THE APPLICATION OF RESILIENCE ENGINEERING TO HUMAN SPACE FLIGHT

C.W. Johnson [(1)], A. Herd [(2)], M. Wolff [(3)],

[(1)] Department of Computing Science, University of Glasgow, Scotland.
http://www.dcs.gla.ac.uk/~johnson, Email: Johnson@dcs.gla.ac.uk
+44 (0)141 330 6053 (Tel.), +44 41 330 4913 (Fax).

[(2)]ESA Operations Safety Unit, D/OPS-H & ESA Independent Safety Office, ESTEC,
Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands.
Fax: +31 71 565 6278, Phone: +31 71 565 6745, Mobile: + 31 650 685425

[(3)] Software Systems Division, Directorate of Technical and Quality Management,
ESTEC Keplerlaan 1, 2201 AZ Noordwijk, The Netherlands.
Fax: +31 71 565 5420, Phone: +31 71 565 3206

ABSTRACT

There is a concern that mishap reporting systems track information about previous anomalies but do little to protect against future failures. In contrast, the term 'resilience engineering' has been coined to describe techniques that identify the thousands of everyday positive actions that prevent accidents from occurring. This change of perspective poses significant theoretical and pragmatic challenges. Just as it can be difficult to identify the causes of previous failures, it is equally difficult to determine what went right and why. We cannot always be sure how close our successes came to failure. Resilience engineering has not previously been applied in any sustained way to space missions. This paper, therefore, uses resilience engineering concepts to analyse the many different ways in which successive crews responded to engineering challenges on the International Space Station (ISS).

1. INTRODUCTION

Space research is a high-risk enterprise. Most missions involve technological innovation through the integration of complex, dynamic systems with finite budgets and timescales. A significant proportion of any investment must, therefore, be devoted to identify and mitigate hazards to people, the environment and mission objectives. Fortunately, mishaps are relatively rare given the operational demands. Space agencies, therefore, integrate 'lessons learned' techniques into their wider processes for safety management. For instance, NASA (2006) has developed Procedural Requirements for Mishap and Close Call Reporting, Investigating, and Recordkeeping (NPR 8621.1B). This document describes the processes that should be followed to safeguard the scene of an incident and then to initiate the investigatory processes that help to identify the engineering and managerial insights that can be derived in aftermath of an adverse event.

The European Space Agency operates a similar approach through the Inspector General's Engineering Knowledge Office. Like NASA, they not only focus on the insights that are to be learned after previous incidents. The Inspector General's remit also includes the dissemination of 'good practice' following successful operations; "Ensuring that the Agency's internal Lessons Learned system reflects the important lessons learned through reviews, enquiry boards and informal contacts with project engineers. The Lessons Learned system is a database on the Agency's intranet where extracts from technical reviews are made accessible to other projects. All Lessons Learned, whether positive or negative, are included together with an explanation of their context of discovery and their likely applicability to future activities" (ESA, 2009).

There are a number of techniques that can be recruited to identify the causes and contextual factors that lead to space mission failures. These include but are not limited to STAMP, Why-Because Analysis, Events and Causal factors analysis (Johnson, 2003). In contrast, there are relatively few techniques that could be employed to provide a systematic framework for analysing those factors that contribute to missions in which safe and successful operations were not compromised. The innovative aim of resilience engineering is to provide the tools to help engineers and managers learn both from failure but also from the lessons of success. However, a number of significant challenges remain. In particular it is not clear how resilience engineering might be integrated with more conventional approaches that are based around hazard analysis and risk assessment, for instance to support assembly, integration and verification. The following pages, therefore, identify both the strengths and the potential weaknesses of resilience engineering for space operations.

2. WHAT IS RESILIENCE ENGINEERING

Resilience can be defined as a self-regulating ability of systems to adjust in response to changes in the operating environment in order to sustain required operations under both expected and unexpected conditions (Hollnagel, Woods and Leveson, 2006). Resilience engineering promotes these self-regulating responses. It is particularly important to encourage those adaptations in performance that promote successful operations in high reliability organisations with a relatively good safety record. In such situations, mishaps are so rare that they provide almost no useful information about everyday operations. In contrast, most activities lead to successful outcomes, from this it follows that it is important to understand the intervention of individuals and teams that promote safe behaviours beyond 'industry standards'. A particular example of this is the manner in which the ESA ATV-1 teams have worked to ensure that there is no sense of complacency in the legacy that they have created for ATV-2 and 3. Such initiatives have been encapsulated within the premises that guide resilience engineering:

1. There is always a degree of uncertainty in the engineering and operation of complex systems. Successful operation, therefore, depends upon individuals and organizations adjusting what they do to match current demands and resources. Because resources and time are finite, such adjustments will inevitably be approximate.

2. Some adverse events can be attributed to the breakdown or malfunctioning of components. In other cases, problems arise in spite of normal system operation as the result of unexpected combinations of performance variability. Most engineers are skilled in identifying and mitigating the first sort of failure. However, relatively little attention is paid to the variability of normal operations until they result in mishaps.

3. Safety management cannot be based exclusively on hindsight, from the investigation of previous failures, nor rely on the calculation of failure probabilities. Safety management must be proactive in strengthening recognized good practices as well as responding to previous mishaps. This in turn requires the development of specific methods for identifying 'good practices' in the first place and, as noted before, it is often hard to determine how close complex missions come to potential failure.

4. Safety cannot be isolated from the core mission objectives or vice versa. Safety is a prerequisite for successful operations and mission success is a prerequisite for safety. Creating organizational or procedural distinctions between these two areas will be counterproductive (Leonhardt et al, 2009).

These four principles build upon the existing 'lessons learned' techniques that have been integrated into the safety management systems of most leading space agencies. In particular, resilience engineering promotes an innovative focus on the variability of 'normal' operations and on the everyday adaptations that people make to 'get the job done'.

3. APPLICATIONS TO SPACE SAFETY

A number of attempts have been made to apply resilience engineering to space based systems. In particular, David Woods (2005), one of the pioneers in this area, advocated the use of this approach in his testimony following the Columbia accident. He argued that the mishap stemmed, in part, from:

- A drift toward failure as defences erode in the face of production pressure. These can be seen as adaptations to launch schedules that eroded the margin of safety.
- The assumption that past success is a reason for increased confidence rather than increased investments to anticipate the changing potential for failure.
- A fragmented problem-solving process that obscured the bigger picture.
- A failure to revise assessments as new evidence accumulates.
- A failure in communication and coordination between organisations that was symptomatic of deeper problems in the programme.

His testimony went on to argue for a new safety organisation that would "use the tools of Resilience Engineering to monitor for 'holes' in organizational decision making and to detect when the organization is moving closer to failure boundaries than it is aware. Together these processes will create foresight about the changing patterns of risk before failure and harm occurs" (Woods, 2003). The following sections describe attempts to realise part of Woods' vision for resilience engineering in space applications by applying the approach to identify the variability of everyday working practices that contribute to safe and successful operations on the International Space Station.

Resilience engineering describes a set of ideas and concepts that have profound implications for safety management. However, it can be difficult to identify the precise tools or techniques that might be used to strengthen the precursors for success. For instance, the

proponents of the approach argue that complex organisations should continually monitor their position along the following dimensions:

- **Anticipatory**: organizations should proactive in identifying evidence of developing problems rather than reacting after problems become significant.

- **Observable:** organisations should monitor safety boundaries and recognize how close they are to 'the edge', for instance by surveying defences and barriers. Information about safety concerns should be widely distributed throughout the organization at all levels and not closely held by a few individuals.

- **Flexible:** organizations should recognise and respond to changes in their safety culture as well as to disruptions, and opportunities in their operating environment.

- **Open to revision: o**rganizations should regularly update their model of vulnerabilities and reassess the effectiveness of countermeasures over time.

Previous sections have mentioned the difficulty of incorporating resilience engineering concepts into existing design and development techniques. For instance, how might a safety manager assess the flexibility of their operations in response to a wide range of organizational pressures? Such attributes are often extremely subjective; one individual might identify an appropriate degree of improvisation in the face of an engineering problem while another identifies a violation of agreed procedures. Similarly, the response to such questions is likely to be highly situated. In some context there may be the time to improve observability through the distribution of information across multi-disciplinary teams. In other situations national security or commercial barriers can prevent the distribution of safety information. One of the reasons for this is the relative novelty of the concepts. The initial ideas for resilience engineering are only 5-6 years old whereas traditional forms of risk assessment and accident analysis have a history stretching back many decades.

The remainder of this paper shows how elements of the resilience engineering approach can be applied to a case study that integrates both the adverse consequences of unexpected failure and also the successful improvisation that characterises many space operations. The urine reprocessing elements of the International Space Station (ISS) Environmental Control and Life Support System provides an appropriate example also because it provides critical lessons about our capacity to cope with infrastructure failures during any future or long-duration human space missions.

4. THE ISS WATER RECOVERY CASE STUDY

The following paragraphs use the attributes of resilience engineering, enumerated in the previous list, to help identify the positive ways in which the crews responded to a series of engineering challenges. In particular, we show how ground-based planning and operations teams cooperated with the crews of STS-126 and the ISS to install and maintain the Urine Processing Assembly (UPA) within the Environmental Control and Life Support System (ECLSS). The development of these systems is essential for long duration space missions but also has applications in areas of the globe that suffer from contaminated water supplies. This paper is, in part, a testimony to the resilient behaviours of these individuals that are often under-valued by the wider engineering community.

Not only does the International Space Station (ISS) support an innovative range of scientific experiments, the on-board infrastructure also demonstrates many innovative technologies. For example, the ISS has two water recovery systems. The Zvezda Service Module at the heart of the Russian Orbital Segment has its own system to process waste water from showers, sinks etc. The output can be drunk in an emergency but is normally used by the Elektron oxygen generation assembly. The US Orbital Segment has its own Water Recovery System. This can handle up to 23 pounds of condensate, crewmember urine, and urinal flush water to produce a purified distillate. The output from the UPA is combined with other wastewater sources collected from the crew and cabin and is processed, in turn, by a Water Recovery System (WRS) to produce drinking water for the crew.

The US Orbital Segment's Water Recovery System was installed during STS-126. This mission launched on the 14th November 2008 with 32,000 pounds of infrastructure equipment intended to support twice the existing ISS crew of six. A further aim was to help sustain operations on the ISS after the point at which the shuttle fleet is retired. By reprocessing water, there would be less need for Soyuz launch mass to be used on delivering drinking water. The programmes to increase the capacity of the ECLSS began some five years before the projected end of the shuttle programme. The testing programme for the water recovery system began at Marshall in May 1990. As we shall see, this provided an important 'margin' of extra time that helped crews face the challenges created by system failure. These initiatives illustrate the way in which the ISS programme worked to **anticipate** future problems rather than react to existing failures.

The Water Recovery System (WRS) in the US Orbital Segment takes waste water, including the crews' urine, and produces water fit for drinking. It consists of:

1. The Urine Processor Assembly. This uses low pressure, vacuum distillation with a centrifuge to compensate for the lack of gravity during the separation of liquids and gasses. Water from the Urine Processor Assembly and from waste water sources are combined to feed the Water Processor Assembly.

2. The Water Processor Assembly. This further filters gasses and solid materials using filter beds and a high-temperature catalytic reactor. The water is then tested by onboard sensors. Any water that does not meet the relevant standards is then recycled through the water processor assembly.

As might be expected, this configuration is the product of a significant development process. Previous shuttle missions, including STS-89 and STS-107, had carried test components but there had been problems with the initial prototype for the compression and distillation units. The eventual design was delivered for installation on STS-126. Brevity prevents a detailed discussion of the design changes that were made after these missions. However, it is sufficient to remark that they arguably illustrated a lack of **flexibility**. This is an important observation – there is inevitably a compromise between many of these attributes in resilience engineering. To show total flexibility and abandon the use of low pressure, vacuum distillation would have threatened the overall mission objectives behind the ELCSS by delaying installation beyond the point at which the shuttle fleet could have been retired.

By 19th November 2008, the crew of STS-126 had worked with the crew of the ISS to install the Reactor Health Sensor and the Catalytic Reactor of the Water Processor Assembly in the Water Recovery System (Harwood, 2008). On the following day, the crews started to install the Total Organic Carbon Analyzer on the front side of the Oxygen Generation System Rack. They also prepared for the initiation of the Urine Processing Assembly by filling a filter tank with pre-treated urine and starting the processing activity. The intention was to activate the process as soon as possible so that STS-126 could return a sample to earth for analysis. This would be used to calibrate the system before a further test phase could begin. NASA managers hoped to collect test data on the urine recycling system for 90 days before a dress-rehearsal in February using the crew of the following shuttle mission to simulate the load imposed on the ISS life support systems when the permanent crew size was increased. This cautious approach was essential to mitigate the risks associated with any potential source of illness that could debilitate more than one member of the ISS crew. These studies were intended to enhance the **observability**, mentioned in the previous discussion of resilience engineering. The stringent set of tests, including bacteriological and taste studies, were intended to ensure that all stakeholders were convinced it would be safe to drink.

Problems started on the 20th November when a test was being conducted on the Urine Processing Assembly. These tests were started towards the end of an Extra-Vehicular Activity (EVA). However, work on the test was interrupted when one of the astronaut's spacesuits showed a build-up of carbon dioxide. Flight controllers told him to return to the airlock as a precaution. On the way back, he experienced further problems in hearing his crewmates and flight controllers. It later emerged that his headset volume control knob had been inadvertently turned down. It was at this time that the remaining crew on the ISS began to hear the alarm associated with the Urine Processing Assembly test. Data was sent back to the ground teams for analysis. Initial concern focussed on the cooling that was provided to the equipment racks. The crews followed operating procedures after such an alarm and de-powered the unit to check for combustible products. The Flight Director commented "This particular time, we were suspicious of the response because we knew the commands we were sending at that time should not have initiated that response. When the crew members confirmed that they had no concerns, no smell of smoke or no odour, especially when they told us the combustion products were all reading zero, we began to think it was a false indication. That was indeed the case." The decision was, therefore, taken to continue water reprocessing without the urine assembly during the evening. The immediate response to the Urine Processing Assembly Alarm provides further examples of **flexibility** during different phases of the response to complex systems failure. In the immediate aftermath of the warning, the crew followed standard operating procedures to ensure that there was no risk of a potential fire. Thereafter, the ground support groups showed a more flexible consideration of the possible causes, including a potentially false indication.

The Urine Processing Assembly would initially work for between two and three hours before shutting down with an alarm. The Flight Director summarized the situation to the crew; "the UPA caution software, we learned some things about it last night in terms of some malfunctions that are paired up with messages that are different than we had thought originally. The UPA hazard message is one you got yesterday and it turns out that message is not critical and it's going to be suppressed on board. So you'll see it, but you won't hear

it. So if you do see that, there's not going to be any actions for that." It could be argued that the decision to suppress the Urine Processing Assembly alarm undermined the **observability** of potential safety problems. However, this decision was only taken after the ground teams were sure that it was a false alarm.

The **observability** of the safety status for the ELCSS was further supported by the ways in which both crews cooperated with the ground teams to identify hypothesis to explain the continuing problems with the Urine Processing Assembly as more data became available. The processing relied on a centrifuge to compensate for the lack of gravity during the separation of liquids and gasses. Monitoring results showed that the centrifuge motor was slowing down and drawing higher than normal current. It seemed as though the internal protection software was intervening to shut the unit down when vibrations began to exceed preprogrammed limits. The ISS Flight Director commented that the symptoms seemed to indicate that something was blocking the spinning motor; "but we really haven't nailed down the exact root cause yet… We did conduct a test overnight where we brought the unit back up and we ran it again and collected some more data and that data is currently being reviewed by the engineers. We won't do anything else until that data is reviewed and we get better understanding of what's going on". Again, these comments can be seen to reinforce the attributes of **observability** and **revision –** the Flight Director admitted the need to gather further data and did not automatically stick with any initial hypothesis about the failure mode.

From 20th to 23rd November, the crew tried different techniques to keep the Urine Processing Assembly in operation. Each time, it worked longer than on previous tests before shutting itself down. One explanation was that thermal expansion occurred after the unit had been running for some time. This could account for the friction or blockage that led to the motor symptoms; including a speed reduction and increased current. It was also possible that the distillation assembly reached an operating frequency that caused the unit to move so that a speed sensor came in contact with the spinning centrifuge. These different hypotheses were considered as were the interactions between both possible causes, again illustrating the resilience principles of **flexibility** and **revision**, which provide key strengths in the face of uncertain failures involving complex applications. Without necessarily committing to either explanation, the ground teams identified a number of possible options:

- Plan A[1]: Delaying the return of STS-126 by one day in the hope that repairs could be made and samples could be gathered for testing back on Earth. If successful, this would support the calibration of the UPA and support the 90 day test period before full operation;

- Plan B: Abandoning the repair attempts and shipping the hardware back to Earth. The Urine Processing Assembly could be repaired and re-launched in February but this would not leave sufficient time to complete testing before it would have to support six crew members.

In evaluating the options, the Flight Director argued that "the longer we can actually perform the checkouts prior to that, the better off we are…The reason we really targeted this flight for performing the analysis, we still have some margin in case something goes wrong and we need to do any re-planning or fly up any additional equipment or consumables on the mission in February. So we do still have some room and some runway ahead of us in this case. If we wait until February, we may not get all the engineering requirements to be sure that all the systems are working as required in order to support six-person crew." Again, this provides an example of the **flexibility** and **anticipation** that characterize effective and resilient responses to uncertainty. The Flight Director did not dismiss either of the two plans that were active during the 21st November but the various teams involved in the repair missions were clear about the consequences of each option.

By the 22nd November and in spite of the 'on-off' interruptions to the UPA, sufficient material had been collected support some calibration after the return of STS-126. The objective had been to obtain a sample ratio of 30:70; processed urine to condensate. By the 22nd, they had achieved a ratio of 10:90. The previous options were also further constrained by the realization that there were no spare centrifuge units on the ground. The hardware on the station had to be repaired to enable the planned crew expansion for May 2009. An additional option was, therefore, considered:

- Plan C: only use the urine processor for short periods between cool downs. The processor was originally designed to run for up to four hours at a time. If the problem causes the system to shut down after every two hours of operation then the unit should only be run for up to 1 hour and 45 minutes before cooling.

[1] The identifiers 'plan A' etc are introduced to assist in the presentation and were not used by the mission teams.

The ISS Flight Director had to acknowledge, however, that the "numbers have not been crunched yet," to determine how many crew members could be supported with this improvised operating technique; “Folks will definitely be going off and studying that." Hence we can see that changing circumstances forced the various team members both to **react**, for instance to the lack of additional centrifuge units, and also to **revise** the initial plans using the hybrid strategy suggested above.

By the 23rd, it was still unclear whether there were interactions between thermal expansion and vibration dampening. Attention began to focus on a number of rubber washers that were used to reduce noise from the centrifuge but might also be allowing sufficient motion from the centrifuge to create harmonic effects that triggered the software alarms. The following actions were, therefore, identified:

- Plan D: to remove the thick rubber washers that provided the vibration damping from the distillation unit's rack mounting system. The unit was then bolted down into the assembly creating a ‘hard mount’ between the shelf and the rack that was intended to minimize any vibrations.

The process was then restarted and ground telemetry showed that it appeared to be working normally even though the crew reported hearing unusual noises from the centrifuge ‘as though something was off-balance’. After operating for two hours on the evening of the 23rd the crew again saw a decrease in the motor speed and a drop in the current just as had been seen on the two previous days. However, unlike previous failures the unit continued to operate. The station commander responded to the news relayed from the ground team; "That sounds dandy news… We've been watching it and actually have the PCS plot function up for the first time in my life and we saw that yeah, we saw it's still going and the current is about one point four. ... So Megan, the big picture plan is to keep processing, and that means I'll probably need to do another fill in about another hour, hour and a half?" The decision was taken to let the process continue and then attempt a refill of the system on the 24th, to which the commander responded "OK, well we have quite a collection (of urine) up here." However, the optimism was premature as the unit then shut itself down once more. This failure occurred after two hours and 52 minutes of operation.

Up to this point, the analysis of resilience during the ISS UPA problems for STS-126 has relied entirely on a taxonomy of attributes identified in previous work. Reaction has been contrasted with anticipation, opacity with observability, stiffness with flexibility and fixation with revision. However, our analysis also illustrates a further attribute of resilience. **Satisficing** refers to the manner in which operators often have to identify and implement plans that are sufficient even though they might identify a more optimal solution if they had additional time or resources. It contrasts with encysting which describes a lack of resilience when individuals and teams lose the ‘big picture’ in seeking more and more detailed solutions for components of a problem. Hollnagel (2009) identifies similar attributes of resilience in his more recent work on the ‘Efficiency-Thoroughness Trade-Off. In this case, plan C and plan D evolved as potential solutions that were intended to be sufficient to support the further operation of the UPA without necessarily rectifying all of the problems that had been experienced in operating the unit.

By the 24th, a further EVA was taking place while other members of the crew again worked on the Urine Processing Assembly. Even though the unit had failed again, there was evidence to suggest that Plan C had partially corrected the problem. The unit had been running for longer than at any other time since STS-126 had transferred the unit to the ISS. However, it was still unclear whether the root cause of the problems stemmed from thermal expansion or unexpected harmonics from the vibration of the centrifuge. With this partial information, the ground teams and the crews developed a further plan of action to get the UPA working again. The dampeners had originally been held in place by six bolts. When plan C had been implemented, only four had been put back after the rubber washers had been removed:

- Plan D-1: continue to operate without the vibration damping (as described in Plan D) but this time insert all six bolts rather than four to further tighten up the ’hard mount’.

It was also decided to push back the return of STS-126 by one day in order to determine whether samples could be provided after this modification had been tested. In preparing to fill the system for the next test, one of the astronauts found that a connection in the water processing rack's cooling system was not fully seated. Although this was not associated with the processing problems, it does indicate the way in which multiple problems can affect complex, space related systems. Focused maintenance activities can introduce additional failure modes and at the same time create opportunities to identify further problems such as the connection issue described here. This reinforces many of the previous points made by the proponents of resilience engineering – it is all too easy to ignore this successful observation and rectification of a potential coupling problem when so much attention was paid to rectifying the UPA fault. However, in the longer term, these successful monitoring behaviors proved to be essential for the

continued operation of the ELCSS after the arrival of the 6-person crews.

The Urine Processing Assembly was restarted at 20:00 on the 24th and continued to run into the early hours of 25th November 2008. Confidence grew that the changes had been effective. Houston reported that "Our regen (regenerative life support system) guys are actually smiling, which is really nice, here in the control center." A few minutes later, the crew noticed a further change in the sound of the centrifuge as well as changes in the motor current but these were short-lived; “Well, not to spoil anything, but I think up here we're feeling the appropriate words are 'yippee!'" However, the ISS flight planners were still considering bringing the distillation unit back to Earth in the shuttle middeck or in the cargo module for repairs if there were further failures. Flight controllers ran the unit for a total of five hours before restarting it after a three-hour cool-down with the intention of running the system all day. This progress helped to justify the decision to extend STS-126 by the additional day. The maintenance of multiple contingency plans even when there was evidence of success is characteristic of resilient organizations. The concurrent work of the planners and controllers illustrates **anticipation** of future problems but also **observability** as both teams monitored each others’ progress. Only in conducting these joint activities is it possible to support the **flexibility** that allows plan **revision**. The down-side is that many of the options that are considered by concurrent planning teams will never be passed to operations if we hope for the best but expect the worst.

By late on the 25th, some six litres (65 pounds) of processed urine and condensate had been obtained for chemical analysis on Earth. The decision had also been taken to keep the processing unit in orbit rather than return it for repair. Mission control reported that “we'll nurse it along the way we have been and learn from the system." However, initial plans were developed for further changes if necessary:

- Plan E: consider altering the mount by introducing additional brackets that would stiffen the structure and further reduce potential vibrations.

The processing unit continued to operate normally into the 26th November and a ‘final run’ was completed before the return of STS-126. The ISS flight director described how "it's pretty amazing to see that we've made it all the way to this point where we actually have in the plan the last water sample to be collected...The sampling plan has changed from what it was pre-flight and it's changed so it'll put the ISS program in a better posture for making a decision about their readiness for six-person crew early next year... So we were glad ... we were able to find a way to get the equipment working and come up with a plan that would accommodate all the samples that had been requested." The UPA was then successfully shut down prior to the departure of the Shuttle; which made a successful landing at 16:35 on the 30th November. The processed urine and condensate were then flown back to the Johnson Space Center in Houston to help the remaining crew calibrate the on-board analyzers.

In the weeks that followed the return of STS-126, the remaining astronauts worked on the ECLSS, installing software upgrades on the Total Organic Carbon Analyzer and performing extensive leak checks on the two Water Recovery System racks etc. However, the Urine Processing Assembly failed again and a new Distillation Assembly was scheduled for delivered by STS-119. In spite of these set-backs the ISS moved to full six-person operation in May 2009 with Expedition 20. Additional water supplies continued to be ferried by Soyuz, the remaining Shuttle missions and the ESA Jules Verne Automatic Transfer Vehicle.

5. CONCLUSIONS

There is a concern that mishap reporting systems track information about previous anomalies but do little to protect against future failures. In contrast, the term ‘resilience engineering’ has been coined to describe techniques that identify the thousands of everyday positive actions that prevent accidents from occurring. This change of perspective poses significant theoretical and pragmatic challenges. Just as it can be difficult to identify the causes of previous failures, it is equally difficult to determine what went right and why. We cannot always be sure how close our successes came to failure. Resilience engineering has not previously been applied in any sustained way to space missions. This paper has, therefore, used resilience engineering concepts to analyse the many different ways in which successive crews responded to engineering challenges on the International Space Station (ISS). The intention is not simply to promote the approach but to identify the strengths and the weaknesses of this novel perspective. A number of significant challenges remain. In particular it is not clear how resilience engineering might be integrated with more conventional approaches that are based around hazard analysis and risk assessment. Some of the attributes of resilience remain difficult to interpret in the context of complex space missions and others contradict accepted engineering practices.

We have illustrated the application of resilience engineering concepts to show how ground-based planning and operations teams cooperated with the crews of STS-126 and the ISS to install and maintain the Urine Processing Assembly within the

Environmental Control and Life Support System. This paper is, in part, a testimony to the resilient behaviours of these individuals that are often under-valued by the wider engineering community. For instance, the **anticipatory** elements of resilience engineering were illustrated by the planning processes that led to the installation of the UPA during STS-126. Rather than reacting to immediate pressures, this mission was the culmination of more than five years work to upgrade the ECLSS in anticipation both of an increase in the ISS crew size and the projected end of the Shuttle missions.

In other areas, we found that the attributes of resilience could have undermined existing engineering processes. For instance, it might be argued that the decision not to change the underlying design of the compression and distillation units illustrated a lack of **flexibility** after the initial prototype had failed. This undermined resilience and exacerbated subsequent attempts to integrate the UPA into the Water Recovery System (WRS). In contrast, we would argue that to abandon the use of low pressure, vacuum distillation would have threatened the overall mission objectives behind the ELCSS by delaying installation beyond the point at which the shuttle fleet could have been retired.

Our analysis based on this initial case study helped to identify a further attribute of resilience; **satisficing.** This occurs when teams have to choose a prompt response based on partial information in situations where they lack the time and resources to exhaustively search for an optimal solution. A key theme in our work has been to stress that resilience often depends upon actions that are determined by the context in which they take place. For instance, the crews' immediate response to the first UPA alarm correctly followed the priorities identified for tasks in Standard Operating Procedures (SOPs) – including the check for combustible products. Only once this had been ruled out, did the ground teams have the opportunity to follow the attributes of resilience engineering; to think more **flexibly** and consider the possible causes of a false alarm.

6. RECOMMENDATIONS & FURTHER WORK

This work presented in this paper represents a first step in identifying the strengths and weaknesses of resilience engineering as a framework to support safety management within human space flight. Our case study suggests that crews may have to rapidly move between predetermined SOPs, identified using conventional engineering techniques, and the more flexible approach advocated by resilience engineering, considering different causal hypotheses for false alarms. Predetermined responses are preferable in situations where safety may be threatened with limited time for decision making. However, the additional flexibility and revision advocated by resilience engineering is strongly preferable where crews have the additional time needed to conduct more considered decision making. These insights are not simply important for the ISS. Our work must be seen within the wider context of long duration human space flight. In particular, any future missions to Mars may have a 9-month cruise phase during which it can take more than twenty minutes for radio communications to be relayed from ground control. During this time the crew will have to develop and implement plans, similar to those presented here, in order to restore critical functions with a future ECLSS. This temporal dimension has not been considered within previous studies of resilience engineering but it is central to future applications in space missions.

Previous sections have characterized the tensions that exist between traditional forms of safety engineering, driven by the analysis of previous failures, and the attributes of resilience engineering, informed by the analysis of successful interactions. Ideally, common tools should support an integrated approach that helps us to gain a clearer overview of the causes of incidents and accidents. These techniques should also help engineers clearly identify the causes of success, given that it can be difficult to determine how close we came to accidents that did not occur. Until such integrated methods are developed then both safety and resilience assessments will continue to rely on the subjective expertise of individual analysts. Leveson, Hollnagel (2009) and their colleagues have suggested how this might be done through the extension of root cause analysis techniques but much work remains to be done.

Finally, questions must be raised about the scope of our work. The water recovery case study provides limited insights. Further research must be conducted to generalise the findings from this study and identify the requirements for future tools/methods. The focus in this paper has been on human space flight; however, the same techniques are arguably of even greater benefit in non-human space operations. The difficulty of establishing and maintaining remote situation awareness adds further levels of complexity as teams work to strike a balance between the need for immediate intervention and more resilient approaches that require the more flexible revision of predetermined plans.

References

European Space Agency Knowledge Engineering Office, September 2009. Last accessed February 2010, available on: http://www.esa.int/SPECIALS/Space_Engineering/SEMTH8KIWZF_0.html

W. Harwood, STS-126/ULF2 Mission Archive, CBS News/Kennedy Space, November 2008. Last accessed February 2010, available on: http://www.cbsnews.com/network/news/space/126/STS-126_Archive.html

E. Hollnagel, D.D. Woods and N. Leveson (eds.), Resilience Engineering: Concepts and Precepts, Ashgate Publishing, London, UK. 2006.

E. Hollnagel, The ETTO Principle: Why Things That Go Right Sometimes Go Wrong, Ashgate, Farnham, UK, 2009.

C.W. Johnson, A Handbook of Accident and Incident Reporting, Glasgow University Press, 2003. Available from http://www.dcs.gla.ac.uk/~johnson/book

J. Leonhardt, E. Hollnagel, L. Macchi, B. Kirwan, A White Paper on Resilience Engineering for ATM, EUROCONTROL, Brussels, Belgium, 2009.

NASA, Procedural Requirements for Mishap and Close Call Reporting, Investigating, and Record keeping. NASA Headquarters, Washington DD, USA (NPR 8621.1B), 2006.

D. Woods, Creating Foresight: How Resilience Engineering Can Transform NASA's Approach to Risky Decision Making David Woods, Testimony on The Future of NASA For Committee on Commerce, Science and Transportation, John McCain, Chair October 29, 2003.

D. Woods, Creating Foresight: Lessons for Enhancing Resilience from Columbia. B. Starbuck and M. Farjoun (eds.), Organisations at the limit: Learning from the Columbia Accident, Blackwell, Oxford, 2005.