Technology Infusion via Certification-based Analysis

Russell Knight

Planning and Execution System Engineering Group, Jet Propulsion Laboratory,
California Institute of Technology, USA
e-mail: russell.knight@jpl.nasa.gov

Abstract

Space mission architects are often challenged with knowing which investment in technology infusion will have the highest return. Certification based analysis gives architects and technologists a means to communicate the risks and advantages of infusing technologies at various points in a process. Various alternatives can be compared, and requirements based on supporting streamlining or automation can be derived and levied on candidate technologies. We present the results of applying certification-based analysis to a real space mission.

1 Introduction

Technology infusion is one of the most daunting challenges facing technologists wishing to further the use of AI for space missions. In general, researchers work in distilled environments where the effect of their technology can be understood and documented. Actual operations are far from distilled, and ad hoc solutions often outperform other contenders due to the inherent knowledge supporting their use and design. But, operations can benefit from using AI applications, e.g., Earth Observing One (EO1) [2], Mars Exploration Rovers (MER) [1], and Orbital Express (OE) [3]. As technologists, we know that our techniques can be beneficial, but as mission operations architects, we need to weigh possible benefits with cost of production, risk, and actual benefit derived in the context of our mission. This paper outlines a technique called Certification-based Analysis (CBA) that helps focus the conversation between technologists and mission architects on those capabilities of the technology at hand that can provide impact to the mission architecture. Our intent is to certify the process thereby avoiding certification of each product.

2 Certification-based Analysis (CBA)

CBA is a technique for analyzing a process and identifying potential areas of improvement. We use the process and analysis products as our lingua franca to communicate between technologists and architects. By process, we mean any of the standard representations of a production flow, one example of which can be seen in Figure 1. Here we see individual steps leading to products, which feed into other steps, until the final product is produced at the end. This sort of process is common for space mission operations, where a set of goals is reduced eventually to a fully vetted command sequence to be sent to the spacecraft. Fully vetting a product is synonymous with certification. For some types of products, this is referred to as verification and validation, and for others it is referred to as checking. Fundamentally, certification is the step in the process where we insure that our product is that intended and contains no flaws.

Candidate technologies are evaluated against a potential area of improvement using criteria such as risk, adaptation cost, adaptation time, reduction in cost, reduction in duration, reduction in risk, and maintainability. Where risk and maintainability are acceptable, and gains in either cost or duration outweigh adaptation costs, then the technology is deemed a suitable candidate. For many technologies, especially artificial intelligence technologies, certification of a technology implies the certification of the process (or process step) that the technology is used for, as compared to certifying the product (using a separate process, which, for space applications is often manual). Certifying the process, and not the product, is the key tenet of CBA.

2.1 Weighing Risks

While each are seemingly orthogonal, risk, adaptation cost, adaptation duration, and maintainability are actually tightly interlinked. Risk is often cited as the key reason that technologies (especially AI technologies) are rejected by space mission architects. Technology has been observed to be brittle, especially in light of quickly changing or off-nominal events. Steps that had required technology introduced risk in that they could not be adapted quickly enough by their designers and implementers to keep up with such situations. Thus, one key measurement of technology is the ease of adaptation and how much knowledge is required for the adaptation. The better the adaptation fits the knowledge of the domain expert, the lower the risk, cost, and duration of adaptation. For example, automated planning languages can be difficult for Power Resource Officers (PROs: power engineers responsible for power systems onboard the International Space Station) to understand, but graphical diagrams of the power systems are not. Fewer mistakes are made when the information is presented to experts in the form they expect, and adaptation costs are reduced [4]. Thus expenditures that reduce adaptation cost by allowing more direct input from domain experts can greatly reduce risk due to uncertainty. This entire discussion of allowing more direct insight and control of the underlying knowledge driving the adaptation of the technology speaks directly to maintainability. In the most extreme case, the adaptation component of the technology is entirely maintainable by domain experts without intercession by technologists.

Of course, maintainability issues such as operating system drift (changes in OS that prompt changes in code), migration, and patching to either the technology or those technologies upon which the technology depends need to be addressed.

To arrive at a metric which is usable for comparison, it would make sense to reduce all costs and durations to monetary expenditures. This is necessary as not all work is done by those of equal pay grade, and not all work requires the same facilities. But in practice, duration is often the preferred metric for comparison.

2.2 Weighing Advantages

Advantages that technologies bring to the table can be substantial. The duration to produce a product can be reduced, the workforce and facilities required for the product can be reduced, and the overall quality of the product can be increased (e.g., fewer entry errors). Fundamentally, these advantages are derived from the automation and streamlining of the process

2.3 Automation

We refer to automation as making a step in the process automatic and migrating the certification steps of the product off-line to be a certification of the process step itself. This usually results in reducing the duration of the step and reducing the cost in terms of work force and facilities, with a low-recurring adaptation and certification cost for the process step. Traditionally, AI technology has been infused as automation.

2.4 Streamlining

We refer to streamlining as changing the structure of the process to avoid high-likelihood or high-cost certification failures. In terms of "bang for the buck", streamlining is typically a high payoff modification, and much of our analysis is bent toward discovering streamlining opportunities. Streamlining often depends on a combination of automation and process restructuring.

Thus, the certification-based analysis process is as follows:

1) Capture the process in a form that represents cost (duration, dollars, etc.) as well as likelihood of failure of certification
2) Identify certifications that can be automated (certifications moved off-line thereby certifying the process, not the product)
3) Identify certifications that can be streamlined (process restructuring and automation)
4) Select from 2 and 3 those actions that fit in our time and budget and do not exceed risk and maintenance requirements, and implement

To make this clear, we look at an example from the Orbital Express mission.

3 Examples from Orbital Express

DARPA's Orbital Express mission demonstrated on-orbit servicing of spacecraft. Two spacecraft were flown: Boeing's ASTRO spacecraft, whose role was that of a doctor, and Ball Aerospace's NextSAT spacecraft, whose role was that of a patient. Experiments included rendezvous and capture, fluid propellant transfer, and on-orbit repair. Most of the planning for the mission was performed by the Boeing team, who also serviced requests from Ball Aerospace. The team was broken up into two units, the Rendezvous Planners who concerned themselves primarily with computing the locations and visibilities of the spacecraft, and the Scenario Resource Planners (SRPs), who were concerned with assignment of communications windows, monitoring of resources, and sending commands to the ASTRO spacecraft.

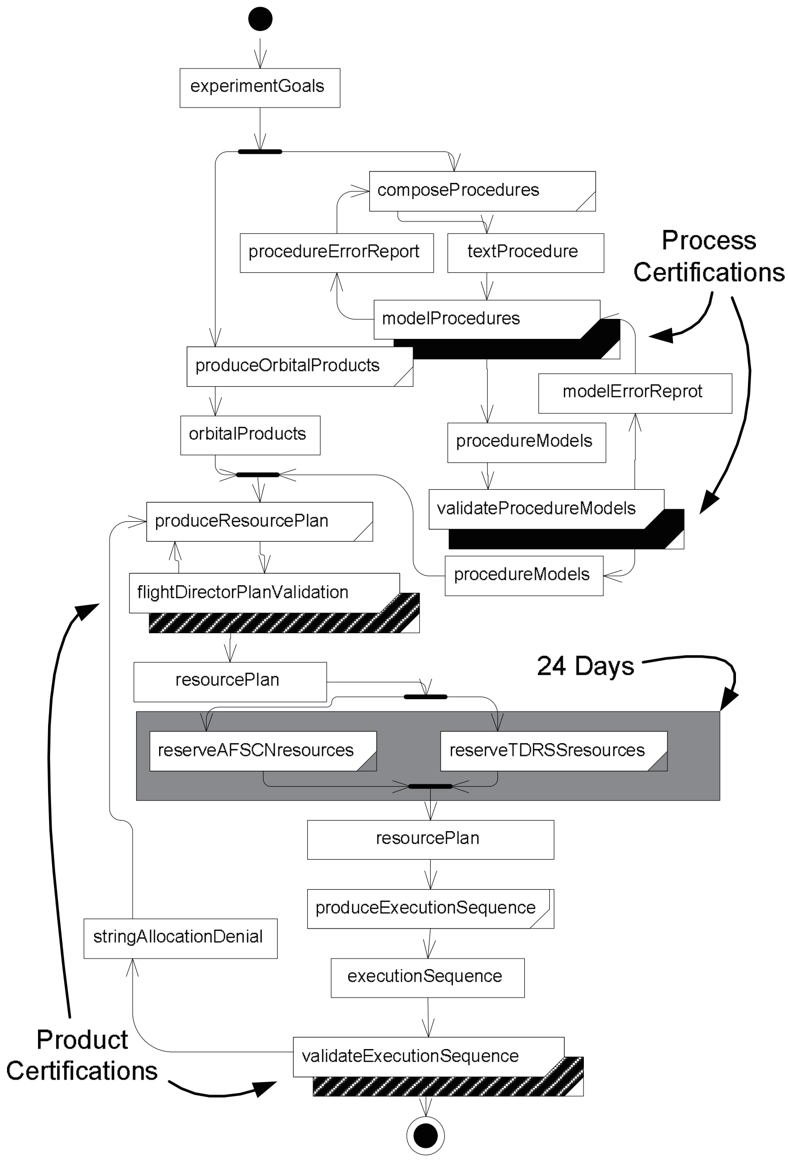

Figure 1. OE initial process

Figure 1 shows the production process used to produce sequences for operating the ASTRO spacecraft. As we can see, certification failure at validateExecutionSequence causes a huge delay in production. If this failure were exceedingly unlikely (i.e. never expected to occur) then we could ignore this possible failed certification. It is important to note that many processes have certification failure as a possibility but they are so improbable as to motivate nearly no spending on their obviation. But, in the case of xxx, this failure occurred quite often. The problem lies in this: the RSC (the facility in Albuquerque that issued commands to the spacecraft) can only allocate a certain number of resources to the Orbital Express mission. Specifically, communications resources that required fast switchovers were very limited. But, our standard approach to operating the spacecraft required fast switchovers so that we could maintain continuous coverage over the United States using our high bandwidth communications (Air Force Satellite Communications Network: AFSCN). Sometimes we could get an exception to use more resources, but that was a facility manager decision, and it couldn't be made until a resource analysis we provided, well after the resources from the various communications networks were reserved.

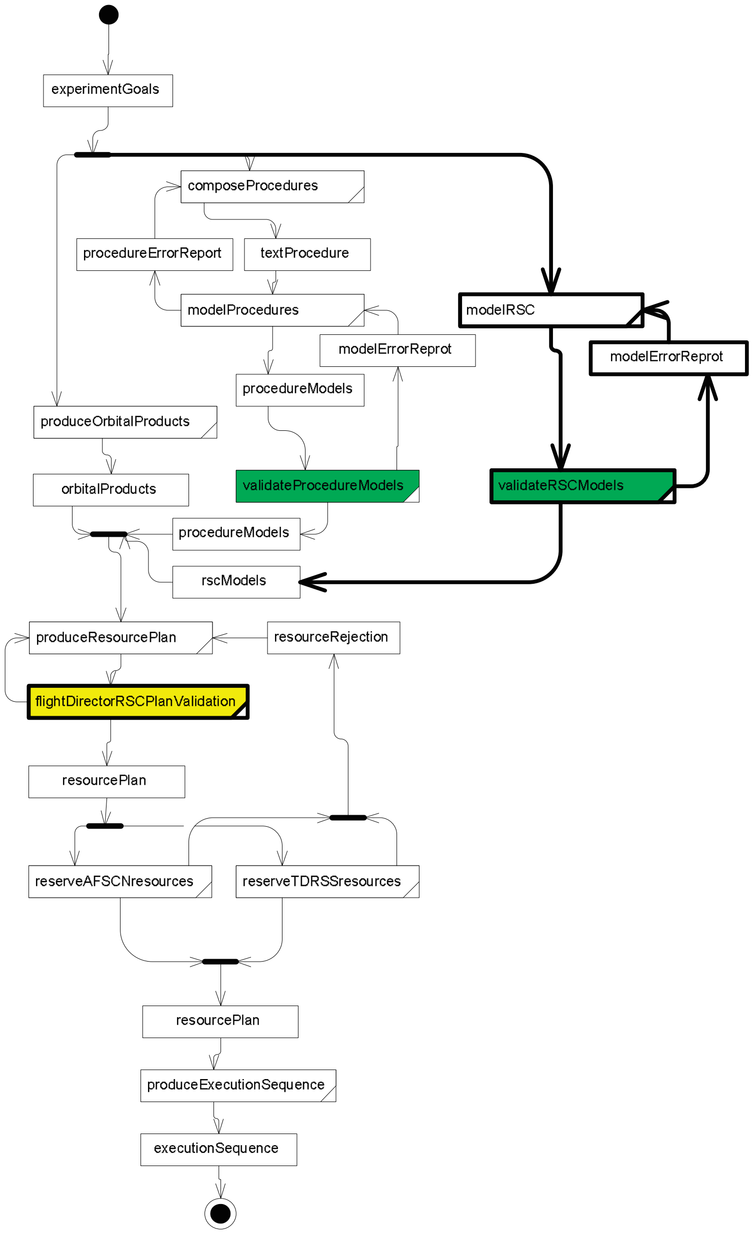

Figure 2. Streamlined RSC certification

We used a two-prong approach to streamline this process. First, we modeled all of the RSC resources in our automated planning tool, moving part of the certification off-line. Now, any product produced in xxx would be guaranteed to pass subsequent RSC manager inspection. But, we still needed to accommodate exceptions where we ask the RSC manager for more resources than are nominally allocated to us. So we include the RSC manager in flightDirectorPlanValidation, resulting in a new process illustrated in Figure 2. Although we still require several weeks to generate a product, we never fail once product generation has cleared flightDirectorRSCPlanValidation.

A second example of streamlining the Orbital Express mission goes to streamlining the communications reservations. As illustrated in our previous example, responsiveness to fundamental changes in the mission results in weeks of delays (for each change). Given that operations experienced off-nominal execution nearly every week; this would result in more than doubling the mission duration. Again, we took a two-pronged approach.

First, we discovered that we didn't need to reserve TDRSS time at all. Our automated system was already flexible enough to re-align any proposed scenario according to whatever TDRSS passes were available. TDRSS reports unused time via the TUT (TDRSS unused time). We implemented a TUT parser and were able to used the "scraps" left behind by other missions in near real-time to accommodate our needs for TDRSS.

Second, we discovered that we didn't need to reserve specific AFSCN passes as long as we had an AFSCN policy where we had a certain number of passes reserved for any given day. As it turned out, we already had requirements to realign AFSCN passes according to changes in availability and orbital drift, thus we could align any scenario to any reasonable profile of AFSCN coverage.

Combining these two gave us our required savings in terms of mission duration reduction.

A final example of streamlining the Orbital Express mission provides as an example of how to handle the "nightmare" scenario for automation: knowledge capture and modeling become integrated into the inline process. For OE, all scenarios and operations were in the form of procedures. These procedures would be painstakingly vetted by hand and be passed to the technologists, who would integrate them into the automated planning system. We discovered during operations that instead of being "carved in stone", these procedures were exceedingly dynamic. Often, procedures would change within hours of delivery of a final sequence that executed the procedure. Sometimes these changes were simple, but often they were fundamental. Instead of making the process shorter, the automation was now at risk of making the process much longer, as illustrated in xxx. To overcome this, we implemented a parser that could read the natural language in the procedure and convert it to the planning language of the automated planning system. In general, this would be a herculean effort, but in the case of OE, it was fairly straightforward. This was due to the formal and limited nature of the language in the procedures. We were not parsing free-form English, but instead a conscripted text that mostly described parameter values, timing, and continuation criteria. The certification of the parser was straightforward. The key requirements were 1) the parser should always output a "broken" model when the parser encountered something that it did not understand, and 2) the parser should never indicate it understands something when it doesn't. In practice, the parser rarely failed and where it failed required a single line of adaptation, instead of the literally thousands of lines of adaptation making up the procedure. This added in general ten to fifteen minutes to the inline process.

4 Future work

Continuing our work on automation and streamlining, we are focusing on Automating and Streamlining ISS Mission Operations (ASIMO). The first component of this work focuses on stowage management and we hope to report results soon.

5 Conclusions

In conclusion, it is not easy to know which investment in technology infusion will have the highest return. Certification based analysis gives architects and technologists a means to communicate the risks and advantages of infusing technologies at various points in a process. Various alternatives can be compared, and requirements based on supporting streamlining or automation can be derived and levied on candidate technologies.

Acknowledgement

We would like to acknowledge those who have contributed to the ideas and given insights that lead to this work. Colin Helms was the ground systems lead for OE, and our conversations with him are what eventually led to a pursuit of formalizing the ideas used in architecting the OE mission. Andrew Mishkin can be attributed with coining the term streamlining as being the modification of a process structure (although streamlining is used in slightly different contexts in the literature).

This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1] M. Ai-Change, J. Bresina, L. Charest, J. Hsu, A. Jónsson, R. Kanefsy, P. Maldegue, P. Morris, K.

Rajan, J. Yglesias, “MAPGEN: Mixed intitive activity planning for the Mars Exploratory Rover mission,” in: Proceedings of Demonstration Systems Track, ICAPS'03, 2003.

[2] S. Chien, R. Sherwood, D. Tran, R. Castano, B. Cichy, A. Davies, G. Rabideau, N. Tang, M. Burl, D. Mandl, S. Frye, J. Hengemihle, J. Agostino, R. Bote, B. Trout, S. Shulman, S. Ungar, J. Van Gaasbeck, D. Boyer, M. Griffin, H. Burke, R. Greeley, T. Doggett, K. Williams, V. Baker, J. Dohm, “Autonomous Science on the EO-1 Mission,” International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2003). Nara, Japan. May 2003.

[3] R. Knight, “Automated Planning and Scheduling for Orbital Express,” International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS 2008). Pasadena, California, USA. February 2008.

[4] R. Knight, S. Schaffer, and B. Clement. “Power Planning in the International Space Station Domain.” 6th International Workshop on Planning and Scheduling for Space. Pasadena, California, USA. July 19th – 21st, 2009.