TruSat White Paper

I. Introduction and Overview

TruSat is an experimental open source, open-sensor system for creating a trusted record of satellite orbits. TruSat is primarily designed to enable the assessment of satellite operations in light of space sustainability standards. Space sustainability is about preserving the use of outer space, and all of its socioeconomic benefits, for present and future generations.[1] As a practical matter, this means mitigating orbital debris, which can render orbits around Earth unusable for generations. Amid projections of a tenfold or more increase in the number of satellites in low Earth orbit ("LEO"), avoiding collisions between satellites is an increasingly urgent focus for space sustainability, and enlightened satellite operators, governments, and civil society are working to define guidelines, best practices, and standards for sustainable orbital operations. A crucial gap in these efforts is the absence of an open, widely-trusted source of data about the orbital position of satellites accessible for use in assessing compliance with these sustainability standards.

To fill this data gap, the TruSat System is designed to enable the emerging space sustainability community to "task" a global network of citizen satellite observers to track satellites of interest, utilizing ubiquitous consumer hardware, and to assemble observations from around the planet into a trusted record of orbital positions suitable for measuring orbital behavior against sustainability standards. The TruSat System comprises three elements:

- Software for prioritizing satellite observations, assisting amateur satellite observers, and processing observations into orbital predictions;
- Observers who make and report satellite observations; and
- An interface with the space sustainability community for aligning the System's satellite observation priorities with sustainability priorities.

Section II of this white paper surveys the space sustainability landscape and situates TruSat within it. It describes an externality in the liability framework for space that leaves the incentives of satellite operators misaligned with the common interest in the long-term sustainability of spaceflight. On the one hand, satellite operators are not generally made to internalize the costs of collisions in orbit. On the other, many sustainability measures—such as maneuvers to avoid a potential collision or for end-of-life disposal—entail a cost to the operator, in terms of the operational life of the satellite. Faced with a probabilistic collision warning, an operator's incentives may point to rolling the dice rather than maneuvering.

Growing concern about space sustainability by the general public and regulators alike, coupled with the emergence of measurable, verifiable standards for sustainable satellite operations, could raise the costs to

1 *See*, e.g., Guidelines for the Long-term Sustainability of Outer Space Activities, UN Doc. A/AC.105/2018/CRP.20 (2018).

operators of unsustainable operations, bringing their incentives more in line with the long-term sustainability of outer space. However, the efficacy of sustainability standards depends on independent assessment of satellite operators' conformity with them. This assessment function, in turn, requires a freely available source of orbital position data trusted by all involved. This element is missing. The trustworthiness of space situational awareness ("SSA") data collected and publicly shared by a government has been widely questioned on account of national interests sometimes at odds with full transparency about satellite positions, and that governments, for national security reasons, neither share the full extent of their SSA data nor the sensor data or algorithms underlying it that would enable independent assessments of its accuracy. The business interests of commercial SSA data providers, which sell SSA data and analysis to satellite operators, raise similar questions about whether their data would be trusted by all parties, particularly in matters concerning present or prospective customers.

TruSat is architected for trust. Whereas trust in existing sources of SSA data depend on trust in the humans and institutions in the loop—that the orbital prediction reflects a competent analysis of the sensor data, free of any institutional interests—TruSat substitutes transparent, verifiable algorithms for institutional input. While humans are involved in generating the initial satellite observation data fed into the System, TruSat's Proof of Satellite software engine is entirely automated, leaving no room for tampering.[2] Unlike existing sources of SSA data, the entirety of the algorithms that translates individual satellite observations into an orbital prediction with a confidence assessment—the confidence factors applied, and their weighting—are transparent, enabling any orbital prediction to be reverse-engineered. Built atop the Ethereum blockchain, TruSat will periodically check its file and data integrity against tamper-evident blockchain records, ensuring that the algorithms in effect at any given time are those approved by the TruSat Community. Additional verifiability of TruSat's orbital predictions comes courtesy of nature: any person may physically observe a satellite in the TruSat catalog to verify the accuracy of the orbital prediction.

Section III provides a detailed technical overview of the Proof of Satellite software engine, illuminating the engineering challenges of creating an orbital location system that is open and permissionless (anyone may feed data into the system), and automated (the system evaluates and weights data, rather than a human administrator), and how each is addressed in the v0.1 release. This section surveys foreseeable attack vectors—from malicious attempts to subvert the system by submitting false observations, to simple mistaken observations—and how the application of confidence factors mitigate each attack vector. Section III likewise explains how the same confidence factors are applied to provide a confidence assessment for each orbital prediction. These confidence factors and the algorithms applying and weighting them will be the subject of ongoing tuning and revision by the TruSat Open Source Community.

Section IV describes the governance arrangements for the TruSat System—who sets the priorities for observation, according to what process—as well as for the TruSat Open Source Community of contributors. The TruSat Charter sets forth the governance arrangements for the initial, experimental phase. The Charter

[2] For a discussion of how TruSat mitigates efforts to tamper with orbital predictions by submitting false observations, *see* Section III.C: Mitigating Attack Vectors.

delineates the commitments of the TruSat Partners—organizations committed to space sustainability—to maintaining and advancing TruSat, and specifies the process by which the Partners may "task" the TruSat System, ensuring alignment of the System's observation priorities with space sustainability priorities. The Charter likewise governs the TruSat Open Source Community, including processes for stakeholder input on feature roadmaps, and open source licensing for contributed code and documentation. The TruSat Charter is publicly available at www.TruSat.org, providing full transparency of the inner workings of the System. As TruSat progresses beyond its experimental phase, the TruSat Partners intend to progressively transition System and Community governance to blockchain-enabled collaboration and voting tools.

Section V lays out a provisional roadmap for further development of the TruSat software over 12-18 months. The version 0.1 release is a relatively limited prototype intended primarily to validate and refine the Proof of Satellite engine. A version 0.2 release later in Q4 2019 will test the "Mission" functionality enabling space sustainability advocates to task the TruSat System to observe the satellites of highest priority for sustainability purposes. These initial alpha releases will rely primarily on the dedicated community of satellite observers testing and feedback. Subsequent releases will add features designed to scale the community of satellite observers—by easing the process of making and reporting satellite observations, and by increasing the variety of observation and input methods accommodated by the software—and fortifying the trust architecture by migrating to Ethereum mainnet and progressively decentralizing the System's governance.

Section VI locates TruSat within the broader mission of ConsenSys Space to diversify, democratize, and decentralize space endeavors, and it's vision for Earth's Space Program. By creating new opportunities to be a part of the solution to the wicked problem set threating the long-term sustainability of outer space, and by removing barriers to participation, TruSat will make important strides in democratizing and diversifying space endeavors. As an experiment in decentralizing space situational awareness, TruSat may well illuminate a new path to building the foundation of trust in government SSA systems necessary for functional space traffic management. As an experiment in global, bottom-up collective action, TruSat will begin to build the collaborative muscle fibers needed to realize the vision of Earth's Space Program.

Section VII outlines opportunities to join and contribute the TruSat Open Source Space Sustainability Community. Transforming TruSat into a powerful tool for the long-term sustainability of outer space, and a global community advancing our future in space, will depend on a diverse range of contributions from a diverse range of people. From observing satellites to contributing code to build out TruSat's feature set, to writing or translating documentation, helping to refine the algorithms of the Proof of Satellite engine, or organizing community meetups, there are a wide variety of opportunities to contribute.

Section VIII explains why ConsenSys, a privately-held corporation, is investing in an open source space sustainability initiative aimed at generating a global public good rather than profit. ConsenSys believes that enabling large-scale collective action is among the most transformative potential applications of the Ethereum blockchain technology, and is investing in solving space-related collective action challenges as R&D to accelerate the realization of this potential.

II. TruSat as a Space Sustainability Tool

This section introduces the challenges to the long-term sustainability of spaceflight, and locates TruSat within an emerging space sustainability ecosystem.

II.A. Transparency and Trust Challenges to Space Sustainability

Space sustainability is about preserving the use of outer space, and all of its socioeconomic benefits, for present and future generations. The primary threat to the long-term usability of space is Earth-orbital debris: non-functional spacecraft, pieces of spacecraft discarded in the course of space missions, and fragments from collisions between spacecraft, and from spacecraft destroyed by weapons tests. While outer space is infinitely large, the orbital positions around Earth suitable for the myriad space applications on which we depend are finite and increasingly congested.

Global public awareness and concern about space sustainability is rising with the deployment of the first of thousands of satellites comprising so-called "mega-constellations" to provide broadband internet service from low Earth orbit ("LEO"). To put these constellations in perspective, the satellites planned for deployment by just three companies would add approximately 16,000-46,000 satellites to the LEO environment in the 2025-2030 time horizon; roughly a ten-to-twenty-five-fold increase above the present population of operational satellites.[3] As LEO grows increasingly congested, so grows the risk not only of collisions—between operational satellites or debris—but also catastrophic, cascading chains of collisions rendering vital swaths of Earth orbits unusable for generations.

In recent years, government regulators, industry, and civil society have identified (and required, in some cases) sustainability measures to lessen the chances of satellite operations generating orbital debris.[4] Many are implemented well before launch, such as spacecraft design and testing regimes intended to prevent "bricks:" defective satellites that are essentially debris upon launch. Other pre-launch sustainability

[3] The lower range of the estimate is based upon regulatory filings and public statements by SpaceX, OneWeb, and Amazon's Project Kuiper. The Federal Communications Commission (FCC) authorized 7,518 of the V-band satellites, and 4,409 of Ku- and Ka-band satellites comprising SpaceX's Starlink constellation. Amazon subsidiary Kuiper, LLC has sought approval of a constellation that "will consist of 3,326 satellites in 98 orbital planes at altitudes of 590 km, 610 km, and 630 km." Whereas OneWeb received FCC authorization for a constellation of 720 satellites in 2017 and indicated it was considering adding an additional 1,972 satellites, more recent public statements suggest the constellation will be fully operational at 648 satellites. The upper range of the estimate accounts for the October 2019 regulatory filings by SpaceX for an additional 30,000 Starlink satellites.

[4] Among the space sustainability measures beyond the scope of this brief introduction are diplomatic efforts to proscribe the use and testing of debris-generating anti-satellite weapons, as well as emerging mission concepts for active debris removal ("ADR"), by which nonfunctional space objects are removed from orbit to mitigate collision risks.

measures go to mission design, including the altitude of operations; satellites at higher altitudes remain in orbit for decades, or even centuries, unless boosted to a disposal orbit. Whereas such pre-launch sustainability measures are vital to the long-term sustainability of space flight, measures whose implementation can be verified on the ground are not the primary focus of TruSat. Instead, TruSat is designed for space sustainability measures that involve orbital operations, including the disposal of a satellite at the end of its operational life, and decisions whether or not to maneuver to avoid a collision risk. Verifying implementation of such measures requires knowledge of satellites' orbital positions at various points in time. Such information about orbital positions, used to predict potential collisions between satellites or debris, is commonly referred to as space situational awareness ("SSA"). [5]

Collision avoidance maneuvers present a particularly acute imperative for standards and verification. Unlike air traffic, which is subject to formal, mandatory collision avoidance regimes, no comparable system exists for space traffic management ("STM"). In orbit, collision avoidance rests on voluntary decisions by satellite operators to maneuver to avoid a potential collision, and relatively loose coordination between operators. A small number of national governments collect SSA data for national security purposes, and screen their orbital object catalogs for "conjunctions:" predictions for objects to pass close enough to be a collision risk. For example, the United States Department of Defense utilizes a variety of sensors to track objects in orbit, and provides conjunction warnings to satellite operators with which it has entered into space situational awareness sharing agreements. Well-resourced satellite operators may supplement government-provided SSA information with commercial SSA services, as well as their own proprietary location data from on-board GPS and ground station telemetry.

A recent, high-profile conjunction between an Earth observation satellite operated by the European Space Agency ("ESA") and a satellite in the Starlink constellation operated by SpaceX offers some insight into the informal, ad hoc, and voluntary nature of space traffic management. In late August, 2019, the U.S. Department of Defense assessed a potential conjunction between ESA's Aeolus satellite and SpaceX's Starlink44.[6] As the probability of a collision came within 1 in 1,000—ten times higher than ESA's internal threshold for collision avoidance maneuvers—ESA operators contacted their SpaceX counterparts, who declined to execute a maneuver.[7] Half an orbit before the potential collision, ESA triggered a series of thruster burns to provide a safe buffer between Aeolus and Starlink44.[8] In a subsequent statement, SpaceX indicated that its decision not to maneuver was based on earlier SSA data placing the collision risk at approximately 1 in 50,000, and that "a bug in our on-call paging system prevented the Starlink operator from seeing the follow on correspondence on this probability increase." [9] Commenting on the incident,

[5] P.J. Blount, Space Traffic Management: Standardizing On-Orbit Behavior, 113 AJIL UNBOUND 120 (2019) (defining space situational awareness as "information about what is in orbit, where it is at a given time, and who (if anyone) controls it.").

[6] European Space Agency, *ESA Spacecraft Dodges Large Constellation* (September 3, 2019).

[7] *Id.*

[8] *Id.*

[9] Jeff Foust, *ESA spacecraft dodges potential collision with Starlink satellite*, SPACENEWS (September 2, 2019).

ESA's Head of Space Safety reflected: "This example shows that in the absence of traffic rules and communication protocols, collision avoidance depends entirely on the pragmatism of the operators involved."[10]

The Aeolus-Starlink44 conjunction illustrates two realities about the present state of space traffic management. First, there is no regulatory authority directing satellite operators to maneuver to avoid collisions; it is up to each operator of a maneuverable satellite to decide whether or not to maneuver. Maneuvers use fuel, thereby diminishing the useful life of the satellite and the revenue it may generate. As such, the operator's decision whether or not to maneuver is essentially a business decision, balancing of the costs of a maneuver with the probability and costs of a collision. Second, the collision risk is uncertain even for operators with access to the highest resolution SSA data and analysis. Whereas the SSA information made publicly available by the U.S. Department of Defense initially placed the Aeolus-Starlink44 collision probability at approximately 1 in a million, ESA had access to more refined information provided by the Department of Defense pursuant to an SSA sharing agreement.[11]

Crucially, as satellite operators evaluate probabilistic conjunction warnings and weigh the costs of a maneuver against the probability and costs of a collision, their incentives are not aligned with the common interest in the long-term sustainability of spaceflight. This is because satellite operators would not be made to internalize the full costs of a collision in orbit. This externality is a function of present legal and practical realities of the space domain. Satellite operators generally obtain insurance coverage for their own economic losses, such as the loss of a satellite damaged in a collision and the corresponding loss of revenue. However, the present implementation of the international legal framework for outer space makes it highly unlikely that private satellite operators would be made to bear the broader costs of a collision: to operators of other satellites damaged in a collision, or from the loss of orbits littered with debris. For one, liability for damage caused by a space object to third parties under international law runs to the national government(s) responsible launching the object into orbit, and few such governments have shown an appetite for indemnification arrangements to pass this liability down to the private operator responsible for a collision.

As a practical matter, third party liability for collisions in orbit is predicated on "fault."[12] While "fault" is not defined in the relevant treaties, establishing that a satellite operator was at fault for a collision would almost certainly require an account of the relevant satellites' orbits trusted by all parties to the dispute.

[10] ESA, supra note 6.

[11] *See* Foust, *supra* note 9 ("[ESA Head of Space Safety] Krag, though, said that while SOCRATES uses publicly available information on spacecraft orbits, known as two-line elements (TLEs), satellite operators like ESA and SpaceX have access to more accurate orbital information provided by the Air Force 'about one order of magnitude better than TLEs.' That, combined with operators' own knowledge of spacecraft positions, yielded a 'more credible' collision probability of 1 in 1,000 that led to the decision to perform the maneuver.")

[12] *See* Convention on International Liability for Damage Caused by Space Objects, Mar. 29, 1972, 961 UNTS 187.

By way of analogy, establishing fault in an automobile collision relies upon eyewitness accounts of the movements of each vehicle prior to the collision. There are few "eyewitnesses" as to the orbital behavior of satellites, and it is doubtful that all parties to a dispute would trust them. Substantial doubts have been raised about the trustworthiness of a national government as a sole source of SSA data, given that national interests may not always be aligned with transparency in orbital behavior.[13] In addition, for national security reasons, agencies such as the U.S. Department of Defense do not share the sensor data or the algorithms underlying its SSA data, limiting the ability for independent validation of results. Similarly, cases can be foreseen in which commercial providers of SSA data and analytics may not be trusted by all parties in a dispute involving one or more of their customers.

To summarize the problem space to which TruSat is addressed: a tenfold increase in the LEO satellite population will increase the frequency of conjunctions (i.e., predictions that satellites will pass close enough to present a collision risk), and thus the frequency of operators' business decisions whether or not to maneuver to avert collisions. The true costs of a collision are not accounted for in operators' decision calculus because it is exceedingly unlikely that a private satellite operator would bear the costs of to third parties of a collision under the international legal framework for space, as presently implemented. Among the practical reasons for this externality are the absence of an accessible source of SSA data trusted by all parties to determine fault in a collision. Given this remote possibility of liability to third parties, insulating private satellite operators from the true costs of collisions, and the costs to operators of maneuvers, operators may find an incentive to "roll the dice," gambling with the long-term sustainability of outer space, rather than to execute a precautionary maneuver.

II.B. The Emerging Space Sustainability Ecosystem

As explained in the foregoing section, the space liability framework does not supply sufficient incentives for an operator faced with a conjunction warning to resolve uncertainty in favor of a precautionary maneuver. At present, the court of public opinion may be closer to minds of satellite operators weighing a maneuver than a court of law or arbitral tribunal. With public awareness of space sustainability challenges rising with the advent of "mega-constellations," constellation operators in particular are touting their sustainability practices, for public relations purposes, and possibly also to preempt the concerns of their regulators.[14] This desire on the part of satellite operators—soon to be competing for internet customers—to positioning their brands as responsible, sustainable stewards of the space environment, may prove to be a more potent lever for bringing the incentives of satellite operators into alignment with the collective interest in the long-term sustainability of outer space.

13 *See, e.g.,* Blount, *supra* note 5.

14 Constellation operator OneWeb has been particularly forward-leaning in positioning itself as a leader in space sustainability. *See, e.g.,* OneWeb's space sustainability website www.responsible.space; Jeff Foust, *OneWeb founder Wyler calls for responsible smallsat operations*, SPACENEWS (August 6, 2019).

Converting public awareness and concern for space sustainability, and the concomitant value to private spacecraft operators of branding their activities as sustainable, into *sustainable orbital operations*, requires two basic elements:

1. **Standards**: Responsible, sustainable space operations must be defined in a measurable, verifiable way.
2. **Assessment:** Operations must actually be graded for conformity with sustainability standards. The assessment function requires both:
    a. a mutually trusted source of data about the orbital behavior of satellites, and
    b. analysis capacity to measure orbital behavior against sustainability standards.

Governments are signaling a policy preference for standards for sustainable orbital operations, but are unlikely to lead in their formulation. In June of 2018 the United States President issued Space Policy Directive-3, National Space Traffic Management Policy, signaling "A STM framework consisting of best practices, technical guidelines, safety standards, behavioral norms, pre-launch risk assessments, and on-orbit collision avoidance services is essential to preserve the space operational environment."[15] A year later, In June of 2019 the United Nations Committee on the Peaceful Uses of Outer Space adopted the Guidelines for the Long-term Sustainability of Outer Space Activities ("LTS Guidelines"),[16] the culmination of nearly a decade of intergovernmental negotiations and representing a global political consensus on the importance of space sustainability. In September of 2019, the European Union announced the formation of a "Safety, Security and Sustainability of Outer Space (3SOS) public diplomacy initiative to "reach a common understanding of all space actors in all parts of the world on responsible and sustainable behavior."[17] The United Nations' LTS Guidelines illustrate why governmental and intergovernmental initiatives are unlikely to produce specific, measurable, verifiable standards for orbital operations, but instead create fertile conditions for other players in the emerging space sustainability ecosystem to do so. Negotiated and adopted by consensus in a global intergovernmental forum, the LTS Guidelines are much too general to enable spacecraft operators or the consuming public to delineate sustainable from unsustainable operations.

Non-governmental actors are beginning to step in to fill this crucial gap. Prominent among them are the recently-formed Space Sustainability Rating ("SSR"), a collaboration of the World Economic Forum, the MIT Media Lab, the University of Texas at Austin, the European Space Agency, and Bryce Space and Technology. SSR seeks to establish standards for a space sustainability certification much like LEED certification in the construction sector, providing brand and other business incentives for meeting sustainability standards.[18] The Secure World Foundation, with a mission of "promoting cooperative solutions for space sustainability," is another example of non-governmental leadership in convening spacecraft operators and civil society

15 For a survey of initiatives to develop STM standards, guidelines, and best practices responsive to the mandates of Space Policy Directive-3, *see* Michael P. Gleason, Establishing Space Traffic Management Standards, Guidelines, and Best Practices, THE AEROSPACE CORPORATION (September 2019).

16 A/AC.105/2018/CRP.20.

17 *See* Jeff Foust, *EU agency starts space sustainability initiative*, SPACENEWS (September 15, 2019).

18 *See A Sustainability rating for space debris,* MIT News (May 6, 2019).

experts toward defining sustainable orbital operations. Most recently, the Space Safety Coalition was formed in September of 2019, with a diverse membership spanning the space industry and civil society, and published the first iteration of its Best Practices for the Sustainability of Space Operations.[19] Whereas the September 16, 2019 iteration of the Best Practices does not prescribe measurable standards for collision avoidance, it does signal that "future efforts may be warranted to… Address maneuver prioritization in the event that two spacecraft with maneuver capability conjunct…." It is foreseeable that voluntary standards for sustainable orbital operations will emerge from these multi-stakeholder initiatives in the near future. There is ample precedent for voluntary technical standards for space operations being subsequently adopted and mandated by national regulators.[20]

The efficacy of voluntary space sustainability standards as an incentive for responsible orbital operations depends on accountability through independent assessment. Whereas many space sustainability practices can be verified on the ground, prior to launch, independent verification of emerging standards for sustainable orbital operations will involve interpretations of SSA information to measure satellite operations against standards. This independent analysis capacity is beginning to emerge; Professor Moriba Jah's ASTRIA Lab at the University of Texas at Austin is a prominent example. With his expertise in astrodynamics and access to the supercomputer resources of the University of Texas, Professor Jah has demonstrated a capacity and willingness to analyze SSA data to call attention to unsustainable operations in orbit.

The missing ingredient in this emerging space sustainability ecosystem is a widely-trusted source of data suitable for assessing orbital operations against sustainability standards. As explained in Section II.A., U.S. Government SSA data suffers the same trust deficit as any SSA data sourced from a single national government, and the Government reserves the right to restrict its use. Private commercial providers of SSA data rely on satellite operators for their revenue and do not have incentives for calling out non-conforming orbital behavior. TruSat is designed to fill this gap through a new approach to SSA uncoupled from government or commercial interests.

II.C. TruSat as a New Approach to Transparency and Trust in SSA

To fill the trusted data gap, TruSat is designed to enable space sustainability advocates to task a global network of citizen satellite observers to track satellites of interest, utilizing ubiquitous consumer hardware, and to assemble observations from around the planet into a trusted record of orbital positions suitable for measuring orbital behavior against sustainability standards. Whereas the initial releases of TruSat are an experiment in producing reliable orbit predictions from amateur visual observations, the TruSat architecture is sensor-agnostic and could support inputs from a range of academic, commercial, or governmental institutions. The TruSat System comprises three elements:

19 www.spacesafety.org; Jeff Foust, *New coalition seeks to improve space safety*, SPACE NEWS (Sept. 18, 2019).

20 *See*, e.g., Gleason, *supra* note 20.

- **Software** for prioritizing satellite observations, assisting amateur satellite observers, and processing observations into orbital predictions. As detailed in Section III, TruSat's Proof of Satellite engine utilizes the physics of orbital mechanics to calculate accurate, actionable orbital positions from amateur visual observations of a satellite from multiple points on Earth. It achieves accuracy through diversity (of observers, geography) rather than relying on exquisite sensors.
- **Observers** who make and report satellite observations. TruSat is sensor-agnostic and can utilize visual observations made with hardware ranging from binoculars and a stopwatch to digital cameras, software-enabled consumer telescopes and internet-remote telescopes (e.g., iTelescope). As outlined in Section V, while the v0.1 release of the TruSat software will require observers to utilize external software to convert visual observations into initial orbit determinations (IODs), subsequent releases will add support for direct inputs from a range of methods (e.g., automated IOD extraction from a digital image).
- An **interface with the space sustainability community** for aligning the System's satellite observation priorities with sustainability priorities. During TruSat's initial, experimental phase, the TruSat Partners—institutions committed to maintaining and advancing TruSat as a space sustainability tool—will be responsible for aligning the System's observation priorities with broader space sustainability goals, in accordance with the procedures set forth in the TruSat Charter.

Community-sourcing orbital object location data through visual observations of satellites is not new. In fact, Operation Moonwatch, a global citizen-science initiative to track the first artificial satellites at the dawn of the space age, was the original SSA network.[21] More recently, the SeeSat community has been tracking satellites for decades—with tools ranging from binoculars and a stopwatch to relatively advanced digital camera equipment—and posting the angles-only visual observations ("IOD" or other reporting formats) to an internet mailing list.[22] TruSat is designed to scale this method of satellite tracking from hobby to a source of truth in orbital behavior trusted by all space actors. What is most fundamentally new about TruSat is its open, decentralized, and automated architecture designed to remedy the trust challenges of existing public sources of SSA data. TruSat's trust architecture is summarized by comparison to existing SSA sources in the subsection that follows, and detailed in Section III.

II.C.1. Leveraging Blockchain Technology in Service of Openness, Transparency, and Trust

The trust deficit of existing sources of SSA data, as applied to space sustainability applications, results from the combination of two attributes. First, as the data is controlled by a single government or business, from

21 *See* W. Patrick McRay, KEEP WATCHING THE SKIES! THE STORY OF OPERATION MOONWATCH AND THE DAWN OF THE SPACE AGE (Princeton University Press, 2008); Operation Moonwatch, WIKIPEDIA (retrieved September 17, 2019); David Dickinson, *Citizen Science, Old-School Style: The True Tale of Operation Moonwatch*, UNIVERSE TODAY (March 15, 2013).

22 SeeSat-L Home Page, http://www.satobs.org/seesat/seesatindex.html.

sensor to analysis, trust in the results is limited to trust in that institution. The application of SSA data to assessing compliance with sustainability standards presents a more difficult trust equation than the case of a satellite operator keeping planning a maneuver. The latter may have built up confidence and trust in given SSA data providers over the course of a long working relationship, and have access to multiple sources of SSA data to verify accuracy. The legitimacy of sustainability standards assessment, by contrast, rests on trust by a much broader, global set of stakeholders. Finding a single institution trusted by all stakeholders in all cases borders on impossible. Irrespective of an institution's track record for accurate results, perceptions of institutional interests in a given case can undermine trust in the data produced by that institution. Compounding this structural trust challenge is the reality that the proprietary nature of SSA systems—with many utilized primarily for national security missions—does not permit independent verification of the results.

TruSat's decentralized architecture removes any institution or individual from the trust equation. In place of an institutional arbiter between sensor data and orbit predictions, TruSat substitutes transparent, verifiable algorithms, which automate the process of determining and refining orbits. Unlike existing sources of SSA data, the entirety of the algorithms that translate individual observations of a satellite into an orbit prediction with a confidence assessment—the confidence factors applied, and their weighting—are transparent, allowing any orbit prediction to be independently assessed.

In the absence of an institution trusted by all stakeholders to the orbital position of satellites, TruSat dispenses with the need to trust any institution. Instead, trust in TruSat's orbital position data derives from the transparency of the observations and algorithms underpinning each orbit prediction, and that any attempt to tamper with the algorithms or output would be evident. This decentralized, tamper-evident architecture is enabled by building TruSat atop the Ethereum blockchain. TruSat will periodically compare its code base and database with tamper-evident blockchain records, ensuring that the algorithms in effect at any given time are those approved by the TruSat Community. Furthermore, the individual data submissions will be secured and logged in "on-chain" transactions, with each observer's ability to provide weighted-contributions to the TruSat dataset be secured by public/private key cryptography through access to on-chain wallet credentials.

The inputs of ConsenSys Space and the TruSat Partners are limited to refining the software and communicating space sustainability observation priorities, in accordance with the TruSat Charter (see Section IV, TruSat Governance). The source code for each update to TruSat's codebase will be available in the TruSat software repositories, ensuring the transparency of the algorithms in effect. While TruSat does not permit human input in processing satellite observations into orbit predictions, the System does rely on human input to make and report satellite observations. Section III.C. details attack vectors ranging from malicious attempts to influence orbital position data to erroneous observations, and how the TruSat software is designed to mitigate the influence of such inputs.

II.C.2. A Toolset for Measuring Orbital Behavior against Sustainability Standards

The easiest sustainability application for TruSat, as presently conceived, is probably verifying post-mission disposal requirements. Many regulatory authorities require that satellites in geosynchronous orbit be raised to a so-called “graveyard” orbit at the end of their useful life, to reduce the potential for collision. More recently, with the advent of “mega-constellations” comprising thousands of satellites in low Earth orbit (“LEO”), regulatory authorities have required operators to place satellites in a disposal orbit--from which the satellite will re-enter Earth’s atmosphere and burn up within one year--at the end of their operational life. Beyond regulatory requirements, post-mission disposal standards are likely to form part of voluntary sustainability standards.

Consider a hypothetical case in which an operator of a LEO constellation announces it has moved a satellite into a disposal orbit, in conformity with its sustainability commitments and/or regulatory requirements, and submits its own information to TruSat for verification. From an orbit prediction derived from the operator’s data, TruSat contributors around the world know where and when to look to observe the satellite, and submit their observations through the TruSat interface. From these observations around Earth, the Proof of Satellite engine generates and orbital prediction and confidence assessment, confirming or disputing the satellite operator’s claims. In this hypothetical, the satellite operator is cooperating with TruSat to obtain the benefit of independent verification of its sustainability practice. But the system does not depend on such cooperation; it could be used to verify post-mission disposal claims without satellite ephemeris data supplied by the operator.

This end of life disposal hypothetical admittedly does not make full use of TruSat’s potential as a sustainability tool. The movement of a satellite from its operational orbit to a disposal orbit is such a large change as not to require precision measurement to verify.

What use cases would make fuller use of a public eyewitness capability in orbit? Consider a conjunction scenario along the lines of the recent Aelous-Starlink44 incident. A conjunction analysis predicts maneuverable LEO satellites operated by Operators A and B have a 1 in 10,000 chance of collision in 72 hours. Assume the presence of a voluntary standard or guideline calling for any maneuverable satellite to execute a maneuver to avoid a collision risk greater than 1 in 10,000. Now imagine TruSat users on multiple continents observing both satellites, creating indelible records of their orbit tracks. Does the world watching, literally, change the decision calculus of Operators A or B? Does it lead either operator to resolve uncertainty in favor of a precautionary maneuver? Does it add to the public relations benefits of a maneuver, and add to the public relations costs of declining to maneuver?

As an experiment, TruSat may well supply answers to some of these questions about how, if at all, introducing an unprecedented public eyewitness capability to orbital operations affects operator incentives in weighing the costs and benefits of maneuvers. As readers are no doubt experienced, the addition of speed cameras or a visible police car have a demonstrable effect on the strictness of adherence to terrestrial rules of the road. Whether this phenomenon extends to “rules of the road” that will emerge for orbital operations remains to be seen.

III. Technical Overview of the Proof of Satellite Software Engine

The technical problem to which the Proof of Satellite software engine ("PofSat") is addressed is proving a satellite's orbital characteristics from individual observations that are not alone trustworthy. By leveraging the immutable characteristics of "on-chain" transactions, the PofSat engine is designed to be a substantially autonomous, decentralized capability, which can produce useful results without central moderation.

A satellite's orbit can only be perturbed by a limited number of phenomena, generally listed in order of effect, from greatest to least[25]:

1. The launch vehicle, or final transfer or insertion stage of the rocket (often called rocket boosters, of which many remain in orbit)
2. The satellite's own propulsion system, or in the future, effects of third-party orbital transfer or satellite servicing
3. Collision with an object or debris particle in orbit
4. Drag effects from the Earth's atmosphere (most prevalent at lower orbits, typically below 1,000km in altitude)
5. Gravitation effects related to the non-spherical nature of the Earth's gravitational field, the Moon, other planets, etc.
6. Solar wind and other space weather effects

The Proof of Satellite engine leverages two features of Earth orbiting objects:

1) The satellite orbits behave according to physics (Newton's Laws), and thusly must be "in-family" with prior observations.
2) Satellites positions and orbits are (theoretically) observable by anyone who knows where or how to look

Because of this, public trust in a prediction of where the satellite will be, according to the physics of orbital mechanics, can be successively verified by observations that the satellite is behaving (orbiting) as predicted. New observations have the potential to improve confidence in the orbit prediction as well as confirm prior estimates of it for purposes of assessing the performance of algorithms, observers and specific techniques contributing to these estimates.

[25] Over time we may expect the space economy to provide an additional perturbation in this list - in the form of "satellite servicing" vehicles, or "another satellite" which would probably exist between #1 and #2 in importance. The Space Shuttle, and the Russian Soyuz and Progress modules have contributed to the reboosting of the MIR Space Station, the International Space Station, and other satellites, such as the Hubble Space Telescope, so far in the history of space exploration.

III.A. Confidence Factors

Beyond leveraging orbital mechanics and successive observations of a satellite from multiple points on Earth, the initial release will utilize the following confidence factors to automatically sort orbit predictions in the catalog into six confidence levels. Unlike the publicly-available SSA data from the U.S. Government—which discloses neither sensor data nor algorithms, citing national security reasons—the factors and evolving algorithms for weighting them in the TruSat catalog will always be transparent. By providing the confidence data and supplying the code which determines it via open-source, TruSat data users receive a confidence assessment, and also the means of understanding how the system reached that assessment. The factors incorporated in these types of algorithmic ranks are anticipated to be subjects of future community governance activities, when the platform is able to accommodate on-chain governance to elect specific versions of new algorithms.

Confidence in the orbit of a particular satellite or space object is influenced by the following factors:

A) Time - Newer observations generally increase confidence. As the data ages, the confidence naturally decays from the influence of perturbing effects on the satellite's orbit.
B) Diversity of Observers: Confidence increases with multiple observations of a satellite, and increases to its highest level when there is geographic and national diversity in the observers.
C) Performance history of the Observer: Observers who have consistently provided good results can advance the confidence of a satellite orbit more quickly than unknown, or inconsistent observers. (see User Rank, below)
D) Traceability, Transparency of the raw data: Raw and derived observation data is cryptographically hashed, and demonstrated to (continue to) exist, and traceable to the orbit determination result. Removal of, or compromise to this data will have a negative effect on the confidence of the satellite orbit.

Based upon these factors, PofSat will automatically sort satellites in the catalog into the following confidence levels:

(5) Highest Confidence
(4) High Confidence
(3) Verified
(2) Plausible
(1) Neutral
(0) Untrusted

Users of the TruSat catalog may sort satellites by confidence level for a variety of purposes. Among them, satellite observers may wish to prioritize their observations to increase confidence in relatively low-confidence predictions by submitting additional observations on satellites rated (3) or below.

III.A.1. Time - Recency of Observation

As introduced above, the confidence in a satellite's orbit decays from the influence of perturbing effects on the satellite's orbit, as well as limitations in the chosen models used to propagate and communicate future predictions of the satellite's orbit. For the initial version of this algorithm, we are implementing the SGP4/SDP4/SGP8[26] algorithms which are optimized for results in the vicinity of the time epoch for which they are produced.

Excluding spacecraft maneuvers from perturbations, orbit knowledge of objects in Low Earth Orbit (LEO) will generally degrade more quickly than objects farther away from the atmospheric and non-spherical gravitational perturbations induced by the Earth itself. As a result, a "High Earth Orbit" object could have orbit knowledge that is more accurate than a low Earth orbit object, with the same time having passed since each of their most recent observations.

For the initial version, the weight factor on "time" will be determined empirically.

III.A.2. Observer Diversity

The PofSat engine depends on access to the physical coordinates from which observations are made, provided directly by the user. The combination of the "angles only" measurements of a satellite, and the location from where they were made are necessary inputs in using these measurements to determine a revised orbit estimate. This cannot be achieved without access to both pieces of data.

As is often the case with measurements of physical properties, a "longer baseline" over which to make these measurements increases the precision of the measurement. Two measurements on opposite sides of the Earth can provide more insight into the orbit than two closely-spaced measurements from Amsterdam are able to. Alternatively, observations of an object by a single observed spaced by a complete orbit provide a higher-fidelity observation, but do not provide diversity in data source which builds additional trust.

In addition to a geographic diversity of observations ability to improve the accuracy of an orbit prediction, when taking into account potential Attack Vectors (described later in this document) on the PofSat engine, we can also benefit from any knowledge of "National Diversity" or more generically, allegiance to a particular faction.

Theoretically, if all the observers in Amsterdam were to conspire to make predictions which benefited themselves, and increased their rank or perpetuate a version of the catalog for their own objectives, we would be motivated to incorporate and provide weight to observations which were linked to allegiance to Amsterdam, or were distinct from it. While to first order, it may seem sufficient to apply this weighting

[26] https://www.celestrak.com/publications/AIAA/2006-6753/

based on these users' geographic coordinates, their ability to acquire observations in remote locales (say, through telescope-for-rent services such as iTelescope, Slooh, or Lightbuckets[27]) could easily evade this weighting. To be able to implement this type of "faction diversity" may require some amount of Know-Your-Customer efforts (KYC) in order to gain third party confidence of the identity and affiliations of an observer, and as a result is scoped for implementation later in the roadmap.

III.A.3. User Rank

Certain efficiencies can be realized by taking into account the performance history of an observer, and an assumed desire on that user's part to obtain a positive ranking within the system, and for the system to provide a beneficial result for the community. The ultimate outcome in evaluating a user's relative ranking is how their observations relate to verifiable accurate results as observed, and subsequently contributed by other users.

If User A contributes an observation resulting in subsequent observations from other users being more accurate, then User A's rank increments positively.

If User A contributes an observation resulting in divergence from subsequent observations, then User A's rank decrements.

Taking into account every contributing user into the system allows for varying weights (both positive and negative) to be applied to observations from known users as appropriate for the available data. In the use-case thus far, observations are made in short batches, and observations by multiple users do not typically overlap within these batches. This provides an opportunity to update the estimate on a user-by-user, observation-batch-by-observation-batch cadence to simplify an initial approach. Due to the immutability of the dataset and its history provided by blockchain, users will not be able to hack the ranking records, but can only influence their own rank by submitting observations which are positively verified by other users.

For the v0.1 release, the specific factors contributing to initial rank are being established, but they are expected to include aspects of:

1) Total number of validated observations contributed
2) Longevity of contributions
3) An empirically-derived weighting factor for #1 and #2 to add weight for users who have contributed recently (this prevents legacy users from having undue influence if they aren't continuing contributors)
4) A factor relating to inclusion/exclusion of a user's observations increasing/decreasing the residuals of other users' observations

The end result of applying user-rank into observation weighting is two-fold:

1) The effect a user's contribution has on the estimate of an object's orbit, and

27 https://www.universetoday.com/93764/roboscopes-real-armchair-astronomy/

2) The speed at which object confidence is elevated per each new observation (low rank users have little-to-no effect, while high-rank observers may allow graduation to the next confidence level immediately)

This rank is anticipated to be updated for every batch of new observations submitted by the user, with the opportunity to update the associated TLE estimates at the same moment.

Details of the user rank algorithm are being finalized for the 0.1 release of the PofSat engine, and will be fully detailed in a 0.1 release of this whitepaper, along with the source code on GIT.

As described in the roadmap at Section V, successive releases will introduce extrinsic incentives for observing satellites.

III.A.4. Observation Trace-ability and Audit

In the post v0.1 roadmap for TruSat will be the ability to accommodate the raw data (position observations of satellites) fusing a facility like IPFS[28] to store it in a distributed and immutable fashion. Semi-automated open source image object identification and observation reporting capabilities already exist, in projects such as sattools[29] and stvid[30], and incorporation of these code-bases and the storage of the raw data in IPFS could allow for third-party verification and machine-learning assistance of verification of observations, as well as contributing to better results as a result of an increasingly large training dataset. Separately, it provides an opportunity for non-observing laypersons to contribute to the effort merely by volunteering to serve as an IPFS node for the TruSat data, similar in style to how “SETI-at-Home” and now the Berkeley Open Infrastructure for Network Computing (BOINC) application allows individuals to donate their spare computing sources towards scientific computation pursuits.

Workflows which include a raw astrometry dataset also provide for the opportunity to extract additional information such as object visual magnitude and variability, which may relate to certain sustainability factors.

Further discussion of the utility of utility of Traceability and auditable data will be discussed in the section relating to Attack Vectors below.

III.A.5. Object Priority

TruSat’s mechanisms for prioritizing specific objects for priority observation are central to its utility as a sustainability tool. To-date, the amateur satellite observation community has loosely coordinated

[28] https://ipfs.io/
[29] https://github.com/cbassa/sattools
[30] https://github.com/cbassa/stvid

observations by email threads, often in connection with new launches of interesting payloads, or with objects that have been a subject of conversation in the community. While TruSat welcomes satellite observations made according to the personal interests and priorities of its observers, it introduces new functionality for communicating the System's highest priorities for observation to its user community, and incentivizing observation of those objects.

Two mechanisms for prioritizing observations are presently contemplated: (1) manual inputs from the space sustainability community (initially via the TruSat Partners) to align observation priorities with that communities data needs; and (2) automated prioritization to increase confidence in orbital predictions.

III.A.5.a. TruSat Missions: Manual Inputs from Space Sustainability Community

TruSat's Mission functionality, which will be introduced for testing in the v0.2 release later in Q4 2019, enables space sustainability advocates to task the TruSat system to obtain data needed for sustainability applications. For example:

- To verify a satellite operator's claim to have placed a satellite in a disposal orbit, in conformity with a sustainability standard.
- To record the tracks of two satellites that are the subject of a conjunction alert.
- To assist in the identification of satellites following a launch of multiple small satellites.

Whereas in the v0.2 release such taskings will be facilitated by the TruSat Partners, pursuant to the procedures prescribed in the TruSat Charter, the TruSat roadmap contemplates a progressive decentralization of the Mission tasking function, in step with the migration to Ethereum mainnet, including the possibility of voting by the broader community of satellite observers.

Each manual tasking creates a Mission: a satellite prioritized for sustainability purposes, with a brief explanation of the significance of the Mission. From this context, individual satellite observers may choose whether or not to participate in the Mission by observing the specified satellite. Whereas the v0.2 release will depend on the intrinsic motivation of observers to contribute to space sustainability ends, subsequent releases may introduce extrinsic motivators, including digital "mission patches" unique to each mission, and other inducements appealing to observer's sporting sensibilities including possible use token-staked bounties which could carry monetary/marketplace value.

III.A.5.b. Automated Prioritization

Separate and apart from "TruSat Missions" functionality, TruSat automatically identifies priorities for observation to increase the confidence of orbital predictions. Below are several factors that TruSat will utilize in developing a priority of an object

- Algorithmically, the Proof of Satellite engine depends on continued observations to address the natural decay in confidence by the passage of time. If the time-since-last-update has passed a threshold for a particular orbit regime or object class, the engine will increment the priority of the object.

- If more geographic diversity is needed to elevate the Object Confidence to the next level, the Object Priority will be incremented.
- If an object has few or no observations by high-ranked users, that object will be prioritized for those users.

TruSat will be able to provide a list of priority-interest objects on a global perspective, as well as provide a priority list which is customized for a specific user's location, rank and capability (such as inferred or specified limiting magnitude sensitivity).

Details of the Object Priority algorithm are being finalized for the 0.1 release of the PofSat engine, and will be fully detailed in the source code in the TruSat software repository and in accompanying documentation.

III.B. Confidence Factors Applied

The scenarios that follow serve to illustrate how the confidence factors are applied by the Proof of Satellite engine to mediate its best estimate of objective orbital truth.

In a traditional curated catalog, expert administrators can review the data as it relates to observer performance for data consistency, technical biases, and other concerns which may influence the final result. That administrator can exercise judgment in "if and how" to incorporate data into their prediction results. TruSat removes any such administrator from the loop, substituting algorithmic applications of confidence factors as follows.

The scenarios in this section illustrate how the TruSat confidence factors are applied to curate data from good faith, competent visual observations in their intended use. The scenarios at Section 1.D.3. illustrate how the same confidence factors may serve to mitigate potential attack vectors, from erroneous observations to malicious attempts to undermine the accuracy of orbital data in the TruSat catalog.

Use Case Summary:

- New observation of a (3) Verified object
- Operator data for a (4) High Confidence object
- New (divergent/poor quality) observation for an object
- Uncorrelated object observation from a well-ranked user
- Uncorrelated object observation from a low-ranked user
- Previously unknown object
- New observation of a low-confidence object
- New low-ranked observation of a "verified" object

III.B.1. Verified Object (3), New Observation

Situation: A demonstrated consistent TruSat user looks for an object ranked “(3) Verified” and is able to find it based on the orbit prediction. She submits one or more observations with the object’s observed location and time. Her observation correlates well with prior observations. She is also located a large distance from other reporting observers, and the object now has observations from three countries over 24 hours.

Resolution: The new observation is used to refine the orbit prediction, object is upgraded to “(4) High Confidence”, and put (lower) in the priority queue for other users to routinely update and maintain confidence.

III.B.2. High Confidence Object (4), New Data from Operator

Situation: A wallet (blockchain-secured account credentials) associated with the operator of a satellite provides its own tracking and telemetry data based on on-board GPS measurements for a “(3) Verified” or higher object. Their data correlates well with other users’ observations.

Resolution: The new observation data is used to refine the orbit prediction with increased weighting, object is upgraded to “(5) Highest Confidence”, and put (lowest) in the priority queue.

III.B.3. Plausible or Higher Object (2+), New (poor quality) Observation

Situation: A new TruSat user submits an observation which appears to correlate to a catalog satellite. However, (as yet) unknown to the system, the user’s clock was wrong, their GPS position was off, or there was otherwise an error in their data which caused the orbit to be made less accurate.

Resolution: The new observation data is used to refine the orbit prediction, object is upgraded to “(3) Verified”, and put in the priority queue for other users to continue to verify and increase confidence. However, subsequent observations reveal this data as an “outlier”, and its weight in future orbit determinations is reduced, and the user confidence is adjusted lower as a result of the poor quality observation.

There is an opportunity in statistical analyses of observations to provide user feedback based on certain observables - in this case, a persistent bias in their observations and potential remedies to correct it. TruSat also provides a new capability for a user to see their observations along-side other users’ observations for direct feedback on their performance relative to other users observing the same object.

III.B.4. Known Object not in TruSat Catalog, New Observer:

Situation: A new TruSat user contributes data which cannot be automatically correlated with an object currently in the database via user-specified NORAD number and/or International Designation. The designation however is on-record with the Celestrak SATCAT which TruSat periodically incorporates.

Resolution: The object is included in the catalog, confidence level is set to “(1) Neutral” and object is flagged for follow-up by other observers.

III.B.5. Known Object not in TruSat Catalog, High Performing Observer

Situation: A well-ranked TruSat user contributes data which cannot be automatically correlated with an object currently in the database via user-specified NORAD number and/or International Designation. The designation however is on-record with the Celestrak SATCAT which TruSat periodically incorporates.

Resolution: The object is included in the catalog, confidence level set to “(2) Plausible” and object is flagged for follow-up by other observers.

III.B.6. Previously unknown object:

Situation: A TruSat user contributes data which cannot be automatically correlated with an object currently in the database based on orbital characteristics derived from 2 or more observation data points, using an “analyst” designation to represent its unidentified (UNID) status.

Resolution: The object is included in the catalog, assigned a unique (hexidecimal) catalog number, recorded at a “(1) Neutral” confidence level and is flagged for followup by other observers.

III.B.7. Unverified (1) or Plausible Object (2), New Observation

Situation: A well-ranked TruSat user looks for an object ranked “(1) Neutral” or “(2) Plausible”, is able to find it based on the orbit prediction and submits an observation.

Resolution: The new observation is used to refine the orbit prediction, and is upgraded to “(3) Verified”, and put in the priority queue for other users to continue to verify and increase confidence.

> Caveat - An observer can’t verify their own observation. If they discovered it, and later (a day, a week) observe it again, the orbit is refined, but the ranking remains “(2) Plausible” until a third party observes it.

III.B.8 New low-ranked observation of a Verified (3) object

Situation: A low-ranked TruSat user submits an object of a verified object. This object may have previously held a higher confidence ranking, but was demoted by the system after passage of too much time.

Resolution: The new observation is used to refine the orbit prediction, but the Object Confidence ranking is not upgraded. The Object is flagged for follow-up by well-ranked users. Note that if this TruSat user once held a higher rank, but was demoted for confirmed inaccurate observations, or the passage of time, the use of well-ranked observers to confirm their observations will assist in increasing their rank.

In the TruSat roadmap, several of these use case may be "upgraded" with the use of source image data from the observation, allowing for third-party or automated confirmation of user observation analyses, and faster escalation of object confidence, or more specific diagnosis of opportunities for improved accuracy in the observations.

Details of the Object Confidence algorithm are being finalized for the 0.1 release of the PofSat engine, and will be fully detailed in the source code in the TruSat software repository and in accompanying documentation.

III.C. Mitigating Attack Vectors

As a decentralized, automated system, the PofSat engine must protect itself against attack vectors ranging from erroneous observations from naïve or inexperienced users to malicious attempts to exploit the open and permission-less nature of the catalog in order to disrupt service, or to skew or corrupt the data.

The following are potential attack vectors, and mechanisms available in the system architecture to mitigate or neutralize these attacks on the TruSat catalog.

<u>List of Attack Vectors currently contemplated</u>

- Single incorrect positive confirmation of a verified object
- Single falsified observation of an object
- Coordinated false reporting on an Object (Positive result)
- Coordinated false reporting on an Object (Negative result)
- Code injection attacks on server
- User Phishing or account take-over attacks
- DDOS attacks

III.C.1. Single Incorrect (or Falsified) Observation of a Verified Object

Situation: A user "falsely reports" an observation. They could do so by making up an observation from scratch, or by submitting a modified version of an observation from another user or attempting to submit another user's observation as their own.

Resolution: The TruSat engine fingerprints all observations on parameters that can uniquely identify observation parameters which directly affect the orbital estimate and would simply "reject" a duplicate observation. This is also protection against accidental re-submission by a user, or an attempt to increase the weight of an observation by submitting it multiple times. The TruSat engine performs a similarly designed fingerprint check against externally-sourced TLEs, in order to prevent duplicate entries, and this feature could also be applied to audit submission of unauthorized 3rd party TLEs for which TruSat has been able to fingerprint the source files.

For observations that are not "identical" to observations already in the system, several checks are applied. For example: After the object is verified to exist in the catalog by NORAD number checks, it is then evaluated for a "nearness" check to the most-recently predicted orbit position for that observation. This check evaluates the fit on location (time across the trajectory), cross-track error, and if multiple observations are available, angular rate. An object in the right location, moving at a plausible rate would fail this test if its direction of travel diverged beyond a threshold. Similarly, if an object is headed in the correct direction, but moving outside a threshold of angular velocity, or deviating too far from its along-track location, the observation would not be correlated with the catalog object on file.

Options from this point would be to simply discard the object, flag the user in an appropriate way, or potentially, put the object into a flow for evaluation as an independent, new object. A demotion of the user rank for observation errors (wrong object specified, observation otherwise flawed) would disincentivize submitting flawed data.

In the TruSat roadmap, incorporation of user-submitted raw-data would help provide more option paths for an object which does not plausibly-map to an existing catalog object; although, then defense against "deep fake observations" would need to be developed.

III.C.2. Single Falsified Observation of a New Object

Situation: A TruSat user "falsely reports" an observation. They could do so by making up an observation from scratch or by submitting an existing observation but change the object number they are reporting against.

This could be from malicious intent (intent to poison data) or a user who is reporting an observation with no correlation to reality, or arbitrary correlation to the expected result.

Resolution: As previously described, a user who attempts to submit an object which may correlate to a different object would have the system rejecting the observation due to its lack of plausible correlation. For a user reporting a “new object” several mitigating actions are available:

A) The object is entered into the catalog with a “(1) Neutral” confidence and flagged per-usual for follow-up.
B) Submission of “new” objects could be restricted to autonomously high-ranked users, or further restricted to users which have passed TruSat community KYC verifications. In this case, the object could be entered into the catalog with an “(0) Untrusted” Object Confidence, but still available for followup - with the designation implying a note of warning to those performing followup.

In either of these cases, well-meaning observers would be unable to verify this object, and their attempts could be logged in the system as “not seen” data points, which could be applied to demote the object’s confidence.

Ill-meaning users could provide follow-up “fake” confirmations, but lacking an independent well-ranked user’s positive confirmation, the object would never escalate above the “(3) Verified” object confidence, and a certain threshold of “not seen” observation could prevent it from reaching this status.

If an Ill-meaning user were to make a number of valid observations of catalog objects for the purpose of gaining enough rank in the system to submit observations for “fake” objects, then they have provided value to the system in the process of doing so. In this case, the positive increments on rank could be offset by significantly larger consequences for rank detriment which would prevent this from happening at a large scale. Limitations on the number of objects a user could verify per time-increment could also be applied to rate-limit this adverse effects. Similar considerations for “51% attacks” in blockchain networks contemplate the value-add that is provided by needing to substantially and positively contribute to the system infrastructure before one gains the capability to “attack it” and then have a rational interest against a self-attack.

Here too, using TruSat’s anticipated roadmap capability of submitting photographic evidence of the observation could increase the challenge of falsifying observations, while increasing the need to protect against falsified image records.

As mentioned previously, geographic diversity and KYC checks on origin or faction could provide unrelated limits to the reach of these nefarious activities.

III.C.3. Coordinated False Reporting on an Object (Positive)

Situation: Extending the case in I.D.3.b to a coordinated effort by multiple users to create a record for an object which does not exist, multiple TruSat users submit self-consistent observations for a fictitious object, or alternatively, self-consistent observations for an existing object in an attempt to skew TruSat’s estimate of its orbit.

Resolution: Here again, as described in the previous case, the low rank of the ill-intended TruSat users would prevent a new object from achieving an Object Confidence greater than “(2) Unverified”. If one or more of them made TruSat contributions in order to achieve rank, their benefit to the system may ultimately outweigh their detrimental (temporary) effects on the system.

A more insidious case is with an attempt to skew the orbit knowledge of an existing object (a known military satellite operated by a dictator regime) by exploiting the open source nature of the project, and estimating the type of observation necessary to (a) pass the correlation and plausibility tests, but (b) result in the observation being worse than would be possible with genuine/accurate data.

Here, the user rank of these nefarious users would have low weight along-side the higher weight observations from well-established users, and the nefarious observations would get successively demoted in user-rank to a weight of zero. This “arms race” could be kept up with perpetual new accounts being established, and perpetual skewed observations being provided, but these patterns may start to present themselves in detectable ways.

Here again, supplying source observation data in image format would increase the difficulty in achieving a “fake.”

III.C.4. Coordinated False Reporting on an Object (Negative)

Situation: This case concerns itself with a ill-intended TruSat user trying to achieve an object being “demoted” from the catalog by submitting poor or negative observation records.

Resolution: In addition to the normal feedback loop from sourcing user rank in the promotion or demotion of object confidence, in this case, any positive confirmation by a well-ranking positive-intent user would quickly re-introduce the demoted object into the catalog. Communications and coordination by the community and these positive representatives in the community could highlight the object as a priority observation target, and the system design could continue to keep the object in the catalog.

III.C.5 Code injection attacks on server

Situation: Leveraging the open-source repositories of the TruSat project, the user infiltrates the operating code and incorporates their malicious version in the operating PofSat engine.

Resolution: The approved/validated code-base would be file-checksummed and logged on the Ethereum blockchain, and would immediately fail a consistency check. The system would automatically roll back to the validated release, or shutdown and notify TruSat Partners to intervene.

As we implement the TruSat roadmap, vulnerability to this type of attack will decrease, as it will become more decentralized, and the attacks would need to take place on a large fraction of the nodes in the system.

III.C.6. User Phishing or account take-over attacks

Situation: Someone with ill-intent attempts to acquire access to a well-ranked user in order to achieve actions only possible with higher-ranking users.

Resolution: The implementation of public/private key cryptography, and the security provided by an Ethereum-wallet for TruSat access make this type of attack increasingly more challenging to achieve. Even if an ill-intended individual were able to obtain access to one or more wallets through a vulnerability, their ability to implement sustained detrimental actions would be mitigated by the safeguards already discussed.

To obtain the private keys of the user, the attacker would need to perform an attack that can access user's local storage, gain access to their email and password, or perform a standard phishing attack. To guard against unauthorized access to local storage, users will be encouraged to switch their account management to Metamask. Finally, a phishing attack can be performed by tricking a user that they are in a safe environment when providing their credentials, this can be addressed by using a unique identifier that the user can check against in order to verify they are safe. Once switched to an established, trusted, environment such as a phone application or decentralized platform, this should no longer be a concern. All the key management will be performed by the user and their device without exposing anything externally.

III.C.7. DDOS attacks

Situation: Ill-intended parties attempt to deny availability of the TruSat service by overwhelming its servers.

Resolution: In the initial release, which relies on Web 2 services, this exists as a vulnerability. Upon transitioning more fully into an Ethereum mainnet solution later in our roadmap, TruSat looks to decentralize the catalog itself, as well as the server processes which interact and add data to the system.

III.D. Data Inputs and Outputs

With its v0.1 release, TruSat will securely store visual position observations, which are fed into the Proof of Satellite engine to generate satellite ephemeris, to be distributed in the two-line element ("TLE") format. To support the workflow, the catalog will also contain basic contextual information retrieved from public sources[31] and well as information required to uniquely identify user accounts, their observation locations,

31 Celestrak.com Satellite Catalog (SATCAT) https://www.celestrak.com/satcat/satcat-format.php and the Union of Concerned Scientists (UCS) Satellite Database https://www.ucsusa.org/nuclear-weapons/space-weapons/satellite-database.

and user performance data to support the automated generation of satellite ephemeris from updated observations.

In its initial release, TruSat will accept observations in IOD, UK or RDE Positional Observation reporting formats[32], although the “COSPAR site” identifying the user observation location in these records will still need to be coordinated through the Seesat-L mailing list until alternative standards can be vetted and established.

Early in our post-initial-release roadmap we anticipate being able to accept TLEs from satellite operators and operator derived state vectors with associated documentation on coordinate frame and time reference. Should the experiment of TruSat be successful, we anticipate being able to expand observations to include radio-doppler measurements such as those made by the SatNOGS[33] network.

Future releases may place greater weight on observations that are accompanied with their raw observation data, to allow third-party verification of the observation result. Other developments in the roadmap are anticipated to enable basic information about the object’s optical properties, including visual magnitude, and related measurements. Additionally, an active community of radio-frequency observers of satellite may provide for additional opportunities to verify satellite orbits, provide data about their RF emissions, or provide insight into objects which cannot be visually measured.

TruSat will not pull data direction from Space-Track.org or any other source that imposes restrictions on use of the data.

Predictions in the catalog will be represented in Two Line Element (TLE)[34] sets, the defacto standard by which the parameters describing a satellite orbit can be shared. The format, originally developed during the punch-card era of computing, is a fixed-field, fixed width partially encoded format for storage efficiency, and has been produced by NORAD for most objects since its establishment. A TLE can be used by the SGP family of orbit propagators to provide reasonably accurate position and velocity of a satellite for a given time. As a TLE is an approximation of the orbit, they are designed to be most accurate near an “epoch” (akin to a manufacture date), and become progressively less reliable the farther away from this date the estimate is used, owing to atmospheric effects, gravitational and other small perturbation, and the approximations of the SGP model itself.

Because of their ubiquitous use, TruSat project will produce predictions in compliance with this standard, and be prepared to follow, and potentially lead in efforts to modernize the format and its use.

32 http://www.satobs.org/position/posn_formats.html
33 Open Source global network of satellite ground-stations https://satnogs.org/
34 https://en.wikipedia.org/wiki/Two-line_element_set

III.E. Tuning the Proof of Satellite Engine

For visual position observations, TruSat is building on the 25 years of publicly available work from the all-volunteer SeeSat Visual Satellite Observer organization[35]. The SeeSat website contains extensive information about the hobby of visual satellite observation, including techniques for finding faint objects, making positional observations, as well as characterizing the visual brightness and time-varying visual properties. The SeeSat community is notable for distributing TLEs for satellites which are not available in the public US STRATCOM satellite catalog, typically classified satellites, operated by the United States and its allies.

The archive of SeeSat data supplied a rich source of data to "train" and tune the Proof of Satellite engine ahead of its initial release. For 25 years, SeeSat users have distributed their visual position observations in one of three machine-readable positional observation reporting formats[36]. The ConsenSys Space team has recovered and organized nearly 400,000 individual observations contributed by this community since December, 1998. Each of the observation data points include information about the apparent sky-location of a known satellite, and the time the observation was made. Additionally, reference information is available about the observer, their location, and the ability to describe observing conditions and the estimated accuracy of the measurement.

This data corpus was an invaluable resource in testing algorithms to support the Proof of Satellite engine, as well as explore UI/UX features to support and grow the community making these observations. Starting TruSat with a "25 year history" of user performance is also a benefit which we hope can continue to serve the Seesat community.

In our initial release, TruSat and its technical advisors will be making choices about the specific algorithm implementations, but we hope that by the 3rd major release of the system, we will be able to transition to a decentralized governance system of the algorithm modules.

ConsenSys Space will continue to tune the Proof of Satellite engine following the initial release, and will make all algorithms publicly available.

III.F. Open Source Algorithm Updates

TruSat is being architected and implemented in an open and modular way. We anticipate that there are likely "better" ways to achieve each portion of the TruSat engine than we are initially implementing, and

[35] SeeSat Website: http://www.satobs.org/
[36] http://www.satobs.org/position/posn_formats.html

there will likely be innovations realized by innovators and contributors throughout the world over time. In addition to supplying data to TruSat, we anticipate the possibility of network users/members to supply "pull requests" for engine algorithm updates. This could be as small as a suggested tuning of a variable, to as complex as a complete rewrite of a particular part of the process.

We hope that the system can allow for the open dataset and open source to be used to demonstrate to the community how their proposed change can improve the system, by demonstrating the potential result on actual data from the system. Furthermore, we anticipate that these types of "upgrades" would be high-value contributions worth of "bounties" that are conceived of the system, allowing individuals or organizations to derive benefits and incentives from the system in ways other than by providing observations.

We anticipate the first opportunity to implement options such as these would present themselves in the second stage of our roadmap, when we are transitioning on Ethereum MainNet, and can take advantage of the security of the "store of value" functionality provided by Ethereum.

III.G. Identity and Privacy

Privacy by default is a core design choice made in architecting the TruSat software. In the default settings, an Ethereum address and some observation history data will be the only information about a contributor viewable by other users of the System. Contributors will have the option to reveal more information, including their username and location (e.g., "India" or "San Francisco, USA") , as well as the option to remove their identity and observation contributions entirely from the network.

To account for a satellite observation, the Proof of Satellite engine must first verify the contributors ownership of an ethereum address, then process the a high-resolution (meters level) position of location from which the observation was made. TruSat utilizes the same public key cryptography used to verify ownership and control of cryptocurrency to verify ownership and control of an observer identity. If an individual wishes to use a different Ethereum address for every contribution or instead to "persist" the same address over time, they will have the option to do so. Using different addresses is likely to result in their observations having less influence on the catalog contents than if they were to choose a persistent identity. These design choices are in step with the long-term goal for the network to be fully decentralized. Ethereum wallets will be a key component of those future plans and early adopters of TruSat will be able to take their first steps toward that end by starting with a more self-sovereign identity.

IV. TruSat Governance

Governance of TruSat in its initial, experimental phase divides into two broad dimensions: System governance, and governance of the open source Community that maintains and advances the TruSat System. Both dimensions will be governed by the TruSat Charter for during TruSat's experimental phase.

The TruSat Charter is an agreement between the TruSat Partners—organizations committed to maintaining and advancing TruSat as a space sustainability tool—and delineates the Partners commitments and roles in this regard. The Charter, which is available at www.TruSat.org/Charter, also governs how the broader community of open source contributors engages with the project. Each iteration of the Charter is time-limited to an "epoch," after which it must be replaced according to the processes it specifies. Beyond the experimental phase, the Partners are committed to progressively decentralizing governance of TruSat by implementing blockchain-enabled governance tools.

At the System level, the primary governance matter is how the System is tasked (i.e., the selection of "Missions" prioritizing satellites for observation) to align System priorities with the needs of the space sustainability community.[37] During the experimental phase, the TruSat Partners will be responsible for facilitating such taskings pursuant to a process specified in the Charter. The public availability of the Charter provides full transparency of the inner workings of the System during this experimental phase. As the System develops and matures, it is foreseeable that the community of satellite observers have a more direct role in setting System-wide priorities for observation.

The TruSat Charter will also specify governance arrangements for the TruSat Open Source Community, in a similar fashion to open source projects such as the OREKIT.[38] Among the matters that will be addressed are the open source licenses applicable to code and documentation contributions, processes for community input on feature roadmaps and priorities, and processes for reviewing and accepting code contributions.

V. TruSat Roadmap

The initial v0.1 software release is a limited prototype for testing and refining the Proof of Satellite software engine. The very-near term roadmap will complete implementation of the initial concept feature-set, and address known-limitations (described below) including refinements to the automation, confidence, priority and user-performance algorithms. The v0.2 release later in Q4 2019 will add Mission tasking and other functionality. Subsequent releases will introduce features designed to enable the system to scale, by easing the process of making and reporting satellite observations, and progressing to a fully decentralized system. The roadmap for these subsequent releases may evolve with user-community feedback introducing new ideas, and focus development priorities.

V.A. Known limitations of initial release

As a prototype, the v0.1 release is far from feature-complete, and has several traditional solutions to common problems, allowing the team to focus on developing and demonstration core innovations in the

37 *See* Section III.A.5.a, supra for a description of Missions.
38 OREKIT Governance Framework, https://www.orekit.org/doc/orekit_governance_v4.pdf.

TruSat system. The initial release relies on a central web server on an AWS instance, as well as a central relational database for storing TruSat data. Whereas subsequent releases will further implement TruSat's decentralized, autonomous architecture, the ConsenSys Space will retain control over the network during alpha testing. The v0.1 release will utilize email-opt-in techniques for verification of identity. The v0.1 release likewise utilizes some relatively dated and simplistic orbital dynamics code, on the theory that utilizing open-source code bases could accelerate the deployment of a basic prototype and accelerate the point at which the community can help to evolve the code base with state-of-the-art algorithms.

V.B. Future Features to Scale Global Contributor Network

The following features and functions are tentatively planned to scale the global network of citizen satellite observers by making it easier and more enjoyable to make and report satellite observations. This roadmap will evolve based on user-community input.

- Features that ease the process of making and submitting satellite observations, such as:
    - A smartphone app that assists an observer in orienting her binoculars at the right point in the sky, and enables her to time the satellite and directly submit an IOD to the system with the press of a single button.
    - A smartphone app leveraging the improved low-light camera capabilities of contemporary phones (e.g,. Pixel 4 and iPhone 11) to place the entire satellite tracking workflow on a single screen.
    - A web app interface enabling satellite observers to directly input the basic observation parameters (time and location), dispensing with the need to utilize third party software to calculate and format an IOD for submission.
    - Automated generation of observation "station."
    - Translation of interfaces and documentation into multiple languages.
    - Integrating radio observations (e.g, "satnogs").
    - Automated image processing and IOD extraction from digital images[39]

- [39] Much of what is needed to integrate automated IOD extraction from digital camera images appears to be available in Cees Bassa's sattools/stvid open-source code):
    - Receive image frame from registered user (known location)
    - (Repeat) Plate-solving
    - Perform feature detection of image / image series, looking for known/expected objects based on user capability profile
    - Optionally, perform further algorithmic exploration of unidentified objects (UNIDs)
    - Extract position / time information, or verify efficacy of user-provided IOD/UK/RDE-formatted position observation
    - For each IOD
    - Cross-Correlate existence of IOD object, with database state vectors (TLE, or more detailed internal format

- Direct tasking of motorized, software-enabled consumer telescopes or scriptable robotic observatories
- Beyond software: grants of satellite observing hardware in service of greater geographic and socioeconomic diversity of satellite observers.

- In addition to making satellite tracking easier, a number of features are contemplated to make the experience more fun and meaningful. Among the possibilities are:
  - Missions linking observations to concrete sustainability outcomes.
  - “Adopt-a-satellite”: assign a satellite to observers around the world, who are responsible for tracking it.
  - “Don’t-brake-the-chain” campaigns and badges.
  - Community-combined or “team” goals: goals that require multiple users to collaborate to achieve an outcome (e.g., track deployment of new microsat cluster launches, communications constellation deployments)

V.C. Web3 Functionality Roadmap

In its initial prototype releases, TruSat will implement three features of blockchain functionality.

A. User identity via wallet-based login, and the associated Web3 onboarding opportunity.
B. Registering of “authenticated” versions of code and algorithm base in the runtime code.
C. Periodic consistency checking of tamper evidence by comparing with blockchain checkpoints for code-base and data-base.

After implementing these basic features, and demonstrating the broader utility and value of the PofSAT engine, a transition to Ethereum mainnet will enable more advanced features:

1. On-chain transactions for (up to) every observation block, and associated TLE prediction update.
2. On-chain stores of value (reputation, incentives, NFTs, etc.).
3. State-channel object storage (TLEs, algorithms, meta-data, etc.).

V.C.1. Decentralizing the System and its Governance

- Migrating to Ethereum Mainnet for decision democracy (e.g., for feature priority).
- Decentralized system for “tasking” the network (i.e., prioritizing satellites for observation).
- Decentralized governance of Proof of Satellite Engine upgrades.
- Decentralized moderation of general TruSat codebase.
  - Community moderated “graduating” privileges of contributors.
- A system for distributing rewards/incentives to contributors (see next section).

V.C.2. Implementing On-Chain Incentives

In its initial releases, TruSat depends on participation (i.e., satellite observations) by individuals motivated by a concern for space sustainability and a desire to be a part of the solution; and individuals who enjoy the "sport" of spotting satellites.

- Relies on limited suited of Web2.0 mechanisms for enhancing the experience of sporting and sustainability-minded users (e.g., analytics and rankings, campaigns linked to sustainability objectives).
- Migrating to Ethereum Mainnet will enable TruSat to implement a fuller set of incentives for participation, including, potentially, distribution of value created by the network to its data contributors.
    - There has never been a tool quite like this, and ascertaining how it is used is part of the experiment. This will inform the universe of possibilities and design of incentive structures for future releases, including potentially a token model.
- Initial release is also dependent on the interests of the TruSat Partners in continuing to operate and develop this sustainability tool. Plan for release on Mainnet is to supply incentives for decentralized maintenance of this tool.

Initially, bug reports and feature requests will be tracked and managed at:
https://github.com/consensys-space/

VI. TruSat as a First Step Toward Earth's Space Program

VI.A. The Guiding Vision of Earth's Space Program

The exploration and development of outer space interests a great many people across our planet and yet is accessible to few. Even as more countries than ever before field space missions, opportunities to participate in space endeavors are concentrated in a small handful of institutions, in a smaller handful of countries, and limited to relatively narrow range of professional disciplines. As a consequence, human space endeavors benefit from only a fraction of the ideas, talents, and financial capital seeking to contribute.

Much of the dazzling progress in space technology since the turn of the century can be attributed to billionaires so impatient with the pace of government space programs that they invested their fortunes and talents in starting their own space programs. Impatience with the pace of progress, good ideas, and a desire to contribute are not limited to billionaires, however; they are but the tip of a much larger iceberg. And yet the capital intensity of space endeavors keeps most would-be contributors beneath the water; doing almost anything in space requires a level of capital only a billionaire or national government could marshal. What most individuals could contribute would, by itself, be such a drop in the bucket to discourage trying.

ConsenSys Space was founded on the belief that creating mechanisms to aggregate and coordinate the human and financial capital wishing to contribute to space endeavors would unlock new solutions to the most stubborn challenges holding back progress in the exploration and development of outer space. The co-founders are well acquainted with these challenges—from the long-term sustainability of spaceflight to the economic hurdles to starting a space venture—having led governmental, intergovernmental, and private space initiatives.

ConsenSys Space is guided by a vision for Earth's Space Program: a global, citizen-led space program accountable to and controlled by its community of contributors. It is a vision for bottom-up collective action on an unprecedented scale, in which individual contributions, modest in isolation, are assembled into space activities on a scale only a handful of governments and billionaires have managed to date.

The Ethereum blockchain presents a new technological foundation on which to build such Earth-scale collective action. Traditionally, individuals contributing time or money to a large endeavor do so on their trust in their fellow collaborators, or a central institution (a non-profit, company, or government). This significantly limits the scale of collaborative projects. At the largest possible scale—accommodating contributions from every country on Earth—it is impossible to know (and thus trust) a significant portion of one's fellow contributors, and near-impossible for such a large, geographically distributed group of stakeholders to agree on a central entity trusted by all. At a fundamental level, Ethereum technology dispenses with the need for a trusted entity, or even to know one's collaborators. In place of trusted individuals and institutions are transparent rules that are guaranteed to execute.

As explained below, TruSat is a first step—a set of experiments in Earth-scale collective action—on the path to Earth's Space Program.

VI.B. Diversifying, Democratizing, and Decentralizing Space Endeavors

Embodied in the vision of Earth's Space Program are imperatives to diversify, democratize, and decentralize space endeavors.

To diversify space endeavors is to harness more of humankind's potential. To bring more and different kinds of people off the sidelines and into the league of space contributors. Diversifying the pool of contributors increases the volume of resources available, as well as the breadth of ideas as perspectives and experiences, including those that have largely been absent to date.

To democratize space endeavors is to make them accessible to anyone with a desire to contribute. This requires the creation of opportunities to participate, such as contributing funds, expertise, in-kind resources, participating in decisions, creating content, or collecting data. It also requires the removal of barriers to participation, beginning with overcoming perceptions that space is only for engineers and scientists with elite credentials. Making participation in space endeavors accessible to a wide swath of Earth's population will require linguistic and technical translation, as well as user interfaces and technological aides that reduce knowledge barriers to making valuable contributions.

To decentralize space endeavors is to apply space mission engineering practices to the institutions that carry them out. No company, government, or other institution is a single point of failure. To engineer a space program for multi-generational progress, it must be decentralized for resilience; to insulate long-term strategic planning on an Earth scale from shifting political and economic winds in individual countries.

VI.C. TruSat as an Experiment in Diversifying, Democratizing, and Decentralizing Space Endeavors

TruSat is an experiment in bottom-up, Earth-scale collaboration to produce a result that heretofore required exquisite, expensive sensor networks. For individuals impatient with the halting pace of intergovernmental cooperation on space sustainability, it will provide opportunities to take direct, concrete action; to be a part of the solution. With its Proof of Satellite software engine and open, transparent architecture, TruSat can assemble individual satellite observations, which would not be of much use in isolation, into a trusted source of truth on orbital location; an unprecedented space sustainability tool.

Geographical diversity—observations of a satellite from multiple points around Earth—enhances the accuracy of orbital predictions. Relying on visual observations of satellites, TruSat advantages participants in rural or less-developed locale, with less light pollution than cities. This inverts the traditional structure of opportunities to participate in space endeavors, which tend to be clustered around large cities.

As an experiment in democratizing space endeavors, TruSat will reduce the knowledge and skill barriers to making and observing visual satellite observations. At present, making a visual observation of a satellite requires the observer to:

- Utilize a web resource such as a Heavens Above, CalSky, or a smartphone app such as SkyView, to determine where and when to look;
- Orient binoculars or DSLR at the appropriate star pattern, and record the time the satellite crosses an imaginary line drawn between two stars;
- Utilize imaging-based software such as sattools/stvid to calculate an Initial Orbit Determination ("IOD") based on the observer's location (GPS coordinates) and the timing of the image-detected satellite observation.

The initial release of TruSat, intended primarily to test and refine the Proof of Satellite software engine, will largely rely upon the above workflow, supplemented by tutorial resources. For some, the sporting element of "hunting" satellites in the night sky contributes to the appeal of the activity, scaling to a global, diverse pool of contributors will likely require a relatively simplified, seamless process for making satellite observations. As described above in the TruSat Product Roadmap, features that ease the process of making and observing satellites will be priorities for subsequent releases.

TruSat's product roadmap contemplates decentralization on multiple levels in service of trust and resilience. At the software level, the TruSat Partners aim to implement TruSat on Ethereum mainnet, with decentralized, on-chain governance. Once completed, priorities for observation, as one example, will be determined by the TruSat community of contributors. In the initial, experimental phases, TruSat will be implemented on the Rinkeby testnet, and governed by the TruSat Partners pursuant to the TruSat Charter.[40]

TruSat's decentralized architecture will help it to surmount the trust challenges of existing SSA networks. Trust in the SSA data generated by TruSat is not dependent on trust in any institution (whether ConsenSys Space or any of its partners). Every aspect of TruSat's software and underlying algorithms will be open and transparent, and its SSA data this subject to verification. Trust but verify.

At an institutional level, TruSat is intended to progressively decouple its longevity and progress from the support of any single institution. This is accomplished by spreading support for TruSat across multiple TruSat Partners, and developing TruSat as on open source project in which ConsenSys Space is one contributor among many.

VII. How to Join the TruSat Open Source Space Sustainability Community

Launched October 21, 2019, TruSat is in its earliest days and will depend on the creativity and efforts of a global community of contributors to progress it from prototype into a powerful system for preserving our spacefaring future. Whether you are interested in tracking satellites, helping to develop the TruSat software, or connecting your community with TruSat, the first step is to join the join the community at www.TruSat.org/Join.

Community involvement in developing the TruSat software will commence in January of 2020, following the version 0.2 release, and amendments to the TruSat Charter to add detailed open source governance arrangements. To catalyze community involvement in the software build, ConsenSys Space will sponsor a TruSat global hackathon in January 2020. By joining the TruSat community you will receive email updates about these opportunities.

If you are interested in becoming a TruSat Partner, please contact space@consensys.net.

[40] *See* TruSat Governance, Section IV *supra*.

VIII. ConsenSys' Business Interest in TruSat

ConsenSys Space, which developed the free, open source software powering TruSat, is wholly-owned by ConsenSys, Inc., a privately-held blockchain technology company. As the TruSat.org domain signifies, TruSat is designed to generate a global public good rather than profit. ConsenSys' business interests underpinning its investments in TruSat and other pioneering space applications of Ethereum blockchain technology merit a brief explanation.

ConsenSys has broad business interests in developing and growing the Ethereum ecosystem. Pioneering applications of Ethereum that illuminate entirely new solution spaces for old, stubborn problems benefit the ecosystem. ConsenSys believes that enabling large-scale, global collective action is among the most transformative potential applications of Ethereum blockchain technology, and is investing in solving space-related collective action challenges as R&D to accelerate the realization of this potential. Space is hard, and that's the point. The challenges of the original moonshot were undertaken "...not because they are easy, but because they are hard, because that goal will serve to organize and measure the best of our energies and skills…"[41] Realizing ConsenSys Space's vision for Earth's Space Program will require a global, open source moonshot, galvanizing developer communities, and drawing a diverse swath of Earth's population—and with them, new perspectives and ideas—into the Ethereum ecosystem. Building the infrastructure to enable all people to contribute and participate directly in space endeavors will necessarily yield innovations that empower a wider range of people across our planet to utilize Ethereum technology, and to develop new, transformative applications of it.

Taking a long view, the potential applications blockchain technology in space exploration and commerce are vast. The legal and physical attributes of space—the only domain of human activity not ordered around territorial sovereignty, in which a diverse range of actors from a growing number of countries must coordinate and transact—presents governance challenges to which Ethereum smart contract functionality supplies natural solutions. ConsenSys' present space R&D therefore represents long-term investments in developing Ethereum's full potential, on Earth and beyond.

[41] President John F. Kennedy, Rice Stadium Moon Speech (September 12, 1962).