Astronomical Workloads

Kees van Berkel

MPSoC 2015,
July 2015, Ventura, CA

Technische Universiteit
Eindhoven
University of Technology

Where innovation starts

Radio Astronomy

Tidal interactions in the M81 group

stellar light distribution

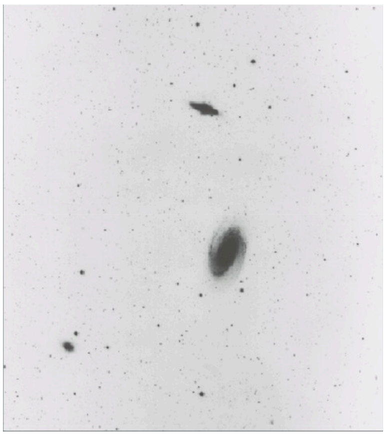

21cm HI distribution

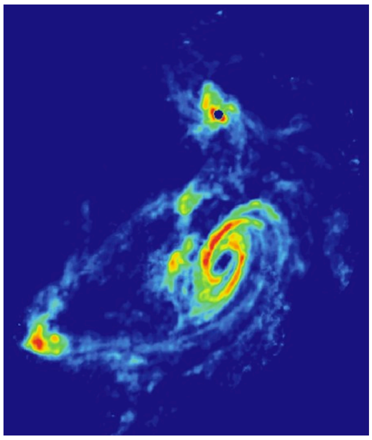

Image courtesy of NRAO/AUI

Technische Universiteit Eindhoven University of Technology

Interferometry

Westerbork Synthesis Radio Telescope:
14 dishes, D=25m, B=3km [NL,1956]

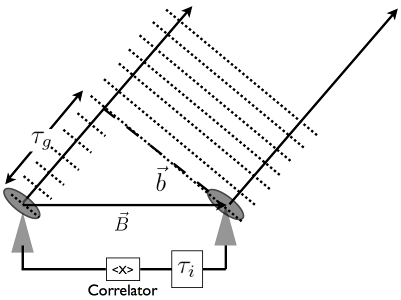

2-element interferometer.
Output of the correlator:

$$V_\nu(\mathbf{r_1}, \mathbf{r_2}) = \langle \mathbf{E}_\nu(\mathbf{r_1}) \mathbf{E}^*_\nu(\mathbf{r_2}) \rangle$$

with v the observation frequency
and * denoting complex conjugation

Technische Universiteit Eindhoven University of Technology

Van Cittert–Zernike theorem [1934-38]

correlator output; sky intensity; speed of light; solid angle; base line vector, separating the 2 antennae

$$V_\nu(\mathbf{r}_1, \mathbf{r}_2) \approx \int I_\nu(\mathbf{s}) e^{-2\pi i \nu \mathbf{s}\cdot(\mathbf{r}_1 - \mathbf{r}_2)/c} d\Omega$$

Adding geometry (assuming "narrow field"):

$$V_\nu(u, v) = \iint I_\nu(l, m) e^{-2\pi i (ul + vm)} \, dl \, dm$$

2D Fourier transform!

where *(l, m)* are sky-image coordinates
and *(u, v)* are coordinates of the base-line vector

[1], [2], [3]

SAN TU/e Technische Universiteit Eindhoven University of Technology

Van Cittert–Zernike theorem [1934-38]

In principle:

(u, v) coverage
(A, φ) at (u,v)

I-FFT →

← FFT

(l,m) image
pixel intensity

v ↑

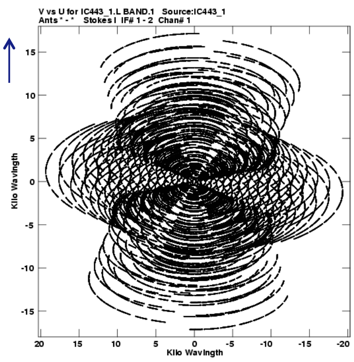

u →

m ↑

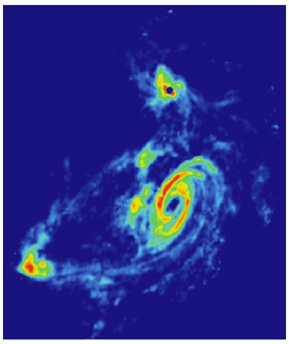

l →

SAN TU/e Technische Universiteit Eindhoven University of Technology

W-projection, W-snapshot [2008/12, Cornwell et al]

However, Van Cittert–Zernike theorem "wide-field"

$$V(u,v,w) = \int \frac{I(\ell,m)}{\sqrt{1-\ell^2-m^2}} e^{-2\pi i[u\ell+vm+w(\sqrt{1-\ell^2-m^2}-1)]} d\ell dm$$

Visibilities are 3D *(u, v, w)*, due to earth' curvature (Fresnel diffraction).

Choose as convolution function $G(\ell,m,w) = e^{-2\pi i[w(\sqrt{1-\ell^2-m^2}-1)]}$
and let *G˜(u, v, w)* be the Fourier transform of *G(l, m, w).*

Then, using the Fourier convolution theorem (W-projection):

$$V(u,v,w) = \tilde{G}(u,v,w) * V(u,v,w=0)$$

W-snapshot
= *W-projection* applied piecemeal to a series of snapshots. [4], [5]

Technische Universiteit Eindhoven University of Technology

Deconvolution (CLEAN, Högbom 1974)

I(l, m) $\xrightarrow{\text{FFT}}$ V(u, v, w=0) $\xrightarrow{**\tilde{G}(u,v,w)}$ V(u, v, w)

Can be computed straightforwardly, but cannot be inverted easily, because *V(u, v, w)* provides only a finite number of noisy samples (and a variety of other reasons, including antenna beam pattern).

CLEAN (next slide) is an iterative *deconvolution* algorithm.
(≈ interpolation in uv plane)

(Under certain conditions CLEAN converges to a solution that is the least-squares fit of the FFT transform.)

[1], [3]

Technische Universiteit Eindhoven University of Technology

Imaging (W-projection/snapshot + CLEAN)

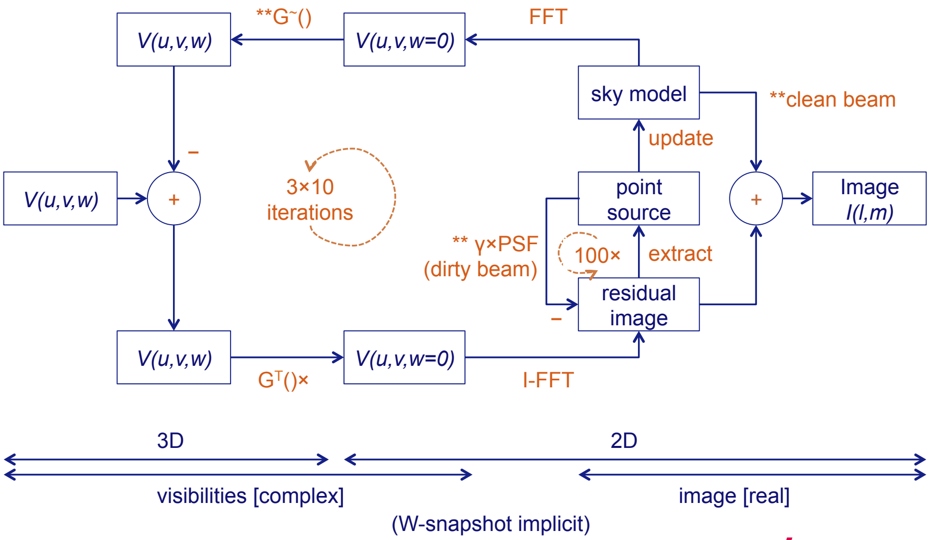

SKA1-mid [South Africa]: science in 2020

Towards a Square Kilometer Array

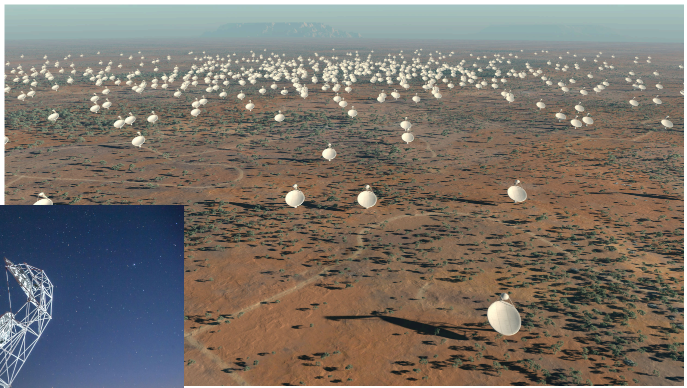

artist impression

photograph

SKA Organisation
/Swinburne Astronomy Productions [6]

Imaging: compute load for SKA1-mid

quantity	unit	^{10}log	note
# base lines		5.5	2^2 ×(#dishes + #stations)^2 = (2×254)^2
dump rate	s^{-1}	1	(integration time = 0.08s)$^{-1}$
observation time	s	3	
# channels		5	"image cube" for spectral analysis
# visibilities / observation		*14.5*	*= input to imaging (≈ 10^{16} Byte)*
# ops /visibility /iteration		4.5	convolution, matrix multiply, (I)FFT
# major iterations		1.5	(3×calibration) × (10×major)
# ops /observation		*20.5*	
# ops /sec	*Hz*	*17.5*	≈ 1 exa-(fl)ops / sec

- #operations/visibility/iteration depends on #snapshots
- calibration loop (3×) around imaging loop
- data type: double|single precision, floating|fixed point?

[7], [8], [9]

SAN TU/e Technische Universiteit Eindhoven University of Technology

Imaging: where is the parallelism?

quantity	unit	^{10}log	note
# ops / sec	*Hz*	*17.5*	*= imaging compute load*
margin (for inefficiencies)		0.5	very aggressive / optimistic
machine	*flop*	*18*	*= 1 exaflops/sec*
# clock frequency	Hz	9	
# channels in parallel		5	☺, all independent data streams!
simd? simt? pipelining?		4	☹, challenging!

Concerns on efficiency:

- data sets are large (≈ 10^{16} Bytes for visibilities),
- and some algorithms are low on compute intensity (high i/o) and or irregular, (e.g. FFT typically 20% efficiency on a CPU | GPU),
- Hence manual optimization of code likely essential.

SAN TU/e Technische Universiteit Eindhoven University of Technology

EXAflops/sec in 2015?

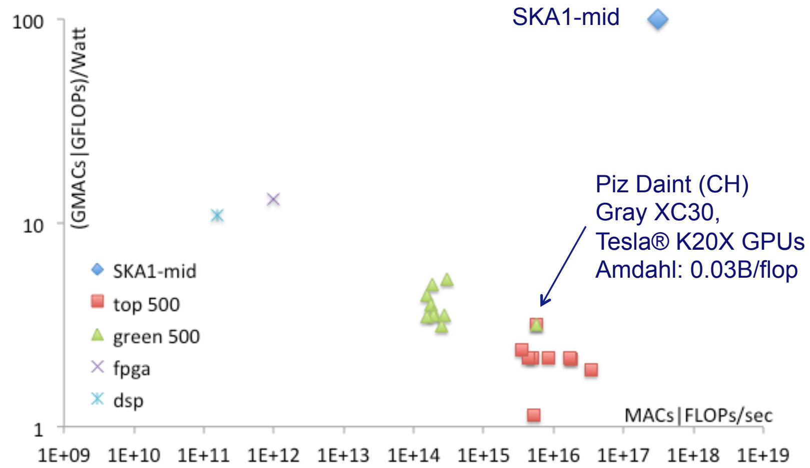

- *net* SKA1-mid computation load "*2020*" versus
- *gross* (peak) compute performance "*2015*"

[10]

TU/e Technische Universiteit Eindhoven University of Technology

EXAflops/sec in 2020?

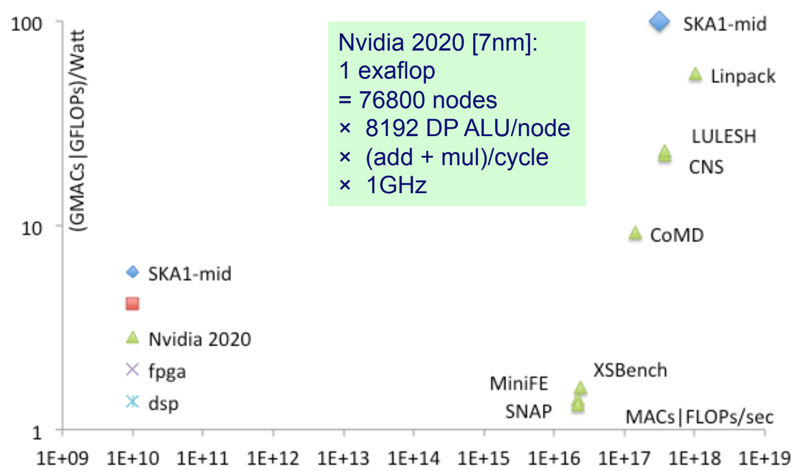

A huge spread per application in achievable FLOPs/sec and GFLOPs/Watt! [11]

TU/e Technische Universiteit Eindhoven University of Technology

Astronomical workloads

Exaflops/sec algorithms?

- Can we expect algorithm innovation beyond w-snapshot+CLEAN?
- Trade lower hot FLOPs (w-snapshot) vs higher cool FLOPs (w-projection)?
- Where can we afford single precision? (Fixpoint?)

Exaflops/sec machines?

- Will GPUs be the obvious accelerator? or will FPGAs or DSPs surprise us?
- Amdahl memory ratio (Byte/flop)?

Exaflops/sec mapping?

- Which forms of parallelism for highest efficiency? (Next to channel ||)
- What levels of efficiency are achievable?

Exaflops/sec requirements?

- When will exaflops/sec SKA1 power consumption be affordable?
- Will SKA2 (>100x) … ?

SAN TU/e Technische Universiteit Eindhoven University of Technology

References

1) G.B. Taylor, C.L. Carilli, and R.A. Perly (eds.), Synthesis Imaging in Radio Astronomy II, ASP Conf Series, Vol. 180, 1999.
2) B.G. Clark, Coherence in Radio Astronomy, pp. 1-10 in [1].
3) Thompson, A., Moran, J., & Swenson, G. 2001, Interferometry and synthesis in radio astronomy (Wiley, New York)
4) T. J. Cornwell, K. Golap, and S. Bhatnagar, "The Non- coplanar Baselines Effect in Radio Interferometry: The W-Projection Algorithm," *IEEE Journal of Selected Topics in Signal Processing*, vol. 2, no. 5, pp. 647–657, Oct. 2008.
5) T.J. Cornwell, M. A. Voronkova and B. Humphreys, Wide field imaging for the Square Kilometre Array, arXiv:1207.5861v1, 25 Jul 2012
6) P.E. Dewdney et al., SKA1 System Baseline Design, tech. report SKA-TEL-SKO-DD-001, SKA, Mar. 2013; www. skatelescope.org/?attachment id=5400.
7) R. Jongerius, S. Wijnholds, R. Nijboer, and H. Corporaal, "End-to-end compute model of the Square Kilometre Array," *IEEE Computer*, Sept. 2014, pp. 48-54.
8) Erik Vermij et al, "Challenges in exascale radio astronomy: Can the SKA ride the technology wave? Intl. Journal of High Performance Computing Applications 2015, Vol. 29(1) 37–50.
9) S. J. Wijnholds, A.-J. van der Veen, F. De Stefani, E. La Rosa, A. Farina, Signal Processing Challenges for Radio Astronomical Arrays,2014 IEEE ICASSP, pp. 5382-86.
10) The Green500 List - November 2014**,** http://www.green500.org.
11) Oreste Villa et al, Scaling the Power Wall: A Path to Exascale, SC14: Intl Conf. for High Performance Computing, Networking, Storage and Analysis, p830-841.

SAN TU/e Technische Universiteit Eindhoven University of Technology

DoE proxy-applications for scientific computing

	Proxy applications for scientific computing		
LINPACK	Linear algebra		
CNS	Compressible Navier Stokes equations		
LULESH	Solve hydrodynamics on a 3D mesh		
CoMD	Molecular dynamics: compute forces; update positions		
SNAP	Solve the neutral, linear Boltzmann transport eqn		
XSBench	Monte Carlo neutronics		
MiniFE	3D finite elements		

[11]

Technische Universiteit Eindhoven University of Technology

Key characteristics of DoE proxy-apps

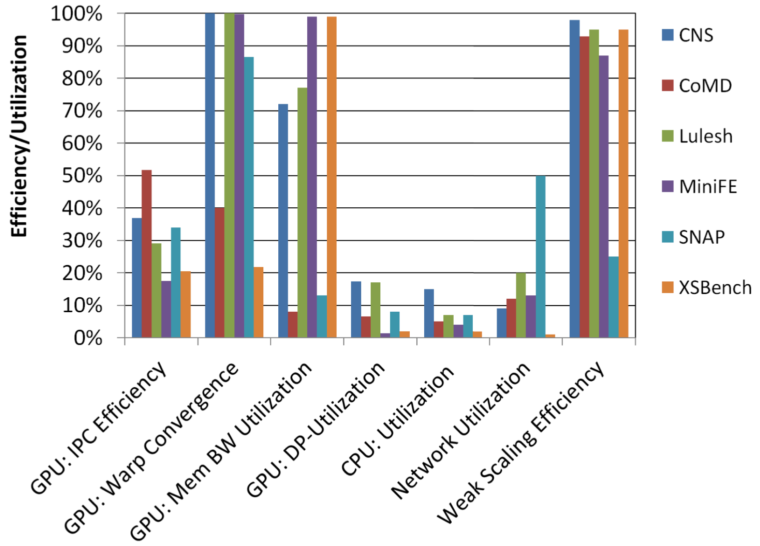

[11]

Technische Universiteit Eindhoven University of Technology