

Examining Customer-Side Engineering Practices for Robotic Systems, CSER 2008

DeWitt T. Latimer IV¹

University of Southern California
941 W. 37th Place, SAL 300, Los Angeles, CA 90089-0781
dlatimer@ieee.org

¹The views expressed in this report are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

Abstract

This reports on the results of a survey conducted on engineers involved in the acquisition of robotic systems from academic, government, and industrial sectors. Respondents indicated their perception of the importance and effectiveness of various types of engineering practices that were employed by these customers in support of their most recent robotic system acquisition. The survey focused on engineering practices from the areas of requirements development, technical management, verification, and validation. Examination of how the opinions and effectiveness of the engineers change over successive experiences give us an indication of what classes of engineering methods are currently being addressed or ignored during the education of the engineers. Ultimately, by examining the responses, we can conclude some areas of engineering that engineers still lack routinely effective methods to help them acquire the right robot to fulfill the organization's need.

Introduction

Successful acquisition of a robot requires many different skills: project management, fiscal planning, legal/contracting, and engineering. However, studies of sufficiency of engineering practice to meet acquisition challenges have not been accomplished. By examining the relatively small acquisition domain of robotic systems, this paper seeks to expose areas in which the engineering discipline needs further elaboration to be effective.

According to (Meyers 2001) acquisition is "the set of activities performed to procure, develop, and maintain a system". The Software Engineering Institute's (SEI) Capability Maturity Model for Acquisition Organizations (CMMI Product Team 2007) has a significant discussion of the engineering and technical tasks that are part of an organization's capability to acquire products and services. Even the US Department of Defense, in its directive for acquisition practice (DoDD 5000.1), outlines extensive engineering goals that acquisition programs need to accomplish. However, all these engineering tasks are presented at the generic system level and do not enumerate specific techniques that engineers would have to undertake to support the acquisition projects. Specialization of the techniques is typically domain (aerospace, automotive, marine, et. al.) specific, and as such we must consider how to engineer our systems of interest, in this case, a robot.

A Robot is a mechanism which physically interact with the world, senses the world, and express some level of autonomy. In this way, a robot is considered an integrated hardware/software mechanism which performs some operations without being explicitly commanded, is physically separate from its human operator, and physically interacts with the world. One example robot would be Northrop Grumman's Global Hawk air vehicle. The Global Hawk air vehicle, while typically remotely piloted, does have enough autonomy to maintain safety of flight during communications interruptions with the operator.

However, a robot alone typically does not satisfy a client's needed for a capability. For this the whole Robotic System must be considered. A robotic system includes the robot, support equipment for the robot (chargers, programming environment, command/control instrumentation, etc.), training materials, repair equipment, or any other items and personnel needed for the robot to provide a capability to a client. In the case of the Global Hawk, the total robotic system would include the air vehicle, ground control unit, spares and maintenance facilities, the training materials for users, and communications infrastructure that enables the system to operate.

Engineers that aid in the acquisition of robotic systems must understand what methods of analysis, specification, and production are appropriate and their impact to cost, schedule, and performance/fitness.

The purpose of this survey was to provide evidence to support the notion that engineering methods aid in the successful acquisition of a robotic system. This is done by rejecting the null-hypothesis that effective engineering does not impact the success of an acquisition. This evidence can be noted by deriving a relationship between completeness of engineering practices and the success of acquiring a robotic system. Further, with the initial relationship of engineering and acquisition, indicators of locations of incomplete methodologies for engineering are given by mismatched importance and completeness as well as the lack of relationship between expected engineering disciplines and project outcomes.

Related Work

The importance of engineering in acquisition has anecdotal support. (Marciniak 1980) attributes several activities to the acquirer/buyer, such as reconciliation of independent verification and validation reports, requirements-cost trade-off decisions, and risk management. However, (Marciniak 1980) mostly noted these as special topics for project management. (Glaseman 1982) noted the need for dedicated technical professionals, specifically for software, to be present in acquisition offices. The lack of embedded expertise, Glaseman argued, prevented acquisition programs from effectively evaluating information coming from prime contract developers, verification and validation authorities, and external think thank analyses.

Recent work in mapping the nature of acquisition tasks have culminated in the Capability Maturity Model (Integration) for Acquisition Organizations (CMMI-ACQ) V1.2 (The CMMI Product Team 2007). Of specific interest here are four process groups that are technically oriented: acquisition requirements development (ARD), acquisition technical management (ATM), acquisition verification (AVER), and acquisition validation (AVAL). ARD covers practices related to the establishment of operational needs, technical requirements, contractual provisions, performing various requirements analyses, and allocating requirements. These ARD activities are distinct from the requirements management functions (such functions seek to establish baselines, ensure traceability, and are essential for the engineering of complex systems). ATM addresses technical exchanges between the acquirer and the developer as well as

the management of external interfaces. AVER and AVAL cover the base practices for the verification and validation of acquisition level processes and products.

(Elm 2007) brought the notion of effectiveness in systems engineering to the forefront by conducting a survey of 64 respondents to determine what correlations exist between CMMI process groups and project success criteria. Of note in this survey, individual correlations may have been based on smaller samples, as not all respondents answered all questions and some responses may have been dropped as outliers. This study was largely based on the version of the CMMI for developers. Due to the survey only asking a few questions about acquisition capabilities to the developers, this survey was unable to establish a clear relationship between acquirer capability and project performance.

Surveys of various sizes have been used to elicit practice of professionals in the field. (Verner 2005) used 42 respondents for determining which software project management practices lead to success. (Lehane 2005) used 43 respondents to determine significant factors in software systems acceptance. (Wojcicki 2006) used 35 respondents to determine the current state-of-practice in verification and validation practices for concurrent software programming. And finally, (Surakka 2007) used only 23 data points to determine which skills were important to software developers in a smaller European country. Viewing engineering as a due-diligence activity, we need only show what a "reasonable" or average engineer would do when encountering similar issues. This view indicates that state-of-practice performance should be able to be determined by only examining a small population, as has been seen in the cited surveys.

Data Collection

Survey Instrument

The survey was conducted in an anonymous fashion via a web service, SurveyMonkey.com. The survey was organized into five (5) web pages: introduction and instruction, respondent background and general experience questions, robotic acquisition project questions, acquisition engineering practices, and survey feedback.

The survey was designed to elicit opinions and perception of completeness in performing engineering analysis tasks by engineers and others involved in the acquisition of robotic systems. Please note that this survey is not an assessment of the respondent's engineering knowledge or skill, nor an assessment of any individual's job performance.

Introduction and Instruction: This section contained no questions for the respondent, but did log into the service provider that a survey had been started. The instructions provided the purpose of the survey, survey instruction (navigation and submission), a definition of acquisition and an example of how and why a robot might be acquired by a university, and the organization of the remaining sections

Respondent Background and General Experience Questions: This section consisted of eleven (11) questions designed to provide control information about the field the respondent works in, their total experience with acquisition and robotics, and education levels. The section also asked for the work location, to track geographic dispersion of respondents.

Robotic Acquisition Project Questions: This section consisted of fourteen (14) questions that asked about the most recently completed acquisition project (or the current project, if none have been completed yet). This "most recent" designation was chosen in an effort to combat

selection bias of only reporting on very successful or failed projects. The first question asked how long ago the most recent project was completed. The next set of questions addressed scale issues: per unit cost, and number of units acquired. Next, quality and technical issues were addressed: anticipated lifespan in the robot's intended environment, human or safety critical designation, commercial-off-the-shelf designation, lifecycle strategy employed, and autonomy level. The final set of questions elicited the respondent's opinion of the overall engineering effort performed in addition to project success criteria of meeting requirements, budget, schedule, and success in terms of fitness for the desired purpose.

Acquisition Engineering Practice Questions: This section consisted of 54 questions, of which 50 were paired. The first question in a pair was about the opinion of how important a type of engineering activity was for the acquisition described above. The second question was about the perception of the completeness of the efforts undertaken, in terms of helping that acquisition. Respondents were instructed that there was no expectation of an explicit or implied correlation between questions; it is acceptable to indicate an activity was important, yet not have been completely effective, or to indicate an activity was unimportant but was successfully completed (in this case, it may have been successful even if only a small amount of effort was expended). Respondents were also instructed that using the entire range for either importance or completeness was not necessary. Further, respondents were instructed to respond based on their perceptions and not the formality of the work. It was possible for activities to have been important and significant effort was expended, but no formal reports or documents were generated. The questions in this section were mostly drawn from the specific practices of the CMMI-ACQ V1.2 (CMMI Product Team 2007), with the addition of two questions that asked about external verification and validation respectively and the omission of a question about peer review methods. The question on peer review was deleted as being redundant with the more general verification and validation practice questions when the survey was evaluated by independent reviewers. Finally, each question allowed for comments and each of the four process groups asked if further practices would be needed.

Note, for all questions, respondents had a 5 point scale for completeness or importance and the option to not answer or indicate that information for that question is not available for any reason.

Survey Feedback: This section asked only one question, if any other engineering practices should have been covered by the survey.

Invitations

The survey invitation was sent to 39 individuals. 36 of the 39 individuals were anticipated respondents. The 3 additional individuals, and 13 of the individuals were asked to propagate the survey within their organizations. Invited groups included: University of Southern California Robotics graduate students and alumni, Stanford University Faculty, Georgia Tech Faculty, Carnegie Mellon Faculty and graduate students, US Department of Commerce researchers, US Department of Defense robotics program managers and engineers, and 4 different commercial (non-defense related).

From this invitation pool, 30 surveys were initiated and 21 were completed, giving a response rate of 53.8%.

Survey Responses

Initial Outlier Detection

Two respondents indicated "No Response" for all engineering practice completeness and importance criteria. Since such a respondent would not be matched in any statistical analysis, these two respondents were dropped from the pool.

Several respondents indicated "No Response" for a number of engineering practices or project success criteria. All these respondents are kept in the pool, but statistical analysis will only be based on the subset of the pool that responded to all factors of interest. This is noted later as the number of cases (or "n") that supports that statistical value.

On examining the responses based on different types of populations, one professional (non-student) respondent indicated low (1-2) completeness and importance for all engineering practices, where all other engineers had higher average responses (3-5). This engineer was reporting for a government, defense-related project, of which the responses were even further non-characteristic. As this respondent was over two standard deviations outside of any other potential matching population, this respondent is a potential outlier and will not be included in the analysis. This gives a total qualified responding pool of 18 respondents. The remainder of the analysis will focus on these 18 respondents.

Given the nature of similar journal articles, especially (Surakka 2007) which used only 23 respondents, and that the relatively new field of robotics is likely considerably smaller than software development as a whole, this should constitute a representative sample from which to attempt to draw conclusions.

Responding Population and Trends

The survey contained largely respondents from California (UK=1, PA=2, MA=1, MD=1, NC=1, VI=1, CA=14). Of this number, 5 of the 14 California respondents were students (2 additional students responded from other areas). This geographic dispersion is similar to the geographic dispersion of the survey invitations. Overall, 7 students and 11 professional (non-students) responded to the survey.

The differences between students and professionals (those with t-test scores under 0.05) included:

- unit cost (students averaged \$5-\$10,000, where professionals were in the \$100,000's per unit)
- acquisition experience (professionals averaging 3 times as much experience)
- robot lifespan (professionals acquire robots for over 1,000 hours of use, student median response was for 25-199 hours)
- designation as safety-critical (professionals acquiring with safety critical concerns)
- level of desired autonomy (students desiring slightly more robotic autonomy)
- opinion of overall engineering effort (professionals reporting strong and complete efforts, students noting fragmentary efforts)
- designation as commercial-off-the-shelf (students tend to buy off-the-shelf robots)
- acquisition strategy (students likely to use a "big-bang" acquisition, professionals were more diverse in their choice)
- acquisition schedule performance (professionals tended to deliver later than students)

There was no statistical difference (t-test scores over 0.05) between students and professionals in terms of:

- years experience with robotic systems
- highest level of education achieved
- age of project
- quantity of robots to be acquired
- if the robot met its requirements
- if the robot was suitable for its purpose.

For engineering practices, students were only statistically different from professionals on the average reported value for the completeness of acquisition technical management practices and on the importance of acquisition verification practices.

For correlations, over all respondents, the overall impression of engineering completeness corresponded more directly to the importances indicated for the specific engineering practices than for the completeness responses of the same individual practices (correlations of 0.55, 0.74, 0.58, and 0.49 to importance verses 0.06, 0.47, 0.24, and 0.09 for completeness of requirements, technical management, validation and verification respectively).

On project success variables, we note that although being fit for purpose and meeting requirements may appear correlated (0.45), the correlation may be random (t-test=0.29, a t-test at or below 0.05 is required to reject the null hypothesis that the variables are unrelated). This may be due to the disconnect between meeting requirements and fitness for purpose with regards to schedule performance. There was a slight negative correlation (correlation -0.27; t-test 0.00) between schedule performance and being fit for purpose, which indicates a slight trend to being more fit for purpose if the robot was delivered late. Further, schedule performance had a slight (0.15) correlation with meeting requirements (t-test=0.00), which indicates a slight trend to meeting requirements when systems are delivered on schedule (no projects in the response pool were delivered early).

The following responses had too little variation to be considered for analysis: if a robot was successfully acquired (16/18 indicated yes or was in progress), if the robot was in budget (16/19 indicated yes or refused to answer).

Outcome Factor Analysis

To determine which variables most correlated with project performance (schedule, requirements satisfaction, and suitability), linear regressions were performed. The process of examining candidate regressions was to attempt initially to use as many engineering and project factors as possible. Then sub-models were considered by deleting or adding parameters to the various base models by examining Mallows's statistic (Cook 1999) at each step. During this addition and subtraction process, the r-squared (amount of variability in the response explained by the model) and p-value (probability that the model will be rejected by finding more data) are considered. The goal is a regression function with a high r-squared, but a p-value below 0.05.

In advance of performing regression analysis, each response variable was centered to its respective average and normalized. This was done so that each response variable has an average of 0 and a standard deviation of 1, which is normal practice to reduce numerical instability due to underlying statistical tool assumptions (most of which assume a mean of zero and a standard deviation of 1) (Cook 1999).

Predicting Schedule Performance: The best linear regression for predicting schedule performance included a positive relation to the completeness of the ARD practices, negative

relation to the importance of AVER, negative relation to experience with robotic systems, and a negative relation to the level of autonomy desired (R-Squared = 0.58, p-value = 0.02, n=17). The only case not used did not provide data about schedule performance. Observe, as previously noted, two of the factors (importance of AVER and level of autonomy) are strongly correlated with student or professional acquisitions. Models that included discriminators for student projects were considered, but rejected (such models had worse F-test scores and significantly worse p-values). Hence, although those factors are correlated, that correlation does not appear to be significant here.

Predicting Meeting Requirements: To truly maximize the predictive power, we consider only projects either fully meeting requirements or those that had the most problems meeting requirements ("partially met specifications, many desired functions available"). Only three projects reported in the middle range ("Partially met specifications, all high-priority desired functions available). Excluding these three middle points (which have normalized response values of 0, so their removal doesn't impact the probability distribution) may aid us in analysis, because the difference in the response levels may not be clear (the value was potentially too qualitative).

The best linear regression for predicting satisfaction of requirements is a positive relation to the completeness of ARD practices, a positive relation to the importance of ATM, and a negative relation to the importance of AVAL (R-Squared = 0.74, p-value = 0.003, n = 14). In addition to the three excluded case, one case did not provide data about meeting requirements.

Predicting Suitability: One respondent, a professional, was determined to be an outlier via a standard outlier test on many of the candidate regressions (outlier test was significant and 4 times other responses in the finally selected regression).

The best linear regression for predicting if a robot would be suitable for its purpose was a positive relation to the completeness of AVER practices, a positive relation to the importance of AVAL, a negative relation to the importance of AVER, and the designation as safety-critical (meaning that being designated was more likely to results in a suitable robot) (R-Squared = 0.80, p-value = 0.016, n=15). Two additional cases were not used, one for not providing information about fitness for purpose and one for not providing information about the completeness of AVER practices.

Conclusions and Future Work

Care must be taken when drawing conclusions from this data. Respondents to the survey were self-selecting and anonymous. This means that respondents may have an opt-in selection bias depending on their prior project experiences. A priori, which type of bias this may take is not known. Further, responses were self-reported and not validated. This means that respondents had to provide no objective proof that their answers were correct. While some respondents may have had experience with CMMI-like questions, based on the few comments received on the practices, at least some of the respondents were not familiar with the formal process language. At best, this means that the responses may be more qualitative, rather than quantitative. At worst, this means that respondents could have made up responses (however, this is unlikely given the personal relationship of the author to most of the respondents). Although not much can be done for the self-selection bias, the qualitative nature of the responses is partially handled by the normalization of the responses and looking at the data trends rather than quantitative correlations.

For schedule performance, the correlation with requirements development and the level of autonomy is not surprising. Likely projects that have more complete requirements can have their scope adequately estimated and more realistic schedules, hence a tendency for those projects to be delivered on-time. Also, higher levels of autonomy would indicate more technologically risky projects, and hence have the potential risk for slipping schedule. The unexpected result was the negative correlation with verification activities and robotics experience. Ideally, the more experienced engineer would be more likely to deliver on schedule, however, this does not play out in this example. There are numerous potential reasons for this, part of which may be the fact that more experienced engineers tend to be placed on higher risk projects. The importance of the verification activities being a negative influence may represent additional management or stakeholder pressures on the engineers to "build the solution right".

Examining requirements performance, the expected result of the completeness of requirements development activities is present. Although the completeness of technical management didn't rate inclusion in the regressions, the importance of technical management may have ensured other (unmeasured) quantities such as access to stakeholders and sufficiency of management support for negotiating interfaces was present where needed. The negative relation to validation importance likely indicates a gap between requirements and validated needs.

Indeed, this plays out in the suitability performance. For mature domains such as automotive and construction, we would hope to see increasing levels of completeness of requirements development or validation activities yielding more suitable systems. In this survey however, different amounts of completeness of these did not explain the different levels of satisfaction. In this case, completeness of verification activities was the most important engineering task. In addition, maintaining the importance of validation activities was important to ensure success, but increased emphasis on verification was a negative influence on suitability. This second importance finding is notable, because as completeness of the verification task rises, a corresponding rise of the importance may be harmful to projects (possibly due to over-emphasis). Finally, safety-critical systems tended to be more suitable than other systems. The additional engineering for safety-critical systems did not show up in the individual engineering factors. In this case, there may be practices (such as certification and accreditation for operation) that influence and ensure the ability of a system to be suitable for its environment that are not captured in the base CMMI-ACQ practices.

Future work will feature examination of the success factors and their indicators in the context of engineering cases to give interpretive value to these observations.

References

- The Defense Acquisition System*, Department of Defense Directive (DoDD) 5000.1, dated 12 May 2003.
- CMMI Product Team (2007). "CMMI for Acquisition, Version 1.2: Improving Processes for Acquiring Better Products and Services", Carnegie Mellon University, Software Engineering Institute Technical Report, CMU/SEI-2007-TR-017, Pittsburgh, PA, USA, November 2007.
- Cook, R. D. and Weisberg, S., *Applied Regression Including Computing and Graphics*, Wiley Interscience, New York, 1999.
- Elm, J., et. al., "A Survey of Systems Engineering Effectiveness - Initial Results", Carnegie Mellon University, Software Engineering Institute Special Report, CMU/SEI-2007-SR-014, Pittsburgh, PA, USA, November 2007.

- Glaseman, S., *Comparative Studies in Software Acquisition*, LexingtonBooks, Lexington, MA, 1982.
- Honour, E., "Understanding the Value of Systems Engineering", <http://www.hcode.com/seroi/documents/ValueSE-INCOSE04.pdf>, accessed 2008-02-01, published 2004.
- Honour, E, "Design of Experiments as Applied to Systems Engineering Return on Investment," *Proceedings of the Conference on Systems Engineering Research*, Hoboken, NJ, 2007.
- Lehane, P. and Huf, S., "Towards understanding system acceptance: the development of an assessment instrument and workpractice". *Proceedings of the 19th conference of the computer-human interaction special interest group (CHISIG) of Australia on Computer-human interaction: citizens online: considerations for today and the future*. (Canberra, Australia, 2005), Computer-Human Interaction Special Interest Group (CHISIG) of Australia, pp 1-9.
- Marciniak, J., Reifer, D., *Software Acquisition Management: Managing the Acquisition of Custom Software Systems*, Wiley Series in Industrial Software Engineering Practice, New York, NY, 1990.
- Meyers, B. C. and O. Patricia . *Managing software acquisition: open systems and COTS products*, Addison-Wesley Longman Ltd., 2001.
- Surakka, S., "What Subjects and Skills are Important for Software Developers?", *Communications of the ACM*, Volume 50, Issue 1, pp. 73-78.
- Verner, Cerpa, "Australian software development: what software project management practices lead to success?", *Australian Software Engineering Conference*, (29 March - 1 April 2005), pp 70-77.
- Wilson, M.; Neal, H., "Diminishing returns of engineering effort in telerobotic systems". *IEEE Transactions on Systems, Man and Cybernetics, Part A*, Volume 31, Issue 5, pp 459 - 465, September 2001
- Wojcicki, M. A. and P. Strooper. A state-of-practice questionnaire on verification and validation for concurrent programs. *Proceeding of the 2006 workshop on Parallel and distributed systems: testing and debugging*. Portland, Maine, USA, pp 1-10.

Biography

Captain DeWitt Latimer IV, USAF earned his PhD in Computer Science at the University of Southern California. He is advised by Prof Barry Boehm and Prof Gaurav Sukhatme. His research focuses on investigating the nature of acquiring autonomous robotic systems. He earned his MS degrees in Robotics (2001) and Civil Engineering (2002) at Carnegie Mellon University. He is a senior member of the IEEE and ACM and a member of INCOSE, ASCE, and AFCEA and was awarded the CSDP credentials from the Computer Society.