

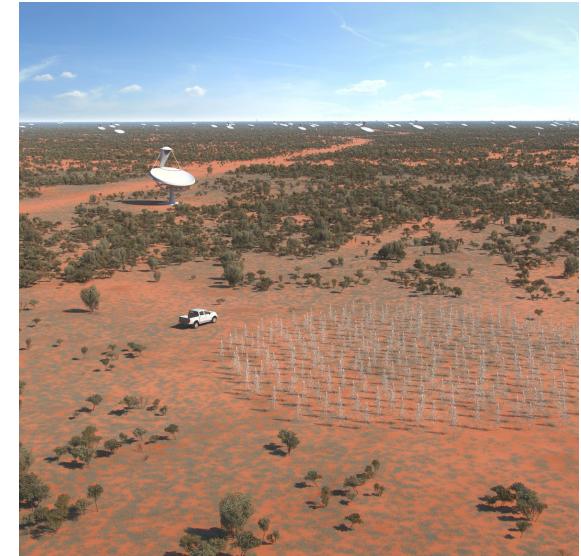
SKA Phase 1 Compute and Power Analysis



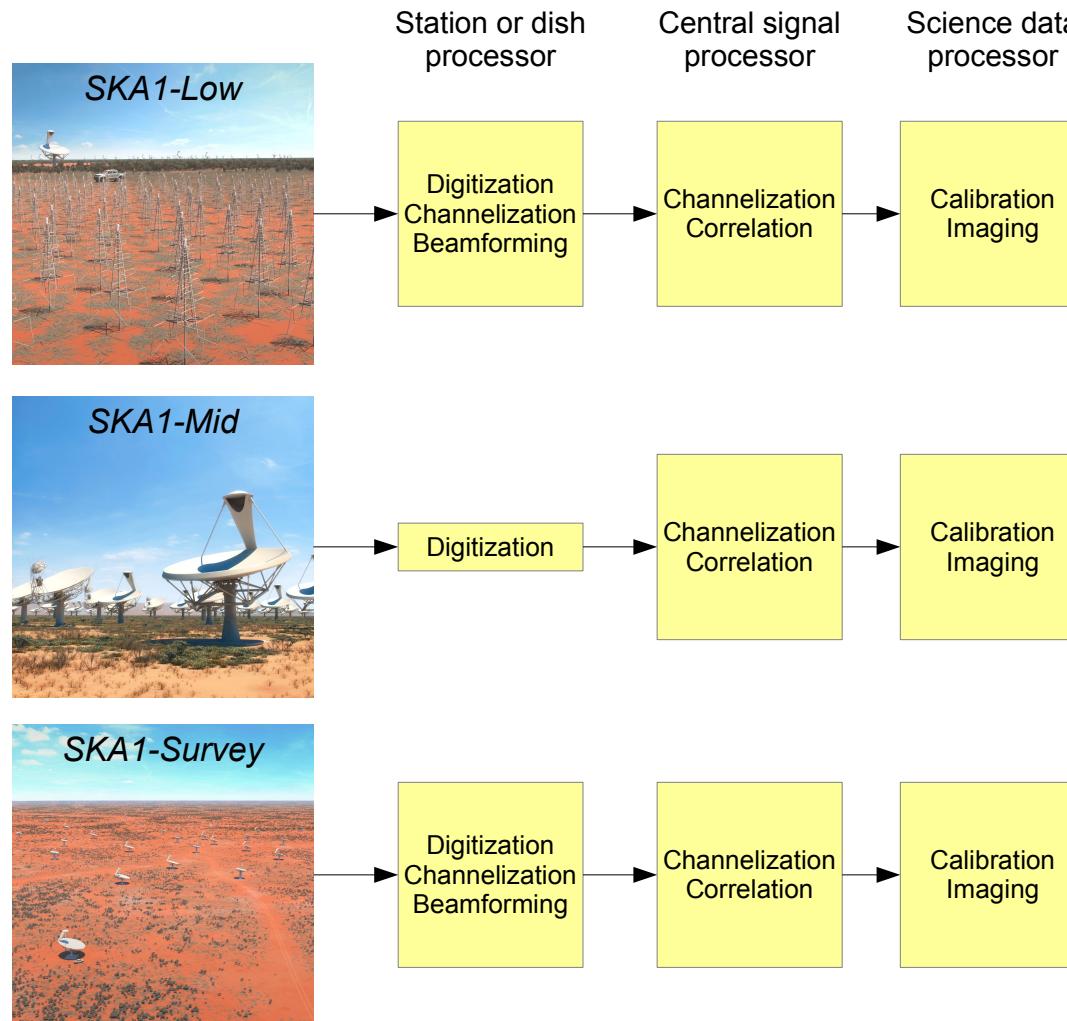
SKA analysis goals



- Understand **compute distribution** for sky imaging with SKA phase 1
- How do these properties relate to the **required compute system**?
 - Flops
 - Energy usage
- Scaling laws retrieved from literature and derived for LOFAR
- Power consumption estimates using Top500 historical data

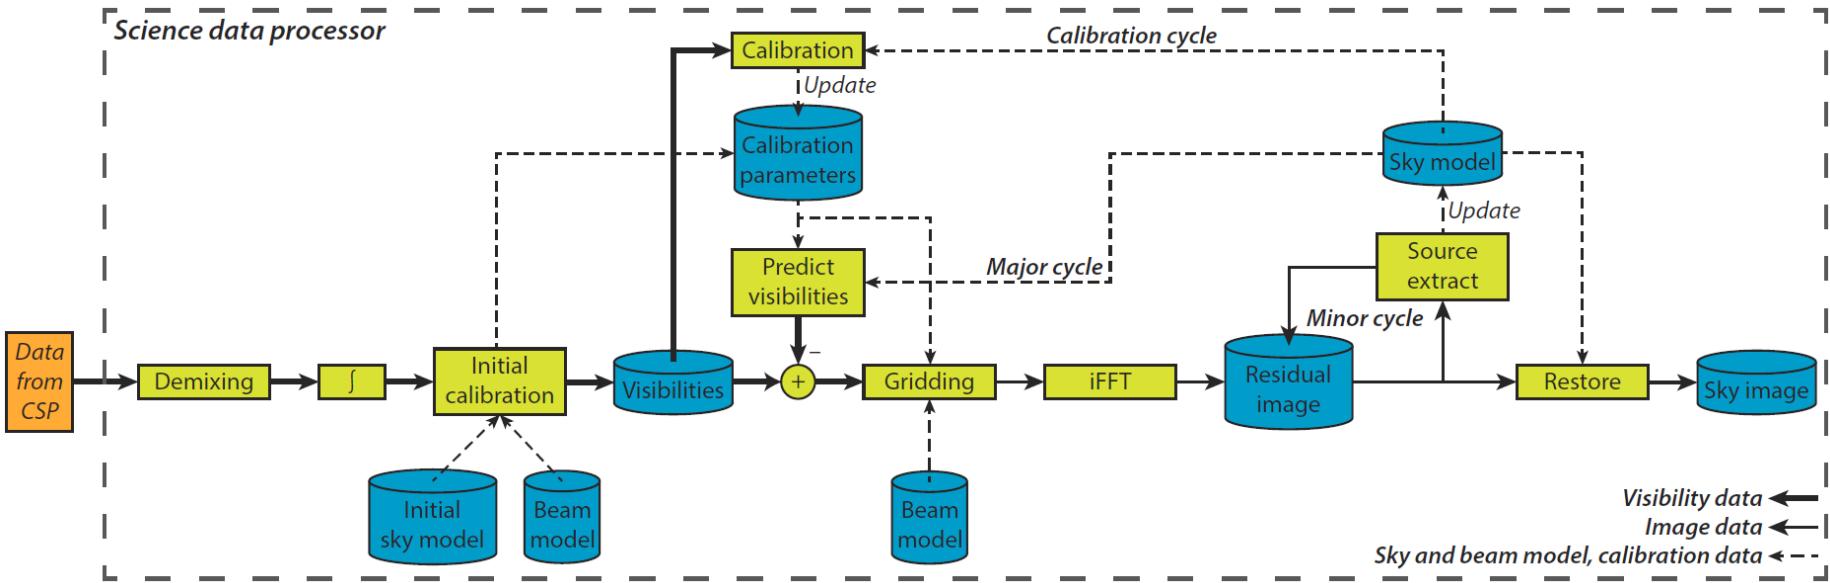
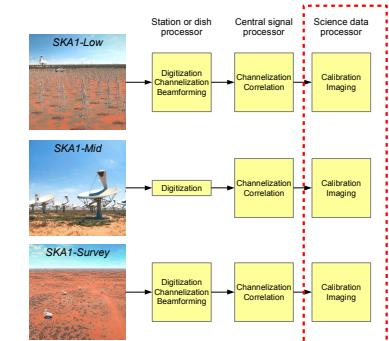


Architecture for digital compute pipeline



- One model for SKA1-Low, SKA1-Mid, and SKA1-Survey

Architecture for the science data processors



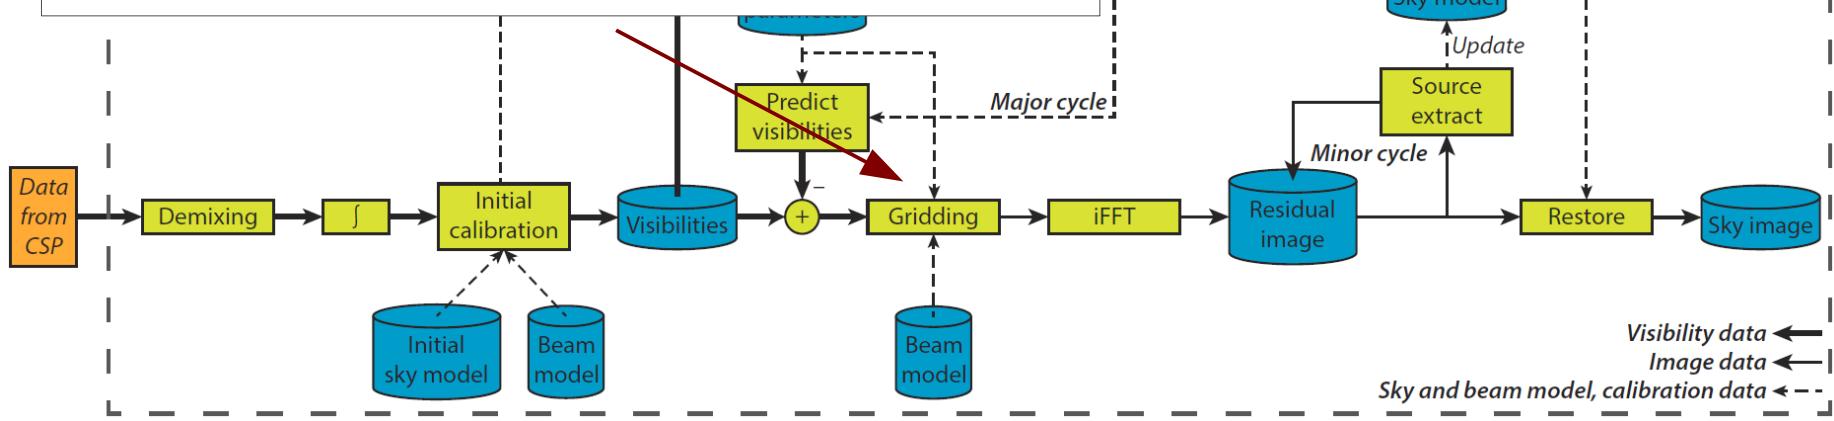
Compute cost of gridding

W-Projection [1] (and other sources):

$$C_{\text{grid}} = N_{\text{op}} 2 V_{\text{chan}} T_{\text{obs}} \sum_{i=0}^{N_{\text{image-chan}} - 1} \left(\left(\frac{w_{\text{rms}}}{w_{\text{max}}} \right)^2 R_{F_i}^2 + R_A^2 \right)$$

where

$$R_{F_i} = \frac{\lambda_i B_{\text{max}}}{D^2}$$



Computing AW-projection kernels

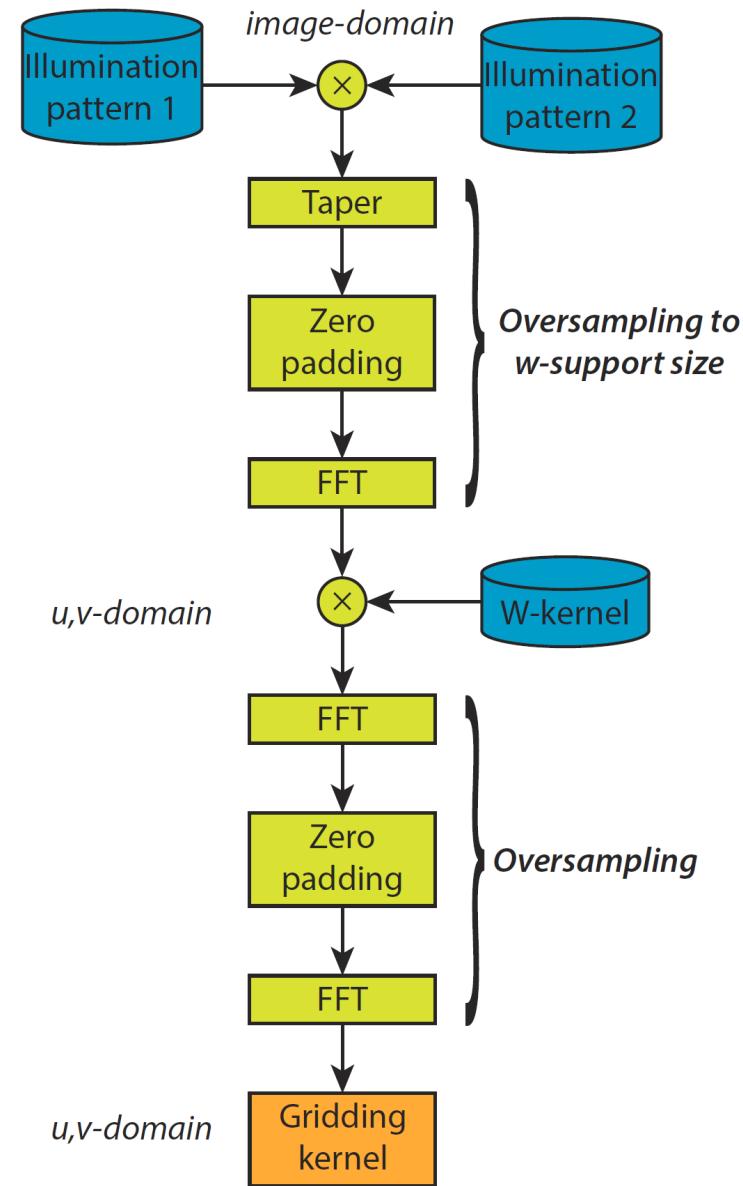
- Last FFT on oversampled data dominates [2]
- Estimate compute based on average kernel size

$$R_{\text{ave},i} = \frac{w_{\text{rms}}}{w_{\text{max}}} \frac{\lambda_i B_{\text{max}}}{D^2}$$

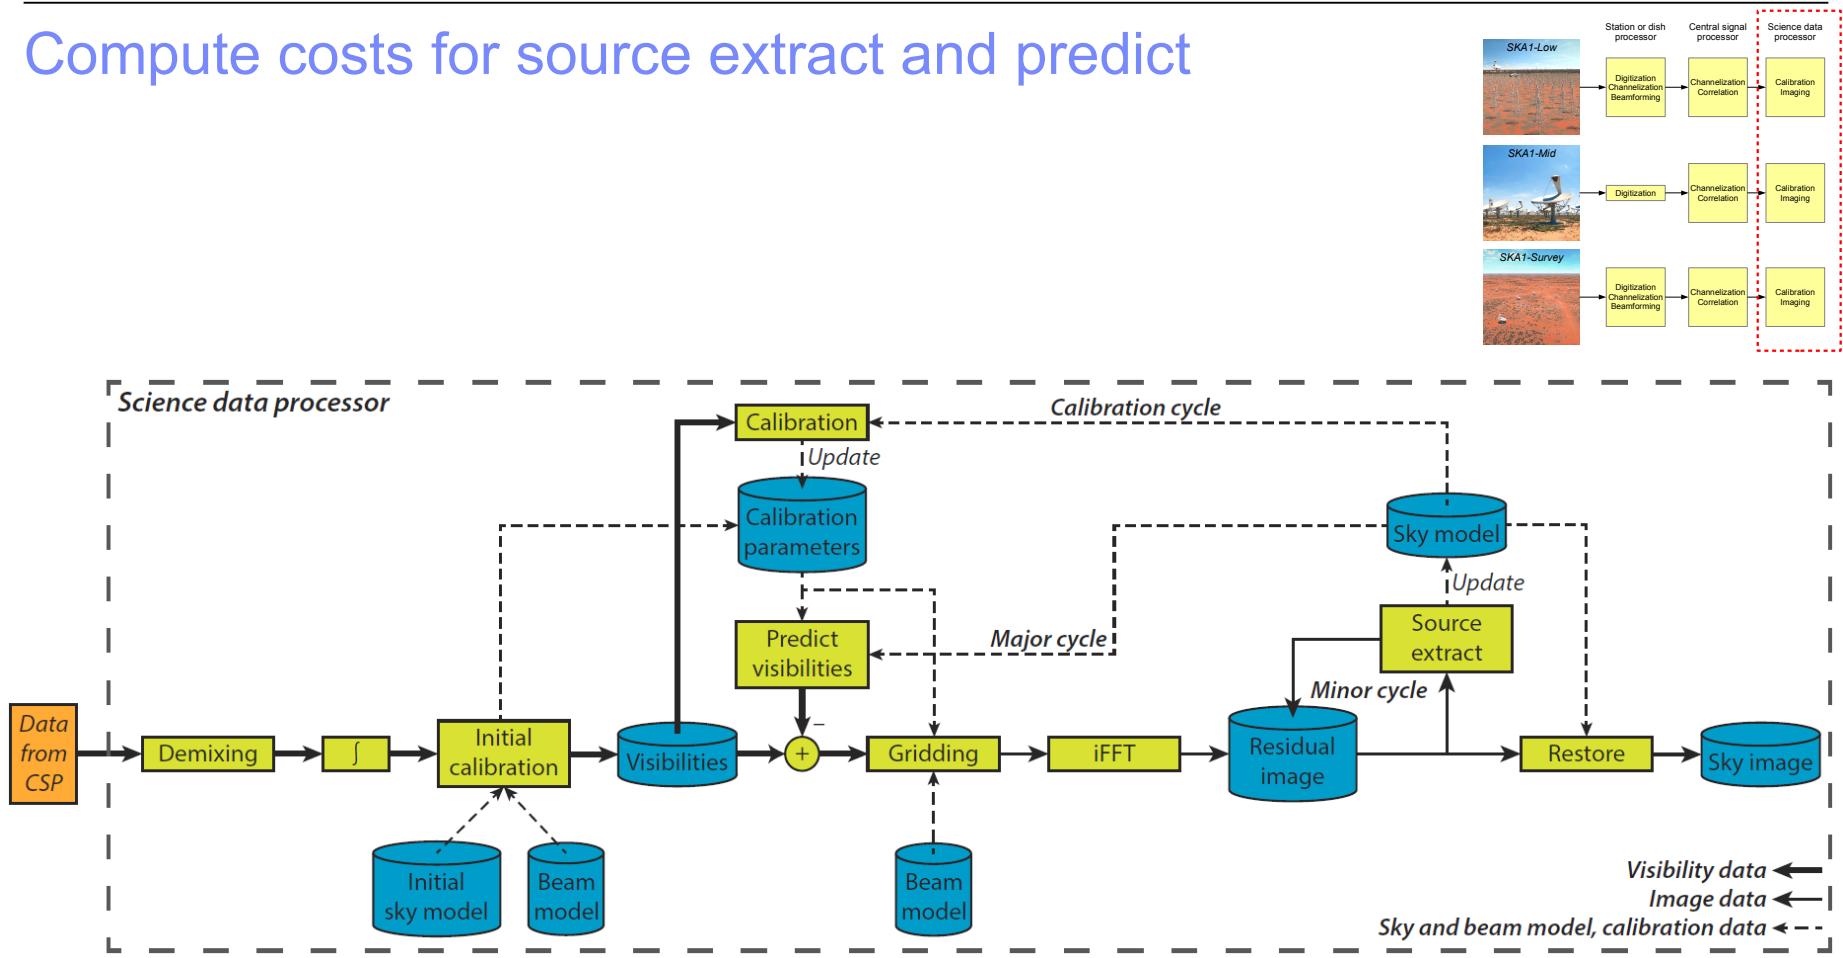
- Compute cost of last FFT

$$C_{\text{a-kernel}} = \frac{T_{\text{obs}}}{\Delta t} \sum_{N_{\text{baseline}}} \sum_{i=0}^{\lceil \frac{\Delta f_{\text{image}}}{\Delta v} \rceil - 1} 5O^2 R_{\text{ave},i}^2 \log_2 (O^2 R_{\text{ave},i}^2)$$

- Time and channel stability Δ_t, Δ_v , oversampling O
- Improve by using actual distribution in w

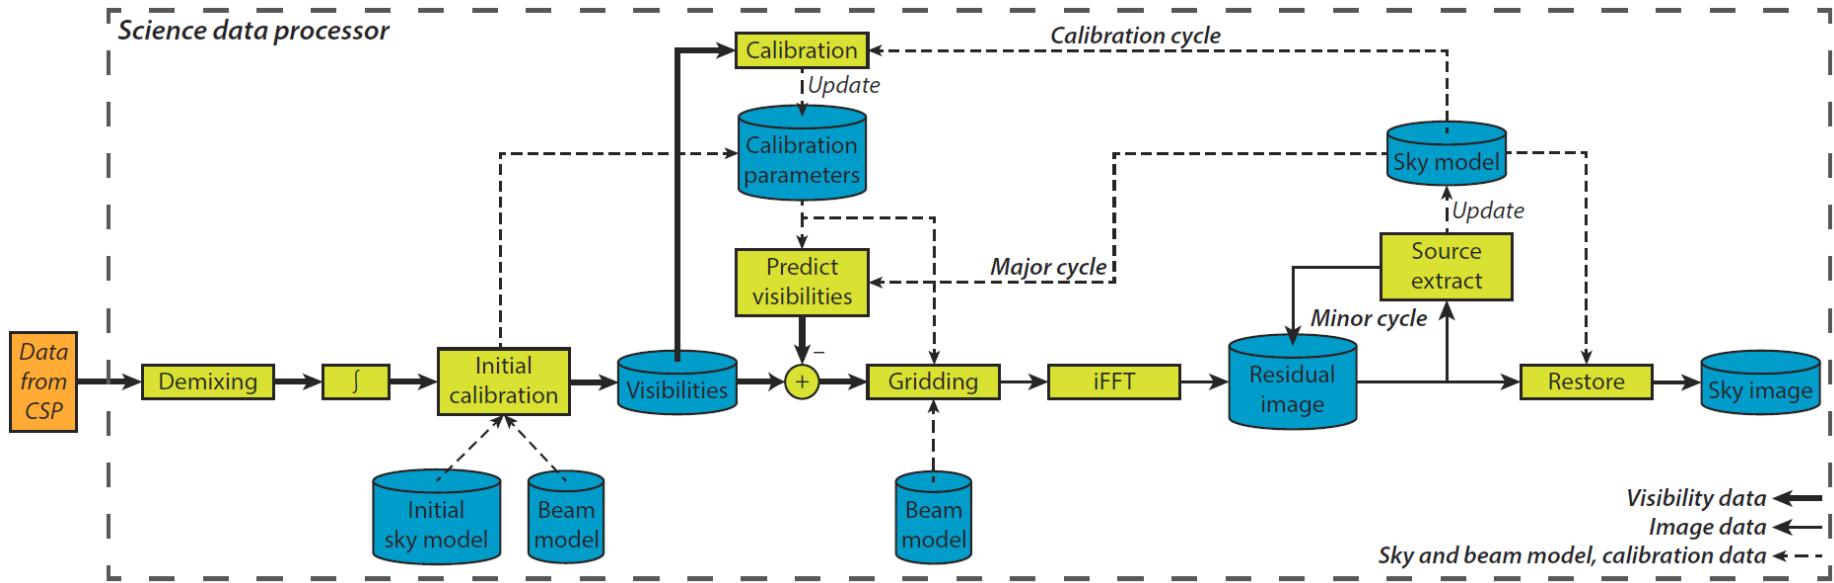
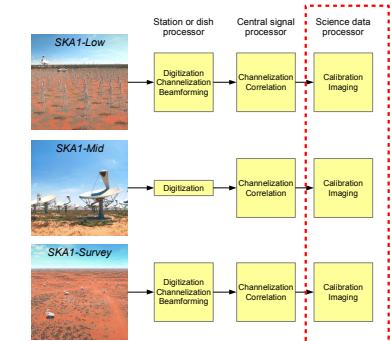


Compute costs for source extract and predict



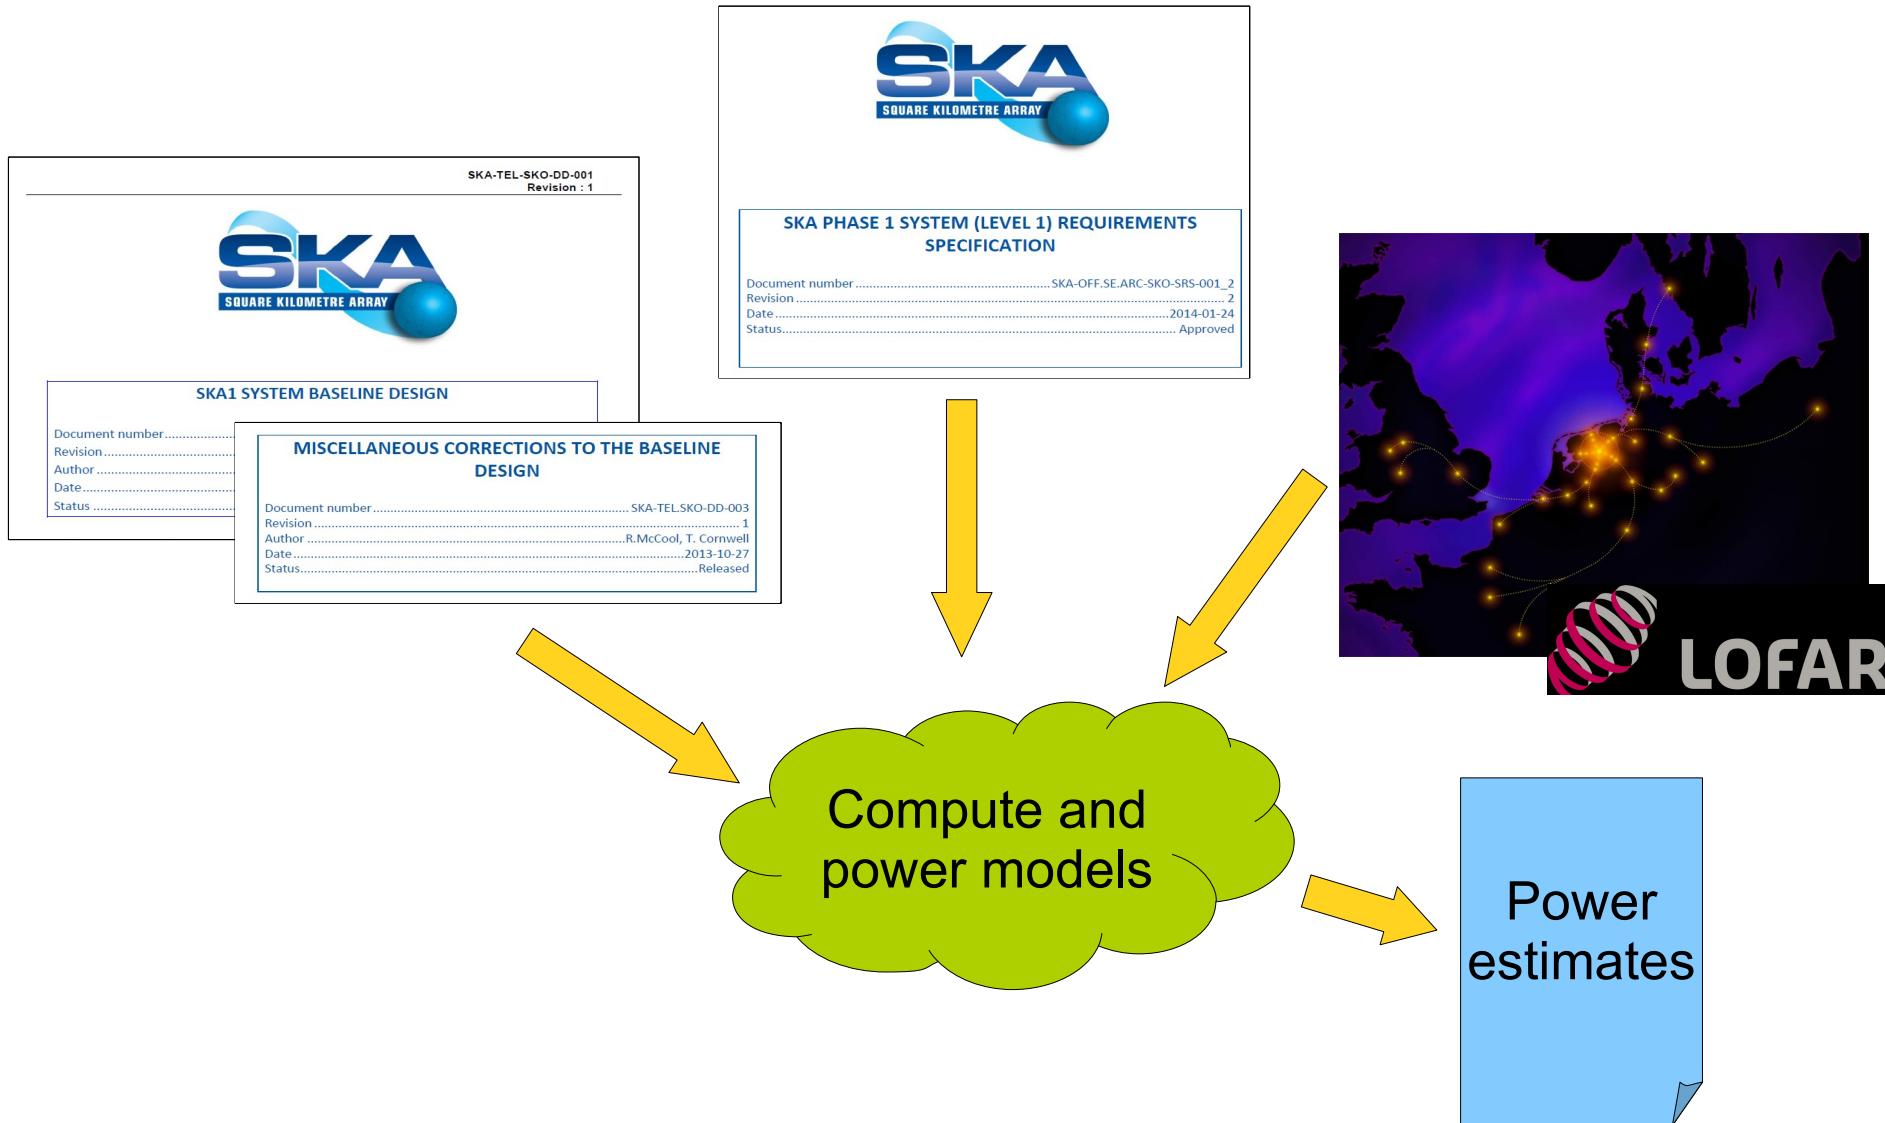
- Source extract: **subtract PSF** of image per source
 - One multiply-accumulate per pixels per source
- Prediction of visibilities with forward FFT and degridding step in major cycle
 - Similar compute requirements as forward gridding

Further model improvements



- No calibration (demixing) routines modeled
 - Compute expected to be dominated by gridding

Model parameters retrieved from various sources



Primary parameters from SKA documentation

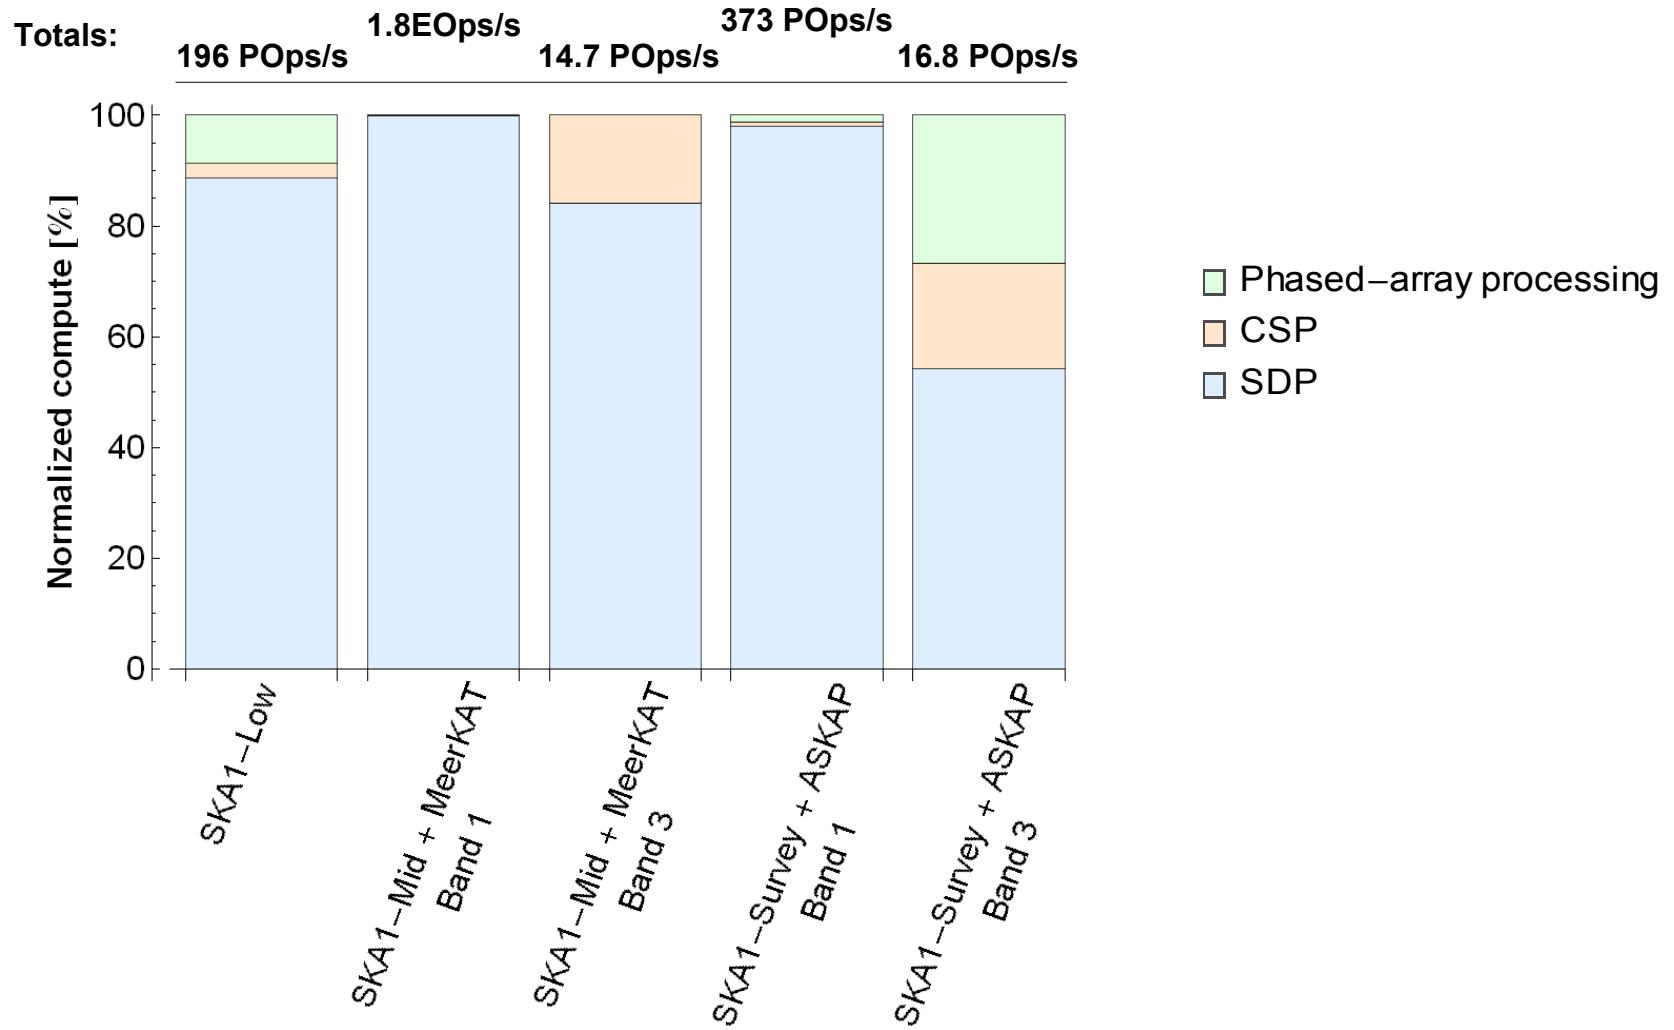


| | SKA1-Low | SKA1-Mid + MeerKAT | SKA1-Survey + ASKAP |
|------------------------------|------------------|--------------------------------------|---------------------|
| Stations or dishes | 1024 | 254 | 96 |
| Antennas | 256 | - | 94 |
| Polarizations | 2 | 2 | 2 |
| Beams | 1 | 1 | 36 |
| Frequency range | 50 MHz – 350 MHz | 350MHz – 13.8 GHz | 350MHz – 4 GHz |
| Bandwidth | 300 MHz | 1 GHz – 2.5 GHz (435 MHz – 2 GHz) | 500 MHz |
| Channels (bands) | 262144 (2048) | 262144 (-) | 262144 (2048) |
| Longest baseline (dump time) | 70 km (0.6 s) | 200 km (0.08 s) | 50 km (0.3 s) |
| Core baseline (dump time) | 6 km (6.6 s) | 9 km (1.6 s) | 10 km (1.2 s) |

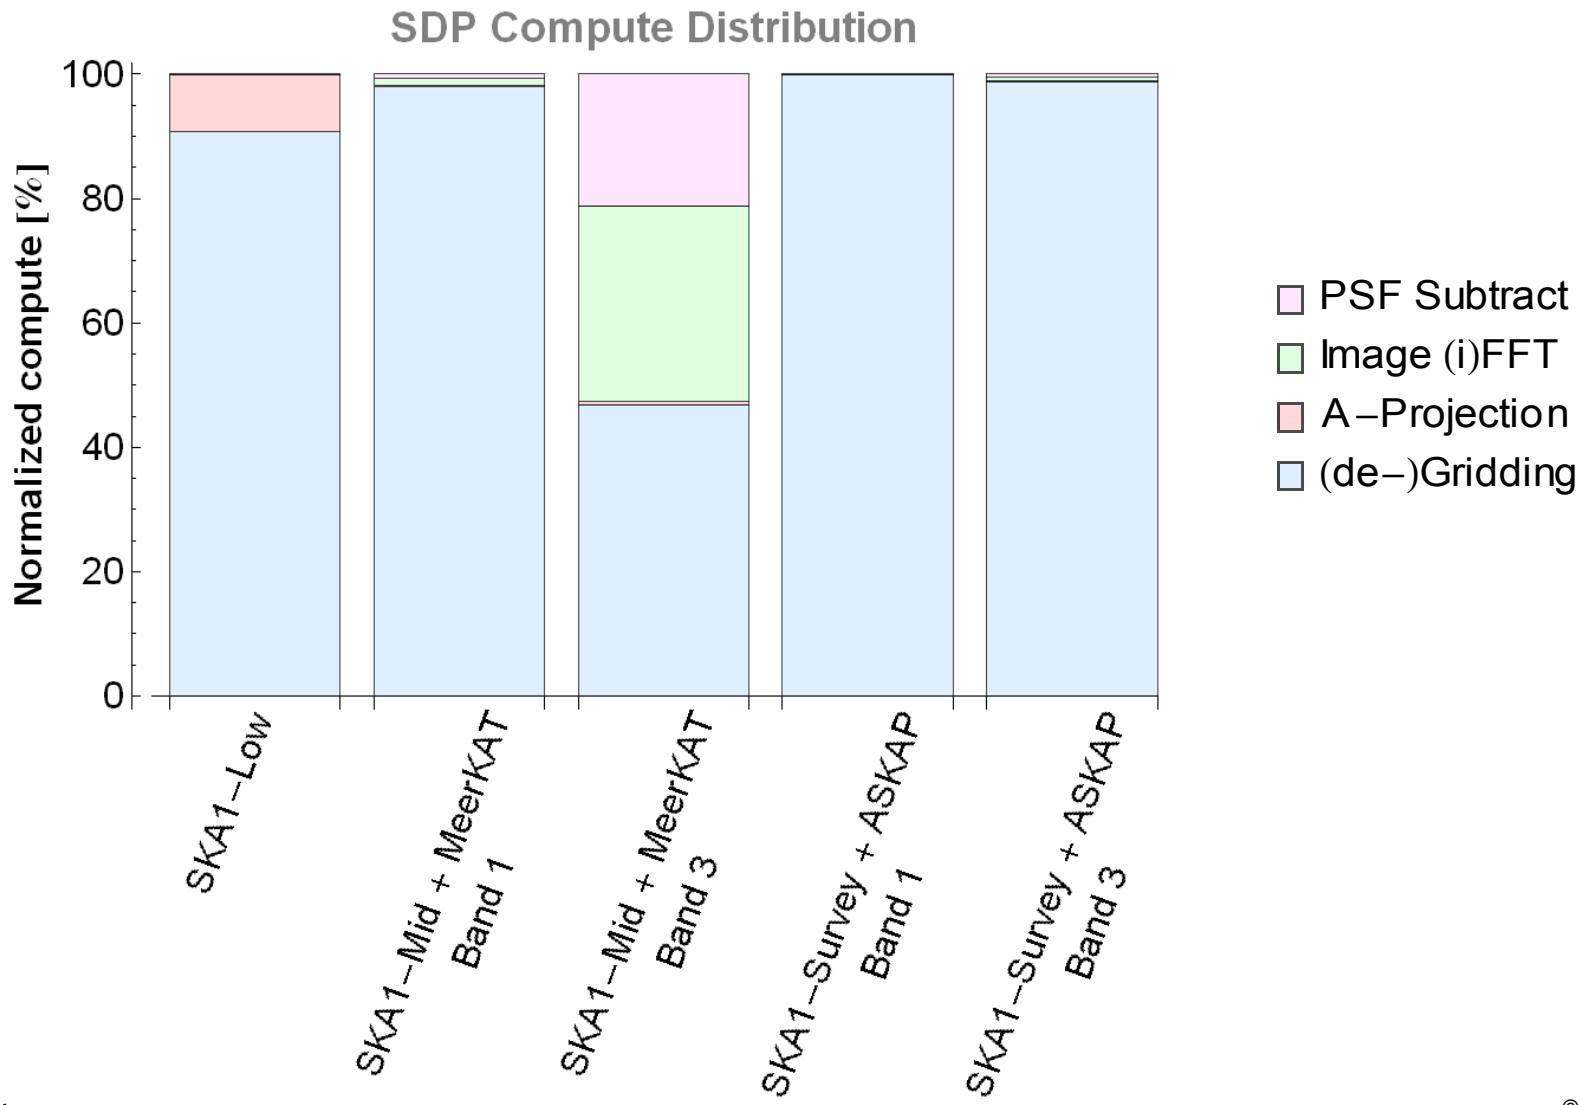
Missing parameters based on experience with LOFAR

- Observation time: 20 minutes (= length of one calibration interval)
 - 3 Calibration cycles
 - 10 Major cycles
 - 100 Minor cycles
- A-projection stability
 - 30 seconds (SKA1-Low), 300 seconds (SKA1-Mid/Survey)
 - 700 kHz

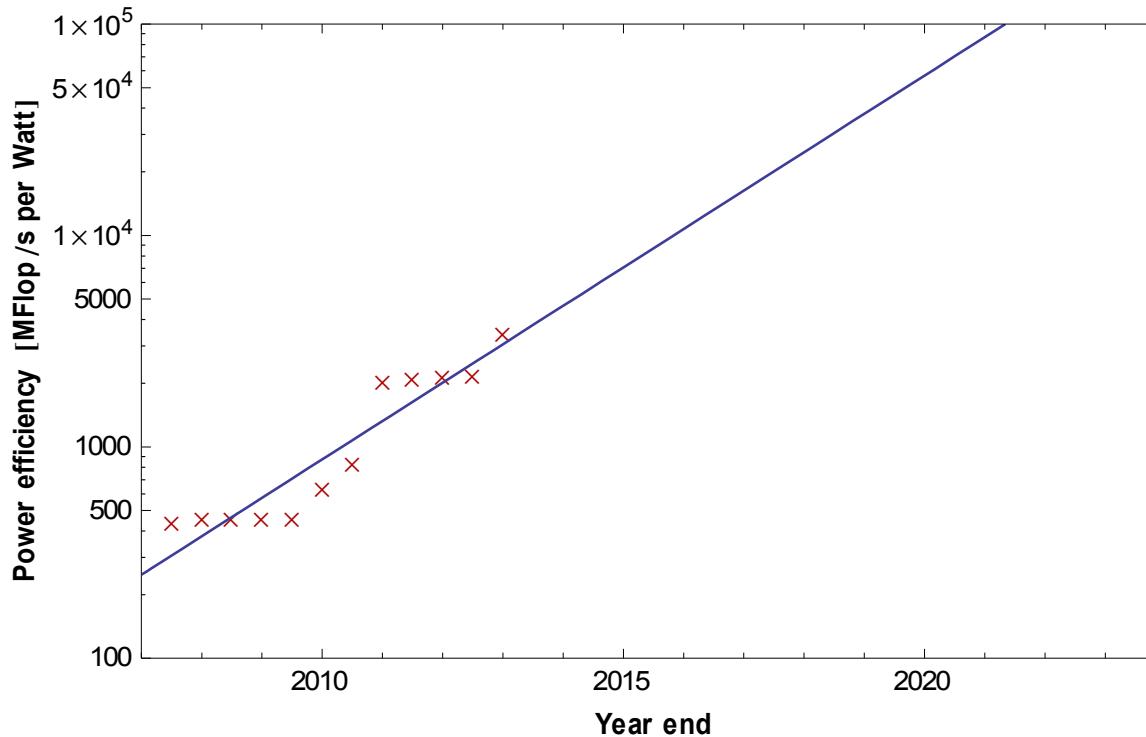
Full array and full bandwidth imaging



Full array and full bandwidth imaging

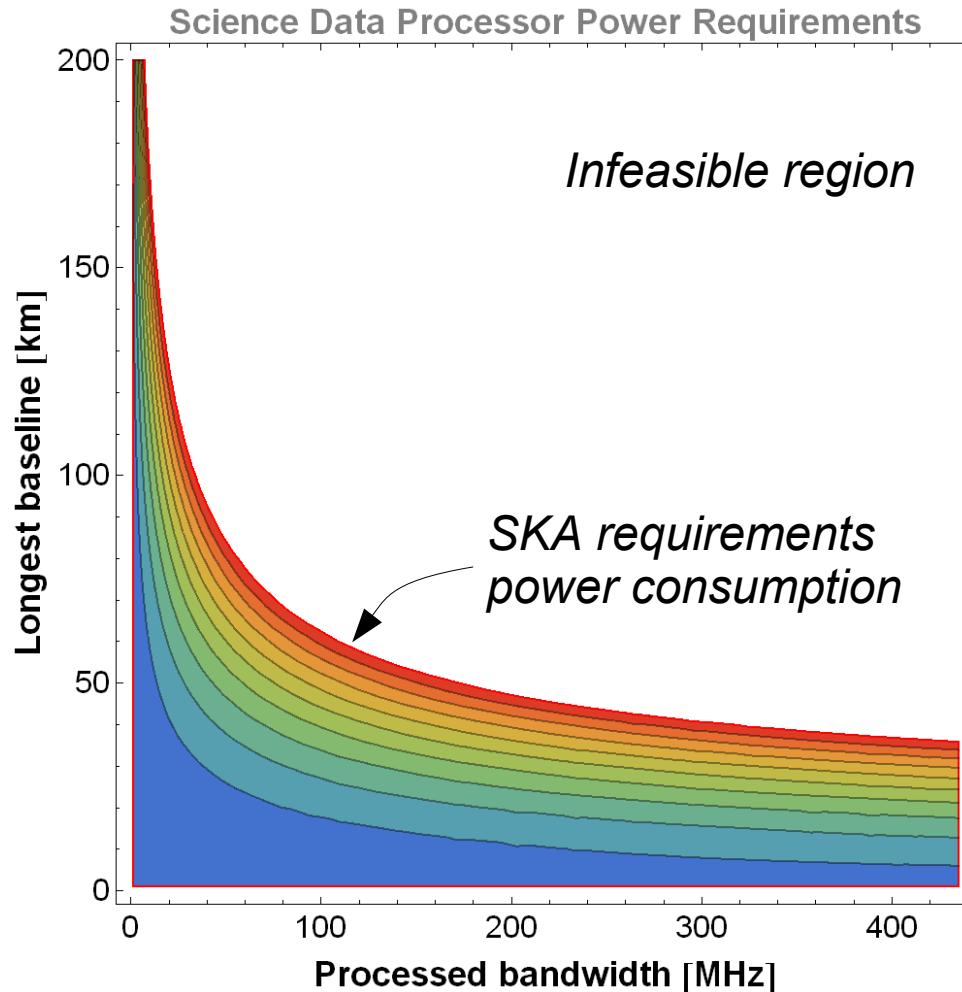


Top 500 power efficiency in 2018

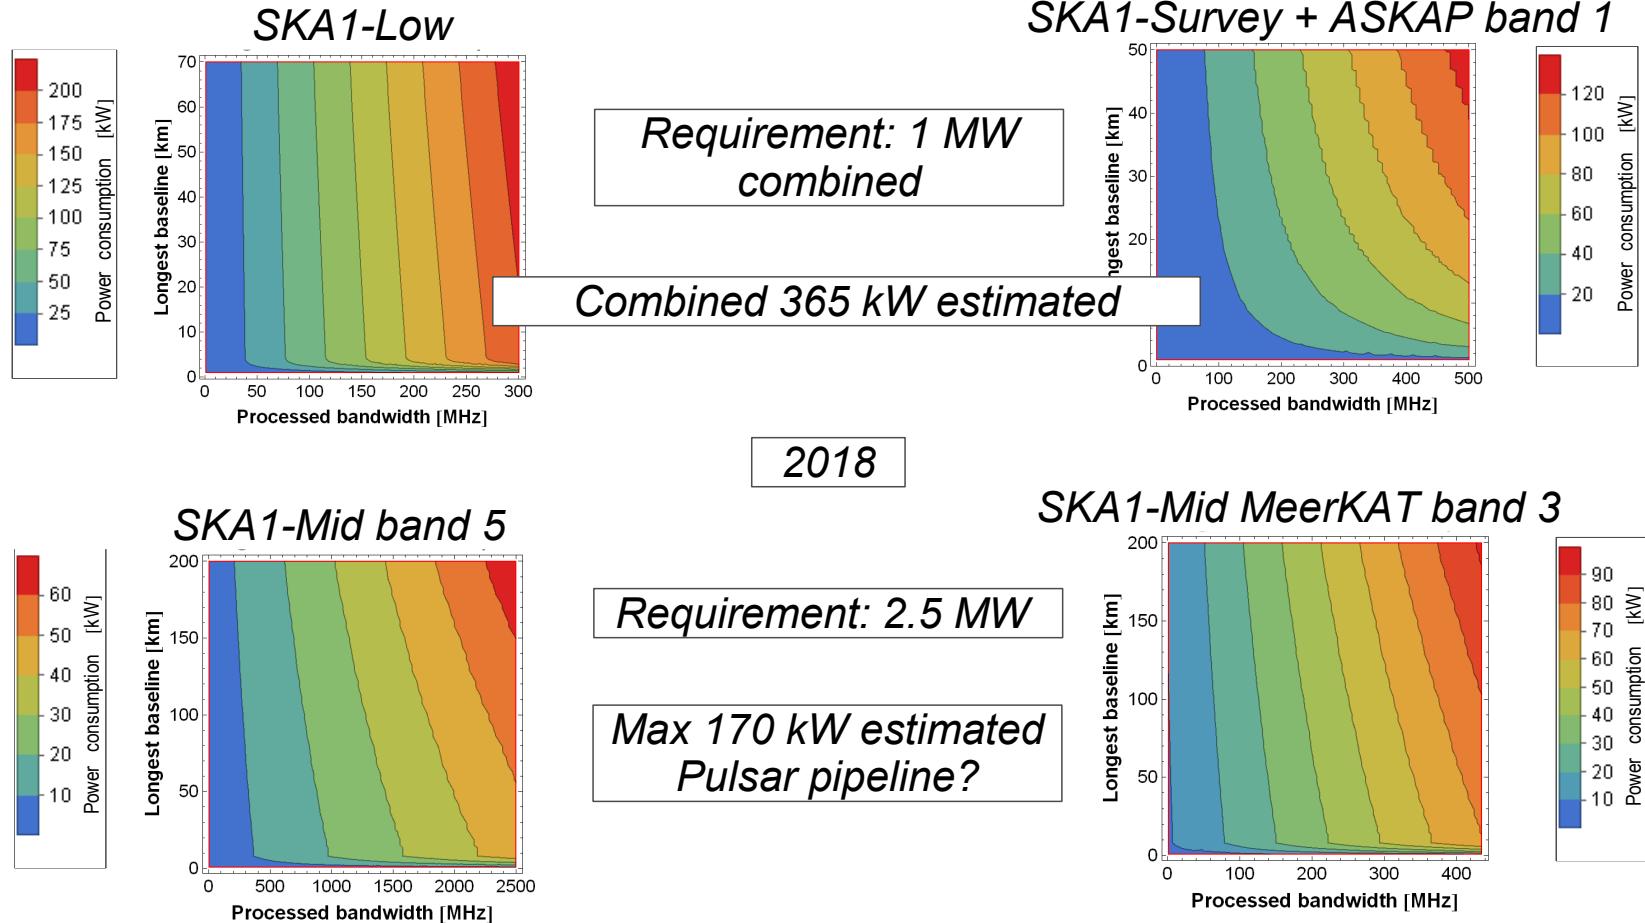


- LinPACK power efficiency
- Best machine from each Top 500 list since June 2008
- Estimate for November 2018 list: ~25 GFlop/s per Watt

Power budget constraints

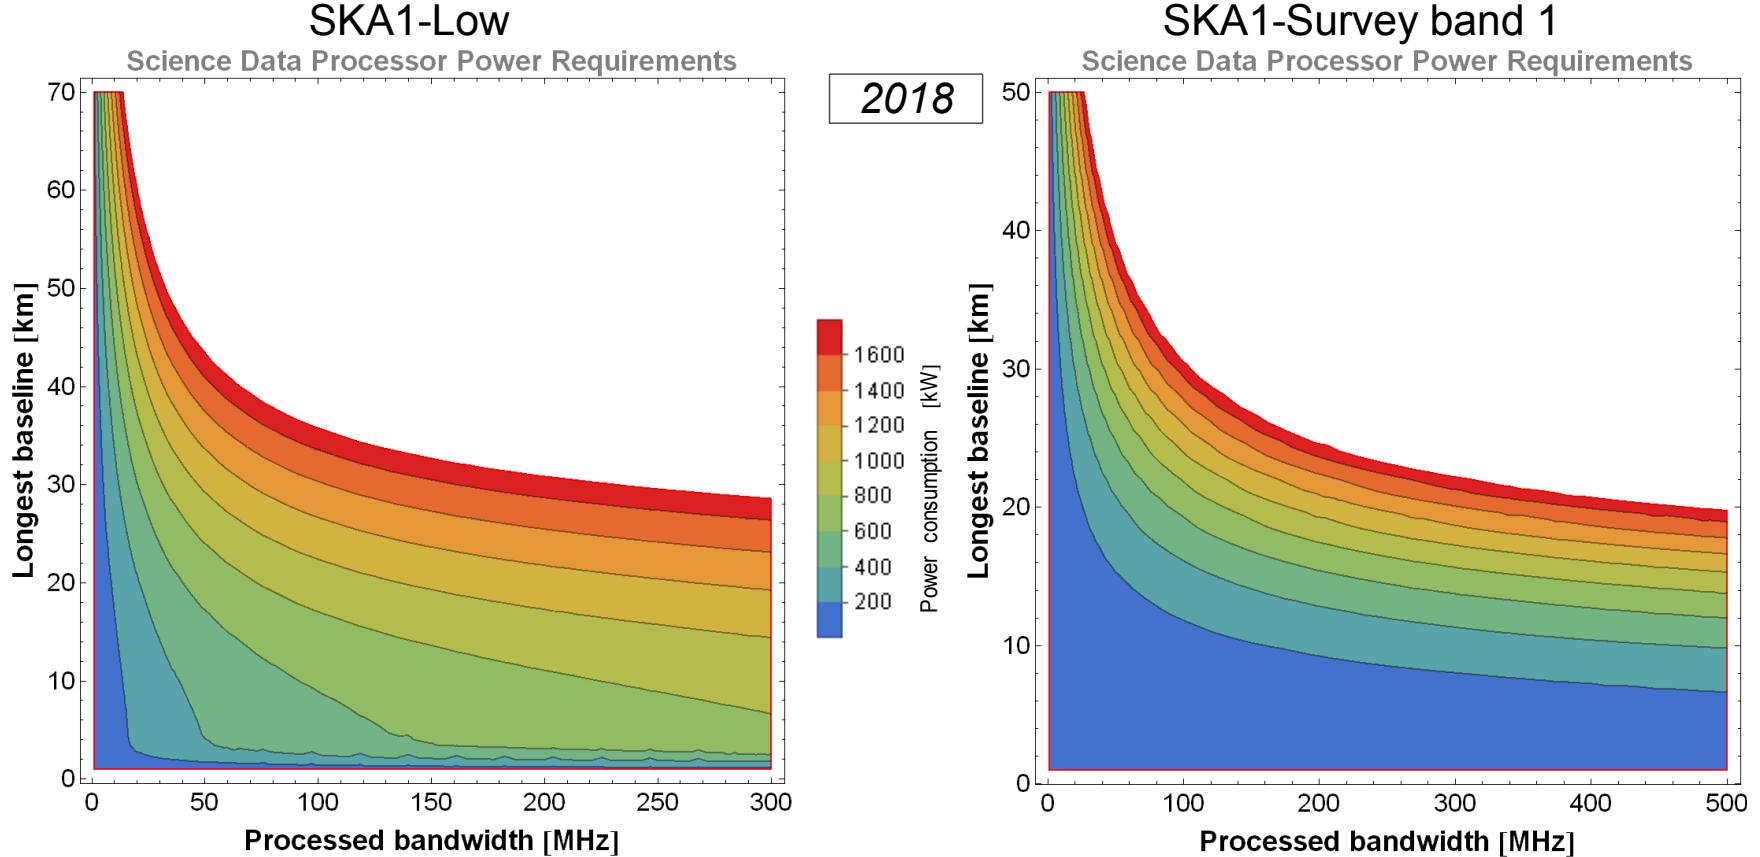


Central signal processor power requirements



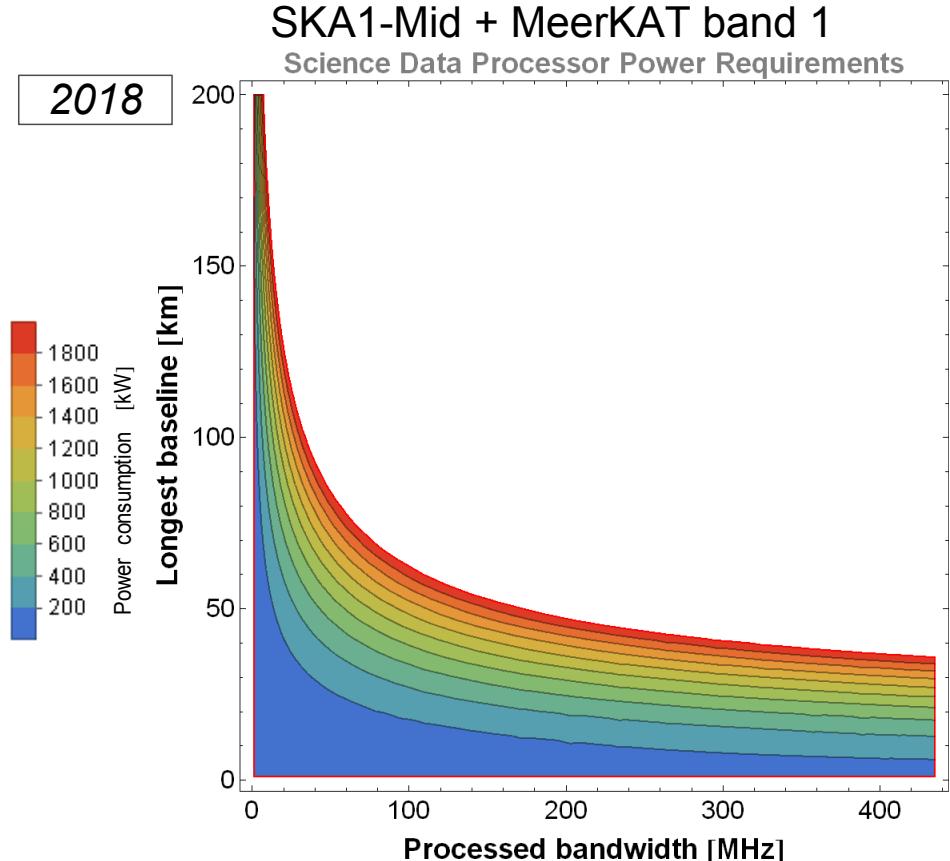
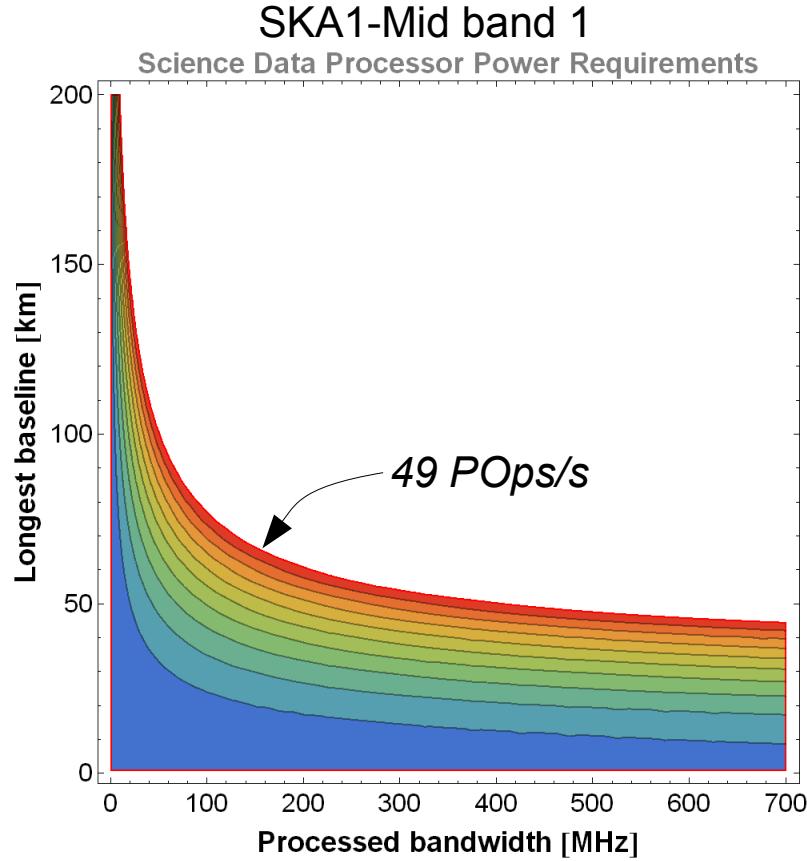
- Power equally distributed over SKA1-Low and SKA1-Survey (simultaneous observations)

SKA1-Low and SKA1-Survey SDP power budget: 3.5 MW



- Power equally distributed over SKA1-Low and SKA1-Survey (simultaneous observations)
- Baseline dependent time/frequency averaging

SKA1-Mid SDP power budget: 2 MW



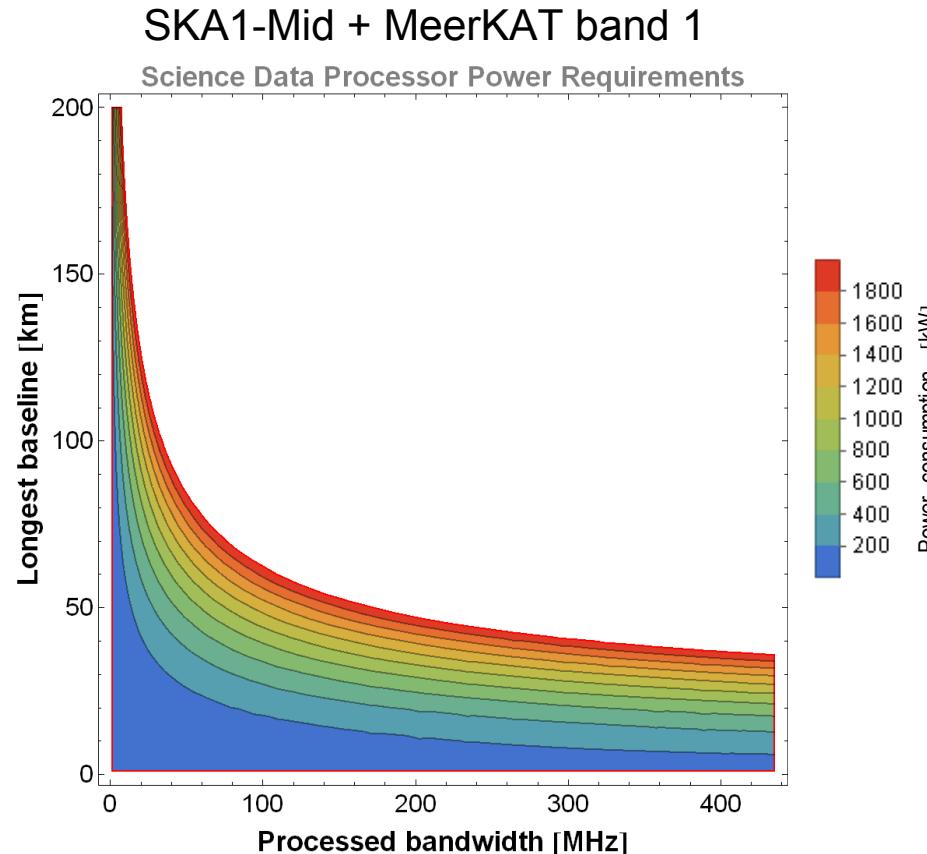
- Baseline dependent time/frequency averaging

Ample power budget for CSP

- Given 1 and 2.5 MW requirement: 635 and 170 kW needed
- A “software” correlator easily fits in the power budget
- Power budget of CSP and SDP askew for sky imaging
- Requirements for other modes? E.g. SKA1-Mid pulsar search?

Power budget constraining SDP

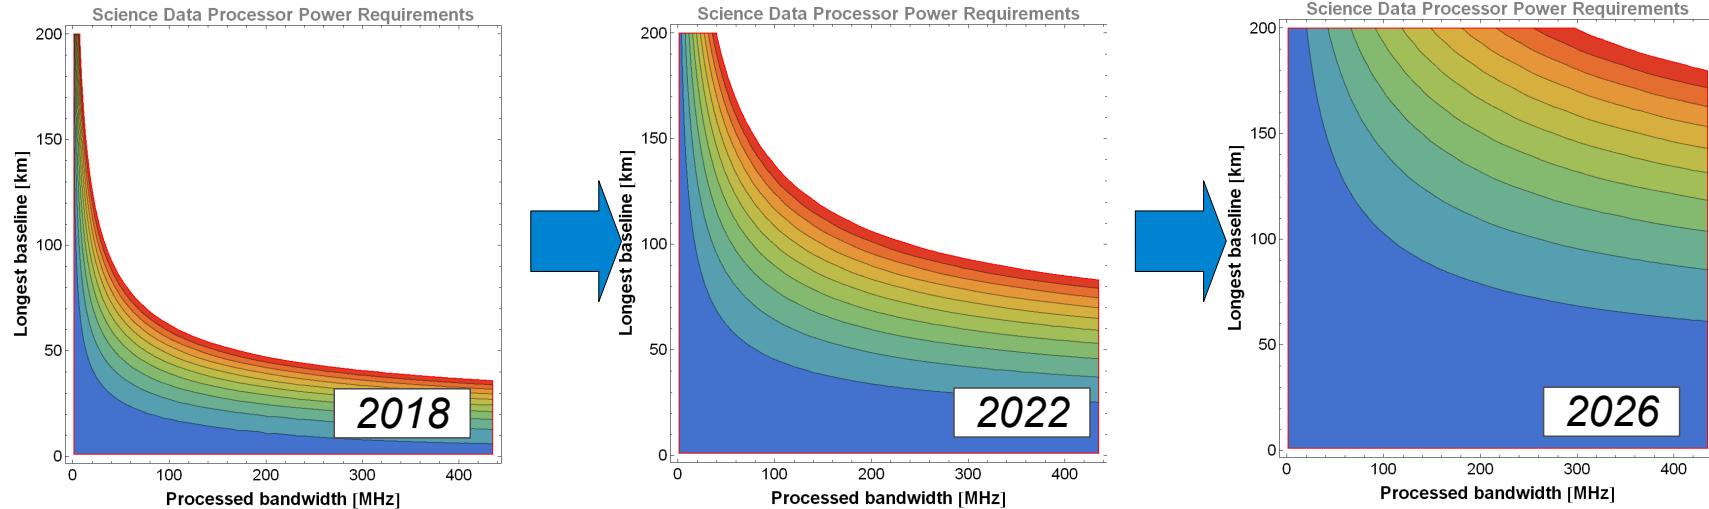
- What do the astronomers want?
 - *Match with SKA phase 1 science cases*



Increase scientific performance without modifying power budget

- Wait for technology to improve power efficiency over time?

SKA1-Mid + MeerKAT band 1

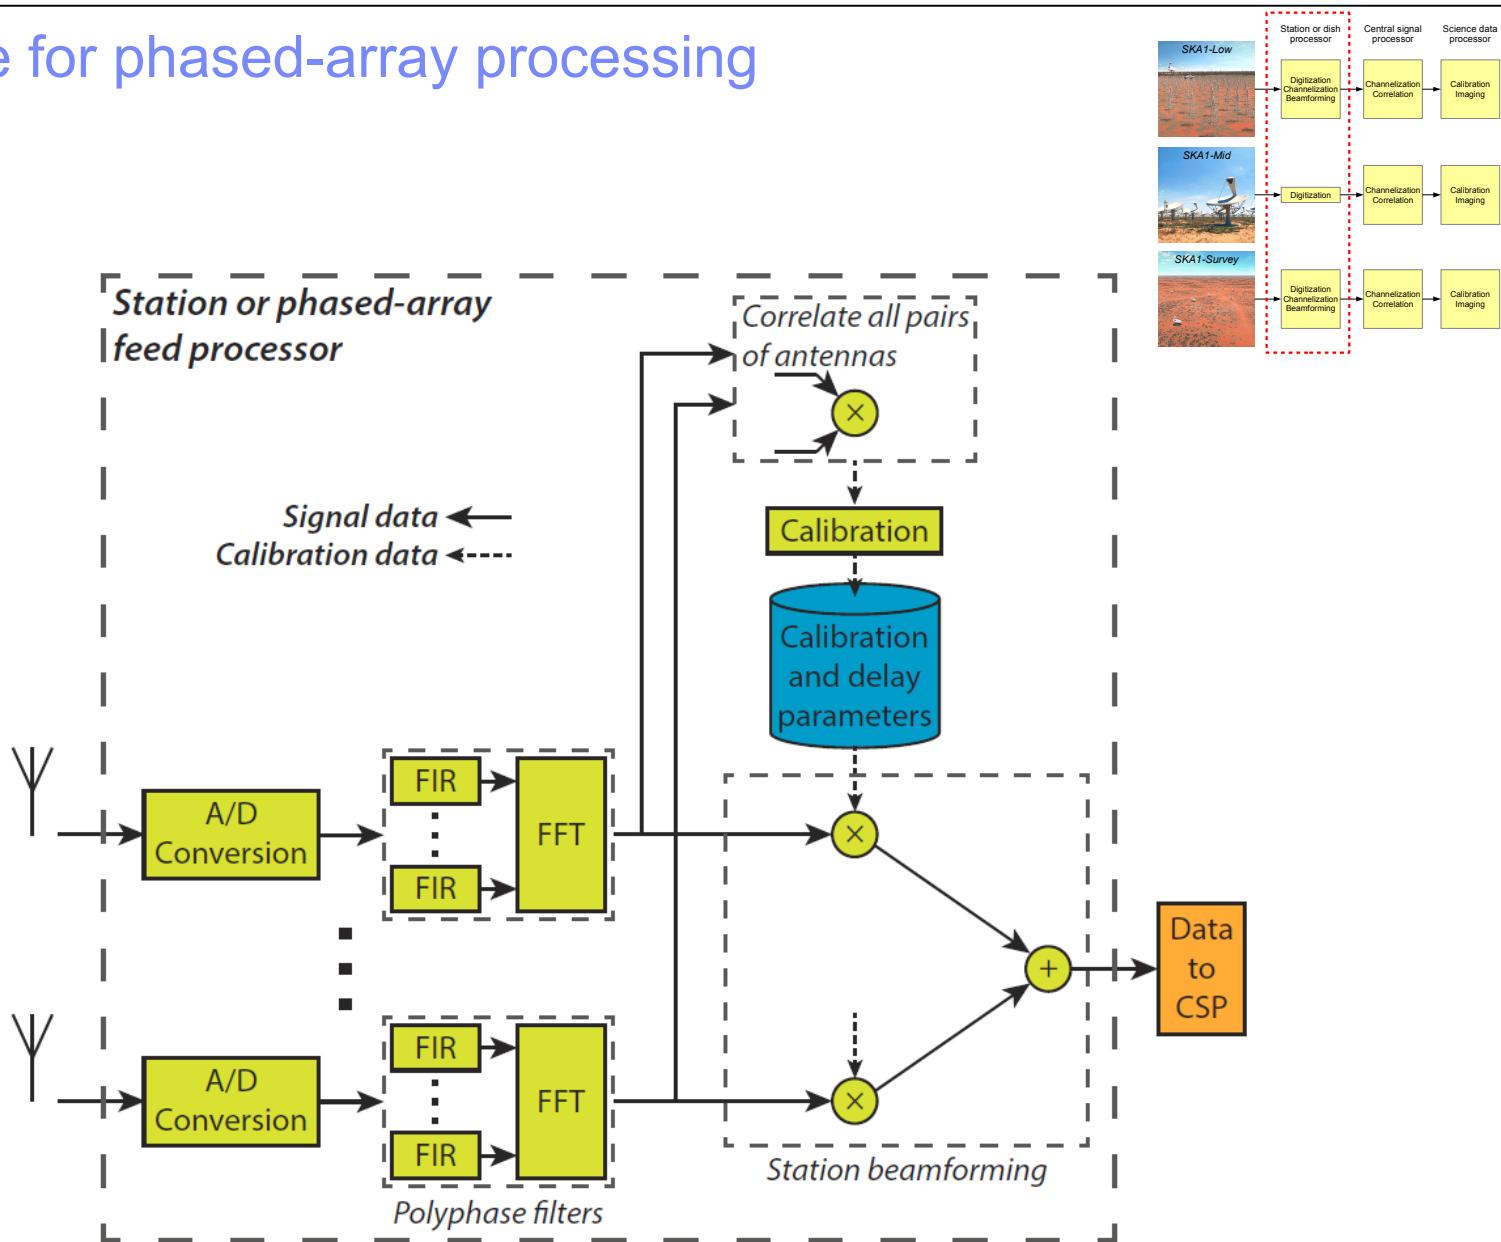


- Top 500 is based on general-purpose hardware such as CPUs and GPUs
 - Use dedicated accelerators for intensive steps such as gridding
 - FPGAs, ASICs,...

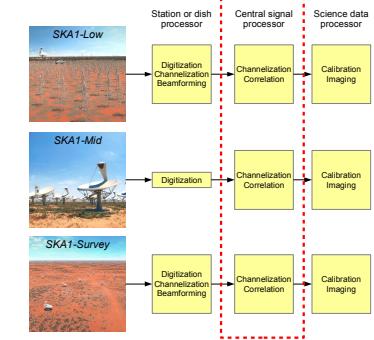
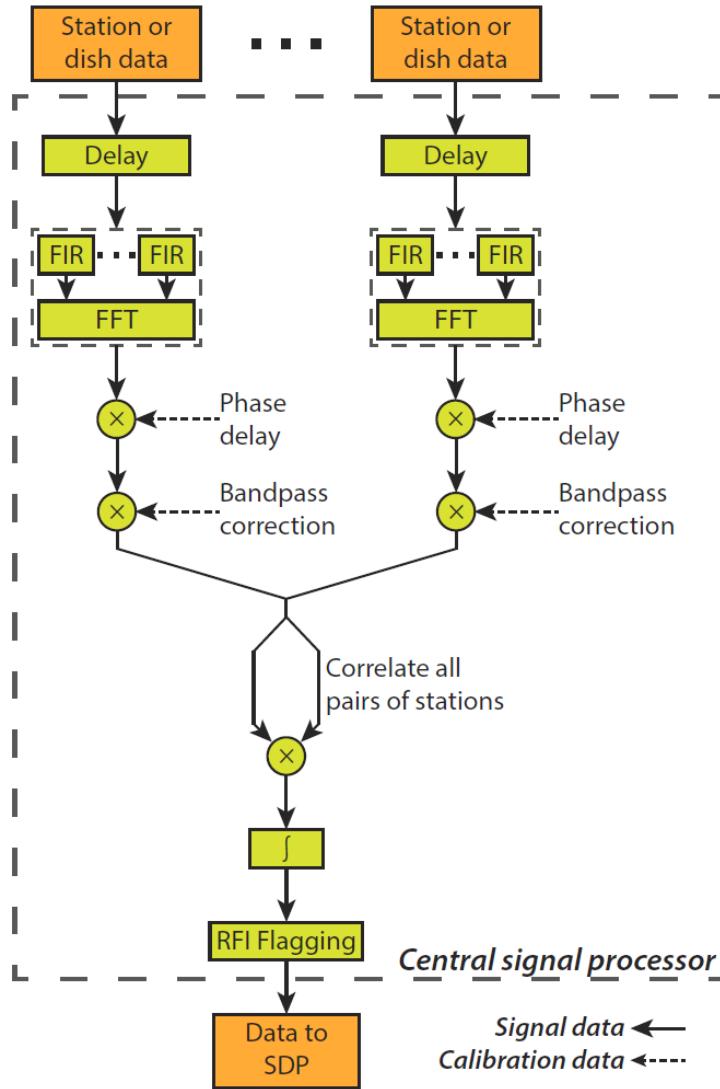
Questions

Backup slides

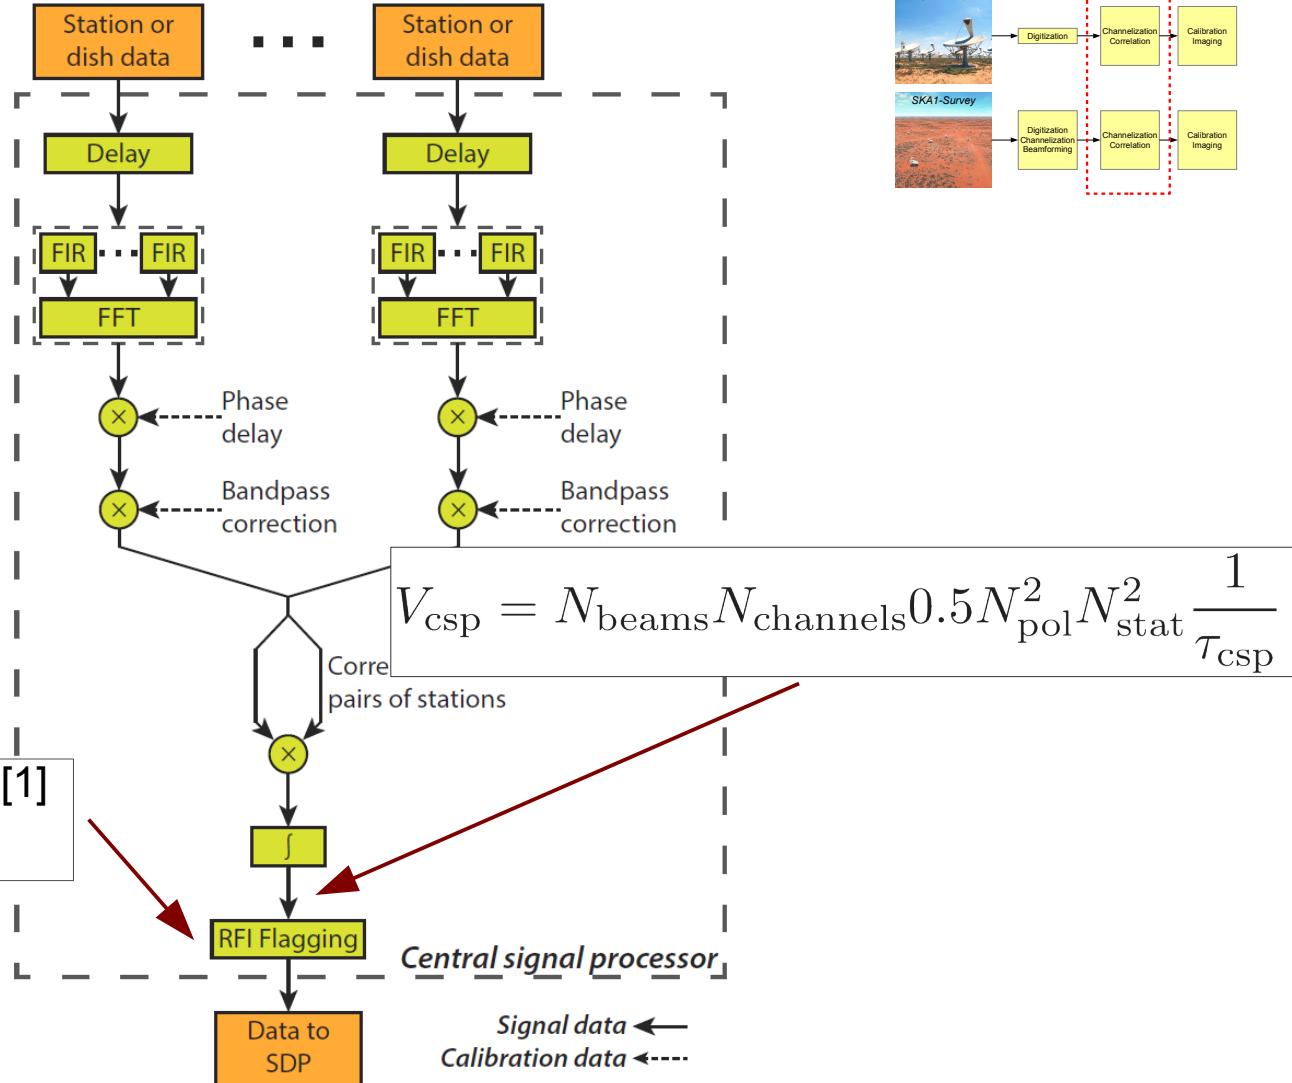
Architecture for phased-array processing



Architecture for the central signal processors



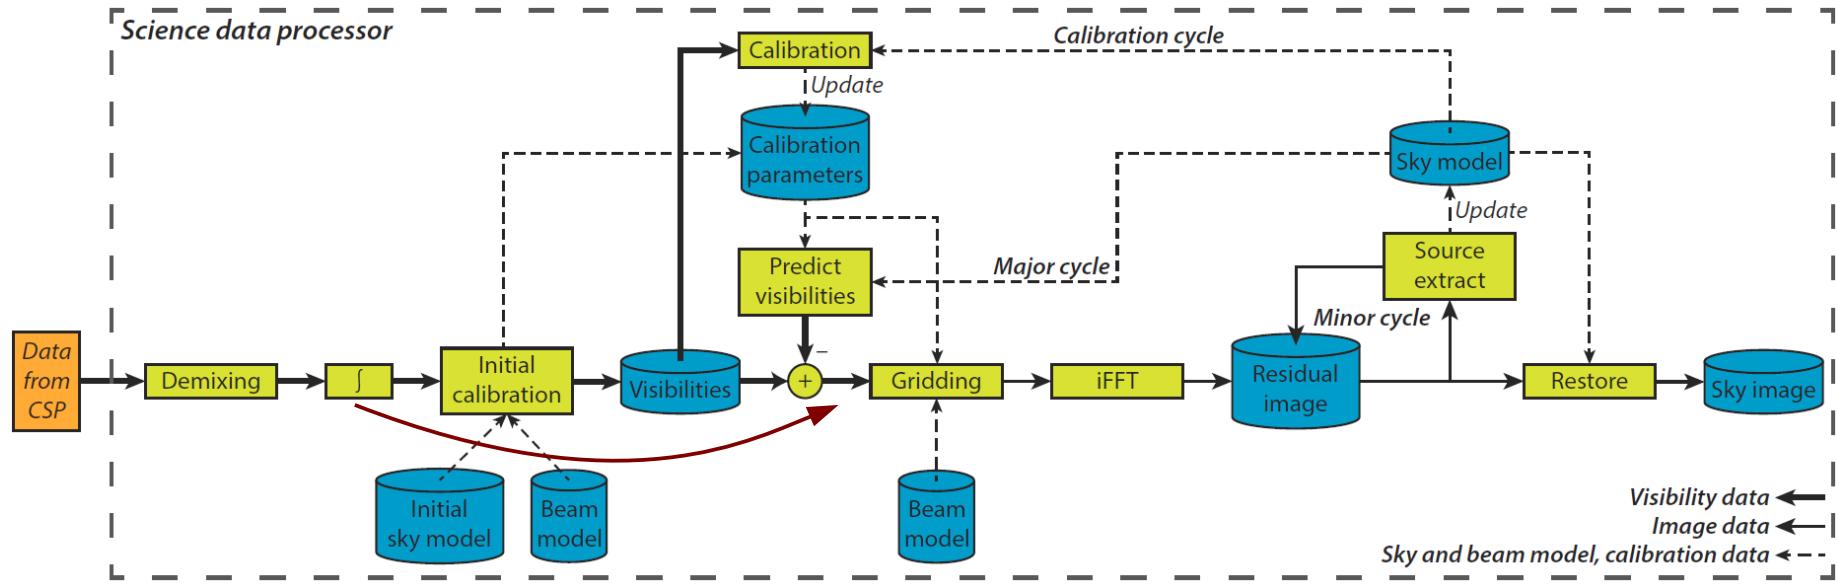
Architecture for the central signal processors



Baseline dependent averaging

| | SKA1-Low 1 zone | SKA1-Low 2 zones |
|-----|-----------------|------------------|
| SDP | 159 POps/s | 884 POps/s |

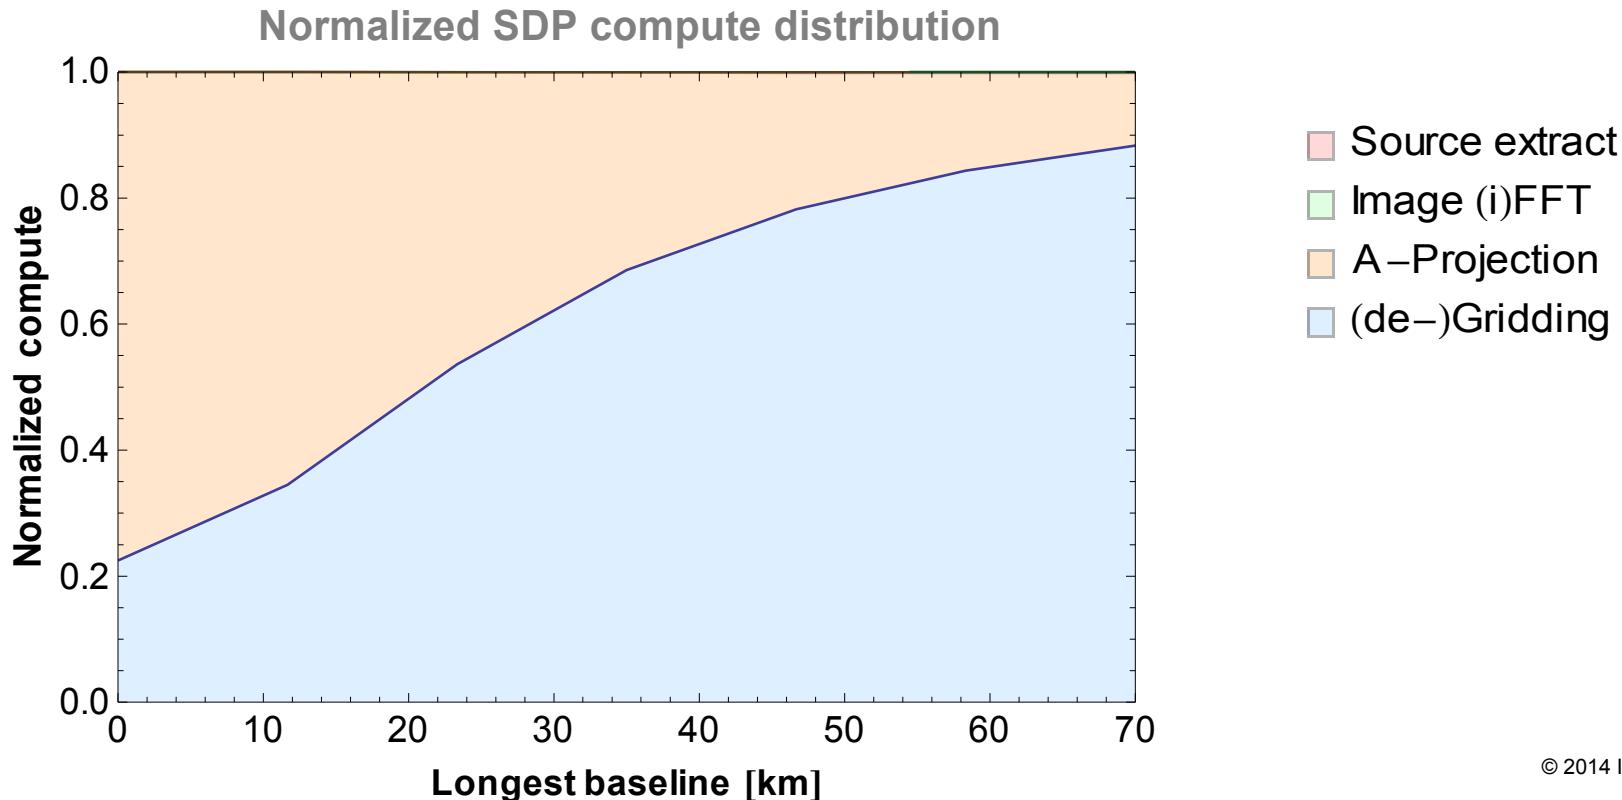
- With one dump time, gridding performs the averaging step on the image
 - Save compute by introducing before gridding
- No impact on presented models as calibration is not included



Gridding complexity vs A-projection complexity

$$\text{Gridding: } \mathcal{O} \left(\frac{\lambda^3 B_{\max}^4}{D^6} \right)$$

$$\text{A-projection: } \mathcal{O} \left(\frac{\lambda^2 B_{\max}^2}{D^4} \log_2 \frac{\lambda^2 B_{\max}^2}{D^4} \right)$$



Full array and full bandwidth imaging

| | | SKA1-Low | SKA1-Mid + MeerKAT | SKA1-Survey + ASKAP |
|-------------------------|----------------------------|------------------|--------------------|---------------------|
| Station/dish processing | One station | 17 TOps/s | - | 46 TOps/s |
| | All stations | 17 POps/s | - | 4 POps/s |
| CSP | Low band | 5 POps/s | 565 TOps/s | 3 POps/s |
| | High band | - | 2 POps/s | 3 POps/s |
| SDP | Low band (% de-/gridding) | 159 POps/s (99%) | 2 EOps/s (98 %) | 366 POps/s (99%) |
| | High band (% de-/gridding) | - | 12 POps/s (47 %) | 9 POps/s (99 %) |

SKA1-Low parameters

▼ Default parameters for instrument SKA1-Low

| Parameter Name | Value |
|---|---------------------|
| Instrument | SKA1-Low |
| Band | |
| Instrument is phased array | True |
| Antenna count | 256 |
| Polarizations count | 2 |
| Beam count | 1 |
| Signal bandwidth of feed [Hz] | 3×10^8 |
| Max signal bandwidth processed [Hz] | 3×10^8 |
| Sub-band count in stations | 2048 |
| Channel count | 262144 |
| FIR tap count in PFBs | 8 |
| Station or dish diameter [m] | 35 |
| Maximum baseline length [m] | 70000 |
| Core baseline length (zone 2) [m] | 6000 |
| CSP dump time [s] | 0.6 |
| Lowest signal frequency in band [Hz] | 50000000 |
| Phased -array correlator dump time [s] | 1 |
| Phased -array calibration table update rate [s] | 240 |
| Observation time [s] | 1200. |
| SDP minor cycle count | 100 |
| SDP major cycle count | 10 |
| SDP calibration cycle count | 3 |
| SDP power constraint [W] | 3.5×10^6 |
| CSP power constraint [W] | $1. \times 10^6$ |
| Station or dish distribution function | SKA1LowStationCount |
| A-projection kernel time stability [s] | 30 |
| A-projection kernel frequency stability [Hz] | 700000 |

▼ Derived parameters

| Parameter Name | Value |
|-----------------------------------|---------|
| Channel bandwidth [Hz] | 1144.41 |
| Station or dish count | 1024 |
| Sub-band bandwidth [Hz] | 146484. |
| SDP integrated channel width [Hz] | 2288.82 |
| SDP integrated channel count | 131072 |
| SDP image channel count | 131072 |
| SDP all baselines dump time [s] | 0.6 |
| SDP core baselines dump time [s] | 6.6 |

SKA1-Mid parameters (band 1)

▼ Default parameters for instrument SKA1-Mid

| Parameter | Name | Value |
|--|------|-------------------|
| Instrument | | SKA1-Mid |
| Band | | Band 1 |
| Instrument is phased array | | False |
| Polarizations count | | 2 |
| Beam count | | 1 |
| Channel count | | 262144 |
| FIR tap count in PFBs | | 8 |
| Station or dish diameter [m] | | 15 |
| Maximum baseline length [m] | | 200000 |
| Core baseline length (zone 2) [m] | | 8000 |
| CSP dump time [s] | | 0.08 |
| Observation time [s] | | 1200. |
| SDP minor cycle count | | 100 |
| SDP major cycle count | | 10 |
| SDP calibration cycle count | | 3 |
| SDP power constraint [W] | | 2000000 |
| CSP power constraint [W] | | 2.5×10^6 |
| Station or dish distribution function | | SKA1MidDishCount |
| A-projection kernel time stability [s] | | 300 |
| A-projection kernel frequency stability [Hz] | | 700000 |
| Lowest signal frequency in band [Hz] | | 350000000 |
| Signal bandwidth of feed [Hz] | | $1. \times 10^9$ |
| Max signal bandwidth processed [Hz] | | $7. \times 10^8$ |

▼ Derived parameters

| Parameter | Name | Value |
|-----------------------------------|------|--------|
| Channel bandwidth [Hz] | | 3814.7 |
| Station or dish count | | 190 |
| SDP integrated channel width [Hz] | | 3814.7 |
| SDP integrated channel count | | 262144 |
| SDP image channel count | | 262144 |
| SDP all baselines dump time [s] | | 0.08 |
| SDP core baselines dump time [s] | | 2.24 |

SKA1-Mid + MeerKAT parameters (band 1)

▼ Default parameters for instrument SKA1-Mid + MeerKAT

| Parameter | Name | Value |
|--|------|-------------------------|
| Instrument | | SKA1-Mid + MeerKAT |
| Band | | Band 1 |
| Instrument is phased array | | False |
| Polarizations count | | 2 |
| Beam count | | 1 |
| Channel count | | 262144 |
| FIR tap count in PFBs | | 8 |
| Station or dish diameter [m] | | 12 |
| Maximum baseline length [m] | | 200000 |
| Core baseline length (zone 2) [m] | | 9000 |
| CSP dump time [s] | | 0.08 |
| Observation time [s] | | 1200. |
| SDP minor cycle count | | 100 |
| SDP major cycle count | | 10 |
| SDP calibration cycle count | | 3 |
| SDP power constraint [W] | | 2000000 |
| CSP power constraint [W] | | 2.5×10^6 |
| Station or dish distribution function | | SKA1MidMeerKATDishCount |
| A-projection kernel time stability [s] | | 300 |
| A-projection kernel frequency stability [Hz] | | 700000 |
| Lowest signal frequency in band [Hz] | | 580000000 |
| Signal bandwidth of feed [Hz] | | $1. \times 10^9$ |
| Max signal bandwidth processed [Hz] | | 4.35×10^8 |

▼ Derived parameters

| Parameter | Name | Value |
|-----------------------------------|------|--------|
| Channel bandwidth [Hz] | | 3814.7 |
| Station or dish count | | 254 |
| SDP integrated channel width [Hz] | | 3814.7 |
| SDP integrated channel count | | 262144 |
| SDP image channel count | | 262144 |
| SDP all baselines dump time [s] | | 0.08 |
| SDP core baselines dump time [s] | | 1.6 |

SKA1-Survey + ASKAP parameters (band 1)

▼ Default parameters for instrument SKA1-Survey + ASKAP

| Parameter Name | Value |
|---|--------------------------|
| Instrument | SKA1-Survey + ASKAP |
| Band | Band 1 |
| Instrument is phased array | True |
| Antenna count | 94 |
| Polarizations count | 2 |
| Beam count | 36 |
| Sub-band count in stations | 2048 |
| Channel count | 262144 |
| FIR tap count in PFBs | 8 |
| Station or dish diameter [m] | 12 |
| Maximum baseline length [m] | 50000 |
| Core baseline length (zone 2) [m] | 10000 |
| CSP dump time [s] | 0.3 |
| Phased -array correlator dump time [s] | 1 |
| Phased -array calibration table update rate [s] | 240 |
| Observation time [s] | 1200. |
| SDP minor cycle count | 100 |
| SDP major cycle count | 10 |
| SDP calibration cycle count | 3 |
| SDP power constraint [W] | 3.5×10^6 |
| CSP power constraint [W] | $1. \times 10^6$ |
| Station or dish distribution function | SKA1SurveyASKAPDishCount |
| A-projection kernel time stability [s] | 300 |
| A-projection kernel frequency stability [Hz] | 700000 |
| Lowest signal frequency in band [Hz] | 350000000 |
| Signal bandwidth of feed [Hz] | $5. \times 10^8$ |
| Max signal bandwidth processed [Hz] | $5. \times 10^8$ |

▼ Derived parameters

| Parameter Name | Value |
|-----------------------------------|---------|
| Channel bandwidth [Hz] | 1907.35 |
| Station or dish count | 97 |
| Sub-band bandwidth [Hz] | 244141. |
| SDP integrated channel width [Hz] | 7629.39 |
| SDP integrated channel count | 65536 |
| SDP image channel count | 65536 |
| SDP all baselines dump time [s] | 0.3 |
| SDP core baselines dump time [s] | 1.2 |

SKA1-Low details

SKA1-Low

Station:

PPF: 14.1312 TOps/s
 Beamforming: 1.2288 TOps/s
 Correlation: 1.31072 TOps/s
 Calibration: 1.78957 GOPs/s

Total compute per station: 16.6725 TOps/s

Total compute for all 1024 stations: 17.0726 POps/s

CSP:

PPF: 60.8256 TOps/s
 Correlation: 5.03316 POps/s
 RFI flagging: 254.72 TOps/s

Total compute: 5.34871 POps/s

Estimated power consumption: 216.226 kWatt (L1 Requirement: 1. MWatt)

SDP:

A-Projection: 8.37162 POps/s
 Gridding: 83.3551 POps/s
 iFFT: 59.2219 TOps/s
 Deconvolution: 94.3718 TOps/s
 FFT: 53.2997 TOps/s
 A-Projection: 7.53446 POps/s
 Degridding: 75.0196 POps/s

Total compute: 174.488 POps/s

Estimated power consumption: 7.05382 MWatt (L1 Requirement: 3.5 MWatt)

▼ Additional data

Gridding:

Minimum Rf: 342.857
 Maximum Rf: 48.9799
 Average Rf: 111.196
 wrms/wmax (estimate): 0.3128
 Ra (estimate): 7
 Corrected minimum w-term: 107.246
 Corrected maximum w-term: 15.3209
 Corrected average w-term: 34.7821

Imaging:

Pixels: 6000

SKA1-Mid details (band 1)

SKA1-Mid Band 1

CSP:

PPF: 48.26 TOps/s
 Correlation: 404.32 TOps/s
 RFI flagging: 46.0395 TOps/s

Total compute: 498.62 TOps/s

Estimated power consumption: 20.1571 kWatt (L1 Requirement: 2.5 MWatt)

SDP:

A-Projection: 897.71 TOps/s
 Gridding: 660.055 POps/s
 iFFT: 6.41213 POps/s
 Deconvolution: 8.38861 POps/s
 FFT: 5.77092 POps/s
 A-Projection: 807.939 TOps/s
 Degridding: 594.049 POps/s

Total compute: 1.27638 EOps/s

Estimated power consumption: 51.5988 MWatt (L1 Requirement: 2. MWatt)

▼ Additional data

Gridding:

Minimum Rf: 761.905
 Maximum Rf: 253.969
 Average Rf: 418.52
 wrms/wmax (estimate): 0.3128
 Ra (estimate): 7
 Corrected minimum w-term: 238.324
 Corrected maximum w-term: 79.4415
 Corrected average w-term: 130.913

Imaging:

Pixels: 40000

SKA1-Mid + MeerKAT details (band 1)

SKA1-Mid + MeerKAT Band 1

CSP:

PPF: 64.516 TOps/s
 Correlation: 449.033 TOps/s
 RFI flagging: 51.1308 TOps/s

Total compute: 564.68 TOps/s

Estimated power consumption: 22.8277 kWatt (L1 Requirement: 2.5 MWatt)

SDP:

A-Projection: 1.54976 POps/s
 Gridding: 916.417 POps/s
 iFFT: 10.2299 POps/s
 Deconvolution: 13.1072 POps/s
 FFT: 9.20694 POps/s
 A-Projection: 1.39478 POps/s
 Degridding: 824.776 POps/s

Total compute: 1.77668 EOps/s

Estimated power consumption: 71.8239 MWatt (L1 Requirement: 2. MWatt)

▼ Additional data

Gridding:

Minimum Rf: 718.391
 Maximum Rf: 410.51
 Average Rf: 536.031
 wrms/wmax (estimate): 0.3128
 Ra (estimate): 7
 Corrected minimum w-term: 224.713
 Corrected maximum w-term: 128.408
 Corrected average w-term: 167.671

Imaging:

Pixels: 50000

SKA1-Survey + ASKAP details (band 1)

SKA1-Survey + ASKAP Band 1

Station:

PPF: 8.648 TOps/s
 Beamforming: 27.072 TOps/s
 Correlation: 10.6032 TOps/s
 Calibration: 241.282 MOps/s

Total compute per station: 46.3234 TOps/s

Total compute for all 97 stations: 4.49337 POps/s

CSP:

PPF: 345.708 TOps/s
 Correlation: 2.70979 POps/s
 RFI flagging: 164.566 TOps/s

Total compute: 3.22007 POps/s

Estimated power consumption: 130.174 kWatt (L1 Requirement: 1. MWatt)

SDP:

A-Projection: 27.6812 TOps/s
 Gridding: 192.154 POps/s
 iFFT: 139.363 TOps/s
 Deconvolution: 204.8 TOps/s
 FFT: 125.426 TOps/s
 A-Projection: 24.9131 TOps/s
 Degridding: 172.938 POps/s

Total compute: 365.614 POps/s

Estimated power consumption: 14.7803 MWatt (L1 Requirement: 3.5 MWatt)

▼ Additional data

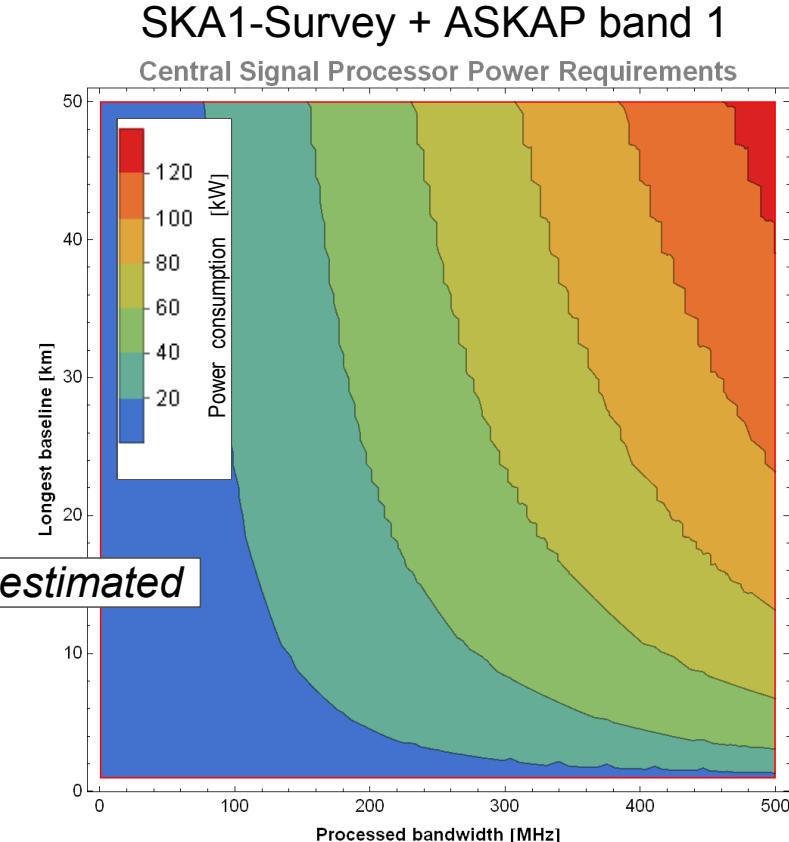
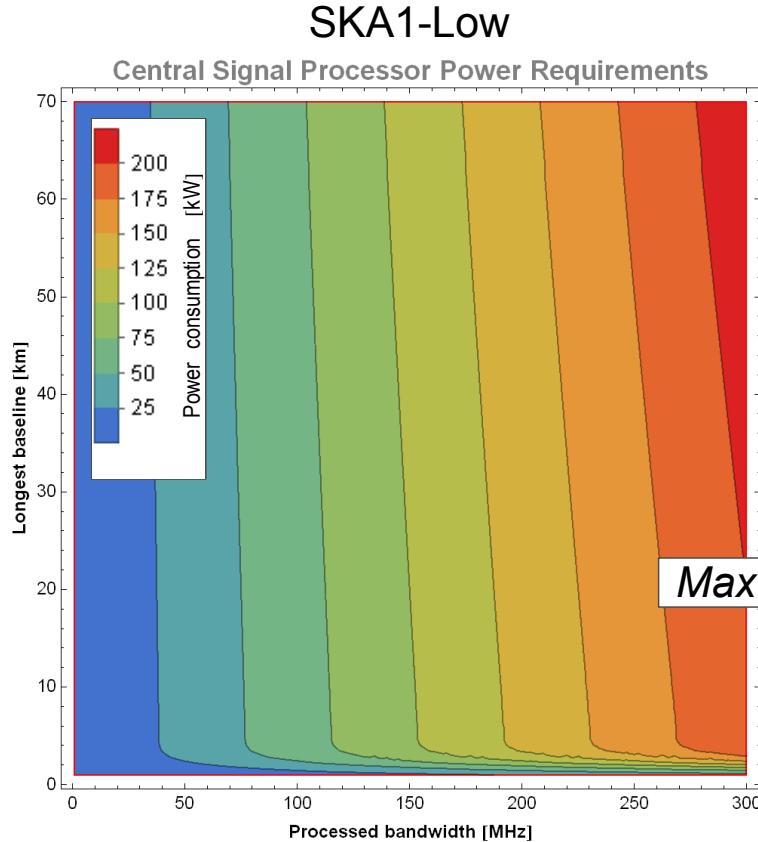
Gridding:

Minimum Rf: 297.619
 Maximum Rf: 122.55
 Average Rf: 184.856
 wrms/wmax (estimate): 0.3128
 Ra (estimate): 7
 Corrected minimum w-term: 93.0952
 Corrected maximum w-term: 38.3337
 Corrected average w-term: 57.823

Imaging:

Pixels: 12500

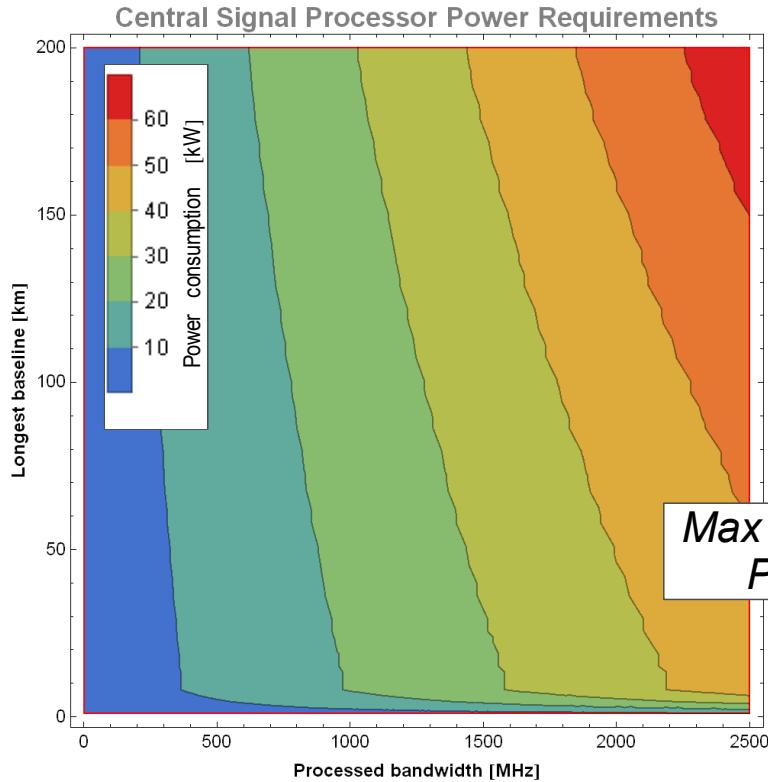
SKA1-Low and SKA1-Survey CSP power budget: 1 MW



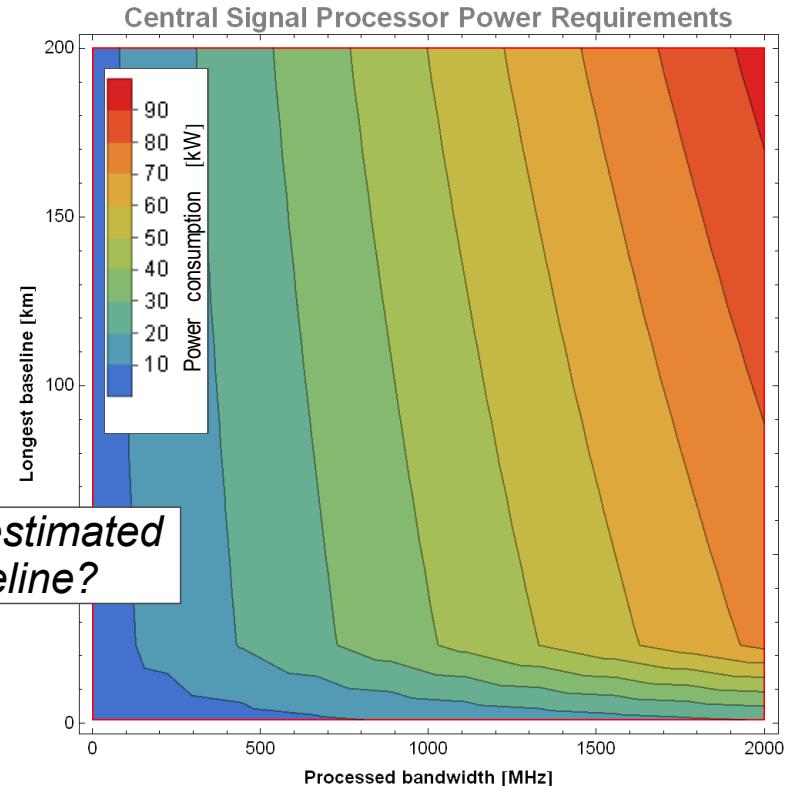
- Power equally distributed over SKA1-Low and SKA1-Survey (simultaneous observations)

SKA1-Mid CSP power budget: 2.5 MW

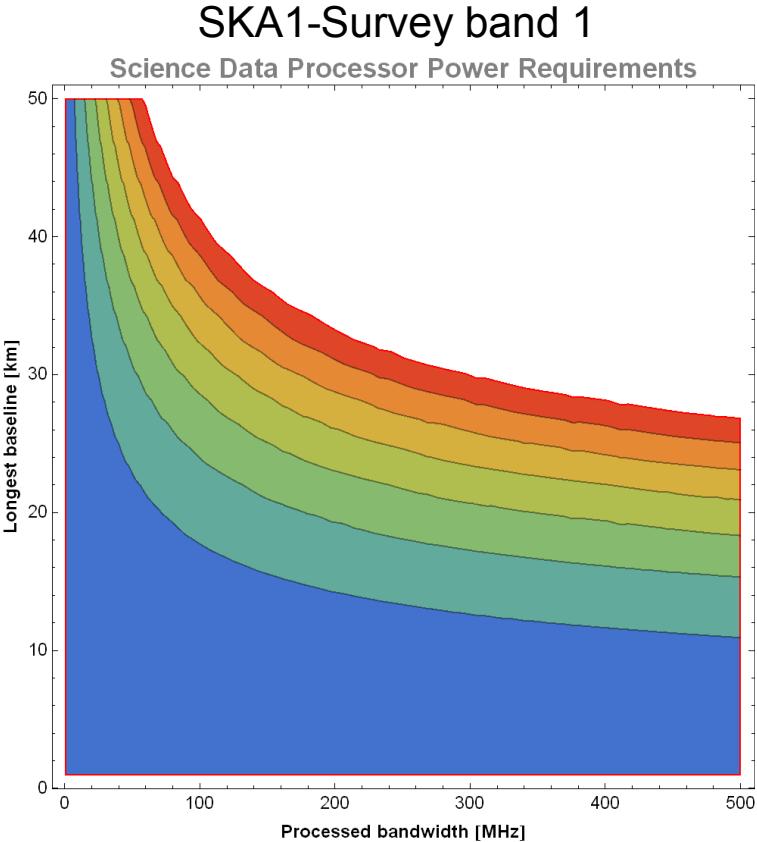
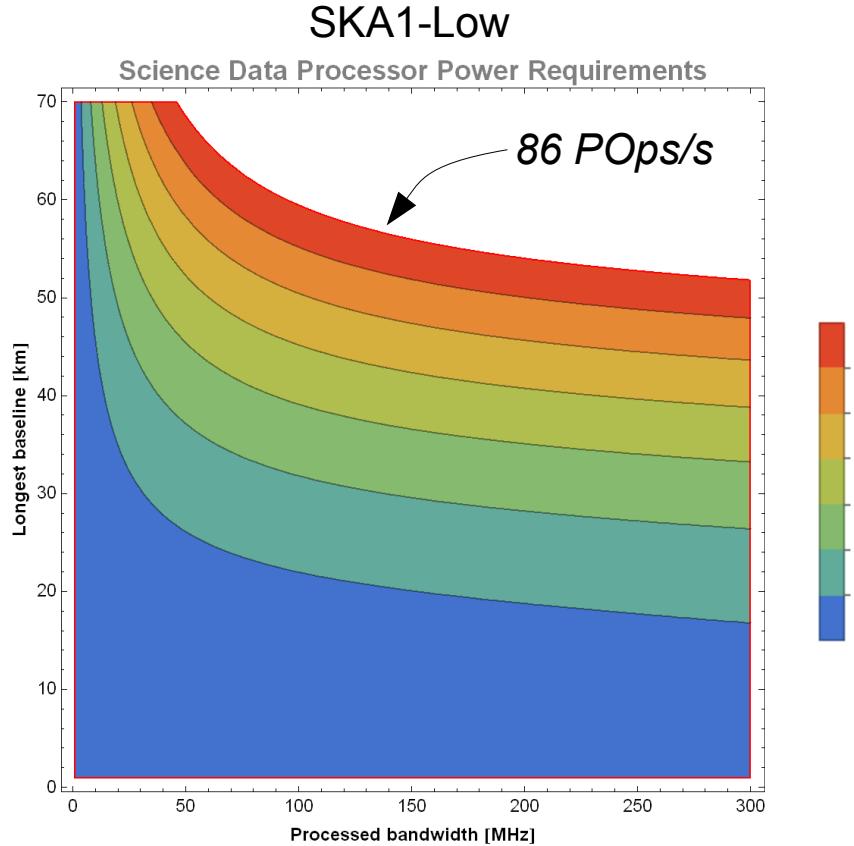
SKA1-Mid band 5



SKA1-Mid MeerKAT band 3



SKA1-Low and SKA1-Survey SDP power budget: 3.5 MW



- 3.5 MW for both SKA1-Low and SKA1-Survey
- Baseline dependent time/frequency averaging