TOWARD ENABLING SAFE EARTH-INDEPENDENT MISSION OPERATIONS

Shu-chieh Wu, PhD(1), Megan Parisi(2), Kaitlin McTigue(2), Tina Panontin, PhD(1), Alonso Vera, PhD(2)

(1) San José State University at NASA Ames Research Center, Moffett Field, CA, 94035, United States, Email: shu-chieh.wu@nasa.gov, tina.l.panontin@nasa.gov
(2) NASA Ames Research Center, Moffett Field, CA, 94035, United States, Email: megan.e.parisi@nasa.gov, kaitlin.r.mctigue@nasa.gov, alonso.vera@nasa.gov

ABSTRACT

To date, much of space operations safety relies on mission control on Earth. How to enable safe Earth-independent mission operations when communication between the flight crew and mission control is delayed or disrupted is a critical and pressing issue for future deep space crewed missions. The present research examines the past and current process by which unanticipated safety-critical vehicle anomalies are resolved, an activity that has traditionally been ground-driven and manpower-intensive. Preliminary results showed that the sense-making process depicted by the Cynefin Framework can be used to describe the mission control anomaly resolution process. Implications of these results for problem-solving, particularly when it is to be done on-board by a small flight crew, are discussed.

1. INTRODUCTION

Even with considerable planning and risk management, low-probability, high-consequence events continue to occur in complex engineered systems, such as space vehicles. From Apollo to the International Space Station (ISS), space vehicles experience frequent system anomalies during missions. Specifically, the ISS experienced 67 high-priority safety-critical anomalies requiring urgent diagnosis and response between 2000-2018, with more than 65% of them occurring early in the operation history (2000-2006) [1]. Despite these anomaly rates, space safety for missions in low-Earth orbit (LEO) has been largely maintained through 24/7 real-time support provided by an army of specialized flight controllers and vehicle system engineers on the ground in mission control centers (MCC) and mission evaluation rooms (MER). For missions beyond LEO, the safety net provided by the constant, near real-time support of ground teams will be severely limited by delayed and intermittent voice communication and reduced data bandwidth. How to enable safe Earth-independent mission operations is a critical and pressing issue for future deep space crewed missions.

A primary purpose of the present research is to examine the past and current process by which unanticipated safety-critical vehicle anomalies are resolved, an activity that has traditionally been ground-driven and manpower-intensive. We begin with an overview of mission control, followed by a review of current understanding of its anomaly response process. We then introduce the Cynefin framework and describe our attempt to apply it to several anomaly response scenarios. Though conceived as a sensemaking device to guide decision making, we found that the Cynefin framework can help provide a high-level view of the generic anomaly response process when applied retrospectively based on decisions and actions made. We discuss the implications of these results for problem-solving, particularly when it is to be done onboard by a small flight crew.

2. ANOMALY RESPONSE PROCESS OF MISSION CONTROL

Mission control is the generic term for the facility and team of people that manage spaceflight missions from launch to landing or until the end of a mission. A primary responsibility of mission control is flight control, tracing back to tasks typically done in flight testing of airplanes (maintaining communications with flight crew, monitoring flight process and problems, and analyzing flight data) but with the additional overarching objective to observe flight rules [2]. Flight rules are preplanned decisions and agreements that govern the execution of a space mission, put in place to protect the flight crew and public; they typically include criteria for space vehicle performance, go/no-go and flight abort decisions, as well as trajectory and guidance guidelines [3]. Mission control is also responsible for ensuring the mission goes as smoothly and successfully as possible so that it can achieve mission objectives [4]. For that, mission control assists the flight crew in activity planning, procedure executions, and anomaly response.

Complexity of the space vehicles and missions leads to the unique social, technical, and technological multi-team makeup of mission control [5]. The social aspect is composed of approximately 16-24 flight controllers (varying in different mission programs) operating consoles in the front room dedicated to subsystems of the spacecraft or supporting specific tasks (e.g., robotics, spacewalks, planning, surgeon); these controllers are assisted by more controllers and system engineers in the back rooms [4]. Operations of the International Space Station (ISS) receive additional technical support from the Mission Evaluation Room (MER), which is staffed by NASA and contractor engineers from Boeing, the primary contractor of the ISS that built the majority of the US ISS hardware [4], Each front room flight controller is a subject matter expert in a specific technical discipline that goes through a lengthy and vigorous training and certification process [6]. Each console is equipped with applications serving the needs of the specific subsystem or task it supports, with its own specialized functional knowledge, data storage, data models, programmatic interfaces, user interfaces, display of telemetry data related to the specific subsystem, and customized business logic [7]. The ability for a flight controller to access and display telemetry data from subsystems outside of their immediate scope of responsibility is supported but there is no visual cue to alert them when relevant information is on display [2]. Rather, a channel-based groupware technology called voice loops is the primary means by which controllers communicate with other personnel and maintain awareness of current events and activities that occur elsewhere in mission control [8, 9]. In the Shuttle era, each controller often monitored concurrent conversations occurring on a minimum of four voice loops while communicating directly on a primary one. Each voice loop was designed to facilitate the flow of communication around a specific subset of people; the composition of those people in turn affects the nature of information communicated (e.g., assessment of a situation vs. discussion of technical details). Since then, the voice loop system capacity has greatly expanded, and so has the number of loops a flight controller monitors. Currently in International Space Station (ISS) operation, a controller must listen to more than a dozen voice loops at the same time, chosen from thousands of available loops [4].

Anomaly response, or anomaly resolution, refers to the mission evaluation task that includes proper identification of anomalies (i.e., deviations from expected or predicted performance of spacecraft or experiments) and determination of the cause and the proper course of actions to be taken both during the mission at the time of anomaly and before subsequent missions [10]. Given the distributed nature of mission control, the anomaly response process of mission control is characterized in part by the coordination between human agents across functions, time, and physical distance to diagnose a system fault or a cascading set of disturbances and, at the same time, to manage contingencies and maintain the safety of the flight crew and integrity of the mission [11]. Some of the process occurs hyper-locally around the actions of a single flight controller. For example, Watts et al. [9] observed what a flight controller did after she noticed an abrupt change in the telemetry data she was monitoring. Through monitoring voice loops, she sought to determine if the flight crew experienced any anomaly aboard the shuttle and if any other controllers observed any anomaly in the telemetry data of their systems. She also gave a "heads up" to another controller monitoring a related system. In addition to maintaining and expanding the awareness of this data anomaly, she discussed with her back room support staff details of the data anomaly and possible causes. For major events, the process can involve a larger-than-normal group of people called "Team 4" (meaning a fourth shift) led by a flight director to work on a time-critical problem "off-console," in addition to the console controllers providing 24-hour coverage in three shifts [4]. Team 4 utilizes all available resources which include facilities and MER engineers as well as hardware providers and contractors. In Team 4, coordination returns to being in-person and synchronized, with team members working together in a dedicated conference room ("War Room") to discuss solutions or perform analysis; they communicate with console controllers using cell phones.

Similarly, Fiore et al. [12] compare anomaly response in mission control to a complex collaborative problem solving process marked by complexity of the problems (a high number of interconnected elements across technical systems, high degree of uncertainty, shifting task priorities, dynamic systems and conditions) and diversity of the individuals across teams (having different skills, knowledge of technical systems, responsibilities, and priorities). Fiore et al. argue that developing a shared problem model is essential to complex and collaborative problem solving, and it is done through individual and team knowledge building processes. Individual knowledge building processes include individual information gathering and individual information synthesis, similar to the actions taken by the flight controller Watts et al. [9] observed. Development of a shared problem model requires team knowledge building processes, which include team information exchange and knowledge sharing, team solution option generation, team evaluation and negotiation of alternatives, and team process and plan regulation, reminiscent of activities of Team 4 and MER [4].

3. THE CYNEFIN FRAMEWORK

Current understanding of the anomaly response process of mission control paints a process performed by experts situated in a specific social and technological setting. Both the layered communications through voice loops [9] and the knowledge building from individuals to teams [12] reflect a manpower-intensive ground-driven operation model made feasible by the availability of constant real-time crew-ground communication, abundant telemetry, and relevant engineering information. To support Earth-independent anomaly response on missions beyond LEO by a small flight crew, it is important to acquire an understanding of the anomaly response process in a generic way that can then be implemented on-board.

To develop this understanding, anomaly resolution processes were studied using the Cynefin framework. Created by David Snowden in 1999, the Cynefin framework is considered a sensemaking device developed to help decision makers make sense of a wide range of unspecified problems [13, 14]. As shown in Fig. 1, the framework consists of four domains (or open spaces) and a fifth central area, the domain of disorder.

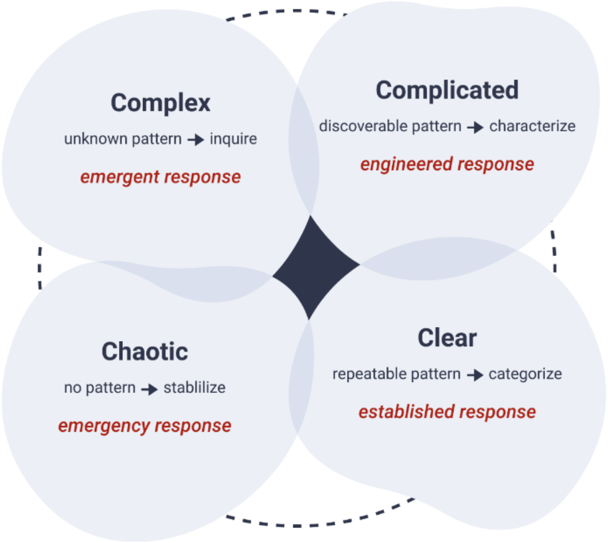

Figure 1. Cynefin framework

On the right are two domains of order. The lower right domain is *Clear* (or *Known*), where the cause and effect related to the problem are known or have been repeated. The appropriate response is readily apparent and already established. The decision model here is to sense and categorize incoming data and respond according to predetermined practice. The upper right domain is *Complicated* (or *Knowable*). Here, stable cause and effect relationships exist but may not be obvious or may be known only by a limited group of people; but the relationships can be discovered (i.e., moving to the Known domain) given resource, capability, and time. The decision model here is to sense and analyze incoming data and then respond according to expert advice or interpretation of the analysis.

On the left are two domains of un-order. The upper left domain is *Complex*, where cause and effect relationships are not directly discoverable through standard methods, possibly hidden among complex interactions that are as yet unknown. The decision model here is to probe and make potential patterns more visible. The lower left domain is *Chaos*, where cause and effect are unknowable given a state of turbulence, and sensemaking cannot be started until the situation is stabilized. The decision model here is to act quickly to reduce the turbulence and then respond accordingly to the reaction of intervention applied.

The Cynefin framework has initially been applied to support decision making on a variety of issues in business management, including knowledge management, product development, strategy, and policy-making [13]. Its application has subsequently been proposed in other domains, including aiding design decisions in the development of large scale complex systems [15], managing risks and uncertainty in cybersecurity [16], and crafting homeland security strategies [17]. Its rich set of considerations for sensemaking contexts and corresponding actions is particularly suited for managing risk and uncertainty associated with anomaly response in space operations. For spaceflight environments, a Chaotic situation aboard the spacecraft would require the crew to rely on their fundamental skills to *stabilize* (safe) the vehicle and maintain flight following emergency procedures. A Complex situation would require the crew to *inquire* and investigate underlying cause and effect relationships and to generate possible solutions. A Complicated (or Knowable) situation would require the crew to analyze and *characterize* the observed cause and effect relationships. A Clear (or Known) situation would simply require the crew to *categorize* a recurring issue of known origin and apply corresponding procedures.

In summary, the Cynefin framework presents a structure for understanding the anomaly response process that can be made generic. In the next two sections, we document our attempt to apply the Cynefin framework to several Apollo and ISS anomaly response scenarios.

4. DATA SOURCE AND ANALYSIS METHOD

Mission Reports and other publicly available articles were analyzed to identify Apollo and ISS anomalies.[1, 4, 18–29]. Anomalies are defined as off-nominal functions of vehicle systems with consequences ranging from benign to life-threatening. Each anomaly selected for inclusion in this report was defined as an anomaly by NASA investigation teams at the time of the incident. Identified anomalies were classified as *significant* incidents if the anomaly affected critical subsystems (e.g., power and life support); depleted essential resources (e.g., atmosphere); and/or involved uncertainty, meaning the event had no set procedure in place for response, required fault tree analysis to discover intermediate and proximate causes, and had potential short-times-to-effect for unwanted consequences. Significance was conferred by an engineering subject matter expert if an anomaly met any of these conditions.

In all, the crewed Apollo missions experienced a total of 362 anomalies across 11 missions. Of these 362 anomalies, 35 were identified as significant incidents requiring urgent diagnosis by a subject matter expert. The

ISS experienced 33 such significant incidents between 2002 and 2019.

Further research was conducted to recreate the sequence of anomaly resolution activities. For Apollo anomalies, available space-to-ground transcripts [30, 31], Public Affairs Office transcripts [30, 31], and Mission Reports [23–27] were reviewed. The space-to-ground transcripts analyzed included instructions from the ground to the crew; crew observations relayed to the ground for analysis; and comments from the ground to the crew on ground analysis currently planned or underway. All utterances captured in the transcripts are timestamped, allowing for the synchronization of events across data sources. Mission Reports, though not timestamped, provide overviews of all in-flight anomalies experienced during a particular mission. These overviews allowed us to fill in technical detail to the specific activities identified through transcript review.

For ISS anomalies, we reviewed publicly available articles detailing anomaly resolution activities [4], as well as ISS on-orbit status reports [32]. We specifically searched for timestamped system information, ground team action items, lines of inquiry presented by ground teams, and any notes on crew actions in order to recreate the sequence of anomaly resolution events.

Using the information described, we recreated the sequence of anomaly resolution activities for one significant ISS anomaly and five significant Apollo anomalies. We then analyzed these sequences using the Cynefin model. For each selected anomaly, each activity in the sequence of events was classified as *stabilization* (Chaotic), *inquiry* (Complex), *characterization* (Complicated), or *categorization* (Clear) using the criteria described below.

Safing activities in response to an anomaly were classified as *stabilization* activities. In assigning these activities, we searched for emergency response actions with the primary purpose of moving the crew and vehicle into a safe configuration. If an activity's primary purpose was to gather information or investigate cause, it was not included in this quadrant. Examples of *stabilization* activities include terminating current operations and completing emergency operating procedures.

We classified activities that served to gather data or analyze information while patterns are still unknown as *inquiry* activities. Here, we searched for troubleshooting actions that produced data for ground analysis. Examples include verification activities (e.g., the ground asking the crew to check if a circuit breaker is closed), and exploratory analyses (e.g., commanding a vehicle system through a series of events and recording the results). We also searched for ground investigations that did not directly reveal predictable patterns, but did provide information for further analysis.

Characterization activities reveal predictable patterns and generate likely hypotheses while the specific origins of the anomaly remain unknown. We classified workaround procedures, fault tree creation, and risks analyses as characterization activities. These activities often narrow the scope of investigations (e.g., the most likely cause is contamination in the abort switch), but if an activity definitively pointed to a cause, it was excluded from this classification.

We classified activities that involved execution of established procedures as part of *categorization* activities in the Clear domain. In this domain, ground controllers applied established procedures to either correct a known recurring problem or as a mitigating option to maintain crew and vehicle safety.

Since anomaly response activity changes as the problem domain changes, domain classification is not a stagnant, one-time deal for a given anomaly. As controllers and crew members gather data, troubleshoot the problem, and uncover patterns, the situation can shift from any one domain to another.

5. RESULTS

As an organization, NASA has made concerted efforts to categorize potential failures for any given mission, and the vast majority of anomalies that do occur during spaceflight have an engineered response already in place. When an anticipated anomaly occurs, ideally, the context shifts quickly from Chaotic directly to Clear as the ground team safes the situation and executes pre-planned procedures. However, even when a problem is anticipated, the established response in place may be inappropriate if the underlying cause is different than anticipated.

Unanticipated anomalies, on the other hand, have no established response in place. For these events, analyses indicate that problem-solving contexts shift as anomaly resolution evolves. Our analysis of anomaly resolution for significant, unanticipated anomalies, did not reveal many instances of resolution that moved sequentially through stabilization, inquiry, characterization, and categorization. A messier overlap of the domains with backtracking and repetition of phases was more common.

However, there were some instances of anomaly resolution in which a pattern with an engineered response did exist within the data, but the ground did not immediately recognize the pattern amidst the slew of downstream effects from the failure. In these cases, the key challenge was to uncover the categorizable pattern within the chaos, at which point resolution followed relatively quickly.

On the Apollo 12 mission, the vehicle experienced a significant anomaly that started out in a state of chaos, but quickly skipped to "clear" when one individual was able to recognize the failure pattern (Figure 2). The Saturn V launch was struck by lightning in the air 36.5 seconds after launch, and again at 52 seconds after launch. From the crew's perspective onboard, alarms indicated complete failure of the electrical power and distribution systems, followed by the primary navigation system. The telemetry stream to Mission Control was severed; the crew read out the onboard cautions & warnings to MCC as the ground assessed the situation. The team determined that the power system and backup navigation systems were functional despite the alarms (1: *stabilize*) and assumed that the problem was a power system malfunction. The ground team continued to analyze the readouts and generate hypotheses (2: *characterize*). After just a few minutes, the EECOM (Electrical, environmental, and consumables manager) flight controller on console recognized a pattern in the jumbled telemetry stream of seemingly random numbers. By chance, he had seen the same pattern of numbers in a past test of the Signal Conditioning Equipment (SCE), and immediately realized that the problem was isolated to the SCE (3: *categorize*). MCC instructed the crew to switch the SCE to its secondary power source, and the onboard instrumentation instantly began receiving normal power again. In summary, this anomaly was resolved when MCC involved the right expert who recognized the pattern quickly due to their extensive experience with the affected subsystem, enabling the team to move through the problem-solving stages quickly and smoothly.

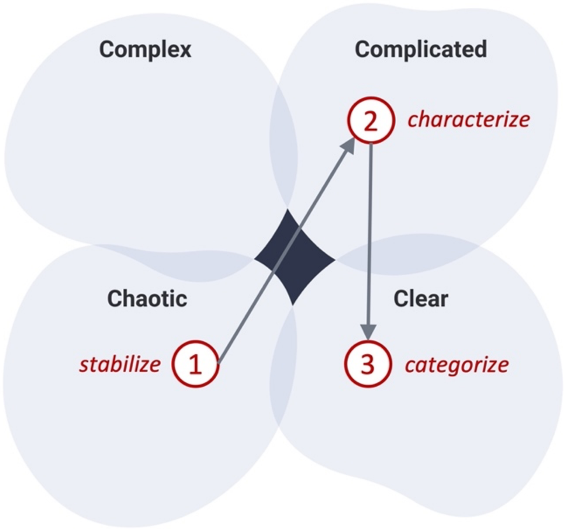

Figure 2: Cynefin model for the Apollo 12 lightning strike anomaly

On the Apollo 15 mission, the Service Module (SM) Propulsion System (SPS) thrust light unexpectedly illuminated with no engine fire command present (Figure 3). MCC dictated the malfunction procedure to the crew (1: *stabilize*) and began troubleshooting the unusual valve position (2: *inquire*). MER engineers determined that the light likely indicated an electrical short rather than a legitimate fire command to the SPS (2: *inquire* (continued)) and a young electrical engineer (working in the MER) named Gary Johnson theorized that the only place to find something floating in a closed cavity (that could cause a short) was the delta-V thrust panel switch, which contained a braided wire (3: *characterize*) [33]. Johnson pitched this theory to Flight Director Christopher Kraft, who directed the team to design a test to prove that the light was illuminating due to a short (4: *inquire*), and not because of a fire command to the SPS (the ground was understandably concerned about the threat of the system firing unexpectedly). Following the procedure, the crew successfully completed the test burn, which gave MCC confidence that that the short had been isolated to the system A delta-V thrust switch (5: *categorize*), and the crew could safely fire the SPS engine and continue the mission. The Apollo 15 SPS anomaly is similar to the Apollo 12 lightning strike anomaly in that the specialized expertise of one engineer contributed to timely diagnosis.

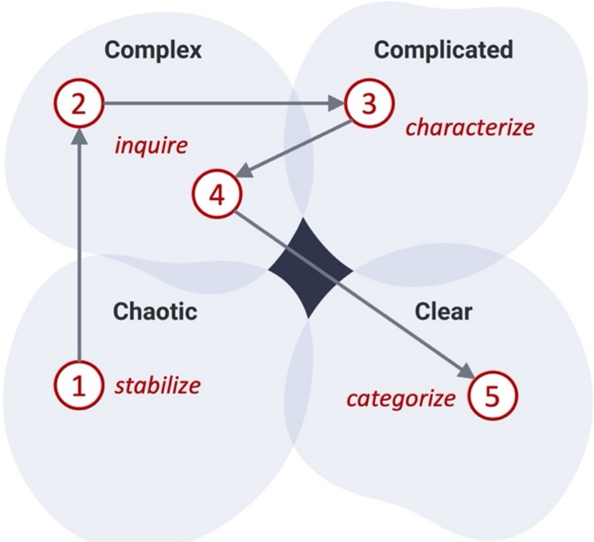

Figure 3: Cynefin model for the Apollo 15 SPS thrust light anomaly

Conceivably, not all anomalies follow such a straightforward trajectory. In many cases, anomaly resolution begins in the chaotic quadrant and then moves back and forth between inquiry and characterization as the investigators begin to gain clarity.

A significant anomaly with this resolution activity pattern took place during Apollo 14 (Figure 4). Around four hours before scheduled powered descent, the ground noticed the abort command in the lunar module guidance computer was set, meaning the descent would have been unintentionally aborted had the ground not noticed the issue. The ground immediately instructed the crew to manually reset the command (1: *stabilize*). Ground teams in MCC began analyzing the event but soon noticed the abort command had been set again. Suspecting a potential hardware issue, the ground asked the crew to attempt to reset the command by physically tapping the panel around the abort pushbutton (2: *inquire*). The ability to physically reset the command, the intermittent nature of the error, and the error being limited to only the lunar module primary guidance computer (and not the abort guidance computer downlink) suggested metal contamination in the abort switch (*3: characterize*). During these troubleshooting efforts, programmers at the Massachusetts Institute of Technology (MIT) were developing preliminary procedures for inhibiting the abort command in the guidance and navigation system (4: *characterize*). Teams at MIT and in MCC ran simulations to verify the procedure but had to search for procedural errors when the simulations ended in crashes of the lunar module multiple times (5: *inquire*). The ground teams reworked the procedure, successfully verified operations in simulations, and sent the completed procedure to the crew (6: *characterize*). During powered descent, the crew manually completed the procedure and landed successfully (7: *categorize*).

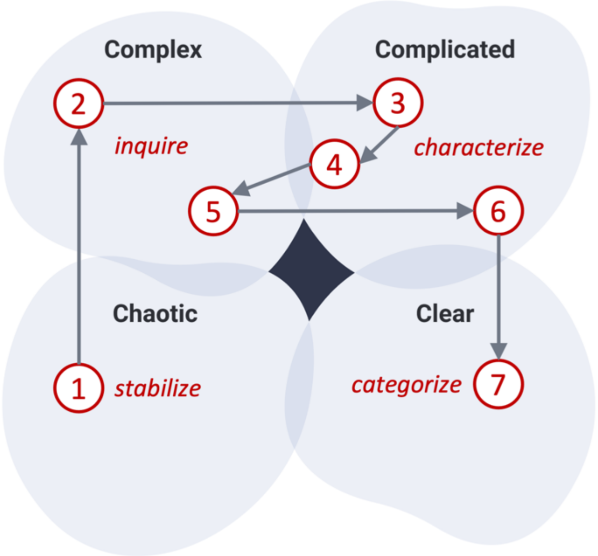

Figure 4: Cynefin model with a resolution path that doubles back: Apollo 14 abort command anomaly

A similar anomaly resolution pattern unfolded for an International Space Station (ISS) anomaly in 2013. Flight controllers in MCC realized cooling loop A had shut down due to an under-temperature warning, meaning half the systems onboard the ISS were in danger of overheating. Ground teams immediately determined which systems could be shut down and which needed to be switched to cooling loop B, while simultaneously beginning procedures to restart the pump in cooling loop A (1: *stabilize*). Temperature data monitored by the ground indicated the Flow Control Valve (FCV) was open but should have been closed. After the pump restarted, the ground commanded the FCV through its full range and noticed a large position offset that could not be overcome (2: *inquire*). Ground teams then conducted analyses to predict the behavior of the ammonia at different temperatures and pressures (3: *characterize*). Teams attempted commanded-from-the-ground work-arounds for warming the loop, including moving the FCV outside of its normal limits, but these attempts were unsuccessful (4: *inquire*). Engineers in the Mission Evaluation Room (MER) generated fault trees to investigate potential proximate causes for the failure (5: *characterize*). Taken together, these inquiries and characterizations suggested the FCV was not fixable via commanded from the ground options. The team decided to replace the pump module with an EVA utilizing established procedures (6: *categorize*).

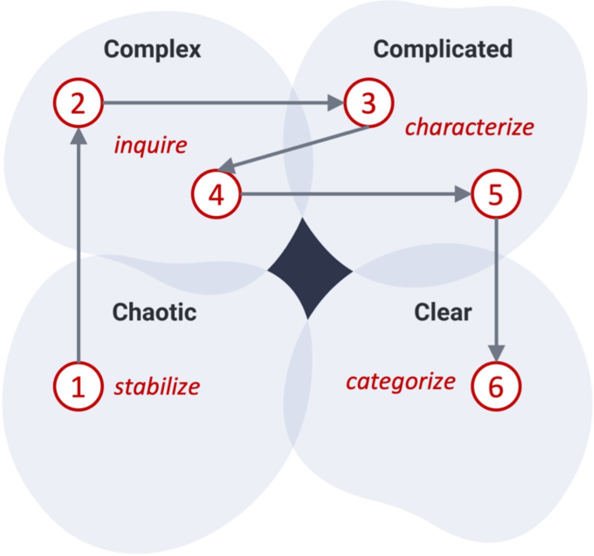

Figure 5: Cynefin model with a resolution path that doubles back: ISS cooling loop anomaly

Perhaps the most infamous anomaly in the history of human spaceflight is the oxygen tank explosion that occurred on Apollo 13, crippling the vehicle, and jeopardizing the lives of the crew. Interestingly, the resolution activity pattern for this event begins with inquiry before moving to stabilization.

The Apollo 13 crew was alerted to the malfunction when they heard a loud bang onboard, accompanied by an undervoltage alarm on the AC bus. A few minutes later the crew witnessed an unidentified gas venting out of the vehicle. MCC got to work monitoring data and troubleshooting to determine if the problem was an instrumentation error (1: *inquire*). After determining that the failure was critical and time-sensitive, MCC began shedding loads to save power (2: *stabilize*) while they performed further troubleshooting to understand the severity of the failure and determine times-to-effect for power loss and other downstream consequences (3: *inquire*), and a separate team worked to uncover the cause of the failure. At this point the EECOM team developed a theory that "we blew a O2 line in one of the fuel cells," and the crew performed a procedure to assess the state of fuel cells under the direction of MCC (4: *characterize*). The team realized that the fuel cells were likely unrecoverable and directed troubleshooting efforts at salvaging the remaining resources from the Command and Service Module (CSM) and transitioning the crew to the Lunar Module (LM) as a "lifeboat" (5: *categorize*). The TELMU (Telemetry, electrical, EVA mobility unit officer) MER team had created an emergency operations plan for the LM lifeboat model several years prior to Apollo 13, but this plan was intended to be used for a much shorter period of time than what was needed, and demanded a lot of alteration and additional planning.

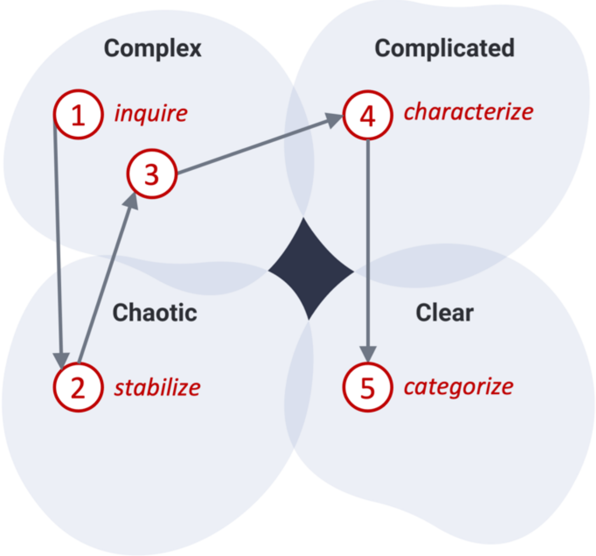

Figure 6: Cynefin model of the Apollo 13 oxygen tank explosion anomaly

An anomaly in the secondary Service Module Propulsion System during Apollo 16 also started with inquiry (Figure 7). While preparing to move the CSM into a circular orbit prior to Lunar Module descent on Apollo 16, the command module (CM) pilot noticed oscillations in the secondary Service Module Propulsion System (SPS) engine gimbal position indicator before a planned circularization maneuver. While the primary SPS engine gimbal position indicator still functioned as expected, a defunct secondary gimbal would eliminate a needed redundancy, and flight rules would necessitate aborting the lunar landing. The CM pilot deactivated and reactivated the gimbal and cycled the gimbal through Auto mode, but the issue persisted (1: *inquire*). The pilot then aborted the planned circularization maneuver (2: *stabilize*). The ground asked the CM pilot to complete a gimbal dive check. The CM pilot moved through each gimbal position and read out the corresponding onboard indications (3: *inquire*). After this inquiry, the ground still could not determine if the gimbal oscillations would cause potentially catastrophic issues should the crew use the secondary gimbal. If the issue could not be characterized during the next five lunar revolutions, the lunar module landing would be aborted. The ground commanded the CM pilot to move the command and service module into a "station-keeping position" next to the lunar module, allowing for Trans-Earth injection should aborting the mission become necessary (4: *stabilize*). The ground recognized that the rotational hand controller (RHC) could be inducing noise into the system. The ground asked the CM pilot to cut all power to the RHC, cut off both alternating and direct current, and repeat the gimbal check (5: *inquire*). This troubleshooting activity eliminated the RHC as a potential cause. While the crew worked on these troubleshooting activities, ground teams ran simulations to determine if the gimbal oscillations would present any structural hazards should the secondary gimbal be used (6: *characterize*). These analyses allowed the ground to determine the issue was stemming from an open circuit. Simulations showed no structural hazards should the secondary gimbal be used, and the ground gave a go for landing (7: *categorize*).

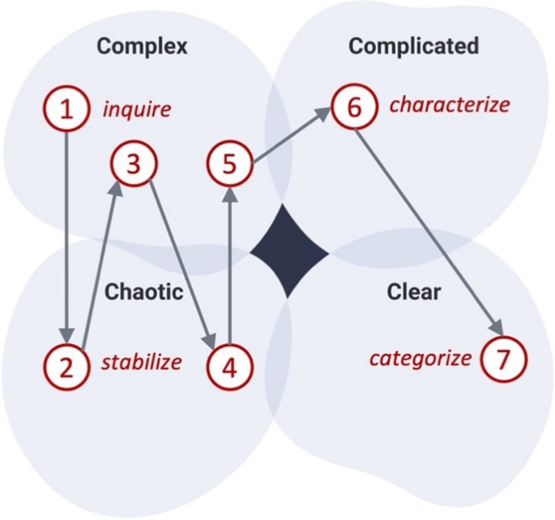

Figure 7: Cynefin model with a resolution path that begins with inquiry: Apollo 16 SM Propulsion System engine gimbal anomaly

6. DISCUSSION

A feature shared by these results is movement between domain boundaries over the course of the anomaly response process. Though not always following a clean path, Cynefin models for most of the incidents analyzed above show a general movement from the Chaos domain to the Complicated domain, and occasionally the Clear. Though the creators of Cynefin emphasize that the framework is used primarily to consider the dynamics of situations for decision making and no one domain is more desirable than any other, they do believe how decision makers think about moving between domains is as important as the way they think about the domain they are in, because moving across boundaries from one domain to another entails a shift to a different context requiring a different model of understanding and interpretation [13]. For example, Kurtz and Snowden [13] describe the movement from the Knowable (Complicated) to the Complex as Exploration, which represents an opening up of possibilities. Conversely, movement from the Complex to the Knowable (Complicated) is described as Exploitation, which often involves refinement of selective stable patterns for ordered representation. This distinction echoes a similar one in organizational learning between the exploration of new possibilities and the exploitation of old certainties [34] and has significant implications on the resilience of an anomaly response process. Watts-Englert, Woods, and Patterson argue that a premature narrowing of options (cf. exploitation) can cause breakdown in anomaly response if evaluation of the options is not revised as the anomaly evolves and new evidence becomes available; conversely, incorporating mechanisms that can help broaden the set of options considered (cf. exploration) in diagnosis and replanning can increase resilience [2].

6.1. Problem Solving and Domain Expertise

Although the Cynefin framework provides guidance for determining the optimal course of action given a problem context, its utilization is predicated on the ability of a problem solver to recognize the context they are in. The ability to recognize patterns in situations has been identified as the key to how experts can quickly make decisions in high-consequence high-risk environments such as firefighting, aviation, medicine, nuclear power, military command and control, oil/gas extraction [35]. Patterns can highlight relevant cues, provide expectancies and suggest typical responses in the respective type of situations [36]. The Cynefin models presented above represent the expert and arguably the optimized versions of the response processes of those anomalies as the responses were carried out by a large team of experts using all conceivable resources mission control could offer [37]. Various training models have been developed to assist the crew in decision making [e.g., 38]. What the results of the Cynefin models suggest is that, in addition to providing onboard data and tools, it is important to support the future Earth-independent crew in recognizing what problem contexts they are in in the first place. Crew members must understand “what they know, what they don’t know, and what they need to know” at any given point and take a systematic approach to their investigation.

6.2. Limitations

In the present research, classification of activities into Cynefin domains was done based on the perceived nature of the actions as described in the data sources. The accuracy of classifications included in the results here is subject to at least two limitations. One limitation is that the data sources may not have captured or presented the most complete picture of the context faced and constraints considered by the crew and the ground team at the time. For example, transcripts provided by the Public Affairs Office (PAO) appear to include commentary on ground team analysis activities not communicated in full to the crew. A second limitation is that there may be more than one way to classify an activity, depending on the scope of context considered. For example, in the analysis for the Apollo 16 SM Propulsion System engine gimbal anomaly (Figure 7), we classified the CM pilot’s action to move the command and service module into a “station-keeping position” as a stabilization activity based on its utility as preparation for a contingency maneuver. It is probably equally valid, however, to classify it as the execution of an established procedure in the Clear domain where the ground recognized (categorized) the situation and what needed to be done. These limitations prompt further considerations and refinements in applying the Cynefin framework to understand the anomaly response process.

7. ACKNOWLEDGEMENT

This work was supported by the Human Factors and Behavioral Performance Element of NASA’s Human Research Program and by NASA Engineering and Safety Center. We thank Dorrit Billman and John Karasinski for comments on an earlier draft of the paper.

8. REFERENCES

1. Wu, S.-C., & Vera, A. H. (2019). Supporting Crew Autonomy in Deep Space Exploration: Preliminary Onboard Capability Requirements and Proposed Research Questions. Technical Report of the Autonomous Crew Operations Technical Interchange Meeting (NASA Technical Memorandum No. NASA/TM—2019–220345). Moffett Field, CA: NASA Ames Research Center. Retrieved from http://ntrs.nasa.gov/search.jsp?R=20190032086

2. Sgobba, T. (2018). Chapter 13 - Human factors in mission control centers. In B. Kanki, J.-F. Clervoy, & G. M. Sandal (Eds.), Space Safety and Human Performance (pp. 569–620). Butterworth-

Heinemann. https://doi.org/10.1016/B978-0-08-101869-9.00013-3

3. Keyser, L. W. (1974). Apollo Experience Report - The Role of flight mission rules in mission preparation and conduct (No. NASA TN D-7822). Houston, TX: Lyndon B. Johnson Space Center.

4. Dempsey, R. C. (Ed.). (2018). The International Space Station: Operating an Outpost in the New Frontier. Retrieved from https://www.nasa.gov/feature/new-nasa-e-book-offers-inside-look-at-space-station-flight-controllers

5. Caldwell, Barrett. S. (2005). Multi-Team Dynamics and Distributed Expertise in Mission Operations. Aviation, Space, and Environmental Medicine, 76(6), B145–B153.

6. Dittemore, G., & Bertels, C. (2011). The Final Count Down: A Review of Three Decades of Flight controller Training Methods for Space Shuttle Mission Operations. In AIAA SPACE 2011 Conference & Exposition. American Institute of Aeronautics and Astronautics. https://doi.org/10.2514/6.2011-7190

7. Trimble, J. P., Dayton, T., & Quinol, M. (2012). Putting the Users in Charge: A Collaborative, User Composable Interface for NASA's Mission Control. In 2012 45th Hawaii International Conference on System Sciences (pp. 674–683). Presented at the 2012 45th Hawaii International Conference on System Sciences, Maui, HI. https://doi.org/10.1109/HICSS.2012.494

8. Patterson, E. S., Watts-Perotti*, J., & Woods, D. D. (1999). Voice Loops as Coordination Aids in Space Shuttle Mission Control. Computer Supported Cooperative Work (CSCW), 8(4), 353–371. https://doi.org/10.1023/A:1008722214282

9. Watts, J. C., Woods, D. D., Corban, J. M., Patterson, E. S., Kerr, R. L., & Hicks, L. C. (1996). Voice loops as cooperative aids in space shuttle mission control. In Proceedings of the 1996 ACM conference on Computer supported cooperative work - CSCW '96 (pp. 48–56). Presented at the the 1996 ACM conference, Boston, Massachusetts, United States: ACM Press. https://doi.org/10.1145/240080.240188

10. Lobb, J. D. (1975). Apollo Experience Report: Flight Anomaly Resolution (NASA Technical Note No. NASA TN D-7968) (p. 74). Houston, TX: Lyndon B. Johnson Space Center. Retrieved from https://ntrs.nasa.gov/api/citations/19750017909/downloads/19750017909.pdf

11. Watts-Perotti, J., & Woods, D. D. (2009). Cooperative Advocacy: An Approach for Integrating Diverse Perspectives in Anomaly Response. Computer Supported Cooperative Work (CSCW), 18(2), 175–198. https://doi.org/10.1007/s10606-008-9085-4

12. Fiore, S. M., Wiltshire, T. J., Oglesby, J. M., O'Keefe, W. S., & Salas, E. (2014). Complex Collaborative Problem-Solving Processes in Mission Control. Aviation, Space, and Environmental Medicine, 85(4), 456–461. https://doi.org/10.3357/ASEM.3819.2014

13. Kurtz, C. F., & Snowden, D. J. (2003). The new dynamics of strategy: Sense-making in a complex and complicated world. IBM Systems Journal, 42(3), 22.

14. Snowden, D. (2002). Complex acts of knowing: paradox and descriptive self-awareness. Journal of Knowledge Management, 6(2), 100–111. https://doi.org/10.1108/13673270210424639

15. Basha, N. S., Kwasa, B., Kanan, H., & Bloebaum, C. L. (2017). Cynefin sensemaking framework for decision support in design and development of large scale complex engineered systems. In Proceedings of the American Society for Engineering Management 2017 International Annual Conference (p. 1).

16. Dykstra, J. A. B. S., & Orr, S. R. (2016). Acting in the unknown: the cynefin framework for managing cybersecurity risk in dynamic decision making. In 2016 International Conference on Cyber Conflict (CyCon U.S.) (pp. 1–6). Presented at the 2016 International Conference on Cyber Conflict (CyCon U.S.). https://doi.org/10.1109/CYCONUS.2016.7836616

17. Bellavita, C. (2006). Changing Homeland Security: Shape Patterns, Not Programs. Homeland Security Affairs, II(3), Article 5.

18. Apollo 7 Mission Evaluation Team. (1968). Apollo 7 Mission Report (Mission Report No. MSC-PA-R-68-15). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a410/A07_MissionReport.pdf

19. Apollo 8 Mission Evaluation Team. (1969). Apollo 8 Mission Report (Mission Report No. MSC-PA-R-69-1). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a410/A08_MissionReport.pdf

20. Apollo 9 Mission Evaluation Team. (1969). Apollo 9 Mission Report (Mission Report No.

MSC-PA-R-69-2). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a410/A09_MissionReport.pdf

21. Apollo 10 Mission Evaluation Team. (1969). Apollo 10 Mission Report (Mission Report No. MSC-00126). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a410/A10_MissionReport.pdf

22. Apollo 11 Mission Evaluation Team. (1971). Apollo 11 Mission Report (Special Publication No. NASA SP-238). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a11/a11MIssionReport_1971015566.pdf

23. Apollo 12 Mission Evaluation Team. (1970). Apollo 12 Mission Report (Mission Report No. MSC-01855). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a12/A12_MissionReport.pdf

24. Apollo 13 Mission Evaluation Team. (1970). Apollo 13 Mission Report (Mission Report No. MSC-02680). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a13/A13_MissionReport.pdf

25. Apollo 14 Mission Evaluation Team. (1971). Apollo 14 Mission Report (Mission Report No. MSC-04112). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a14/A14MRntrs.pdf

26. Apollo 15 Mission Evaluation Team. (1971). Apollo 15 Mission Report (Mission Report No. MSC-05161). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a15/ap15mr.pdf

27. Apollo 16 Mission Evaluation Team. (1972). Apollo 16 Mission Report (Mission Report No. MSC-07230). Houston, TX: Manned Spacecraft Center. Retrieved from https://www.hq.nasa.gov/alsj/a16/A16_MissionReport.pdf

28. Apollo 17 Mission Evaluation Team. (1973). Apollo 17 Mission Report (Mission Report No. JSC-07904). Houston, TX: Lyndon B. Johnson Space Center. Retrieved from https://www.hq.nasa.gov/alsj/a16/A16_MissionReport.pdf

29. Gore, B. F., Vera, A. H., Marquez, J. J., Holden, K. L., Amick, R. Z., Dempsey, D. L., & Munson, B. A. (2021). Evidence Report: Risk of Inadequate Human-System Integration Architecture (p. 254). Houston, TX: NASA Johnson Space Center. Retrieved from https://ntrs.nasa.gov/api/citations/20210012912/downloads/HSIA_Evidence_Report_Final_04-27-21.pdf

30. Apollo Lunar Surface Journal. (2018). Retrieved October 11, 2021, from https://history.nasa.gov/alsj/

31. Apollo Flight Journal. (2019). Retrieved October 11, 2021, from https://history.nasa.gov/afj/index.html

32. Keeter, B. (2021). ISS On-Orbit Status Report. Retrieved October 14, 2021, from https://blogs.nasa.gov/stationreport/

33. Johnson, G. W. (2018). Lessons Learned from 50+ Years in Human Spaceflight and Safety (Contractor Report No. JSC-E-DAA-TN62036). SAIC. Retrieved from https://ntrs.nasa.gov/citations/20190028301

34. March, J. G. (1991). Exploration and Exploitation in Organizational Learning. Organization Science, 2(1), 71–87.

35. Orasanu, J. (2005). Crew Collaboration in Space: A Naturalistic Decision-Making Perspective. Aviation, Space, and Environmental Medicine, 76(6), B154–B163.

36. Klein, G. (2008). Naturalistic Decision Making. Human Factors, 50(3), 456–460. https://doi.org/10.1518/001872008X288385

37. McTigue, K. R., Parisi, M. E., Panontin, T. L., Wu, S.-C., & Vera, A. H. (2021). Extreme Problem Solving: The New Challenges of Deep Space Exploration. Presented at the Space CHI: Human Computer Interaction for Space Exploration, A Workshop at ACM CHI 2021. Retrieved from https://human-factors.arc.nasa.gov/publications/McTigue_Parisi_Panontin_Wu_Vera_SpaceCHI.pdf

38. Martin, L., O’Keefe, W., Schmidt, L., Barshi, I., & Mauro, R. (2012). Houston, We Have a Problem Solving Model for Training. Journal of Organizational Psychology, 12(1), 57–68.