34th Annual INCOSE
international symposium
hybrid event
Dublin, Ireland
July 2 - 6, 2024

BLACK HOLE CINEMA: APPLICATION OF SYSTEMS ENGINEERING METHODS TO EXPAND AND ENHANCE AN EARTH-SIZED TELESCOPE

Abstract. On April 19, 2019, billions of people around the world caught a glimpse of infinity for the first time. The Event Horizon Telescope (EHT) released the first image of a black hole, shining a light on one of the darkest, most mysterious objects in the universe. To do so required the linking together of existing radio telescopes all over the world to create a "virtual" telescope array with the highest angular resolution of any telescope humanity had ever built. The result was an image that appeared on the front page of nearly every major newspaper on the planet. This iconic image was truly a breakthrough in astronomy. It is considered one of the most widely viewed images in science history. But rather than being a culmination of a decades-long effort, this image represents the beginning of a whole new era in astrophysics and in humanity's ability to use the extreme environment surrounding a black hole as a laboratory to understand the fundamental nature of space-time. To build on the effort and the momentum generated through its public impact, a team of EHT scientists and engineers is looking ahead to the next horizon: movies of black holes. This requires operating at a larger scale and faster pace than before, and a project team capable of designing and implementing a complex construction project in multiple countries simultaneously. It requires an investment of tens of millions of dollars and rigorous yet flexible project management controls and processes. In short, realizing the ambitious science goals of the next-generation EHT (ngEHT) project and managing all the complex interactions that come with those goals requires an organized, lean, efficient, and systematic approach.

Keywords. black hole, EHT, astronomy

Introduction

On April 19, 2019, billions of people around the world caught a glimpse of infinity for the first time. The Event Horizon Telescope (EHT) released the first image of a black hole, shining a light on one of the darkest, most mysterious objects in the universe. To do so required the linking together of existing radio telescopes all over the world to create a *virtual* telescope array with the highest angular resolution of any telescope humanity had ever built, deploying specialty instrumentation subsystems in remote environments, developing imaging algorithms and techniques that had not been invented before, and an international team of over 200 people working together from across the globe. The result was an image that appeared on the front page of nearly every major newspaper on the planet.

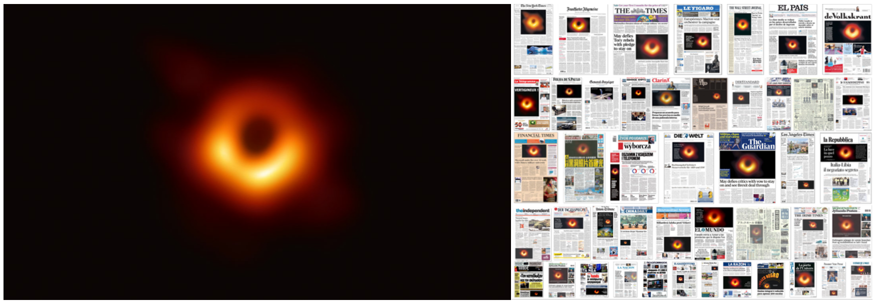

Figure 1. (Left) The first image of the supermassive black hole in the M87 galaxy. Credit: Event Horizon Telescope Collaboration. (Right) A collection of newspapers announcing the breakthrough M87 image.

This iconic image was truly a breakthrough in astronomy. It is considered one of the most widely viewed images in science history. But rather than being a culmination of a decades-long effort, this image represents the beginning of a whole new era in astrophysics and in humanity's ability to use the extreme environment surrounding a black hole as a laboratory to understand the fundamental nature of space-time. To build on the effort and the momentum generated through its public impact, a team of EHT scientists and engineers is looking ahead to the next horizon: movies of black holes. This is the aim of the Next Generation Event Horizon Telescope (ngEHT), a program to explore and define a long-term plan to enhance the EHT to realize a new set of transformative science goals (Doeleman, et al., 2023). Funded through the National Science Foundation's Mid-Scale Research Infrastructure (MSRI) program, the ngEHT Program, which has completed its initial stages of design, is poised to implement transformative upgrades to the array and help evolve the international collaboration for this new era.

To make scientifically interesting movies and probe the dynamics of matter as it falls into a black hole, an entirely new array—and operating model—is needed. Whereas the EHT succeeded in creating an image by piecing together a loose collaboration of institutions and pointing their existing telescopes at the same target for a few days' time, a movie requires new sites in new geographic locations to fill in gaps in the uv plane, newly designed antennas with high surface accuracy, four times more observing time, and the ability to process an order of magnitude more data through high speed instrumentation and data pipeline systems. It requires the evolution of an international collaboration to operate at a larger scale and faster pace than it has before, and a project team capable of designing and implementing a complex construction project in multiple countries simultaneously. It requires an investment of tens of millions of dollars and rigorous yet flexible project management controls and processes. In short, realizing the ambitious science goals of the ngEHT project (Johnson, et al., 2023) and managing all the complex interactions that come with those goals requires an organized, lean, efficient, and systematic approach.

This paper discusses how the ngEHT team has created and implemented a systems engineering process, team, and culture to help deliver the next transformational result in black hole imaging. While the ngEHT presents a unique concept, the system-related challenges are common to many projects facing similar stages of development and growth, building on past successes and legacy systems, and designing systems and products to meet a compelling, ambitious long-range vision.

Natural Evolution of a Breakthrough Concept

Decades of technical development and precursor observations preceded the first image of the M87 supermassive black hole. For much of that time, many people still thought actually imaging a black hole was impossible. The astronomical technique employed, called Very Long Baseline Interferometry (VLBI), is not new. It has been used by radio astronomers since the 1960s to push the limits of angular resolution by linking distinct geographic sites together, effectively creating a "virtual" telescope array that can be much larger than any single dish or single geographic location. But no one had attempted VLBI at high enough frequencies and at a large enough scale to image and resolve a black hole.

Pioneering efforts in the 1990s started to show that various technical challenges could be overcome. VLBI experiments demonstrated that observations at the required frequency, 230 GHz, were possible (Padin, et al., 1990) (Krichbaum, et al., 1998). Purpose-built back end and recording systems deployed to early EHT arrays (Doeleman, et al., 2008, 2012) confirmed that imaging the two largest supermassive black hole sources on the sky was feasible, leading to grassroots efforts to build a community, and eventually a formal collaboration, around black hole imaging. A collaboration initially numbering in the dozens slowly grew to over 200 worldwide, laying all the technical, logistical, organizational, and analytical groundwork needed along the way.

The effort proceeded organically, without a single, significant, sustained funding source, from a variety of contributions from American, European, Asian, and South American research grants. Further, unlike other large science collaborations like the Laser Interferometer Gravitational-Wave Observatory (LIGO), the EHT, and the very technique of VLBI, takes advantage of existing infrastructure that has been built for other science purposes, only carving out enough observing time as necessary for VLBI-specific observations (VLBI typically accounts for a small percentage of overall radio astronomy observations). Purpose-built VLBI systems are limited to specific instrumentation rather than full astronomical facilities. In other words, the EHT probably inevitably started the way many startups do: pieced together with innovative efforts of a few visionaries, held together with metaphoric (and actual, in some cases) duct tape, and a sparse funding runway only long enough to eke out the next incremental result leading up to the first image.

Could the M87 image that shocked the world have been created if the team had developed the EHT array using systems engineering tools and techniques from the beginning? A centralized, systematically-designed array and organization not only would have carried a prohibitively high price tag, but it likely would have faced an uphill battle gaining support within academic communities and funding agencies before the momentum earned by the collaboration through a globally-acclaimed result. In the case of large scientific collaborations without established formal engineering processes, a rigorous, systematic process to design can also seem overly burdensome, especially to a highly capable team that has already achieved global fame. However, as with startups scaling up and maturing into a larger enterprise, what led to the initial growth and success is not usually the same recipe for continued evolution. Higher complexity calls for the adoption of techniques to manage that complexity to deliver more ambitious goals. In this way, systems engineering—when done properly and implemented at the right stage of a project, as in the case of EHT—can be used to bring about a new phase in a team's growth.

System Concept

The EHT's success is the result of collaboration between many institutions globally that have deployed the needed subsystems for VLBI observations and devoted a modest amount of observing time (~two weeks/year). The foundation of the ngEHT concept is to expand the observing window to up to 40 days/year and augment the existing EHT array with new, modest-diameter (10-meter class) radio antennas in strategic locations around the globe. These new sites fill in critical gaps in the uv plane by providing both long and short baselines. The long baselines—stretching across the entire hemisphere of the Earth—support higher resolution that enables observation of finer scale structure at the black hole's event horizon, while the short baselines support wider field-of-view imaging to connect the dynamics at the event horizon to the larger scale relativistic jets emanating from the the poles of spinning black holes.

Additionally, to date, the EHT has primarily observed at a single frequency, 230 GHz, with some observations at a subset of sites observing non-simultaneously at 345 GHz (this is on the upper end of the radio frequency bands of the electromagnetic spectrum that radio astronomers are mostly interested, which ranges from 3 kHz up to about 900 GHz). The ngEHT project focuses on expanding the observing frequency range of most of the EHT array to three simultaneous frequency bands: 86, 230, and 345 GHz. Combined with larger diameter dishes in the existing array on the other end of new baselines, and the expanded frequency range and higher bandwidth added through new receivers and back end subsystems, the ngEHT concept will dramatically enhance sensitivity and Fourier spatial frequency coverage of the EHT array. For a detailed discussion on the process to design, architect, and implement the ngEHT concept, see Reference Array and Design Consideration for the next-generation Event Horizon Telescope (Doeleman, et al., 2023).

To summarize, the ngEHT concept brings:

- Multi-frequency enhancements to existing site instrumentation, enabling simultaneous observation at up to three frequency bands and increasing weather availability through a technique called *frequency phase transfer*
- New, dedicated sites with new antennas that fill in key gaps in the uv plane and provide redundancy for critical baselines, leading to greater and more reliable imaging fidelity
- A robust, modern data analysis pipeline and correlation capability designed for an order of magnitude more data and streamlined operations for the next decade

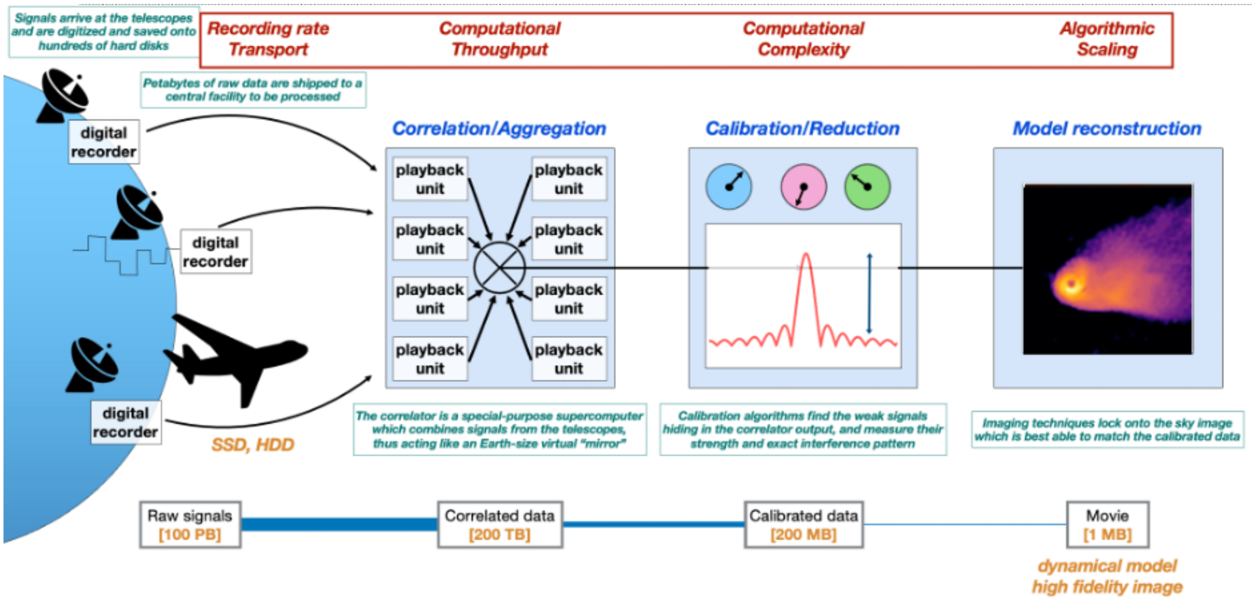

Figure 2. Data pipeline concept of operations for the ngEHT project. Radio frequency signals are collected from antennas all over the world, digitized on-site, and stored on hard disk drives. Those drives are shipped to a central correlation center where the signals are correlated. Next, a calibration step is performed to extract and measure specific parts of the signal, further reducing the data to a few hundred MB. Finally, imaging algorithms employing models based on underlying physics are used to create a high fidelity image. Credit: Lindy Blackburn, EHT Collaboration.

A System of Systems Challenge

Designing a next generation global telescope array that spans the globe and literally uses the entire spinning Earth to help fill in imaging data, built on a backbone of one of the most successful and iconic science results of all time, presents a fascinating systems challenge. The mix of both technical and non-technical challenges make for a perfect storm: a multi-variable complex problem that pushes the limits of both technology and organization.

On the technical side alone, there are several aspects that make black hole cinema using a VLBI array challenging, summarized as follows:

- **Unconstrained system boundaries**: The only major constraint to this VLBI concept is the actual size of the Earth. As a ground-only array, the largest possible baseline (and therefore the maximum achievable angular resolution) is set by the max distance where two points on the globe can both see the same source on the sky at the same time; effectively, the hemisphere facing the source at any point. [Note: This is not to mention a separate concept to expand the EHT into space through an orbital space telescope. This concept, called the Event Horizon Explorer (EHE) is also being pursued by members of the EHT and many of the same people involved in the EHT are also part of the ngEHT Program].
- **Complex system optimization problem**: There are a myriad of parameters that make up the overall "tradespace model," as shown in Figure 3 below. These parameters include architecture decisions on one side and the performance and programmatic attributes to optimize for on the other side. The tradespace seeks to answer questions like: how large do the antennas need to be? How

many sites? Where will they be located? Is the weather good enough there during the expected observation window to allow co-visibility with enough baselines to pass a critical threshold in data collection? This framework is used to define and test system concepts through simulations (Doeleman, et al., 2023).

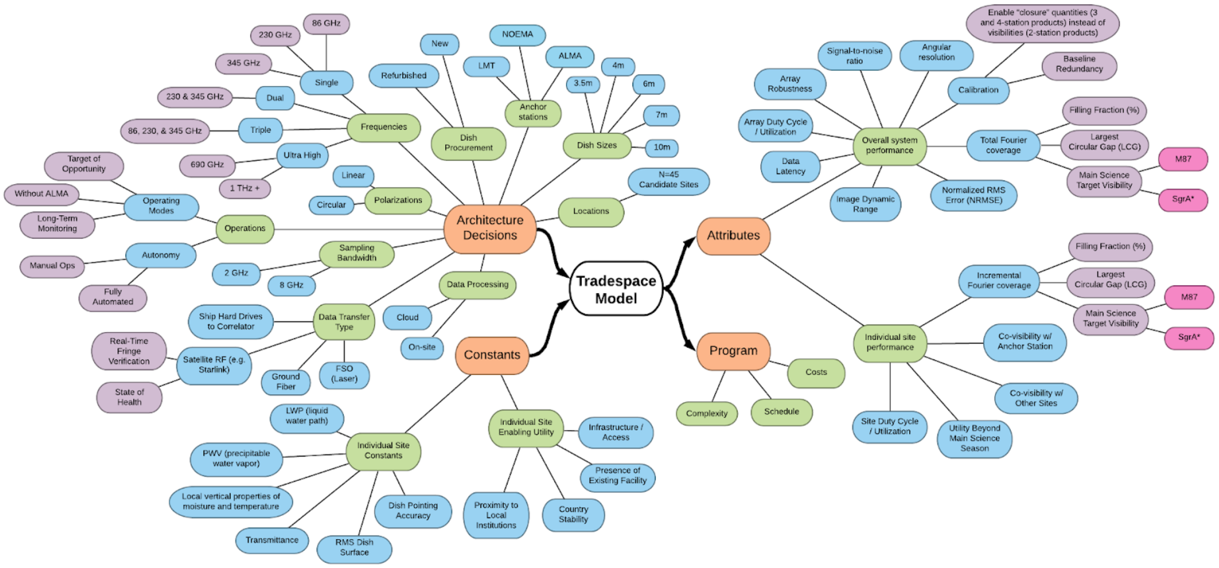

Figure 3. Tradespace model of the ngEHT program. Major architecture decisions and relevant constants are shown as inputs on the left; attributes on the overall system and individual site performance, as well as programmatic attributes, are the outputs on the right. The "model" is used as a conceptual framework for analyzing the tradespace, as opposed to being a singular parametric model that can be simulated in one software envi-ronment.

- **Simultaneous tri-band observing**: Today, the EHT has conducted the majority of its observations at a single frequency band, 230 GHz, with some additional non-simultaneous observations at 345 GHz. The ngEHT system concept relies on simultaneous tri-band observing at 86, 230, and 345 GHz to achieve science goals. Simultaneous tri-band receivers have been built at lower frequencies, but never before at these frequencies. Implementing these tri-band receivers is one of the most significant technical engineering challenges of the program.
- **4x increase in observation time**: The EHT currently observes for about two weeks a year in either March or April when the weather across the array is optimal for high frequency observations. For movies of M87, which changes on a scale of a few days, an observation window of three months is required, with observations taking place on a cadence of once every three days within those three months. Not only does this increase in observation time stretch the available time at those observatories for VLBI science, but it is also technically challenging to coordinate the logistics for media recording, retrieval, shipping, and processing at a central correlation facility or facilities.
- **4x increase in recording bandwidth**: Science goals require higher sensitivity, which can be achieved through increased system equivalent flux density (SEFD) received at antennas (i.e. bigger antennas), or through higher recording bandwidth. The ngEHT system concept increases recording bandwidth from 32 Gbps per frequency band today to 128 Gbps per frequency band through the deployment of higher speed back end subsystems to digitize at this rate, and higher speed recording subsystems to record, offload, and ship media at a turnaround rate that keeps up with observing cadence.

- **10x increase in data throughput**: More observing time plus higher bandwidth equals more data. The EHT currently processes on the order of 10 PB of data annually. The ngEHT concept increases this to about 100 PBs annually. This increased volume of data must be managed efficiently throughout the data pipeline.
- **Several bottlenecks in existing data pipeline**: Aside from an increase in data volume, the current data pipeline suffers from several technical bottlenecks resulting in backlogs of data from previous campaigns. The time-to-science on average in the current EHT is about two years from data collection to science result. This needs to be reduced to six months to efficiently operate multiple observing campaigns annually.

While the above-listed technical aspects are significant, the non-technical factors of the ngEHT Program are perhaps even more challenging.

- **Managing an engineering design and construction project within an academic environment**: University labs and organizational processes are typically designed for smaller scale instrumentation projects: bespoke systems deployed to singular *field* locations. Scientists, engineers, and support staff are hired for, trained on, and gain experience from these types of projects. Larger engineering and construction efforts are routinely carried out by academic institutions, and the EHT benefits from the experience of many members at institutions that have been involved in radio telescope projects from scratch. However, the project management controls and processes needed for a complex international construction effort will be largely new for the ngEHT project team, and will likely stretch the various partner institutions in new ways.
- **Structure of international collaboration and operating model**: The EHT is largely a loose confederation of international institutions. Observatories provide in-kind contributions of personnel time to conduct VLBI observations, and the array is only convened as an array for a distinct two-week period every year (in addition to a dress rehearsal that typically takes place in January of each year). The collaboration handles elements ranging from membership to data rights to publication policy to operations planning. The overall structure, organization, governance, and operating model will likely evolve over time as the EHT grows from a dedicated experiment into more of a facility-type model.
- **Complex stakeholder landscape**: The EHT could present an interesting case study in stakeholder management alone . There are existing sites and partners in the EHT today, some with their own joint operating agreements to operate as a single site that need to be considered when operating as part of a wider array. There are international funding agencies with competing priorities, timelines, and funding opportunities. For new sites, there are regulatory approval agencies with unique requirements in each country where those sites are proposed for development. Additionally, local stakeholders in the communities surrounding each site, new and existing, need to be considered as the program weighs environmental and cultural impacts throughout the lifecycle of the array.
- **Desire to take advantage of the "splash" momentum of the first black hole image**: The first image of a black hole produced a big splash. It garnered national attention and personnel who made it happen received numerous awards. There is a motivating sentiment within the EHT community to take advantage of this momentum and use it to realize a transformational next step for this array. This desire and the ambitions associated with it create a unique set of challenges to set the bar high enough to be commensurate with another big splash. This also sets high expectations and greater pressure on the project team to meet those expectations.

All of these challenges add up to a truly awesome system of systems opportunity. The ngEHT Program is a chance to design a set of stations specifically and primarily for VLBI observation. The sections that follow provide a glimpse into how systems engineering tools and techniques have been applied on the ngEHT Program, and how the program is positioned to employ systems engineering processes in the next phase of implementation.

Project Organization

System Context

The Array is the heart of the EHT. Its purpose is to perform VLBI observations and supply the resulting data for further scientific analysis. Conceptually, it is a collection of radio telescopes distributed around the world along with supporting systems to operate it and process the collected data. Yet, this collection is an abstract construct: each of the radio telescopes comprising the EHT Array today is itself a separate system, designed and built for other millimeter and submillimeter requirements, and operated in a distinct manner according to each station's unique set of stakeholders. To date, the EHT has only assembled for EHT observations during a two-week window each year, with specific VLBI-hardware and software systems installed at each telescope to enable these observations. The concept of "the Array" can therefore be confusing to draw a boundary around; systems are both a part of the virtual telescope that comes together for a part of the year and their own observatory that exists outside the context of the EHT. For the purposes of this system design project, both contexts are important views. Figure 4 shows one view of the system context, focusing on the most critical relationships and stakeholders influencing the Array as it exists in both EHT and non-EHT environments.

The ambition of this program presents yet another new aspect of complexity: new sites with new partners that are purpose-built for the VLBI science identified through stakeholders of the ngEHT Program and the EHT Collaboration. Whereas to date, only existing telescopes have joined the EHT, in the near future, new sites will be built to meet key science goals of the ngEHT and EHT communities. This added complexity means the Array will be further heterogeneous, comprising a mix of not only many different existing telescopes, but also new sites that are fully built through the implementation phase of this program. From a systems engineering perspective, this hybrid mix of "existing plus new" places even more importance on defining "the system" accurately. The tools and language used to define and manage requirements, architecture, and behavior need to account for this hybrid mix.

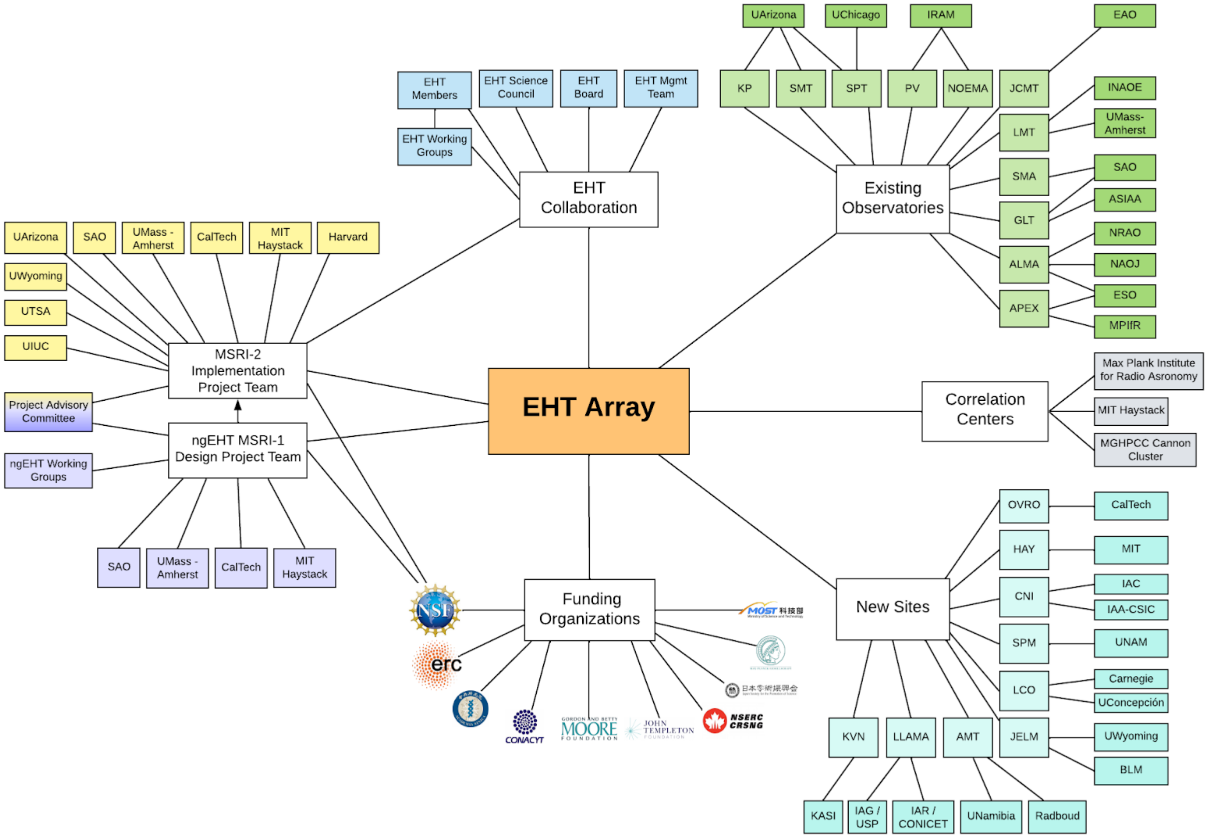

Figure 4. System context diagram for the EHT Array, showing major relationships between the array and the stakeholders influencing the design, development, and implementation of new array capabilities.

The figure above models the EHT Array at the center as the “system of interest,” with each line representing a major relationship between the EHT Array and external elements. Each of these external elements is a grouping of stakeholders. The outer ring of this group contains the primary stakeholders that influence the design, development, and implementation of the EHT Array. These groupings, and their relationships with the EHT Array, are defined in more detail below.

EHT Collaboration

The EHT Collaboration is an established international collaboration that sets the governance, policies, and processes to use the EHT to produce black hole science. Relevant stakeholders include all **EHT members**, who, by joining the EHT, agree to follow the codes of conduct, data access and publication policies, and overall direction of EHT leadership. Many EHT members are a part of **EHT working groups**, which are teams of experts oriented around specific domain expertise, working collaboratively to produce science results and methods. The **EHT Science Council** (SC) is another body of stakeholders composed of EHT members. The SC provides scientific guidance to EHT leadership, and occasionally to the publication committee. This group has also, historically, served as a conduit for concerns about scientific culture and climate and ethics within the collaboration. The **EHT Management Team** manages the execution of activities within the EHTC (observations, logistics, near-term instrumentation upgrades, interactions with sites, etc.). Finally, the **EHT Board** is a committee of EHT members appointed by the collaboration to represent the interests of EHT members and each institution of the collaboration, and to oversee the activities within the EHTC as a whole.

Existing Observatories

The term "Existing Observatories" applies to the telescopes in the EHT Array today that join EHT observations. Existing observatories may be either single dish telescopes or array telescopes. The primary stakeholder of these telescopes are typically the institution that operates each telescope (a full context view of all stakeholders of each telescope would be too complex to fit in a single diagram; therefore only the operating institution is shown in Figure 4). These observatories and their abbreviated codes are listed below:

- Single Dish Telescopes
    - Kitt Peak (**KP**), Kitt Peak, Arizona
    - Submillimeter Telescope (**SMT**), Mt. Graham, Arizona
    - South Pole Telescope (**SPT**), Antarctica
    - IRAM 30m Telescope, Pico Veleta (**PV**), France
    - James Clerk Maxwell Telescope (**JCMT**), Mauna Kea, Hawaii
    - Large Millimeter Telescope (**LMT**), Volcán Sierra Negra, Mexico
    - Greenland Telescope (**GLT**), Thule Air Force Base, Greenland
    - Atacama Pathfinder EXperiment (**APEX**), Atacama, Chilé
- Array Telescopes
    - Atacama Large Millimeter / Submillimeter Array (**ALMA**), Atacama, Chilé
    - Submillimeter Array (**SMA**), Mauna Kea, Hawaii
    - Northern Extended Millimeter Array (**NOEMA**), Plateau de Bure, France

Correlation Centers

Crucial to the data processing of the EHT Array are the Correlation Centers, which take in raw data from the telescopes, filter out the noise, and find coherence signals between pairs of telescopes, or baselines. EHT data today is correlated in two centers: the **MIT Haystack** VLBI Correlator and the **Max Plank Institute for Radio Astronomy** VLBI Correlator. An additional correlation site at the Massachusetts Green High Performance Computing Center (**MGHPCC**) Cannon Cluster is proposed as a new addition to help process an order of magnitude more raw data with the enhanced EHT Array.

New Sites

Expanding the capabilities of the current EHT Array to meet the ambitious science goals outlined in the Science Traceability Matrix (STM) requires filling in the uv plane with new telescopes in strategic locations around the world. These new locations are defined as "New Sites" here. Extensive study has determined the best locations of these new sites (Doeleman, et al., 2023), factoring in the existence of an active observatory for ease of construction and lower cost, and optimal science return. New sites include: 1) sites with already-built antennas that will join the EHT Array with new instrumentation, 2) sites with antennas and instrumentation which are currently under construction by existing partners of the EHT, and 3) sites with newly-built antennas that are new partners to the EHT. These sites and their abbreviated codes are listed below:

- Sites with already built antennas that will join the EHT Array with new instrumentation:
    - Owens Valley Radio Observatory (**OVRO**), Owens Valley, California
    - Haystack 37m (**HAY**), MIT Haystack Observatory, Massachusetts
- Sites with antennas and instrumentation which are currently under construction by existing partners of the EHT

  - African Millimeter Telescope (**AMT**), Gamsberg, Namibia
  - Large Latin American Millimeter Array (**LLAMA**), Atacama, Argentina
  - Korean VLBI Network (**KVN**), comprising three telescopes: KVN Yonsei Radio Observatory (KY), KVN Ulsan Radio Observatory (KU), and KVN Tamna Radio Observatory (KT); Korea
- Sites with newly-built antennas that are new partners to the EHT
  - Canary Islands (**CNI**), Teide Observatory, Tenerifé, Spain
  - San Pedro Mártir National Astronomical Observatory (OAN-**SPM**), Baja California, Mexico
  - Las Campanas Observatory (**LCO**), Atacama, Chilé
  - Mount Jelm (**JELM**), Wyoming Infrared Observatory, Wyoming

Application of Systems Engineering

The ngEHT Program has been designed following a systems engineering approach that emphasizes rigorous analysis of the interactions and interfaces between components and traceability from high-level requirement to design detail. A dedicated systems engineering team serves as the "connective tissue" between engineering disciplines and the science team, ensuring shared understanding of the needs and goals of the system and well-documented decision making. Maintaining a singular and consistent "source of truth" for requirements, architecture, and V&V will be critical as the program moves into an implementation phase with many different subsystems developed by many different partners.

Systems-Driven Process to Define Key Science Goals

A primary goal of the systems engineering process employed by the project team is to understand the nature of the relationships between top-level science goals, the requirements hierarchy, and elements of the system architecture. Armed with this understanding, and with tool-assisted capabilities for analyzing these relationships, it is possible for the team to:

- Understand which science goals and associated performance metrics are associated with specific requirements and system functions
- Rapidly identify the downstream impacts of proposed changes in science goals or requirements
- Develop V&V test suites that can objectively show coverage of specific science goals

But first, the team needed to define its key science goals that would drive performance targets for the array. Though the stated vision ("from still images to movies of black holes") was a compelling rallying cry and a logical next step in the high level output of the array, there was no real consensus at the start of the project around the specific measurement and observations required to produce a scientifically compelling movie. What defines a worthwhile movie that justifies a potentially one hundred million dollar investment of purpose-built VLBI infrastructure? Three frames of an image? Ten? Thirty? How do the changing astrophysical parameters surrounding a black hole collide with the very real operational constraints of weather at every site of the array, observation cadence and duration, performance, and logistics? Further, the study of horizon-scale dynamics only represents a subset of the possible science that can be done with imaging analyses. What kind of science was possible from a ground-based imaging array for additional black hole sources, beyond M87 and the Sagittarius A* black hole in the center of our Milky Way galaxy?

To begin, the project team organized the community of interested scientists into eight different working groups, and tasked them with coming up with a prioritized list of science goals. Within each working group, scientists made a case for particular goals (and associated array requirements needed) through presentations, white papers, and used supporting simulations of array performance to understand and convey what was possible within a reasonable parameter space. Defining that parameter space was a product of a parallel effort from a small group of scientists and engineers within the project team. This was a critical step to give the science working groups constraints—even if they were merely informed guesses at first—to have a clear place to start from.

Guided by the systems engineering process shown in Figure 5 below, the project team facilitated an iterative process to refine science goals and gain community consensus on priorities. This involved three international science meetings, each of which was open to the full physics and astronomy communities (i.e., no prior membership in EHT was required to participate), and many rounds of analysis.

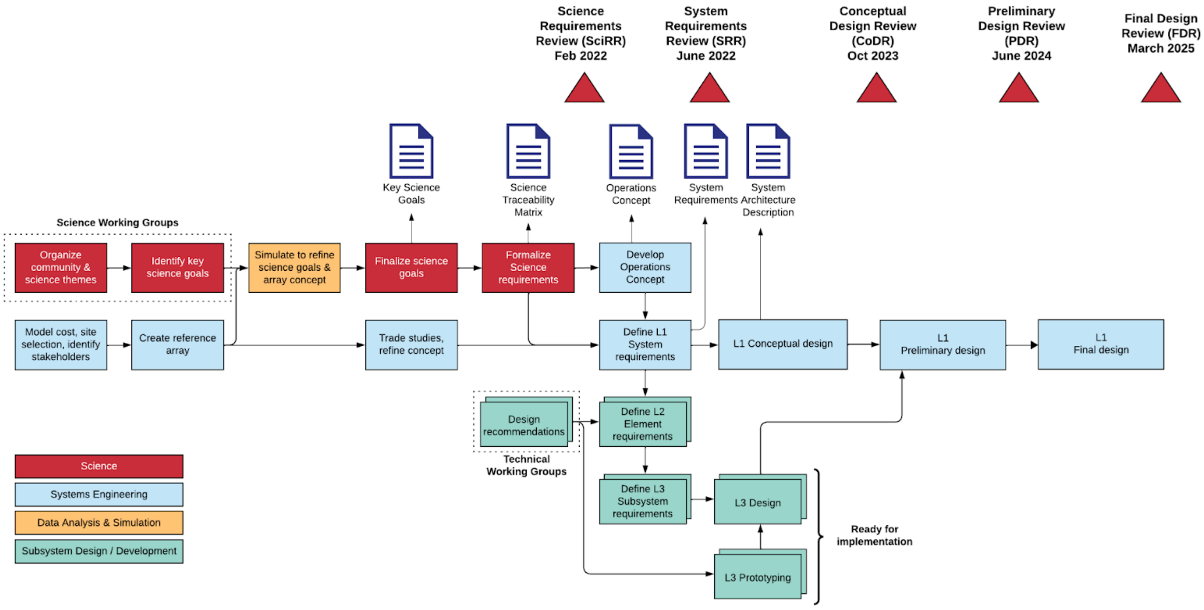

Figure 5. ngEHT systems engineering design process flow, from science goals through final design

This process ultimately resulted in the definition of a series of Key Science Goals, organized by science theme:

Fundamental Physics:

1. †Establish the existence and properties of black hole horizons
2. †Measure the spin of a Super Massive Black Hole (SMBH)
3. Constrain the properties of a black hole's photon ring
4. Constrain ultralight boson fields

Black Holes & their Cosmic Context:

1. †Reveal Black Hole-Galaxy Formation, Growth and Coevolution
2. Determine how supermassive black holes (SMBHs) merge through observations of sub-parsec binaries
3. Connect SMBHs to high-energy and neutrino events within their jets

Black Hole Accretion:

1. †Reveal how black holes accrete material using resolved movies on event horizon scales

2. †Observe localized heating and acceleration of relativistic electrons on astrophysical scales
3. Detect frame dragging within the ergosphere of a rotating black hole

Jet Launching:

1. †Determine whether jets are powered by energy extraction from rotating black holes
2. †Determine the physical conditions and launching mechanisms for relativistic jets

Transients:

1. Measure the inner jet structure and dynamics in black hole X-ray binaries
2. Detect the kinetic power, physical structure, and velocity in extragalactic transients

New Horizons:

1. Detect proper motions and secular (CMB) parallaxes of Active Galactic Nuclei (AGN) up to ~80 Mpc distances
2. Leverage AGN accretion disk megamasers to measure their AGN host properties

Each of these key science goals is classified as either "Threshold" or "Objective." Threshold science goals define the minimum target. Objective science goals are additional major science opportunities or stretch targets. Here, a dagger (†) indicates a "threshold" science goal.

The Science Traceability Matrix

For each of the Key Science Goals, the highest level operational and technical performance targets for the array are presented in detail in a dedicated special issue of Galaxies with an associated summary publication (Johnson et al., 2023). To aid the decomposition into structured requirements, the key science goals and their performance targets are formally captured in the ngEHT Science Traceability Matrix (STM).

This NASA-specified tool was selected for use due to the concise way in which it shows how science goals and objectives trace to system requirements. The systems engineering and science teams worked closely together to assimilate the community-sourced science objectives into explicit statements that link science goals to measurement requirements, operational configuration, and array performance. Figure 6 provides a high-level summary of the full STM.

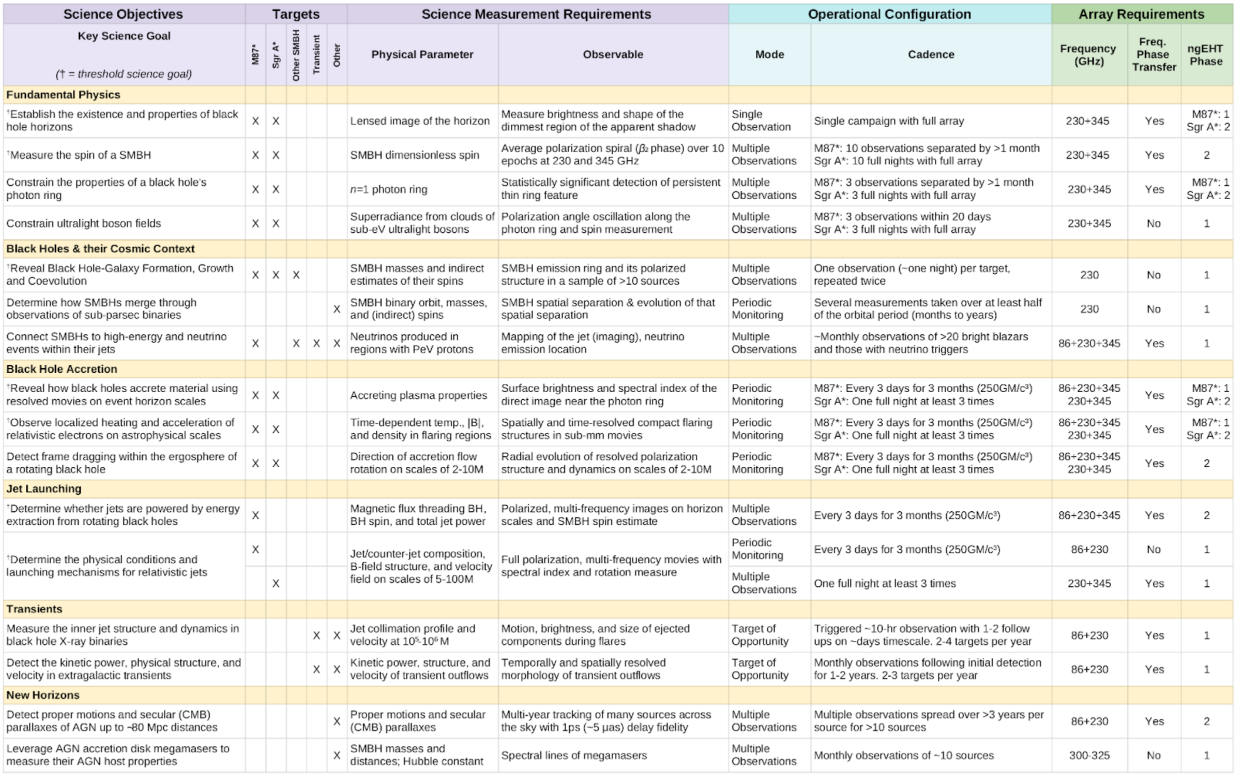

Science Objectives	Targets					Science Measurement Requirements		Operational Configuration		Array Requirements		
Key Science Goal († = threshold science goal)	M87*	Sgr A*	Other SMBH	Transient	Other	Physical Parameter	Observable	Mode	Cadence	Frequency (GHz)	Freq. Phase Transfer	ngEHT Phase
Fundamental Physics												
†Establish the existence and properties of black hole horizons	X	X				Lensed image of the horizon	Measure brightness and shape of the dimmest region of the apparent shadow	Single Observation	Single campaign with full array	230+345	Yes	M87*: 1 Sgr A*: 2
†Measure the spin of a SMBH	X	X				SMBH dimensionless spin	Average polarization spiral (β_2 phase) over 10 epochs at 230 and 345 GHz	Multiple Observations	M87*: 10 observations separated by >1 month Sgr A*: 10 full nights with full array	230+345	Yes	2
Constrain the properties of a black hole's photon ring	X	X				n=1 photon ring	Statistically significant detection of persistent thin ring feature	Multiple Observations	M87*: 3 observations separated by >1 month Sgr A*: 3 full nights with full array	230+345	Yes	M87*: 1 Sgr A*: 2
Constrain ultralight boson fields	X	X				Superradiance from clouds of sub-eV ultralight bosons	Polarization angle oscillation along the photon ring and spin measurement	Multiple Observations	M87*: 3 observations within 20 days Sgr A*: 3 full nights with full array	230+345	No	1
Black Holes & their Cosmic Context												
†Reveal Black Hole-Galaxy Formation, Growth and Coevolution	X	X	X			SMBH masses and indirect estimates of their spins	SMBH emission ring and its polarized structure in a sample of >10 sources	Multiple Observations	One observation (~one night) per target, repeated twice	230	No	1
Determine how SMBHs merge through observations of sub-parsec binaries					X	SMBH binary orbit, masses, and (indirect) spins	SMBH spatial separation & evolution of that spatial separation	Periodic Monitoring	Several measurements taken over at least half of the orbital period (months to years)	230	No	1
Connect SMBHs to high-energy and neutrino events within their jets	X		X	X	X	Neutrinos produced in regions with PeV protons	Mapping of the jet (imaging), neutrino emission location	Multiple Observations	~Monthly observations of >20 bright blazars and those with neutrino triggers	86+230+345	Yes	1
Black Hole Accretion												
†Reveal how black holes accrete material using resolved movies on event horizon scales	X	X				Accreting plasma properties	Surface brightness and spectral index of the direct image near the photon ring	Periodic Monitoring	M87*: Every 3 days for 3 months (250GM/c^3) Sgr A*: One full night at least 3 times	86+230+345 230+345	Yes	M87*: 1 Sgr A*: 2
†Observe localized heating and acceleration of relativistic electrons on astrophysical scales	X	X				Time-dependent temp., |B|, and density in flaring regions	Spatially and time-resolved compact flaring structures in sub-mm movies	Periodic Monitoring	M87*: Every 3 days for 3 months (250GM/c^3) Sgr A*: One full night at least 3 times	86+230+345 230+345	Yes	M87*: 1 Sgr A*: 2
Detect frame dragging within the ergosphere of a rotating black hole	X	X				Direction of accretion flow rotation on scales of 2-10M	Radial evolution of resolved polarization structure and dynamics on scales of 2-10M	Periodic Monitoring	M87*: Every 3 days for 3 months (250GM/c^3) Sgr A*: One full night at least 3 times	86+230+345 230+345	Yes	2
Jet Launching												
†Determine whether jets are powered by energy extraction from rotating black holes	X					Magnetic flux threading BH, BH spin, and total jet power	Polarized, multi-frequency images on horizon scales and SMBH spin estimate	Multiple Observations	Every 3 days for 3 months (250GM/c^3)	86+230+345	Yes	2
†Determine the physical conditions and launching mechanisms for relativistic jets	X					Jet/counter-jet composition, B-field structure, and velocity field on scales of 5-100M	Full polarization, multi-frequency movies with spectral index and rotation measure	Periodic Monitoring	Every 3 days for 3 months (250GM/c^3)	86+230	No	1
		X						Multiple Observations	One full night at least 3 times	230+345	Yes	1
Transients												
Measure the inner jet structure and dynamics in black hole X-ray binaries				X	X	Jet collimation profile and velocity at 10^5-10^6 M	Motion, brightness, and size of ejected components during flares	Target of Opportunity	Triggered ~10-hr observation with 1-2 follow ups on ~days timescale. 2-4 targets per year	86+230	Yes	1
Detect the kinetic power, physical structure, and velocity in extragalactic transients				X	X	Kinetic power, structure, and velocity of transient outflows	Temporally and spatially resolved morphology of transient outflows	Target of Opportunity	Monthly observations following initial detection for 1-2 years. 2-3 targets per year	86+230	Yes	1
New Horizons												
Detect proper motions and secular (CMB) parallaxes of AGN up to ~80 Mpc distances					X	Proper motions and secular (CMB) parallaxes	Multi-year tracking of many sources across the sky with 1ps (~5 μas) delay fidelity	Multiple Observations	Multiple observations spread over >3 years per source for >10 sources	86+230	Yes	2
Leverage AGN accretion disk megamasers to measure their AGN host properties					X	SMBH masses and distances; Hubble constant	Spectral lines of megamasers	Multiple Observations	Monthly observations of ~10 sources	300-325	No	1

Figure 6. Condensed Science Traceability Matrix for the ngEHT Program, highlighting the connection between key science goals, operational configuration, and array

Science Requirements Modeling

Though the STM is highly organized, providing an efficient and powerful view of the needs of science stakeholders, it does not easily integrate into modern requirements management tooling. To achieve this integration, and therefore gain traceability from science to system requirements within the program's requirements management system, the most constraining performance required from the science cases in each science theme was identified. Stakeholder requirements were then defined, again for each science theme, using appropriate requirements language; what results are the L0 Science Requirements, since L0 is defined as the top level of the model hierarchy; see the later section on "The System Model".

These requirements have been further refined through evaluation by an external panel of experts during a formal Science Requirements Review and a subsequent System Requirements Review. They join other L0 Program and Operations requirements to form the complete set of L0 Stakeholder requirements for the program.

The System Model

The system model is the single source of the definitions of many of the elements that need to be managed by the systems engineering process. This covers the entire V-model including requirements, architecture elements (e.g., systems, interfaces, interactions), and V&V (verification and validation). Other activities such as risk management are part of the systems engineering process, but do not have direct reflections in the system model.

Requirements are maintained using Jama Connect. They have been organized into logical groupings based on the type of requirement (functional or non-functional) being specified. These requirements drive the architecture of the system and its constituent lower-level systems, subsystems, and components. The rest of the system model is maintained using Cameo Systems Modeler and model-based Systems Engineering methods. Verification and validation artifacts are created in these tools to maintain linkage with requirements at each level.

The system model is organized into four main levels:

- **Level 0 – System of Systems**
  The level at which stakeholders have influence on or are affected by the ngEHT program
- **Level 1 – Systems**
  The top-most elements of the architecture that have system requirements defined
- **Level 2 – Elements (components of the Systems)**
  These arise from initial functional analysis of the L1 systems where related functions are grouped into elements and identifying where natural or beneficial interfaces exist
- **Level 3 – Subsystems (components of Elements)**
  These are a further decomposition of L2 elements which are considered too complex or risky to specify as a single architectural element

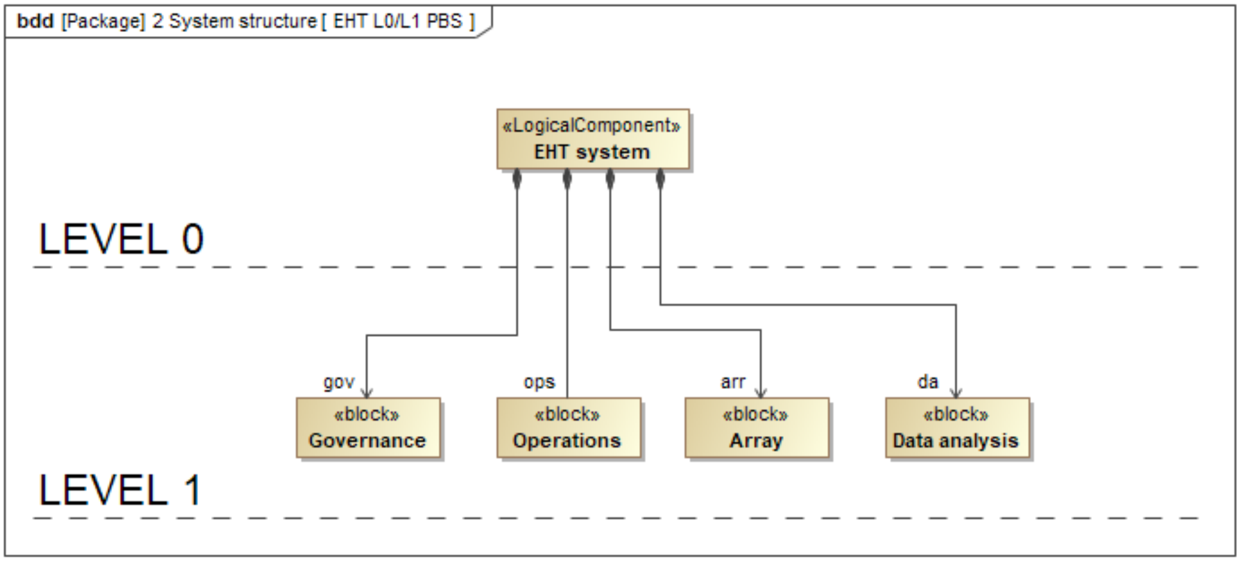

Figure 7. EHT system-of-systems Product Breakdown Structure (PBS)

While a single black-box system solution could have been considered, the decomposition into separate systems is considered to have advantages over that approach. Analysis of the types of things exchanged among the external interfaces of the overall system (the frame of Figure 8) shows that there are some distinct elements that could be better handled by separate systems.

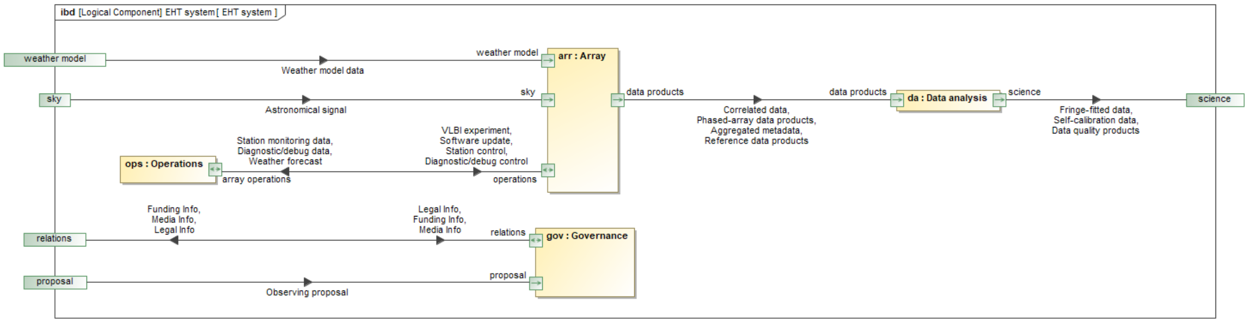

Figure 8. EHT system-of-systems solution architecture

The Array is captured as a separate system from Operations to separate the concerns (and therefore requirements) of operating the array to perform an observation from the functional, performance, and other concerns of performing the observations.

Likewise, the Governance system is treated separately to avoid specifying requirements on the Array for determining which observations should be performed. Additionally, the nature of how the EHT engages with the public and other entities is undefined at this stage of development. Separating the specification of the Array from these matters is intended to reduce the development risk of the Array.

The Array and Data Analysis systems are separate due to the expectation that the way data from specific observations is analyzed may differ based on the science goal or other factors. It is possible that not-yet-developed analysis techniques may be used or required to meet science objectives, but that do not require the Array to do VLBI observations any differently. As with Governance, this separation allows the Array to be specified and designed in isolation from the specific details of post-observation data analysis.

Figure 9 shows a precise cross-discipline logical architecture of the L1 Array system. This architecture captures the logical subsystems of the system of interest and specifies interfaces (not shown in the figure) to determine how these subsystems shall inter-operate with one another and integrate into a whole. Defining the logical subsystems serves the purpose of establishing the work-breakdown structure (WBS) in the systems engineering project.

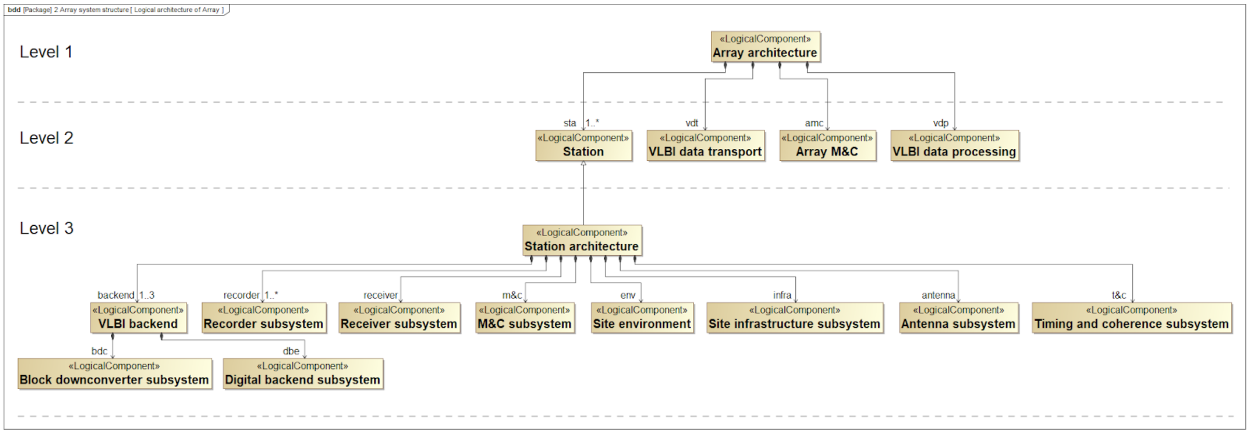

Figure 9. Product Breakdown Structure (PBS) showing system decomposition of the L1 Array system

Traceability

Much of the power of MBSE methods arises from the relationships established between elements of the system model. Manual and tool-based analysis of these relationships enables immediate and low-effort evaluation of the impact of changes and the requirements coverage of verification suites. The model acts as the single source of truth for the definitions of the requirements and the system design. To achieve this, a relationship schema needs to be defined for the system model. Figure 10, taken from the Jama Connect requirements management tool, shows the schema employed for the ngEHT Program.

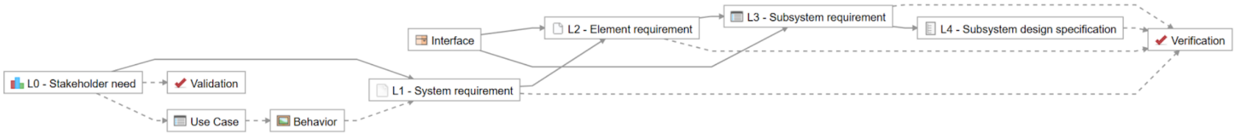

Figure 10. System model element relationships

Direct access to a system model, while powerful, does not meet the wider needs of the project. Only a small fraction of the team are practitioners of MBSE and able to efficiently work directly with the MBSE tools. To make the information in the system model accessible to the larger engineering and science teams, document artifacts are generated. Figure 11 provides a snapshot of the document hierarchy at the current stage of the project, while Figure 12 shows examples of these interconnected relationships within the model. Note that the relationships between document types mirrors the relationships between the elements of the system model.

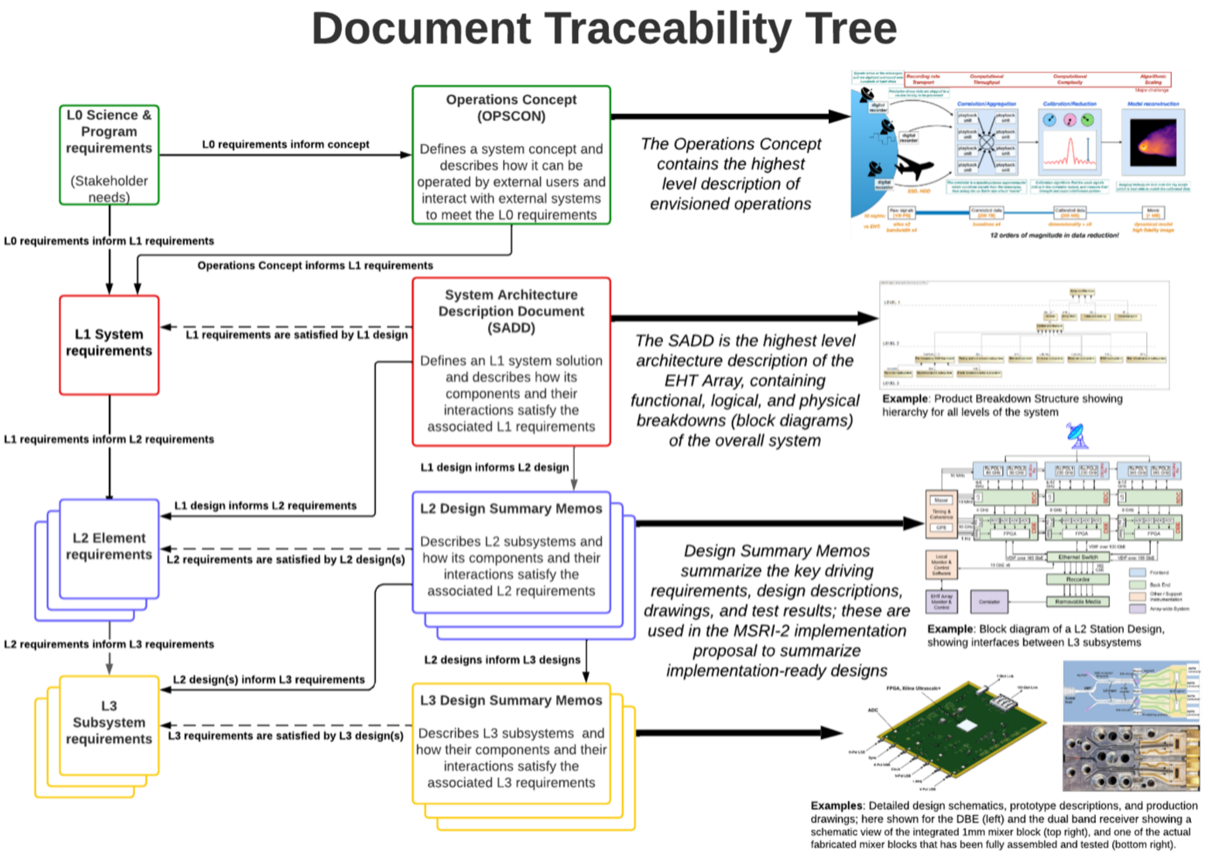

Figure 11. Document traceability tree with examples of project artifacts at each level of abstraction

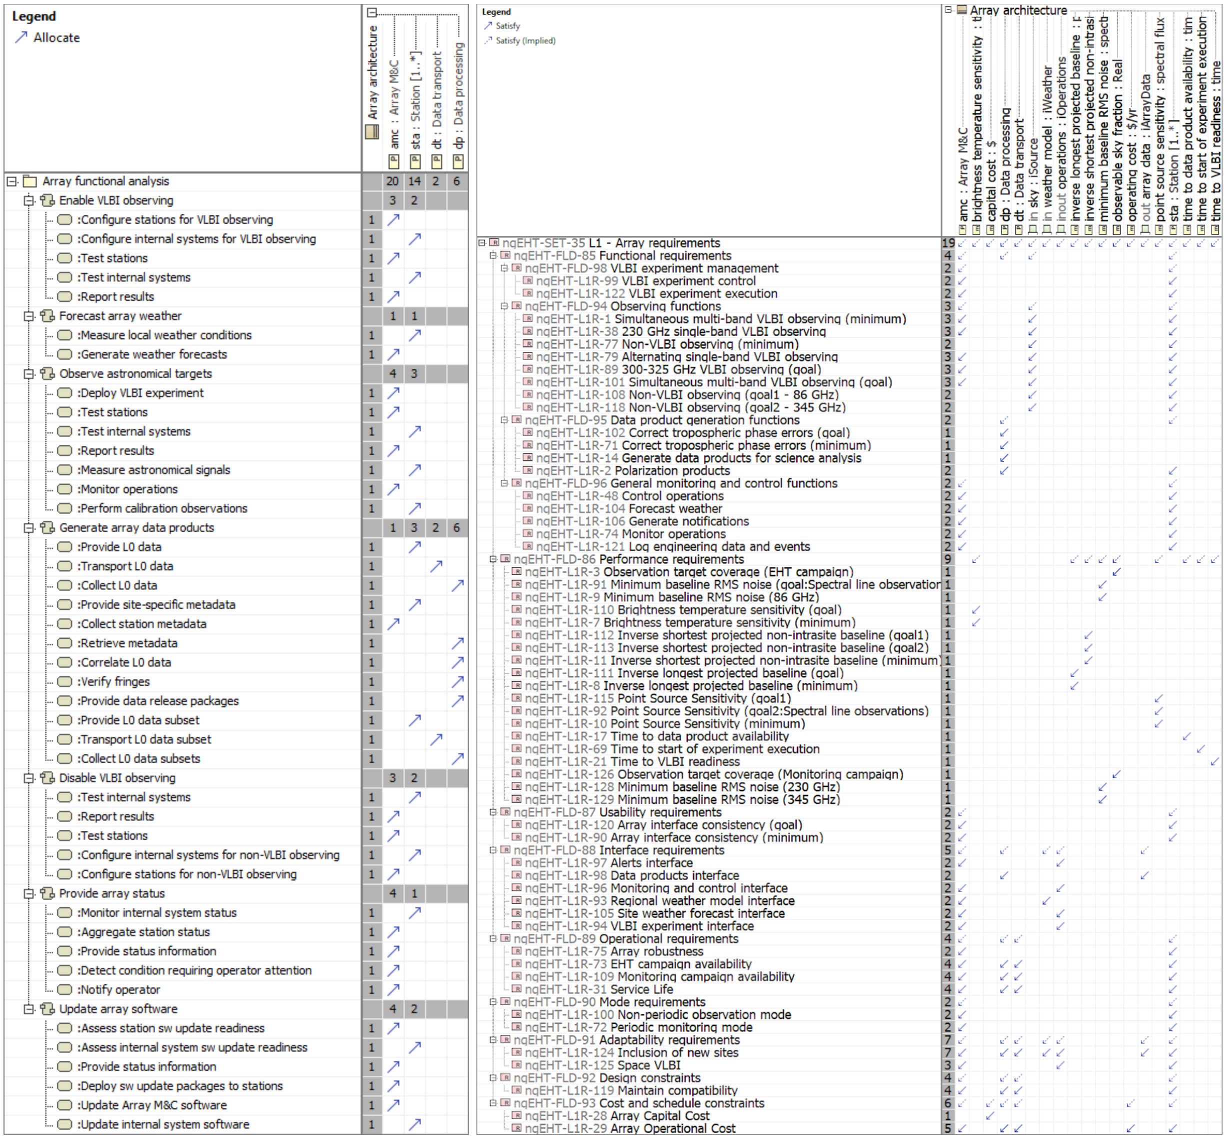

Legend

↗ Allocate

	Array architecture	amc : Array M&C	sta : Station [1..*]	dt : Data transport	dp : Data processing
Array functional analysis		20	14	2	6
Enable VLBI observing		3	2		
:Configure stations for VLBI observing	1	↗			
:Configure internal systems for VLBI observing	1		↗		
:Test stations	1	↗			
:Test internal systems	1		↗		
:Report results	1	↗			
Forecast array weather		1	1		
:Measure local weather conditions	1		↗		
:Generate weather forecasts	1	↗			
Observe astronomical targets		4	3		
:Deploy VLBI experiment	1	↗			
:Test stations	1	↗			
:Test internal systems	1		↗		
:Report results	1	↗			
:Measure astronomical signals	1		↗		
:Monitor operations	1	↗			
:Perform calibration observations	1		↗		
Generate array data products		1	3	2	6
:Provide L0 data	1		↗		
:Transport L0 data	1			↗	
:Collect L0 data	1				↗
:Provide site-specific metadata	1		↗		
:Collect station metadata	1	↗			
:Retrieve metadata	1				↗
:Correlate L0 data	1				↗
:Verify fringes	1				↗
:Provide data release packages	1				↗
:Provide L0 data subset	1		↗		
:Transport L0 data subset	1			↗	
:Collect L0 data subsets	1				↗
Disable VLBI observing		3	2		
:Test internal systems	1		↗		
:Report results	1	↗			
:Test stations	1	↗			
:Configure internal systems for non-VLBI observing	1		↗		
:Configure stations for non-VLBI observing	1	↗			
Provide array status		4	1		
:Monitor internal system status	1		↗		
:Aggregate station status	1	↗			
:Provide status information	1	↗			
:Detect condition requiring operator attention	1	↗			
:Notify operator	1	↗			
Update array software		4	2		
:Assess station sw update readiness	1	↗			
:Assess internal system sw update readiness	1		↗		
:Provide status information	1	↗			
:Deploy sw update packages to stations	1	↗			
:Update Array M&C software	1	↗			
:Update internal system software	1		↗		

Legend

↙ Satisfy

⇙ Satisfy (Implied)

	Array architecture	amc : Array M&C	brightness temperature sensitivity : t	capital cost : $	dp : Data processing	dt : Data transport	in sky : iSource	in weather model : iWeather	inout operations : iOperations	inverse longest projected baseline :	inverse shortest projected non-intrasi	minimum baseline RMS noise : spect	observable sky fraction : Real	operating cost : $/yr	out array data : iArrayData	point source sensitivity : spectral flux	sta : Station [1..*]	time to data product availability : tim	time to start of experiment execution	time to VLBI readiness : time
ngEHT-SET-35 L1 - Array requirements	19	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙	⇙
ngEHT-FLD-85 Functional requirements	4	⇙			⇙		⇙										⇙			
ngEHT-FLD-98 VLBI experiment management	2	⇙															⇙			
ngEHT-L1R-99 VLBI experiment control	2	↙															↙			
ngEHT-L1R-122 VLBI experiment execution	2	↙															↙			
ngEHT-FLD-94 Observing functions	3	⇙					⇙										⇙			
ngEHT-L1R-1 Simultaneous multi-band VLBI observing (minimum)	3	↙					↙										↙			
ngEHT-L1R-38 230 GHz single-band VLBI observing	3	↙					↙										↙			
ngEHT-L1R-77 Non-VLBI observing (minimum)	2						↙										↙			
ngEHT-L1R-79 Alternating single-band VLBI observing	3	↙					↙										↙			
ngEHT-L1R-89 300-325 GHz VLBI observing (goal)	3	↙					↙										↙			
ngEHT-L1R-101 Simultaneous multi-band VLBI observing (goal)	3	↙					↙										↙			
ngEHT-L1R-108 Non-VLBI observing (goal1 - 86 GHz)	2						↙										↙			
ngEHT-L1R-118 Non-VLBI observing (goal2 - 345 GHz)	2						↙										↙			
ngEHT-FLD-95 Data product generation functions	2				⇙												⇙			
ngEHT-L1R-102 Correct tropospheric phase errors (goal)	1				↙															
ngEHT-L1R-71 Correct tropospheric phase errors (minimum)	1				↙															
ngEHT-L1R-14 Generate data products for science analysis	1				↙															
ngEHT-L1R-2 Polarization products	2				↙												↙			
ngEHT-FLD-96 General monitoring and control functions	2	⇙															⇙			
ngEHT-L1R-48 Control operations	2	↙															↙			
ngEHT-L1R-104 Forecast weather	2	↙															↙			
ngEHT-L1R-106 Generate notifications	2	↙															↙			
ngEHT-L1R-74 Monitor operations	2	↙															↙			
ngEHT-L1R-121 Log engineering data and events	2	↙															↙			
ngEHT-FLD-86 Performance requirements	9		⇙							⇙	⇙	⇙	⇙			⇙		⇙	⇙	⇙
ngEHT-L1R-3 Observation target coverage (EHT campaign)	1												↙							
ngEHT-L1R-91 Minimum baseline RMS noise (goal:Spectral line observatior	1											↙								
ngEHT-L1R-9 Minimum baseline RMS noise (86 GHz)	1											↙								
ngEHT-L1R-110 Brightness temperature sensitivity (goal)	1		↙																	
ngEHT-L1R-7 Brightness temperature sensitivity (minimum)	1		↙																	
ngEHT-L1R-112 Inverse shortest projected non-intrasite baseline (goal1)	1										↙									
ngEHT-L1R-113 Inverse shortest projected non-intrasite baseline (goal2)	1										↙									
ngEHT-L1R-11 Inverse shortest projected non-intrasite baseline (minimum)	1										↙									
ngEHT-L1R-111 Inverse longest projected baseline (goal)	1									↙										
ngEHT-L1R-8 Inverse longest projected baseline (minimum)	1									↙										
ngEHT-L1R-115 Point Source Sensitivity (goal1)	1															↙				
ngEHT-L1R-92 Point Source Sensitivity (goal2:Spectral line observations)	1															↙				
ngEHT-L1R-10 Point Source Sensitivity (minimum)	1															↙				
ngEHT-L1R-17 Time to data product availability	1																	↙		
ngEHT-L1R-69 Time to start of experiment execution	1																		↙	
ngEHT-L1R-21 Time to VLBI readiness	1																			↙
ngEHT-L1R-126 Observation target coverage (Monitoring campaign)	1												↙							
ngEHT-L1R-128 Minimum baseline RMS noise (230 GHz)	1											↙								
ngEHT-L1R-129 Minimum baseline RMS noise (345 GHz)	1											↙								
ngEHT-FLD-87 Usability requirements	2	⇙															⇙			
ngEHT-L1R-120 Array interface consistency (goal)	2	↙															↙			
ngEHT-L1R-90 Array interface consistency (minimum)	2	↙															↙			
ngEHT-FLD-88 Interface requirements	5	⇙			⇙			⇙	⇙						⇙					
ngEHT-L1R-97 Alerts interface	2	↙							↙											
ngEHT-L1R-98 Data products interface	2				↙										↙					
ngEHT-L1R-96 Monitoring and control interface	2	↙							↙											
ngEHT-L1R-93 Regional weather model interface	2	↙						↙												
ngEHT-L1R-105 Site weather forecast interface	2	↙							↙											
ngEHT-L1R-94 VLBI experiment interface	2	↙							↙											
ngEHT-FLD-89 Operational requirements	4	⇙			⇙	⇙											⇙			
ngEHT-L1R-75 Array robustness	2	↙															↙			
ngEHT-L1R-73 EHT campaign availability	4	↙			↙	↙											↙			
ngEHT-L1R-109 Monitoring campaign availability	4	↙			↙	↙											↙			
ngEHT-L1R-31 Service Life	4	↙			↙	↙											↙			
ngEHT-FLD-90 Mode requirements	2	⇙															⇙			
ngEHT-L1R-100 Non-periodic observation mode	2	↙															↙			
ngEHT-L1R-72 Periodic monitoring mode	2	↙															↙			
ngEHT-FLD-91 Adaptability requirements	7	⇙			⇙	⇙		⇙	⇙						⇙		⇙			
ngEHT-L1R-124 Inclusion of new sites	7	↙			↙	↙		↙	↙						↙		↙			
ngEHT-L1R-125 Space VLBI	3	↙							↙								↙			
ngEHT-FLD-92 Design constraints	4	⇙			⇙	⇙											⇙			
ngEHT-L1R-119 Maintain compatibility	4	↙			↙	↙											↙			
ngEHT-FLD-93 Cost and schedule constraints	6	⇙		⇙	⇙	⇙								⇙			⇙			
ngEHT-L1R-28 Array Capital Cost	1			↙																
ngEHT-L1R-29 Array Operational Cost	5	↙			↙	↙								↙			↙			

Figure 12. On the left, the allocation of functions to array subsystems is shown as derived from functional analysis. On the right, how elements of the array satisfy L1 array system requirements.

Conclusions

At the time of this paper's submission, the project has just completed its Conceptual Design Review (CoDR) in October 2023 and is now hard at work developing a proposal that will fund the implementation of the project vision. The systems engineering approach has been instrumental in providing clarity to stakeholders, confidence in the array design, and the appropriate level of process and structure to manage the complexity of this very ambitious project.

The landscape of the ngEHT Program is in some ways unique and poses challenges not seen in systems engineering efforts geared towards more traditional projects such as the development of a commercial product. However, the tools and methods used to manage complexity in an academic environment are common to any project entering into a stage of growth from fast-paced, bootstrapped breakthrough to enterprise-level, multi-dimensional implementation scale. The approach to identify stakeholders and harness community buy-in through tools that provide a system-level viewpoint (e.g. system context diagram, science traceability matrix, system architecture model) can be applied on just about any project where the connection between goals and designs is not immediately apparent, and long-range visions require some translation between ambition and actionable plans.

When done properly, systems engineering can become something that is not only begrudgingly accepted but actively championed by scientists and other stakeholders if approached on common grounds. Scientists are highly analytical, objective, and logical. The logic and analysis of a systems approach can appeal on these common grounds. One approach that has been used effectively on the ngEHT Program is to convey the need to apply a corresponding level of rigor and discipline to designing both the project and the system itself as is needed to achieve breakthrough scientific results. In other words, complex science that pushes the limits of an exciting new field requires a system and processes to manage that complexity that is commensurate with the challenge at hand.

References

Doeleman, S.S.; Weintroub, J.; Rogers, A.E.E.; Plambeck, R.; Freund, R.; Tilanus, R.P.J.; Friberg, P.; Ziurys, L.M.; Moran, J.M.; Corey, B.; et al. (2008). Event-horizon-scale structure in the supermassive black hole candidate at the Galactic Centre. Nature 455, 78-80. https://doi.org/10.1038/nature07245

Doeleman, S.S.; Fish, V.L.; Schenck, D.E.; Beaudoin, C.; Blundell, R.; Bower, G.C.; Broderick, A.E.; Chamberlin, R.; Freund, R., Friberg, P.; et al. (2012). Jet-Launching Structure Resolved Near the Supermassive Black Hole in M87. Science, 338, 355. https://doi.org/10.1126/science.1224768

Doeleman, S. S., Barrett, J., Blackburn, L., Bouman, K. L., Broderick, A. E., Fish, V. L., et al. (2023). Reference Array and Design Consideration for the Next-Generation Event Horizon Telescope. Galaxies, 11(5), 107. MDPI AG. Retrieved from http://dx.doi.org/10.3390/galaxies11050107

Johnson, M. D., Akiyama, K., Blackburn, L., Bouman, K. L., Broderick, A. E., Cardoso, V., Fender, R. P., et al. (2023). Key Science Goals for the Next-Generation Event Horizon Telescope. Galaxies, 11(3), 61. MDPI AG. Retrieved from http://dx.doi.org/10.3390/galaxies11030061

Krichbaum, T.P., Graham, D.A., Witzel, A., Greve, A., Wink, J.E., Grewing, M., et al. LBI observations of the galactic center source SGR A* at 86 GHz and 215 GHz. (1998) A&A, 335, L106-L110

Padin, S., Woody, D.P., Hodges, M.W, Rogers, A.E.E., Emerson, D.T., Jewell, P.R., et al. (1998). 223 GHz VLBI Observations of 3C 273. ApJL 1990, 360, L11. https://doi.org/10.1086/185800