End-to-end compute model of the Square Kilometre Array

Rik Jongerius*†, Stefan Wijnholds‡, Ronald Nijboer‡, Henk Corporaal†
*IBM Research, The Netherlands
†Eindhoven University of Technology, The Netherlands
‡ASTRON, The Netherlands

***Abstract*—Scientific requirements for next-generation radio telescopes significantly influence compute requirements of their processing facilities. For future instruments, such as the Square Kilometre Array, seemingly minor changes can easily push compute requirements into the exascale domain. We present a model for engineers and astronomers to understand these relations and to make trade-offs in radio telescope designs.**

I. Introduction

The future Square Kilometre Array (SKA) radio telescope will provide more sensitivity and *survey speed* than any present-day instrument [1]. The telescope, to be built in South Africa and Australia, will be constructed in two phases. Construction of phase one starts in 2018 and early science results are expected around 2020. The telescope will be extended to its full size in phase two.

For SKA phase one, more than 250,000 dual dipole antennas and 350 dishes are placed in the field, resulting in a big data challenge: after digitization, the antennas alone generate raw data at a rate of more than 2.5 Pbit/s. As it is infeasible to store this data for long observations, this will be processed in real-time. Two major challenges exist: how can we provide the compute power, at low energy consumption, to process the data stream and how can we built the infrastructure capable of transporting this large amount of data towards the distributed processing facilities.

The design process of a radio telescope is a complex affair. Radio astronomers drive the design by posing constraints based on astronomical science. These parameters translate into requirements for compute performance and data transport. For a feasible telescope, a balanced design is needed between, on one hand, affordable compute and bandwidth requirements, and, on the other hand, a telescope fulfilling the astronomers' requirements. In this article, we present models for compute and bandwidth requirements of an end-to-end digital processing system for radio telescopes and derive a lower bound on energy consumption for the most demanding processing facilities. System designers can use these models to assess telescope designs and understand the impact of proposed changes.

This work is conducted in the context of the joint ASTRON and IBM DOME project and is funded by the Netherlands Organisation for Scientific Research (NWO), the Dutch Ministry of EL&I, and the Province of Drenthe.

II. Square Kilometre Array—Phase 1

We will use our models to assess the three instruments constructed in phase 1 of the SKA [1], [2]. The first instrument, SKA1-Low, receives signals in the band between 50 and 350 MHz and uses phased-array technology. A total of 1024 35-m stations will be constructed, each consisting of 256 dual-polarized antennas. The signals from the antennas are beamformed per station to point the telescope on the sky. The maximum *baseline* length—the distance between a pair of stations or dishes—is 70 km.

The SKA1-Mid array uses 15-m dishes with single-pixel feeds to receive signals between 350 MHz and 13.8 GHz. The signal processing back-end is dimensioned to process up to 2500 MHz out of the full receiver bandwidth. Currently, the MeerKAT telescope, with 12-m dishes, is built on the SKA1-Mid site as a precursor to SKA phase 1. Both instruments can be combined for observations with up to 254 dishes and baselines up to 200 km.

The third instrument, the dish-based SKA1-Survey, receives 500 MHz signal bandwidth between 350 MHz and 4 GHz. In contrast to SKA1-Mid, its 15-m dishes are equipped with phased-array feeds. Each phased-array feed contains 94 individual antennas and points 36 beams on the sky to increase the field of view of the instrument. Combined with the ASKAP 12-m dishes, 96 dishes are available for baselines of up to 50 km.

III. End-to-end compute model

The digital processing pipeline of radio telescopes for imaging science cases can be coarsely divided in three processing steps: 1) Local to the phased arrays, for example SKA1-Low or SKA1-Survey, signals are digitized, channelized—the frequency resolution is increased—, and beamformed to point the telescope at the sky. 2) Data from beams and dishes with single-pixel feeds (SPFs) is sent to a central signal processor (CSP) for further channelization and correlation. Correlating two data streams, and integrating over a short time span, yields a *visibility*, a sample of the Fourier-transformed sky. 3) Visibilities are sent to a science data processor (SDP), where the instrument is calibrated and a sky image is created.

Existing literature on compute systems for the SKA, or other radio telescopes, focuses on parts of the processing pipeline. Several studies map the algorithms used in the stations and the CSP onto platforms like dedicated ASICs and determine the

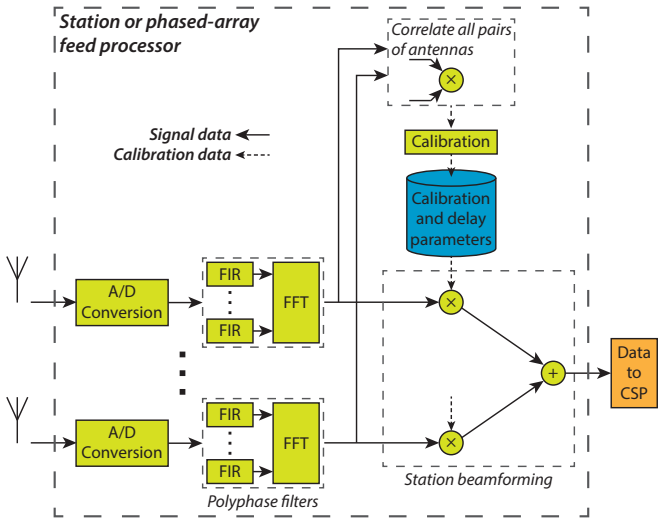

Fig. 1. A block diagram for phased-array processing. Antenna signals are channelized using polyphase filters, calibrated, and delay-compensated before they are beamformed. Beamformed samples are transported to the central signal processor.

impact of telescope requirements onto system size and power requirements [3]. Algorithms in later stages of the pipeline, for example calibration [4] and imaging [5], are actively researched and developed by the radio astronomy community to improve the quality of the images and to reduce the compute load.

The compute and bandwidth models derived in this work build upon this prior work from the community by using their computational scaling rules for several of the algorithms. Together with the scaling rules we derive, we construct an end-to-end model for a radio telescope. The imaging pipeline of the low-frequency array (LOFAR) telescope, a pathfinder instrument for the SKA, is used as a basis for our models [6].

A. Station processing

In phased-array telescope systems, antennas are grouped in *stations* and data is first processed close to the phased array as depicted in Figure 1. After digitization, signals from the antennas are channelized using polyphase filter banks. The polyphase filters consist of finite impulse response (FIR) filters followed by a fast Fourier transform (FFT). Channelization increases the frequency resolution by distributing the signal over *sub-bands*.

Each phased array contains N_{ant} antennas with each N_{pol} polarizations. Given the signal bandwidth Δf_{signal}, the signals are Nyquist sampled at a rate of $2\Delta f_{\text{signal}}$ Hz. An N_{tap}-tap FIR-filter performs approximately $2N_{\text{tap}}$ operations, and a real-to-complex N-point FFT $2.5N\log_2(N)$ operations. The number of points in the FFT is twice the number of sub-bands. The sub-band bandwidth and sub-band sampling rate are both Δf_{band} Hz for $N_{\text{bands}} = \frac{\Delta f_{\text{signal}}}{\Delta f_{\text{band}}}$ sub-bands. Together, all polyphase filters in a station require

$$\begin{aligned} R_{\text{ppf}} &= 2N_{\text{tap}}N_{\text{ant}}N_{\text{pol}}2\Delta f_{\text{signal}} \\ &\quad + 5N_{\text{bands}}\log_2(2N_{\text{bands}})\,N_{\text{ant}}N_{\text{pol}}\Delta f_{\text{band}} \end{aligned} \tag{1}$$

real operations per second.

Beamforming points the phased array at the sky by delaying antenna signals depending on the beam direction and adding channels from different antennas together. A complex gain implements the time delay and includes multiplication with calibration parameters. For each of the N_{beams} beams, the complex multiplication and the complex addition for beamforming, require 8 operations, resulting in a compute rate of

$$R_{\text{bf}} = 8N_{\text{beams}}N_{\text{bands}}N_{\text{ant}}N_{\text{pol}}\Delta f_{\text{band}}. \tag{2}$$

Calibration parameters vary slowly with respect to the sampling rate, and, as a result, correlation and calibration is performed every few minutes on a subset of the data. The cost of computing all N_{pol}^2 correlations per pair of antennas is similar to the compute cost of the CSP correlator, derived later in this work. Compute cost of calibration is not included, as it is small compared to the other algorithms.

Beams are transported to the CSP for further processing. The output data rate of a single station in bits per second is

$$B_{\text{stat}} = N_{\text{beams}}N_{\text{pol}}N_{\text{bands}}\Delta f_{\text{band}}b_{\text{beam}}, \tag{3}$$

where b_{beam} is the complex sample size.

B. Central signal processor

The central signal processor correlates channelized beams from the phased arrays, or signals from dishes with single-pixel feeds. For the SKA, three CSPs are constructed, each receiving data from N_{stat} stations or dishes. Several steps are performed before the actual correlation, as is shown in Figure 2.

A time delay buffer corrects for geometric delays between stations, after which a second polyphase filter (further) increases the frequency resolution of the signal. The polyphase filters generate complex *channels* with a bandwidth and sampling rate of Δf_{chan} Hz. The compute rate of the polyphase filters is similar to the compute rate in Equation 1 for filters in the stations, however, it is larger for phased-array beams, due to the complex nature of the input samples. The channelized signals are phase-delayed, a fine-grained correction for the geometric delay, and corrected for an artifact introduced by the stations' polyphase filters.

Depending on the science case, $N_{\text{channels}} = \frac{\Delta f_{\text{image}}}{\Delta f_{\text{chan}}}$ channels are correlated, with $\Delta f_{\text{image}} \leq \Delta f_{\text{signal}}$. Per baseline, a pair of stations or dishes, N_{pol}^2 correlations are calculated. Including autocorrelations, the total number of visibilities is $0.5N_{\text{stat}}^2N_{\text{pol}}^2$ per integration time and the compute rate is

$$R_{\text{cor}} = 4N_{\text{beams}}N_{\text{channels}}N_{\text{stat}}^2N_{\text{pol}}^2\Delta f_{\text{chan}}. \tag{4}$$

Visibilities contaminated with radio-frequency interference (RFI) are flagged and removed. A model derived by Offringa for the AOFlagger [7], used by LOFAR, estimates 278 operations per second per visibility to remove RFI. The rate at which visibilities are generated depends on the dump time, or integration time, τ_{csp} and equals

$$V_{\text{csp}} = N_{\text{beams}}N_{\text{channels}}0.5N_{\text{pol}}^2N_{\text{stat}}^2\frac{1}{\tau_{\text{csp}}}. \tag{5}$$

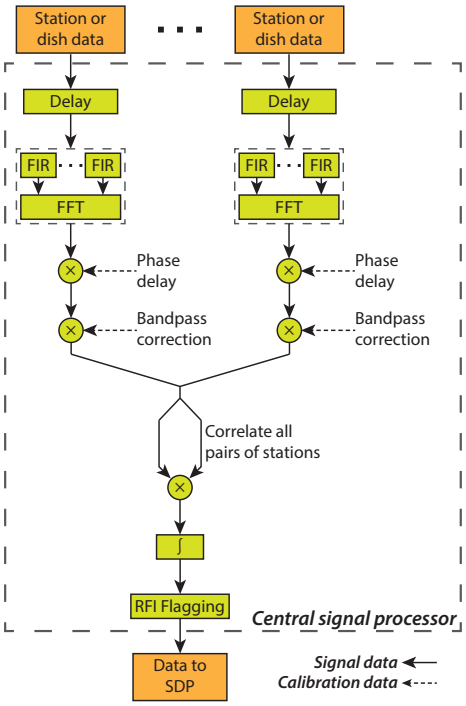

Fig. 2. An overview of the central signal processor for sky imaging. The CSP aligns data in time, performs further channelization and correlates data between pairs of stations or dishes. The resulting integrated visibilities are checked for RFI and transported to the science data processor.

The output data rate of the CSP in bits per second is the visibility rate multiplied with the size of the complex visibilities. We neglect data reduction due to RFI flagging, as, in practice, only a few percent of the data is contaminated with RFI.

C. Science data processor

Visibilities, received from the central signal processor, are imaged for each beam in the science data processor. For details about the imaging process, we refer to Taylor et al. [8]. While the exact calibration strategy is still unknown and actively researched with simulations and the precursor instruments, a potential strategy, based on LOFAR, is shown in Figure 3. The process contains three cycles: the *minor*, *major*, and *calibration cycle*. We estimate the operation count required for each algorithm, and determine the minimum compute load for near-real-time operation based on the iteration count. Our model focuses on the algorithms in the major and minor cycle, as it is expected that those are dominant in terms of computational load. Compute costs for demixing, calibration, and foreground subtraction are not included.

Demixing removes the contributions of bright radio sources, for example Cassiopeia A or Centaurus A, from the visibilities at the maximum time and frequency resolution. After demixing, time and frequency resolution can be further reduced based on science requirements.

Due to earth rotation, the required frequency and time resolution increases with baseline length. We consider different dump times for long and short (array core) baselines to reduce computational load. Given a fraction F_{short} of short baselines, between 0.5 and 0.9 for the SKA instruments, and two dump times τ_{short}, and τ_{long}, the visibility rate per channel and per beam, after integration, is given by

$$V_{\mathrm{chan}} = 0.5 N_{\mathrm{pol}}^2 N_{\mathrm{stat}}^2 \left(\frac{F_{\mathrm{short}}}{\tau_{\mathrm{short}}} + \frac{(1 - F_{\mathrm{short}})}{\tau_{\mathrm{long}}} \right). \quad (6)$$

The frequency channels received from the CSP are further averaged into channels of Δf_{sdp} Hz wide. The imager creates one image per frequency channel, resulting in $N_{\mathrm{image-chan}} = \frac{\Delta f_{\mathrm{image}}}{\Delta f_{\mathrm{sdp}}}$ images per beam.

Based on an initial model of the sky, containing a set of known sources, a first estimate on calibration parameters is made. Visibilities are corrected and stored in a buffer. The imaging process improves image quality by iteratively extending the sky model in the major and minor cycles, and optimizing calibration parameters in the calibration cycle.

Known sources are subtracted from the visibilities and the result is *gridded* onto a Fourier plane. Gridding of visibilities can be achieved using the W-projection algorithm [5]. The number of operations to grid the visibilities per beam is [9]

$$C_{\mathrm{grid}} = N_{\mathrm{op}} 2 V_{\mathrm{chan}} T_{\mathrm{obs}} \sum_{i=0}^{N_{\mathrm{image-chan}}-1} \left(\left(\frac{w_{\mathrm{rms}}}{w_{\mathrm{max}}} \right)^2 R_{\mathrm{F}}^2 + R_{\mathrm{A}}^2 \right), \quad (7)$$

where

$$R_{\mathrm{F}} = \frac{\lambda_i B_{\mathrm{max}}}{D^2}. \quad (8)$$

Here, N_{op} is the number of operations required to grid one visibility to one grid point, T_{obs} the observation time, λ_{i} the observed wavelength per channel, B_{max} the longest baseline, D the station or dish diameter, and R_{A} the size of the kernel to taper the field of view. The factor $\frac{w_{\mathrm{rms}}}{w_{\mathrm{max}}}$ is related to the distribution of different baseline lengths, and is beyond the scope of this article, more information can be found in Humphreys et al. [10]. Based on experience with LOFAR, we set $N_{\mathrm{op}} = 10$, $R_{\mathrm{A}} = 7$, and $\frac{w_{\mathrm{rms}}}{w_{\mathrm{max}}} \approx 0.3$. We assume that the gridding kernels are precomputed. However, the kernels have to be initialized and experience with LOFAR learns us that this, when time-dependent effects are taken into account, requires a significant amount of resources.

The Fourier grid is transformed using a 2D inverse FFT and generates a residual image cube with $N_{\mathrm{image-chan}}$ image planes. The number of pixels along each axis of an image is approximately $N_{\mathrm{pixels}} = \frac{3 B_{\mathrm{max}}}{D}$, such that we need

$$C_{\mathrm{IFFT}} = N_{\mathrm{image-chan}} 5 N_{\mathrm{pixels}}^2 N_{\mathrm{pol}}^2 \log_2 \left(N_{\mathrm{pixels}}^2 \right) \quad (9)$$

operations for transforming one image cube.

The residual image cube is used in the minor cycle to improve the sky model. Iterative *source extraction*, or deconvolution, identifies sources in the residual image and subtracts the *point spread function (PSF)*—the telescope's response to a point source in the sky—for each source from the residual image. For one iteration of the minor cycle, subtraction of the PSF requires a multiply and a subtract operation per pixel, resulting in a compute requirement of

$$C_{\mathrm{source-extract}} = 2 N_{\mathrm{image-chan}} N_{\mathrm{pixels}}^2 N_{\mathrm{pol}}^2 \quad (10)$$

operations.

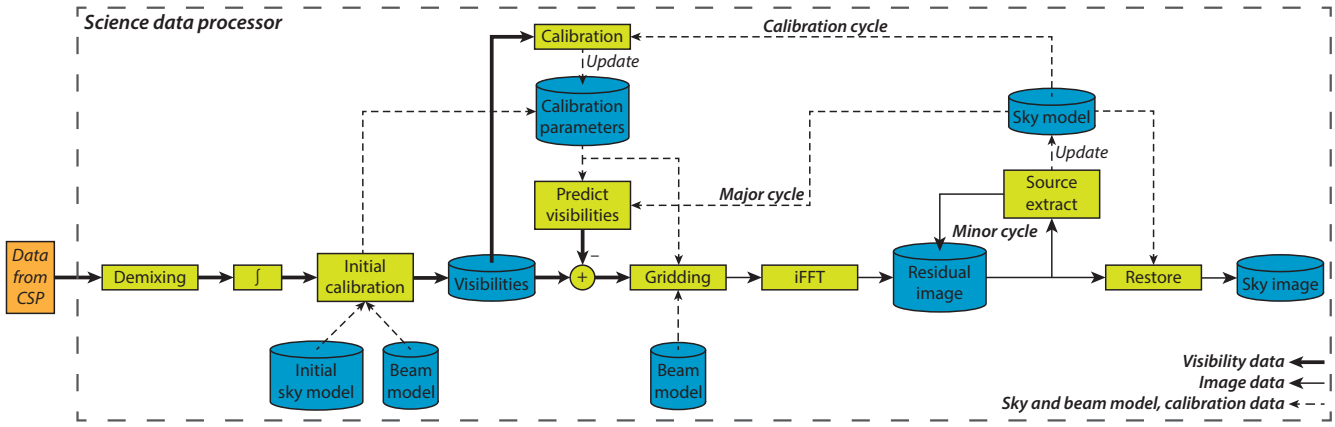

Fig. 3. A potential calibration and imaging strategy in the science data processor for the SKA. Visibilities are processed in three nested loops: the calibration, major, and minor cycle. An image cube is generated from the visibilities, after which the sky model is updated using several iterations of the major and minor cycles. With the updated sky model, new calibration parameters are determined and a new iteration is started to refine the sky model.

Subtraction of the PSF is not perfect in the image domain, making it impossible to find weak sources. Therefore, after several iterations of the minor cycle, a new major cycle is started. The contribution of sources found thus far is predicted, which can be achieved with a forward 2D FFT and degridding step with similar compute requirements as forward gridding and the 2D inverse FFT. After source subtraction, a new residual image cube is created for further deconvolution, and subtracted from the visibilities.

After several major cycles, new calibration parameters are found using the improved sky model and a new calibration cycle is started to further enhance the sky model. After a predefined set of iterations, based on best practices, the final sky image is created by restoring the sources found in the residual image.

IV. Compute and Bandwidth Estimates

We use the model to estimate minimum compute load, in multiply or add operations per second, and data bandwidths for continuum imaging observations in one band of the phase 1 instruments. For SKA1-Low only one frequency band exists, for SKA1-Mid and SKA1-Survey we select the lowest frequency band. While the signal bandwidth in band 1 of SKA1-Mid is 1 GHz, only 435 MHz is processed which overlaps with MeerKAT's UHF band. For phased arrays, 2048 bands are generated and station correlation and calibration is performed every 4 minutes. For all three instruments, 262,144 channels are generated in the CSP. Baseline-dependent time resolutions and frequency resolutions are retrieved from the baseline design, however, for SKA1-Low and SKA1-Survey channels are integrated in the SDP to respectively 2.3 kHz and 7.6 kHz based on bandwidth smearing requirements. For calibration and imaging, we set three calibration cycles per 20 minutes of observing, ten major cycles per calibration cycle, and 100 minor cycles per major cycle, based on experience with LOFAR.

Table I shows the results of the analysis for the three instruments. Station processing is only required for SKA1-Low and SKA1-Survey, the phased-array instruments. Channelization dominates for SKA1-Low due to the large amount of antennas, while beamforming of 36 beams is the main contributor for SKA1-Survey. The station processor requires a sustained compute throughput of 17.1 POps for SKA1-Low and 4.4 POps/s for SKA1-Survey. Especially for SKA1-Low, station processing significantly reduces the data rate from the antennas.

For the central signal processor, correlation is the most compute-intensive step for all three instruments. However, for both SKA1-Low and SKA1-Survey, more compute is needed for station processing than for the CSP. It should be noted that the SKA1-Mid correlator load is higher for wider bands, however, we selected band 1 due to the higher compute load of the imager. The data rate decreases after correlation for SKA1-Survey due to a long enough integration time. In contrast, the data rate after correlation is higher than the input data rate for both SKA1-Low and SKA1-Mid.

The science data processor has the highest compute load for all three instruments, and is dominated by gridding and prediction of visibilities. SKA1-Mid has by far the highest requirement with 1.8 EOps/s. Compute requirements for SKA1-Low and SKA1-Survey are of the same order, around 150–300 POps/s. Primarily the 200-km maximum baseline for SKA1-Mid strongly influences the high compute load.

V. Energy Consumption

The results show that more than 90 % of the required operations are performed in the SDP. We can find a lower bound on energy consumption by estimating the power consumption of compute operations in the SDP. We assume that the SDP uses double-precision floating-point arithmetic. Keckler et al. estimates energy consumption of a double-precision fused multiply-add (DFMA, 2 operations) at 8.7 pJ in 2017's 10-nm technology [11]. Construction of the SKA starts in 2018, so we scale the power consumption to 2018 technology based on CV^2 scaling from the ITRS road map [12].

We estimate the energy of one DFMA at 7.95 pJ in 2018. Based on these numbers, power required for the DFMA logic only, excluding, for example, data movement and instruction

TABLE I

MINIMUM COMPUTE THROUGHPUT AND REQUIRED DATA BANDWIDTH FOR CONTINUUM IMAGING WITH THE THREE INSTRUMENTS IN THE LOWEST FREQUENCY BAND. IMAGING WITH SKA1-MID AND SKA1-SURVEY INCLUDES THE PRECURSOR INSTRUMENTS MEERKAT AND ASKAP. NO ASSUMPTIONS ARE MADE ON INTEGER OR FLOATING-POINT ARITHMETIC, AS THIS DOES NOT INFLUENCE THE ACCURACY OF THE INSTRUMENT.

Processing stage	**Processing step**	**SKA1-Low**		**SKA1-Mid band 1 and MeerKAT UHF band**		**SKA1-Survey and ASKAP band 1**	
	Raw data stream	*One station* 307 GB/s	*1024 stations* 314 TB/s	*One dish* 4 GB/s	*254 dishes* 760 GB/s	*One dish* 188 GB/s	*96 dishes* 18 TB/s
Station processing	Channelization	14.1 TOps/s		-		8.6 TOp/s	
	Beamforming	1.2 TOps/s		-		27.1 TOps/s	
	Correlation and calibration	1.3 TOps/s		-		10.6 TOps/s	
		One station	*1024 stations*	-		*One dish*	*254 dishes*
	Total station compute	16.7 TOps/s	17.1 POps/s	-		46.3 TOp/s	4.4 POps/s
	Output bandwidth	1.2 GB/s	1.2 TB/s	-		72 GB/s	6.9 TB/s
Central signal processor	Channelization	60.8 TOps/s		64.5 TOps/s		342.1 TOp/s	
	Correlation	5.0 POps/s		449.0 TOps/s		2.7 POps/s	
	RFI Flagging	254.7 TOps/s		51.1 TOps/s		161.2 TOps/s	
	Total CSP compute	5.3 POps/s		564.7 TOps/s		3.2 POps/s	
	Output bandwidth	7.3 TB/s		1.5 TB/s		4.6 TB/s	
Science data processor	Gridding	83 POps/s		916.4 POps/s		154.6 POps/s	
	Inverse 2D FFT	59.2 TOps/s		10.2 POp/s		139.4 TOps/s	
	Deconvolution	94.4 TOps/s		13.1 POps/s		204.8 TOps/s	
	Predict visibilities	74.7 POps/s		824.8 POps/s		139.1 POps/s	
	Total SDP compute	158.0 POps/s		1.8 EOps/s		294.2 POps/s	

overhead, ranges from 628 kW for SKA1-Low to 7.2 MW for SKA1-Mid in the SDP. As DFMA energy for CPUs was only 3 % of the total energy per instruction in 2010, the potential power consumption of the SDP can easily be several tens of times higher than the DFMA power [11].

VI. CONCLUSIONS

Based on the models and the analysis, several conclusions can be made about the SKA phase 1 design. The SDP will not be located on the same site as the instrument itself and, as the central signal processor increases the data rate for both SKA1-Low and for SKA1-Mid, the CSP should be physically located close to the SDP to reduce long-distance bandwidth requirements. A second solution is to move demixing and integration from the SDP to the CSP: baseline-dependent averaging significantly reduces the required bandwidth.

The science data processor is by far the most demanding with compute throughput requirements up to 4 EOps/s and a high power consumption. In particular, (de-)gridding imposes high requirements on the system, and research should focus here to reduce the compute load and, ultimately, power consumption. Both telescope calibration and gridding is still an active area of research, and better algorithms or reducing the number of cycles may have a significant impact on compute load. An other option is to construct the SDP in phases and have technology improve compute throughput and power consumption. Such a strategy would imply that the array cannot be used to its full potential immediately.

The model and analysis presented here give an understanding of the impact of science requirements and telescope parameters on the compute platform for sky imaging. Such models provide insight in compute and bandwidth hot spots for future radio telescopes, such as the Square Kilometre Array. In future work, the model should be extended to include calibration and different operating modes, such as the beamforming modes. Furthermore, if the model is used to optimize the compute load, telescope quality metrics should be added as there are no trivial configurations.

REFERENCES

[1] P. E. Dewdney, W. Turner, R. Millenaar, R. McCool, J. Lazio, and T. J. Cornwell, "SKA1 system baseline design," www.skatelescope.org/?attachment_id=5400, SKA, Tech. Rep. SKA-TEL-SKO-DD-001, March 2013.

[2] R. McCool and T. J. Cornwell, "Miscellaneous corections to the baseline design," SKA, Tech. Rep. SKA-TEL.SKO-DD-003, October 2013.

[3] L. D'Addario, "Low-power architectures for large radio astronomy correlators," in *General Assembly and Scientific Symposium, 2011 XXXth URSI*, 2011, pp. 1–4.

[4] S. Wijnholds, S. van der Tol, R. Nijboer, and A.-J. van der Veen, "Calibration challenges for future radio telescopes," *Signal Processing Magazine, IEEE*, vol. 27, no. 1, pp. 30–42, 2010.

[5] T. Cornwell, K. Golap, and S. Bhatnagar, "The noncoplanar baselines effect in radio interferometry: The w-projection algorithm," *Selected Topics in Signal Processing, IEEE Journal of*, vol. 2, no. 5, pp. 647–657, 2008.

[6] M. P. Van Haarlem, M. W. Wise, A. W. Gunst *et al.*, "LOFAR: The Low-Frequency Array," *A&A*, vol. 556, Aug. 2013.

[7] A. R. Offringa, A. G. de Bruyn, S. Zaroubi, and M. Biehl, "A LOFAR RFI detection pipeline and its first results," in *Proceedings of Science, RFI2010*, 2010.

[8] G. B. Taylor, C. L. Carilli, and R. A. Perley, *Synthesis Imaging in Radio Astronomy II*. Astronomical Society of the Pacific, 2003.

[9] S. J. Wijnholds and R. Jongerius, "Computing cost of sensitivity and survey speed for aperture array and phased array feed systems," in *AFRICON 2013*, 2013.

[10] B. Humphreys and T. Cornwell, "Analysis of convolutional resampling algorithm performance," SKA, Memo 132, January 2011.

[11] S. Keckler, W. Dally, B. Khailany, M. Garland, and D. Glasco, "GPUs and the future of parallel computing," *Micro, IEEE*, vol. 31, no. 5, pp. 7–17, 2011.

[12] International Roadmap Committee. (2012) International Technology Roadmap for Semiconductors. http://www.itrs.net/Links/2012ITRS/Home2012.htm.

APPENDIX

Rik Jongerius is a predoctoral researcher at IBM Research and a PhD candidate at the department of Electrical Engineering at the Eindhoven University of Technology. His research

interests include high-performance computing, computer architecture, and accelerators. He received his MSc degree in Electrical Engineering from the Eindhoven University of Technology. He is a student member of the IEEE. Contact him at r.jongerius@ieee.org.

Stefan Wijnholds is a system researcher with the R&D department of ASTRON, the Dutch institute for radio astronomy. His research interests are in the area of calibration and imaging algorithms and system design of next-generation radio telescopes. He received his MSc in Astronomy and his MEng in Applied Physics (both cum laude) from the University of Groningen and his PhD (cum laude) from Delft University of Technology. He is a senior member of the IEEE. Contact him at wijnholds@astron.nl.

Ronald Nijboer leads the Computing group within the R&D department of ASTRON. His research interests are in the area of science data processing for radio telescopes. He received his MSc in Mathematics from the Eindhoven University of Technology and his PhD in Plasma Physics from the VU University Amsterdam. Contact him at rnijboer@astron.nl.

Henk Corporaal is Professor in Embedded System Architectures at the Eindhoven University of Technology. He has gained a MSc in Theoretical Physics from the University of Groningen, and a PhD in Electrical Engineering, in the area of Computer Architecture, from Delft University of Technology. His current research projects are on low power single and multi-processor architectures, their programmability, and the predictable design of soft and hard real-time systems. For further details see his website at corporaal.org.