



# Large Array Signal Processing



Roger Cappallo  
MIT Haystack Observatory  
2001.7.12





# Roadmap

- Moore's law vis a vis correlators
- XF - past, present, and future
- FX possibilities
- Interconnection
- Transient data buffers
- Soft(ware) conclusions

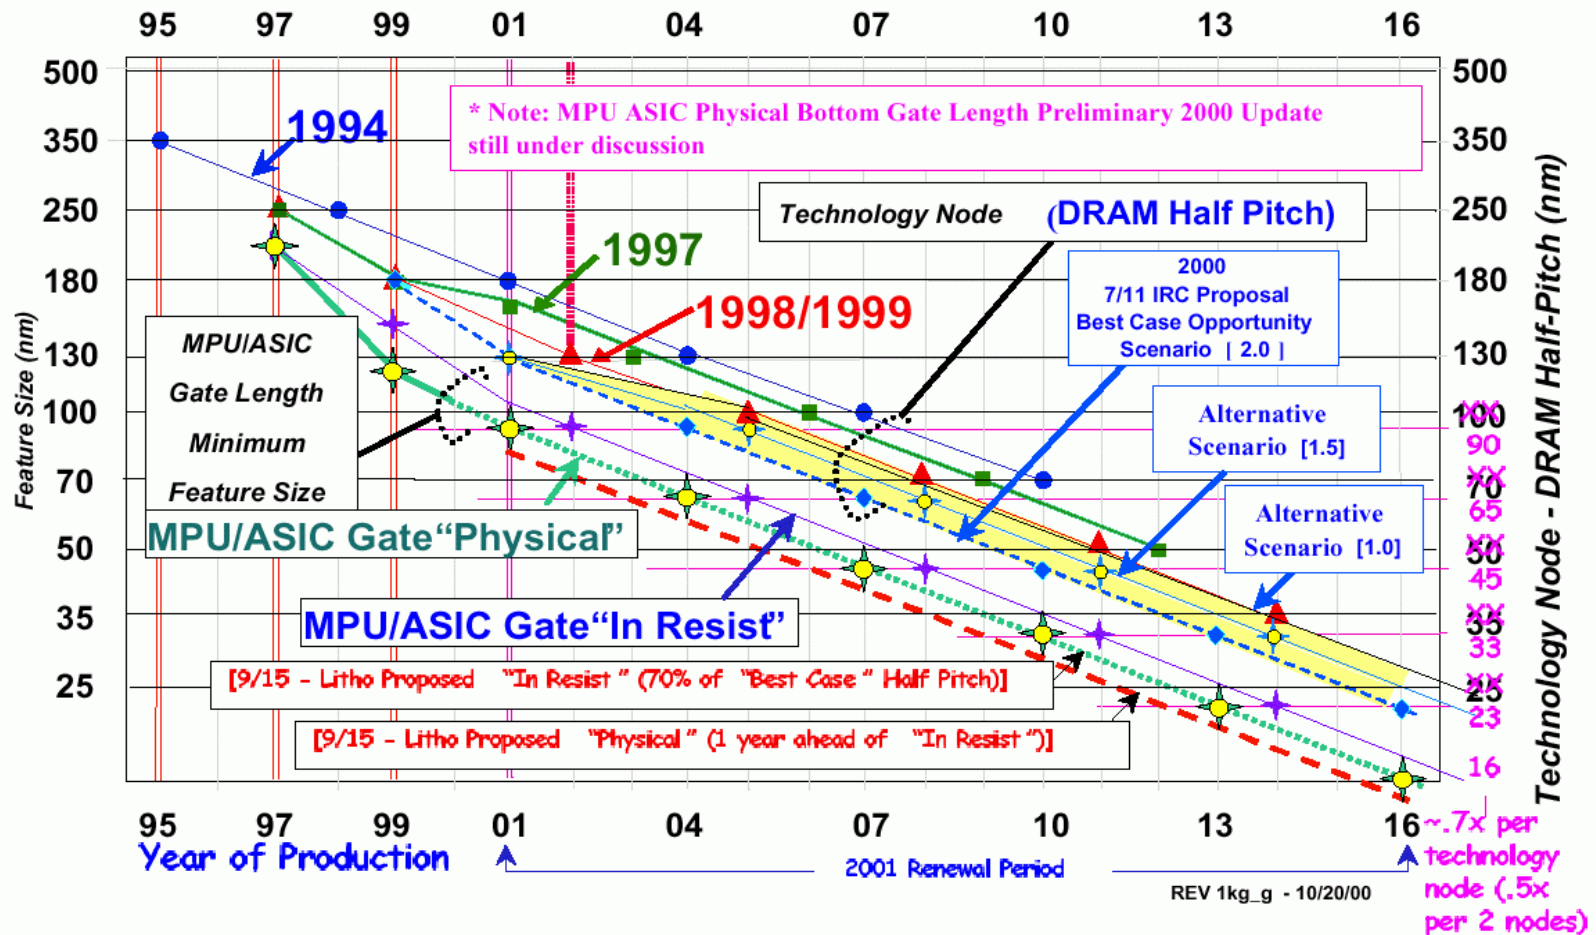


# International Tech. Roadmap

- 2000 update
- detailed projections (guidelines) for semiconductor industry over a 15 year period
- mixture of specific enabling technologies & “Moore’s Law” extrapolations

# ITRS Roadmap Acceleration

ITRS Roadmap Acceleration Continues...  
(Including MPU/ASIC "Physical Gate Length" Proposal)



# International Technology Roadmap

Year	Line Widths (nm)	ASIC trans/chip (Millions)	ASIC local freq (MHz)	ASIC Cross-chip freq (MHz)	DRAM size (Gb)
2011	50	4570	11050	1595	17

**Note: Mark 4 correlator chip uses ~700K transistors to implement 512 lag registers and 8 delay & phase trackers, at a 64 MHz clock rate**

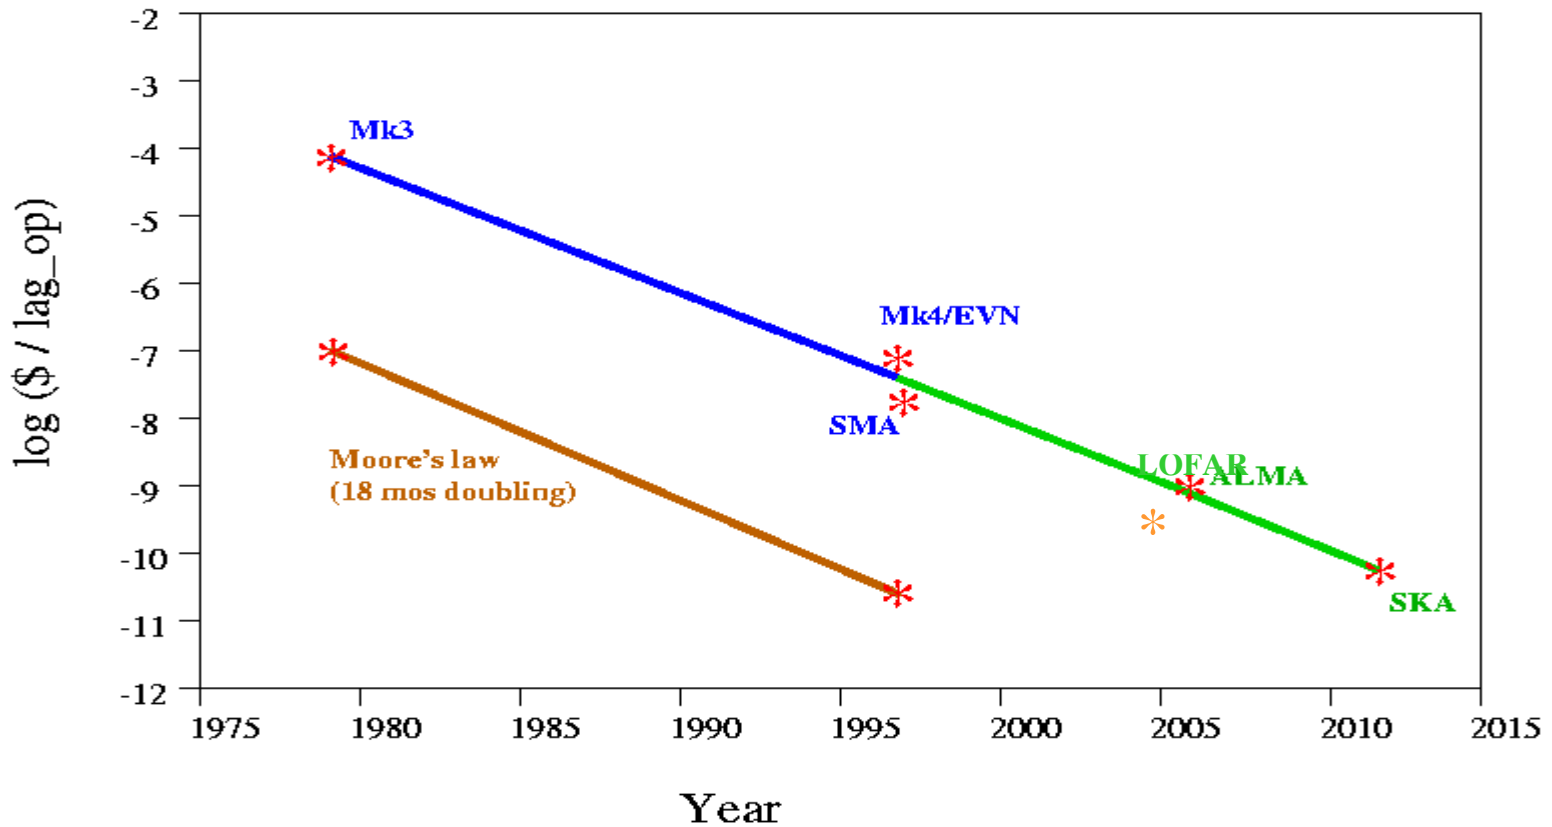
# Historical Extrapolation

- Moore's Law (variation #17): *The cost per given IC function is halving every 18 months.*
- Moore's Law was based on observations of the mass-market for IC's, but a similar relationship seems to hold for radio-astronomy correlators as well
- for comparing XF correlators, a useful *figure of merit* is the number of lag operations (multiply & accumulate) per second

# XF Correlator Throughput

	Mark 3	Mark 4	SMA	ALMA	SKA
stations	8	16	8	64	1000
baselines	28	120	28	2016	499500
boards	84	16	96	-	-
basebands/board	1	128	128	-	-
basebands	84	2048	12288	16128	499500
lags/baseband	8	64	128	128	64
counters/lag	2	2	1	1	2
speed (Msamp/s)	4	32	53	4000	4000
bits/sample	1	2	2	2	2
polarizations	1	1	1	2	2
correlator beams	1	1	1	1	3
throughput (lag_ops/sec)	5.4E+09	8.4E+12	8.3E+13	1.7E+16	1.5E+18

# Cost per XF lag\_operation

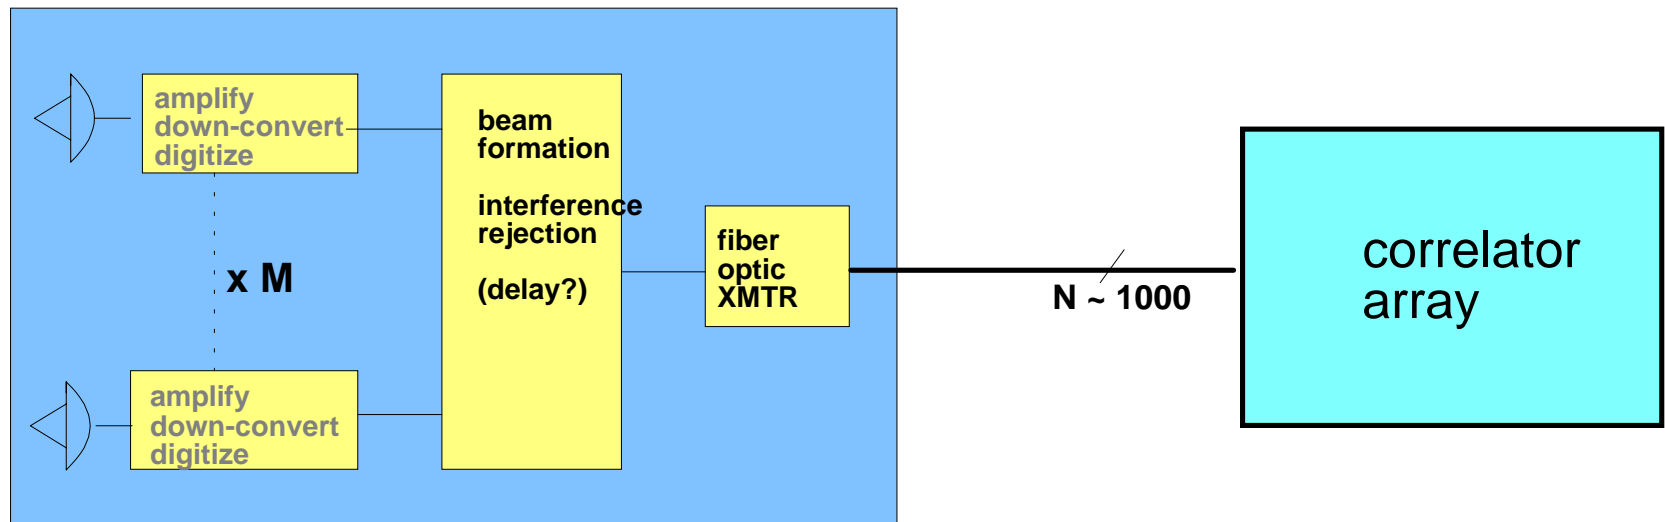


# XF Cost Projection

Mk4 → SKA extrapolation

- assume manufacture in 2012
- assume 20 month Moore's Law for Mk4/SMA → SKA
- extrapolated cost per function:  $5 \times 10^{-11}$  \$ / lag\_op
- extrapolated cost for strawman SKA: \$75 M

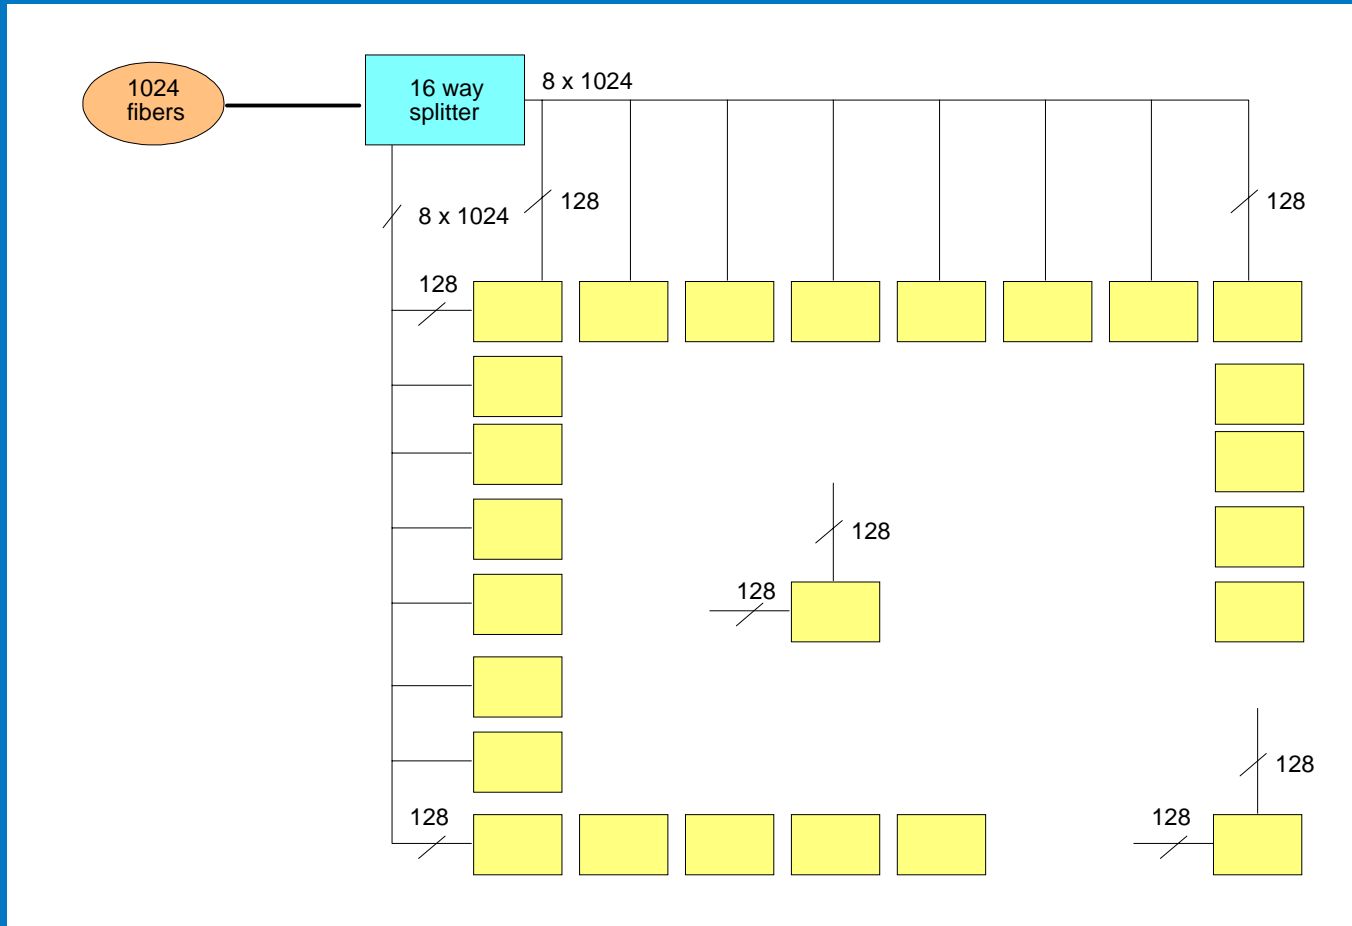
# SKA Signal Hierarchy



(1 of N stations)

# Correlator Array

## w/ Optical Signal Distribution



# Correlator Array Characteristics

- ~500K baselines are distributed through 64 correlator racks, each rack having about 8k baselines
- the 8k baselines to be correlated are half of the 16k product baselines in the off-diagonal racks
- the racks along the diagonal only have 8k baselines to correlate, since both sets of 128 signals are the same

# Correlator Array Characteristics

- there are 256 optical fibers entering each rack, terminating directly on the appropriate circuit boards (perhaps through an optical backplane connector)
- in order not to duplicate the delay stages for correlator beam formation, the delay function could take place at each station

# Correlator Array Characteristics

- if the rack were comprised of 512 correlator chips, such as the Mark 4 rack, each chip would have to perform  $4.5 \times 10^{13}$  lag\_ops / s
- a 1 GHz ASIC correlator chip would need to have 45,000 lags, a factor of 90 greater than the Mark 4 chip, or about 6 times the ALMA chip
- the 2011 ITR projection for transistors/chip would predict a capacity of about 3 Mlags/chip

# FX vs. XF Correlators

- XF correlators: time samples are cross-correlated by baseline, summed in hardware (sec time-scale), then Fourier transformed into cross-power spectra via software
- FX correlators: time samples are Fourier transformed by station, the resultant spectra are cross-multiplied (ms timescale) and summed

# FX vs. XF Correlators

- except for station deformatting and possibly fringe-rotation, XF correlators do everything for  $\sim N^2 / 2$  baselines
- FX correlators do only  $N$  FFT's, and  $\sim N^2 / 2$  cross-multiplies

# FX vs. XF Correlators

- in FX correlators, there is little penalty for large spectra...
- ...*but* FX correlators multiply data rate at FFT output by perhaps a factor of 4 - 32, increasing demands on signal routing
- e.g. - NRAO chose XF for the ALMA correlator, citing costs of interconnection and 2 custom chip designs (FFT, Xmult)

# FX Strawman based on DSP's

- year 2000 implementation using Analog Dev. Tiger-SHARC (0.9 Gflop/s; 1024 pt. FFT in 69  $\mu$ s; ~\$150)
- need total of ~ 9 million TS's, dominated by cross-multiply & accumulate per baseline
- current cost of 9 million TS's: ~ \$1300M!
- in 2012 ITR predicts equivalent DSP cost of \$5.2M in 12K chips

# FX Strawman based on FPGA's

- Xilinx Virtex II chip:
  - High-speed I/O channels; 3.5 Mb on-chip RAM
  - embedded Power PC
  - 600 billion MAC/s; 1024 pt. FFT in  $< 1 \mu\text{s}$
- Would require  $\geq 72,000$  Virtex II's at a current cost of  $\sim \$70\text{M}$

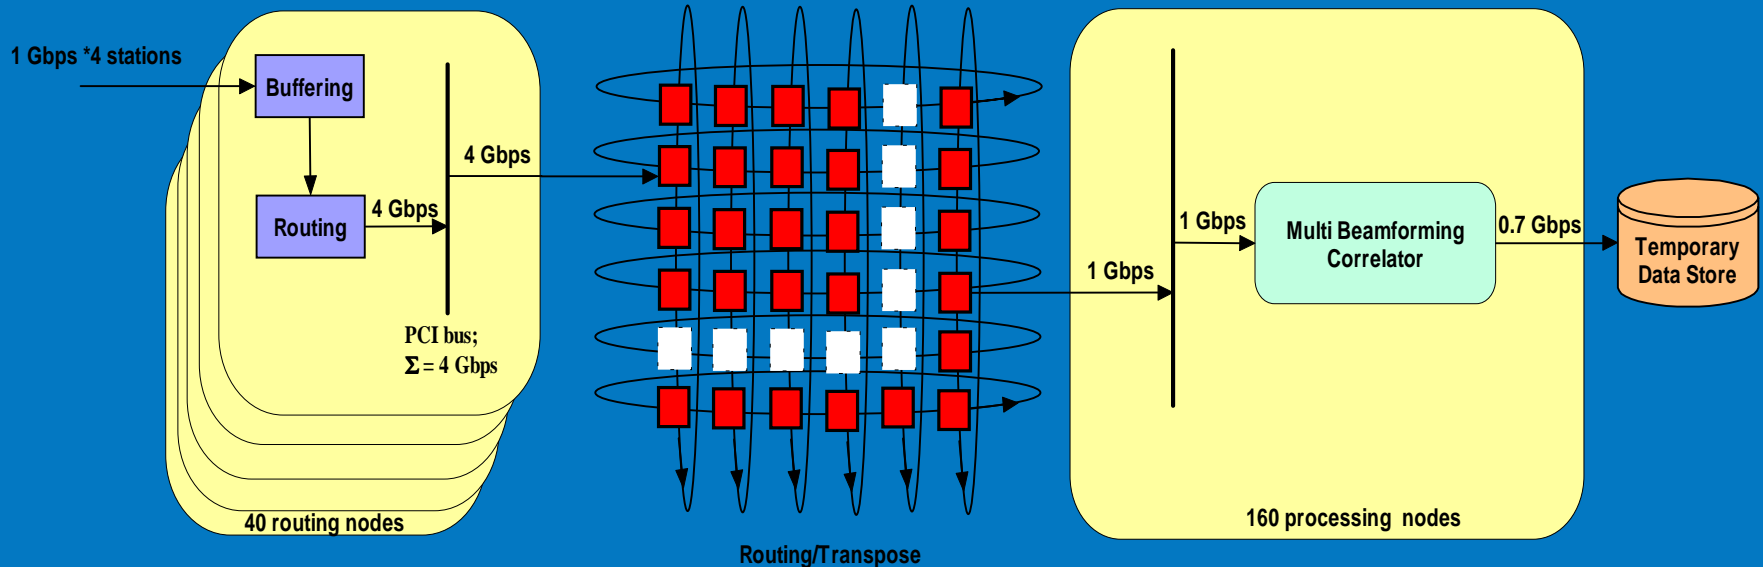
# Interconnection Developments

- **ATA** suggested dealing with interconnection problem via IP packetization, using COTS routers and PC's with FPGA-based signal processing boards
- **LOFAR** using array of several hundred PC's for central correlator, communicating with SCI protocol across a routing torus

# LOFAR Routing Torus

## LOFAR Central Processing Transpose function and partitioning

(input: 16 Beams \* 4 MHz = 64 MHz total)



**160+40=200 Node 2D Torus:**  
 $\Sigma = 160 \text{ Gbps}$

# LOFAR Transient Data Buffer

- Primary FOV about 1 steradian
- Always running, except when frozen by triggering events:
  - external (priority) request (e.g. LIGO or GRB)
  - auto-detection
- *freeze epoch* for flexibility
- Large RAM, cyclically written

## TDB - Size and Cost

- 165 stations, 81 el., 2 pol., 4 MHz BW, 2 bytes/sample = 428 GB/s
- for 100 s: 43 TB for 165 stations
- Currently 256 MB RAM costs \$40
- \$7M for the full array
- in 5 years: \$700K (thanks to St. Moore)

# Buffer Download Options

- In background, data transport rate determines duty cycle of buffer
- High data volume capture might be compacted to single station beam before download
- Save to disks upon receipt at processor, for intensive post-analysis
- Data may be combined into interferometer beam to decrease size (as much as 40 TB)

# SKA Transient Comparison

- Smaller primary FOV (could use subarrays)
- higher sample rate
- fewer stations
- greater RFI rejection

# Software

- Studies of software technology gains show marginal productivity improvement
- Mk3: ~250K lines of Fortran & assembly language; ~20 man-years
- Mk4: ~150K lines of C; ~15 man-years
- modest reduction in software effort, but for substantial increase in functionality
- conclusion: user demand is a perfect gas

# Conclusions

- a 1000 station (probably FX) SKA correlator looks buildable, albeit expensive
- in this decade, some combination of FPGA's and/or DSP's may become an attractive alternative to ASIC's, especially for large FX correlators
- an SKA Transient Data Buffer is feasible, scientifically interesting, and merits further study
- More\$ Law: timescale for software function/cost doubling approx. a decade