U.S. Radio/Millimeter/Submillimeter Science Futures II

Correlator for Next Generation VLA

Larry D'Addario

2016 August 3

Top Level Architecture and Some Notation

(Pure) FX Architecture

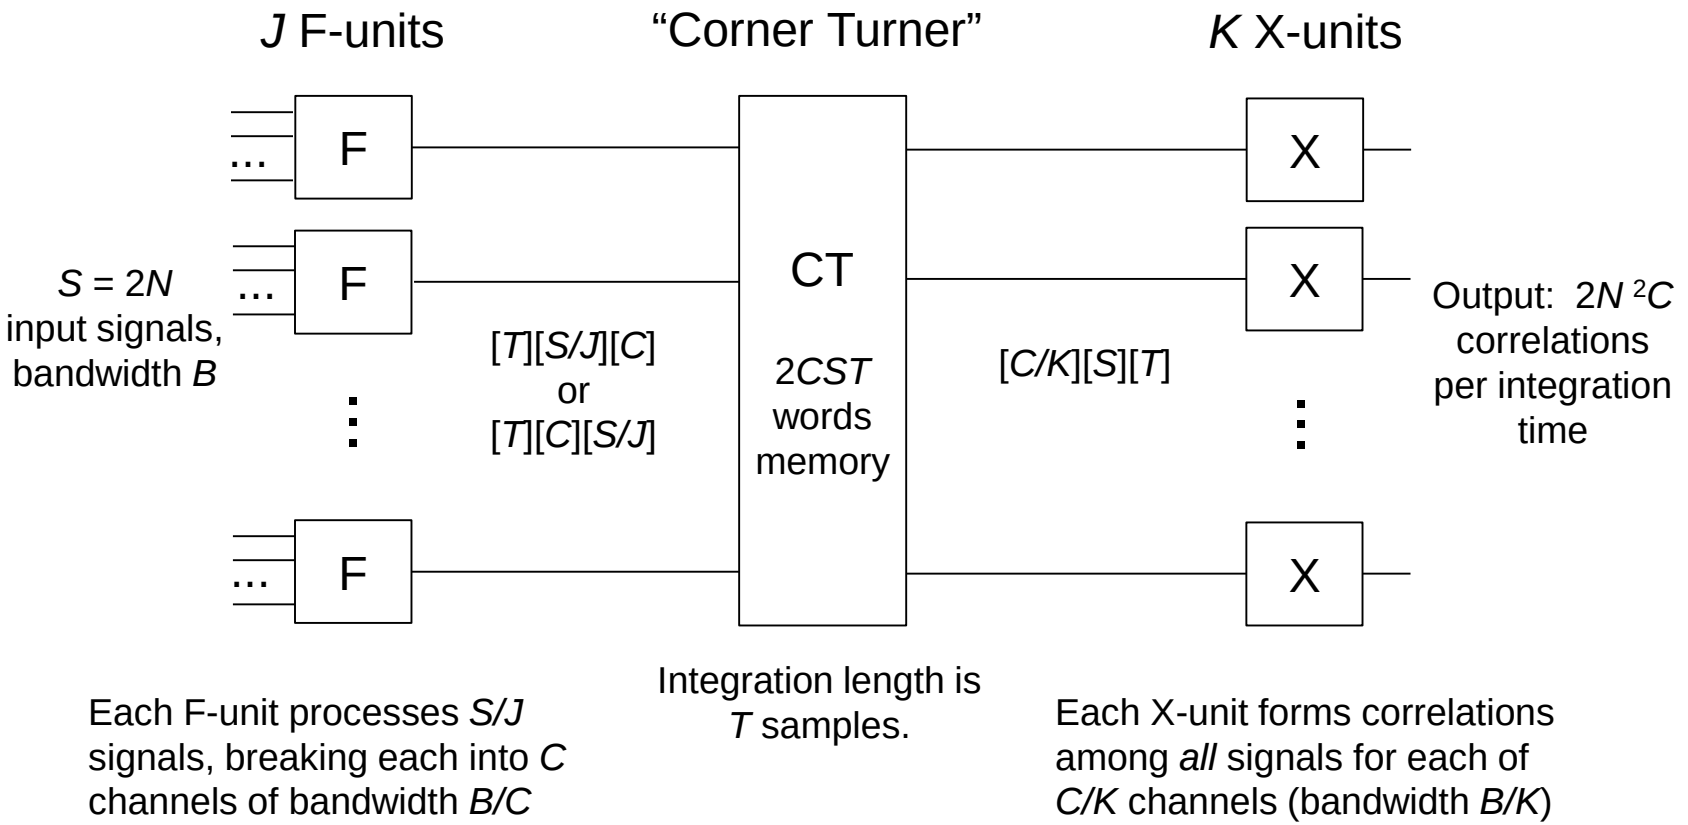

Main Specifications (worst case?)

- N = 355 dual-polarization antennas (710 input signals)
- B = 50 GHz processed bandwidth (1.2-48 GHz or 70-120 GHz)
  - will this really be supported by the front ends?
  - is it really useful?
  - is it worth the cost?
- C = (TBD) spectral channels – see discussion later

Resources Needed

To construct a correlator (or any other digital signal processing machine) the main resources needed are:

- Computation (arithmetic: add, subtract, multiply)
- Memory (working storage for intermediate results)
- Input and Output (external and internal)

The minimum *computation rate* required is independent of architecture and technology.

The external *input and output rates* are also independent of architecture, but *internal I/O* requirements depend on architecture.

The *memory* (quantity and bandwidth) needed depends strongly on architecture.

Computation Rates

Correlator size cannot be reliably estimated from computation rate alone!

Table 1: Computation Rates

		future	*future*	*existing*	*existing*	*re-baselined*		*original*
		NGVLA	*DSA-low*	*JVLA*	*ALMA*	*SKA1-low*	*SKA1-mid*	*SKA1-mid*
Main specifications								
antennas *N*		355	2048	32	64	512	197	254
bandwidth *B*	Hz	5.00E+10	7.00E+07	8.00E+09	8.00E+09	3.00E+08	5.00E+09	5.00E+09
channels *C*		262,144	14,000	16,384	8,192	65,536	65,536	262,144
Computation rates								
Rf1	FIR/s	**1.42E+14**	1.15E+12	2.05E+12	4.10E+12	1.23E+12	7.88E+12	1.02E+13
Rf2	BF/s	**1.60E+14**	9.87E+11	1.79E+12	3.33E+12	1.23E+12	7.88E+12	1.14E+13
Rx	CMAC/s	**1.26E+16**	5.87E+14	1.64E+13	6.55E+13	1.57E+14	3.88E+14	6.45E+14

Ratios

		NGVLA/ JVLA	*NGVLA/ ALMA*	*NGVLA/ SKAmOrig*
Rf1	FIR/s	69.34	34.67	13.98
Rf2	BF/s	89.15	48.00	13.98
Rx	CMAC/s	769.20	192.30	19.53

More Specifications

- To work out the I/O and memory requirements, we need more system-level specifications:
  - **Channel bandwidth** (affects memory)
  - **Minimum integrating time** (affects output bandwidth)
  - **Maximum integrating time** (affects memory)
- From u,v plane smearing (10% visibility loss at 3 dB point of primary beam):

  d = 18m, D = 300 km → T_{max} = 0.41 s

  d = 18m, f_{min} = 70 GHz → b_{max} = 1.75 MHz (70-120 GHz) → C = 28,571

  d = 18m, f_{min} = 8 GHz → b_{max} = 200 kHz (8-48 GHz) → C = 200,000

  d = 18m, f_{min} = 1.2 GHz → b_{max} = 30 kHz (1.2-8 GHz) → C = 226,666
- If $T_{min} < T_{max}$, there is little effect on the correlator internally, but the output rate increases and this affects post-correlation processing. For now, assume $T_{min} = T_{max}$.
- Do we need a different correlator for each band?
  - Using b_{max} = 30 kHz for mid or high → C = 1,666,667.

Memory and I/O (*N*=355, 8-48 GHz band)

I/O rates		NGVLA	JVLA	ALMA	SKAmOrig			
integration, max	s	0.4	1	1	0.08			
ws	b	4	3	3	8	Ratios		
wi	b	8	4	4	8	*NGVLA/ JVLA*	*NGVLA/ ALMA*	*NGVLA/ SKAmOrig*
wo	b	32	64	32	64			
in	b/s	2.27E+14	3.07E+12	6.14E+12	4.06E+13	73.96	36.98	5.59
F to X	b/s	2.27E+14	2.05E+12	4.10E+12	2.03E+13	110.94	55.47	11.18
out	b/s	5.29E+12	2.1E+09	2.1E+09	2.706E+13	2461.43	2461.43	0.20
Memory								
antenna diameter	m	18	25	12	15			
longest baseline	m	300,000	34,000	7,500	150,000			
lowest freq	Hz	8.00E+09	4.30E+10	9.50E+11	1.42E+10			
bandwidth smearing	Hz	2.00E+05	1.32E+07	6.34E+08	5.92E+05	[1]		
rotation smearing	s	0.414	5.069	11.030	0.689	[1]		
min channels		199840	607	13	8444			
integration length	samples	61035	488281	976563	1526			
reordering memory	b	9.09E+13	2.05E+12	4.10E+12	1.63E+12	44.38	22.19	55.91
internal memory	b	3.47E+08	1.25E+08	5.00E+08	6.20E+06	2.77	0.69	55.91

[1] Bandwidth smearing and rotation smearing correspond to 90% visibility at primary beam -3dB point

Simple-Minded Scaling From SKA1-mid

- SKA1-mid correlator-beamformer as of PDR, Dec 2014:
  - 512 FPGAs (14 nm technology) — 341 adjusted
  - 12.3 M€ hardware — 8.2 M€
  - 8.9 M€ development labor — 5.9 M€
  - 140 kW — 93 kW
- Roughly 1/3 of the hardware supports the pulsar beamformers; adjusting for this gives the numbers on the right.
- These numbers do not include reserves, travel, infrastructure or pre-construction design work.
- Scaling hardware cost and power by 14x (ratio of computation rates), labor by 1x, and adding 10% reserve gives for ngVLA:
  - 4774 FPGAs (14 nm)
  - 133 M€ = 148 M$ today
  - 1.3 MW
- **This ignores the memory problem**, which makes this whole scaling exercise invalid. If the SKA numbers are right, then the scaled ngVLA numbers are *minimums*.
- Things will get somewhat better by 2022, but by how much?

Correlator Building Blocks

For N=256

3 versions:
a. 4096
b. 65,536
c. 262,144 ch

1 version:
3.3 Tb/s
in and out
512 x 950

1 version:
M_{CT} = 3.1 Gb
M = 63 Mb
M_{LTA} = 1.3 Gb
n = 64
f = 300 MHz

1.2-8 GHz	1 segment	F ver (c), J = 512	K = 725
8-48 GHz	5 segments	F ver (b), J = 256	K = 950
70-120 GHz	6 segments	F ver (a), J = 128	K = 950

Totals: 2,560 F units 1 CT 5,700 X units

Advantages of Segmented Bands

- The filter banks and correlator have a natural bandwidth of one segment, with larger bandwidths built up by duplication.
- If correlating 50 GHz is too expensive, it's easy to build only as many segments as we can afford.
- From the digitizers onward, all bands use identical hardware (well, mostly, see slide 9).
- If the front end designs change (e.g., different band edges and/or number of bands), it has little effect on the design of the back ends.
- A single segment tunable over the front end bandwidth is sufficient for many observations.
- Two independently tunable segments gives a powerful and flexible instrument. Is there science that would benefit significantly from more than this?

FPGAs, ASICs, GPUs in 2016

X part only, $N = 256$, $B = 50$ GHz, 2016 technology:

	K	P
GPUs:		
Extrapolation from SKA1-low	4,100*	1681 kW
FPGAs:		
Xilinx Ultrascale 14 nm	5,000	930 kW
ASICs (4 units per PCB):		
JPL chip, 32 nm, current version	12,800	40 kW
JPL chip, 32 nm, w/faster I/O	4,600	35 kW

* Per SKA1-low design, 2 GPUs per unit. NVIDIA Pascal GPUs.

But what about in 2022?

"Prediction is hazardous, especially about the future." Old Danish proverb, sometimes attributed to Niels Bohr, Mark Twain, or Yogi Berra.

- Moore's law ain't what it used to be. (Source: *ITRS*, 2013)
  - Transistor count doubling time was 1.5 years in the 1980s, slowed to 2 years in the 1990s, and has been 3 years since 2007. It is predicted to remain 3 years through 2019 then slow to 4.5 years.
  - Maximum die area of 130 mm^2 will remain constant through 2028.
  - Clock frequency is growing only 4% per year.
  - Operating voltage is decreasing only 2%/year (0.74V in 2021).

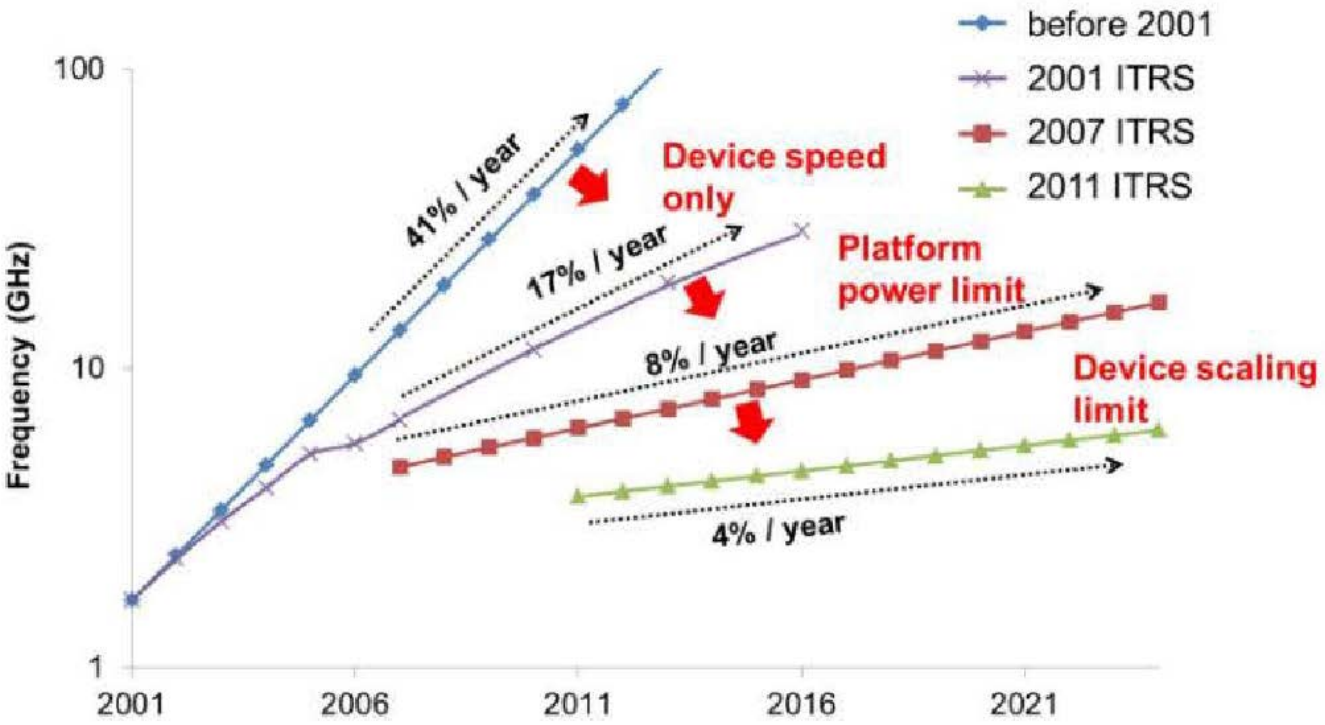

Backup Slides

Backup slides follow

Some Strawman Per-Unit Numbers

For $n = 64$ and $f = 300$ MHz:

			(SKA1-mid original)
X units	5333		525
M_{CT}	15.4 Gb	(3.1 Gb)	3.1 Gb
M	313 Mb	(63 Mb)	6.2 Mb
M_{LTA}		(206 Mb)	(4.1 Gb)
r_i	38.4 Gb/s		38.7 Gb/s
r_1			
r_2	9.8 Tb/s		9.8 Tb/s
r_3		(2.6 Gb/s)	
r_o	515 Mb/s		51.5 Gb/s

GPU (2016) calculation

Ref: SKA1-low PDR design by Curtin U. (SKA-TEL-CSP-000054, 10/28/2014).

Two NVIDIA Pascal GPUs and one host PC per unit (sec 5.3.2.2.2, p 27)

- Max computing 10 GFLOPS/GPU, 20 TF/unit.
- Used for SKA1-low X part: 9.88 TF/unit
- Same units are shared with F part.
- Assume 14 TF/unit available for an X-alone design.
- N=1024 B=1250 Hz (1 channel of SKA1-low) needs 19.3 GF.
  - Since this is 2.621 CMAC/s, the design uses 7.3636 F/CMAC
  - So 14 TF/unit => 1.901e12 CMAC/s/unit
- N=256 B=50 GHz is 6.55e15 CMAC/s => 3445 units needed

Input bandwidth ~50 Gb/s per unit (sec 6.6.1.2, p 37)

- N=256 B=50 GHz needs 2.05e14 b/s => 4100 units

Output bandwidth ~150 Gb/s per unit (sec 6.2.1.5, p 37)

- N=256 T=0.4s C=256K needs 2.75e12 b/s => 19 units

Power: 410 W/unit (sec 6.1.2, p 32). 4100 units => 1.681 MW.

X Unit: Internal Architecture

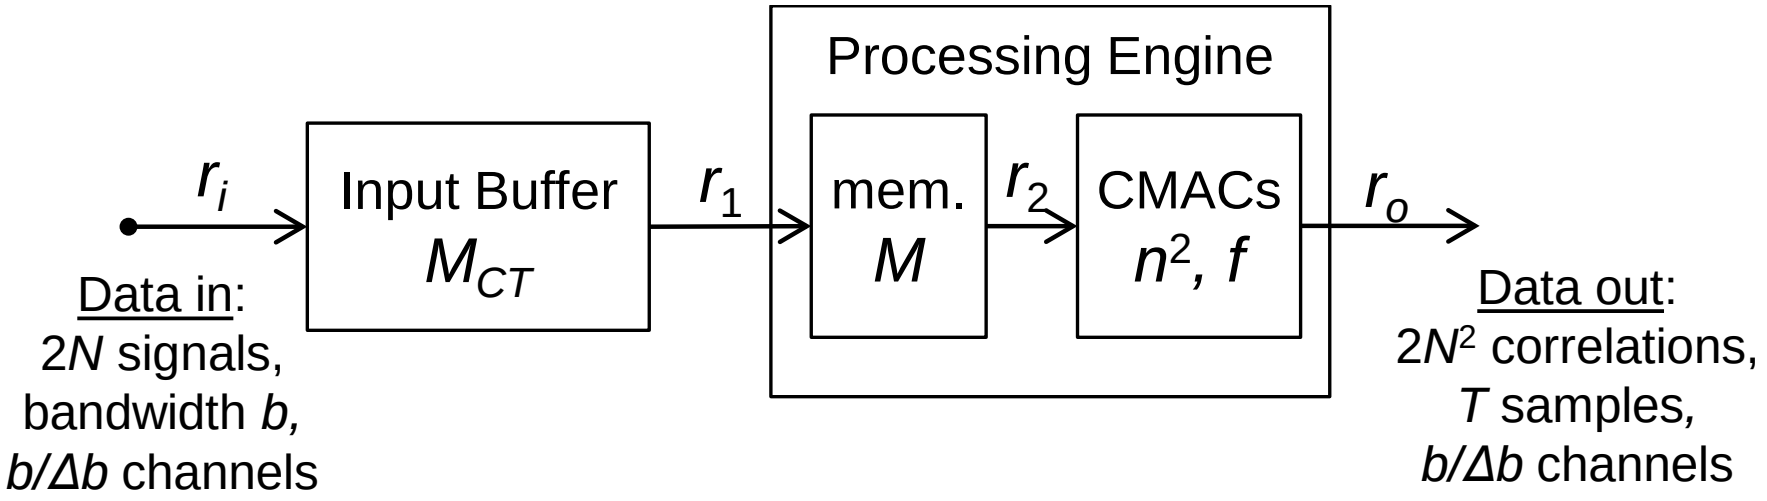

- The input buffer re-orders the data; logically part of the corner turner.
- The processing engine's memory allows re-use of the same data for computing more than n^2 correlations.

X Unit

n = 64, f = 300 MHz →
K = 5699 units, each 46 ch = 9.2 MHz, τ = 0.4 s (T = 80,000)

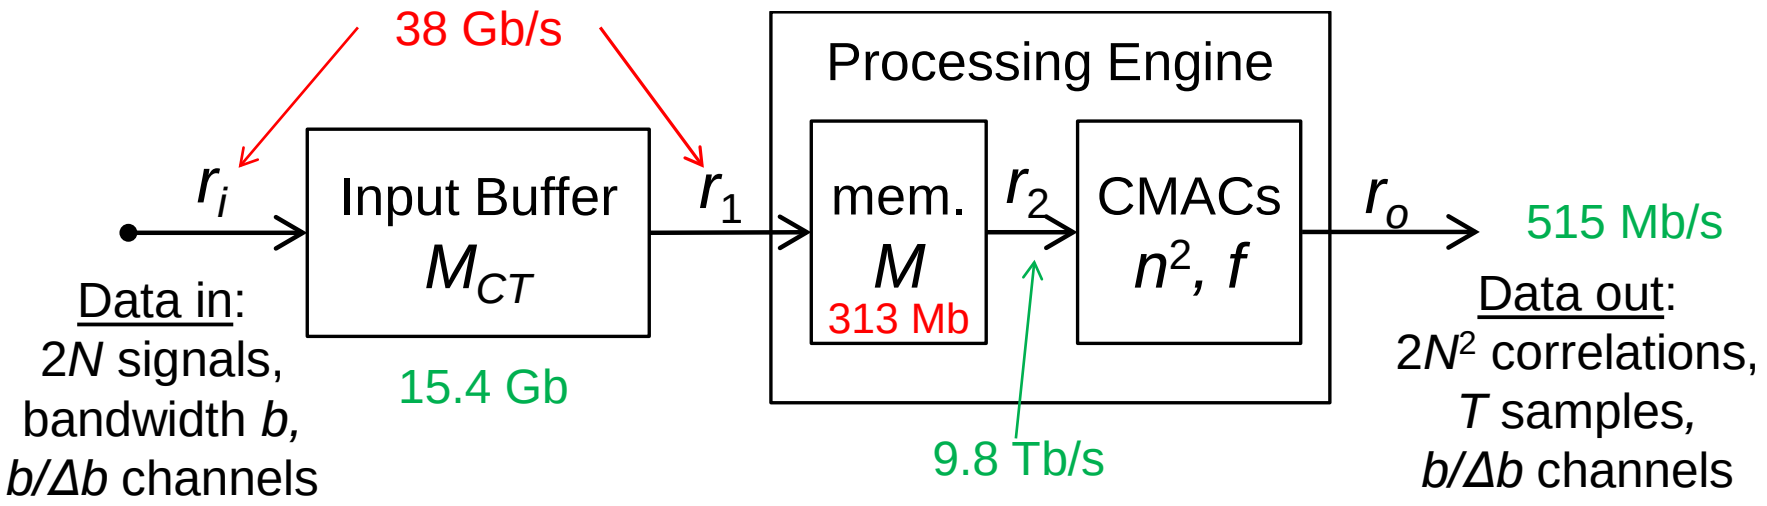

X Unit with LTA

n = 64, f = 300 MHz →
K = 5699 units, each 46 ch = 9.2 MHz, τ = 80 ms (T = 16,000)

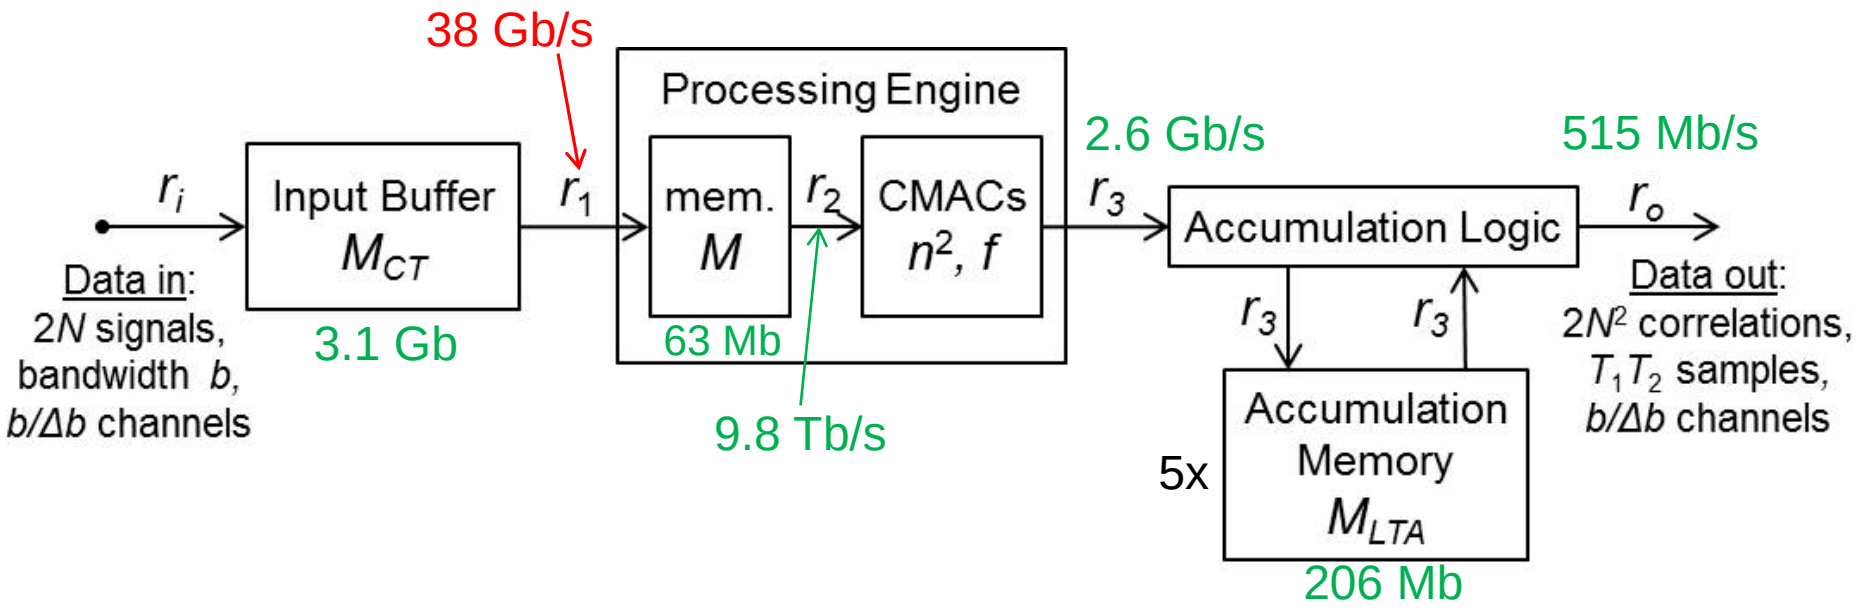

16.1 Gb/s 3.2 Gb/s

3.1 Gb 8.5 Mb 1.3 Gb

Low band (1.2-8 GHz):
K = 725 units, each 309 ch = 9.2 MHz, τ = 0.4 s (T = 10,376)

Corner Turners (1 of 2)

Basic 3D corner turner

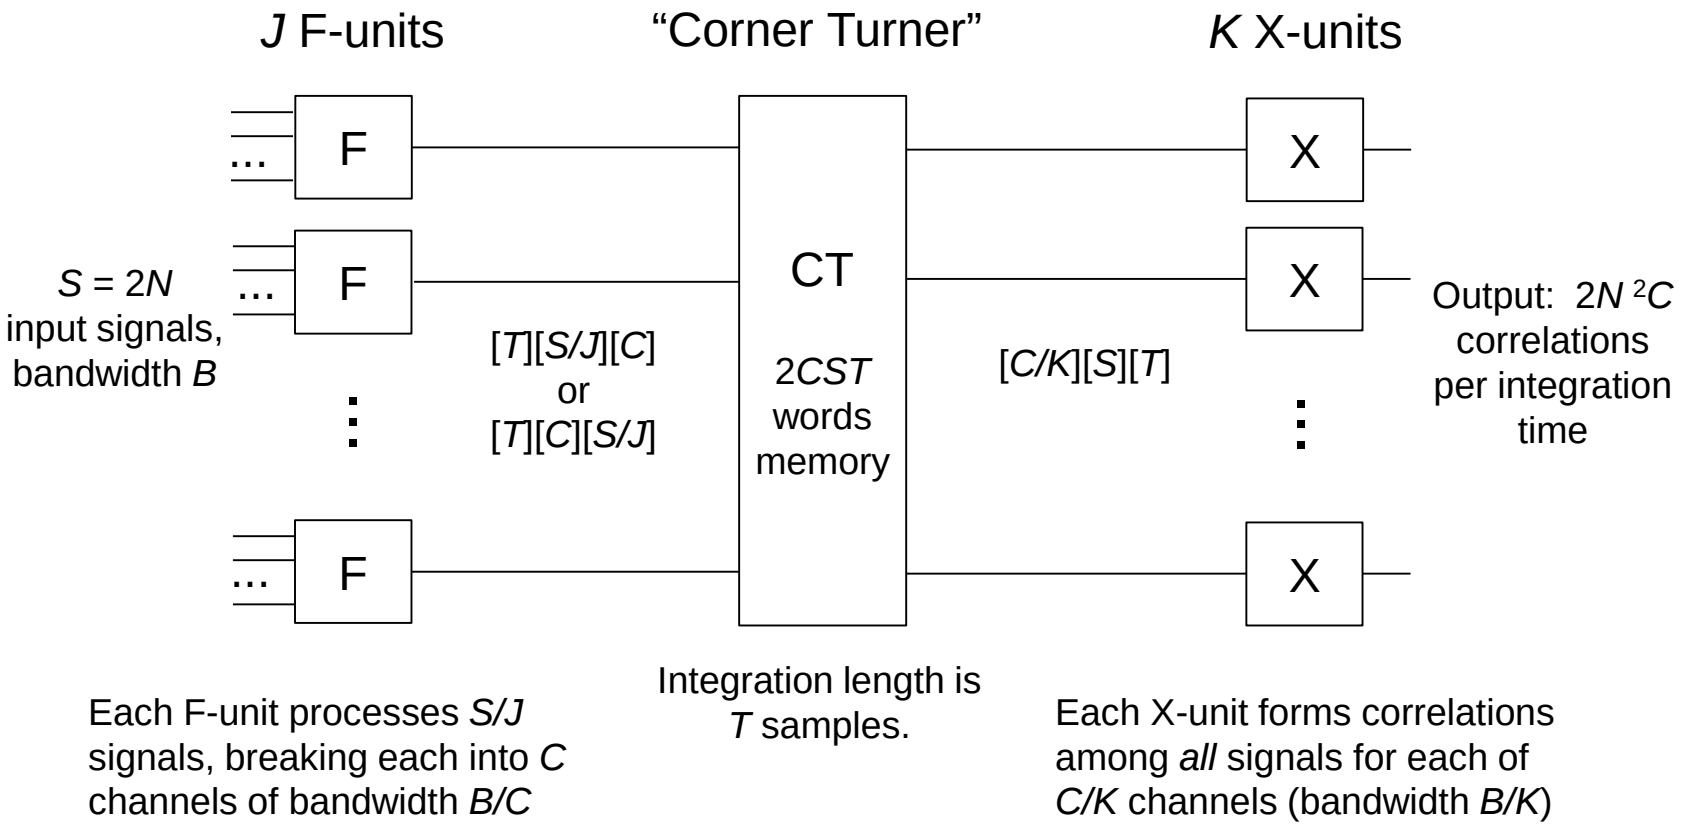

Corner Turners (2 of 2)

3D corner turner implemented as interconnect and multiple transposers

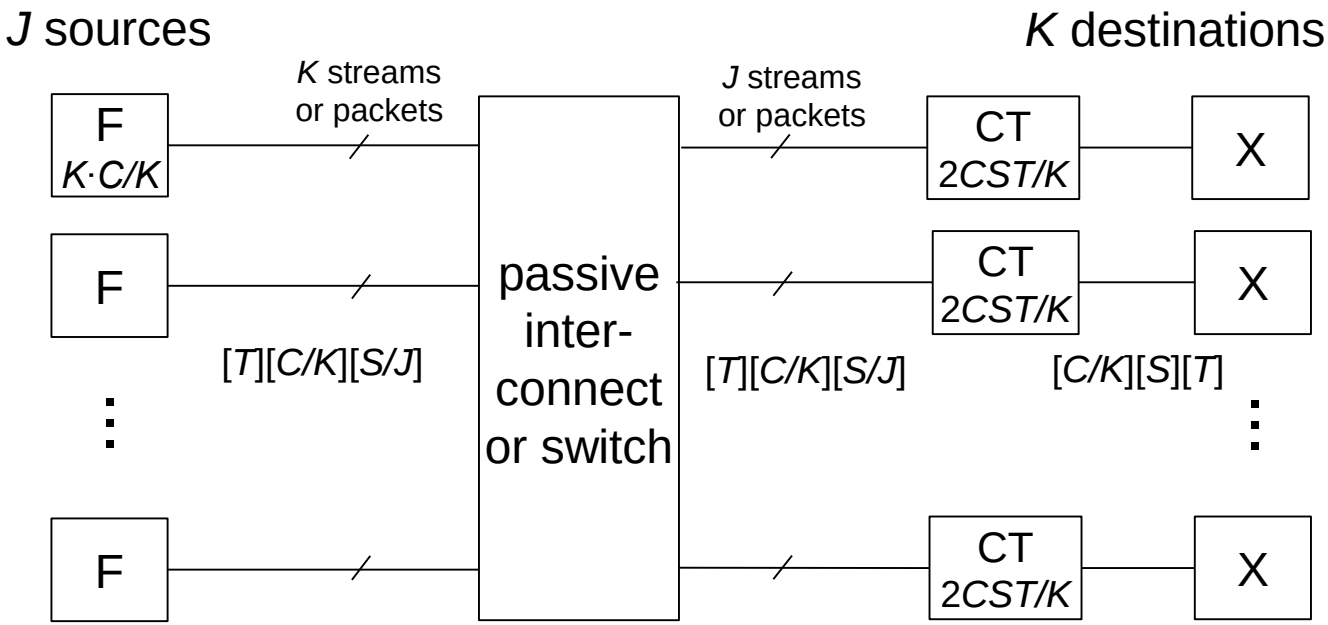

Time series of matrices of size *S* x *C*, transmitted on *KJ* streams or packets in row-first order. Each stream carries 1/*K* of the columns and 1/*J* of the rows.

Memory: $JC + 2CST$ samples

Timescale For New Technology Node Introduction

Source: ITRS update 2012, summary, Fig 1a

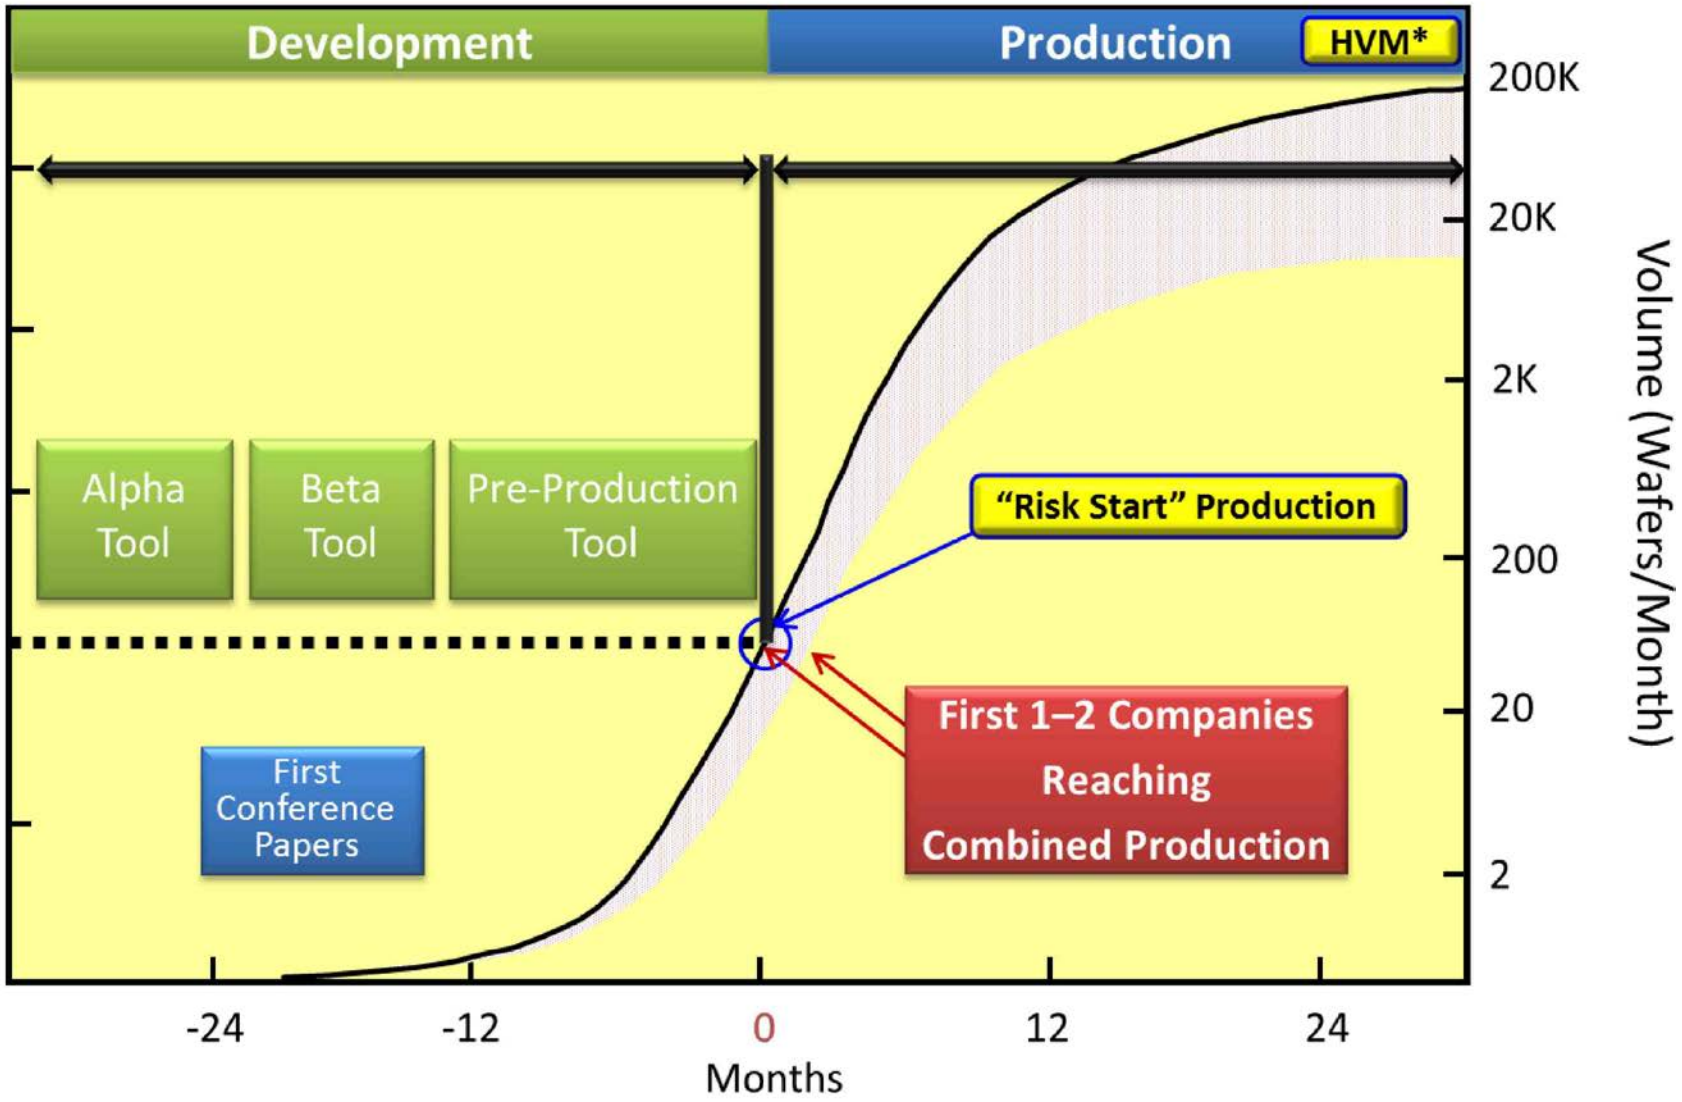

Digitizers and Band Segmenting

- It is currently difficult to digitize more than 5 GHz of bandwidth as a single signal because ADCs with sampling rates above 5 GHz are not readily available. [A 5 GHz channel can be digitized at baseband using 2 ADCs on quadrature (I/Q) versions of the signal.]
- Much faster ADCs are technically feasible, but existing ones are mostly embedded in products unrelated to our application (oscilloscopes, high speed digital receivers).
- Only moderate growth in commercially-available ADC bandwidth is expected by 2022, perhaps 2x (10 GHz).
- If we really want to digitize 50 GHz all at once, we should plan to develop custom digitizer chips. But is this a good idea?
- A better approach is to break the wide bands into segments using analog downconverters and filters. Consider having all ngVLA bands use segments of the same size, say 8 GHz:
  - 1.2 – 8 GHz is covered in one segment
  - 8 – 53 GHz is covered in 6 segments
  - 70 – 120 GHz is covered in 6 (or 7) segments

Moore's Law for Cost and Power

- Each new generation requires a new "fab"
  - Cost of a fab was ~300M$ in 1970's
  - TSMC announced starting on a new fab in 2010 at 9.3B$.
  - Samsung announced staring on a new fab in 2014 at 14.7B$.
  - List price of the largest Xilinx 7-series (28 nm) FPGA is $5,451
  - List price of the largest Xilinx Ultrascale (20 nm) FPGA is $15,188
  - The 20nm chip has only about 1.5x more RAM and DSP blocks.
- Power
  - Dynamic energy per operation is decreasing slowly.
  - Static power ("leakage") is increasing.
- Conclusion
  - Dramatic improvements in construction cost and power consumption cannot be expected from technology advancement alone.
  - From 2015 to 2022, a reasonable estimate is 2x improvement in each measure. Is is unrealistic to expect 3x or better.
  - Strategy: Design now in today's technology, then extrapolate.

Secondary Features

- Certain features beyond the basic specifications are frequently desired.
  - Pulsar gating (easy); pulsar binning (not so easy)
  - Zoom frequency resolution
  - Subarray support
  - Phased array output (VLBI and time-domain astronomy)
  - Baseline-dependent integrating time
  - “Flexibility” (grossly overrated)

Proposed Arrangement (each polarization)