ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-FR1BCO03

SKA PROJECT STATUS UPDATE

N. P. Rees*, SKA Observatory, Jodrell Bank, UK
on behalf of the The SKA Software Collaboration†

Abstract

The SKA Project is a science mega-project whose mission is to build the world's two largest radio telescopes with sensitivity, angular resolution, and survey speed far surpassing current state-of-the-art instruments at relevant radio frequencies. The Low Frequency telescope, SKA-Low, is designed to observe between 50 and 350 MHz and will be built at Inyarrimanha Ilgari Bundara, the CSIRO Murchison Radio-astronomy Observatory in Western Australia. The Mid Frequency telescope, SKA-Mid, is designed to observe between 350 MHz and 15 GHz and will be built in the Meerkat National Park, in the Northern Cape of South Africa. Each telescope will be delivered in a number of stages, called Array Assemblies. Each Array Assembly will be a fully working telescope which will allow us to understand the design and potentially improve the system to deliver a better scientific instrument for the users. The final control system will consist of around 2 million control points per telescope, and the first Array Assembly, known as AA0.5, is being delivered at the time of ICALEPCS 2023.

INTRODUCTION

This is the third in a series of SKA status papers presented at ICALEPCS meetings and together they provide a picture of the evolution of the SKA Software. At the time of the 2021 paper formal construction had just commenced, and we were in the process of issuing the first contracts. This paper covers the evolution since this time, but it starts with a brief historical overview.

HISTORY

Early Years

While the SKAO was officially created as an international observatory on 15 January 2021, the concept dates back to the early 1990's when astronomers proposed the idea of tracing the history of the Universe by mapping its most abundant element: Hydrogen (see, for example, Wilkinson [1]). Specifically, they proposed to observe the H1 emission line from it's rest frequency of 1420 MHz to red-shifts of more than 10, thereby observing it from just after the beginning of the universe to the present day.

To achieve this goal required a telescope of unparalleled sensitivity, with a collecting area approaching a square kilometre, and a frequency range from less than 100 MHz to a few GHz. Of course, once these basic parameters were described, scientists realised this instrument would be capable of observing a huge number of other phenomena. This significantly enhanced the science case and scientific interest in the project, but it also complicated the design, and necessitated the construction of two telescopes - one optimised for mid frequencies between 0.35 and 15 GHz (being built in South Africa) and the other between 50 and 350 MHz (being built in Australia).

The main technical difference between the two telescopes is that at the higher frequencies, it is still economic to provide the collecting area with traditional dishes - in the baseline design (known as Array Assembly 4) there are 133 15 m diameter dishes that will be integrated with the existing 64 13.5 m MeerKAT dishes on the same site to create an array of 197 dishes.

For the SKA-Low telescope, it is uneconomic to build dishes large enough that will be efficient at collecting at the longer wavelengths (50 MHz is about 6 m), so the SKA-Low design has 512 "stations" each of which is a collection of 256 phased wire log-periodic antennas. Each station is about 40 m in diameter, and is roughly equivalent to a dish antenna of the same diameter, but is much cheaper to build.

The project has always been seen as international in scope, and in 2013 it was agreed to establish an Inter-governmental Organisation to manage construction and operate the telescopes. This was about the same time when detailed design began, and so from 2013 to the start of construction in 2021 development followed two paths - one technical to deliver the design and the other governance, to deliver the organisation and funding structure.

Pre-Construction

Up until 2019 the technical development was known as "pre-construction" and international consortia developed detailed designs on different aspects of the telescopes. These culminated in a series of Critical Design Reviews coordination by the SKA Organisation. Between the completion of these CDR's and the formal start of construction was a "bridging" phase. In software, the design phase was far from ideal because it was largely a paper exercise, and the design consortia did limited prototyping, and consequently limited practical design validation. This changed in the bridging phase, where we pivoted to the development of processes with the adoption of the Scaled Agile Framework [2], and code creation, but we are still evolving the design as we learn about the system.

Construction

At the time of the last paper construction had only just commenced. The plan is to build the two telescopes simultaneously, and each is to be delivered in stages, as successively larger and more powerful interferometers. Each stage is known as an "Array Assembly (AA)", and the basic parame-

* Nick.Rees@skao.int
† See Appendix

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-FR1BC003

Table 1: Basic Parameters for SKA Major Milestones (The Operations Readiness Review)

Milestone Event (earliest)	Array Size	SKA-Mid (end date)	SKA-Low (end date)
AA0.5	4 dishes 6 stations	2025 May	2024 Nov
AA1	8 dishes 18 stations	2026 Apr	2025 Nov
AA2	64 dishes 64 stations	2027 Mar	2026 Oct
AA*	144 dishes 307 stations	2027 Oct	2028 Jan
Operations Readiness Review		2028 Apr	2028 Apr
End of Staged Delivery programme		2028 Jul	2028 Jul
AA4	197 dishes 512 stations	TBD	TBD

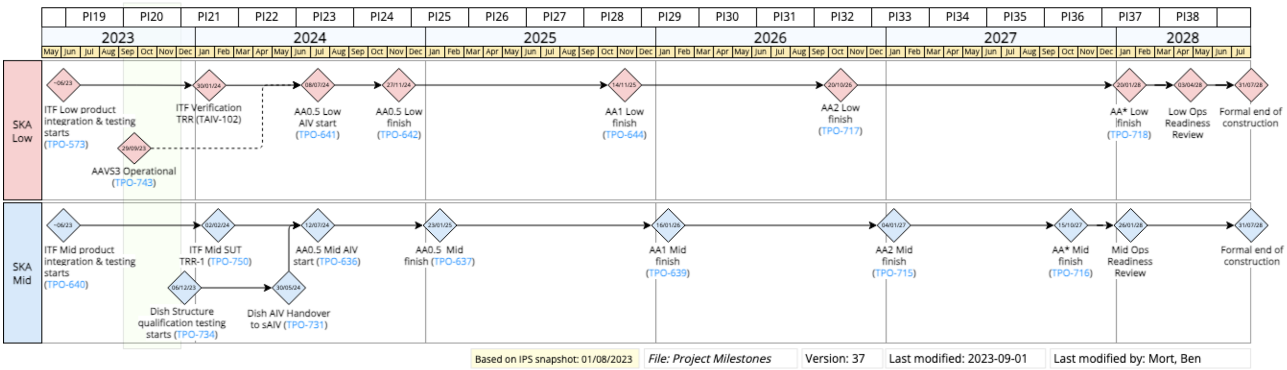

Figure 1: Graphical representation of the SKA project schedule as at 1 August 2023. Milestones dates do not include any allowance for schedule contingency, so the actual dates are anticipated to be delayed in a managed way.

ters of each array assembly is shown in Table 1. An outline of the major delivery milestones for both telescopes is shown in Figure 1. The original implementation plan had 4 Array releases (AA1 to AA4), but since then we have introduced a Minimum Viable System, which is smaller than AA1 (hence AA0.5) and funding is currently limiting the final size (hence the replacement of AA3 and AA4 by a funding limited AA*). Irrespective of this, the goal of the observatory is to implement the full AA4 sized system as soon as funds allow it.

CURRENT STATUS

As of August 2023, 70 contracts and significant purchase orders had been awarded for a total commitment of approximately €575M. These contracts ranged across the major infrastructure works in each country as well as the technical subsystems (such as software) and services required. Recently, for example, contracts were endorsed for the main access, airstrip and site accommodation for the SKA-Low telescope in Australia, and delivery of 64 Dish Structure Systems for the SKA-Mid telescope in South Africa. Most software contracts were raised early in construction and are time and materials professional service contracts.

The progress measures indicate 15.2 % of the project work has been completed, against 17.0 % planned and 15.8 % spent. The key performance indicators show that the project is behind schedule and over budget for the achieved work to date; the over budget is the result of advanced payments required for the execution of the SKA-Low Infrastructure work and will be recovered as progress is made.

Since the start of construction, significant project contingency funds have been spent to manage the Observatory-level risks associated with the global economic situation driven by the pandemic and regional war as well as delays in deliverables external to the project but impacting its execution (e.g. schedule delays in land access, membership-driven changes in planning, host country deliverables). The management of the Observatory level issues, through a Management Reserve, was recommended at the level of 10 % (€111M in 2022) within the Construction Proposal, approximately half of which has been realised in the execution of the current contracts with the remainder of the procurement related issues tracking to the overall envelope. Without the availability of this Observatory funding, the additional scope would be added to the project with an outcome projected to be only at the 17 % level for probability of success. The Management Reserve allows a recovery of the project contingency which restores the probability of success to approximately 84 % as of the July 2023 data.

There are currently no significant changes to either performance or scope.

With regards to practical progress on the ground, at the moment the main focus is building the construction camps at both sites and accepting the first deliveries of equipment - at the low site in particular (See Figs. 2 and 3).

Software Status

A project of the SKA scale poses a unique set of political challenges and we were required to spread software development over 8 countries (Australia, China, India, Italy, South Africa, Switzerland, The Netherlands and the United

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-FR1BC003

Figure 2: The SKA-Low construction camp site. Bottom centre is the “fly camp” or the camp built to enable other construction camp activities. The clear area to its left is the site of the final construction camp, and the clear area to the top right is the lay-down area.

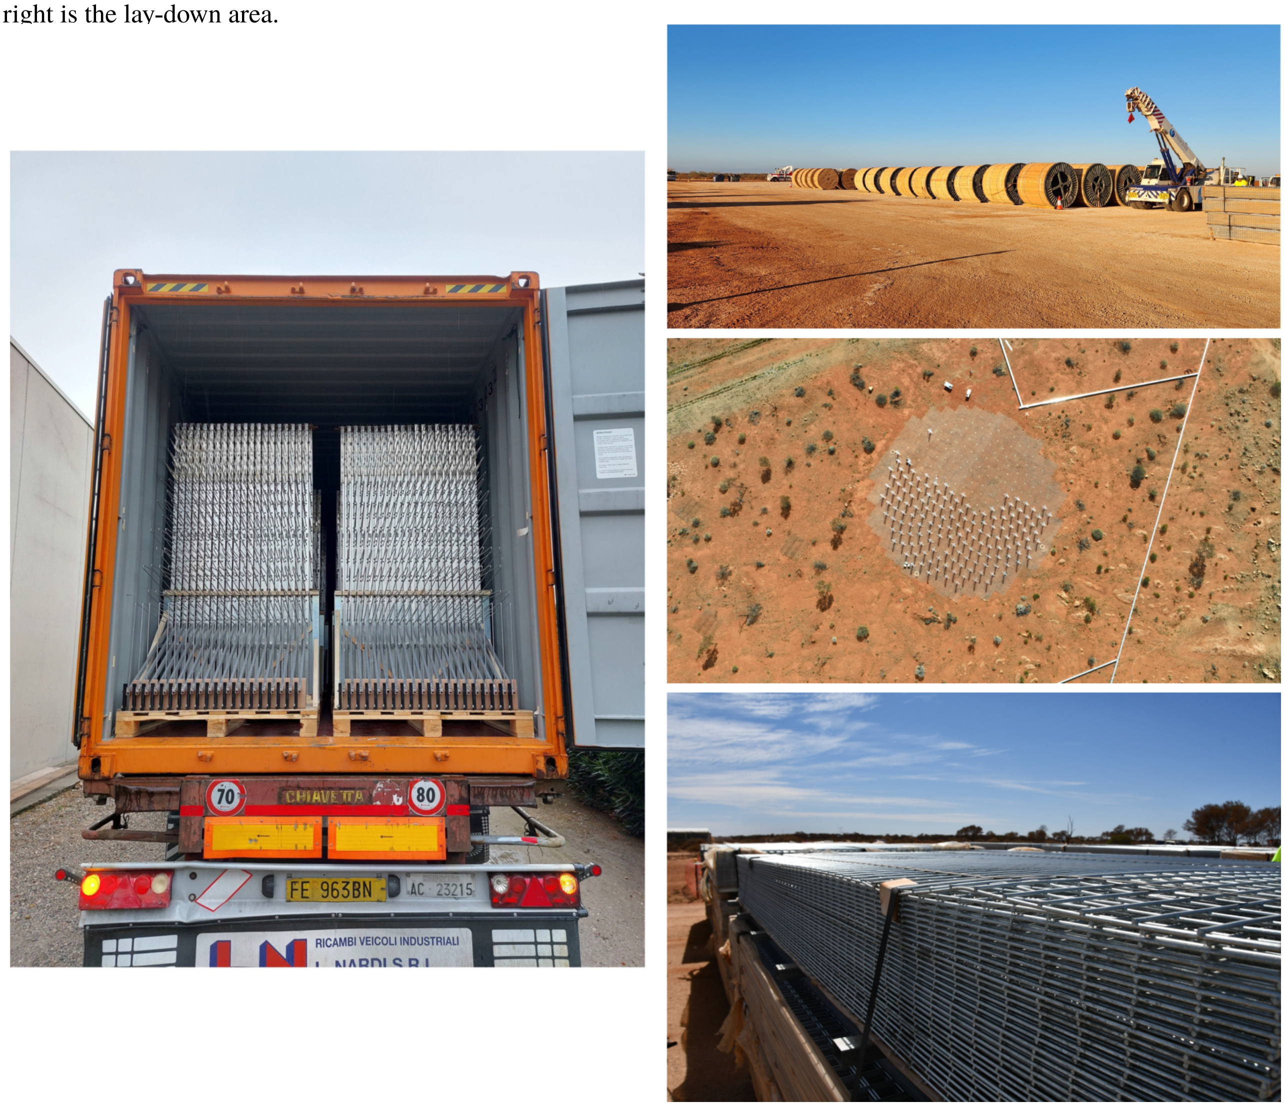

Figure 3: Recent construction pictures from the SKA-Low site. Left: A container of antennas; Right-top: Cable spools in the lay-down area; Right-middle: The partially completed AAVS3 station; Right-bottom: Wire mesh in the lay-down area.

Kingdom) with strict monetary allocations and around 150 contracted developers. Also, in order to allow for ongoing change and optimisation, we adopted the Scaled Agile Framework as the basis for our development processes. In order to avoid contractual silo's impacting the development process we decided to adopt highly relational contract structure based on the Vested [3] approach rather than traditional transactional contracts. This foregoes the traditional tightly defined contract specification in the belief that this is counterproductive if the work has a large number of unknowns, and instead replaces it with a general Statement of Intent describing how contractors must behave, and a commitment to the deliver on the goals of the programme. This, along with short term legal commitments but with long term assurances of continued work if standards are met, plus the range of companies involved reduces the overall risk and ensures the development teams work closely together with a goal of system level optimisation.

At the time of writing the main focus of software development is:

- Working with the assembly, integration and Verification (AIV) teams to deliver suitable monitoring and control systems for AA0.5 needs. AA0.5 is a very limited interferometer (4 SKA-Mid Dishes and 6 SKA-Low stations), with no scaling issues. We have working examples of most code components, but still struggle with system level integration.
- Supporting the rollout of a new test station for SKA-Low (known as AAVS3).
- Supporting the installation and testing of the first dishes for SKA-Mid.
- Developing the scaled processing needed for later Array Releases (AA2 and AA*)
- Developing the basic observing modes needed - such as interferometric pointing calibration for the SKA-Mid dishes.
- Integrating the signal chains from digitization through to science data processing.
- Improving the team's ability to deliver software efficiently.

Software is integrated and tested in a number of sites. We make a heavy use of the Continuous Integrated (CI) capabilities of the GitLab platform, and so we can run a CI pipeline wherever we can execute a GitLab runner. Most pure software integration happens in cloud systems, with the bulk of it currently running on private Science Cloud systems generously provided by the UK, but we also have the ability to use commercial cloud systems, if needed. Hardware integration happens either close to the hardware system developers and integrators (i.e. Canada, Australia and The Netherlands), or at the specialised SKA integration and test facilities in the host countries (specifically Cape Town in South Africa and Geraldton in Western Australia). At present, the current status of the software is that we can demonstrate:

- Preliminary end-to-end software for both SKA-Mid and SKA-Low
- Some operational modes of both telescopes, including setting up and running a simple (simulated) observation
- Handle (simulated) real time data ingest and some pointing calibration
- Run a simple pipeline to process simulated imaging
- Support future observation through sensitivity calculator
- Pulsar Search Software and other Data Processing pipelines are being prepared but will be integrated at a later stage.

CHALLENGES

Scaling

The biggest challenge for the SKA Software Collaboration is scaling. This is a multidimensional challenge, encompassing:

- Compute scaling of the data processing.
- Complexity scaling, particularly in the control system.
- Software development scaling - trying to maintain agility and flexibility in a large development team.
- Geographical and cultural scaling across time-zones, cultures and nationalities.

Each of this brings their own challenges. The first is dealt with by extensive testing of the code in scaling situations. The second relies on architectural encapsulation, and the ensuring that interfaces are kept simple and testable. The third is a human problem, and inherently relies on having aligned but autonomous teams, so the coordination and communication in manageable. However, this relies on the architecture mapping to the teams in a reasonable way. Finally, the last is an inevitable result of the international nature of the SKAO - it can be mitigated by asynchronous communication and reducing the communication needed between teams widely separated in time-zones, cultural understanding but some working out of hours is inevitable.

Continuously Working Product

A cornerstone of any agile development is a continuous integration system that ensures that a continuously working and tested set of products is always available. In a single large system, this is extremely difficult to guarantee, because the integration challenges ar significant. In this, we have found without the focusing drive of stakeholders who are using the system everyday it is easy to regress.

Software vs System Engineering

The SKAO has a strong System Engineering culture. In his excellent paper Maier [4] highlights a conflict between traditional System Engineering and Software Engineering - System Engineering is based on an assumption that the decomposition of products has a hierarchical structure with each sub-product having an "*is-part-of*" relationship with a parent product. This simplifies the concept of the system Engineering V model of decomposition design and verification. Maier points out that software product relationships are usually *is-used-by*, not *is-part-of*, which can either generate a flat structure of (ideally independent) services or a

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-FR1BC003

layered structure of modules. Since these don't adhere to the system engineering hierarchy, important software products that are used by many parts of the system often aren't recognised in the Product Breakdown structure used for System verification.

Architectural Evolution

In any large system, architectural evolution is difficult. Because of the extended, siloed, pre-construction phase where no system level code development was done, many of the design aspects were immature and suffered from the effects on Conway's Law [5]. The primary interface behaviours were untested, and few, if any, "unhappy path" behaviours were unspecified. However, other aspects of the design within sub-systems were quite detailed. This is contrary to the Agile ideal which would see a Minimum Viable Product built that had little functionality within the sub-systems, but was still representative of aspects of the whole system.

CONCLUSION

Development of the SKA Software is one of the most complex software undertakings of any scientific project. At the time of writing, it is at an exciting stage where we are starting to roll out software for use in the initial testing, commissioning and verification phase of the project. This presents many challenges, but we have organised the project to address these challenges as they arise and adapt to the changing needs of the project.

APPENDIX

The SKA Software Collaboration for at the time of this paper comprises:

- T. Dijkema, S. Van Der Tol, S. Wijnholds – ASTRON, Dwingeloo, The Netherlands;
- P. Osório, D. Regateiro, B. Ribeiro, H. Ribeiro – Atlar Innovation, Pampilhosa da Serra, Portugal;
- M. Lukkezen, C. Salvoni – CGI Nederland B.V., Rotterdam, The Netherlands;
- D. Acosta, G. Berriman, S. Ellis, C. Gallacher, N. Quinn, J. Sawdy, T. Swain – CGI IT UK Ltd, Leatherhead, United Kingdom;
- I. Novak – Cosylab Switzerland GmbH, Brugg, Switzerland;
- M. Colciago, M. Droog, J. Taylor – Covnetics Limited, Nuneaton, United Kingdom;
- M. Paulo, M. Santos – Critical Software S.A., Coimbra, Portugal;
- G. Hodosan – Csillagászati és Földtudományi Kutatóközpont, Budapest, Hungary;
- D. Mitchell, S. Ord – CSIRO Space & Astronomy Business Unit, Marsfield, Australia;
- D. Maia, B. Morgado – Faculdade de Ciências da Universidade do Porto, Porto, Portugal;
- J. Cantos, W. Gauvin, A. Jameson, W. van Straten – Fourier Space Pty Ltd, Kew East, Australia;
- F. Wang Guangzhou University, Guangzhou, P.R.China;
- J. Carrivick, C. Gray, N. Steyn, J. Strauss, R. Tobar – International Centre for Radio Astronomy Research, Crawley/Perth, Australia;
- G. Brajnik – Interaction Design Solutions, Udine, Italy;
- V. Alberti, C. Baffa, M. Canzari, M. Di Carlo, E. Giani, G. Marotta – Istituto Nazionale di Astrofisica (INAF), Roma, Italy;
- P. Thiagaraj, R. Uprade – NCRA-TIFR, Pune, India;
- A. Child, A. Clemens, E. Scott – Observatory Sciences Ltd, St Ives, United Kingdom;
- A. Dayanand, A. Deolalikar, A. Dubey, N. Gupta, J. Kolatkar, M. Lalwani, T. Phadtare, V. Shelake, O. Verma – Persistent Systems Limited, Pune, India;
- B. Mcilwrath, C. Pearson, N. Thykkathu – Rutherford Appleton Laboratory, Didcot, United Kingdom;
- S. Valame – Sanikaizen Solutions, Pune, India;
  C. Christelis, P. Dube, B. Ojur, D. Petrie, J. Venter – South African Radio Astronomy Observatory (SARAO), Observatory, South Africa;
- J. Guzman – SKA Observatory, Kensington, Australia;
- K. Ngoasheng – SKA Observatory, Observatory, South Africa;
- T. Ambaum, A. Avison, J. Banner, R. Barnsley, M. Bartolini, R. Beer, R. Bolton, R. Brederode, A. Bridger, A. Brown, M. Caiazzo, A. Clarke, J. Collinson, M. Deegan, D. Fenech, P. Griffiths, J. Hammond, P. Harding, D. Hayden, A. Holt, M. Jeavons, R. Joshi, T. Juerges, R. Laing, R. Leadbetter, P. Lewis, S. Lloyd, E. Mampuru, L. Mann, J. Masih, M. Miccolis, J. Mooneyan, B. Mort, J. Muller, A. Noutsos, P. Prior, A. Reader, N. Rees, S. Rice, R. Schofield, P. Shepherd, J. Stoddart, W. Swart, L. Tirone, S. Twum, S. Ujjainkar, S. Vrcic, B. Wallace, E. Wang, M. Waterson, P. Wortmann, U. Yilmaz – SKA Observatory, Macclesfield, United Kingdom;
- H. Groot – Science and Technology Experts Pool B.V., Delft, The Netherlands;
- G. Mant – The Science and Technology Facilities Council (STFC), Swindon, United Kingdom;
- M. Ambekar, S. Bajare, A. Dange, Y. Kamble, J. Kumbhar, V. Kumthekar, T. Nanaware, D. Pandey, B. Parakh, M. Patil, A. Patkar, M. Sharma, M. Shirore, A. Sonawane, R. Tayade Tata – Consultancy Services Limited, Pune, India;
- M. Nijhuis TriOpSys b.v., Utrecht, The Netherlands;
- L. Bartlett, A. Biggs, T. Kenny, P. Klaassen, V. Pursiainen, S. Williams – UK Astronomy Technology Centre, Edinburgh, United Kingdom;
- V. Allan, J. Allotey, M. Ashdown, J. Coles, F. Dulwich, Q. Gueuning, C. Lu, B. Nikolic, V. Stolyarov – University of Cambridge Astrophysics Group, Cambridge, United Kingdom;
- R. Braddock, J. Harvey, L. Levin-Preston, S. Melhuish, M. Mickaliger, R. Oberland, B. Shaw, B. Stappers –

ISBN: 978-3-95450-238-7 ISSN: 2226-0358 doi:10.18429/JACoW-ICALEPCS2023-FR1BCO03

University of Manchester Department of Physics and Astronomy, Manchester, United Kingdom;
- K. Adamek, S. Etemadi Tajbakhsh, G. Jagwani, A. Karastergiou, A. Naidu, A. Taqi, C. Williams, D. Wright – University of Oxford, Oxford, United Kingdom;
- V. Mohile – Vivek Satishchandra Mohile, Pune, India;
- K. Govender, F. Graser, K. Kirkham, A. Odendaal, D. Schaff – Vivo Resources CC t/a Vivo Technical, Somerset West, South Africa.

REFERENCES

[1] P. N. Wilkinson, "The Hydrogen Array," in *Radio interferom.: Theory Tech. Appl.: Proc. 131st IAU Colloq.*, Socorro, NM, USA, vol. 19, Jan. 1991, pp. 428–432.

[2] *SAFe - Scaled Agile Framework*. https://www.scaledagileframework.com

[3] *Vested®*. https://www.vestedway.com

[4] M. W. Maier, "System and software architecture reconciliation," *Syst. Eng.*, vol. 9, no. 2, pp. 146–159, 2006. doi:10.1002/sys.20050

[5] M. E. Conway, "How do committees invent?" *Datamation*, vol. 14, no. 4, pp. 28–31, 1968.