Configuration Management as a Common Factor in Space Related Mishaps

C.W. Johnson, Ph.D.; Department of Computing Science, University of Glasgow, Scotland.
L.L. Fletcher, PhD, US Air Force Advanced Space Operations School, Colorado Springs, CO, USA.
C. Michael Holloway, NASA Langley Research Center, 100 NASA Road, Hampton, VA, USA.
C. Shea, M.Ed., Ph.D., ESR Technology Ltd, Warrington, Cheshire, UK.

Keywords: Configuration Management, Space Mishaps, Accident Analysis.

Abstract

Configuration management is essential for system safety. It helps to ensure that requirements and constraints, which are identified in previous stages of development, are preserved through subsequent modifications. It has, therefore, been recognized as a core component for standards ranging from IEC 61508 through to ISO 9000. Unfortunately, the pressures on development teams to respond to changing demands in dynamic environments and the complexity of many safety-critical systems combine to undermine these processes. This paper, therefore, uses the insights gained from three recent space-related mishaps to identify the threats to configuration management: the NOAA N-Prime fabrication incident, the loss of the X43-A and the 'hard landing' of the Genesis mission. The conclusions illustrate the importance of configuration management from conceptual and mathematical modeling through to operational deployment. They also illustrate the increasing importance of configuration management techniques in helping accident investigators identify the state of complex systems in the aftermath of adverse events. Poor configuration management not only increases the likelihood of mishaps, it also frustrates attempts to learn lessons from any failures that do occur.

Introduction

Configuration management helps to ensure that requirements and constraints, which are identified in previous stages of development, are preserved through subsequent modifications. Within this general description there are a range of more specific concerns – for example, one aspect of configuration management focuses on the maintenance of well defined interfaces between system components. Other areas focus more on security and authentication – ensuring that any changes to a system are correctly authorized. More broadly, configuration management consists of procedures and processes that are intended to ensure the consistency of a product with both functional and non-functional requirements throughout the development and operational lifecycle.

The importance of configuration management has been recognized through its incorporation within most of the main safety standards including IEC 61508 and DO-178B, as well as more general quality guidance such as the CMMI and ISO9000. It has been applied to software but also more widely to hardware and to systems as a whole, for instance through MIL-HDBK-61A(SE) which explicitly provides 'Configuration Management Guidance' for the US Department of Defense. This defines configuration management to be 'a management process for establishing and maintaining consistency of a product's performance, functional, and physical attributes with its requirements, design and operational information throughout its life' [1]. Figure 1 shows how the DoD guidance places configuration management at the heart of systems engineering.

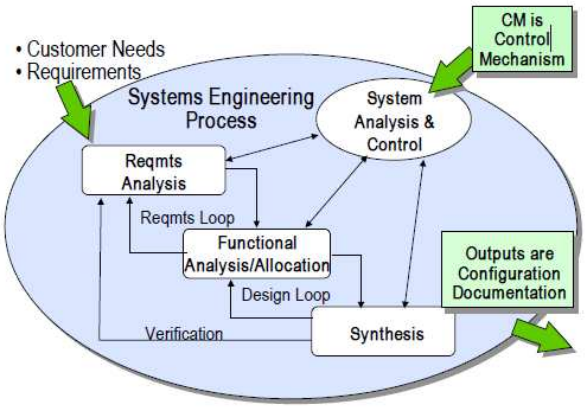

Figure 1 — Relationship between Configuration Management and Systems Engineering (Ack: MIL-HDBK-61A)

Although configuration management has been recognized as a principle component of systems engineering, it is often perceived to impose unnecessary burdens on the engineering of complex systems. It can be bureaucratic and time consuming. Often the benefits of following particular procedures are not gained by those who complete the configuration management documentation but by colleagues who must subsequently modify or maintain complex safety-critical systems. In addition, a growing number of mishaps have been caused by the sheer complexity of configuration management tasks in safety-critical applications. The following pages draw upon the lessons learned in the aftermath of three recent space-related incidents to illustrate these arguments. The insights from these failures are then used to inform an eight-stage process that is intended to reinforce the configuration management processes that are embedded within existing guidance and international standards.

NOAA N-Prime 'Turn-Over Cart' mishap

The first mishap involved the Television Infrared Observational Satellites (TIROS) National Oceanic and Atmospheric Administration (NOAA) N-Prime satellite [2]. The platform was intended to provide a polar-orbiting satellite for collecting environmental data about the Earth's atmosphere, cloud cover, radiation, atmospheric ozone, aerosol distribution, sea surface temperature etc. The mishap occurred during assembly as the satellite was being repositioned. The aim was to rotate and tilt the vehicle from a vertical to a horizontal position when it slipped from a Turn-Over Cart (TOC). There were no injuries although there was considerable damage to the system hardware. At first sight, the causes for the failure were relatively obvious; 24 bolts that secured the satellite adaptor plate to the TOC were missing. The proximate cause was, therefore, that the "operations team failed to follow procedures to properly configure the TOC, such that the 24 bolts that were needed to secure the TOC adapter plate to the TOC were not installed". However, a more sustained analysis revealed numerous operational and managerial issues that contributed to the mishap and ultimately undermined configuration management for the satellite and the cart.

Configuration and Communication between Systems Teams: The N-Prime accident stemmed from the difficulty of ensuring that teams reconfigure shared resources between complex operations. This can be traced back to a 1973 National Space Policy study when the US Office of Management and Budget identified advantages from bringing together aspects of the DOD and NOAA operational weather satellite programs. NOAA was directed to use the DOD's Defense Meteorological Satellite Program (DMSP) Block 5D spacecraft. The decision to create cost savings by bringing the two programs together led to the development of facilities in different parts of the same building. The two programs also shared ground support equipment because of the similarity of the vehicles that they used. However, each satellite required unique test configurations, hence the need to develop specific adaptors for Turn-Over Carts. Communication between the Integration and Test managers on each program was relatively informal. This can be seen as an advantage in overcoming restrictive organisational barriers, for instance by allowing each other to share equipment. In retrospect, however, these informal practices also served to undermine necessary configuration management and other SOPs that were specifically intended to safeguard operations in each program.

Dangers of Partial Reconfiguration: In the days before the accident, the Defense Meteorological Satellite Program decided to use the NOAA TOC because their own was 'red tagged' with a fault. However, work to reconfigure the NOAA TOC for the DMSP satellites was interrupted part way through the process of removing the NOAA adapter ring. The DMSP team discovered that the NOAA TOC was also red tagged. It was, therefore, easier for DMSP to clear the fault with their own TOC. This change in plan left the NOAA/TIROS adapter ring on the NOAA TOC with its 24 attachment bolts removed. No red tag was added this TOC to indicate an incomplete configuration nor were the partial changes in configuration communicated by the DMSP to the NOAA/TIROS development teams. This followed the philosophy that the TOC was not cleared for use until its configuration had been verified before each operation was started. It is clear, however, that the red-tagging of faulty equipment helped to mask the partial reconfiguration of the TOC.

Configuration Management and Complexity in Maintenance Operations: NOAA/TIROS personnel cleared the red tag on their TOC, but they did not realise that their adaptor had been partially removed by the DMSP team. This omission can be explained by the complexity of the repair to their TOC. The red-tag related to a failed jack. They did not have a direct replacement with a sufficient rating and so a lower powered jack was used. This could not be used to move the TOC with the spacecraft attached. As a result, the NOAA/TIROS teams had to reconfigure the lift and mating procedure for the satellite. On the day of the accident, the TOC was removed from the common Ante Room into the NOAA bay and prepared to support the N-Prime vehicle.

Configuration Management, Supervision and Training: The Responsible Test Engineer, a lead technician, a technician, and the Technician Supervisor were present as they began to reposition the satellite. The Responsible Test Engineer had been involved in eight previous TOC configuration operations. However, the investigation team argued that he had previously relied on a lead technician who was not working on the day of the accident to perform the TOC configuration. The Test Engineer was also amongst the least experienced of those involved in the operation to move the satellite. Hence, he may have interpreted a comment by the Technician Supervisor about the 'empty bolt holes' to refer to 44 of the 88 bolts in the payload adapter that was intentionally not installed rather than the 24 missing bolts from the reconfigured TOC. The Mishap Investigation Board concluded that the Technician Supervisor 'lacked the knowledge to recognize the problem' and that the rest of the team, 'due to complacency and channelized attention, failed to pursue the apparent warning' [2].

Configuration Management and Expertise in Oversight: The NOAA acting Integration and Test (I&T) Manager was present as an observer during the spacecraft lift, but was present as preparations were made to use the TOC. A Government Quality Assurance Representative was required during the operation. Although he had witnessed the turnover procedure 20 to 25 times, he had never witnessed the reconfiguration of the TOC. This was consistent with SOPs; since it was outside his remit for Ground Support Equipment. Instead, the company QA inspector was required to 'sign off' that the TOC was correctly configured. They arrived after the operation has begun. This deprived the operations team of critical input during the preparation phase. It also illustrates the competing demands between operational pragmatism and the care demanded by configuration management processes. The QA inspector had to decide whether to delay the operations while he checked the status of the TOC following SOPs. Alternatively, he could trust the work of his co-workers and sign-off the preparatory phase. He chose to approve the configuration – following a practice that seems to have been implicitly tolerated by management given the numerous Corrective Action Reports (CARs) that already documented 'stamping violations' [2].

Configuration Management and Degradation over Time: The Mishap Investigation Board also argued that the operations team failed to follow procedures because there was complacency with respect to spacecraft handling. This complacency exacerbated poor communication and poor coordination between the teams. They also suffered from inadequate procedures. The paperwork describing the repositioning operation was developed from the Program Directive and included a single, hand-written instruction in the Log of Operations and four steps in the Instrument procedure, TI-MHS-327820 that violated SOPs. These comments raise a host of issues in relation to configuration management of safety-critical systems. In particular, the criticism that complacency led to the erosion of procedures has parallels in many similar accidents [3]. The procedures that led to the N-Prime mishap were continually described as 'routine' even though they involved moving the satellite through angles that threatened the integrity of the platform and which created risks for those involved in the assembly processes.

Configuration Management of Configuration Management: Many other accidents also stem from a failure to apply configuration management processes to the SOPs and other documents that are intended to guide configuration management. The NASA investigators found that the operations teams involved in the N-Prime accident had to follow standard operating procedures that contained ambiguous terminology, such as 'assure', which provided insufficient guidance on the degree of care to be taken in verifying that a requirement had been met. They also identified semi-formal modifications, such as the use of red underlining to show that an operation only had to be conducted once. These practices were all identified as preconditions for the eventual mishap.

Configuration Management, Audit and Enforcement: Inadequate oversight and the failure to correct known problems contributed to the N-Prime mishap. The Responsible Test Engineer and the Integration & Test managers were seen to have violated configuration management procedures, in particular those that related to the monitoring of operations by their crews. These problems were compounded by the late notification of government inspectors, poor test documentation including configuration data etc. It was also argued that these practices stemmed from the organisation and management of the parent organisation and by relevant government agencies. This led to inadequate resources for safety and for quality assurance functions. The Mishap Investigation Board concluded that there was an "unhealthy mix of a dynamic integration and testing climate with a well-established program and routine operations" [2]. This provides a succinct summary of the ways in which changing requirements serve to undermine the processes and procedures that should safeguard the principles of configuration management. A specific example was provided by the way in which the in-house Government Quality Assurance Representative (acting on behalf of the Defense Contract Management Agency) waived a Mandatory Inspection Point during the operation. In the aftermath of the accident, it was argued that the waiver indicated the failure of external oversight to provide barriers against adverse events. Over time, complacency had reduced the effectiveness of inspection and audit to the point where it became compliance driven rather than proactive in identifying deficiencies. Even when government representatives became aware of specific deficiencies in configuration management and policy enforcement they did not report them to the NASA project sponsors.

X43A Hyper-X Research Vehicle

The N-Prime mishap shows how configuration management issues can threaten established procedures with relatively well understood technologies. The Mishap Investigation Board focused on the role that complacency can play in such circumstances. The loss of the X-43A Hyper-X Research Vehicle (HXRV) provides a complete contrast to the previous incident and demonstrates how uncertainty and complexity also affect configuration management in leading edge research projects [4]. NASA created the Hyper-X initiative to be a 'high-risk, high payoff' programme. The intention was to improve hypersonic air-breathing propulsion to the point where it could be taken from the laboratory into flight. Air breathing engines should have significantly better specific impulse while within the atmosphere than rocket engines. In order to move beyond the theoretical models, it was necessary to develop and operate a scramjet aircraft. Traditional turbojet engines use a gas turbine driven fan to further compress the air intake. This gives greater power at low speeds and increases the efficiency of the engine. However, the turbines are increasingly complex and the temperature tolerances of the turbine section limit thrust at high speed. In contrast, scramjets use a tapered intake and the forward motion of the engine itself to compress the air intake without the need for a turbine. In consequence, the engines are far simpler and have a much higher potential power to weight ratio that traditional designs. Clearly, however, the engines much reach an initial speed before the air intake is sufficiently fast for compression and combustion to occur efficiently. A scramjet requires supersonic speeds to maximise the efficiency of combustion. However, this technology offers the theoretical potential to reach Mach 24; this compares with a top speed of Mach 4 for conventional air-breathing manned vehicles. The risks associated with the project should not be underestimated. While short suborbital scramjet test flights have been successfully completed, most vehicles have not survived the test phase. In consequence, orbital scramjets have been described as 'the hardest way to reach orbit', while the proponents of the approach continue to spend millions of dollars to develop underlying technology.

The X-43A Hyper-X Research Vehicle (HXRV) was based on a hybrid design using a rocket propelled Hyper-X Launch Vehicle (HXLV) to accelerate the aircraft to the minimum speeds necessary for scramjet performance. The HXRV was attached to the launch vehicle by an adapter that provided services to the HXRV before separation. The HXLV, HXRV and the adapter were collectively known as the X-43A stack. They were flown under a B-52

aircraft until they reached the launch area. Three flights were planned for hypersonic speeds; greater than Mach 5. During the first mission, the X-43A stack was released from the B-52 at 0 seconds mission time. The HXLV solid rocket motor ignition occurred 5 seconds later and the mission proceeded as planned. Around 11.5 seconds into the flight, the X-43A stack began to experience a unexpected roll oscillation during a planned pitch-up manoeuvre. This increased until 13 seconds when the launch vehicle rudder actuator ceased to respond to commands from the autopilot [4]. This caused the stack to diverge rapidly from the planned trajectory increasing the loading on the starboard elevon. The vehicle was terminated by ground control at 49 seconds after release.

Configuration Management and the Pressures for Innovation. Although the X-43A project is radically different from the N-Prime mishap, there are some similarities. In both cases, the hardware platforms were developed from legacy applications. In the case of the N-Prime incident, the TOC adaptor had to be developed to accommodate differences between the NOAA/TIROS platforms and the common DSMP Block 5D architecture. In the case of the X-43A project, the HXLV rocket was a modified form of the Pegasus launch vehicle, stage one. As mentioned, this was intended to accelerate the stack to the speed and altitude necessary to initiate the scramjet. However, this required considerable reconfiguration of the Pegasus rocket given that the trajectory for the initial launch was at a lower altitude and a higher dynamic pressure than for more conventional applications of the Pegasus. The difficulty of modelling the impact of these changes and then correctly configuring the components of the stack were identified as the root causes of this mishap. Subsequent analysis identified that problems stemmed both from the design of the launch vehicle but also from inaccuracies in the models that had been inherited from earlier Pegasus missions, which significantly over-estimated the safety margins for critical operating parameters.

Understand the Complexity of Configuration Management. Neither the initial roll oscillation nor the problems with the rudder actuator were predicted in the pre-flight analysis. One reason for this was the need to integrate a number of very diverse models using innovative analytical techniques that informed configuration management for the X-43A launch stack. These models were intended to represent the interactions between specific components within the systems architecture; they were also intended to identify problems leading from boundary conditions, from variations in nominal data and so on. In particular, the control system was based on models that provided no means of anticipating and then responding to the roll oscillation during transonic flight. This led to the failure of the rudder actuator because other models failed to anticipate the loading that might be placed on this component during abnormal conditions. In consequence, the failure of the rudder increased deflections to the point where the integrity of the stack could not be maintained.

Configuration Management and Multiple Model Integration: Configuration management is one of the foundations of systems engineering. It helps to ensure that complex, safety-critical systems architectures meet many different functional and non-functional requirements. Increasingly, the heterogenous nature of the components within complex systems ensures that configuration management must be informed by mathematical and engineering models that are come from many different technical disciplines. The need to integrate the results from these techniques to guide the development and operation of applications often stretches projects into areas that are poorly understood or where there are considerable uncertainties about the reliability of the eventual results. In terms of the X-43A mishap, problems in fin actuation stemmed from discrepancies in modelling the electronic and mechanical fin actuator system components. This involved the integration of two different engineering approaches and, in part, prevented accurate predictions being made about the ability of the actuators to meet the control requirements being placed upon them. There were further inaccuracies in the mathematical models that were used to make predictions about the aerodynamic performance. These did not stem from configuration management in the integration of physical components as in the N-Prime satellite TOC but from problems in the integration of wind tunnel data into the mathematical abstractions. There was insufficient wind tunnel data to support the extrapolations that were fed into the aerodynamic predictions. There were also problems in configuration management of the mathematical models themselves – outer mold line changes that were associated with the thermal protection of the stack had not been incorporated into the aerodynamic abstractions. These limitations extended to insufficient analysis of the variation that might affect the parameters used to describe many system components and hence gauge the uncertainty associated with the results for aerodynamic, fin actuation and control system models [4].

Configuration Management and Multiple Modelling Failures: A common feature between the N-Prime and X-43A mishaps is that both investigation teams identified systemic causes that went deeper than the behaviour of individuals and teams on the day of the incident. One consequence of this is that the deeper problems in both programs led to

numerous different issues in the configuration management of the different safety-related applications. In other words, the investigations argued that even if the N-Prime and X-43A mishaps **had not failed** in the way that they did then it is possible that other weaknesses may have led to other incidents in the future. For instance, the difficulty in integrating multiple theoretical models and in using data from complex mathematical systems not only revealed itself in the higher level aerodynamic aspects of the HXRV project, it also led to errors in the modelling of dynamic aerodynamics and aeroservoelasticity as well as mass.

Configuration Management and the Interaction with Uncertainty: Previous sections have referred to the complexity of research projects, such as X-43A. This explains the difficulty of ensuring configuration management techniques are exploited in many safety-critical applications. At the same time, complexity also increases the need to exploit appropriate configuration management techniques in order to ensure that critical details are not missed. Even all of the errors in the mathematical models that were produced to guide the configuration and architecture of the X-34A stack cannot on their own explain why this mishap occurred. For instance, the aileron gain margin was reassessed and the revised models showed a reduction from the predicted pre-flight level of 8 dB down to less than 2 dB once the errors in the calculations had been corrected. This was significantly less than the 6 dB 'safe' gain margin. The investigation team concluded that although this came close to creating instability, the revised prediction was still stable [4]. Non-linear time history predictions had to be used to account for the behaviour of the X-34A during the incident. In particular, this had to be revised to consider parameter uncertainty. It was argued that "no single contributing factor or potential contributing factor caused this mishap. The flight mishap could only be reproduced when all of the modelling inaccuracies with uncertainty variations were incorporated in the system level linear analysis model and nonlinear simulation model" [4].

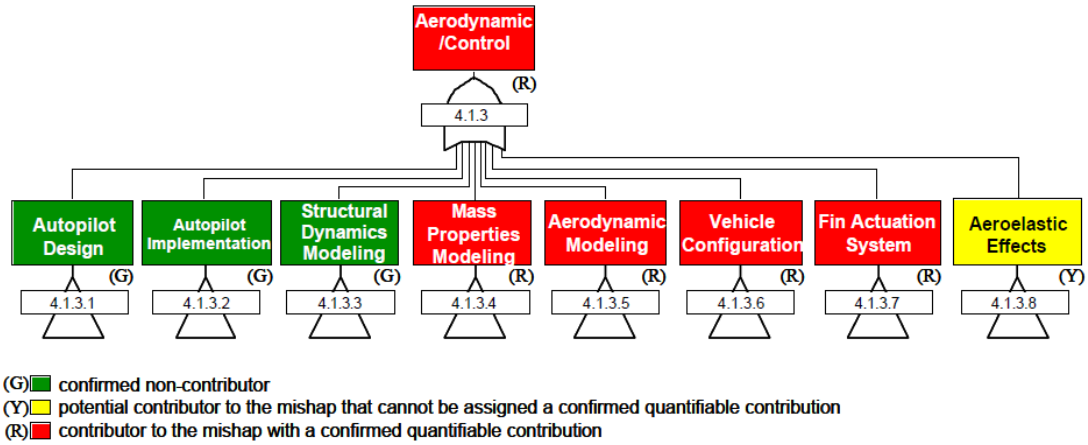

Figure 2 — Critical Fault Tree Brach (Ack: X-43A Mishap Investigation Board [4])

Configuration Management and Accident Investigation: In many accidents that stem from poor configuration management, it can be difficult and sometimes impossible to determine the state of the system before the failure occurred [3]. Inadequate record keeping and a failure to follow SOPs not only contribute to mishaps, they also frustrate attempts by accident investigators to identify what went wrong. These factors do not seem to have been significant barriers to the X-43A report. However, it is clear that significant resources of time, skill and expertise were required to piece together the complex interactions between the modelling problems and configuration of the vehicle prior to launch. The team not only had to look at the HXLV parameters but also the legacy models from the Pegasus project, keeping in mind that it flew a revised profile from previous missions using this hardware. Wind tunnel testing was also required and so configuration management principles became just as important for the investigatory team as they are for development projects. They had to ensure that the data derived from the mishap simulations could actually be used to inform their analysis of the incident. Figure 2 provides an overview of the role that configuration management issues played in this mishap. It also shows how techniques such as Fault Trees provide a documentation tool that must be controlled and managed in the investigation of a failure, just as it must also be during the risk analysis for development projects.

Genesis

The Genesis mission was launched in August 2001 [5]. Its aim was to enable scientists to study the formation of the solar system. In order to do this, an ambitious programme was developed to collect samples of solar wind and return them to Earth where they could then be analysed in greater detail to provide insights into the chemical and isotopic composition of the materials that were collected. The mishap occurred in September 2004 when the return capsule's drogue parachute failed to deploy during the Entry, Descent and Landing phase of the mission over the Utah Test and Training Range. The loss of stability that should have been provided by the drogue parachutes as the vehicle slowed to transonic speeds led to a tumbling motion. The capsule crashed into the Test Range triggering the operation of a contingency response plan as teams began to ensure that the remains of the vehicle were safe and so that recovery could begin.

The subsequent investigation focussed on the failure of G-sensors within the Aviation Units that were intended to trigger parachute deployment. A G-switch sensor is used to measure acceleration. When correctly mounted, an internal plunger compresses a spring until the gravitational force is sufficient to make an electrical contact. The intention behind the design of the Genesis Sample Return Capsule (SRC) was that g-forces would gradually increase on the vehicle due to drag from the atmosphere. At 3-G's, the plunger in the G-Switch would close a circuit and arm a sequencer. When the SRC began to slow and the gravitational force fell below 3-G's, contact would be lost as the plunger was pushed away from the contact by the internal spring. The loss of contact was intended to start the sequencer that not only triggered the various pyrotechnic events for parachute deployment but also activated a GPS transceiver and UHF beacon that were to help crews recover the vehicle. The design relied on a total of four G-switches within two duplicated avionics units. A positive firing from either of these units was sufficient to trigger the sequencer. There were also two G-switches within each avionics unit so that a signal had to be received from both to ensure agreement. This was intended to avoid premature firing of the parachute system. The signals derived from the G-switches were also filtered to ensure that the sequencers were not inadvertently triggered by transient signals generated from the buffeting in early stages of re-entry.

Consistent Configuration Management but Incorrect Design: The subsequent investigation were able to use the project's configuration management systems to demonstrate that the board assembly drawings, flight board closeout photographs, and the flight A and B side AU's were consistent and that the G-switches were installed in the manner indicated in the design drawings [5]. However it was apparent that the Genesis G-sensors were in an inverted position when they were compared with similar systems from the Stardust mission. Stardust successfully returned comet dust to earth in 2006 from a mission that was launched prior to the failure of Genesis in 1999. The inverted position of the G-sensors in Genesis made it impossible for the SRC to detect the increases in G-force as drag increased from the atmosphere during re-entry. There was insufficient force for the plunger to make the connection against the pressure from the internal spring, hence the sequencer was neither armed nor triggered, the pyrotechnics were not fired and the parachutes did not deploy. This analysis provides numerous further insights into the role of configuration management within systems safety. Firstly, the configuration management of the project seems to have been very good in that it was possible to use the relevant records to establish the precise configuration of a complex sub-system after the mishap. It was even possible to cross-reference these documents with the records of the configuration of the previous Stardust mission and hence to identify the potential anomaly. This again reiterates the importance of configuration management as a foundation for accident investigations involving increasingly complex applications. If these records had been missing or were incomplete then it would have been far harder for investigators to determine the true configuration of the SRC during Entry, Descent and Landing. Secondly, the Genesis mishap illustrates how configuration management showing consistency between development documents is of little benefit if an error is introduced early in the development process – the error will reliably be propagated into the final application. This justifies a greater focus on the interaction between requirements engineering and configuration management within the development lifecycle.

Configuration Management and Requirements Engineering: Configuration management helps to ensure that functional and non-functional requirements are addressed in a consistent manner throughout the design of a complex system. From this it follows that configuration management will help to ensure that design errors are accurately propagated into an eventual implementation. This observation makes it critical that engineers and managers conduct sufficient inspections using the resources provided within a configuration management system to identify underlying

problems that may be propagated between different stages in the development lifecycle. The Genesis mishap board make this clear when they describe how the Genesis requirements included the constraint that the SRC avionics subsystem provide functionality to deploy a drogue. This included the phrase *"descending X axial deceleration"* [5]. However, it did not describe any system or AU-level coordinate scheme to indicate the direction within this axis. Following the mishap, it was argued that the systems engineering teams assumed the requirement would be interpreted in the same way that it had been during the development of the Stardust mission, where a centrifuge had been used to test that the requirement was implemented in the appropriate components. In this case it can be argued that configuration management should have considered not only **internal consistency** with requirements for the Genesis mission but also **external consistency** with requirements that had been inherited from the Stardust programme. In retrospect it is clear that the saving from heritage software and hardware can only be achieve if there is a detailed engineering understanding of the constraints that affect those designs. The Genesis report argued that "to reach this level of understanding, for a new set of requirements and without the original designer, the heritage design must be reviewed as thoroughly as new hardware" [5].

Configuration Management and the Interaction with Systems Engineering/Project Management: There are further parallels between all three of the case studies in this paper. Common to all was the recommendation to improve discipline and enforce conformance with existing SOPs across the many different organisations that were involved in these mishaps. In all of the examples, the investigatory agencies argued that the root causes might have been addressed by increasing the level of project management support to Systems Engineering. This finding is particularly surprising given that the same observations had been made in many previous space related mishaps that had involved several of these organizations ranging from the Mars Climate Orbiter, the Mars Polar Lander, the Thermosphere Ionosphere Mesosphere Energetics and Dynamics (TIMED) and Comet Nucleus Tour (CONTOUR) missions. The Genesis investigation, therefore, argued that there were systemic weaknesses in the rigour with which fundamentally sound configuration management and other engineering processes were being applied to these missions. It was proposed that future projects should include reviews of the engineering processes as part of the normal management control gates – the effective application of these processes would become a formal requirement for continuation towards the next control gate. The report reiterated points made in the previous paragraphs – contingency management relies upon the effective application of systems engineering, however, it is not an end in itself; "Focus should not be solely on process and plans, such as the Systems Engineering and Management plan and the Configuration Management plan, but also on review of the technical products (reports, trade studies, requirements, verification results, etc.) that the Systems Engineering team has produced" [5].

Configuration Management and the Review Processes: The Genesis mishap investigation board recommended a thorough review of all "project Systems Engineering progress, plans, and processes as part of existing major milestone reviews". The intent behind this finding was to stress that configuration management is not an end in itself. In other words, it is no good maintaining accurate records to ensure conformance between stages in the development cycle if engineering teams do not **use** those documents to determine whether or not successive stages of design will actually meet high level project requirements. It can be argued that the Genesis project management teams had not ensured sufficient oversight of the contractor's activities. This reiterates comments that were made about the involvement of the in-house Government Quality Assurance Representative and other project management groups following the TOC failure during the N—Prime incident. In the case of the Genesis incident, however, the investigators argued that greater involvement by systems engineering groups in the project management team might have identified 'key process errors' during the design, test and review of the spacecraft even if they might not have directly identified the failure scenario that eventually led to the mishap [5]. Specific errors were identified in the design review processes, the verification process and the Red Team review of the SRC; red team reviews are intended to promote a challenging analysis by independent experts. These safeguards had become 'superficial and perfunctory'. They had also failed to follow the guiding principles of *test as you fly;* in other words the G-sensors should have been assessed at each stage as they progressed from design through fabrication to installation and operation.

The Interaction between Configuration Management, Resources and Test Documentation: The failure to test the G-switches using the centrifugal techniques developed for Stardust can be explained in several ways. In particular, there were considerable time pressures; the scheduled delivery of the SRC to the launch platform had slipped by several months because of design problems associated with drive motors in the Avionics Units. The pressures

created by these delays are difficult to under-estimate given the fear that a cost-capped project might be cancelled. The changes in the Avionics Unit design at the same time increased the need for and scope of a centrifugal test while at the same time eating into the limited resources that remained to finance such verification. At the same time, the engineering teams were preoccupied with the 'spoofing' that might occur during re-entry if buffeting during descent and entry created an erroneous signal from the G-switches. This may have diverted the amount of attention that might otherwise have been paid to the orientation of the components. The centrifuge procedure used on Stardust was instead replaced by a 'quick-lift' test to ensure that all of the G-switch sensors made contact – this was done by manually lifting the circuit box to ensure that the contacts were made. There was no requirement for this procedure to test orientation and alignment hence the installation problems with the G-switches were not detected. The belief that this 'quick lift' test was adequate to replace the centrifuge procedure seems to have been based on a cryptic note stating "SRC-AU 3-G test approach validated; moved to unit test; separate test not required" [5]. However, Project Management and Systems Engineering did not question the meaning of this bullet point when it was presented to them. In addition, there was no documentation of the change in verification methods for instance through a specific Change Request or Technical Memorandum. The investigation board concluded that "It remains unclear if a Change Request was required by the Configuration Management process at that time, but a Technical Memorandum was clearly appropriate" [5].

Configuration Management and Engineering Competence: Previous sections have argued that one of the greatest problems in configuration management across the space industries is to ensure the competence and experience of the individuals and teams that must implement the associated processes and procedures. Many of the individuals involved in TOC operations during the N-Prime mishap had limited experience in leading or auditing these maneuvers. The innovative nature of the X-43A programme meant that there were problems in identifying individuals with the appropriate expertise to integrate the multiple models required to predict the behavior of the stack. Similarly, by relying on the 'quick lift' test and a manual inspection against the Stardust drawings, the Genesis teams were relying on considerable insight from those conducting the verification; 'This approach could have been successful if it had been performed by an experienced Mechanical Engineer or guidance, navigation, and control; however, the drawing inspection was performed by … an Electrical Engineer who lacked the necessary mechanical experience, but apparently did not realize his limitation' [5]. The Genesis systems engineering verification processes did not consider a verifier's qualifications nor did it explicitly require confirmation from several different individuals.

Configuration Management and Contingency Planning: Previous sections have argued that a beneficial side effect of rigorous configuration management is that it becomes easier to derive 'lessons learned' from those adverse events that might occur in the future. The Genesis mishap illustrates a further critical role for configuration management in the aftermath of an incident. In particular, it demonstrated the need for the same policies and procedures that guide the documentation of development and operational requirements to also be applied to contingency planning. For Genesis, the contingency plans did not adequately document the procedures to be used during ground recovery from a 'hard landing'. The main objectives were to ensure safety during the recovery by checking for dangerous gases, such as HCN, CO, and SO2 , within the vehicle and then purging nirtrogen from the science canister. This had to be accomplished before removing thermal close-out panels to gain access to disconnect and remove the battery. There also requirements to liaise with external organisations and follow lines of communication within wider Federal contingency management processes. Unfortunately, the Genesis Project did not maintain a central resource or document containing all recovery contingency plans. Those plans that did exist were not available to everyone in the recovery teams. In consequence, there were inconsistencies between the different documents that were available to teams at different levels in the organisational structure. There were also deficiencies in the training that was conducted to support ground recovery teams following a contingency. The Mishap Investigation Board, therefore, recommended that future missions "assemble and maintain all recovery contingency documentation in a single binder with configuration-controlled copies deployed to appropriate elements of the recovery team" [5].

The Interaction between Configuration Management and Policy: Several members of the investigation teams involved in the incidents described in this paper, stated that poor configuration management not only increases the likelihood of a mishap, it also undermines attempts to learn lessons from those failures that do occur. One consequence of this greater significance for configuration management is that it requires support from the highest levels of leadership. The Genesis project formed part of NASA's 'Faster, Better, Cheaper' philosophy. This

encouraged teams to accept increased risk as a means of reducing cost. One way of doing this was for programme management teams to rely more on subcontractors for quality control and review. This increased the risk that lessons learned in previous missions might not be successfully used within the Genesis development.

Conclusions and Further Work

This paper forms part of a wider, international collaboration to look at the role that configuration management plays in incidents across both the military and civil space industries. We have identified many different ways in which inadequate configuration management contributes both to the immediate and longer term causes of failures. The ubiquitous need to document research and modeling through to requirements and fabrication, through to verification and operational deployment was not a great surprise given the prominence of configuration management in a host of hardware, software and process management standards. We were, however, more surprised to find that 'lessons learned' and accident investigation methodologies increasingly depend upon the principles and processes of configuration management. These arguments have been illustrated from three recent mishaps: the NOAA N-Prime fabrication accident, the loss of the X43-A and the 'hard landing' of the Genesis mission. A companion paper draws similar conclusions from several different military, space-related missions.

References

1. US Department of Defense, Military Handbook: Configuration Management Guidance, MIL-HDBK-61A(SE), Washington DC, 7 February 2001.
2. NASA, NOAA N-Prime Mishap Investigation Board, Final Report, Washington DC, 13th September 2004.
3. C.W. Johnson. *Failure in Safety-Critical Systems: A Handbook of Accident and Incident Reporting*, University of Glasgow Press, Glasgow, Scotland, ISBN 0-85261-784-4, 2003.
4. NASA, Report of Findings into the X-43A Mishap By the X-43A Mishap Investigation Board, Washington DC, 5th August 2003.
5. NASA, Genesis Mishap Investigation Board Report: Volume 1, Washington DC, 30th November 2005.

Biography

Chris.W. Johnson, DPhil, MA, MSc, FBCS, Ceng, CITP, Dept of Computing Science, Univ. of Glasgow, Glasgow, G12 8RZ, Scotland, UK.
Telephone +44(141)3306053, Fax +44(141)3304913, Johnson@dcs.gla.ac.uk, http://www.dcs.gla.ac.uk/~johnson

Chris Johnson is Professor of Computing Science at the University of Glasgow in Scotland. He heads a small research group devoted to improving the reporting and analysis of incidents and accidents across safety-critical domains ranging from healthcare, to the military to aviation and rail.

Louis L. Fletcher, PhD, Dean Advanced Space Operations School, Colorado Springs, CO, USA.
Telephone +1(719)593-8794 x320, Fax +1 (719) 637-9007, louis.fletcher@afspc.af.mil

Dr. Fletcher is currently the Dean of Air Force Space Command's Advanced Space Operations School. He was the first Chief of Safety for the Space Innovation & Development Center where he was responsible for verifying the safe operation and testing of ground-based radars, flight test launched Intercontinental Ballistic Missiles, on-orbit space systems and terrestrial command and control equipment.

C. Michael Holloway, NASA Langley Research Center, 100 NASA Road, Hampton, VA 23681-2199, USA.
Telephone - (757) 864-1701, facsimile - (757) 864-4234, e-mail - c.m.holloway@nasa.gov.

C. Michael Holloway is a senior research engineer at NASA Langley Research Center. His primary professional interests are system safety and accident analysis for software intensive systems. He is a member of the IEEE, the IEEE Computer Society, and the System Safety Society

Christine Shea, M.Ed., Ph.D., Prinicipal Consultant, ESR Technology Ltd, Whittle House, 410 The Quadrant, Birchwood Park, Warrington, Cheshire, UK.
Telephone: +44(1925)843 472, Fax: +44(1925)843500, christine.shea@esrtechnology.com

Christine is a principal consultant in safety and risk management with ESR Technology. Her work involves the management of risk in complex, safety-critical domains including aviation, rail, the petroleum industry and health care. Her research interests include the management and organisation of work in safety critical domains, safety culture, the development and implementation of incident reporting systems and human error.