Proceeding of the 2004 American Control Conference
Boston, Massachusetts June 30 - July 2, 2004

Operational Reconfigurability in Command and Control[1]

N. Eva Wu† and Timothy Busch‡
†Dept. of ECE, Binghamton Univ., Binghamton, NY 13902, evawu@binghamton.edu
‡AFRL/IFSB,525 Brooks Road, Rome, NY 13441, timothy.busch@afrl.af.mil

Abstract

In this paper, a modern military air operation is modeled as a hybrid feedback control system. The focus of our study is on the impact of the redundancy architecture on the overall air operation, by which the functionality of the controller residing in a Command and Control Center is supported. The resilience of the architecture that reflects the quality of monitoring and management of redundancy is measured by operational reconfigurability. Maintenance policy and capability are also considered as parameters in the effort to quantify the availability of the controller.

1 Introduction

Our objective is to meet the demand of the ever decreasing cycle length in military air operations, which is the sum of the times required for planning, tasking, execution, and evaluation. This objective, when projected onto the expectation for a command and control (C2 hereafter) center, implies a more swift operation that involves information gathering, information processing, decision-making, and command issuing. The swiftness in turn requires a high level of C2 system availability. Availability of a system can be generally thought of as the fraction of the system uptime divided by the sum of the uptime and downtime. Our ultimate goal is to achieve nearly uninterrupted C2 operations.

It is apparent that a C2 center plays the role in an air operation as the controller in a feedback system. It carries out the functional mapping from information to decision in the feedback loop. A study conducted at the Draper Labs[6] that focuses on the effect of frequency of loop closure in air operations concludes that ability to close the loop at a higher rate (4 hour cycle v.s. 24 hour cycle), among other benefits, significantly shortens the time to achieve air campaign objectives. Other endeavors to enhance the C2 functional capability using different criteria and formalizations have also been reported[4,5,8,9].

The underlying assumption so far has been that the structure which supports the functional mapping in a C2 center is always intact. In reality, however, a typical C2 center has grown to be a large and complex system, and this system is imperfect. Many subsystem failures can occur for many different reasons. For example, a miscarriage in information flow can be attributed to a broken link, a faded or jammed channel, a power outage, a failed sensor, an impaired storage device, a crashed processor, a human operator error, etc. In general, failures that disable the C2 functional mapping can be related to subsystems designated to perform data storage, transmission, processing, or interpretation. They impact information availability, integrity, and decision making in C2 centers. The current status in the effort to address these issues is still in the very early stage of installing monitoring tools. There is a severe lack of consideration in tackling the more fundamental issues of redundancy architecture and an appropriate level of automation for failure accommodation. The latter is important to mitigate unnecessary human errors and delays.

In our view, the concern over the loss of C2 system availability could be effectively addressed by a conscious effort of modification to the existing architecture to eliminate all single point failures. The term *C2 system* has been and will be used in the following development to represent the network of subsystems and components that host and support the functional mapping performed by the controller in a C2 center. The most efficient way to achieve the modification is to make use of and effectively manage the redundancy likely already in existence in the C2 systems. The rest of the paper will explore such a possibility, and quantify the benefit of doing so to the overall air operation.

The paper is organized as follows. In section 2, C2 system modeling is discussed for the purpose of availability analysis. The notion of operational reconfigurability is introduced to describe the effective level of redundancy. Section 3 discusses the assessment of C2 system availability under variable conditions such as subsystem failure rate, effectiveness of redundancy management, maintenance policy, and restoration rate. Section 4 discusses the effect of C2 system operational reconfigurability on the outcome of an air operation. Section 5 concludes the paper.

[1]The first author acknowledges the support by the US Air Force Research Laboratory under Contract F30602-020-C-0225.

2 Operational reconfigurability

The first part of this section discusses qualitative availability modeling for a C2 system. Based on the model, the need for a measure on the effectiveness of redundancy management is argued, and the notion of operational reconfigurability introduced.

Availability[7] is the probability that a system is performing its required function at a given point in time when used under stated operating conditions. Among many definitions of availability, steady state availability will be considered, which represents the situation that the failure-restoration cycle has entered a steady state. Such a steady state definition will be assumed elsewhere in the paper without further explanation. The availability value of a system is determined by the following factors
(i) reliability[1] distributions of individual subsystems and the functionalities of the subsystems in relation to the overall system;
(ii) the policy and capability by which the system is maintained[2], such as the decision on the restoration of failed subsystems and the distribution of the time required to do so;
(iii) the methods and the likelihood of success in management of existing redundancy, which are heavily influenced by our ability to monitor and diagnose subsystem failures, and to reconfigure the system upon identifying the failures.

Figure 1 shows a hybrid model[10] of an air operation. The discrete state strategic model at the top will be further explained in Section 4. The tactical model is represented in Figure 1 by a continuous state closed-loop control system which governs the execution of an air operation. The forward loop contains a model of battle dynamics. The tactical state vector may contain in its components, for example, the strengths of Blue assets, the rates of change of the strengths, their geographic locations, rates of change of locations, etc. The functional mapping carried out by the controller in a C2 center is represented by the two blocks in the feedback loop. It is responsible for generating two sets of signals. One is an estimate of the strategic state $\hat{x}$, and another is a corresponding control $u_{\hat{x}}$ to drive the tactical state ξ to wherever desirable. Availability of various subsystems in a C2 system is required in order to ensure that $\hat{x}$, and the parameters associated with $\gamma_{\hat{x}}$ and the estimator are available and correct, η and r are available and current, and finally $\hat{\xi}$ and $u_{\hat{x}}$ are current and correct.

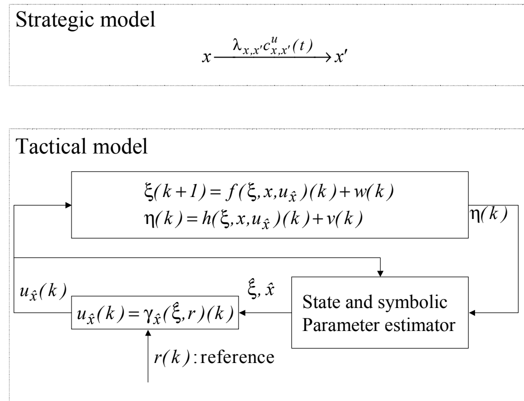

Figure 1 A two-level model of an air operation

An example of a functional decomposition of a C2 system is given in Figure 2 where the blocks marked TS (tactical and strategic sensors), DL (I/O control modules and data links), SM (storage media), CP (critical processors), and CS (critical software) represent some of the functional units. A functional unit is defined as a subsystem of a particular functionality that is necessarily available in order for the C2 system to be available. Each functional unit can be a complex interconnection of many subsystems. Considerable effort is usually necessary to arrive at a functional decomposition. Let A_{C2} denote the availability of the C2 system, and A_i is the availability of the ith functional unit. Then, the availability of a C2 system with N functional units is given by

$$A_{C2} = A_1 \times A_2 \times \cdots \times A_N. \tag{1}$$

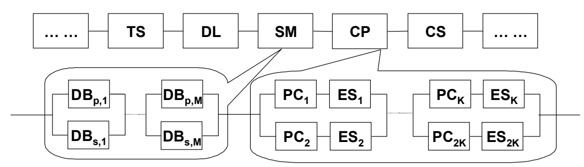

Figure 2 Some functional units in a C2 system

In the current C2 systems, vast opportunities exist for availability improvement without hardware addition, and without overburdening the subsystems in terms of processing speed, memory space, bandwidth, etc. The opportunities can be seized by, for example, assigning multiple tasks to multiple subsystems rather than assigning a single task to a dedicated subsystem within a functional unit, or using multiple copies of smaller data set for recursive processing rather than using a single copy of larger data set for batch processing. The expanded portions of Figure 2 show two examples of proposed architectural change for availability improvement. The original SM unit contains M storage media holding independent databases. These subsystems are named primary $\mathrm{DB}_{p,i}$ for $i = 1, \cdots, M$. Each primary subsystem is now appended with a redundant "cache", called a secondary $\mathrm{DB}_{s,i}$, using leftover storage space elsewhere to hold the most critical

[1] Reliability is the probability that a (sub)system will perform a required function for a given period of time when used under stated operating conditions

[2] Maintainability is the probability that a failed system will be restored to a specified condition within a period of time when maintenance is performed in accordance with prescribed procedures

and immediately needed data. The original CP unit contains $2K$ processor cell-Ethernet switch pairs with non-overlapping tasks. Each pair is now equipped with the necessary (software) tools of one other pair, and a K series redundant CP unit is formed.

This paper, however, is not intended to explore innovative ways to raise the level of redundancy, but to reason the significance and to assess the benefit of having an adequate redundancy level. A portion of the C2 system above will be used as a vehicle for our intended purpose. This portion is shown in Figure 3.

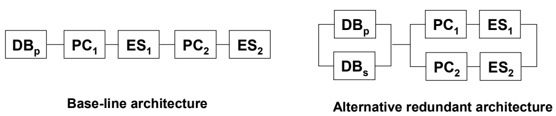

Figure 3 A glimpse of architectural change in C2

Most functional units within C2 are themselves complex interconnections of components. Statistical modeling[1] of the failure process of individual subsystems must follow carefully designed experiments of data collection, parameter (or distribution) estimation, and goodness-of-fit tests. Based on our initial investigation, failure rates (number of failures per unit time) of many individual subsystems are below 10^{-5}/hour, when intermittent failures are excluded. Therefore, subsystems are reliable. There is no doubt that intermittent failures will reduce the C2 availability. Modeling of intermittent failrues is our ongoing effort. In addition, both diagnosis and restoration for permanent failures of some subsystems can be lengthy processes (hours to tens of hours). The most fundamental reason for need of redundancy is the fragility of an architecture that allows single point failures. Individual subsystems do fail and can fail at an unfavorable time. The consequence to an air operation can be detrimental, as will be seen in Section 4. It is a fact to be kept in mind that statistical model development results from our lack of knowledge of the physical processes leading to a failure. As a consequence, we can only infer from our sample of failure data to the general population, and our predictions tell little concerning an individual system or failure occurrence. Therefore, a non-redundant architecture with highly reliable subsystems is robust but fragile[3].

In order to measure the non-fragility. the notion of operational reconfigurability (OR hereafter) is introduced. In this paper, OR is specific to characterize the C2 system survivability with respect to single point failures. Consider a canonical redundancy architecture with a parallel-to-series interconnection where each parallel interconnection in the outmost layer is considered a functional unit. The right side of Figure 3 shows a two-layer canonical interconnection. It is degenerated in the sense that there are no parallel interconnections in the inner layer. Suppose there are N functional units in a single layer canonical decomposition, and each has M_n $(n = 1, \cdots, N)$ subsystems. Let $c_{m,n}$ denote the coverage of the mth subsystem in the nth functional unit, where $c_{m,n}$ is define as Prob(n^{th} unit operates | its m^{th} subsystem has failed). Evaluating $c_{m,n}$'s is not a trivial task[11] because of its association with monitoring, diagnosis, and redundancy management policy. Its value also depends on how many remaining operating subsystems are in the functional unit. In general, the larger the number (M_n), the larger the value of $c_{m,n}$ due to reduced risk in redundancy management.

Operational reconfigurability OR for a single-layer parallel-to-series interconnected system is given by

$$OR \equiv \min_{n \in \{1, \cdots, N\}} \frac{1}{M_n} \sum_{m=1}^{M_n} c_{m,n}. \tag{2}$$

In a multi-layered parallel-to-series interconnection scheme, the expression would contain layers of minimum-average operations. OR points to the weakest functional unit in terms of its ability to manage the redundancy for covering its first failure. In particular, since $c_{m,n} = 0$ whenever $m_n = 1$, $OR = 0$ for any system that contains a non-redundant functional unit. This is a measure without the influence by a priori subsystem failure distributions, which is precisely needed to reflect the non-fragility.

In the representative C2 system of Figure 3, let OR_b denote the OR for the baseline architecture, and OR_r denote the OR for the alternative redundant architecture. Then $OR_b = 0$, and $OR_r = 1$, if $c_{m,n} = 1 \; \forall \; m, n$. In general, $0 \leq OR \leq 1$. OR essentially measures the available redundancy in a C2 system and how it is managed. Because of its dependence on coverage, it is reflective of monitoring and supervisory control performance. Such performance indicates the C2 ability to allow restoration of system function via reconfiguration upon subsystem failures. Reconfiguration can mean the removal of a failed subsystem, the switch-on of a spare subsystem, rescheduling of jobs, rerouting of the information flow, redistribution of the information storage, etc.

3 OR and availability

This section discusses availability modeling in a more quantitative manner. Its relation to OR, and other parameters such as restoration rate and maintenance policy, is of particular interest. Two simplifying assumptions are made here. (i) All subsystems in Figure 3 have exponential failure time distributions. (ii) All restoration time distributions are also exponential. Through out the section the representative C2 system of Figure 3 is used.

Viewed as a canonical form, the baseline architecture in Figure 3 contains only one type of functional unit. On the other hand, the alternative redundant architecture carries two types of functional units. The composite availability expression in (1) allows us to solve for the availability of individual units independently. A complete solution for availability requires the specification of the maintenance policy. The following policies are considered in our study:

(i) restoration to as good as new in one step with a prescribed restoration rate independent of the failure state when a (functional) unit level failure occurs;

(ii) restoration to as good as new in one step with the lowest restoration rate for the failure state when a unit level failure occurs;

(iii) restoration to as good as new in one step with a restoration rate determined by the sum of average restoration times of all failed subsystems associated with the failure state when a unit level failure occurs;

(iv) restoration to as good as new in multiple steps with a restoration rate determined by the criterion of quickest unit recovery or of the most important subsystem recovery (e.g., primary v.s. secondary) when a unit level failure occurs;

(v) restoration to as good as new in one step with a prescribed restoration rate independent of the failure state when any subsystem failure occurs;

(vi) restoration to as good as new in one step with the lowest restoration rate for the failure state when any subsystem failure occurs;

(vii) restoration to as good as new in one step with a restoration rate determined by the sum of average restoration times of all failed subsystems associated with the failure state when any subsystem failure occurs;

(viii) restoration to as good as new in multiple steps with a restoration rate determined by the criterion of most speedy unit recovery or of most important subsystem recovery when any subsystem failure occurs.

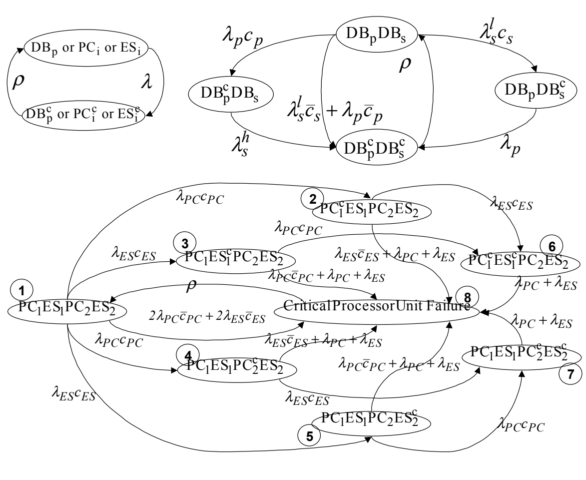

Figure 4 Rate transition diagrams of three types of functional units

The rate transition diagrams corresponding to the maintenance policy stated in (i) are shown in Figure 4 for the three types of functional units in Figure 3. Introduction of additional states is necessary under most of the other maintenance policies. The notations used in Figure 4 are as follows. λ denotes a failure rate, ρ denotes a restoration rate. A subsystem name appearing in a section of a state name indicates that the subsystem is up. A subsystem name with a superscript c appearing in a section of a state name indicates that the subsystem is down. c with appropriate subscript denotes coverage, and $\bar{c} = 1 - c$. Superscript l means a low value indicating, for example, that a subsystem is in standby mode, and superscript h means a high value (when the subsystem is no longer in standby). The state marked by "Critical Processor Unit Failure" has aggregated all CP unit failure states.

In general, for each of the three cases above, a set of state transition probabilities can be solved from the forward Kolmogorov equation $\dot{P}(t) = P(t)Q,\ P(0) = I$ which can be established directly from balancing the probability flow[2] from a rate diagram at each state. Therefore, transition rate matrix Q completely determines the set of transition probabilities. From the transition probabilities, any state probability can be easily calculated by setting appropriate initial state conditions. When the number of the states becomes large, numerical techniques and approximations must be sought to solve for the interested state probabilities directly from $\dot{\vec{p}}(t) = \vec{p}(t)Q, \vec{p}(0) = \vec{p}_0$, where $\vec{p}(t) = [p_1(t) \ \cdots \ p_n(t)]$ is the state (row) vector for an n-state functional unit.

Since our interest is in the steady state availability, the problem is much simplified. The steady state *unavailability* can be obtained by solving from the algebraic equation

$$\vec{p}_s Q_s = \begin{bmatrix} 0 & \cdots & 0 & 1 \end{bmatrix} \qquad (3)$$

for steady state probability vector $\vec{p}_s$, where Q_s is obtained by replacing the state equation involving the derivative of unit failure state by $\sum_{i=1}^{n} p_i = 1$. Arranging the states for the redundant critical processor unit in the order as marked in Figure 3, and let $\lambda_c = \lambda_{PC}$, $\lambda_e = \lambda_{ES}$, $c = c_{ES} = c_{PC}$, $\rho = \lambda_{PC} + \lambda_{ES}$, and ρ^l (ρ^h) denotes a restoration rate.

$$Q_s = \begin{bmatrix} -2\rho & \lambda_c c & \lambda_e c & \lambda_c c & \lambda_e c & 0 & 0 & 1 \\ 0 & -\rho - \lambda_e & 0 & 0 & 0 & \lambda_e c & 0 & 1 \\ 0 & 0 & -\lambda_c - \rho & 0 & 0 & \lambda_c c & 0 & 1 \\ 0 & 0 & 0 & -\rho - \lambda_e & 0 & 0 & \lambda_e c & 1 \\ 0 & 0 & 0 & 0 & -\lambda_c - \rho & 0 & \lambda_e c & 1 \\ 0 & 0 & 0 & 0 & 0 & -\rho & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\rho & 1 \\ \rho^{l,h} & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

The unavailability of the functional unit is given by the state probability corresponding to the unit failure state. Figure 5 shows the results of evaluation the unavailability reduction factor,

$$URF \equiv \frac{1 - A_{C2}(OR_b)}{1 - A_{C2}(OR_r)} = \frac{1 - A_b}{1 - A_r}, \qquad (4)$$

the ratio of the unavailability of the baseline to that of redundant architecture under maintenance policy (iii) for the SM unit (top), and for the CP unit (bottom), respectively. About a $20 \sim 95$ time reduction in unavailability in both units is observed for the range of failure rates indicated in Table 1, and for $\rho^l = 1/24$ hr^{-1}. Numerical comparisons are also made with respect to different maintenance polices listed above using two different restoration rates[3]. The results are not shown due to space limit.

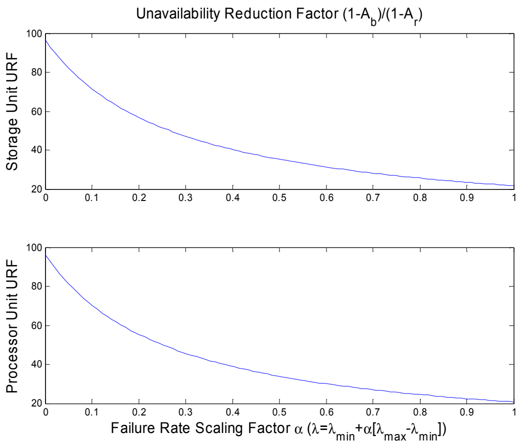

Figure 5 Unavailability reduction factor (URF) for SM and CP units due to redundancy architecture change

$\lambda_{PC} = \lambda_c$ (hr^{-1})	$9.0 \times 10^{-6} \sim 10^{-4}$	ρ^l (hr^{-1})	1/24
$\lambda_{ES} = \lambda_e$ (hr^{-1})	$7.4 \times 10^{-6} \sim 10^{-4}$	ρ^h (hr^{-1})	1/4
λ_p (hr^{-1})	$5.0 \times 10^{-6} \sim 10^{-4}$	OR_b	0
λ_s^l (hr^{-1})	$5.0 \times 10^{-7} \sim 10^{-5}$	OR_r	0.99
λ_s^h (hr^{-1})	$1.5 \times 10^{-5} \sim 10^{-3}$		

Table 1. SM and CP units failure, restoration rates, and OR numbers

4 OR and air operations

It is time to turn to the overall air operation, and investigate the benefit of a higher OR to the winning probability of Blue. This is an understandably difficult problem because of the conceivable complexity in establishing the linkage between the availability of the controller in a C2 center and the success of an air operation, though we have just established a definitive relationship between the availability and the OR. Fortunately, an earlier development[10] in our effort has provided the right framework to encourage an attempt. A brief review of the framework is in order.

A simple representation of a strategic model is shown at the top of Figure 1. It refers to the mathematical description of the evolution of a strategic plan in an air operation. It takes the form of a discrete state and continuous time Markov process. The model is specified by (i) a state space $\{\mathcal{X}\}$, (ii) a set of initial state probabilities $\{p_x(0), x \in \mathcal{X}\}$, and (iii) a set of state transition rates $\{\lambda_{x,x'}(t), x, x' \in \mathcal{X}\}$ from the current state x to the next state x'[2]. Figure 6 shows a low resolution example of a strategic model composed of 4 binary states: (Blue threatened, Blue defeated, Red targeted, Red defeated). A state of (True, False, True, False) can be represented by $x = 1010$ in binary. The meanings of the remaining states can be similarly explained. States 0000, 1000, and 1010 in Figure 6 are *transient states* and states 0100, 0101, and 0001 are *absorbing states*. Depending on whether preserving the Blue assets besides destroying the Red assets is also part of the mission of an air operation, the set of desirable outcomes for Blue can be one of $\{0101, 0001\}$ and $\{0001\}$.

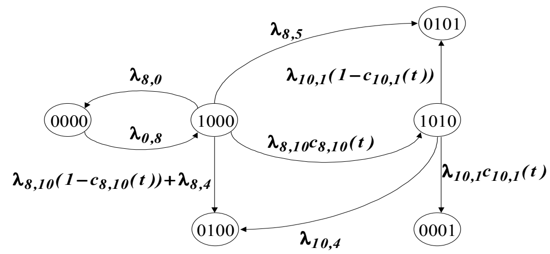

Figure 6 A low resolution strategic model

An important set of parameters introduced in our earlier work is the set of transition coverage values[10]. A transition coverage associated with a transition from x to x' is the conditional probability that the intended transition in fact occurs given that a triggering event has arrived. It is denoted by $c^u_{x,x'}(t)$, where u indicates its dependence on the control policy used in the tactical operation produced by the controller in a C2 center. It can be seen that a transition coverage serves to effectively modify the corresponding transition rate via $\lambda_{x,x'} c^u_{x,x'}(t)$. The transitions that have transition coverage attached to them are called controllable transitions. Blue's control objective is to maximize the transition coverage under the constraints of its resources and battle dynamics. The presence of C2 availability naturally modifies the originally defined transition coverage[10] for any controllable transition from x to x' in the following manner.

$$\mathcal{C}^u_{x,x'} = c^u_{x,x'} A_{C2},\ \bar{\mathcal{C}}^u_{x,x'} = c^u_{x,x'}(1 - A_{C2}) + (1 - c^u_{x,x'}) \quad (5)$$

where

$$\mathcal{C}^u_{x,x'} + \bar{\mathcal{C}}^u_{x,x'} = 1 \quad (6)$$

forms the Poisson decomposition[10] of the associated transition rate $\lambda_{x,x'}$. The original Poisson decomposition becomes a special case when C2 availability is perfect. It is obvious that the introduction of an imperfect C2 system availability reduces the effectiveness of the controller in the C2 center.

We now examine the average effect as well as the real-time effect of an imperfectly available C2 system on the air operation model of Figure 6. The following data are

[3]The authors would like to thank Ms. Xiaoxia Wang for her help in carrying out some of the availability calculations

used in producing the result in Figure 7 and Table 2. $\lambda_{0,8} = 0.2$, $\lambda_{8,0} = 0.02$, $\lambda_{8,4} = 0.04$, $\lambda_{8,5} = 0.001$, $\lambda_{8,10} = 0.4$, $\lambda_{10,4} = 0.005$, $\lambda_{10,1} = 0.05$, $c_{8,10} = .95(1 - .5e^{-t/5})$, and $c_{10,1} = .95(1 - .5e^{-t/10})$. Modifications in transition coverage are as follows.

$$\mathcal{C}^u_{8,10}(t) = c^u_{8,10}A_{C2},\ \ \bar{\mathcal{C}}^u_{8,10} = c^u_{8,10}(1-A_{C2})+(1-c^u_{8,10}),$$

$$\mathcal{C}^u_{10,1} = c^u_{10,1}A_{C2},\ \text{and}\ \bar{\mathcal{C}}^u_{10,1} = c^u_{10,1}(1-A_{C2})+(1-c^u_{10,1}),$$

where various cases of A_{C2} considered are listed in Table 2. Function $1(t)$ in Table 2 denotes the unit step function. The 4 hour and the 24 hour time slots of unavailable C2 system correspond to the restoration rates used in the calculation of the previous section. The results are shown in Figure 7. Final winning probabilities are also summarized in Table 2.

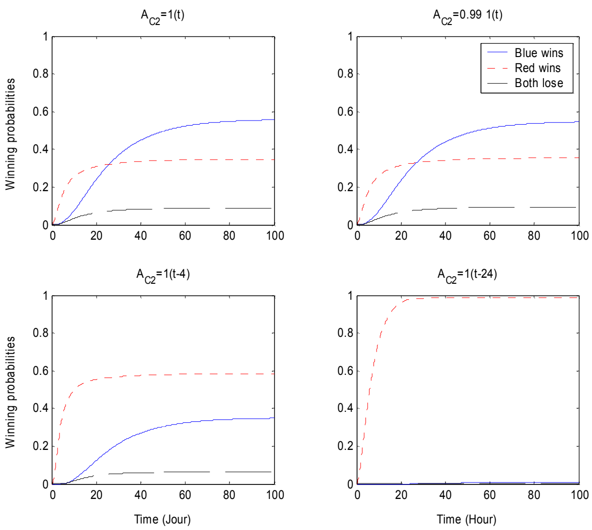

Figure 7 Winning probabilities up to the 100th hour of a military air operation

C2 availability	Blue wins	Red wins	Both lose
$A_{C2} = 1(t)$	0.56	0.35	0.09
$A_{C2} = 0.99 \times 1(t)$	0.54	0.36	0.10
$A_{C2} = 1(t-4)$	0.35	0.58	0.07
$A_{C2} = 1(t-24)$	0.01	0.99	0.00

Table 2 Winning probabilities at $t = 100$ hours when transient state probabilities have died out.

It can be seen from Figure 7 and Table 2 that a slight reduction in C2 availability has a limited effect on the outcome of the air operation on average. However, when the real-time unavailability of a C2 system falls within a critical period, the outcome can be disastrous. The latter case is shown in the two plots at bottom of Figure 7, and two items at the bottom of Table 2. These show where the fragility lies. An enhanced operational reconfigurability can reduce the unavailability and hence fragility by 2 orders of magnitude as shown in the example of the previous section. The reduction is achieved by filling in the periods of operation interruptions with a fairly unsophisticated usage of existing redundancy.

5 Conclusions

This paper delineated the importance and the potential of being able to provide and manipulate redundancy in the command and control system of a military air operation. The effort boils down to modification of the C2 system architecture so as to raise the system operational reconfigurability. An enhanced OR helps reduce the fragility of an otherwise robust system. The cost of reduction of fragility is the extra complexity of the system which must include diagnosis and management of redundancy (or supervisory control). The complexity introduced, however, is a miniature increment of a more costly and less carefully studied effort in setting up monitoring tools within C2 centers. Some simple but quantified case studies were presented to support our argument. Our ongoing effort is focused on more detailed availability modeling.

References

[1] G. Casella, and R. L. Berger, *Statical Inference*, 2nd Edition, Duxbury, 2002.

[2] C. G. Cassandras, and S. Lafortune, *Introduction to Discrete Event Systems*, Kluwer Academic Publishers, 1999.

[3] J. M. Carlson, and J. Doyle, Highly optimized tolerance: robustness and design in complex systems, *Physical Review Letters* vol.84, pp.2529-2532, 2000.

[4] J. B. Cruz, Jr., M. A. Simaan, A. Gacic, H. Jiang, B. Letellier, M. Li, and Yong Liu, Game-theoretic modeling and control of a military air operation, *IEEE Transactions on Aerospace and Electronic Systems*, vol. 37, pp.1393-1405, 2001.

[5] M. Curry, J. Wohletz, C.G. Cassandras, and D. Castanon, Modeling and control of a joint air operation environment with imperfect information, *Proc. SPIE 16th Annual Intl. Symposium*, vol.4716, pp.41-52, 2002.

[6] Draper Laboratory, *Closed-Loop, Hierarchical Control of Military Air Operations*, Final Report CSDL-R-2914 to DARPA Joint Forces Air Component Commander Program, 2001,

[7] C.E. Ebeling, *An introduction to Reliability and Maintainability Engineering*, McGraw-Hill, 1997.

[8] D. Ghose, J.L. Speyer, and J.S. Shamma, A game theoretical multiple resource interaction approach to temporal resource allocation in an air campaign, Annals of Operations Research, pp.15-40, 2002.

[9] J. M. Wohletz, D. A. Castanon, M. L. Curry, Closed-loop control for joint air operations, *Proc. American Control Conference*, pp.4699-4704, 2001.

[10] N.E. Wu, and T. Busch, Strategic reconfigurability in air operations, *Proc. IEEE Conference on Decision and Control*, 2003.

[11] N.E. Wu, Reliability of fault tolerant control systems, part i & ii, *Proc. IEEE Conference on Decision and Control*, 2001.