15th ICCRTS

The Evolution of C2

Paper ID 093

Topic 5 Experimentation and analysis

Alternative COA Selection Methodologies: The Quantum Command and Control Theory

Steven L. Forsythe

POC: Steven L. Forsythe
Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Road
Laurel MD 20723-6099
Phone: 937-912-5601
Email: Steven.Forsythe@jhuapl.edu

The Quantum Command and Control Theory

Abstract

This paper presents results of a Johns Hopkins University Applied Physics Laboratory (JHU/APL) effort on the fundamental theory of C2. A standard method for Course of Action (COA) selection is the Weighted Numerical Comparison (WNC) method. There are numerous issues with the WNC method. One of these issues is that the expected results are frequently used to select a course of action (COA). The Quantum C2 Theory, like the physics based quantum theory, treats outcomes not as point estimates but rather as probability distributions. Understanding and employing these probability distributions can aid commanders in shaping and selecting courses of actions that are the best tradeoff between mission success and risk. Several alternatives to WNC are presented which mitigate some of the issues with WNC and provide improved support the commander's COA selection.

Introduction

Many military activities have an underlying theoretic body of knowledge. Search Theory, for example, provides a basis for algorithms and situations involving searching, such as locating an enemy sub or trying to find a downed pilot. A theoretical foundation serves many useful purposes by enabling analysts to discuss and analyze the problem in a mathematical way. Competing methods can be evaluated based on the theory to identify the relative merits of each. Command and Control (C2) is a critical military activity that has to date not been adequately described by an underlying computational theory. One of the more successful models for C2 is the Find, Fix, Track, Target, Engage, and Assess (F2T2EA) model. The limitation of this model is that it is not general enough: it assumes that there is a target to be engaged. This paper presents a more generalized C2 model that focuses on the course of action (COA) selection by the commander.

Current methods for COA selection involve constructing a matrix with criteria to evaluate each COA against. Figure 1 shows a typical COA decision matrix. Courses of Action for this example are Conventional Ballistic Missile (CBM), B-2 (bomber aircraft), and Tomahawk Land Attack Missile (TLAM). The criteria for the evaluation are listed on the left. Each COA is scored against each criteria and then the score is multiplied by a weight to get the weighted results. The weighted results are then summed and the highest sum is presumed to be the best COA. This approach has several limitations. What are the best criteria? How are appropriate weights chosen? Is expected value really the best statistic to use?

COA Selection Decision Matrix

Criteria	Weight	CBM	B-2	TLAM	Weighted Results		
					CBM	B-2	TLAM
Pd	20	0.6	0.9	0.8	20	60	40
Execution Time	20	36	128	300	60	40	20
Collateral Damage	20	Low	High	Med	60	20	40
Fratracide Risk	20	Low	Med	Med	60	40	40
Economy of Force (Risk to our Forces)	20	Low	Med	Low	60	40	60
Anti-Access (Ability to Defeat)	20	High	High	Med	60	60	40
Threats (against this option)	20	Low	Med	Med	60	40	40
Total					380	300	280

Recalculate

Figure 1. Example COA Selection Matrix

Note that the assessments of the COAs in Figure 1 are point estimates. These are the expected values but there is no insight into the extremes. In addition, Figure 1 assumes that the results are all independent. In reality if the attack misses the target, the risk of collateral damage goes way up and the probability of kill goes way down. A revised method is needed that takes into account these issues.

Background

Military models of C2 take on many forms, each with particular strengths and weaknesses. With the revolution in communications fueled by the internet, new C2 paradigms have been proposed (e.g. net-centric) which propose a significantly different perspective on the organization and function of facilities. Today's military missions typically require the participation of allies and non-government organizations (NGOs). This complex environment makes it challenging to choose the best COA. A model of C2 should be helpful in developing a methodology for COA selection.

Numerous conceptual models of C2 exist, such as John Boyd's OODA Loop (Observe, Orient, Decide, Act). This model was developed to explain fighter aircraft engagements but has been applied to a wide range of situations. Similar models include MAPE (Monitor, Access, Plan, Execute), SDA (Sense, Decide, Act) and MAAPPER

(Monitor, Access, Analyze, Predict, Plan, Execute, Report). There is also the Lawson C2 model as well as the Enterprise Theory of C2. [3]

A representation of the OODA Loop is shown in Figure 2.

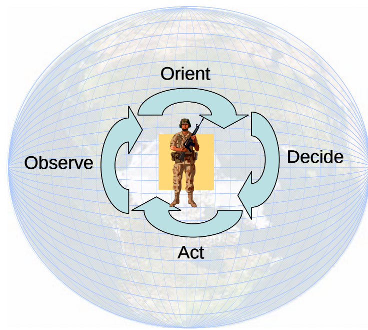

Figure 2. Representation of the OODA Loop

Other approaches such as Builder's Command Concepts focus on the art of command, suggesting that the essential communications up and down the chain of command can (and should) be limited to disseminating, verifying, or modifying command concepts. The ideal command concept is one that is so prescient, sound, and fully conveyed to subordinates that it would allow the commander to leave the battlefield before the battle commences, with no adverse effect upon the outcome. [3]

The C2 Problem

The Department of Defense defines Command and Control as: "The exercise of authority and direction by a properly designated commander over assigned and attached forces in the accomplishment of the mission. Command and control functions are performed through an arrangement of personnel, equipment, communications, facilities, and procedures employed by a commander in planning, directing, coordinating, and controlling forces and operations in the accomplishment of the mission." [1] For the purposes of this paper we will adopt a slightly broader definition of C2 which encompasses both military and civil organizations. We shall define C2 as an allocation of resources by a leader over a set of opportunities to achieve a set of objectives while managing risks. The solution to the C2 problem is a COA and the commander's challenge is to choose the best COA.

The C2 problem, then, addresses:

1) how to define the set of opportunities,
2) how to best allocate resources among the opportunities, and
3) how best to carry out those decisions.

In order to analyze how a leader allocates resources over a set of opportunities, we must consider a plethora of different factors. These factors relate to:

1) desired outcomes: the objectives we are trying to achieve and,
2) unwanted outcomes: those effects that we are trying to avoid or minimize.

Usually, there is a tradeoff to be made between achieving the desired outcomes and avoiding unwanted outcomes. The classic example is to neutralize a military facility but not to harm the hospital nearby. With today's precision weapons, airstrikes have an excellent chance of destroying a target with minimal collateral damage, unlike in WWII, when bombing campaigns caused massive collateral damage because there were no other options. Numerous issues relate to the commander's willingness to trade a risk of negative effects in order to achieve a greater likelihood of mission success. This tradeoff between mission success and risk is the key to C2 and the art of command. The science of C2 is to use modern capabilities to inform the commander of the best set of options that span the mission success vs. risk tradespace, so that the commander can apply his training and experience in selecting the best option from a set of good alternatives.

Figure 3 illustrates the key functions of the C2 process. In this figure we define the boundary of the C2 process as the information interface with the rest of the world. Information flows in and the C2 system evaluates the information, selects opportunities to pursue, assigns and schedules resources against those opportunities and then generates orders and other communications to execute the plan.

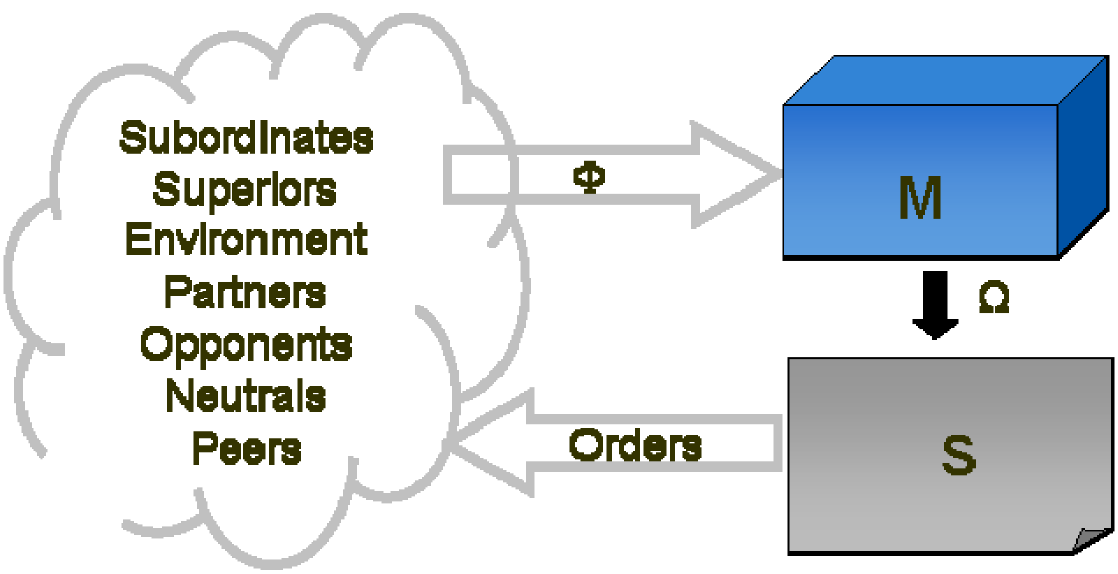

Figure 3. The C2 Problem

In Figure 3:

- The cloud represents those things outside the C2 system
- Φ, the set of information received by the commander (and staff)
- M, “Machine” that takes Φ as an input and produces Ω.
- Ω, set of opportunities that resources may be assigned against.
- S, “Machine” that takes Ω as input and produces orders.
- Orders, set of communication that indicates the Course of Action (COA) selected.

COA Selection

One of the reasons military decision-making is so challenging is that there is always a tradeoff between possible unwanted outcomes (casualties and collateral damage) and achieving the desired outcomes of the mission objectives. It is clearly the commander’s prerogative to make this risk vs. reward tradeoff in choosing the best COA. Assuming the commander doesn’t quantify ahead of time the explicit tradeoff (very unlikely) then the C2 problem becomes a multiple objective optimization problem.

In situations where there are multiple criteria to optimize, there isn’t one optimal solution; there is an efficient frontier of solutions. These solutions maximize the chance of success for a given amount of risk. Other solutions, which do not lie on the efficient frontier, are said to be dominated solutions as they are strictly inferior to one or more other solutions.

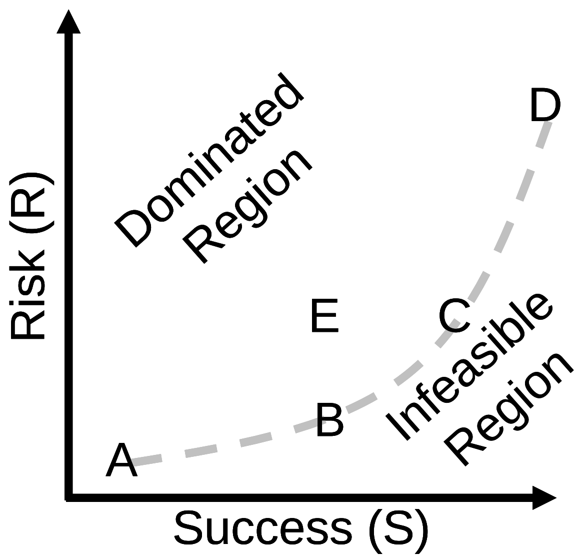

Figure 4. The Efficient Frontier

In Figure 4, A, B, C, and D represent solutions on the efficient frontier. Each represents the maximum chance of success for a given level of risk. Solution E is dominated by solutions B and C. For the same level of success, B offers lower risk than E, while C offers greater rewards for the same level of risk as E.

How does one measure risk and reward for military situations? For each of these measures a utility (or value) function is created, that assigns each solution a value based on the probability of negative outcomes or positive outcomes happening. Casualties and loss of assets are clearly negative outcomes that a commander would like to minimize, while positive outcomes are the accomplishment of various mission objectives.

Factors for Desired Outcomes and Unwanted Outcomes

Figure 5 presents an example of the factors that could be used to calculate a COA's rating for desired outcomes and unwanted outcomes [6] [8]. The mission objectives define the desired outcomes and are usually stated in the commander's intent. Also, rules of engagement and constraints on operations are also defined by the commander and define the unwanted outcomes that the commander wants to minimize while maximizing the desired outcomes.

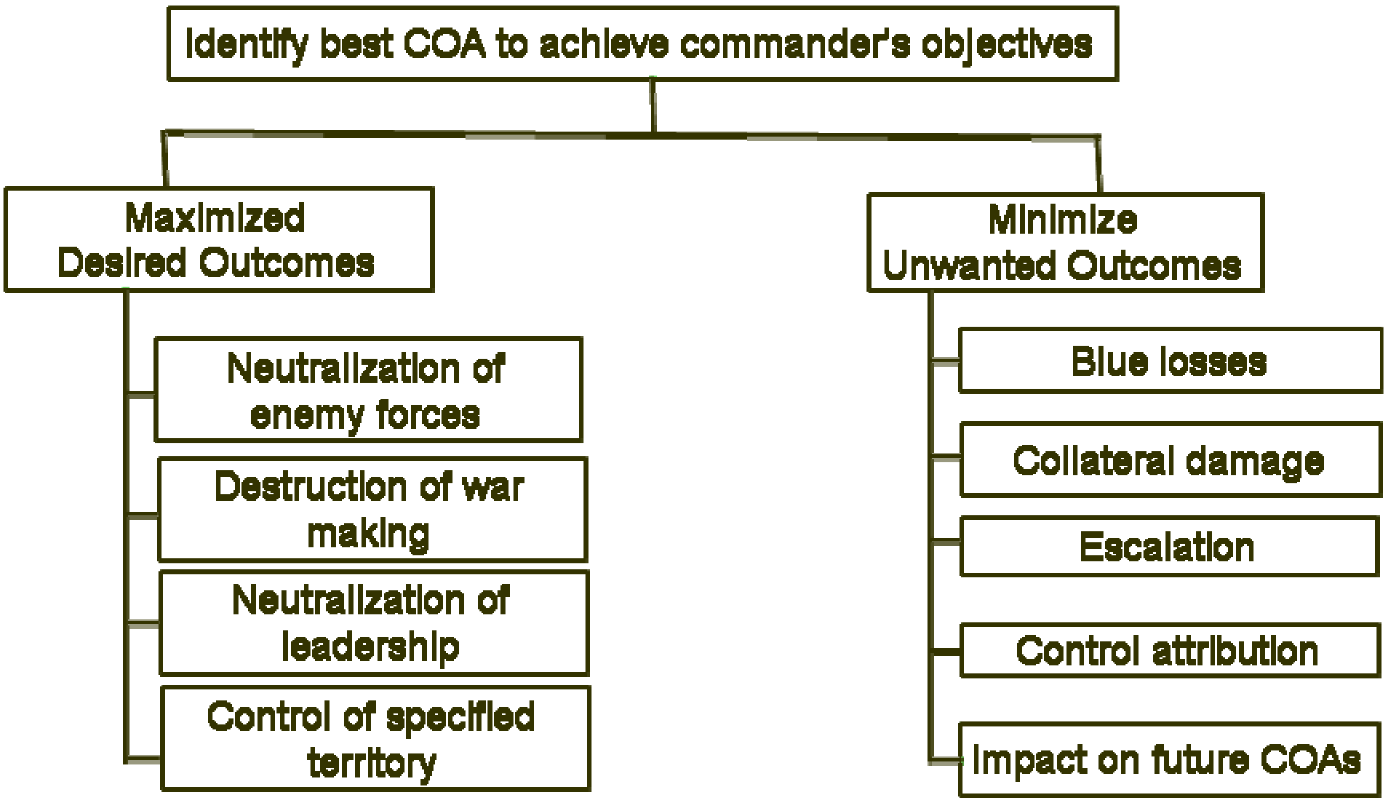

Figure 5. Sub-objectives for Desired Outcomes and Unwanted Outcomes

Figure 6 illustrates the two dimensional COA outcome space. Ideally, a commander would like to accomplish the desired outcomes with few losses, low collateral damage, etc. Therefore, the lower right quadrant is preferred since most or all of the desired outcomes are achieved with few, if any, unwanted outcomes. The upper left quadrant represents the worse possible outcomes: failing to accomplish the mission yet suffering significant unwanted outcomes. The upper right represents a victory but at a high cost. This is the very definition of a Pyrrhic victory. Another example of choosing a COA that is likely to fall into this quadrant would be the urban bombing campaigns of WWII. With the weapons of the day, there was no way to destroy the war making capabilities of the enemy without significant collateral damage. An example of the lower left quadrant is where there is reduced effectiveness but also low risk. An example might be an air campaign prior to ground combat in Desert Storm. Bombing was generally done from high altitude so as to minimize attrition even though this reduced the effectiveness of any given sortie. The mission requirements were such that moderate effectiveness at low risk was preferred to higher effectiveness at medium risk.

Figure 6. Two-Dimensional COA Outcome Space

Any given COA has some probability of being in any of these quadrants. Figure 7 shows a probability distribution mapped to the two dimensional COA outcome space.

Figure 7. Probability Distribution in the Two-Dimensional COA Outcome Space

Notice that the distribution need not be unimodal. Figure 7 shows a bimodal distribution. The most likely result of this strike is to achieve high target damage and low collateral damage. The results in the upper left indicate that there is also some likelihood of achieving only a few of the desired outcomes while suffering significant unwanted outcomes.

In the case of attacking a single target, where the desired outcome is the percentage of the target and the unwanted outcome is collateral damage, the two dimensional probability outcome space might look like figure 8.

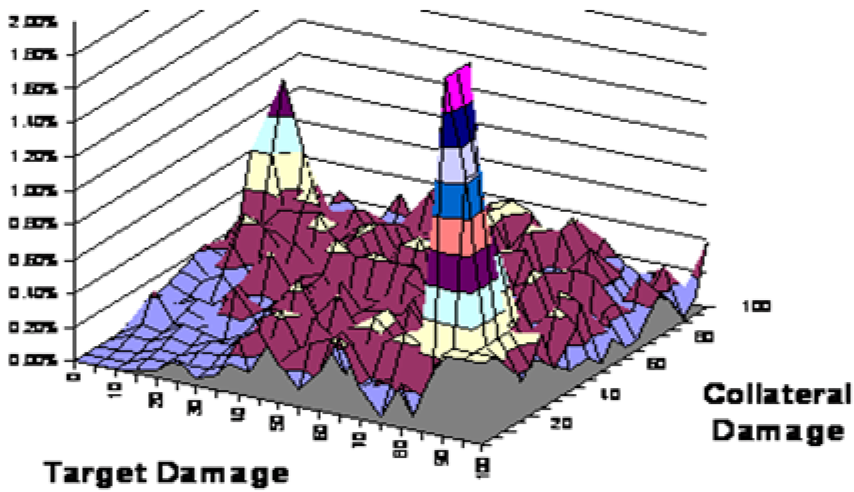

Figure 8. Example Probability Distribution in the Two-Dimensional COA Outcome Space

Combat simulations often have bimodal distributions where some key early event causes the results to shift from one combatant to the other. It is this type of unusual distribution that makes treating the COA as a probability distribution so important. Just estimating the mean of these values does not convey enough information.

How to Evaluate COAs

There are several ways that the probability distribution can be used in COA selection. Depending on the commander's comfort with quantitative methods, different methods are available ranging from simple (but still very useful) to more complex. This approach allows a commander to choose the most appropriate approach for a given mission thus maximizing the use of the science of C2 in support of the commander's art of C2.

A commander specifies information such as the maximum risk rules of engagement and minimum conditions for success in the commander's intent. This information can be combined with a probability distribution such as the one shown in Figure 9 to define the acceptable region. One measure of COA goodness is the probability that the COA will have a satisfactory outcome. In many ways this is the key question that drives a COA selection.

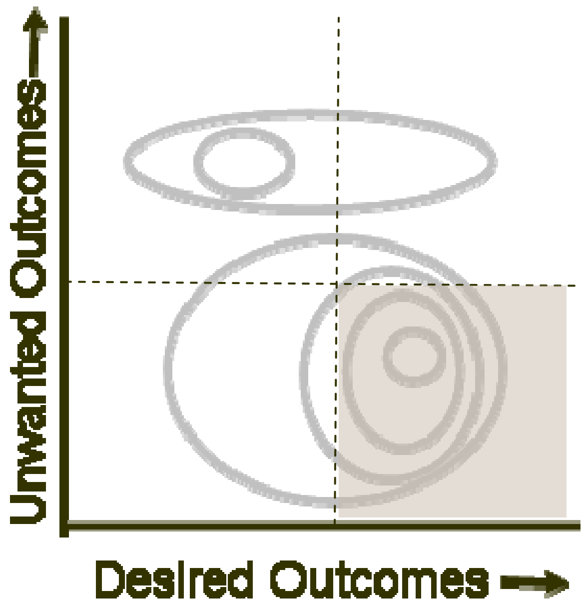

Figure 9. Method One for Defining the Acceptable Outcome Space

The shaded region in Figure 9 is the acceptable outcome space. The probability that this COA will result in an acceptable set of outcomes is equal to the percentage of the distribution that lies within the shaded region. Other COAs would have different distributions and so each COA could be evaluated based on the probability of achieving a set of acceptable outcomes.

One limitation to this method is that a commander might be willing to have higher unwanted outcomes if the desired outcomes are achieved. A second method begins by asking the commander to refine his statements on risk. What is the maximum unwanted outcomes that are acceptable given minimum acceptable desired outcomes? What is the maximum unwanted outcomes that are acceptable given total success in all desired outcomes?

If a commander has numerous sub-objectives for desired outcomes and unwanted outcomes and doesn't want to use a utility function to combine them into the two dimensional decision space, an alternative method would be to define a n-dimensional acceptable region where there are n sub-objectives. Each COA would have a probability of achieving a result that satisfied all n of the sub-objectives. The advantage of this approach is that it avoids the subjective weights of the utility functions. The disadvantages of this approach are that it treats all outcomes in a binary fashion: all acceptable results are equally good and all unacceptable results are equally bad. This n-dimensional method also doesn't provide the commander with a risk vs. mission success tradeoff.

Figure 10 shows how the shape of the desired outcome space is modified based on this information. Figure 10 clearly shows that a tradeoff exists between the desired outcomes and the unwanted outcomes.

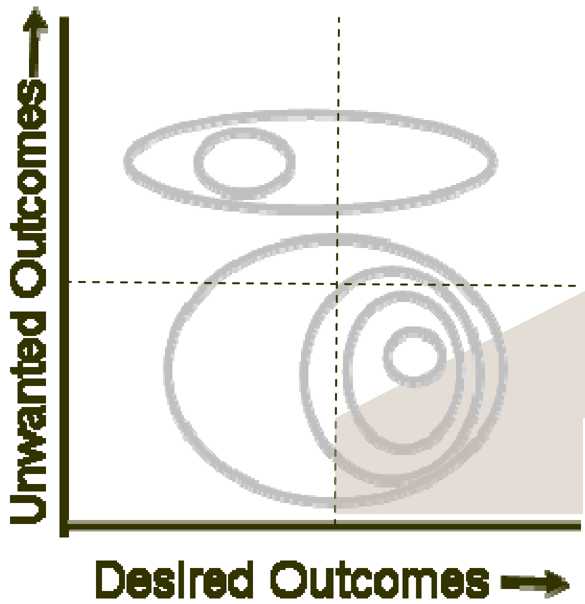

Figure 10. Method Two for Defining the Acceptable Outcome Space

Military decision makers often use green/yellow/red connotations to quickly convey information. If the acceptable outcome space is defined to be "green" then a portion of the outcome space could also be defined as "red." This red space could be thought of as disastrous results or other such descriptions based on the commander's intent. A "Hail Mary" play in football, for example, has a high probability of getting a great result but it also has a high probability of being intercepted. It tends to be an all-or-nothing choice. Other strategies might have a lower chance of getting an acceptable outcome (getting a first down) but they avoid the disastrous results of a turnover. Using a Green/Red space such as displayed in Figure 11 would allow a commander to compare the likelihood of achieving acceptable results vs. the likelihood of the results being very, very bad.

An n-dimensional approach can be taken with this method as well. The n sub-objectives would have a range of "green" values which are acceptable and each objective could have a "red" range of values that the commander wants to avoid at all costs. In this method, each COA would be evaluated to calculate the probability of it achieving and outcome where all the sub-objectives are green as well as the probability that any of the sub-objectives were red. The COAs could then be plotted on an X-Y with the X values being the probability of green outcome and the Y value being the probability of a red outcome. Once again, the commander could then make a risk vs. mission success tradeoff.

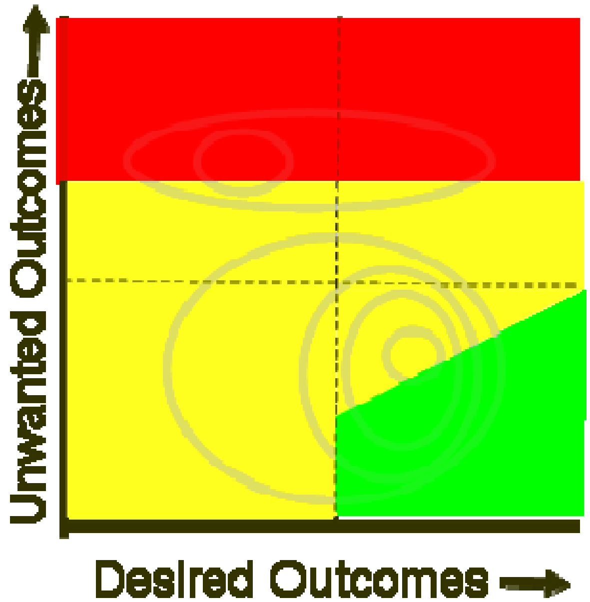

Figure 11. Method Three for Defining the Acceptable Outcome Space

Figure 12 shows how the results of using the third method can be displayed for a commander. The COAs are plotted on an XY plot with the X axis being the probability that the outcome will have a green result for all criteria. The y axis is the probability that the outcome will have a result were one or more criteria are red. Once again the commander is given a tradeoff between increasing the likelihood of mission success (all green) vs. risk (any red results).

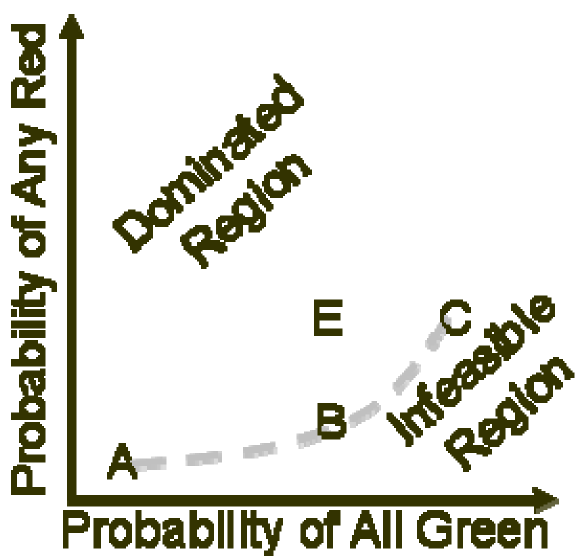

Figure 12. Methodology Number Three Mission Success vs. Risk Tradeoff

A fourth and final method can be used if the commander or his staff can define a set of relative values for each part of the outcome space. Such a set of relative values (called utility values) represents how much the commander values achieving the specified outcomes. Figure 13 shows a notional set of values with the best possible outcome being valued 1000 and the worst possible valued at zero. A simple polynomial surface can be fit to a set of values to allow extrapolation between the points. The probability distribution of each COA is then scored against the utility values of the commander.

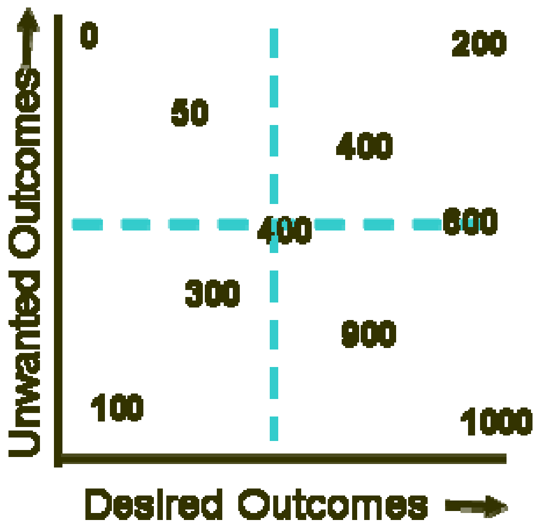

Figure 13. Commander's Utility Values

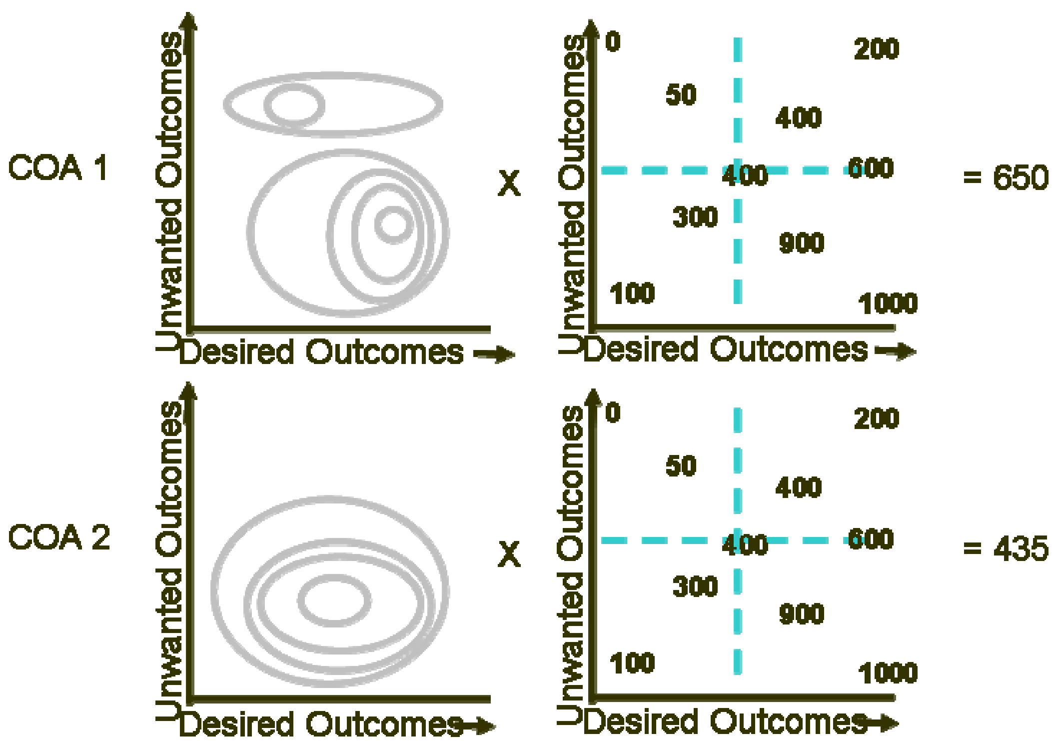

Figure 14. Calculating Expected Utility Value for Two COAs

Figure 14 shows how the expected utility value for each COA is calculated. By applying values to each possible outcome, this method uniquely values each possible outcome. This Utility Theory method captures the commander's concern for unwanted outcomes as well as the importance of achieving the mission. This method has the most stringent requirements for data: utility values that the commander has confidence in. By using the utility data, the method allows for the calculation of the highest expected utility. This method therefore takes into account the entire probability distribution and the commander's utility for outcomes associated with all parts of the outcome space.

Conclusion

This paper provides an approach to C2 COA selection which is an alternative to the current COA selection methodology. The Quantum C2 theory treats the COA as a probability distribution that needs to be evaluated to determine the best COA based on the commander's guidance. Four methodologies are discussed each of which requires further testing and evaluation but may be an improvement to current COA evaluation methodologies. The simplest method only requires knowledge of the minimum acceptable desired outcomes and unwanted outcomes. The most sophisticated approach uses Utility Theory to calculate the COA with the maximum expected utility. Thus the Quantum C2 approach can accommodate different commanders with different approaches to COA evaluation.

References

[1] Defense Technical Information Center, *DoD Dictionary of Military and Associated Terms*. Joint Publication 1-02. http://www.dtic.mil/doctrine/jel/doddict/data/c/01085.html (June 2007)

[2] Builder, Carl H., Bankes, Steven C., and Nordin, Richard, *Command Concepts, A Theory Derived from the Practice of Command and Control*, RAND 1999

[3] North, Paul D., Forsythe, Steven L., "A Process Decomposition Approach for Evaluating Command and Control (C2) Functional Performance" 2006 Command and Control Research and Technology Symposium

[4] Alberts, David S., Hayes, Richard E., *Understanding Command and Control*, CCRP Publication, 2006

[5] Alberts, David S., Hayes, Richard E., *Planning: Complex Endeavors*, CCRP Publication, 2007

[6] Loerch, Andrew G, Rainey, Larry B., *Methods for Conducting Military Operations Analysis*, MORS Publication, 2007

[7] North, Paul D., "Experimentation via the Use of an Executable Workflow Model to Evaluate C2 Decision Quality" 2008 Command and Control Research and Technology Symposium

[8] Keeney, Ralph L., *Value Focused Thinking*, Harvard University Press, 1992