SKA Phase 1:
Costs of Computation

Duncan Hall

CALIM 2010

2010 August 24, 27

Outline

- Motivation
- Phase 1 in a nutshell
- Benchmark from 2001 [EVLA Memo 24]
- Some questions

Amdahl’s law overrides Moore’s law!

- Let T_S be the time spent on all operations and moves in serial
- Let p be the number of processors operating in parallel
- Let f be the fraction of operations performed in parallel
- Then the time for processing in parallel, T_P, is given by:

$$T_P \geq T_S \times [\,(1-f) + f\,/\,p\,]$$

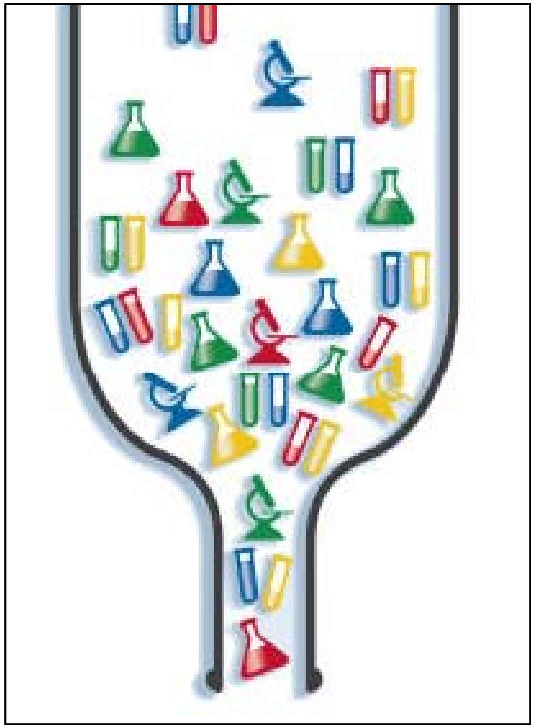

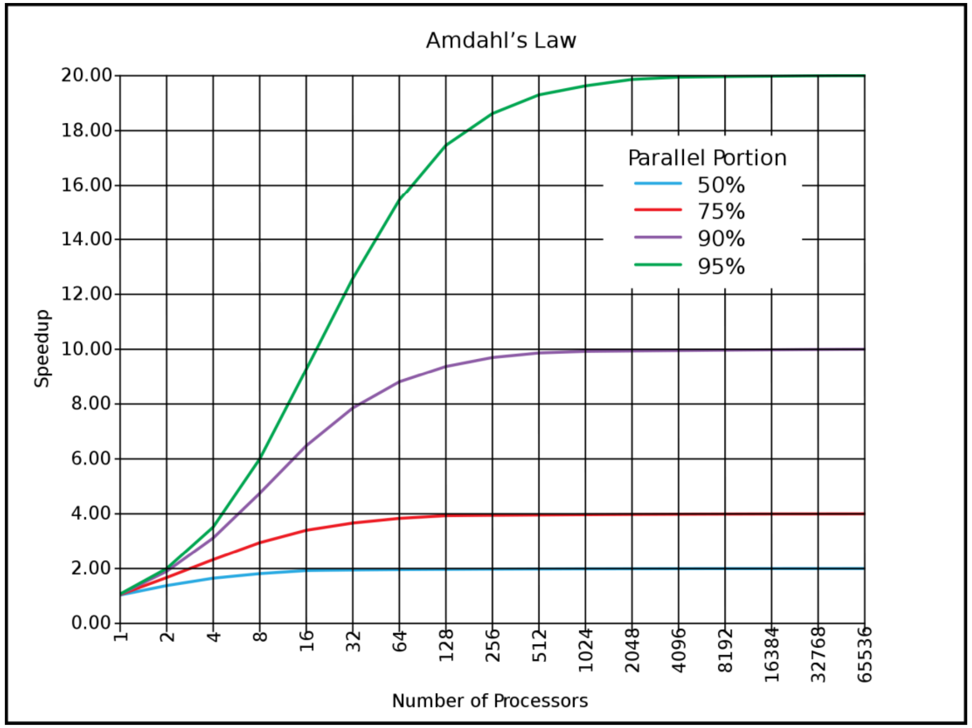

How much is an Exaflop?

3 x 10^{11} stars ➔

1 Exaflop = 10^{18} 32-bit floating point operations <u>per second</u>

10^{18} ≈ number of stars in **<u>3 million milky way galaxies</u>**

Pushing the Flops envelope:

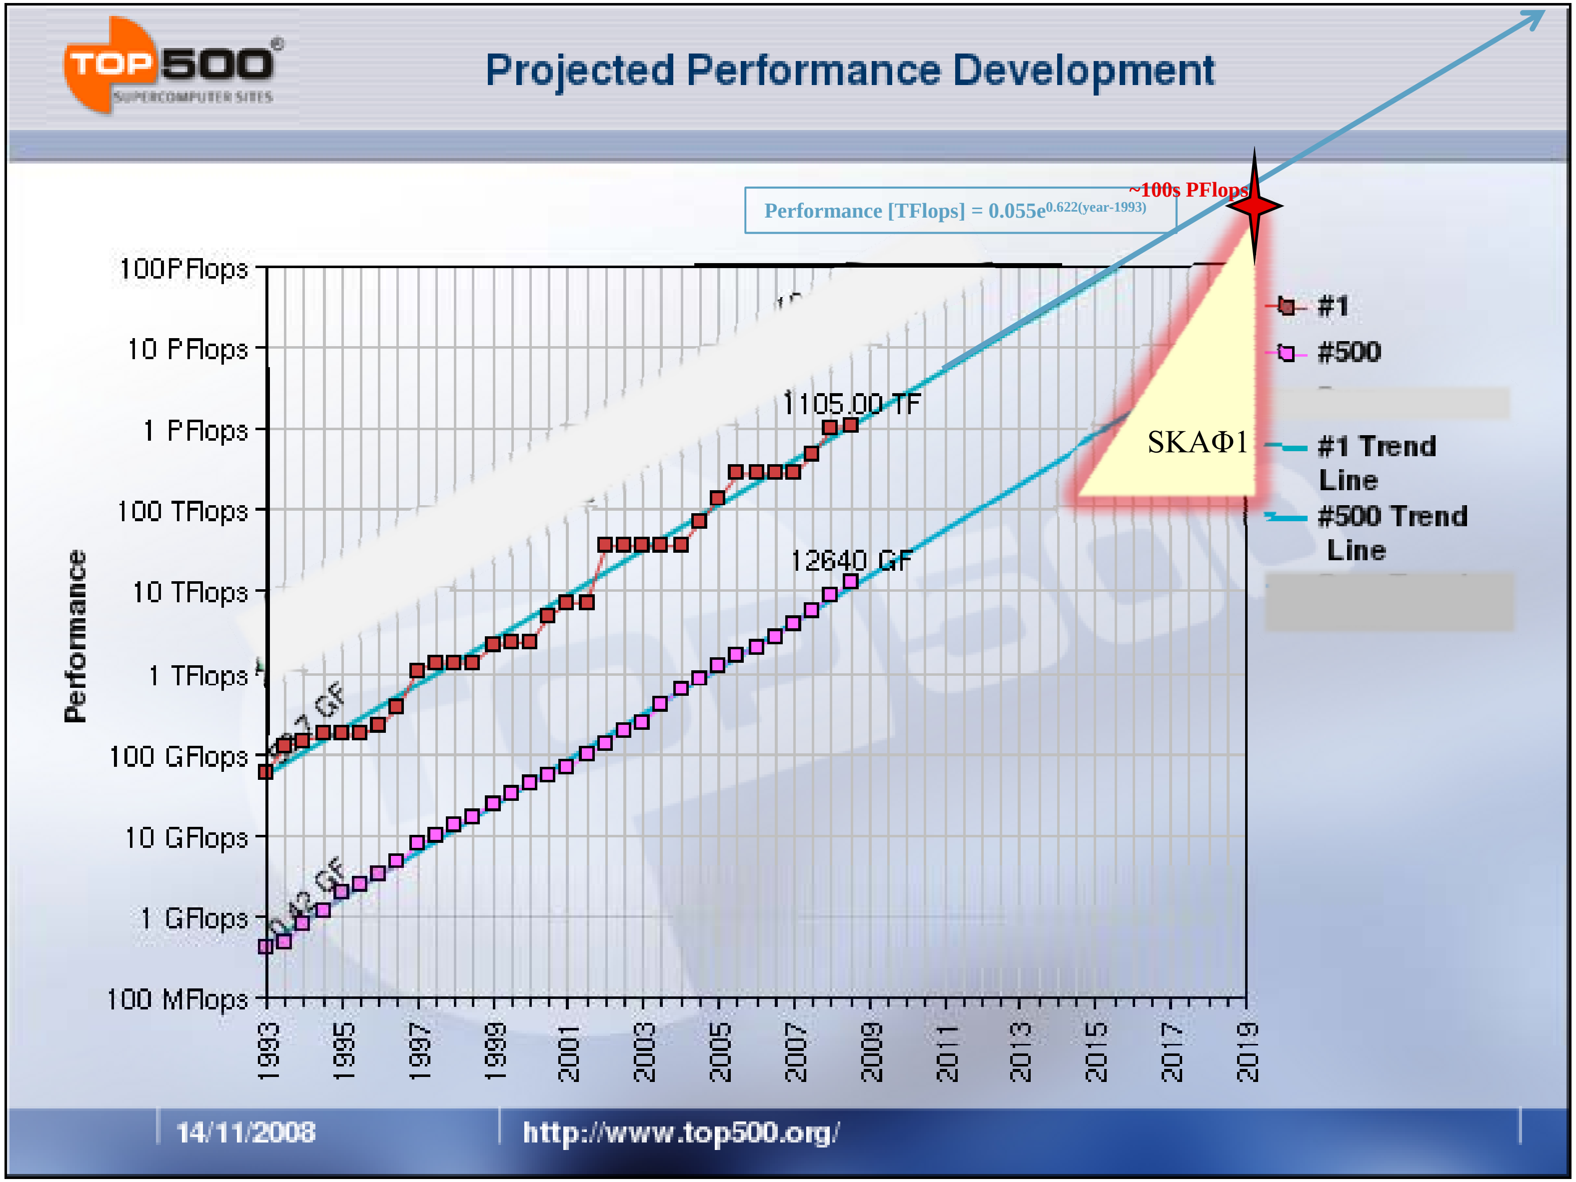

Cornwell and van Diepen "Scaling Mount Exaflop: from the pathfinders to the Square Kilometre Array" http://www.atnf.csiro.au/people/Tim.Cornwell/MountExaflop.pdf

Notes: [1] MACs ≠ FLOPs;
[2] Lines for ASIC and FPGA are for devices only

CPU cabinets for ~1 petaflop Cray Jaguar occupy 560+ square metres

Satellite view of data centre building: chillers on roof; ~1,000 square metres per floor

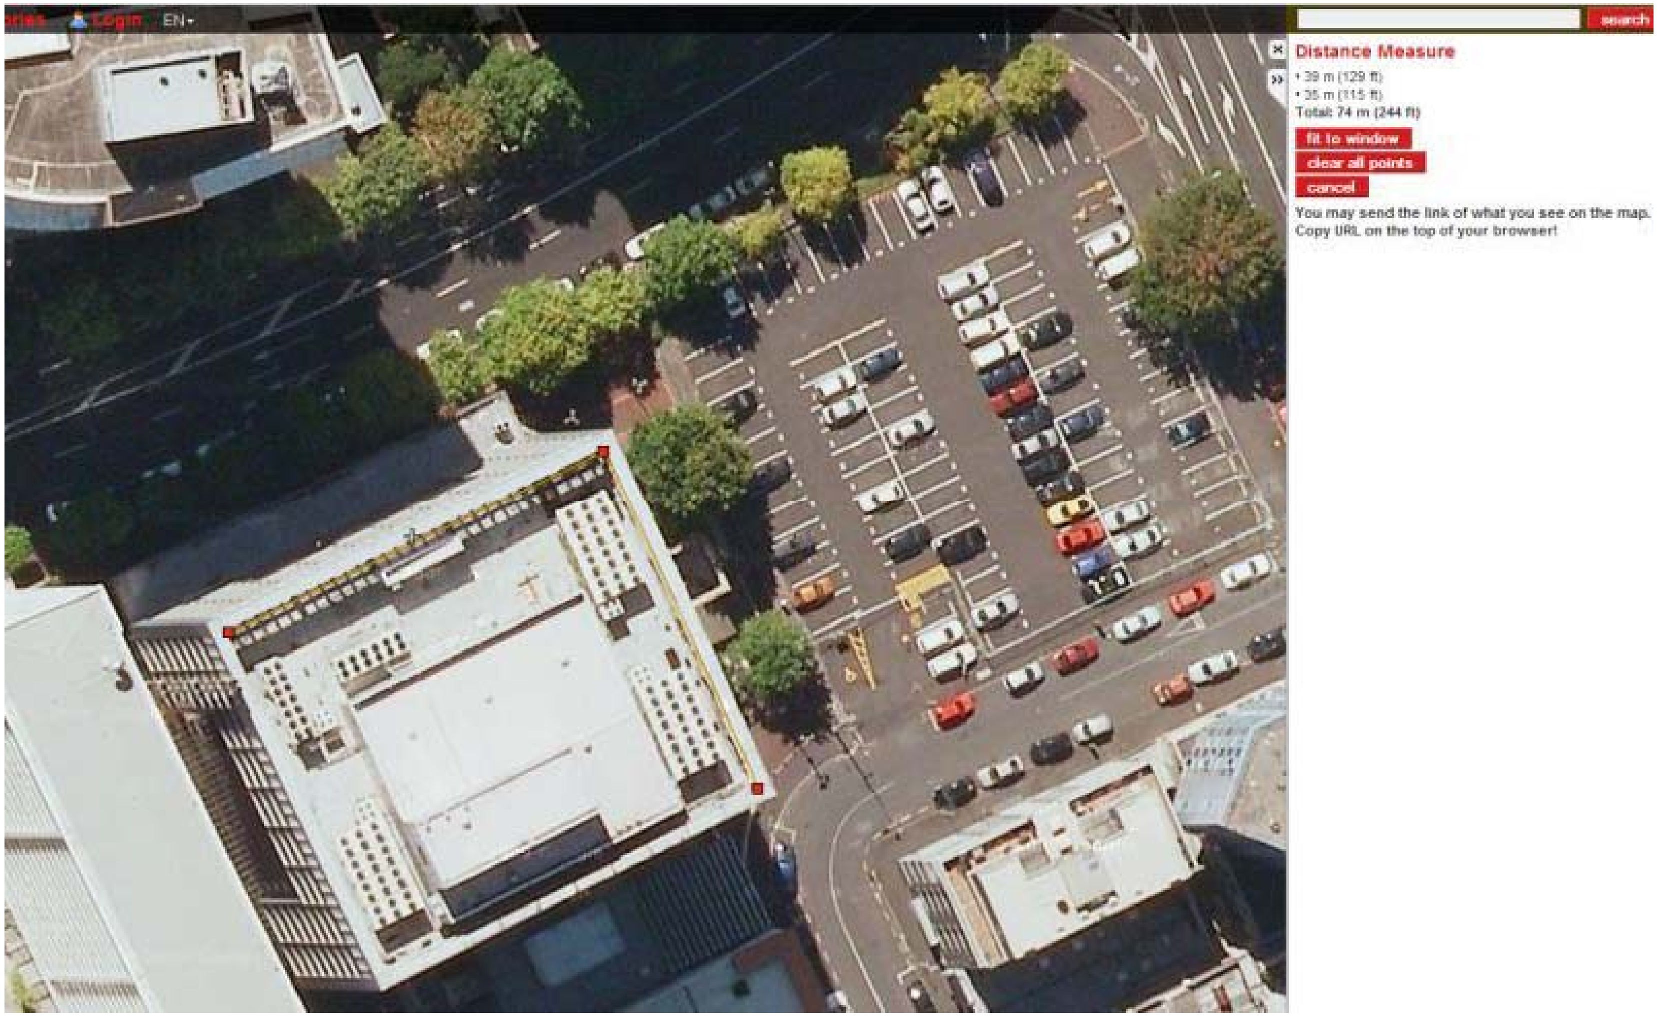

Building cost for data centres: ~€10,000 per square metre

Include power, (-H)VAC, data storage, telecommunications, security ...

Outline

Motivation

Phase 1 in a nutshell:

3,000,000 : 1 dynamic range in ~2018

Benchmark from 2001 [EVLA Memo 24]

Some questions

Dynamic range: historical progress and target for SKA Phase 1:

Noordarm *et al* (1982)	3C84	WSRT 1.4 GHz	10,000:1
Geller *et al* (2000)	1935-692	ATCA 1.4 GHz	77,000:1
de Bruyn & Brentjens (2005)	Perseus	WSRT 92 cm	400,000:1
de Bruyn *et al* (2007)	3C147	WSRT 1.4 GHz	1,000,000:1

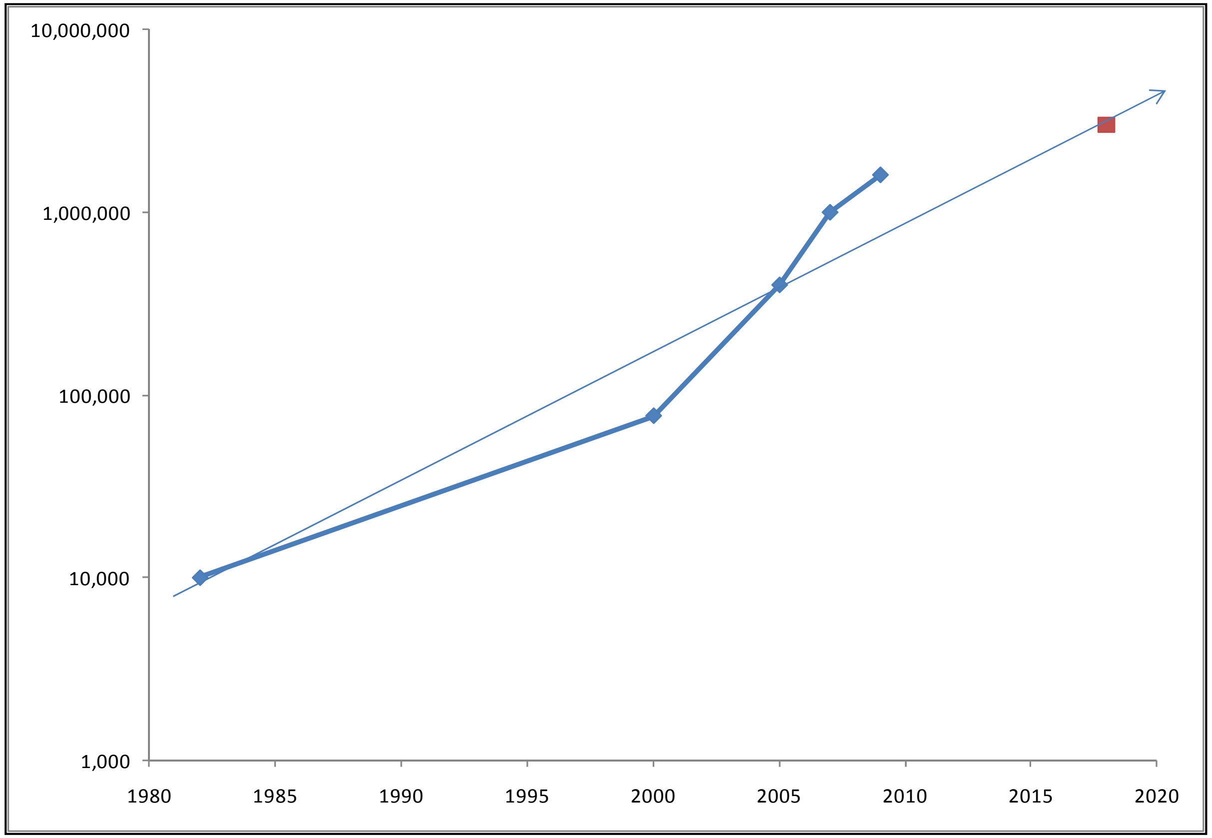

Kemball: "Array Calibration" SA SKA 2009
Smirnov: "Luxury Problems of High Dynamic Range Imaging" SKA 2010

Outline

- Motivation
- Phase 1 in a nutshell
- Benchmark from 2001 [EVLA Memo 24]
- Some questions

2001 algorithm performance:

Peak data rate	25 MB/s
Data for Peak 8-hr observation	700GB
flops per float	100 - 10000
Peak compute rate	5Tflop
Average/Peak computing load	0.1
Average compute rate	0.5Tflop
Turnaround for 8-hr peak observation	40 minutes
Average/Peak data volume	0.1
Data for Average 8-hr observation	70GB
Data for Average 1-yr	80TB

Table I: Typical and peak data and computing rates for the EVLA

T. Cornwell "EVLA Memo 24: Computing for EVLA Calibration and Imaging", 2001 January12

Outline

- Motivation
- Phase 1 in a nutshell
- Benchmark from 2001 [EVLA Memo 24]
- Some questions

At first order, only a few key parameters define Phase 1 computing:

Description	Assumption or Derivation	Reference	Units	Dishes	D+WBSPFs	Sparse Aas	Sum
Maximum baseline length	2 x maximum radius of 100 km	SKA_phase1_definition_v0 1	metres	200.0E+3		200.0E+3	
Dish or station diameter		SKA_phase1_definition_v0 1	metres	15		180	
Number of dishes or stations	*n*	SKA_phase1_definition_v0 1		250		50	
Number of unique baselines	Calculated: $n(n-1)/2$			31,125		1,225	
Maximum frequency of operation		SKA_phase1_definition_v0 2	Hertz	2.0E+9		450.0E+6	
Minimum frequency of operation	Only one Feed available at a time	SKA_phase1_definition_v0 2	Hertz	1.0E+9		70.0E+6	
Fractional bandwidth		Astro2010; DRM		1.0		1.0	
Instantaneous bandwidth	(Max freq - Min freq) x Fractional bandwidth	SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+9		380.0E+6	
Frequency resolution		SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+3		1.0E+3	
Number of frequency channels		SKA_phase1_definition_v0 2		67.0E+3		67.0E+3	
Number of beams formed per dish or station		SKA_phase1_definition_v0 1		1		480	
Number of polarisation products				4		4	
Number of floats per complex float				2		2	
Calculated parameter for use in Smearing	(Maximum baseline length) / (Dish or station diameter)			13.3E+3		1.1E+3	
		SKA_phase1_definition_v0 2	Hertz	5.0E+0		250.0E-3	

Estimated hardware for Phase 1 ranges into hundreds of petaflops

Description	Assumption or Derivation	Reference	Units	Dishes	D+WBSPFs	Sparse Aas	Sum
Maximum baseline length	2 x maximum radius of 100 km	SKA_phase1_definition_v0 1	metres	200.0E+3		200.0E+3	
Dish or station diameter		SKA_phase1_definition_v0 1	metres	15		180	
Number of dishes or stations	*n*	SKA_phase1_definition_v0 1		250		50	
Number of unique baselines	Calculated: $n(n-1)/2$			31,125		1,225	
Maximum frequency of operation		SKA_phase1_definition_v0 2	Hertz	2.0E+9		450.0E+6	
Minimum frequency of operation	Only one Feed available at a time	SKA_phase1_definition_v0 2	Hertz	1.0E+9		70.0E+6	
Fractional bandwidth		Astro2010; DRM		1.0		1.0	
Instantaneous bandwidth	(Max freq - Min freq) x Fractional bandwidth	SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+9		380.0E+6	
Frequency resolution		SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+3		1.0E+3	
Number of frequency channels		SKA_phase1_definition_v0 2		67.0E+3		67.0E+3	
Number of beams formed per dish or station		SKA_phase1_definition_v0 1		1		480	
Number of polarisation products				4		4	
Number of floats per complex float				2		2	
Calculated parameter for use in Smearing	(Maximum baseline length) / (Dish or station diameter)			13.3E+3		1.1E+3	
		SKA_phase1_definition_v0 2	Hertz	5.0E+0		250.0E-3	
	Assume pipeline processing in near realtime						
Dump rate in floating point numbers	All visibilities have the same limiting dump rate		floats/sec	83.4E+9		78.8E+9	162.2E+9
Required flops per float - optimistic	Assume can achieve 10^7 dynamic range (?)	Advice from ASTRON, CSIRO, TDP-CPG		100,000		100,000	
Required flops per float - pessimistic	Assume can achieve 10^7 dynamic range (?)	Advice from ASTRON, CSIRO, TDP-CPG		400,000		400,000	
Required flops - optimistic				8.3E+15		7.9E+15	16.2E+15
Required flops - pessimistic				33.4E+15		31.5E+15	64.9E+15
Estimated HPC efficiency - optimistic	Refer to [A] at bottom of this colmn	20091116 news release from Cray		50%		50%	
Estimated HPC efficiency - realistic	Refer [B] at bottom of this column	Hoisie et al; DOI: 10.1177/109434200001400405		10%		10%	
Required HPC flops - optimistic	Calculated			16.7E+15		15.8E+15	32.4E+15
Required HPC flops - pessimistic	Calculated			333.7E+15		315.2E+15	648.8E+15

CPG Memo 3 (2009-11-6) confirms requirements for extreme scale computing:

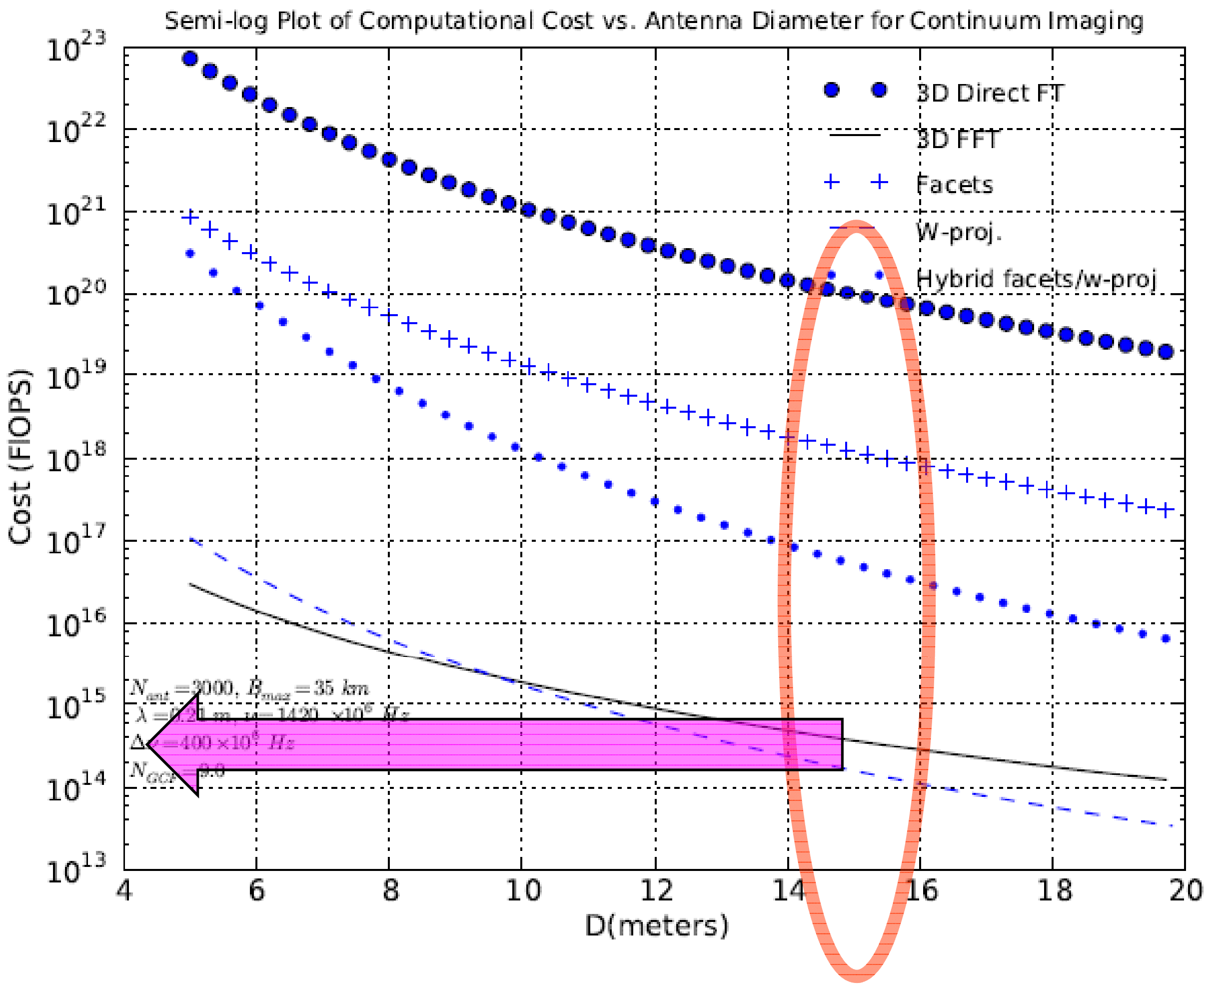

Figure 1: Semi-log y plots of computational costs (without consideration of deconvolution and parallel computing efficiency η) vs. antenna diameter D for continuum imaging for the 3-D direct FT, 3-D FFT, facets, w-projection, and hybrid facets/w projection imaging algorithms.

One driver: smearing <2%

Description	Assumption or Derivation	Reference	Units	Dishes	D+WBSPFs	Sparse Aas	Sum
Maximum baseline length	2 x maximum radius of 100 km	SKA_phase1_definition_v0 1	metres	200.0E+3		200.0E+3	
Dish or station diameter		SKA_phase1_definition_v0 1	metres	15		180	
Number of dishes or stations	*n*	SKA_phase1_definition_v0 1		250		50	
Number of unique baselines	Calculated: *n* (*n* -1)/2			31,125		1,225	
Maximum frequency of operation		SKA_phase1_definition_v0 2	Hertz	2.0E+9		450.0E+6	
Minimum frequency of operation	Only one Feed available at a time	SKA_phase1_definition_v0 2	Hertz	1.0E+9		70.0E+6	
Fractional bandwidth		Astro2010; DRM		1.0		1.0	
Instantaneous bandwidth	(Max freq - Min freq) x Fractional bandwidth	SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+9		380.0E+6	
Frequency resolution		SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+3		1.0E+3	
Number of frequency channels		SKA_phase1_definition_v0 2		67.0E+3		67.0E+3	
Number of beams formed per dish or station		SKA_phase1_definition_v0 1		1		480	
Number of polarisation products				4		4	
Number of floats per complex float				2		2	
Calculated parameter for use in Smearing	(Maximum baseline length) / (Dish or station diameter)			13.3E+3		1.1E+3	
		SKA_phase1_definition_v0 2	Hertz	5.0E+0		250.0E-3	
	Assume pipeline processing in near realtime						
Dump rate in floating point numbers	All visibilities have the same limiting dump rate		floats/sec	83.4E+9		78.8E+9	162.2E+9
Required flops per float - optimistic	Assume can achieve 10^7 dynamic range (?)	Advice from ASTRON, CSIRO, TDP-CPG		100,000		100,000	
Required flops per float - pessimistic	Assume can achieve 10^7 dynamic range (?)	Advice from ASTRON, CSIRO, TDP-CPG		400,000		400,000	
Required flops - optimistic				8.3E+15		7.9E+15	16.2E+15
Required flops - pessimistic				33.4E+15		31.5E+15	64.9E+15
Estimated HPC efficiency - optimistic	Refer to [A] at bottom of this colmn	20091116 news release from Cray		50%		50%	
Estimated HPC efficiency - realistic	Refer [B] at bottom of this column	Hoisie et al; DOI: 10.1177/109434200001400405		10%		10%	
Required HPC flops - optimistic	Calculated			16.7E+15		15.8E+15	32.4E+15
Required HPC flops - pessimistic	Calculated			333.7E+15		315.2E+15	648.8E+15

Where does the "smearing <2%" come from?

1.4. System-Level Design Specifications

During the course of the development of the Design Reference Mission, it was recognized that several of the design parameters transcend specific science cases. These parameters are considered to be "system-level" specifications, which is also consistent with the system-level approach recommended by the SKA International Engineering Advisory Committee (IEAC) and the design approach adopted under the aegis of PrepSKA. Table 1.1 summarizes these and discussion of each follows.

Table 1.1. SKA-mid and SKA-lo System-Level Design Specifications

Parameter	Value
Fractional Instantaneous Bandwidth	~ 1 (continuum observations)
Spectral Baseline	Sufficiently flat to enable spectral line observations
Correlator Integration Time	Sufficient to mitigate time-average smearing
Spectral Resolution	Sufficient to mitigate bandwidth smearing
Survey "On Sky" Time	2 yr
Deep Field Integration Time	1000 hr (~ 3 Ms)

SKA DRM v. 1.0 2010 March 16

The DRM asserts that smearing shall be <2%

1.4.3. Correlator Integration Time

The integration time provided by the correlator should be sufficient so that time-average smearing at the edge of the field of view is not larger than 2%.

1.4.4. Spectral Resolution

The spectral resolution provided by the correlator should be sufficient so that bandwidth smearing at the edge of the field of view is not larger than 2%. The spectral resolution required to enable identification and excision of radio frequency interference (RFI) may be higher, and it may be site dependent. Therefore, we do not consider RFI considerations in the remainder of this document.

Example SKA Phase 1 dish configurations: 0.3 ~ 3 dumps/s ?

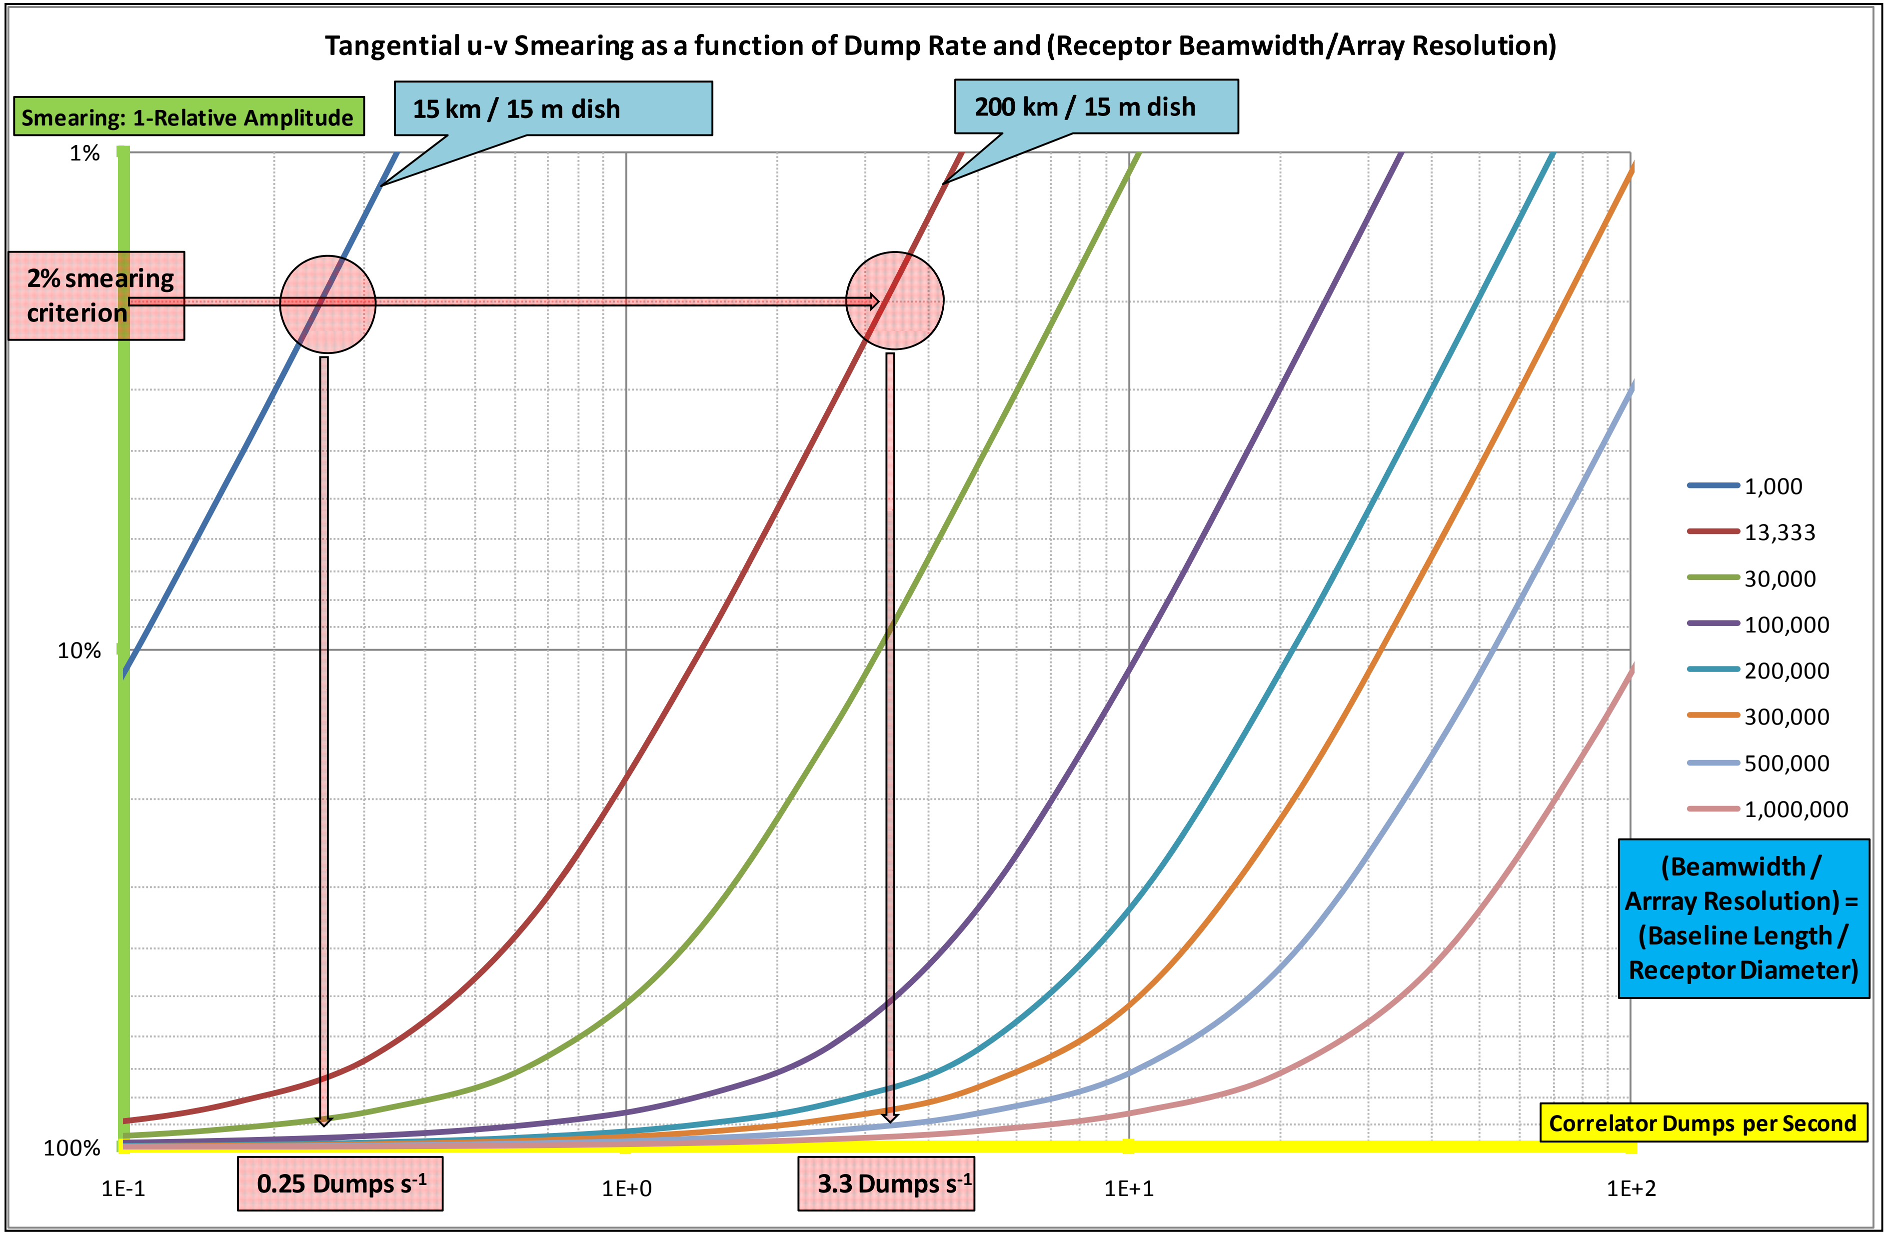

Bridle and Schwab's approximations:

functions applied to the data. For a synthesis image of a source near the North or South Celestial Pole, the average fractional reduction in amplitude $\overline{\langle R_\tau \rangle}$ produced by time averaging for a source a distance θ from the phase-tracking center can therefore be written in the simple form

$$\overline{\langle R_\tau \rangle} \approx 1 - \frac{\alpha\pi^2}{12}\,\omega_e^2\tau_a^2\left(\frac{\theta}{\theta_{\mathrm{HPBW}}}\right)^2 , \qquad (18\text{–}40)$$

which is valid in the regime of small intensity losses. We now evaluate the constant α for a few simple cases:

Square coverage, without tapering For square (u,v) coverage of side A (see Eq. 18–14, the beam is given by Equation 18–15, so θ_{HPBW}=1.206/A. For this case, $\overline{L_X^2 + L_Y^2} = A^2\lambda^2/6$, i.e., $\alpha = \frac{1.206^2}{6} = 0.2424$. The average intensity loss factor for a circumpolar point source is therefore

$$\overline{\langle R_\tau \rangle} = 1 - 1.05 \times 10^{-9}\left(\frac{\theta}{\theta_{\mathrm{HPBW}}}\right)^2 \tau_a^2 , \qquad (18\text{–}41)$$

assuming that the loss due to time-average smearing is small.

Circular coverage, without tapering For circular (u,v) coverage of diameter D, the beam has $\theta_{\mathrm{HPBW}} = 1.410/D$. For this case, $\overline{L_X^2 + L_Y^2} = D^2\lambda^2/8$, i.e. $\alpha = \frac{1.410^2}{8} = 0.2485$. The average intensity loss factor for a circumpolar point source is therefore

$$\overline{\langle R_\tau \rangle} = 1 - 1.08 \times 10^{-9}\left(\frac{\theta}{\theta_{\mathrm{HPBW}}}\right)^2 \tau_a^2 , \qquad (18\text{–}42)$$

assuming that the loss due to time-average smearing is small.

Circular coverage with Gaussian tapering If the array produces a Gaussian beam with FWHM θ_{HPBW} (Eq. 18–22), the (u,v) distribution must approximate its transform (Eq. 18–21), so that $\overline{u^2 + v^2} = \gamma^2/\pi^2\theta^2_{\mathrm{HPBW}}$ and $\alpha = \gamma^2/\pi^2 = 4(\ln 2)/\pi^2 = 0.2810$. The average intensity loss factor for a circumpolar point source is therefore

$$\overline{\langle R_\tau \rangle} = 1 - 1.22 \times 10^{-9}\left(\frac{\theta}{\theta_{\mathrm{HPBW}}}\right)^2 \tau_a^2 , \qquad (18\text{–}43)$$

again assuming that the loss due to time-average smearing is small.

Bridle and Schwab 1999: “Bandwidth and Time Average Smearing”; Synthesis Imaging in Radio Astronomy II, pp. 380-381

Emil Lenc's online calculator:

Widefield VLBI Calculator

Emil's Homepage

Widefield VLBI Calculator

Frequency	1000	MHz	Wavelength = 30.00 cm
Bandwidth	550	MHz	
Channels	67000		Channel Bandwidth = 8.21 KHz
Maximum Baseline Length	200	km	Synthesized Beam = 309.4 mas
Averaging Time	0.3	s	
Phase Centre Offset	41	arcmin	
Telescope Diameter	15	m	Primary Beam HWHM = 41.9 arcmin

Calculate

Results for a phase centre offset of 41.00 arcmin

Effect	Effect Error	Cumulative Error
Primary Beam	48.51%	48.51%
Bandwidth Smearing	0.21%	48.62%
Time-Average Smearing	0.69%	48.98%
Non-coplanar	100.00%	100.00%

Return to Emil's Home Page

Parse error: syntax error, unexpected $end, expecting ')' in **/home/elenc/public_html/bbclone-0.4.7/var/access.php** on line **150**

http://astronomy.swin.edu.au/~elenc/Calculators/wfcalc.php

But is <2% smearing sufficient for DR = 65dB for SKA Phase 1 ?

1.4.3. Correlator Integration Time

The integration time provided by the correlator should be sufficient so that time-average smearing at the edge of the field of view is not larger than 2%.

1.4.4. Spectral Resolution

The spectral resolution provided by the correlator should be sufficient so that bandwidth smearing at the edge of the field of view is not larger than 2%. The spectral resolution required to enable identification and excision of radio frequency interference (RFI) may be higher, and it may be site dependent. Therefore, we do not consider RFI considerations in the remainder of this document.

Outline

- Motivation
- Benchmark from 2001 [EVLA Memo 24]
- Phase 1 in a nutshell
- Some more questions

The flops per *uv* float question:

Description	Assumption or Derivation	Reference	Units	Dishes	D+WBSPFs	Sparse Aas	Sum
Maximum baseline length	2 x maximum radius of 100 km	SKA_phase1_definition_v0 1	metres	200.0E+3		200.0E+3	
Dish or station diameter		SKA_phase1_definition_v0 1	metres	15		180	
Number of dishes or stations	*n*	SKA_phase1_definition_v0 1		250		50	
Number of unique baselines	Calculated: *n* (*n* -1)/2			31,125		1,225	
Maximum frequency of operation		SKA_phase1_definition_v0 2	Hertz	2.0E+9		450.0E+6	
Minimum frequency of operation	Only one Feed available at a time	SKA_phase1_definition_v0 2	Hertz	1.0E+9		70.0E+6	
Fractional bandwidth		Astro2010; DRM		1.0		1.0	
Instantaneous bandwidth	(Max freq - Min freq) x Fractional bandwidth	SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+9		380.0E+6	
Frequency resolution		SKA1_Concept_Definition_SSEC_draft.pdf	Hertz	1.0E+3		1.0E+3	
Number of frequency channels		SKA_phase1_definition_v0 2		67.0E+3		67.0E+3	
Number of beams formed per dish or station		SKA_phase1_definition_v0 1		1		480	
Number of polarisation products				4		4	
Number of floats per complex float				2		2	
Calculated parameter for use in Smearing	(Maximum baseline length) / (Dish or station diameter)			13.3E+3		1.1E+3	
		SKA_phase1_definition_v0 2	Hertz	5.0E+0		250.0E-3	
	Assume pipeline processing in near realtime						
Dump rate in floating point numbers	All visibilities have the same limiting dump rate		floats/sec	83.4E+9		78.8E+9	162.2E+9
Required flops per float - optimistic	Assume can achieve 10^7 dynamic range (?)	Advice from ASTRON, CSIRO, TDP-CPG		100,000		100,000	
Required flops per float - pessimistic	Assume can achieve 10^7 dynamic range (?)	Advice from ASTRON, CSIRO, TDP-CPG		400,000		400,000	
Required flops - optimistic				8.3E+15		7.9E+15	16.2E+15
Required flops - pessimistic				33.4E+15		31.5E+15	64.9E+15
Estimated HPC efficiency - optimistic	Refer to [A] at bottom of this colmn	20091116 news release from Cray		50%		50%	
Estimated HPC efficiency - realistic	Refer [B] at bottom of this column	Hoisie et al; DOI: 10.1177/109434200001400405		10%		10%	
Required HPC flops - optimistic	Calculated			16.7E+15		15.8E+15	32.4E+15
Required HPC flops - pessimistic	Calculated			333.7E+15		315.2E+15	648.8E+15

How big should m x m be?

A.L. Varbanescu et al. / Building high-resolution sky images using the Cell/B.E. 117

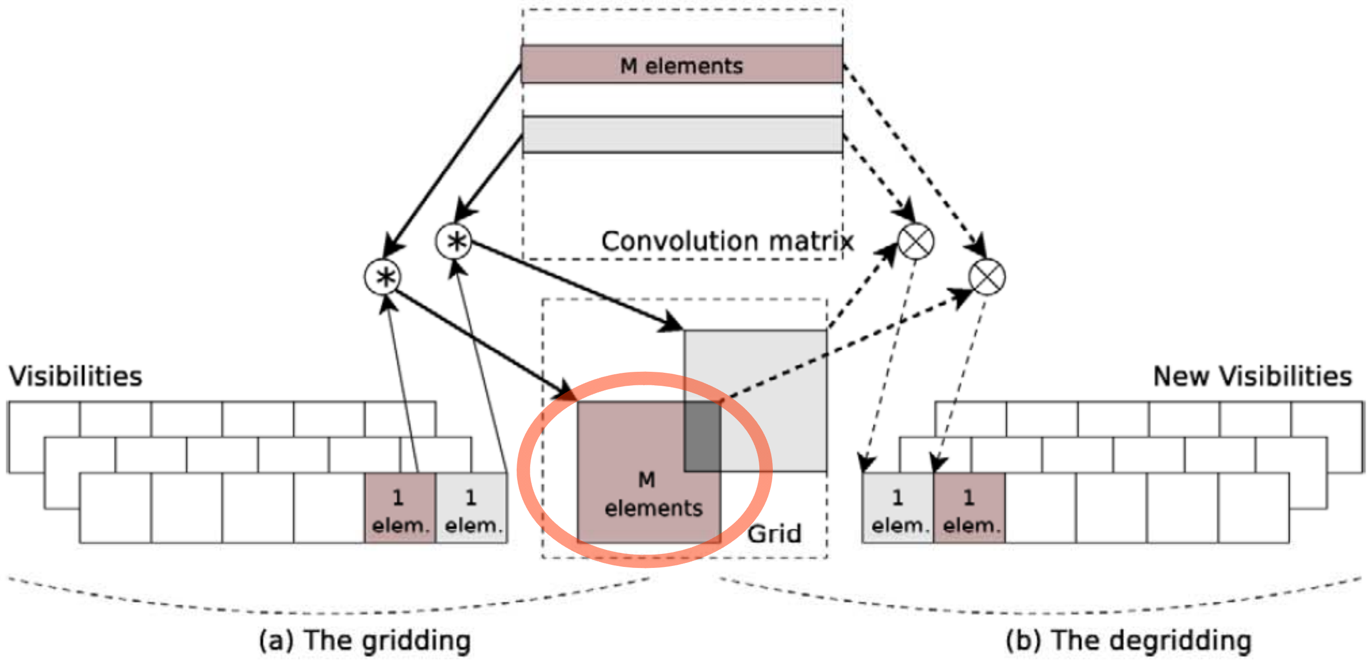

Fig. 4. The two kernels and their data access pattern: (a) gridding (continuous lines) and (b) degridding (dotted lines). Similarly shaded regions are for a single visibility sample. For gridding, newly computed data that overlaps existing grid values is added on. For degridding, newly computed visibilities cannot overlap.

More questions about the 65 dB challenge:

How much “over sampling” is required?

How many “major cycles” are required, worst case?

Alternative algorithms for gridding irregularly spaced samples?

Empirical work for asymmetric side lobes?

Faint sources that may be indistinguishable from imaging artefacts?

Automatic “flagging” and removal of RFI etc.?

Other questions:

- Amdahl’s law ... I/O data rate – e.g. “memory bandwidth?”
- Data cache memory requirements?
- Energy efficiencies of computation and data movement?
- ... ?

Even more questions ...