

How To Implement SKA Digital Signal Processing So That It Uses Very Little Power

Larry R. D'Addario

Jet Propulsion Laboratory, California Institute of Technology

Workshop on

Power Challenges of Mega-Science: the example of SKA

Moura, Portugal

2012 June 20

©2012 California Institute of Technology. All rights reserved.

- The SKA power budget used for the recent site selection activities was [1]
 - On site: 67.4 MW
 - Off-site computing center: 40 MW
 - Annual cost estimate: 124 M€ in Australia or 50 M€ in South Africa [2]
- The on-site power includes about 0.1 MW for general infrastructure and 9 MW for mechanical things:
 - 1 kW x 3000 for dish pointing motors
 - 2 kW x 3000 for dish receiver cryocoolers
- The remaining on-site power, 58.3 MW (86.4%) is for signal-processing electronics. This value is far too high, and can be made at least an order of magnitude smaller by appropriate design choices with no sacrifice in performance.

Sources:

[1] Georginia Harris, "Power for the SKA." Presentation at PrepSKA WP2 meeting, October 2011

[2] James Moran *et al.*, "Report and Recommendation of the SKA Site Advisory Committee (SSAC)," 16 February 2012, p 138.

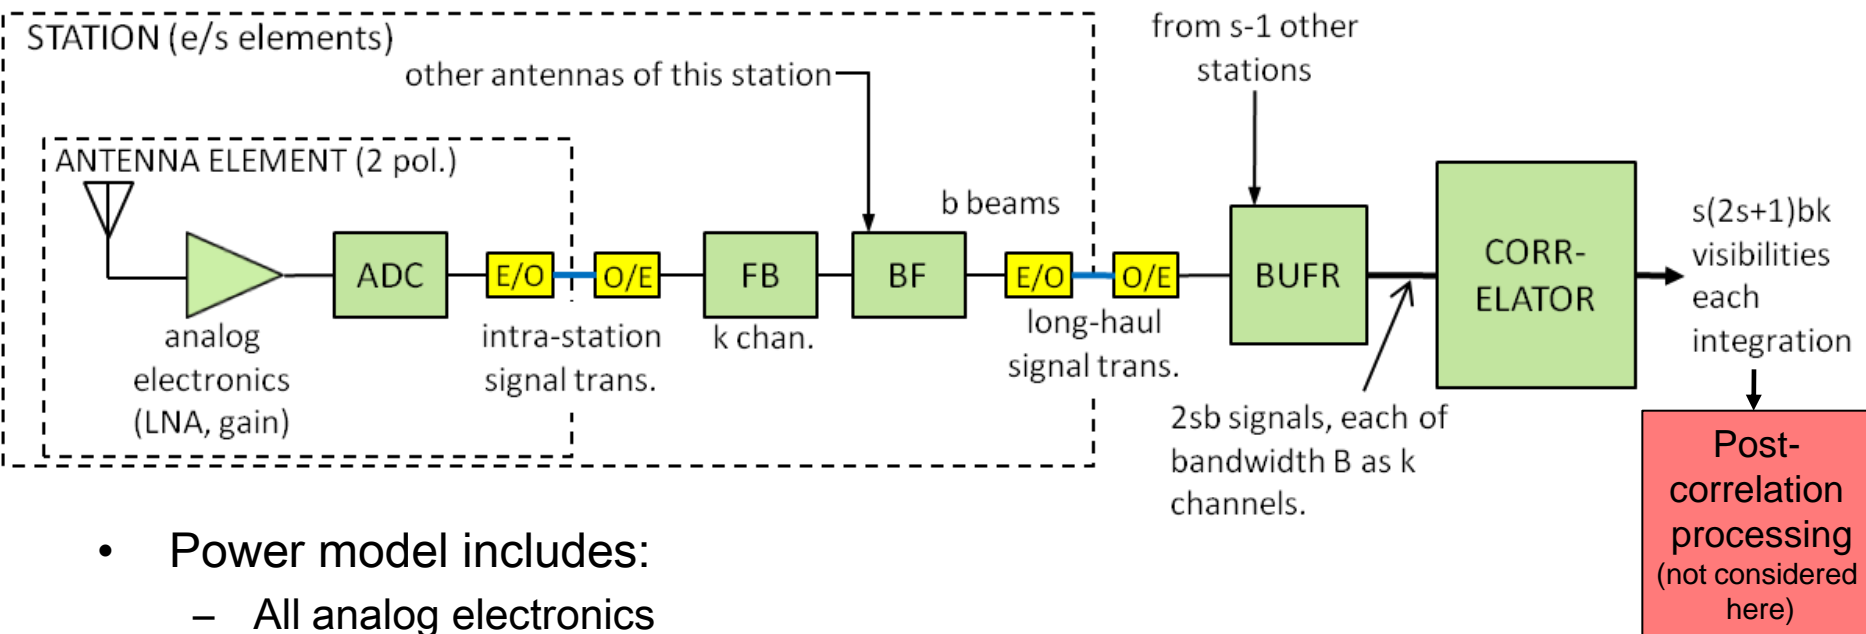
- No worries: Moore's Law will save us – *No!*
 - Some of the CoDR estimates already assumed 6-12 years of Moore's Law improvement (but here we assume only 2012 technology).
 - Moore's Law applies to transistors per IC, not directly to power.
 - Transistors/IC doubling time has slowed from 1.5 years in 1990s to 2 years in 2000s to ~3 years today. It is predicted* to remain 3 years through 2026.
 - Growth in internal clock speed has slowed dramatically; 4% annually is predicted* through 2026 (doubling time of 18 years).
 - Over the next 12 years, power per gate is expected* to decrease by a factor of only 4.5, with the maximum practical power/chip remaining unchanged at about 100W. So functionality/power will improve by only 4.5x, not by 16x as might be expected from transistor count. That is, *we can't actually use all those transistors at full speed because the chip gets too hot!*
 - Conclusion: Moore's Law is not enough, and it's less than some people think.

* Source:


International Technology Roadmap For Semiconductors, 2011 Edition.

- Optimize the architectures of
 - Each major component (filter banks, beam formers, correlators)
 - End-to-end system (antenna configuration, order of signal processing operations)
- Use purpose-built hardware
 - Energy/operation decreases dramatically as hardware is made more specialized:
General purpose (CPU/GPU) → Programmable (FPGA) → Dedicated (ASIC)
typical power reduction: 30x 10x
 - Three ASICs needed: Filter Bank, Beam Former, Correlator.
- Quantize as coarsely as possible
 - $2b+2b$ per sample except for intermediate values inside a processing IC
 - Re-quantize after channelizing in frequency and after beamforming
- Re-consider whether the AA-mid component can be cost-effective, considering that it dominates the power use (and construction cost).
 - We may be forced to eliminate it, or make it $\sim 10x$ smaller than now envisioned.

Model applies to all 4 SKA components: AA-low, AA-mid, Dishes/SPF, and Dishes/PAF



- Power model includes:
 - All analog electronics
 - Digitizers (ADCs)
 - All I/O for moving signal samples from one IC to another
 - Short-haul signal transmission within a station (<300m)
 - Long-haul signal transmission from station to center
 - Realistic models of IC power consumption
 - Filter banks and ADCs from published results
 - Inter-chip communication from published results
 - Correlator based on high-level design and analysis of new ASIC



FILTER BANK ASIC:
Polyphase FB with length-8 prefilter.
Many FBs per chip, depending on required bandwidth.
Same ASIC usable for all arrays of SKA.

BEAMFORMER ASIC:
Contains a large array of beamformer elements (delay + add), configurable as needed.

If there are fewer beams than elements per station, then fewer FBs are needed if they are placed after the BFs. But this requires time-domain BFs, which are more complex than frequency-domain BFs.

CORRELATOR ASIC:
Internal architecture chosen for low power consumption (SKA Memo 133). There is wide agreement that at least this ASIC is needed for SKA Phase 2.

									Current (2012) Technology				
SKA Component		Stns	SPDO Budget		CoDRs				New Model		Optimized		
			Each	Total	Each	Total	SP Part	Notes	Each	Total	Stns	Each	Total
Stations	AA Low (SAA)	250	40.0	10,000	65.8	16,451	16,451	[1]	12.73	3,184	186	7.79	1,448
	AA Mid (DAA)	250	123.0	30,750	192.0	47,996	47,996	[2]	97.64	24,411	234	39.38	9,214
	Dishes with SPFs	3000	4.0	12,000	4.1	12,366	453	[3]	0.01	17		0.03	17
	Dish PAFs (increment)	2270	2.0	4,540	3.0	6,802	6,802	[3]	0.39	883		0.64	883
	Dish remote stns (incr.)	25		0	47.1	1,177	1,177	[9]					
Central	AA Low (SAA)			10,000		706	706	[4]		74			178
	AA Mid (DAA)					3,271	3,271	[5]		356			1,248
	Dish SPFs					250	250	[4]		28			28
	Dish PAFs					8,003	8,003	[6]		687			687
	Pulsar/transients							[7]					
	Infrastructure			100		100							
Off Site	Computing center			40,000		40,000		[8]					
	GRAND TOTAL			107,390		137,121	85,109			29,639			13,704
	WITHOUT AA MID			76,640		85,854	33,842			4,873			3,242

"Each" and "Total" columns are powers in kW. Further details given in backup material.

RED: >10% OF TOTAL

[1] A. Faulkner, "SKA AA Implementation." AA CoDR, April 2011, p 13.

[2] Scaled from [1] in proportion to bandwidth and number of beams.

[3] Dish array estimates not given in CoDRs; values here are a reasonable breakdown of the SPDO budget. "SP Part" excludes mechanical power for motors and cryogenics, estimated at 3 kW/station.

[4] Median of values presented at Signal Processing CoDR, April 2011. Actual values varied widely.

[5] AA Mid correlator not estimated at CoDRs. Value scaled from AA Low median by bandwidth and number of beams.

[6] Dish PAF correlator not estimated at CoDRs. Value scaled from Dish SPF correlator by assuming 32 beams.

[7] Pulsar/transient processor not estimated at CoDRs. Omitted from models.

[8] Off-site computing power not estimated in CoDR; budgeted amount used here.

[9] Beamformers for remote stations omitted from models, but receivers included; BFs expected to use very little power.

Mostly from Signal Processing Concept Design Review, April 2011


Rough attempt to compare power and cost of correlators designs in concept papers.

Designs did not all stick to the Memo 130 specs, and they varied considerably in their projections of future technology.

Component	Source	Technology	Raw	Burdened	Notes	
			W	W		
SKA1-low correlator (50,480)	Ford, Memo 139	GPU, 2013	960,000	1,411,765	1920+1920 Maxwell cards @250W	
	Ford, SP CoDR	GPU, 2017	280,000	411,765	560+560 Tesla GPU cards @ 250W	
	Szomoru, SP CoDR	FPGA, 2011	3,300,000	4,852,941	9414 Uniboards, only 160 beams	
	Kapp, SP CoDR	FPGA, 2015	2,976,000	4,376,471	ROACH4, 16/beam	
	Bunton, SP CoDR	FPGA, 2011	64,000	94,118	320 2-FPGA boxes @ (guess) 200W	
	Carlson, SP CoDR	ASIC, 2011, ?nm	10,800	15,882	3360 ASICs @ 3W.	
	D'Addario, SP CoDR	ASIC, 2007, 90nm	8,863	13,034	1842 ASICs, 75 stations	
	D'Addario, latest	ASIC, 2012, 32nm	5,430	7,985	7200 demo ASICs, 50 stations	608
SKA1-mid correlator (250,1)	Szomoru, SP CoDR	FPGA, 2011	134,608	197,953	384 Uniboards	
	Kapp, SP CoDR	FPGA, 2015	428,000	629,412	ROACH4 x 2298	
	Bunton, SP CoDR	FPGA, 2011	6,400	9,412	16 4-FPGA boards	
	Carlson, SP CoDR	ASIC, 2011, ?nm	1,512	2,224	504 ASICs, 7 boards	
	D'Addario, SP CoDR	ASIC, 2007, 90nm	704	1,035	121 ASICs, N=300	
	D'Addario, latest	ASIC, 2012, 32nm	271	399	256 demo ASICs, N=250	1,578
SKA2-low correlator (250,480)	Carlson, SP CoDR	ASIC, 2011, ?nm	732,800	916,000	N=256, b=1000, 300MHz	
	D'Addario, SP CoDR	ASIC, 2007, 90nm	337,920	496,941	N=300, extrap from SKA1-mid.	
	D'Addario, latest	ASIC, 2012, 32nm	50,264	73,917	N=50, X only, 55680 demo ASICs	12.4
SKA2-mid correlator (3000,1)	Carlson, SP CoDR	ASIC, 2011, ?nm	370,000	462,500	N=3072. PS eff incl. in raw.	
	D'Addario, SP CoDR	ASIC, 2007, 90nm	25,771	37,899	N=2025, X only	
	D'Addario, latest	ASIC, 2012, 32nm	20,322	29,885	N=3000, X only, 10752 demo ASICs	15.5

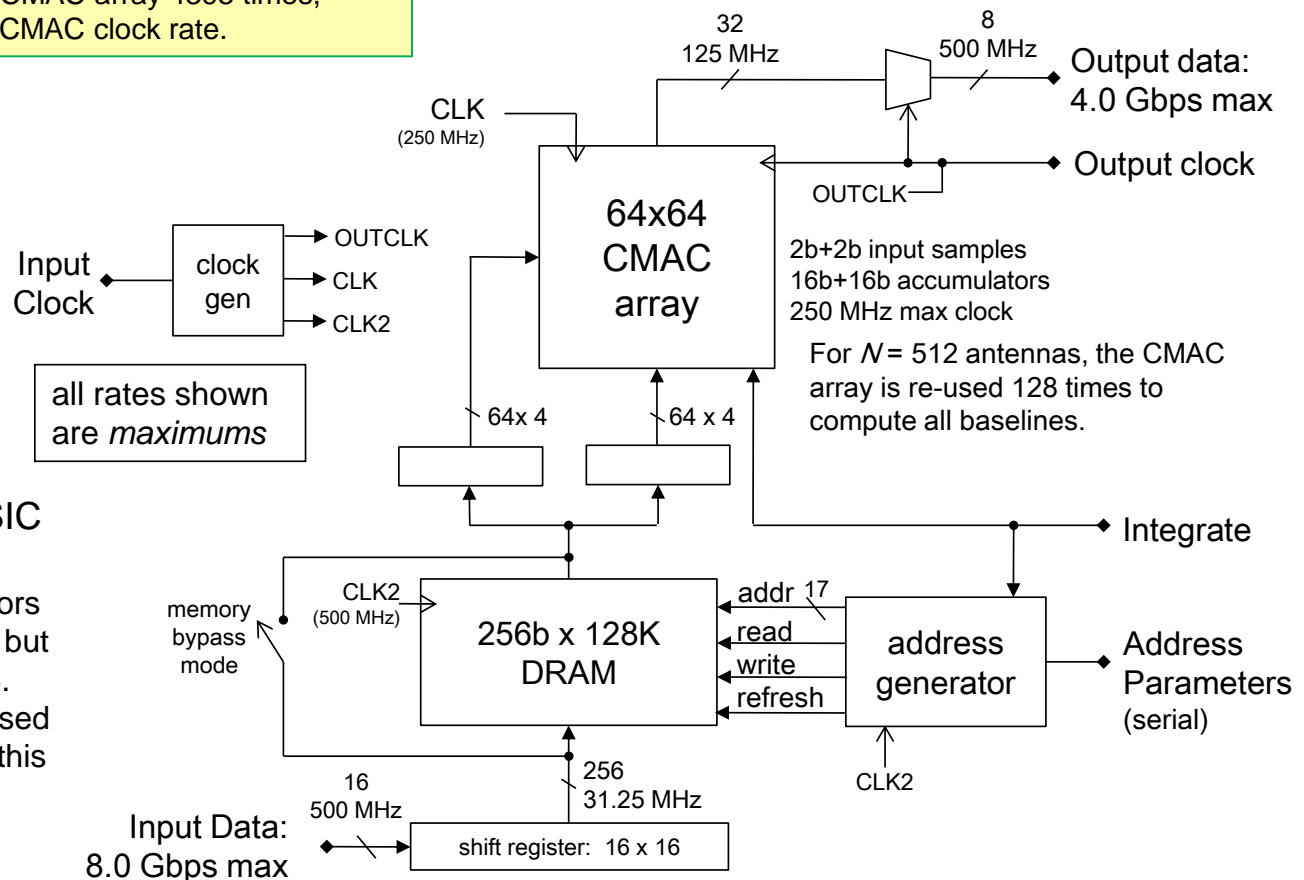
"Burdened" includes power supply efficiency at 85% and adds 25% for cooling relative to "Raw" estimates.

Concept



Example: $N=3000$ requires 18 million baselines. Can be done with $n^2 = 4096$ CMACs by re-using the CMAC array 4395 times, achieving a bandwidth of $1/4395$ of the CMAC clock rate.


This architecture uses less power than alternatives by a factor of 2 to 13. See SKA Memo 133.




Block diagram of correlator ASIC being designed at JPL

This device can be used in correlators for almost any number of antennas, but is most efficient for $N=32$ to 1024. Correlators based on this chip are used for all the power estimates given in this presentation.

SKA1 Low (s=50,b=480,e=560000)




SKA1 Mid (s=250,b=1)



Totals on this and subsequent "pie chart" slides do NOT include power supply efficiency (85%) nor cooling (+25%), but they are included in the tabulated estimates (pp 6, 7, 21).


SKA1 Low ($s=50, b=480, e=560000$)




SKA1 Mid ($s=250, b=1$) 1.3 kW




SKA2 Low ($s=250, b=480, e=2.8e6$)




SKA2 Dishes SPF ($s=3000, b=1$)




SKA2 AA-Mid ($s=250, b=1000, e=16.8e6$)



SKA2 PAF ($s=2270, b=32:128$)



SKA2 AA-Low ($s=250, b=480, e=2.8e6$)




SKA2 Dish-SPF ($s=3000, b=1$)




32.4 kW

Area of each pie chart is proportional to total power required.

SKA2 AA-Mid ($s=250, b=1000, e=16.8e6$)



SKA2 Dish-PAF ($s=2270, b=32$)




995 kW

- Make AA-low and AA-mid smaller while maintaining survey speed (next slide)
- Develop lower power analog circuitry (LNAs, gain blocks)
 - care needed to maintain low noise and high dynamic range
- Integrate ADCs with filter bank ASICs
- Obtain much lower power short-haul optical transmitters
- Use 40 Gb/s or 100 Gb/s per fiber for long-haul data transmission
 - 10 Gb/s per lane and COTS devices were assumed here
- Use larger correlator ASIC with on-chip multi-Gb/s I/Os
 - 4096 CMACs with parallel I/O was used here
 - At least a 5x larger chip is feasible, but it would not be useful until SKA2.
- Wait. We will get some help from Moore's Law (though not enough).
 - Models here are based on current (2012) technology.
 - Technology for SKA2 does not need to be frozen for at least 8 years.
 - Expect about 2.5x reduction in power for digital computation, less for I/O, less for analog.


- Current nominal designs for arrays of dipole-like elements have these survey speeds at 130 MHz:
 - SKA1 AA-low : $A^2\Omega = 8.90E10 \text{ m}^2\text{-sr}$ (at 130 MHz with 480 beams)
 - SKA2 AA-low : $2.22E12 \text{ m}^2\text{-sr}$ (at 130 MHz with 480 beams)
 - SKA2 AA-mid: $1.25E10 \text{ m}^2\text{-sr}$ (at 1 GHz with 1000 beams)
- Also:
 - SKA2 Dish-PAF: $1.83E09 \text{ m}^2\text{-sr}$ (at 1 GHz with 32 beams)
- If we maintain the same survey speeds but minimize the end-to-end signal processing power (including both stations and central correlator), we can substantially reduce the required power.

SKA2 AA-mid: Power for Signal Processing at Constant Survey Speed




SKA2 AA-Mid ($s=250, b=1000, e=16.8e6$)

16.8 MW



Optimized $s=234, b=5296, e=1.39e6$

7.1 MW



- LNAs
- ADCs
- BFs
- FBs
- Buffers
- Correlators
- Transmission, short
- Transmission, long
- I/O

- Analog circuitry assumed to use 0.13W per signal (same as AA CoDR).
- Digital circuitry scaled to 32 nm CMOS.
- *ADCs*: Based on published results*.
- *Correlator ASICs*: Power derived from simulations and analysis of demonstration chip now being designed at JPL. 64x64 CMAC array and input memory. Architecture is described in SKA Memo 133. Includes power for on-chip memory board-level FPGA for high-speed serial I/O.
- *Filter bank ASICs*: Based on published results*.
- *Beamformer ASICs*: Derived from CMAC simulation and scaling results (see SKA Memo 133).
- *I/O between chips*: Based on published results* for 10 Gb/s links.
- *Signal transmission*: COTS devices*
 - 10 Gbps multi-mode fiber within stations
 - 10 Gbps per lane single-mode fiber from stations to center
- * Further details and references are given in a backup slide.

- A criticism of ASICs is that they are difficult and costly to develop.
 - At today's process nodes (<40 nm gate lengths), it costs several M€ to bring a single ASIC to production-ready status (including fabrication and testing of prototypes). Complex devices can require more than one design/fab cycle, bring the cost to the 5 to 10 M€ range.
 - Many more people know how to program FPGAs and CPUs, making them more popular.
- However, delivering the budgeted power to the SKA requires:

	ANZ (€M)	RSA (€M)
Power-Capital Cost	331,478,500	131,968,000
Power-Operation and Maintenance Cost per Year	124,000,000	49,798,000

Source: James Moran *et al.*, "Report and Recommendation of the SKA Site Advisory Committee (SSAC)," 16 February 2012, p 138.

- Thus, if we spend 30 M€ on ASIC development (10 M€ for each of the 3 devices recommended here) it would correspond to
 - 9% to 23% of the capital cost of the budgeted power, or
 - 3 to 7 months of the operating cost of the budgeted power.
- The actual power requirement is at most 40% of budgeted amount (see slide 3).
- *We should spend much less on power and more on improving the electronics.*

- Proceed with design and prototyping of the three ASICs that are needed to implement SKA signal processing. (Each of the SKA's four telescope types needs all three.)
- Invest in development of much lower-power analog signal processing: LNAs and gain blocks.
- Keep up with developments in industry toward higher data rates per optical channel at similar power (40 Gbps and 100 Gbps channels).
- Study ways to minimize power as a function of layout for the AA telescopes, while keeping measures of scientific return constant.
- Other potential power savings not considered here:
 - Modern cryocoolers can use far less power than budgeted (6 MW)
 - Off-site computing power budget might be excessive (40 MW). As with everything else, cost-effectiveness, not just desirability, must be considered.

- Site selection work was based on an SKA power budget of 107.4 MW, with **67.4 MW** on-site.
- Estimates in CoDRs suggest that the current nominal design requires on-site power of **97.1 MW**, of which **85.0 MW** is for signal processing electronics. Some CoDR estimates assumed future technology.
- Careful designs, using purpose-built components rather than generic ones, require only **29.6 MW** for all signal-processing electronics using current (2012) technology, based on a detailed model.
- If AA-low and AA-mid are configured for minimum power at the same survey speed, total signal processing power becomes **13.7 MW**.
- In all cases, SP power is dominated by AA-mid. Without that component, total SP power becomes **3.2 MW**.
- By 2020, when SKA2 designs might need to be frozen, technology advances are predicted to produce 60% reduction in power for digital signal processing. Assuming no reduction in analog electronics power, total SP power becomes **1.65 MW** without AA-Mid, or **7.36 MW** including AA-Mid.
- The cost of development of power-efficient electronics is far more than offset by the saving in the *capital* costs of power production. The huge saving in *operating* costs is additional.

End of Presentation

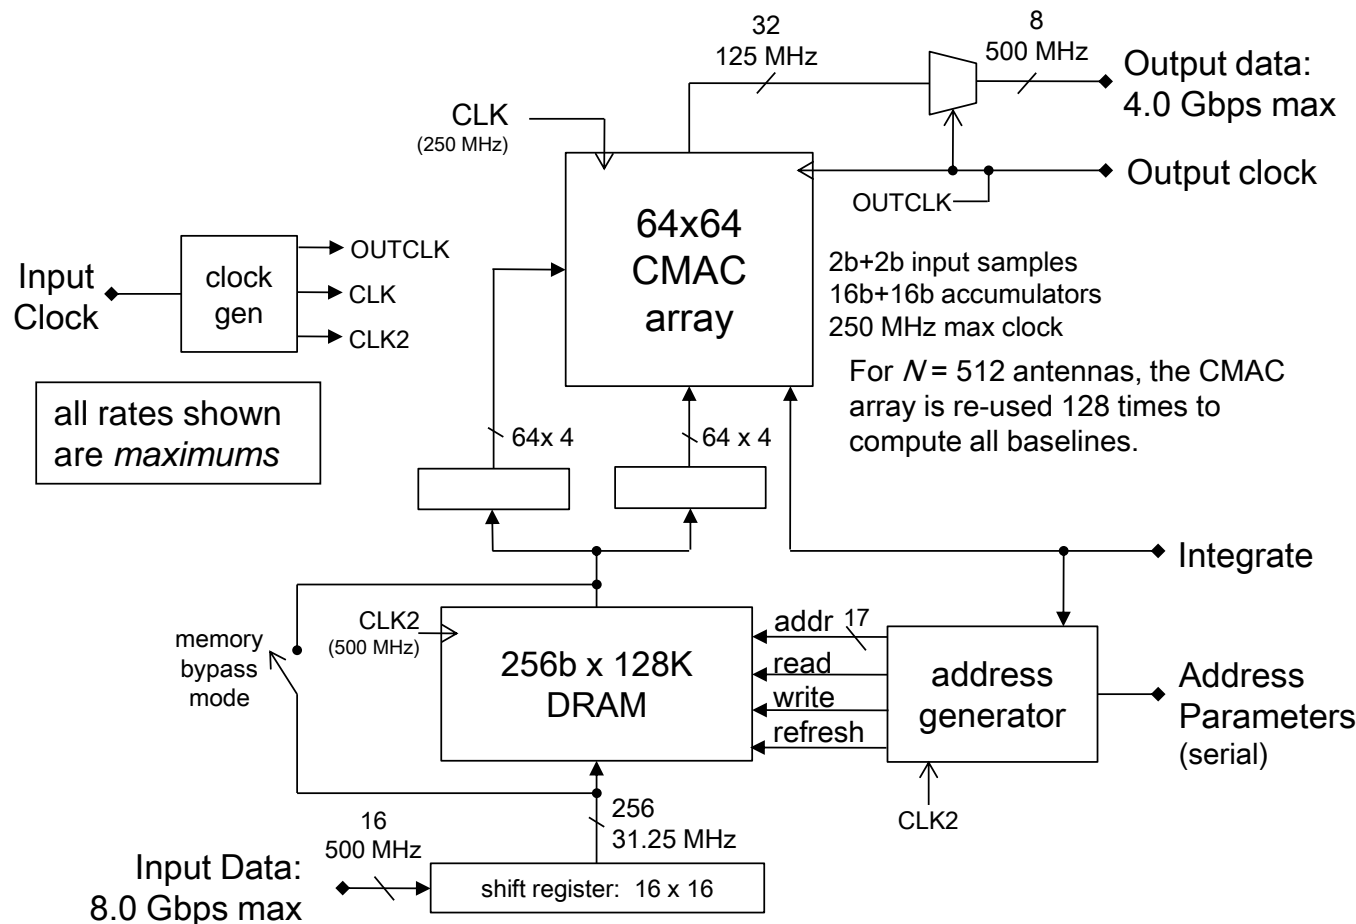
Backup Slides Follow

Basis of Power Models

Sym.	Description	Units	Value	Ref
ce1	power per LNA (indep of bandwidth)	W	0.13	[1]
ce2	energy per sample digitized (ADC)	J	9.450E-11	[2][3]
cf	energy per Filter Bank operation (FFT radix 2 butterfly)	J	2.226E-11	[4][5]
cbf	energy per frequency-domain beamformer operation, W/Hz	J	4.267E-12	[5]
cbt	energy per time-domain beamformer operation, W/Hz	J	2.844E-11	[5]
ct1	energy per transmission (one sample, short haul: element to station)	J	2.391E-10	[6]
ct2	energy per transmission (one sample, long haul: station to center)	J	6.015E-10	[8]
ci	energy per I/O (one sample, chip to chip)	J	1.510E-12	[9]
cc	energy per CMAC	J	2.700E-12	[10]
cm	energy per Read+Write to RAM (one sample)	J	4.800E-11	[11]

All values scaled to 32nm CMOS when possible.


- [1] A. Faulkner, "SKA AA Implementation." AA CoDR, April 2011, p 13.
- [2] M. Choi et al., "A 6-bit 5-GSample/s Nyquist A/D converter in 65nm CMOS," 2008 IEEE Symp. on VLSI Circuits.
- [3] Derived from [2] as described in [5], but scaled to 32nm and doubled for complex/Nyquist.
- [4] B. Richards et al., , "A 1.5GS/s 4096-Point Digital Spectrum Analyzer for Space-Borne Applications." IEEE Custom Integrated Circuits Conference, September, 2009.
- [5] L. D'Addario, "Low-Power Correlator Architecture For the Mid-Frequency SKA," SKA Memo 133, 2011 March 21.
- [6] Avago AFBR-703SDZ SFP+ transceiver, 10GbE 850nm 300m 594 mw -> 59.4 pJ/b SERDES added.



STATUS:


- Functional design is complete. Preliminary data sheet available on request.
- Foundry, process, and IP selected.
- Detailed logical design is underway.
- Final physical design, prototype fabrication, and testing require additional funding.

SKA2 AA-low: Power for Signal Processing at Constant Survey Speed




SKA2 Low ($s=250, b=480, e=2.8e6$)

2,215 kW



SKA2 Low Optimized ($s=186, b=2621, e=689700$)

1,106 kW



Comparison of Power Estimates

- SPDO Budget
- Estimates given at Concept Design Reviews
- New estimates based on model described here.

	Qty		No.		SPDO Budget		CoDRs				New Estimate	
	per stn	stns	Each	Total	Each	Total	Ref	SP items	Each	Total	Each	Total
AA Low Station (n=11200, b=480, B=0.5GHz)			kW		W	kW			W	kW		
LNAs, etc.	22,400	250			0.13	728.0	[1]	728.0	0.13	728.0		
ADCs	22,400	250			0.10	560.0	[1]	560.0	0.04	238.1		
Intra-station transmission	22,400	250			0.32	1,764.0	[1]	1,764.0	0.11	602.6		
Filter Bank	22,400	250			1.17	6,552.0	[1]	6,552.0	0.08	433.6		
Beam former	960	250			6.07	1,456.0	[1]	1,456.0	0.08	97.5		
Station-to-center transmission	960	250			0.53	126.4	[2]	126.4	0.27	65.0		
Power supply loss and cooling (47%)				incl		5,264.2		5,264.2		1,018.7		
TOTAL			40	10,000		16,450.6		16,451		3,183.5		
AA Mid Station (n=27520, b=1000, B=1GHz)												
LNAs, etc. plus analog BFs (clusters of 4)	220,160	250			0.125	6,925.0	[1]	6,925.0	0.16	8,541.1		
ADCs	55,020	250			0.19	2,160.6	[1]	2,160.6	0.09	1,301.8		
Intra-station transmission	55,020	250			0.32	2,823.3	[2]	2,823.3	0.24	3,294.0		
Filter Bank	55,020	250			0.30	12,695.8	[1.5]	12,695.8	0.18	2,453.2		
Beam formers	55,020	250			0.26	7,374.2	[1.5]	7,374.2	0.05	708.3		
Station-to-center transmission	2,000	250			1.32	658.4	[2]	658.4	0.60	300.8		
Power supply loss and cooling (47%)						15,358.8		15,358.8		7,811.4		
TOTAL			123	30,750		47,996.1		47,996.1		24,411		
Dishes with SPFs												
Motors	1	3000			900.0	2,700.0	[3]					
Cryocoolers	1	3000			1,800.0	5,400.0	[3]					
SPF LNAs and IFs	2	3000			0.13	0.8	[3]	0.8	0.13	0.8		
SPF downconverters (incl LO)	2	3000			50.00	300.0	[3]	300.0	1.00	6.0		
SPF ADCs	2	3000			0.25	1.5	[3]	1.5	0.25	0.3		
SPF filter banks	2	3000								1.1		
SPF signal transmission	2	2400			1.32	6.3	[2]	6.0	0.27	3.6		
Power supply loss and cooling (47%)						3,957.0		145.1		5.5		
TOTAL				12,000		12,365.6		453.3		17.3		
PAFs (increment)												
PAF LNAs and IFs (128 PAF el, 32 beams)	256	2270			0.13	75.5	[3]	75.5	0.13	75.5		
PAF downconverters (incl LO)	256	2270			1.00	581.1	[3]	581.1	0.25	145.3		
PAF ADCs	256	2270			0.25	145.3	[3]	145.3	0.25	54.9		
PAF filter banks	256	2270								0.14	84.1	
PAF beamformers	64	2270			25.00	3,632.0	[3]	3,632.0	0.10	14.0		
PAF signal transmission, intra-station	64	2270							0.96	139.0		
PAF signal transmission, to center	64	2270			1.32	191.3	[2]	191.3	0.60	87.4		
Power supply loss and cooling (47%)						2,176.6		2,176.6		282.4		
TOTAL				4,540		6,801.8		6,801.8		882.6		
Remote stations (increment)												
Remote station beamformers	64	25			500.0	800.0	[3]	800.0				
Remote station signal transmission	2	25			1.3	0.1	[2]	0.1				
Power supply loss and cooling (47%)						376.5		376.5				
TOTAL		25				1,176.6		1,176.6				
Central processing												
Dish SPF correlator (N=3000, b=1, 1 GHz)						250.1	[4]	250.1		19.3		
Dish PAF correlator (2270, 30, 1 GHz)						8,003.2	[5]	8,003.2		467.1		
AA-low correlator (250, 480, 480 MHz)						706.5	[4]	706.5		50.3		
AA-mid correlator (250, 1000, 1 GHz)						3,270.7	[7]	3,270.7		242.3		
Pulsar/transient processor				100			[6]					
Power supply loss and cooling included above												
TOTAL				10,000		12,230.5		12,230.5		778.9		
Infrastructure				100		100.0						
Off-site computing		1	40,000	40,000		40,000.0	[6]					
GRAND TOTAL				107,390		137,121		85,109		29,273		

[1] A. Faulkner, "SKA AA Implementation." AA CoDR, April 2011, p 13 for AA-low and p 24 for AA-mid.

[1.5] For these items, numbers in [1] p 24 are unreasonable. Values here scaled from p 13 by bandwidth, beams, elements.

[2] C. Shenton, "Signal Transport And Networks." Presentation at STaN CoDR, June 2011, pp 8,23.

Assumed 400 MSa/s and 8b for AA-low, 1000 MSa/s and 8b for AA-mid and dishes.

[3] My estimated breakdown of budget

[4] Median of values presented at Signal Processing CoDR; see separate table.