



UvA-DARE (Digital Academic Repository)

On optimising cost and value in compute systems for radio astronomy

Broekema, P.C.; Allan, V.; van Nieuwpoort, R.V.; Bal, H.E.

DOI

[10.1016/j.ascom.2019.100337](https://doi.org/10.1016/j.ascom.2019.100337)

Publication date

2020

Document Version

Final published version

Published in

Astronomy and Computing

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Broekema, P. C., Allan, V., van Nieuwpoort, R. V., & Bal, H. E. (2020). On optimising cost and value in compute systems for radio astronomy. *Astronomy and Computing*, 30, [100337]. <https://doi.org/10.1016/j.ascom.2019.100337>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.



Full length article

On optimising cost and value in compute systems for radio astronomy

P. Chris Broekema^{a,d,*}, Verity Allan^b, Rob V. van Nieuwpoort^{c,e}, Henri E. Bal^d^a *ASTRON, the Netherlands Institute for Radio Astronomy, Netherlands*^b *University of Cambridge, United Kingdom of Great Britain and Northern Ireland*^c *University of Amsterdam, Netherlands*^d *Vrije Universiteit Amsterdam, Netherlands*^e *Netherlands eScience Center, Netherlands*

ARTICLE INFO

Article history:

Received 1 April 2019

Accepted 16 October 2019

Available online 30 October 2019

ABSTRACT

Large-scale science instruments, such as the distributed radio telescope LOFAR, show that we are in an era of data-intensive scientific discovery. Such instruments rely critically on significant computing resources, both hardware and software, to do science. Considering limited science budgets, and the small fraction of these that can be dedicated to compute hardware and software, there is a strong and obvious desire for low-cost computing. However, optimising for cost is only part of the equation; the value potential over the lifetime of the solution should also be taken into account. Using a tangible example, compute hardware, we introduce a conceptual model to approximate the lifetime relative science value of such a system. While the introduced model is not intended to result in a numeric value for merit, it does enumerate some components that define this metric. The intent of this paper is to show how compute system related design and procurement decisions in data-intensive science projects should be weighed and valued. By using both total cost and science value as a driver, the science output per invested Euro is maximised. With a number of case studies, focused on computing applications in radio astronomy past, present and future, we show that the hardware-based analysis can be, and has been, applied more broadly.

© 2019 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Modern large-scale science instruments critically rely on specialised data-intensive computer technologies, to turn instrument data into useful science results. Considering limited science budgets, of which only a small fraction can be dedicated to computing, there is a strong desire to use these expensive systems in an optimal way. The design of such an optimised system is heavily influenced by experience from previous installations. For instance, the design priorities of the GPU-based correlator and beamformer system for the LOFAR radio telescope, in particular its focus on an I/O optimised design, borrowed heavily from previous experiences with Blue Gene based systems (Broekema et al., 2018).

In this paper we discuss both the cost and value of computing technologies, and how to optimise the combination of these two for maximum science impact. Since these are difficult to measure for the complex combination of hardware, middleware and software that are generally required, we focus our detailed analysis on hardware. We enumerate some of the factors that impact the total cost of a system. However, we propose that

total cost over the lifetime of a system is only part of the equation: the computational and scientific performance of different solutions may radically differ for the applications in question, depending on system and application characteristics. A more valuable metric would look at the useful output of a system per invested Euro. For example, the Distributed ASCI Supercomputer (DAS) (Bal et al., 2016) consortium tracks the effectiveness of its distributed cluster infrastructure via the number of awarded PhDs per cluster generation, as shown in Table 1.¹ Considering the nearly constant budget for these systems, between 1.2 and 1.5 M€, discounting inflation, the cost per supported PhD has dropped considerably over the lifetime of the DAS consortium. Alternatively, we can argue that the relative science value per invested Euro has dramatically increased.

In this paper we study a number of cases in radio astronomy, a computationally- and data-intensive science that has been using high-performance computing technologies since the very early days of computing to achieve scientific results. We show how the methodology proposed in this paper has informally been used in the past. The main contributions in this paper are:

* Corresponding author.

E-mail address: broekema@astron.nl (P.C. Broekema).

¹ Source: <https://www.cs.vu.nl/das4/phd.shtm>, <https://www.cs.vu.nl/das5/phd.shtm> and historical data.

Table 1
Awarded PhDs per Distributed ASCI Supercomputer generation.

	Year	PhDs	€/PhD	Research agenda
DAS-1	1997	7	€ 214.285	Wide-area computing
DAS-2	2002	22	€ 68.181	Grid computing
DAS-3	2006	36	€ 41.666	Optical grids
DAS-4	2010	33	€ 45.454	Clouds, diversity, green IT
DAS-5	2015	40	€ 37.500	Harnessing diversity & complexity

- the introduction of the concepts *relative science value* and *total value of ownership*, including two potential ways to estimate total value of ownership over the lifetime of a system,
- the introduction of a way to reason about compute system technology beyond just cost,
- a number of case studies that show practical trade-offs between cost and value in radio astronomy.

Although we present a number of equations in this paper, it is not our intention that these are used to generate a numeric merit value for a particular system or technology. Rather, they are intended to illustrate which components contribute to the cost and merit of a system and as a starting point for a more detailed discussion on the relative value of various compute systems. With these components, and some examples of cost and value past and present in this paper in mind, system designers and architects have the tools needed to better balance their designs, and evaluate their design choices within this framework.

The intent of this paper is to show how compute system related design and procurement decisions in data-intensive science projects should be weighed and valued. By using both total cost and science value as a driver, the science output per invested Euro is maximised. While the general concepts discussed in this paper are known in systems engineering, we hope to introduce them to a broader audience of scientific decision makers, principal investigators, and system architects and designers.

2. Compute systems for large-scale science

The study of Physics, in particular Astrophysics, has relied on state-of-the-art computer science and high-performance computing. Modern aperture synthesis radio astronomy in particular was made possible by the development of the Fast Fourier Transform (FFT) (Cooley and Tukey, 1965)² and computers fast and cheap enough to use them at scale. For example, the One-Mile Telescope, built at the Mullard Radio Astronomy Observatory, Cambridge, in 1964, relied on the computing advances of the EDSAC II and TITAN computers, as is illustrated in our Case Studies in Section 7.1. This telescope, and others, like the Half-Mile Telescope at Cambridge and the Westerbork Synthesis Radio Telescope in the Netherlands, depended on the abundant and increasingly cheap computation available to develop the new scientific technique of aperture synthesis, which unlocked new science and ultimately won a Nobel Prize.

More recently, the range of applications that benefit from large-scale computing has increased dramatically with the rise of Data Science, and the ease with which high-performance (if not world-leading) compute infrastructures have become available via Cloud Computing. This paper is thus presented at a timely moment, to provide decision makers, principal investigators and designers of new compute systems and applications with a framework to help evaluate and guide their design choices.

² Ryle in his Nobel lecture credits Dr. David Wheeler with the invention of the FFT in 1959 (Ryle, 1992).

3. On relative science value

In the previous section we have argued that modern data-intensive science relies heavily on computing. Given the high cost of such resources, there is an obvious desire to maximise their usefulness, or minimise their cost. We introduce a system's *Relative science value*, defined as its value per invested Euro over its lifetime, as a measure for the merit of a system over its lifetime. The definition of *value* will be discussed in Section 4.

The computational systems supporting modern data-intensive science are often a complex collection of hardware, middleware and software. Quantifying the cost and relative value of such a complex integrated system is nearly impossible. To start our exploration we will focus on one of the more tangible components: hardware.

By first exploring ways to quantify hardware cost and value, we reduce the complexity of the system under investigation without impacting the value of the analysis. In Section 7 we show that the methodical hardware-based analysis can be applied more broadly, as similar considerations can be used to evaluate other system costs, such as software development, maintenance and power consumption.

The relative usefulness of a hardware system, its relative science value (M_S), depends on its total aggregate value accrued over time (total value of ownership, *TVO*) and aggregate cost over the lifetime of the system (total cost of ownership, *TCO*):

$$M_S = \frac{TVO}{TCO} \quad (1)$$

Total Cost of Ownership is a well known concept, both as a tool to inform purchasing decisions in general (Ellram, 1995), and in computer science. In this paper we give our own definition of the Total Cost of Ownership of a system. We introduce the generic concept of Total Value of Ownership in this paper.

From Eq. (1) it is obvious that there are two ways to maximise the relative science merit of a compute system: reduce Total Cost of Ownership, or increase Total Value of Ownership of a system. In practice, a carefully considered combination of the two is likely to produce the optimal result. Obviously, total cumulative value *TVO* is not easy to quantify, and we note that the time over which value is accumulated may extend well beyond the lifetime of the system. In the next sections we will explore the components that make up *TVO* and *TCO*.

4. Total value of ownership

Whereas the concept of Total Cost of Ownership is well known and established, the same cannot be said for its value counterpart; we shall therefore introduce this first. In economic terms, we are interested in the return on investment, which we will refer to as Total Value of Ownership (*TVO*) in this paper, to contrast to Total Cost of Ownership. While this is an essential question to ask during the definition phase of a project, the answer is seldom easy to quantify. The success of science projects is generally measured in the importance of its scientific results, often expressed in the number of published peer-reviewed papers produced. However, from a system design perspective, it is attractive to use a more easily measured metric, such as compute power, throughput or storage capacity, to describe the value of a system. While such metrics are convenient and may be useful in their own right, we argue that these do not necessarily provide an accurate reflection of how the system will be used. Furthermore, these do not necessarily take computational efficiency, scientific impact, or average required capacity per accepted paper into account. In this section we propose two measures for a system's *TVO* that are designed to more accurately reflect the actual scientific usefulness of a

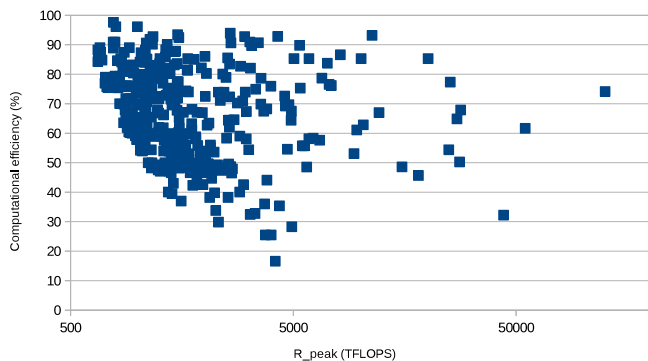


Fig. 1. HPL computational efficiency in the Top500 (November 2017).

system: total lifetime computational value (V_c), and total lifetime scientific value (V_s). While we provide equations, these are not designed to be used to model TVO; but rather to capture the relationship between some of the various elements that define system value.

Total performance, computational or otherwise, of a system can be a useful measure for the value of a (hardware) system. However, even this can be difficult to quantify beforehand. Whereas peak computational performance is relatively easy to determine, often only a small fraction of this can be achieved in practice. The same can be said for other metrics like peak network and storage performance. The fraction of the computational resources that can effectively be used by an application is determined by its computational efficiency. A discussion on the factors that impact computational efficiency is beyond the scope of this paper, but we note that these factors should be foremost in the mind of a hardware system architect. To illustrate this point, we look at the yearly Top500 list of the fastest supercomputer in the world for the HPL benchmark.³ Computational efficiencies of these systems, shown in Fig. 1, range from 15.6% to 97.6%, which shows that the impact of unexpectedly low computational efficiency may be catastrophic.

By taking into account the target applications for a specific system, we introduce an estimate for its total lifetime computational value (V_c , in FLOP), as shown in Eq. (2).

$$V_c = T_l A_o \sum_{p=0}^P (f_p R_{max,p}), \text{ with } \sum_{p=0}^P f_p \leq 1 \quad (2)$$

Here, we take the total lifetime of the system, T_l , and its availability as a fraction of total lifetime, operational availability (A_o), to get the effective time the system is usefully available over its lifetime. For each application p , its maximum achieved performance on the target system ($R_{max,p}$), and the fraction of operational time it is expected to be run (f_p) are taken to get a value for the average maximum achieved performance over all applications to be run on the system. Combined, these two components make up the system's total lifetime computational value.

Eq. (2) thus takes computational efficiency into account over all target applications, and considers both system lifetime and operational availability. Similar analyses could be done for other performance metrics, such as network bandwidth and storage system performance. We do not measure average performance of applications, rather we determine the total effective performance over the lifetime of the system. However, the eventual goal of a compute system is not the delivery of capacity per se, but rather to facilitate science. A discussion on appropriate

metrics for scientific output is well out of scope for this paper, for the purpose of the discussion in this paper we use *scientific publication* as a placeholder. This is most easily measured in peer-reviewed journal or conference publications; however, one may also consider monographs and PhD theses, or even awards (see Section 7.1).

To illustrate these points, we introduce a system's *Total lifetime scientific value*, which in our example is based on its previously introduced total lifetime computational value. Since not all science requires the same amount of resources, processing power or other, per scientific publication, we add the *average computational resource required per scientific publication*. An appropriate *impact factor*,⁴ which is not necessarily the same as a journal impact factor, may be added to differentiate potential Nobel prize winning research from more generic projects. Notably, this impact factor may be highly time sensitive, in the sense that groundbreaking projects generally have very high impact factors (see Section 7.2 for an example). We note that these two factors may be subjective, highly sensitive, and may have significant political implications.

$$V_s = T_l A_o \sum_{p=0}^P \left(f_p \frac{R_{max,p}}{C_{cpp,p}} I_p \right), \text{ with } \sum_{p=0}^P f_p \leq 1 \quad (3)$$

Total lifetime scientific value V_s is defined in Eq. (3) by the maximum achieved performance of the application associated with science case p on the system under investigation $R_{max,p}$, divided by the average amount of resources required per scientific publication for that science case $C_{cpp,p}$. This results in the number of scientific publications per unit of time for that science case and system. Multiplied by some impact factor per science case, I_p , and summed over all science cases targeted by the system P and normalised using the fraction of time each application is expected to consume (f_p), gives us a measure for scientific impact per measure of time for that system. Multiplying that by the total lifetime of the system T_l and the fraction of that time the system is actually available (operational availability A_o) gives us the total scientific value of a system, in a unitless *scientific impact*. For convenience we use computational resources, in floating point operations (FLOP), as a measure for resources required per scientific publication in this model, but other metrics (such as bandwidth, storage capacity, etc.), or a combination of such metrics, may be used instead.

The two value measures introduced in this section are by no means the only ones that can be defined. They are intended to start the discussion and offer an initial indication of the processes and thinking involved. Characterising the performance of a compute system in a single number is notoriously difficult, which has been studied in some detail. Previous work suggested the use of harmonic means of runtime of a number representative benchmarks to express the useful performance of a computer (Smith, 1988), which expresses performance in terms of the total runtime of a set of benchmarks. While benchmarks certainly have their place, and runtime is an appropriate measure for performance of a system, this is not necessarily equivalent to value. However, we note that the V_c is equivalent to the weighted harmonic mean suggested by Smith et al. multiplied by $T_l A_o$. Essentially, instead of computing average performance, we focus on aggregate performance over the effective lifetime of the system, taking system lifetime and operational availability into account.

⁴ We are aware that constructing a useful impact factor has many pitfalls. See Lariviere and Sugimoto (2018) for a discussion about journal impact factors as an example.

³ www.top500.org.

5. Total cost of ownership

Having looked at various ways to define the value potential of a system, we now turn to more familiar ground: cost. The aggregate cost of a system over its lifetime is usually referred to as its Total Cost of Ownership. While the definition of TCO is relatively easy to give, calculating it a priori may not be as simple, in particular in large-scale science installations. The lifetime of a particular system may be unpredictable, and the often non-conventional use of such systems may lead to unexpectedly large operational costs. Furthermore, complex and highly integrated systems make for difficult deployment and integration, which is hard to plan and budget for. Having said that, TCO can be defined as a combination of capital investment (C_{cap}), engineering cost (C_{eng} , often called non-recurring expense, or NRE), installation, deployment and integration cost (C_{int}), development cost (C_{dev}), recurring operational cost (C_{ops}) over the lifetime of the system (T_l) and miscellaneous costs not covered elsewhere C_{misc} , as shown in Eq. (4).

$$TCO = C_{cap} + C_{eng} + C_{int} + C_{dev} + \sum_{t=0}^{T_l} C_{ops} + C_{misc} \quad (4)$$

The one time investment to acquire a system is referred to as its capital cost, C_{cap} . This includes all readily available hardware required to install and commission the system. Capital cost is usually either capped, or relatively easy to estimate. We note, however, that even capital cost becomes highly uncertain when predicted several years in advance, due to fast moving markets and uncertain performance characteristics and pricing of newly developed components. Models often resort to extrapolation from existing systems using some form of Moore's law scaling to estimate future cost and performance (see for instance the SDP costing for the SKA telescope (Alexander et al., 2013)). While this has historically been somewhat accurate, the demise of Dennard scaling (Dennard et al., 1974) around 2005 has made modelling much more complicated. This uncertainty is exacerbated by an erratic market that is increasingly dominated by single players without significant competition.

When a system requires engineering investment in order to be usefully employed, this is engineering cost, C_{eng} . This may involve custom cooling solutions, or other non-standard equipment specific to the system (see for an example the LOFAR GPU-based correlator and beamformer (Broekema et al., 2018)). Costs associated with certification of a custom solution may also be considered engineering cost. General purpose systems generally have no or very little engineering cost, but in more specialised systems this may be a significant cost driver.

Any investment needed to integrate and commission the system into an existing infrastructure is captured in integration and commissioning cost, C_{int} . Note that in software systems, especially if the source code of this software is available, integration, commissioning and development may be closely related.

It is unlikely that the application software of a science instrument or experiment remains static over the lifetime of the instrument. Part of the software evolution will be to add additional functionality or implement advances in algorithmic or scientific understanding of the problem. Another part of this development will be to adapt existing code to run (efficiently) on a newly installed platform. The cost of this particular development effort is the development costs of that system (C_{dev}). Such costs may be small (e.g., porting code to a newer system with the same or a similar architecture), or very large, for example, porting functionality from a CPU cluster to a GPU-based system, as was done for LOFAR correlator (Broekema et al., 2018). These costs may be difficult to predict during the design phase of a long-lived

Table 2

The offered solutions, with detailed cost, lifetime and availability information.

	Cheap	Inefficient	Ops	Custom	Specialised
C_{cap} (€)	250.000	350.000	350.000	300.000	400.000
C_{eng} (€)	–	–	25.000	–	25.000
C_{int} (€)	25.000	–	25.000	–	25.000
C_{dev} (€)	750.000	600.000	1.250.000	1.250.000	1.000.000
C_{ops}/yr (€)	50.000	25.000	75.000	25.000	25.000
C_{misc} (€)	25.000	25.000	25.000	25.000	25.000
T_l (yr)	5	5	5	5	5
A_o	0,9	0,95	0,85	0,95	0,95
TCO (€)	1.300.000	1.100.000	2.050.000	1.700.000	1.600.000

instrument, which, in the LOFAR case, was a decade earlier. It is likely that not all development effort is expended before the system is deployed, and C_{dev} may extend significantly into the lifetime of the system. Furthermore, added development effort may have a significant impact on computational efficiency, with a corresponding effect on Total Value of Ownership. There is a direct coupling between the development costs, and thus Total Cost of Ownership, and Total Value of Ownership. Conversely, if a system performs well enough, there is no need to expend more development effort to improve performance, unless this opens opportunities for, for instance, additional science cases.

Whereas all previously mentioned costs, with the exception of Development costs, are expended before the system becomes operational, Operational cost (C_{ops}) is a recurring line-item during the lifetime of the system. This includes costs associated with energy consumed, infrastructure cost (i.e. rack space, network connectivity, both physical links and bandwidth, heat dissipation, etc.), maintenance and system administration. We have simplified our model by using a single operational cost component; reality is often more complex, especially in a hosted environment where the components mentioned above are provided by different entities or organisations. While we have opted to keep operational cost in our model flat over the lifetime of the system, this is again a simplification. Operational cost in the initial phase of the system may be higher both due to early failure of hardware and staff unfamiliarity and training. Near the end of the operational lifespan of the system, often after four or five years in general purpose computing, an increase in hardware failures may be observed, which may increase operational cost, depending on the chosen service model. Furthermore, operational cost may depend on inherently volatile pricing of, for instance, electricity. Energy costs are often estimated using the previously mentioned extrapolation using Moore's law scaling, while staffing levels and costs may be based on industry standard fractions of FTE per rack or PetaByte (Graser and Taylor, 2015).

Finally, staff costs not included in the components above, such as those required to secure funding, acquire the system (e.g. writing tender documentation and evaluating responses) and to decommission the system after its useful lifetime, as well project management and support other than system administration, are included in miscellaneous cost (C_{misc}).

The remainder of this paper takes the concepts introduced, and shows, using artificial and real-world case studies taken from radio astronomy past and present, the value of this structured approach to compute system design.

6. A synthetic instructive example

In the previous sections we identified a metric that we can optimise for: total relative science value as defined in Eq. (1), M_s , but its definition is (deliberately) ambiguous. While it is not our intention to advocate numeric values for the total relative science values for eScience technologies, we can use the equations introduced above to identify ways to optimise their usefulness.

Table 3
Application characteristics and performance per offered solution.

	f_p	C_{cpp}	I_p	Cheap	Inefficient	Ops	Custom	Specialised
A	0,04	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
B	0,08	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
C	0,02	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
D	0,02	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
E	0,40	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
F	0,11	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
G	0,07	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
H	0,08	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$
I	0,02	$1 \cdot 10^4$	100	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$10 \cdot 10^8$
J	0,16	$1 \cdot 10^4$	5	$2 \cdot 10^8$	$1 \cdot 10^8$	$5 \cdot 10^8$	$4 \cdot 10^8$	$2,5 \cdot 10^8$

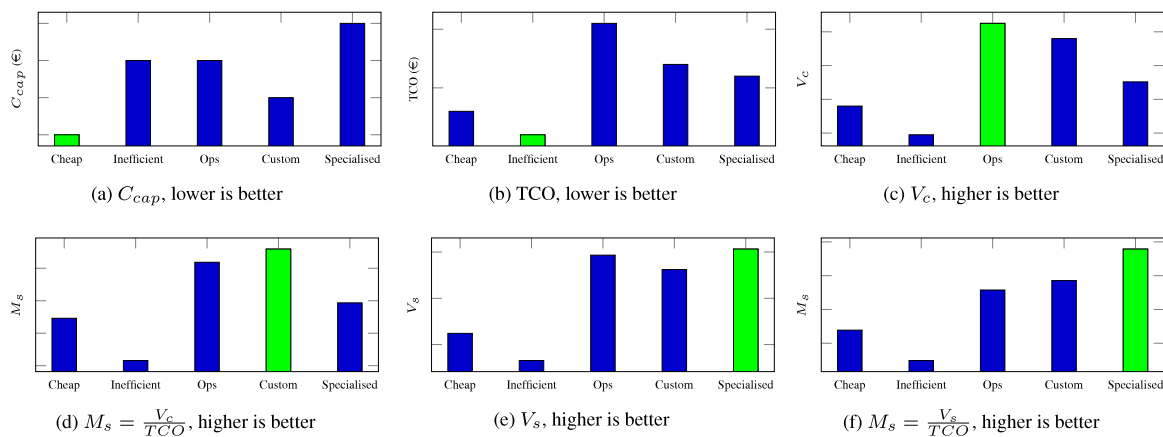


Fig. 2. The offers evaluated against six cost and value measures. The superior offers for each measure are shown in green. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

In this section we illustrate the value of the proposed methodology using a thought experiment. We have constructed an example that is obviously manipulated to show the desired results. However, using this example we show that, depending on the value measure selected, any of the offered solutions can be judged superior to the others.

Table 2 describes responses to a hypothetical request for tender for the replacement of a key computer system. A set of ten key applications was identified that cover the lifetime of this system, and performance of each system for these applications was measured, as shown in Table 3.

Each of these offers were evaluated using the model introduced in this paper, the results of which are shown in Fig. 2. For each value measure, the superior solution is shown in green.⁵ While the offers are fictional and the use-case is obviously constructed, it is clear that, depending on the chosen selection criterion, a different solution wins, highlighting both the power and importance of the concept introduced in this paper. More importantly, this example shows the dangers of selecting the wrong value measure for convenience or not carefully considering all possible components that make up the selected value measure.

In Section 1 we postulate that the useful (scientific) output of the system per invested Euro is the most useful value metric of a system. Not using such a metric, and instead focusing solely on total cost of ownership, would, in this example, lead to the selection of the far inferior *Inefficient* solution.

7. Case studies

To further illustrate the value of the conceptual model introduced in this paper, three radio astronomy use cases will be

discussed: the use of the TITAN computer for one of the first operational radio interferometers, the LOFAR array and the SKA Science Data Processor. We also highlight the variability of value over the lifetime of a compute system using the performance impact of a recent hardware vulnerability as an example.

7.1. The TITAN computer and the Mullard Radio Astronomy Observatory

The One-Mile Telescope, sited at the Mullard Radio Astronomy Observatory (MRAO) near Cambridge, was an early aperture synthesis telescope, and was the first designed to use the Earth rotation aperture synthesis technique. It was conceived when the EDSAC II computer at Cambridge University was in operation, and was completed in 1964, as the TITAN computer came online. TITAN was then used by the One-Mile, the Half-Mile and Interplanetary Scintillation Array (IPSA) telescopes, until TITAN was decommissioned in 1973. The One-Mile was explicitly designed to use the improved computing resources provided by TITAN, first to provide the control tapes for the telescope and then using the Fast Fourier Transform (FFT) to power the data analysis (Elsmore et al., 1966; Ryle and Neville, 1962). As Wilkes recalls:

One day, Ryle came to me to say that he was planning the erection of a much larger telescope and to ask whether the Mathematical Laboratory could undertake to provide the computing support required. Wilkes (1985), p.193

TITAN was a ground-breaking computer itself, with hardware procured from the Ferranti company, and software developed by staff at the University of Cambridge, mostly from the Mathematical Laboratory. The Mathematical Laboratory was, unusually for the time, already running as an effective computing service,

⁵ All underlying data and analysis used in this paper are available here: <https://doi.org/10.5281/zenodo.2270842>.

where users applied for time with their projects (as is common with HPC resources today) (Nash, 1990; Wilkes, 1985). This differed from companies such as Ferranti and IBM, which were producing computers for the commercial market, and other universities, which were producing computers primarily as a way of investigating computers themselves (the purpose of instruments such as the Manchester Baby and CSIRAC), without explicit support for scientific research (Croarken, 1990).

Although the University wished to buy a new computer to replace EDSAC II, it did not have a large capital budget available. Thus they bought a heavily-discounted Ferranti Atlas (usual cost £2 million; price actually paid: £350,000 (approximately £6–7 million today⁶) with an additional £75,000 for a large disk store (Ahmed, 2013)). However, the University now had to spend a lot of money on salaries to develop the software, but this cost was not explicitly tracked by the University, and their decisions were made purely on the C_{cap} .

The performance of TITAN, combined with David Wheeler's FFT algorithm (Ryle, 1992), allowed TITAN to do the calculations necessary for the first Earth-rotation aperture synthesis observations with the One-Mile telescope, and then to produce the first maps of the radio sky (Ryle and Neville, 1962). It was also used to support IPSA, which was used by Dame Jocelyn Bell Burnell to discover the first pulsar.

These scientific breakthroughs, backed by TITAN, won Tony Hewish and Sir Martin Ryle their joint Nobel prize for their innovative telescope design work (Ryle et al., 1964). Furthermore, at least 30 PhD theses using the One-Mile, and the subsequent Half-Mile and IPSA, used TITAN-processed data, or used TITAN for theoretical modelling.⁷ It is unfortunately not possible to track all the papers that were produced with TITAN, as it was not comprehensively tracked at the time and not all authors note their use of TITAN. Therefore there are aspects of TITAN's value that are not captured.

Nevertheless, TITAN delivered exceptional TVO extending well beyond its lifetime. It was not only used in radio astronomy, although radio astronomy made unique use of its capabilities, but also in computer science (applications included one-way functions for storing passwords, timesharing systems, computer language research, early version control systems (Ahmed, 2013)), crystallography (another field that used the FFT), statistics, Computer Aided Design, agronomy, and quantitative economic methods (for which one TITAN user, Sir Richard Stone, won the Nobel Prize for Economics) (Stone, 1966)⁸ (Needham, 1992).

Many of the people who designed, programmed for, and used, TITAN were or became leaders of their fields, bringing rewards (both financial and reputational) to their institutions in the subsequent decades; thus TITAN provided a TVO that far outweighed cost of purchasing, developing, and running the system. There is a significant "long tail" to TITAN's value, exemplified by Dame Jocelyn Bell Burnell's receipt of the Royal Society Royal Medal in 2015, and the Special Breakthrough Prize in Physics (2018), both of which specifically cite her work on pulsars. To illustrate the disparity between the lifetime of the TITAN system and its value, we have plotted major prizes won by TITAN users in radio astronomy, as compared to TITAN's lifespan, in Fig. 3.

Awards were not confined to the radio astronomy community. Eleven TITAN users have been elected Members of the Royal




Fig. 3. Awards given to TITAN radio astronomy users over time. The blue bar indicates when TITAN was active.

Society.⁹ TITAN users have won many other major awards in their fields, including: the Royal Statistical Society Guy Medal in Silver,¹¹ the Gold Medal of the Royal Astronomical Society,¹² the London Mathematical Society De Morgan Medal,¹³ the Faraday Medal,¹⁴ the IEEE John von Neumann medal¹⁵ and the Karl G. Jansky Lectureship, which "is an honor established by the trustees of Associated Universities, Inc., to recognise outstanding contributions to the advancement of radio astronomy".¹⁶ The precise role of TITAN in these awards is difficult to quantify; however, having TITAN available clearly provided important support and enablement for people at all stages of their careers – precisely the purpose of a scientific computing resource. Moreover, this lists only the very highest achievers amongst TITAN users; there are shallower network effects from the existence of TITAN, which are even harder to account for, but which indicate that resources such as TITAN are vital for the scientific community. Although a significant investment, both capital, as well as engineering and development, was required, TITAN's ten-year lifespan and high-impact and long-lasting contributions make its relative science value exceptional. Even if the full list price of the hardware had been paid, the capital outlay would still have been justified by its scientific success, which far outshone other contemporaneous systems.

7.2. LOFAR

LOFAR, the LOW Frequency ARray (van Haarlem et al., 2013), is a modern low-frequency large-scale distributed radio telescope in the Netherlands, with international stations in various European countries. The concept and design of the LOFAR telescope, which started in the late 1990s, is a study in trading off value and cost. A number of early papers discussing the telescope concept and initial design (Bregman, 1999, 2000), as well as some retrospective analysis of the design considerations (Bregman, 2012), make this a particularly interesting instrument to study.

As discussed above, modern radio interferometry was made possible by the availability of abundant and affordable compute resources. In LOFAR, this concept is taken even further by replacing a small number of large parabolic reflectors with many simple, cheap and omni-directional dipole antennas and software-based digital beamforming. Essentially, many simple

⁹ <https://royalsociety.org/>

¹⁰ Frank Yates, Donald Lynden-Bell, David George Kendall, Maurice Wilkes, Sir Martin Ryle, Peter Swinnerton-Dyer, Malcolm Longair, Brian Pippard, Roger Needham, John Baldwin and Dame Jocelyn Bell Burnell. "The Royal Society is a Fellowship of many of the world's most eminent scientists and is the oldest scientific academy in continuous existence", and members must have made "a substantial contribution to the improvement of natural knowledge, including mathematics, engineering science and medical science".

¹¹ Won by Geogle Kendall and MJR Healy.

¹² Won by Donald Lynden-Bell, Sir Martin Ryle, and Jeremiah P. Ostriker.

¹³ Won by D. G. Kendall.

¹⁴ Won by Ryle, Maurice Wilkes, and Roger Needham.

¹⁵ Won by Wilkes.

¹⁶ Awarded to Bernie Fanaroff and Dame Jocelyn Bell Burnell, both of whom used TITAN during their PhDs.

⁶ <http://www.nationalarchives.gov.uk/currency-converter/>

⁷ One of us (VA) checked the archived PhD theses of the Astrophysics Group, Cavendish Laboratory, University of Cambridge. TITAN may have been used for PhD theses in other groups; however, it is difficult to locate all of these 40 years later.

⁸ A copy of this work is held in the Library of the Department of Computer Science and Technology, University of Cambridge, classmark V75-14.

antennas are combined, by coherent addition, into a single virtual receiver. Early design concepts for this low-frequency array that could act both as a technology demonstrator for the future Square Kilometre Array, as well as scientifically open a relatively unexplored frequency range, identify a “processing window of opportunity”. This early concept predicted that, while computational cost for the processing required for this low-frequency array was at the time infeasibly large, it would become affordable, assuming Moore’s law continued to apply, after 2003.

In further work, instrument sensitivity was defined as the key value parameter (and thus a measure for the TVO of the instrument) for the design trade-offs in this instrument (Bregman, 2004), although other value measures such as survey speed and resolution were also taken into account. In order to achieve optimal performance over cost, all main constituents of the complete LOFAR system were designed to have a similar marginal performance over cost ratio.

This analysis shows that both TVO and TCO for the LOFAR telescope in general, and the digital processing systems in particular, were carefully considered early on in the conceptual design phase of the instrument. A clear choice was made to use sensitivity over other technical or scientific metrics, such as survey speed or resolution, as a measure for the total value of the instrument. We note that this implicitly assumes this technical measure translates to scientific value. Regardless of this technical measure, suitability for a small number of key science projects was also a key design consideration in the development of LOFAR. Furthermore, the cost of the digital processing system was analysed, and, more importantly, judged to become affordable at some point in the mid future. This realisation allowed development of the instrument, and its associated software infrastructure, to start before the required compute capacity became financially feasible.

Since its opening in June 2010, one measure of LOFAR’s science value, the number of peer-reviewed scientific publications using LOFAR produced data, has been monitored.¹⁷ A different way to express the value, or in this case more accurately the return on investment of a science instrument, is to evaluate how much of the invested money is reinvested in the local (national) economy. A Dutch research institute that specialises in research on the impact of science, Rathenau, recently studied the LOFAR telescope and the Dutch contribution to three other major science instruments: CERN,¹⁸ ESRF¹⁹ and ITER.²⁰ They defined a return coefficient (R) as the capital reinvested in the national economy, divided by the Dutch contribution in that instrument. The results, published in Dutch (Tai et al., 2019), are summarised in Table 4. While it is difficult to compare four completely different instruments, this work shows that the financial return of the LOFAR telescope for the Dutch economy has been excellent. This is due to the fact that most, if not all, of the IP was developed in the Netherlands, and therefore production of those components, even for international stations, is likely to occur there as well.

The hierarchical and modular nature of the LOFAR system has allowed several dedicated systems to be added to the telescope to increase its scientific value at modest cost. While some, like Dagnet (described in Section 7.2.3), were just plug-in systems that required little to no additional engineering to add to LOFAR, others, like AARTFAAC (see Section 7.2.2), require raw antenna data not available in standard LOFAR observation modes. We will explore the cost and value considerations of some of these components in the following sections.

7.2.1. LOFAR correlator and beamformer systems

A key signal processing component of the instrument, the LOFAR correlator and beamformer, and specifically its hardware evolution, is relatively well described. This part of the instrument is algorithmically simple and the required functionality is fairly constant. Therefore, for this specific example, cost (with all its different components), operational availability, and lifespan mostly determine the relative science value of the correlator and beamformer. Early concepts for the LOFAR central processor show a 1600 node hybrid cluster compute system that uses conventional processors and dataflow co-processors to process the data (de Vos et al., 2001; van der Schaaf et al., 2003). While feasible, the considerable size of this compute concept meant that a bespoke supercomputer was a viable and, more importantly, cost-effective alternative. In 2003, an IBM Blue Gene/L, briefly the fastest supercomputer in Europe, was installed as the central correlator and beamformer for LOFAR (Romein et al., 2006). This was upgraded to a much smaller, IBM Blue Gene/P in 2008 (Romein et al., 2010), that was not only more powerful, but also considerably more energy efficient. Whereas the total lifetime computational and scientific value of this new system was similar, its reduced operational costs, as well as improved software environment made its relative science value considerably higher than the previous Blue Gene/L. However, supercomputers are inherently expensive, so research into more cost-effective solutions continued (van der Schaaf and Overeem, 2004; van Nieuwpoort and Romein, 2010). This eventually resulted in the procurement and commissioning of a much smaller and more affordable GPU-based correlator and beamformer platform, Cobalt (Broekema et al., 2018). A more capable second generation of this system, Cobalt 2.0, started operations in 2019.²¹

The timeline of the LOFAR correlator and the construction of the instrument as a whole is interesting to study. As mentioned above, the telescope was opened in 2010 and at that time the initial Blue Gene/L correlator and beamformer has already been replaced. While it is not accurate to say Blue Gene/L was never used in production as the LOFAR correlator and beamformer, it is clear that it was procured and installed early. Arguably, its cost was considerable (although the actual investment was never made public), and its value limited. However, the strategic alliance and collaboration agreement between ASTRON and IBM was an important consideration in securing sufficient construction funding for LOFAR. Furthermore, spare computational capacity was made available to other scientific users. Therefore, while the total lifetime scientific value of the Blue Gene/L correlator for the LOFAR telescope was low, its general value for the LOFAR telescope was extremely high and its total lifetime scientific value for the wider community was comparable to other high-performance computing systems. Nevertheless, the Blue Gene/L system was never used to its full potential in the LOFAR telescope, and even the Blue Gene/P system was significantly under-utilised for most of its lifetime. These systems did however provide extremely valuable experience that was essential to the success of Cobalt and was used to great effect in the hardware design of the SKA Science Data Processor. Whether this was worth the significant initial capital investment is beyond the scope of this paper.

When the LOFAR Blue Gene/P was nearing the end of its service life, a feasibility study into possible upgrades was undertaken (Holtjes, 2012). Four drop-in replacement options (Blue Gene/Q, an FPGA-based Uniboard system, a CPU-based cluster with GPU accelerators, and a CPU-based cluster) were evaluated for risk, development effort, cost, power consumption and scalability. It is clear from these selected criteria that various cost

¹⁷ <https://old.astron.nl/radio-observatory/lofar-science/lofar-papers/lofar-papers>.

¹⁸ <https://home.cern/>.

¹⁹ <https://www.esrf.eu/>.

²⁰ <https://www.iter.org/>.

²¹ <https://old.astron.nl/cobalt20-sets-stage-fully-multitasking-lofar>.

Table 4

Return coefficients for the Dutch economy for four large scale science infrastructure projects.

Source: Rathenau institute (Tai et al., 2019).

	CERN	ESRF	ITER		LOFAR	
	Average per year	Average per year	2008–2017 construction incl. grants	2008–2015 construction grants only	2004–2013 construction	2014–2017 operations per year
Total investment	1,104M CHF	90M €	6,120M €	4,581M €	92M €	4.3M €
Total Dutch investment	50,9M CHF	2.7M €	161M €	120M €	81.2M €	3.4M €
Dutch contribution	4.61%	2.97%	2.63%	2.63%	88.3%	77.8%
Total expenditure	343M CHF	57M €	4,330M €	101M €	92 €	4.3M €
To the Netherlands	8.7M CHF	0.58M €	7.9M €	4.3M €	89.2 €	4.1M €
Dutch ROI	2.54%	1.01%	0.07%	4.18%	97%	96.5%
Return coefficient	0.55	0.34	0.07	1.59	1.10	1.24

components were carefully considered, while value was expected to be equal among the contenders considering any new system was expected to be replicated the functionality of the existing Blue Gene/P based correlator and beamformer. A cluster with GPU accelerators was judged to be the most cost-effective solution, based on low cost and power consumption, good scalability, and relatively little development effort required. By extension this was therefore also the option with the highest relative science value, and selected for implementation as the Cobalt correlator and beamformer (Broekema et al., 2018).

7.2.2. AARTFAAC

While the LOFAR correlator and beamformer described above are integral parts of the original LOFAR design, AARTFAAC is an add-on system that was added to increase functionality and enable additional science cases. The Amsterdam–ASTRON Radio Transients Facility and Analysis Center (AARTFAAC) system (Prasad et al., 2016) is a real-time all-sky transient detection system. Data from a subset of LOFAR antennas is duplicated during normal LOFAR operations and processed independently into all-sky images of the low-frequency radio sky that can subsequently be monitored for bright transient events. This is a significant advance over the capabilities of the original LOFAR system, which was only possible due to investments made early in the LOFAR project to over-provision both the bandwidth of the LOFAR station ring network and the LOFAR Wide Area Network. For this specific addition, a custom shim was added the station data transport ring: Uniboard-RSP Interface boards. These duplicate raw antenna data, normally beamformed in RSP boards, to AARTFAAC Uniboards. Furthermore, the FPGA firmware on the LOFAR core stations that take part in the AARTFAAC system had to be modified to generate the additional AARTFAAC packets. Data from the AARTFAAC system is transported to dedicated processing nodes located in the same central processing facility as the LOFAR correlator and beamformer, sharing spare network capacity.

While AARTFAAC adds undoubtable value to the LOFAR telescope, its addition required significant additional engineering and manufacturing. In particular the additional firmware programming requires the use of scarce resources that are generally overcommitted. We will not discuss the relative merits of this addition over others, or whether the investment was valuable or not. However, we do note that additional investment in the development of data spigots at the LOFAR station during construction would have made the development of AARTFAAC much cheaper and easier. This was considered during design, but technology had not progressed sufficiently; the additional cost would have been significant and the idea was shelved.

7.2.3. DRAGNET

Whereas AARTFAAC is a real-time transient monitor that operates in UV-space, the DRAGNET cluster (Bassa et al., 2017a) is a non-real-time pulsar and transient search system that operates in the time domain. It takes beamformed data from the Cobalt correlator and beamformer and uses blind coherent de-dispersion to identify fast transients and millisecond pulsars. This system has demonstrated its value by the discovery of the second fastest-spinning pulsar to date, and one of the first at such low observing frequencies (Bassa et al., 2017b).

The DRAGNET system consists of 23 nodes, each of which has 4 NVIDIA Titan X GPUs that provide the bulk of the processing capacity. Its source data is produced by the LOFAR Correlator and Beamformer, Cobalt. Data is stored locally and processed non-real-time, resulting in a pulsar and/or transient candidate list for further analysis.

Since DRAGNET uses a standard LOFAR data product as input, only limited modifications were necessary to integrate the system into the LOFAR telescope. The only major investment, apart from the cluster and dedicated software for DRAGNET itself, was the integration of the system into the LOFAR monitoring and control system. DRAGNET has added significant capability to the LOFAR telescope: the ability to search for extremely fast-spinning pulsars, and a way to detect fast transient events that would be missed by the original LOFAR telescope. This adds significant additional value to the instrument, since it allows new science cases to be explored. The majority of the additional investment was in the actual cluster and the software development needed to process the data, with limited investment needed to modify the existing system.

7.2.4. International LOFAR stations

International LOFAR stations are not just valuable parts of the International LOFAR Telescope (ILT), these can also operate in *local or standalone* mode. In this mode, station data is not sent to the central LOFAR correlator and beamformer, but instead redirected to a local system and can thus act as a fully functional telescope in its own right. The comparatively small size of these stations, and the low observation frequency, make them relatively unsuited for imaging observations, so most effort has gone into local transient and pulsar search. The ARTEMIS backend (Armour et al., 2012) was developed as a real-time GPU accelerated suite of software to search for these events in data from modern radio telescopes. Four international stations are equipped with such systems (Serylak et al., 2012).

Changing an international LOFAR station to stand-alone mode is, from a high level, as easy as changing destination IP number and MAC address of the receiving nodes. The ability to use these international stations in this mode can be partly attributed to the extensive use of standardised protocols and interfaces, as well as the modular nature of the LOFAR telescope. This means that LOFAR is potentially a large collection of independent instruments.

One international station, the French station near Nançay, differs significantly from any other antenna field in the LOFAR instrument. Apart from the low- and high-band antennas as in every international station, an unused third analog data path in the LOFAR station hardware is used to add a cluster of 96 mini-arrays, each of which consists of 19 antennas sensitive from 10 to 87 MHz (Zarka et al., 2015). The resulting giant extension of LOFAR, NenuFAR, while not as large as the LOFAR telescope, adds a similar number of low-band antennas to the instrument as all other stations combined (1938 vs. ~2700). In stand-alone mode, NenuFAR, currently under construction and accepting early science proposals,²² intends to support a wide range of data products, very similar to those produced by the LOFAR telescope. This shows that NenuFAR is a powerful instrument itself, especially for pulsar and radio transient science. A dedicated correlator and beamformer, based on the newly commissioned Cobalt 2.0 correlator and beamformer, is currently being installed.

This extension to the French international station was made possible by the availability of an unused analog data path in the LOFAR station hardware. This data path was intended for a third receiver type, eliminated early in the design process for cost reasons. In Dutch LOFAR stations this data path is used to connect half of the low-band antennas.

Finally, a LOFAR station was constructed in Lapland, Finland, near Kilpisjärvi, well above the arctic circle (McKay-Bukowski et al., 2014). This station, KAIRA, is not part of the International LOFAR Telescope (ILT) and not connected to the rest of the LOFAR network. Instead it is used exclusively in stand-alone mode, primarily for atmospheric imaging using reflected transmissions from a number of remote radar sites. Experiments have shown fringes on recorded data between it and the international LOFAR station at Effelsberg in Germany, proving that for exceptional experiments it is possible to add the station to the LOFAR array, albeit not in real-time.

7.2.5. Retrospective

The LOFAR concept design identified a period in time where the relatively high impact of ground-breaking radio astronomical research in a relatively unexplored frequency range, combined with dropping costs for computing, would result in an instrument with optimal relative science value. During its design and operational lifetime, the LOFAR correlator and beamformer in particular has benefited from continued development of cost-optimised solutions to improve the relative science value of an already successful and cost-effective instrument. The modular nature of the LOFAR telescope enabled the addition of additional systems to the instrument, further increasing its science value.

We note that the cost and value analysis of these additional systems was not as rigorous as that done for the original LOFAR system. While the engineering challenges of such add-on system were generally considered, the operational impact was often under-estimated and (un)availability of critical development resources lead to significant slippage in project schedules. Within ASTRON this has led to a more formal and structured application process for funding and the adoption of more rigorous systems engineering practices. For the second phase of LOFAR we are considering the establishment of a LOFAR architecture team to centralise and formalise the responsibility for the considerations on cost and value for the instrument. Modern distributed radio telescopes are, due to their inherent modular nature, exceptionally adaptable and extendable. Taking possible extensions into account during the design of a new instrument will make the addition of such extensions easier and thus cheaper.

When looking at radio telescope systems as a whole, instead of just the compute systems they rely on, scientific value is the better understood factor while the sum of all costs is often not fully appreciated. This is, at least to some degree, a result of the funding model for scientific instruments. Funding proposals are evaluated on scientific merit first, and cost second. Furthermore, costing an addition to a complex distributed sensor system, like the LOFAR telescope, is exceedingly complex and prone to overseeing non-trivial component costs.

7.3. Spectre and Meltdown: how value of an existing system may change unexpectedly

We argue in this paper that we can try to estimate the total lifetime computational value of a hardware system beforehand. However, value is not constant over time and may be impacted by external factors beyond the control of the user. In January 2018 a number of critical and widespread flaws in the hardware design of current generation processors were published (Lipp et al., 2018; Kocher et al., 2018). These unparalleled hardware vulnerabilities hit virtually every installed compute system currently in operation. While many software bugs may cause temporary performance issues, or cause delays in achieving top performance, the mitigations implemented to address these unprecedented flaws in processor design caused a completely unexpected and major reduction in performance of current systems, including otherwise well-performing systems. Due to the nature of these flaws, critical separation failures in performance-critical speculative execution, mitigation efforts in processor microcode and operating system kernel, have resulted in significant performance impacts, thus reducing the value of existing compute systems. In particular I/O heavy workloads, such as those encountered in the LOFAR correlator and beamformer, that cause large numbers of context switches are expected to see performance reduced by very significant amounts²³. For the Linux kernel, the dominant operating system in both high-performance computing, as well as distributed computing applications, these are known as Kernel Page Table Isolation (KPTI). These are kernel level fixes, that can be activated or deactivated at boot-time with a kernel boot parameter.

We illustrate the performance impact of these mitigating efforts in Fig. 4. We test three Linux kernels, one released just before the announcements mentioned above (4.13.16), one that includes the initial mitigating patches (4.14.14) and one more recent kernel (4.19.1) in which the mitigations have been in place for some months. Since a key task in the correlator and beamformer systems in LOFAR involves receiving large numbers of UDP/IP streams, we measure performance impact, and therefore the hit on value, by trying to receive as many UDP/IP packets as possible on a CPU-bound system with a 40 GbE device. Results are normalised to the performance of the oldest kernel, which, for reference, achieved around 1,65 million packets per second.

This measurement shows that the value of a system has the potential to change over time (here between 5% and 10%), and may be affected by factors and risks outside its operators' and designers' control. In this particular case, most of the performance impact may be avoided by turning off page table isolation (`nopti`) and `retpoline` (`nospectre_v2`) at boot time, at the cost of accepting that the system is trivially exploitable (which may be acceptable for a dedicated cluster behind a firewall).

²² <https://nenufar.obs-nancay.fr/en/astronomer/>.

²³ <http://www.brendangregg.com/blog/2018-02-09/kpti-kaiser-meltdown-performance.html>




Fig. 4. Maximum UDP/IP packet receive performance for three kernels, normalised to the oldest kernel. Blue shows the default configuration, green when Spectre and Meltdown v2 mitigations are turned off. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

7.4. SKA SDP

The Square Kilometre Array (SKA) is a next-generation radio telescope, currently in the design phase. It will consist of two telescopes, a low-frequency telescope (50–350 MHz) consisting of 130,000 antennas in over 500 stations in Western Australia, and 133 mid-frequency dishes (4350 MHz–14 GHz) in South Africa, which latter will be joined to the existing MeerKAT telescope. These telescopes will be constructed and deployed in phases across a 5-year period. The SKA is designed to achieve exceptional scientific value, and to enable potential Nobel Prize-winning research²⁴.

A key component of this instrument is the Science Data Processor (SDP), where instrument data, produced by specialised correlator hardware, is turned into science-ready data products, such as radio astronomy images, using high performance general-purpose compute systems (Alexander et al., 2019b). There will be data centres in Perth and Cape Town, which will host an SDP for each country, where the data will be received, processed, and transmitted for use by astronomers. The data rates from SKA will be extremely large: each telescope will output up to 3.1Tb/s from the correlator. The main function of the SDP is thus to perform a data reduction, outputting data products that are able to be used by scientists, but which are also somewhat easier (and cheaper) to store and transmit. The (in)ability of the SDP to perform this function may impact the science that can be performed by the SKA as a whole: if it takes too long to reduce the data, or the SDP cannot reduce the data by a sufficient factor, less data- or compute-intensive observations will have to be scheduled (Alexander et al., 2019b). Thus the design of the SDP is critical to the scientific value of the telescope.

In order to maximise its relative science value, the SDP will use a mix of custom-designed software components and off-the-shelf software. In order to reduce TCO, the SDP will make extensive use of existing technologies. A platform management system is envisioned to provision and organise its compute resources. Such a system allows for the automation of compute deployment, at the cost of a mild computational overhead. This saves on operator time, and allows for reliable and reproducible deployment of operating systems and other support services. The reduced operator time needed and increased reliability drive down operational costs (C_{ops}); reproducibility renders it easier and quicker (hence cheaper) to detect bugs. OpenStack, an open source platform management product in use by HPCs and data

centres, is a candidate solution for this, in part because SKA is already working with CERN on improving OpenStack technologies. The SKA will save cost of development (C_{dev}) and ongoing maintenance costs by using this off-the-shelf open source software, rather than writing their own suite of complicated software for the same purpose. The viability of this approach has already been prototyped (Taylor, 2019).

In addition, in order to improve TVO, a new suite of astronomy data processing software will be developed, focusing on a highly reusable modularised architecture. The principal idea is to create low-level software modules that can be reused by many data processing pipelines (Alexander et al., 2019a). However, rather than using existing code from existing telescopes, these modules will be newly implemented for two reasons: scalability in parallel environments and maintainability over the expected 50-year telescope lifetime.

Providing modules that can easily be connected for use in large clusters is key for the SDP, as, without taking advantage of the inherent parallelism available in a lot of astronomy data processing, it will be difficult to achieve the data throughput necessary. This modularisation not only allows designing for an embarrassingly parallel processing environment, it also permits programmers to quickly and easily provide new modules for optimised use with new hardware, and implement new algorithms for new science without rewriting other parts of the software infrastructure. This is explicitly to reduce C_{dev} , by anticipating the need to port code to new, potentially very different, architectures, in contrast to the issues LOFAR experienced, as noted in 7.2. This also allows the SDP to run on generic COTS hardware, while also allowing for future software-hardware co-design for key algorithmic components. Similarly, modularisation of code handling hardware interfaces allows for pivoting to new technology – an inevitability in a long-lived project.

Maintainability is also a key driver for writing new code: technical debt accrues in software projects over time, as programmers can end up prioritising writing code quickly, rather than writing it well, or with an eye to help reduce maintenance costs. Some of the commonly-used radio astronomy code, such as CASA, has parts that are nearly 40 years old, and which were not designed to be used in highly parallel compute systems. Thus the SKA has the opportunity to reduce its total lifetime costs by investing in new code that is designed to be more easily maintainable, especially around providing new algorithms and pipelines for its highly parallel environment. Proofs of concept of this approach have similarly been prototyped, to verify ease-of-use and explore the scalability required for SKA (Allan et al., 2019; Cornwell et al., 2019).

This requires a significant up-front investment for rewriting code – SKA SDP software accounts for approximately 8.2% of the SKA construction budget, compared to 7.1% for the VLT, 5.7% for ALMA and 4.3% for ASKAP (Kemball and Cornwell, 2004; Guzman et al., 2014). We note that for SKA SDP this is processing software only, excluding telescope manager – functionality that is included in the figures for VLT, ALMA and ASKAP. However, this will improve TVO, by making it easier to make efficient use of the data processing hardware, by making it easy to implement new algorithms, and by isolating where code changes to support those algorithms are needed. This should thus reduce some of the maintenance cost of the SDP, and improve its ability to unlock new science across the lifespan of the telescope, albeit at an increase in upfront development cost (C_{dev}). The SDP will also undertake a phased hardware deployment, to provide compute when it is needed to support the increasing number of antennas and dishes on the ground, which will both keep capital costs lower, and reduce overall operational costs (C_{ops}). Furthermore, the deployment of hardware later on in the project allows hardware to be tailored to the software and vice versa, similar to

²⁴ <https://www.skatelescope.org/news/the-road-to-key-science-observations-with-the-ska-kicks-off-in-stockholm/>

the Cobalt correlator and beamformer in LOFAR, improving total relative science value of the resulting system. The SDP is for the SKA thus deliberately considering and trading off in different areas, the TCO and TVO of the system, with some decisions made to manage cost, and others to maximise total lifetime scientific value.

8. Related work

This work is a form of hardware–software co-design, as practised in the design of compute systems for large-scale science instruments. However, up to now, hardware–software co-design has focused mostly on more easily measured metrics, such as cost, power consumption and peak performance. Furthermore, while the literature often speaks of the importance of application co-design, the metrics used are agnostic and described mostly in terms of cost functions and constraints in energy and capital. In this paper we explore what these systems are really built for, and what a suitable measure for their performance would be.

This work can be considered a specialisation of general cost–benefit analysis in economics. Whereas cost–benefit analysis normally evaluates the social or financial benefit of a certain investment, this paper focuses on the scientific benefit in particular. There is research that introduces the concept of total value of ownership (Wouters et al., 2005) in accounting, but this is introduced as potential future research as an extension to TCO based decision making and not expanded upon. In that paper it is claimed that TVO builds on the concept of value as described in marketing literature.

Total value of ownership, also referred to as total value of opportunity, is also a metrics-based methodology for measuring and analysing the business value of enterprise IT investments (Apfel and Smith, 2003). This is an extension of TCO analysis, where both cost and any benefits of the proposed investment, tangible or intangible, are considered.

Value Engineering, Value Management and Value Analysis in Systems Engineering describe processes to achieve an optimal solution (Walden et al., 2015). This optimal solution is based on stakeholder value metrics; the processes are agnostic to these. In this paper we take the stakeholder view, describing and enumerating the value metric, while not considering the detailed processes required to optimise these.

Some work was done to analyse the societal impact of the High-Luminosity Large Hadron Collider (HL-LHC) upgrade of the LHC (Florio et al., 2016; Bastianin and Florio, 2018), predicting a larger than 90% chance of positive net economic benefit to society based on Monte Carlo simulations. These simulations estimate the economic returns from diverse benefits such as value of training for students, technological and industrial spillover, cultural effects for the public and academic publications. A comparison of the impact of the upgrade to the LHC with the non-upgraded instrument was also presented. An impressive effort is made to estimate the total cost of the current LHC, a difficult task even though all CERN expenses are well documented, due to the many in-kind contributions by member and non-member states. Societal impact analysis are very useful for funding agencies to gauge the value of an instrument to society using *scientific* and *objective* criteria. However, analysis of the methodology found many ambiguities and the scientific benefits of the LHC is given as less than 2% of the total impact of the instrument using this method: a drastic underestimation (Schopper, 2016). Furthermore, it was found that the societal impact of CERN's mission, “promote science and bring nations together”, was impossible to measure, since no way has been developed to measure in economic terms the success of the second objective. In comparison, the concepts

introduced in this paper look at more immediate impact, computational or scientific, and attempt to be more directly useful when making design choices.

Recent work on design optimisation of low-frequency telescopes using cost constraints (Boonstra and Nijboer, 2018) takes a slightly different and more domain specific approach. Here, an attempt is made to model both cost and scientific performance of a radio telescope using Lagrange multipliers. Scientific performance, defined by two instrumental figures of merit – sensitivity and survey speed, is optimised using both models and an assumed fixed capital budget. The LOFAR architecture as built, and the SKA phase 1 baseline design are analysed using the introduced model and variants optimised for survey speed and sensitivity are proposed. This methodology focuses on receiver and front-end optimisation and mostly ignores the cost required for compute capacity or how this scales with the number of stations and length of baselines. While the cost model does include a central correlator and beamformer, its model is exceedingly simple. Calibration, imaging and other post-processing costs, as well as long-term storage of data products, monitoring and control and operational costs are not modelled. Furthermore, we note that the chosen degrees of freedom in this paper, number of stations and number of antennas, have an enormous impact on required compute capacity for calibration and imaging. In this paper we take a more generic approach that is not limited to radio astronomy and that focuses on the cost and value of the compute systems that are not considered by Boonstra et al.

9. Summary and conclusions

In this paper we introduced a more formal way to reason about cost and value of compute resources, both hardware and software. We suggested that a focus on minimising cost alone is not sufficient to design an optimal solution. The introduction of several new concepts, total value of ownership, total lifetime computational value, total lifetime scientific value and relative science value, gives us the *vocabulary* to effectively discuss routes towards more optimal solutions. Although both total lifetime computational value and especially total lifetime scientific value are difficult to quantify, and we do not expect anyone to do so using the formulas given in this paper, we do show a number of components that allow us to reason effectively about this metric.

We provided a number of case studies in which we demonstrate the concepts introduced in this paper. We can see the utility of explicitly considering a metric of total lifetime scientific value, as the TITAN computer sought only to minimise capital cost (which happily led it to deliver truly exceptional value), whereas the SKA designers are explicitly allowing for relatively high costs in some areas to maximise total scientific value. In the LOFAR use case we noted the explicit trade-off made between high-impact science and dropping cost for computing, which led to an identified “processing window of opportunity” some years in the future where relative science value was perceived to be optimal. Some of the later additions to LOFAR were discussed, each adding their own value to the complex machinery that is the LOFAR telescope. Finally we showed, using a recent highly publicised processor flaw and its mitigating patches, that the total computational value of a system may potentially change over a system's lifetime. Together, these concepts and case studies provide a framework for decision makers, principal investigators, designers, and engineers of computing solutions to reason about the optimal solutions, in hardware or software, for their applications.

Acknowledgements

The authors would like to thank Yan Grange, Ágnes Mika, Bram Veenboer and Cees Bassa at ASTRON for their valuable feedback on early versions of this manuscript. We would like to thank Dr Elizabeth Waldram and Professor Malcolm Longair from the Cavendish Laboratory, Cambridge, for their help with understanding the impact of the TITAN Computer, and the Librarian of the Department of Computer Science Library at the University of Cambridge for access to the TITAN archival material.

References

- Ahmed, Haroon, 2013. *Cambridge Computing: The First 75 Years*. Third Millennium Publishing Limited, London, UK.
- Alexander, P., Allan, V., Badenhorst, U., Broekema, C., Cornwell, T., Gounden, S., Graser, F., Kirkham, K., Mort, B., Nijboer, R., Nikolic, B., Simmonds, R., Taylor, J., Wicencac, A., Wortmann, P., 2019a. SDP System Module Decomposition and Dependency View. Technical report, SDP Consortium.
- Alexander, Paul, Broekema, Chris, Ratcliffe, Simon, Bolton, Rosie, Nikolic, Bojan, 2013. SDP Element concept. Technical report, SDP Consortium.
- Alexander, P., Nikolic, B., Allan, V., Broekema, C., Deegan, M., Wortmann, P., 2019b. SKA1 SDP High Level Overview. Technical report, SDP Consortium.
- Allan, V., Nikolic, B., Farreras, M., Cornwell, T., Lyon, R., 2019. SKA1 SDP Execution Frameworks Prototyping Report. Technical report, SDP Consortium.
- Apfel, Audrey, Smith, Michael, 2003. TVO methodology: Valuing IT investments via the Gartner business performance framework. *Strateg. Anal. Rep.*
- Armour, Wesley, Karastergiou, Aris, Giles, Michael, Williams, Chris, Magro, Alessio, Zagkouris, Kimon, Roberts, Sarah, Salvini, Stefano, Dulwich, Fred, Mort, Ben, 2012. A GPU-based survey for millisecond radio transients using ARTEMIS. In: *Astronomical Data Analysis Software and Systems XXI*. In: ASP Conference Series, vol. 461, pp. 33–36.
- Bal, H., Epema, D., de Laat, C., van Nieuwpoort, R., Romein, J., Seinstra, F., Snoek, C., Wijshoff, H., 2016. A medium-scale distributed system for computer science research: Infrastructure for the long term. *Computer* 49 (5), 54–63.
- Bassa, C.G., Pleunis, Z., Hessels, J.W.T., 2017a. Enabling pulsar and fast transient searches using coherent dedispersion. *Astron. Comput.* 18, 40–46.
- Bassa, C.G., Pleunis, Z., Hessels, J.W.T., Ferrara, E.C., Breton, R.P., Gusinskaia, N.V., Kondratiev, V.I., Sanidas, S., Nieder, L., Clark, C.J., Li, T., van Amesfoort, A.S., Burnett, T.H., Camilo, F., Michelson, P.F., Ransom, S.M., Ray, P.S., Wood, K., 2017b. LOFAR discovery of the fastest-spinning millisecond pulsar in the galactic field. *Astrophys. J. Lett.* 846 (2), L20.
- Bastianian, Andrea, Florio, Massimo, 2018. Social cost-benefit analysis of HL-LHC.
- Boonstra, Albert-Jan, Nijboer, Ronald, 2018. Radio telescope design optimization using costing constraints: Extrapolating LOFAR costing to the Square Kilometre Array. *Radio Sci.*
- Bregman, Jaap D., 1999. Design concepts for a sky noise limited low frequency array. In: Smolders, A.B., van Haarlem, M.P. (Eds.), *Perspectives on Radio Astronomy – Technologies for Large Antenna Arrays*.
- Bregman, Jaap D., 2000. Concept Design for a Low-Frequency Array. In: *Proc. SPIE*, vol. 4015, pp. 4015 – 4015 – 14.
- Bregman, Jaap D., 2004. System optimisation of multi-beam aperture synthesis arrays for survey performance. *Exp. Astron.* 17, 365–380.
- Bregman, Jaap D., 2012. System Design and Wide-field Imaging Aspects of Synthesis Arrays with Phased Array Stations. PhD thesis. Rijksuniversiteit Groningen.
- Broekema, P. Chris, Mol, J. Jan David, Nijboer, Ronald, van Amesfoort, Alexander S., Brenjens, Michiel A., Loose, G. Marcel, Romein, John W., 2018. Cobalt: a GPU-based correlator and beamformer for LOFAR. *Astronom. Comput.* 23, 180–192.
- Cooley, James W., Tukey, John W., 1965. An algorithm for the machine calculation of complex Fourier series. *Math. Comput.* 19 (90), 297–301.
- Cornwell, T.J., Wortmann, P., Nikolic, B., Farnes, J., 2019. SKA1 SDP Algorithm Reference Library (ARL) Report. Technical report, SDP Consortium.
- Croarken, Mary, 1990. *Early Scientific Computing in Britain / Mary Croarken*. In: Oxford Science Publications, Clarendon Press.
- de Vos, C.M., Schaaf, K. van der, Bregman, J.D., 2001. Cluster computers and grid processing in the first radio-telescope of a new generation. In: *Proceedings of the 1st International Symposium on Cluster Computing and the Grid, CCGRID '01*. IEEE Computer Society, Washington, DC, USA, p. 156.
- Dennard, Robert H., Gaensslen, Fritz H., Rideout, V. Leo, Bassous, Ernest, LeBlanc, Andre R., 1974. Design of ion-implanted mosfet's with very small physical dimensions. *IEEE J. Solid-State Circuits* 9 (5), 256–268.
- Ellram, Lisa M., 1995. Total cost of ownership: an analysis approach for purchasing. *Int. J. Phys. Distrib. Logist. Manage.* 25 (8), 4–23.
- Elsmore, B., Kenderdine, S., Ryle, Sir, 1966. The operation of the Cambridge one-mile telescope. *Mon. Not. R. Astron. Soc.* 134, 87.
- Florio, Massimo, Forte, Stefano, Sirtori, Emanuela, 2016. Forecasting the socio-economic impact of the Large Hadron Collider: A cost-benefit analysis to 2025 and beyond. *Technol. Forecast. Soc. Change* 112, 38–53.
- Graser, Ferdl, Taylor, John, 2015. SKA SDP Costing Basis of Estimate. Technical report, SDP Consortium.
- Guzman, Juan Carlos, Chiozzi, Gianluca, Bridger, Alan, Ibsen, Jorge, 2014. The cost of developing and maintain the monitoring and control software of large ground-based telescopes. In: *Software and Cyberinfrastructure for Astronomy III*, Vol. 9152. International Society for Optics and Photonics, p. 91521P.
- Holties, H.A., 2012. BG/P Replacement Options. Technical report, ASTRON.
- Kemball, Athol J., Cornwell, T.J., 2004. A simple model of software costs for the Square Kilometre Array. *Exp. Astron.* 17 (1–3), 317–327.
- Kocher, Paul, Genkin, Daniel, Gruss, Daniel, Haas, Werner, Hamburg, Mike, Lipp, Moritz, Mangard, Stefan, Prescher, Thomas, Schwarz, Michael, Yarom, Yuval, 2018. Spectre attacks: Exploiting speculative execution. ArXiv e-prints.
- Larivière, Vincent, Sugimoto, Cassidy R., 2018. The journal impact factor: A brief history, critique, and discussion of adverse effects. In: Glanzel, W., Moed, H.F., Schmoch, U., Thelwall, M. (Eds.), *Springer Handbook of Science and Technology Indicators*. Springer International Publishing, Cham, Switzerland.
- Lipp, Moritz, Schwarz, Michael, Gruss, Daniel, Prescher, Thomas, Haas, Werner, Mangard, Stefan, Kocher, Paul, Genkin, Daniel, Yarom, Yuval, Hamburg, Mike, 2018. Meltdown. ArXiv e-prints.
- McKay-Bukowski, Derek, Vierinen, Juha, Virtanen, Ilkka I., Fallows, Richard, Postila, Markku, Ulich, Thomas, Wucknitz, Olaf, Brentjens, Michiel, Ebbendorf, Nico, Enell, Carl-Fredrik, et al., 2014. KAIRA: The Kilpisjärvi atmospheric imaging receiver array—System overview and first results. *IEEE Trans. Geosci. Remote Sens.* 53 (3), 1440–1451.
- Nash, Stephen G., 1990. A history of scientific computing. In: *ACM Press History Series*. Addison-Wesley.
- Needham, Roger M., 1992. Later developments at Cambridge: Titan, CAP, and the Cambridge ring. *IEEE Ann. Hist. Comput.* 14.
- Prasad, Peeyush, Huizinga, Folkert, Kooistra, Eric, van der Schuur, Daniel, Gunst, Andre, Romein, John, Kuiack, Mark, Molenaar, Gijs, Rowlinson, Antonia, Swinbank, John D, et al., 2016. The AARTFAAC all-sky monitor: System design and implementation. *J. Astron. Instrum.* 5 (04), 1641008.
- Romein, John W., Broekema, P. Chris, van Meijeren, Ellen, Schaaf, Kjeld van der, Zwart, Walther H., 2006. Astronomical real-time streaming signal processing on a Blue Gene/L supercomputer. In: *ACM Symposium on Parallel Algorithms and Architectures (SPAA'06)*, Cambridge, MA, pp. 59–66.
- Romein, John W., Broekema, P. Chris, Mol, Jan David, van Nieuwpoort, Rob V., 2010. The LOFAR correlator: Implementation and performance analysis. In: *ACM Symposium on Principles and Practice of Parallel Programming (PPoPP'10)*, Bangalore, India, pp. 169–178.
- Ryle, Martin, 1992. Radio telescopes of large resolving power. In: Lundqvist, Sven (Ed.), *Nobel Lectures, Physics 1971-1980*. World Scientific Publishing Co., Singapore.
- Ryle, M., Elsmore, B., Neville, Ann C., 1964. Observations of radio galaxies with the one-mile telescope at Cambridge. *Nature* 207, 1024–1027.
- Ryle, M., Neville, A.C., 1962. A radio survey of the North Polar region with a 4.5 minute of arc pencil-beam system. *Mon. Not. R. Astron. Soc.* 125, 39.
- Schopper, Herwig, 2016. Some remarks concerning the cost/benefit analysis applied to LHC at CERN. *Technol. Forecast. Soc. Change* 112, 54–64.
- Serylak, Maciej, Karastergiou, Aris, Williams, Chris, Armour, Wesley, Giles, Michael, LOFAR Pulsar Working Group, et al., 2012. Observations of transients and pulsars with LOFAR international stations and the ARTEMIS backend. *Proc. Internat. Astron. Union* 8 (S291), 492–494.
- Smith, James E., 1988. Characterizing computer performance with a single number. *Commun. ACM* 31 (10), 1202–1207.
- Stone, Richard, 1966. Input-output and demographic accounting: A tool for educational planning. *Minerva* IV (3).
- Tai, S.Y., Tjong Tjin, van den Broek, J., Deuten, J., 2019. De Impact van Grootchalige Onderzoeksinfrastructuur – Een Meetmethode Voor de Return on Investment van Internationale Onderzoeksfaciliteiten. Rathenau Instituut, In Dutch.
- Taylor, J., 2019. SKA1 SDP Performance Prototype Platform (P3-Alaska) Prototyping Report. Technical report, SDP Consortium.
- van der Schaaf, K., Bregman, J.D., de Vos, C.M., 2003. Hybrid cluster computing hardware and software in the LOFAR radio telescope. In: *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications, PDPTA, Vol. 2, Las Vegas, Nevada, USA*, pp. 695–701.
- van der Schaaf, Kjeld, Overeem, Ruud, 2004. COTS correlator platform. *Exp. Astron.* 17 (1), 287–297.
- van Haarlem, M.P., Wise, M.W., Gunst, A.W., et al., 2013. LOFAR: The low-frequency array. *Astron. Astrophys.* 556.
- van Nieuwpoort, Rob V., Romein, John W., 2010. Building correlators with many-core hardware. *IEEE Signal Process. Mag.* 27 (2), 108–117, (special issue on “Signal Processing on Platforms with Multiple Cores: Part 2 – Design and Applications”).

- Walden, David D, Roedler, Garry J, Forsberg, Kevin, Hamelin, R Douglas, Shortell, Thomas M, 2015. *Systems Engineering Handbook: A Guide for System Life Cycle Processes and Activities*. John Wiley & Sons.
- Wilkes, M. V. (Maurice Vincent), 1985. *Memoirs of a Computer Pioneer* / by Maurice V. Wilkes. In: *The MIT Press Series in the History of Computing*, MIT Press, Cambridge, Mass. ; London.
- Wouters, Marc, Anderson, James C., Wynstra, Finn, 2005. The adoption of total cost of ownership for sourcing decisions—a structural equations analysis. *Account. Organ. Soc.* 30 (2), 167–191.
- Zarka, P., Korso, Mohammed Nabil El, Boyer, Remy, Larzabal, Pascal, et al., 2015. NenUFAR: Instrument description and science case. In: *International Conference on Antenna Theory and Techniques - ICATT2015*. pp. 1–6.