(B rack 5 or 6)
sb-n to Quadrant 4
(B rack 7 or 8)
64, 4-wafer
“short” cables
x16
x16
x16
64, 4-wafer
“long” cables
Station Rack #1: Stations 1-4

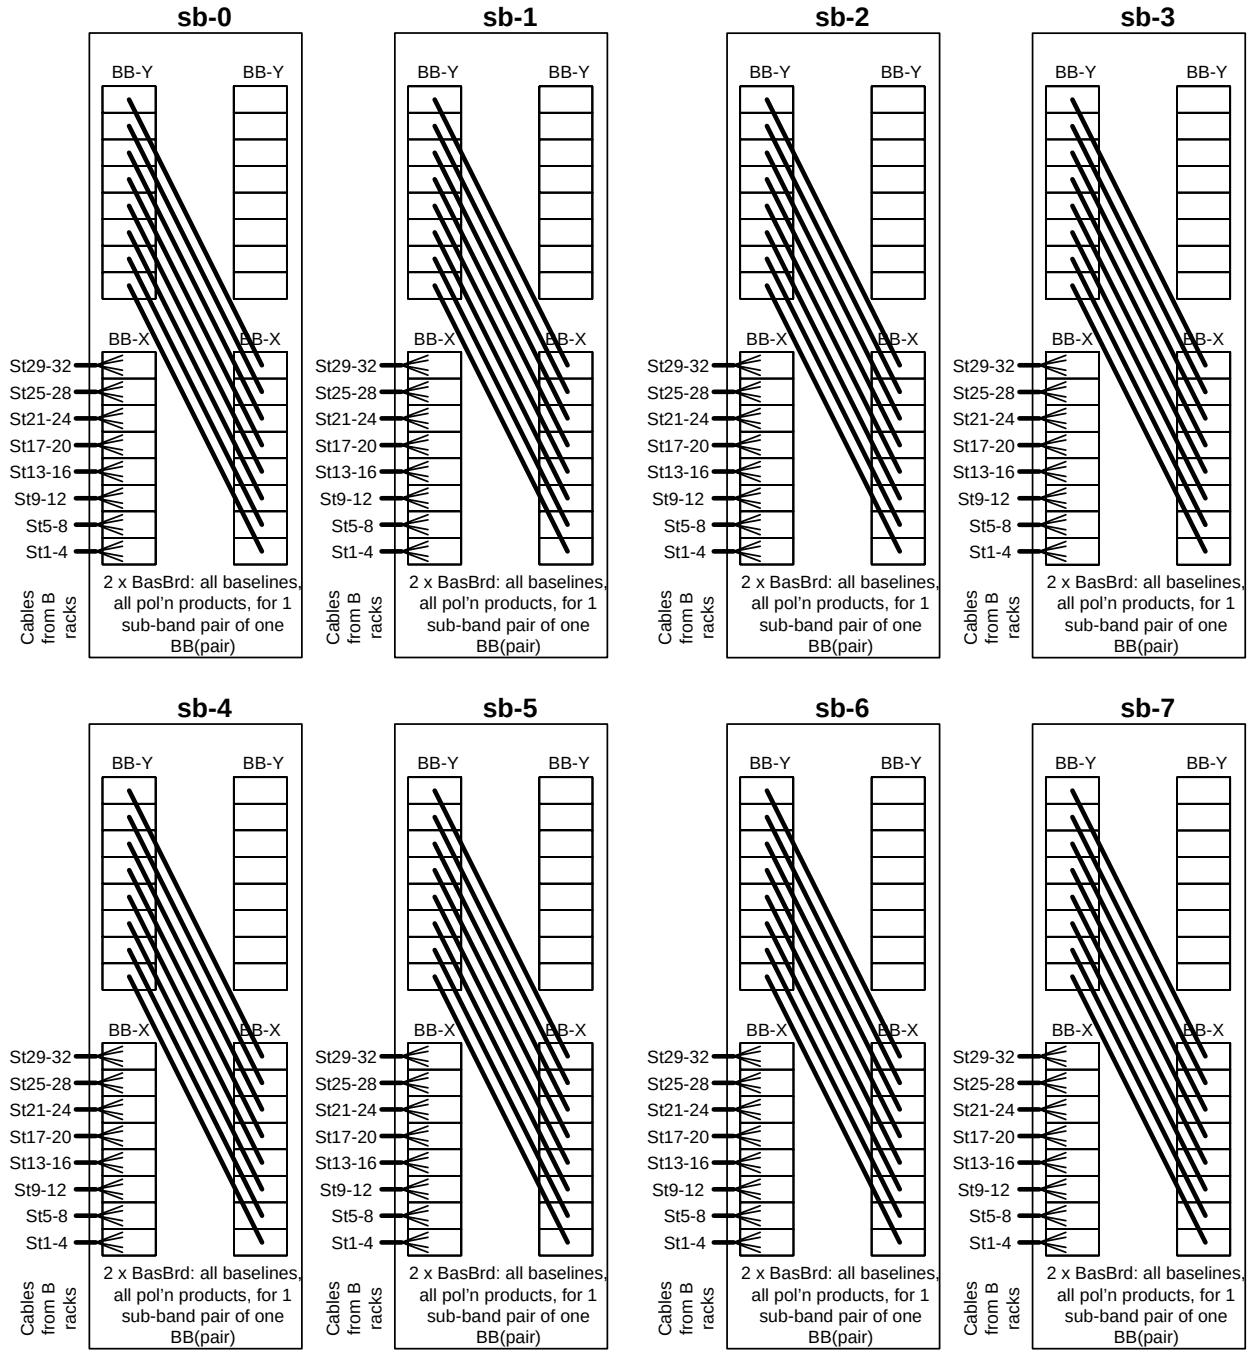

Baseline Rack #1: ½ quadrant: all baselines for sb-0 to sb-7 of one BB pair (2 GHz total)

Baseline Board #1: 1st half of baselines, all pol'n products, 1 sub-band pair

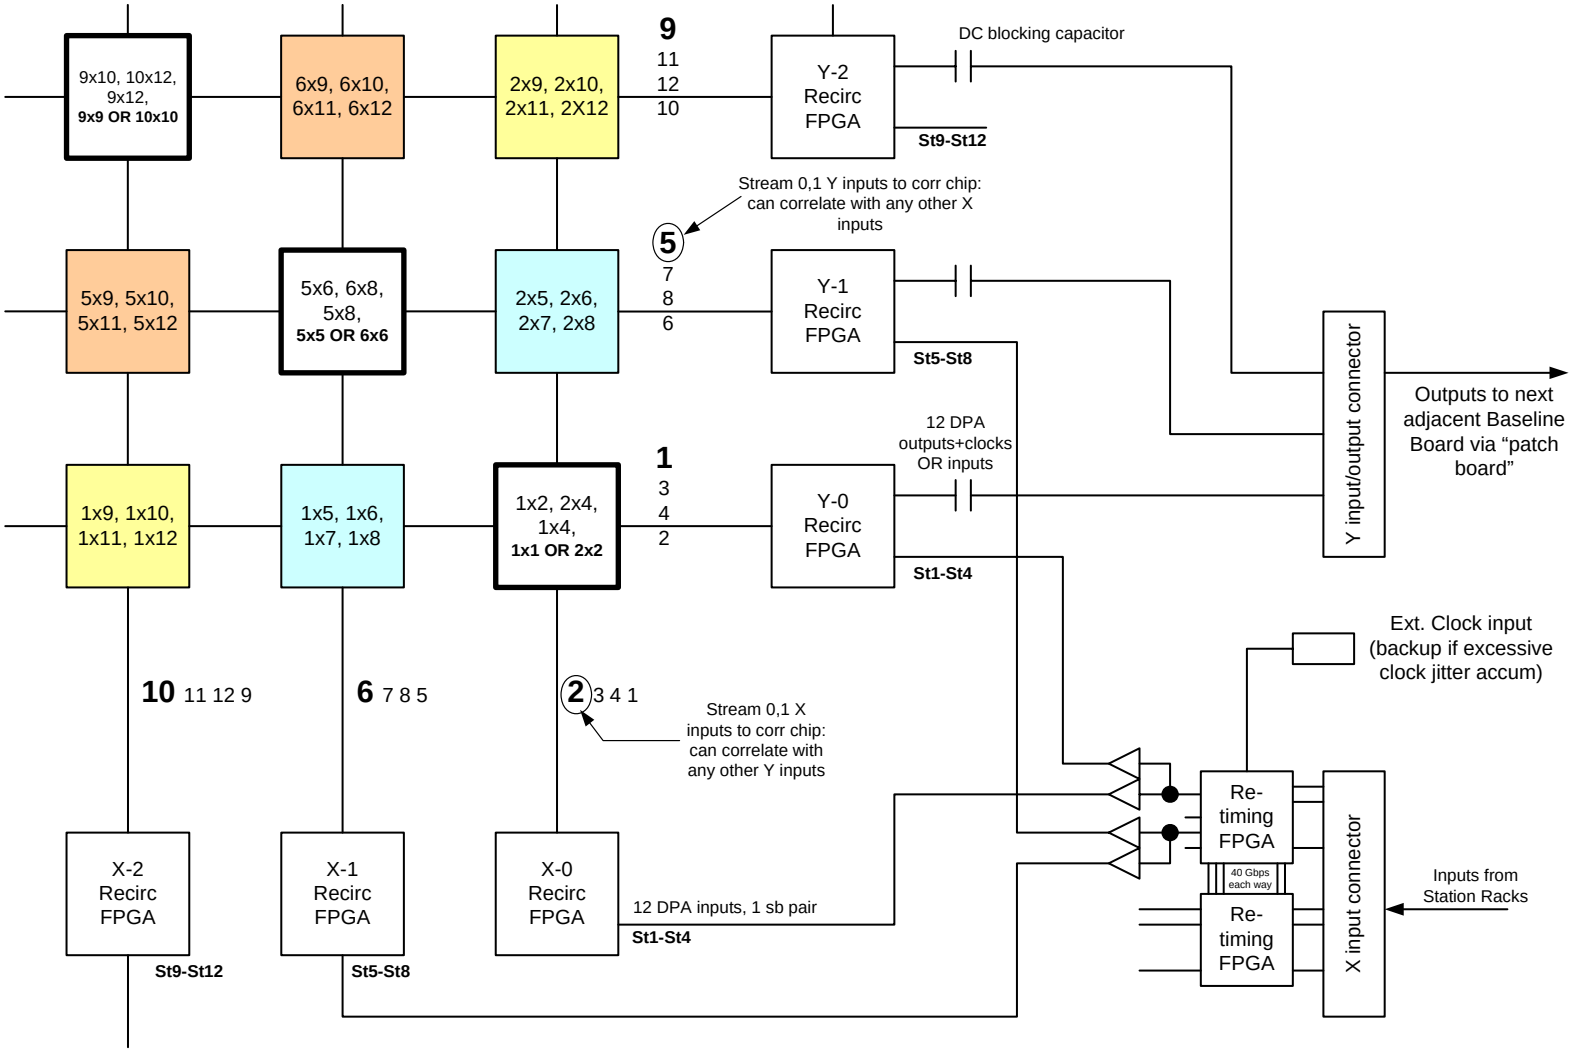

Baseline Board #2: 2nd half of baselines, all pol'n products, 1 sub-band pair

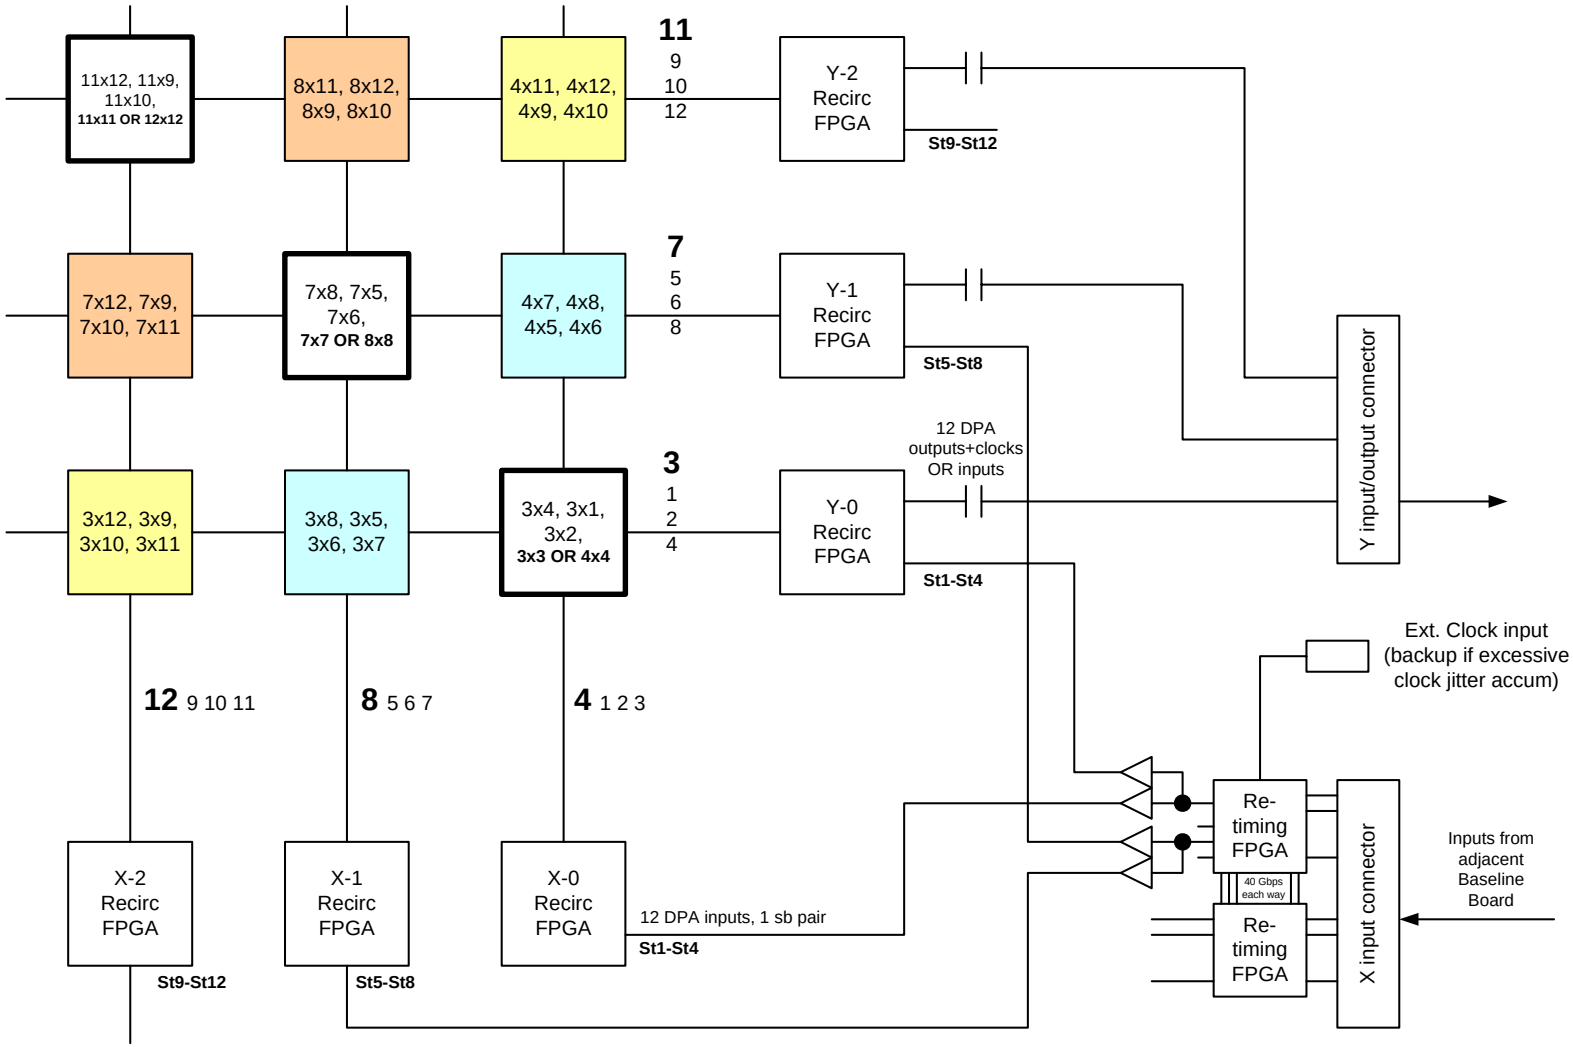

Example: Corr Chip detail

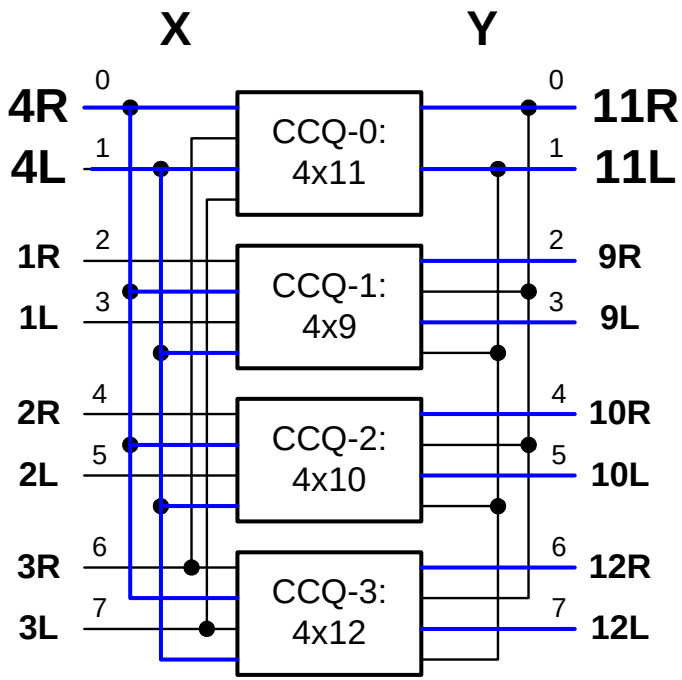

4x11, 4x12, 4x9, 4x10

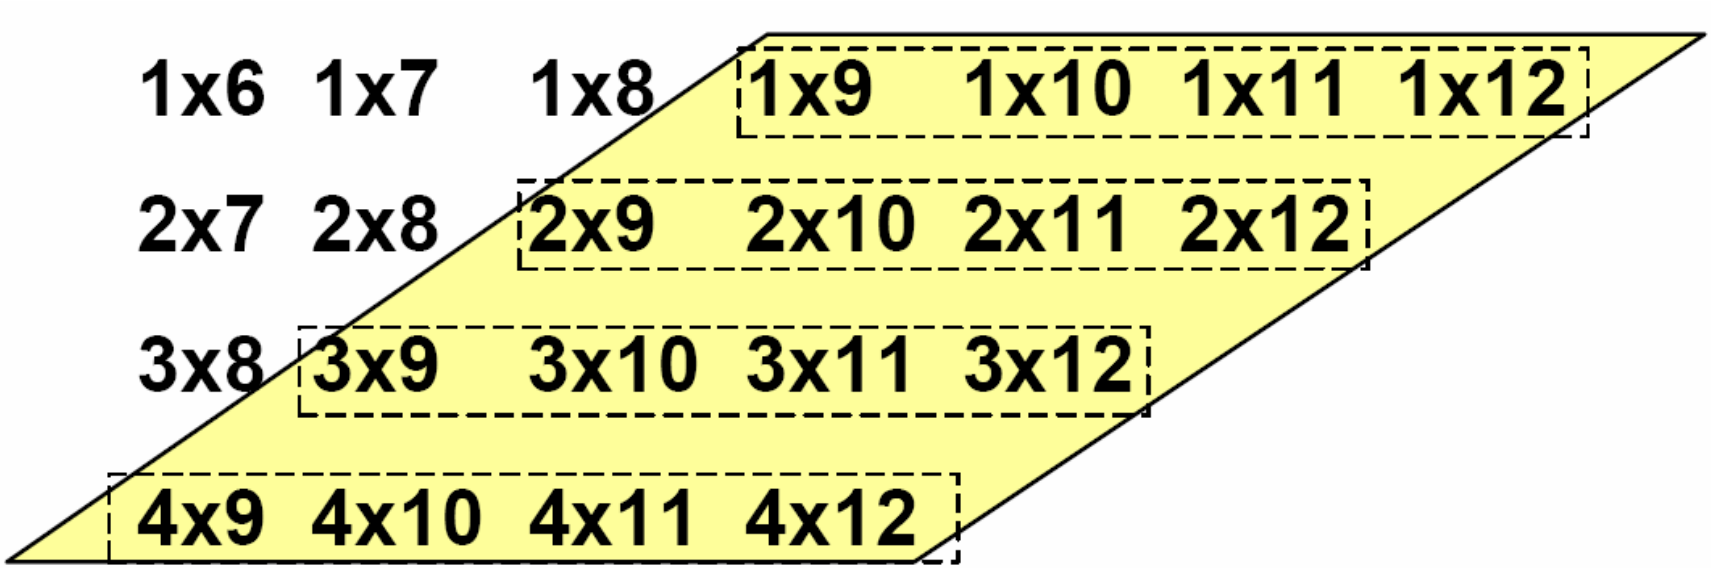
1x6 1x7 1x8 1x9 1x10 1x11 1x12
2x7 2x8 2x9 2x10 2x11 2x12
3x8 3x9 3x10 3x11 3x12
4x9 4x10 4x11 4x12

Overview/motivation for new scheme

- Preliminary informal review of new scheme held in Socorro on April 24, 25.

- Additional requirements:
  - Use the “re-timing FPGA” as a full cross-bar switch to overcome limitations of sub-arrays…done…sub-arrays are flexible, in quanta of <=4 antennas.
  - Study ways in which all autocorrelation products can be produced.
    - Station Board already acquires all lag-0 auto-products with 10 msec time resolution.
    - Use unused capacity on Baseline Board (Sonja, NRC-EVLA Memo# 029)
    - Do in re-timing FPGA or X-bar FPGA if required.

Overview/motivation for new scheme

- Since April, all required connections, and then some, on the Baseline Board have been completed…still some "cleanup".
  - Signal integrity analysis and of course, proto testing still required.

- Logic design of re-timing FPGA with full cross-bar finished and PAR'd. Still needs functional testing.
  - 80 lines of 512 Mbps DDR running between FPGAs. 1.8 V logic.
  - Analysis indicates speed can be met with expected $C_{loading}$

- Design enhancements to the Recirc FPGAs, GigE FPGAs (pins only) now complete and PAR'd.

Overview/motivation for new scheme

- H/W impacts:
  - Station Board – no changes.
  - Fanout Board – gone.
  - X-bar Board – need to develop; 2 FPGAs, power supplies, on 6U x 160 mm standard board (maybe other h/w to support doing something with phased data). Set connections via HM Gbps CTRL channel, from Station Board.
  - Baseline Board – 2 new FPGAs and input section; no changes to corr chip/LTA array and GigE FPGA.
  - Hi-speed Meritec cabling…~1000 cables…waiting for go-ahead to place order.
  - Simpler baseline racks. More complex station racks (8 X-bar Boards), but infrastructure already in place.
  - Requires monolithic GigE switch from Baseline Boards to CBE…~$130k.

Overview/motivation for new scheme

- S/W impacts:
  - Station Board – minimal…registers in "Timing Chip" for HM Gbps CTRL channel to X-bar switch commands.
  - Baseline Board – minimal…additional MAHs and GUIs for re-timing FPGA (also to support phasing).
  - CBE – minimal to none…still get packets of lag frames…still have to assemble them…might want to accept phased data packets?
  - MCCC – mapper…major…but not much work done on it…new scheme is simpler to understand…
  - CPCC – simpler…fewer racks/modules to control…each CPCC now fits in one rack-mount PC or CPCI crate.

"Modes"

- From NRC-EVLA Memo# 028
- Useful for discussion…not meant to place restrictions on how the correlator can be configured.
- Mixture of "modes" possible.

64SB-4PP:3IQ/4RQ:FB:NR

- Full bandwidth, all sub-bands, all products.
- Crux of new scheme since requires largest number of independent cross-corrs.
- Refer to previous BB figures.

64SB-2PP:3IQ/4RQ:FB:NR

- Full bandwidth, all sub-bands, except only 2 pol'n products.
- Same as last mode, except 2X spectral channels per product.

64SB-4PP:3IQ/4RQ:RB:RC

- All I/Fs, all sub-bands active.
- Reduced bandwidth per sub-band (or not).
- Recirculation active…could be different bandwidth and factor for each sub-band.
  - Subject to “DUMPTRIG” gen limitations on the Station Board.
- Up to 16 k channels per product, 4 Mchannels/baseline.

64SB-2PP:3IQ/4RQ:RB:RC

- Same as previous, except 2 pol’n products.
- More “real lags” per product; 2X BW at 16k channels each.

64SB-1PP:3IQ/4RQ:RB:RC

- All sub-bands active, 1 pol'n product.
- Use external recirculation RAM only…262,144 channels per product. 4X as many "real lags" vs 4 pol'n products.

16SB-4PP:3IQ/4RQ:FB:SRC

- Process only one I/F pair (4 GHz/antenna).
- 1 quadrant in use…other quadrants, via X-bar switch, can be used to provide 4X+ spectral channels using "static" recirculation.
- 256 channels minimum per product (w/o recirc, 128 MHz sb BW).

16SB-2PP:3IQ/4RQ:FB:SRC

- Same as previous, except 2 pol'n products.
- 2X as many channels.

16SB-4PP:3IQ/4RQ:RB:RC

- <4 GHz per antenna, 16 sub-bands active.
- Each quadrant fed the same data…one pol'n product per quadrant.
- Use dynamic recirculation with external RAM for up to 262,144 channels per product (4M channels/baseline).

16SB-4PP:3IQ/7RQ:RB:RC

- 7-bit re-quantization on 16 sub-bands, 1 I/F pair.
- Each quadrant processes 4 cross-products, one pol'n product.
- 8 k channels/baseline (64 MHz sub-bands); up to 16 k channels per product at reduced sub-band bandwidth.
- Up to 262,144 channels/product if 1 pol'n product.

16SB-4PP:8IQ/7RQ:RB:RC

- 8-bit initial sampling…different pol'n on each Station Board.
- Use the X-bar to route sub-band pairs from different input wafers to one output wafer.
  - Requires PHASERR switching (more logic…no problem).
  - PHASEMOD switching…but easiest to have SB CMIB gen all phase models for each pol'n.
- Otherwise, same as previous 7-bit processing.

Related Topics—Phasing

- Possibility of phasing in re-timing FPGAs now very real.
  - Can phase-up all available bandwidth all the time.
  - Existing S30 FPGA should allow 1 sub-array per sub-band (possibly 2), with output via GigE FPGA (S30 FPGA is 55% used with re-timing and cross-bar). *Possibly output to HM Gbps…next talk.*
  - If S60 FPGA used, then should get 2 sub-arrays per sub-band…possibly 4. Also, output via GigE FPGA and HM Gbps.
  - PAR is done with S30, with migration path to S60…still need to do the phasing design and PAR. S30 may be sufficient for 2 sub-arrays per sub-band.
  - Various options on what to do with the data…not black sheep anymore!
    - Need to decide to nail-down X-bar board requirements, sys requirements.

Baseline Board Phasing Connections

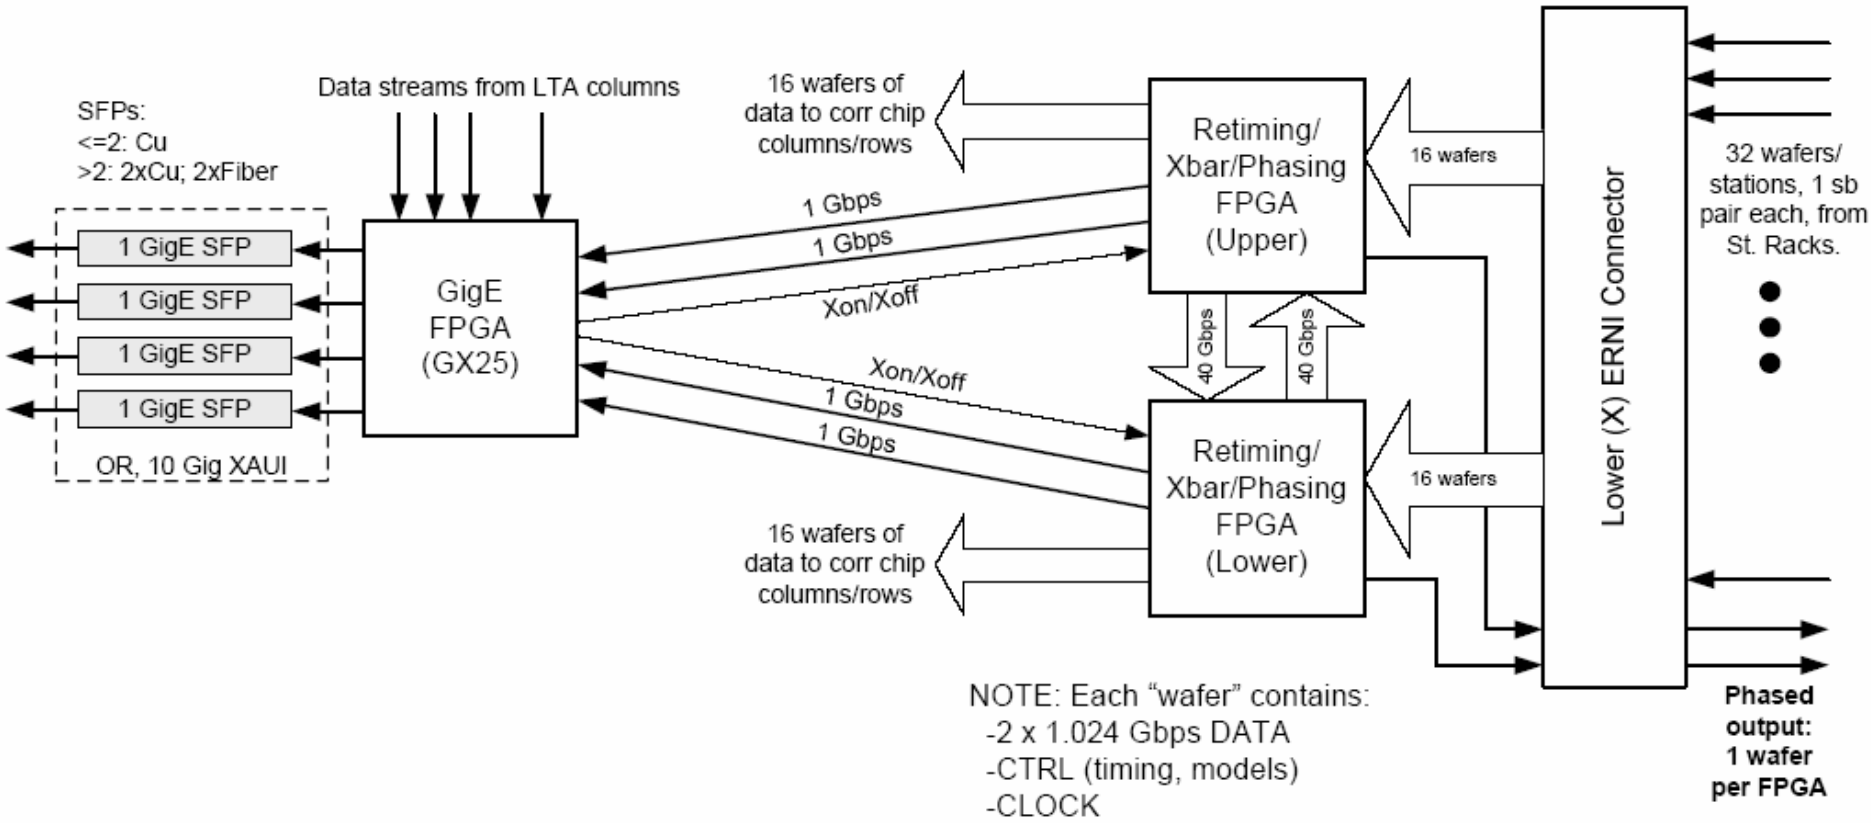

- What to do with ***HM Gbps*** phased output?
  - Gotta do something with it…outputs are LVDS so must go into X-bar board for signal amplification at least.
  - Use the X-bar board to switch/select phased outputs of interest to go to VLBI recorder…but then what?
    - 1 X-bar board in each baseline rack.
    - Design logic, connectors, 10 GigE XAUI on X-bar to handle it? Populate only the required number of boards with this stuff. This is a separate build…mfg tooling, testing etc. Best to have CMIB on as well…more development effort.
  - How about feeding the X-bar board "switched phase outputs" from multiple baseline racks into one or more "special purpose" Station Boards?
    - VSI inputs to Xilinx FPGAs are supposed to be able to handle 1 Gbps mux data (requires higher speed grade than used for Filter Chip)…need to test.
    - Use the filters to do additional filtering…go to VSI outputs. Need external board to go to 10 GigE…"iBob" board?
    - Can keep the X-bar board simple…2 FPGAs, drivers, power supplies.

X-bar Board Concept–Phased outputs

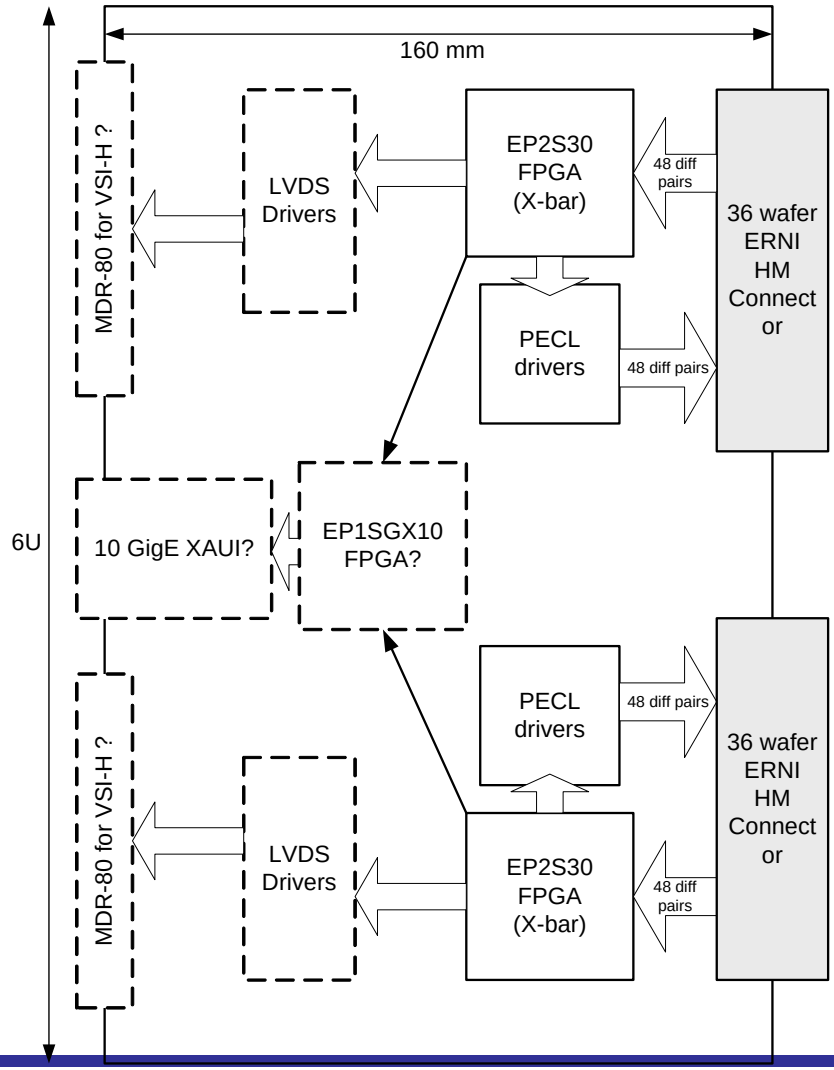

X-bar Board Concept–Simple

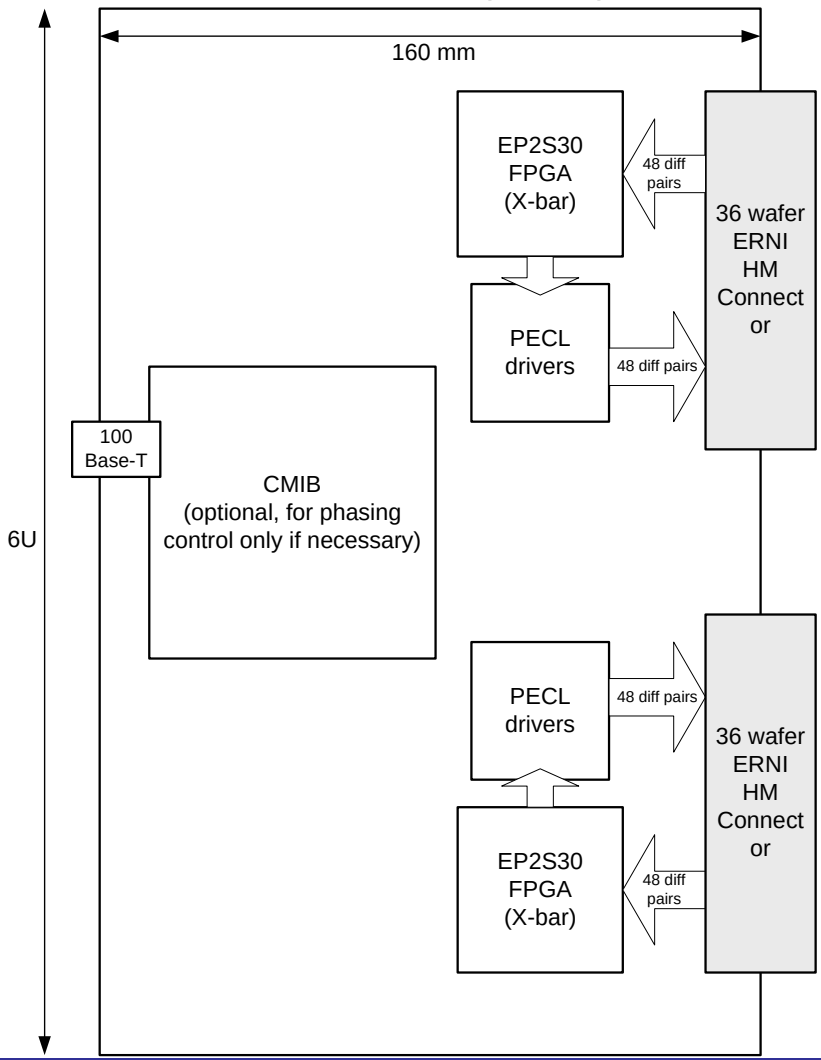

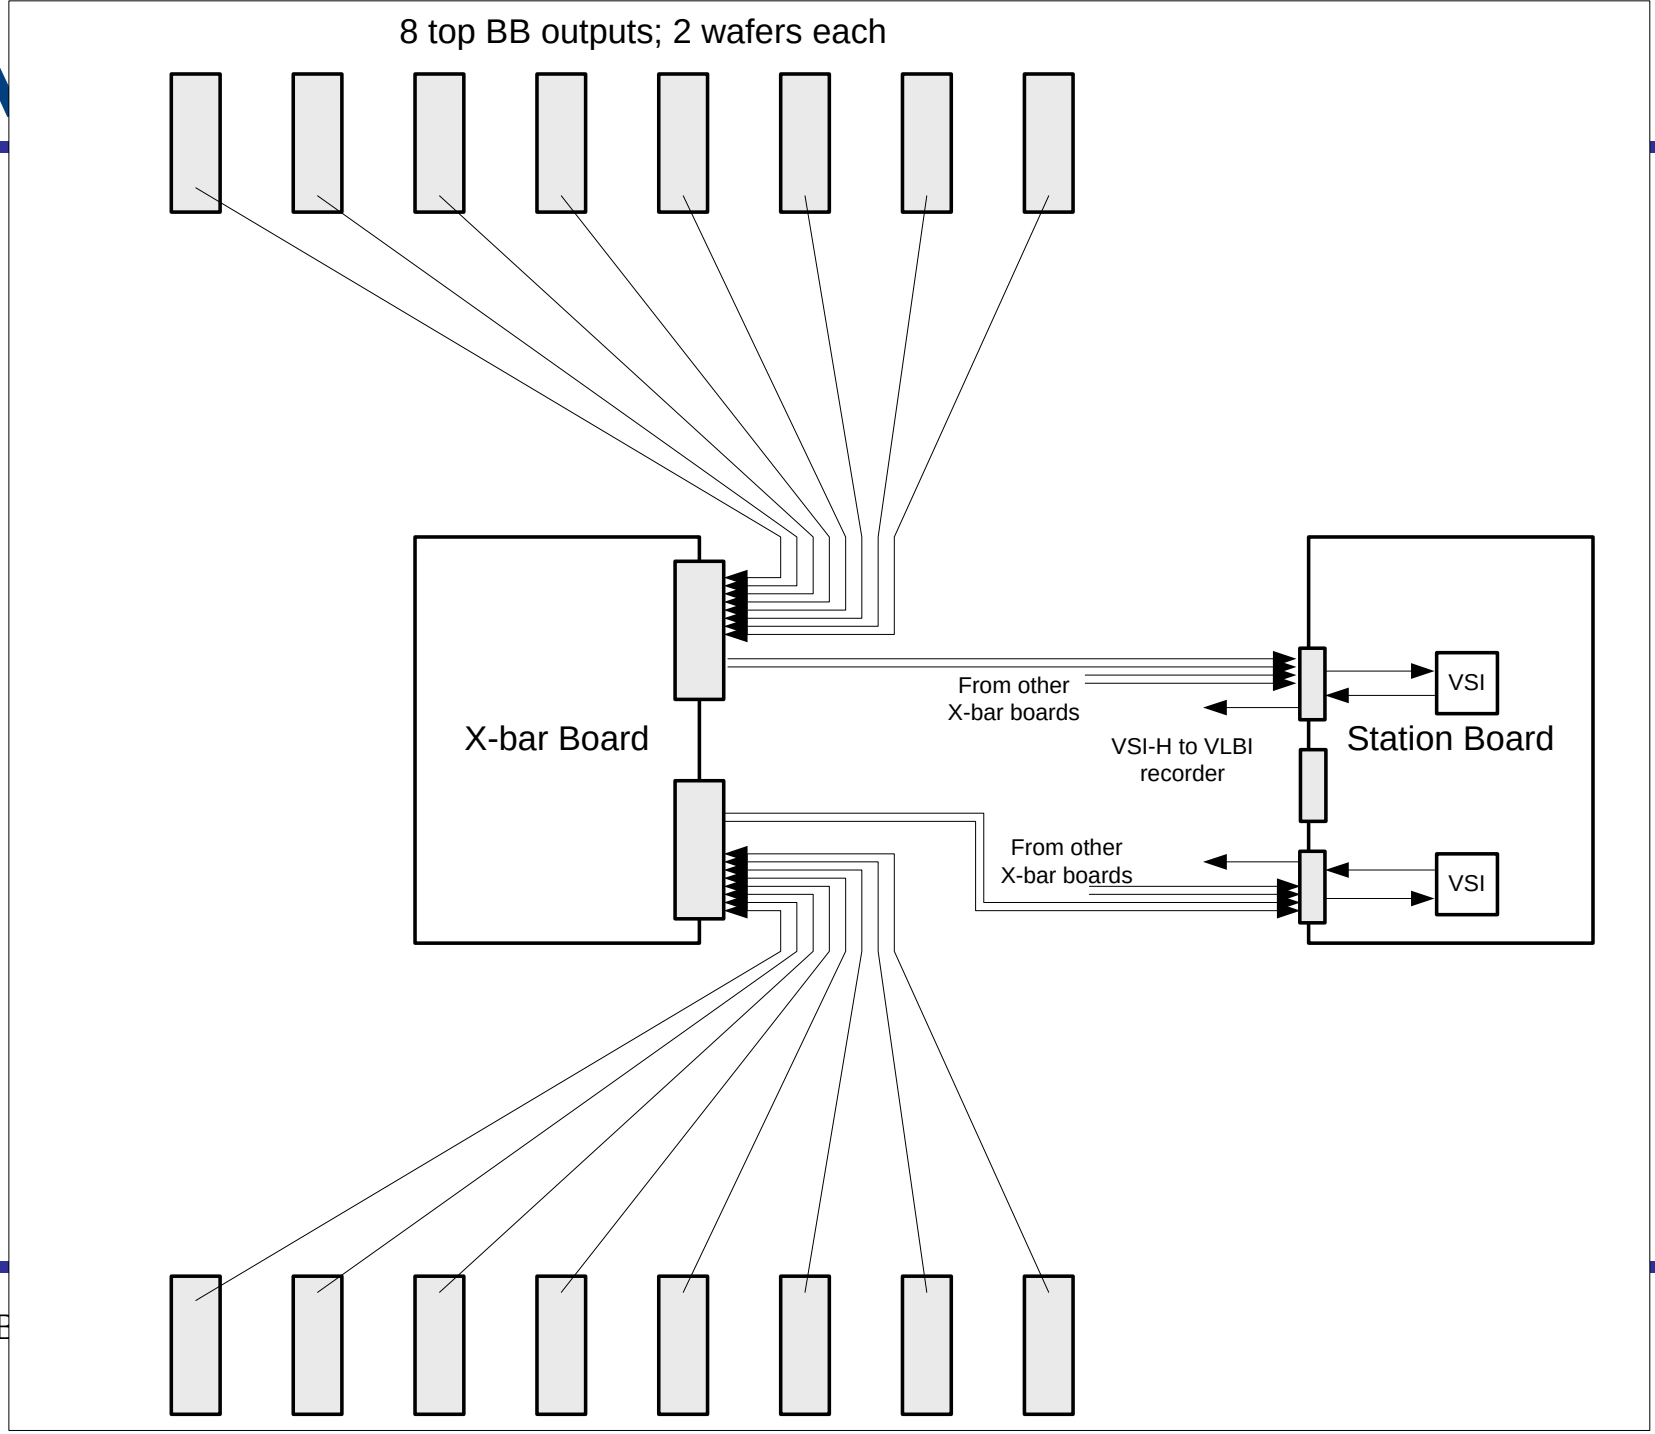
8 top BB outputs; 2 wafers each
X-bar Board
From other
X-bar boards
VSI-H to VLBI
recorder
Station Board
VSI
From other
X-bar boards
VSI

Phased output/VLBI recorder(s)

- 1 Station Board can handle BW for 2 VSI outputs:
  - Each VSI output is 32 diff pairs, or 16 2-bit streams at up to 128 MHz BW per stream (16 x 128MHz = 2.048 GHz).
  - So, 1 Station Board is enough BW for 4 GHz…but has only 32 sub-band pairs in to select from…need to decide X-bar topology to allow dynamic access to desired sub-bands (e.g. all sbs of one BB pair, and some sbs of others?).
  - Likely need X-bar board in front of Station Board to repeat/amplify signal.
- Where to put the Station Board?
  - Could use unused Station Boards in existing racks?
  - Separate rack with X-bar board, and VLBI recorders?

Related topics – auto-corr

- Diagonal corr chips obtain baselines as well as auto-corrs.
- Leftover bits can do some, but not all auto-corrs:
  - E.g. 1Rx1R, 2Lx2L, 1Rx1L, 2Lx2R
  - Can do only two auto-corr products per antenna.
- Instead/in addition, for interference detection purposes, could put additional or unused Baseline Boards into auto-corr mode:
  - Each BB can do 64 independent auto-corrs, 2k channels each.

Dump Control

- Each Recirc FPGA handles 4 antennas.
- Right now, one dump control signal is selected and used.
- Could upgrade to one dump control signal per antenna…but dump signals at lowest level must be harmonically related.
- Each sub-band could be different…subject to limitations in Station Board "Timing Chip" and real-time S/W.

Expansion

- Can expand beyond 32 antennas…8 station increments.
- Could expand 1, 2, 3, or 4 I/Fs.
- Y connector on Baseline Board can be input instead of output with different (already existing and compiled) Y Recirc FPGA.

VLBA Correlator

- Separate standalone correlator. 2 racks=16 stations, 2 GHz/pol'n, 8 k channels/baseline at full bandwidth, more with recirculation.
- Peel-off 1 station rack from EVLA, and 16 additional Baseline Boards.
- Additional cost over not doing it is ~$250k…gotta buy all the corr chips anyway (although, reduces spares).
- Still leaves 28 stations for EVLA.

VLBA Correlator

- At a meeting on April 25th, it was decided:
  - Continue to leave VLBA corr options open in our development.
  - NRAO to proceed in parallel with feasibility/construction of s/w VLBA correlator.
  - Once WIDAR components finalized and production schedule solidified, decide on whether to build standalone WIDAR VLBA correlator, or continue with s/w correlator.
    - Wrinkle: iSine needs to know NOW how many chips to build. 12k chips is ok for EVLA, eMERLIN qtys x 1.05 x 1.03 (9984 chips…leaving 2k spares), but is a bit thin if another 16/17 boards built (1024 chips).
    - If corr chips ordered LATER, MOQ=3500. NOW, 1k chips costs $161k.
    - We're about to place order for production-qty FPGAs…money spending requirements.

Summary

- Reviewed old scheme…some nagging problems.
- Overview/motivation of new scheme.
  - Better capability…sub-bands independent…phase all bandwidth.
  - Buy back contingency…simpler, more reliable system.
  - Can't do all auto-corr products on Baseline Board.
- Modes.
- Related topics: phasing, auto-corr, dump control, expansion, VLBA corr.