Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)

Assuring Safety of NextGen Procedures

Cody H. Fleming, Nancy G. Leveson, M. Seth Placke
System Safety Research Lab
Massachusetts Institute of Technology
Cambridge, USA
msplacke@mit.edu

***Abstract*— This paper introduces an innovative approach to analyzing safety in the next generation of air traffic management systems. The proposed method is based on systems and control theory and is able to capture system design and component interaction causes that are increasingly frequent in accidents. The new methodology is applicable during the entire design lifecycle from early concept selection through final certification. Hazard analysis of a completed NextGen concept, In-Trail Procedure, is demonstrated as well as use in the early concept development of Trajectory Based Operations.**

Hazard Analysis; Safety Guided Design; Safety of Air Traffic Management Concepts

I. INTRODUCTION

In both the United States and Europe, planning efforts are underway to increase air traffic capacity and efficiency in anticipation of rising demand levels. Common goals of these efforts include increased efficiency, lessened environmental impact, and increased safety across the air transportation system through the implementation and integration of increasingly complex technologies [3, 4].

It has been recognized that traditional risk-based modeling techniques do not adequately account for human error and software related accidents [2]. Software is increasingly an important part of systems and allows enormously more complex and tightly coupled systems to be constructed. The potential for accidents arising from unsafe interactions among non-failed components, i.e., unplanned system and software behavior, is increasing. As automation becomes an increasingly key component of air traffic management (ATM), human controllers will begin to shift from direct control to supervision of automation, which can complicate human decision-making. Like software, the changing roles of pilots and ground controllers introduces the potential for new causes of accidents that are not well handled by today's failure-oriented and hardware-oriented approaches.

To deal effectively and efficiently with these new accident causes in the next generation air traffic management schemes, more powerful risk management tools are needed. This paper describes a new approach to hazard analysis, called STPA (System-Theoretic Process Analysis), which is based on systems and control theory rather than reliability theory. In STPA, safety is treated as a control problem rather than a failure prevention problem, allowing not only consideration of the causes of the component failure accidents that were predominant in the past but also the new causality factors that are increasingly important today. New technology is introducing new types of accident causes and the causality models of the past, upon which risk management is based, must be extended to include these new accident causes.

STPA can be used in all phases of system development. The earlier that safety is part of the decision making process, the easier the final certification will be and, hopefully, the safer the final system will be. In this paper, we describe STPA and the new, extended model of causality on which it is based called STAMP. We then illustrate how STPA can be used for different purposes by describing the hazard analysis of a relatively complete design called the In-Trail Procedure (ITP) and how STPA can be used in the very early concept development of Trajectory Based Operations (TBO).

II. STAMP

STAMP (Systems Theoretic Accident Model and Process) is a new accident causality model that accounts for the non-linear, indirect and feedback relationships among events in identifying potential causes of accidents. In this way, the traditional "chain of failure events" model is extended to consider new types of accident causality brought about by component interactions (rather than just component failures), cognitively complex human mistakes, management and organizational errors and software errors (particularly requirements errors). Accidents or unacceptable losses can result not only from system component failures but also from interactions among system components, both physical and social, that violate system safety constraints.

In systems theory, emergent properties associated with a set of components are related to constraints upon the degree of freedom of those components' behavior [7]. Since system safety is an emergent property, it may be treated at the system level as a control problem rather than at the component level as a reliability issue. The controls may be managerial, organizational, physical, operational or manufacturing. When these control mechanisms are not adequate to mitigate component failures, external disturbances and/or dysfunctional interactions among system components, then an accident or loss will occur.

This research was partially funded by NASA Contract NNL10AA13C

In a systems-theoretic view of safety, the emergent safety properties are controlled or enforced by a set of safety constraints related to the behavior of the system components. Safety constraints specify those relationships among system variables or components that constitute the non-hazardous or safe system states: for example, the power must never be on when the access door to the high-power source is open; two aircraft must never violate minimum separation requirements; pilots in a combat zone must be able to identify targets as hostile or friendly; and the public health system must prevent the exposure of the public to contaminated water and food products. Accidents result from interactions among system components that violate these constraints—in other words, from a lack of appropriate constraints on component and system behavior.

III. STPA

STPA is a hazard analysis technique built on STAMP. As described above, accidents are viewed in STAMP as resulting from inadequate enforcement of constraints on system behavior. Figure 1 shows a generic safety control structure in place to enforce safety constraints. Each hierarchical level of the control structure represents a control process and control loop with actions and feedback. Two control structures are shown; system development (on the left) and system operations (on the right), both of which have different responsibilities with respect to enforcing safe system behavior. The reason behind the inadequate enforcement may involve classic component failures, but it may also result from unsafe interactions among components operating as designed or from erroneous control actions by software or humans.

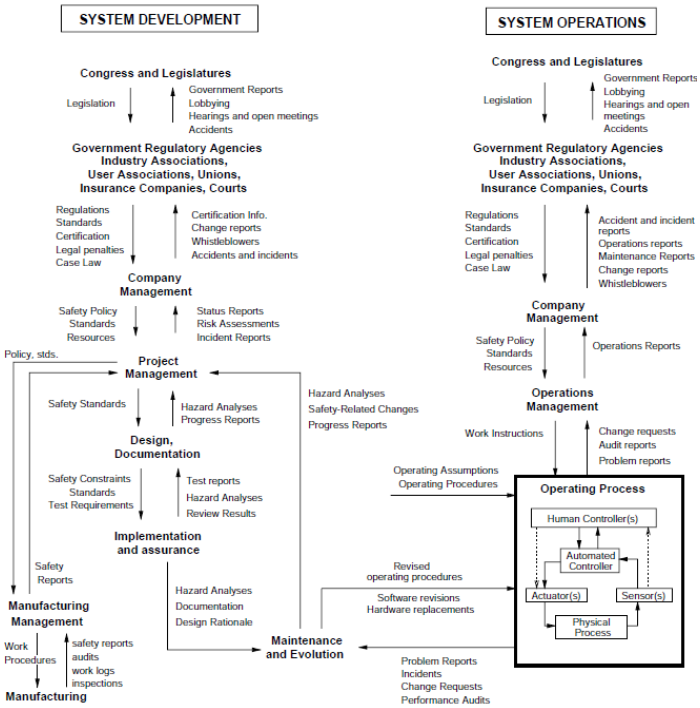

Figure 1 - Socio-Technical Safety Control Structure

Human and automated controllers use a process model (usually called a mental model for humans) to determine what control actions are needed for various states of the controlled process (Figure 2). The process model contains the controller's understanding of 1) the current state of the controlled process, 2) the desired state of the controlled process and 3) the ways the process can change state. Software and human errors often result from incorrect process models, e.g., the software thinks the spacecraft has landed and shuts off the descent engines. Accidents can therefore occur when an incorrect or incomplete process model causes a controller to provide control actions that are hazardous. While process model flaws are not the only cause of accidents, they are a major contributor for accidents involving software and human errors.

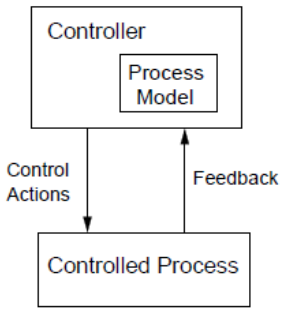

Figure 2 - Simple Control Loop and Process Model

There are four types of hazardous control actions that need to be eliminated or controlled to prevent accidents:

1) *A control action required for safety is not provided*
2) *An unsafe control action is provided that leads to a hazard*
3) *A potentially safe control action is provided too late, too early, or out of sequence*
4) *A safe control action is stopped too soon or applied too long*

Identifying the potentially unsafe control actions for the specific system being considered is the first step in STPA. These unsafe control actions are used to specify safety requirements and constraints on the behavior of both the system and its components. Additional analysis can then be performed to identify the detailed scenarios leading to the violation of the safety constraints and used to generate more detailed safety requirements. As in any hazard analysis, these scenarios are the basis for designing controls and mitigation measures for the hazards. Any hazards that cannot be adequately controlled at the system level must be allocated in the form of behavioral requirements on the lower-level system components.

IV. STPA for Design Evaluation

In this section, the use of STPA to evaluate an existing design is illustrated by using the NextGen "In-Trail Procedure in Oceanic Airspace." A complete version of this analysis can be found in [6].

A. ITP Background

ITP [9] is being implemented to increase operational efficiency and throughput in oceanic airspace. Due to the lack of radar coverage in such remote airspace, air traffic controllers have used conservative minimum separation rules along predefined flight paths and organized tracks to ensure safe passage. Because Air Traffic Control (ATC) has limited capability for monitoring the exact positions or separations of aircraft in an oceanic sector, the separation requirements are often much larger than those in continental sectors with sufficient surveillance. These concerns have precluded passing maneuvers that would facilitate different cruising speeds within one track and flight level changes aimed at increasing fuel efficiency as aircraft experience weight change through flight (fuel burn).

The ITP will allow many of these previously blocked flight level changes to occur. ITP enables either leading or following "Same Track" aircraft to perform a climb or descent to a requested flight level through an intervening flight level. Level changes are currently restricted to two flight levels. The crew will use information derived on the aircraft to determine if the criteria for executing the ITP are met with respect to one or two Reference Aircraft at intervening flight levels. Note that the standard separation minimum between aircraft does not hold during the ITP maneuver but the ITP equipment provides information to the flight crew to ensure that the ITP-defined reduced separation minimums are observed.

B. STPA Step 0

As in the traditional hazard analysis, the STPA process starts by identifying hazards, although hazards are not equated to failures as is often the case. Instead, a hazard is defined as a system state that under worst-case environmental conditions will lead to a loss or accident. This definition encompasses undesired states resulting from many causes, including but not limited to component failures. The general hazards for aircraft include:

- H-1: A pair of controlled aircraft violate minimum separation standards
- H-2: Aircraft enters unsafe atmospheric region
- H-3: Aircraft enters uncontrolled state
- H-4: Aircraft enters unsafe attitude (excessive turbulence or pitch/roll/yaw that causes passenger injury but not necessarily aircraft loss)
- H-5: Aircraft enters a prohibited area

For the introduction of ITP, hazard H-1 is the most relevant and leads to the high-level system safety requirement/constraint: "The ITP must not cause a pair of controlled aircraft to violate minimum separation standards." The hazard analysis identifies ITP system and component requirements necessary to enforce this constraint.

STPA works on a functional control model of the system and is performed in two steps. Figure 3 shows the safety control structure for ITP.

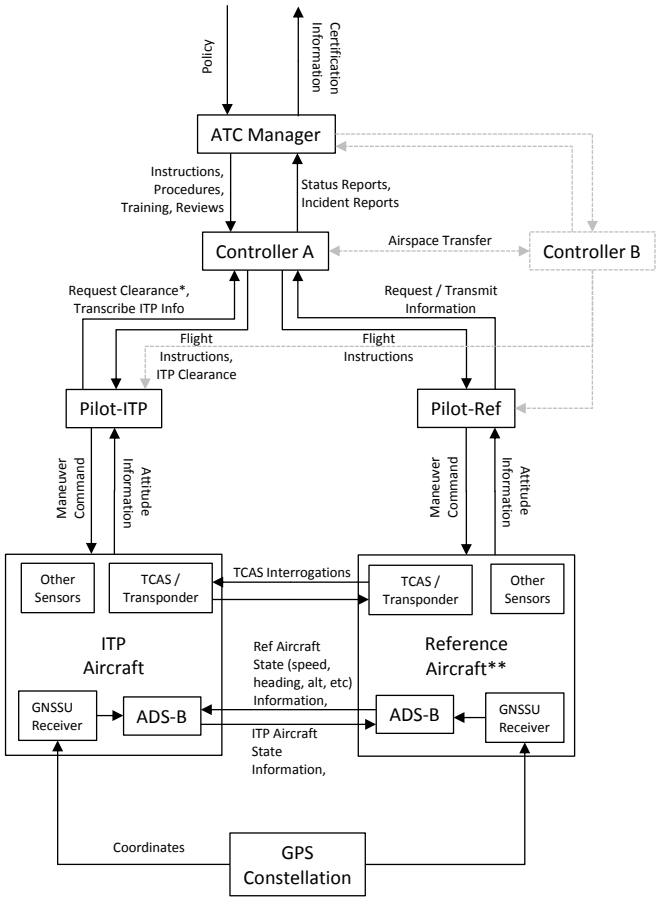

Figure 3 – ITP Control Structure

C. STPA Step 1

STPA Step 1 identifies hazardous control actions for each component that can lead to one or more of the defined system hazards. These identified unsafe control actions are used to refine the high-level safety constraints/requirements into more detailed safety requirements. The four general types of unsafe control actions were shown above and a table can be used to document the hazardous control actions identified, as in Table 1. The hazardous control actions can then be translated into high-level system and component safety requirements and constraints.

To produce the table, each potential entry is evaluated to determine whether that control action can lead to the system hazard (violation of minimum separation assurance). Consider the potential control action "Flight Crew Executes ITP" in Table 1. If the flight crew does not provide that control action, hazard H-1 does not result and the table entry is empty. On the other hand, there are several conditions under which providing the control action (execute ITP) could lead to the hazard, namely: executing the ITP procedure when it is not approved; executing it when the ITP criteria are not satisfied; and executing it with incorrect parameters (e.g., an incorrect climb rate or final altitude).

TABLE I. POTENTIALLY UNSAFE CONTROL ACTIONS FOR THE FLIGHT CREW

Controller: Flight Crew	Not Providing Causes Hazard	Providing Causes Hazard	Wrong Timing/Order Causes Hazard	Stopped Too Soon/Applied Too Long
Execute ITP		ITP executed when not approved. ITP executed when criteria are not satisfied. ITP executed with incorrect climb rate, final altitude, etc.	ITP executed too soon before approval. ITP executed too late after reassessment.	ITP aircraft levels off above requested FL. ITP aircraft levels off below requested FL.
Abnormal Termination of ITP	FC continues with maneuver in dangerous situation	FC aborts unnecessarily. FC does not follow regional contingency procedures while aborting.		

Once the tables are created, the identified unsafe control actions are rewritten as high-level system safety constraints. The constraints are refined further, in a top-down system engineering process, during STPA Step 2. The identified constraints are used not only in the refinement to more detailed requirements, but also to provide traceability back to the unsafe control actions from the more detailed requirements and constraints in order to document where these requirements came from and therefore why they are needed.

The constraints on the ITP flight crew are:

1) SC-FC.1 - The flight crew must not execute the ITP when it has not been approved by ATC.

2) SC-FC.2 - The flight crew must not execute an ITP when the ITP criteria are not satisfied.

3) SC-FC.3 - The flight crew must execute the ITP with correct climb rate, flight levels, Mach number, and other associated performance criteria.

4) SC-FC.4 - The flight crew must not continue the ITP maneuver when it would be dangerous to do so.

5) SC-FC.5 - The flight crew must not abort the ITP unnecessarily. (Rationale: An abort may violate separation minimums).

6) SC-FC.6 - When performing an abort, the flight crew must follow regional contingency procedures.

7) SC-FC.7 - The flight crew must not execute the ITP before approval by ATC.

8) SC-FC.8 - The flight crew must execute the ITP immediately when approved unless it would be dangerous to do so.

9) SC-FC.9 - The crew shall be given positive notification of arrival at the requested FL.

The second step identifies potential causes for the violation of these safety constraints.

D. STPA Step 2

The second step of STPA examines each control loop in the safety control structure to identify potential causal factors for each hazardous control action, i.e., the scenarios that could lead to providing one of the unsafe control actions. Some causal scenarios can be eliminated. Those that cannot must be controlled or mitigated during system design and operation. This process may lead to design changes and more detailed behavioral (functional) requirements on the system components to ensure that all the components operating together cannot create a system hazard, in this case H-1 (violation of minimum separation standards). The system and component-level requirements are used to design controls and to certify the safety of the system and of its components. This process refines the high-level safety requirements identified in Step 1. Figure 4 shows a generic control loop that can be used to guide this step. While STPA Step 1 focuses on the provided control actions (upper left arrow in the loop), STPA Step 2 expands the analysis to consider causal factors along the rest of the control loop.

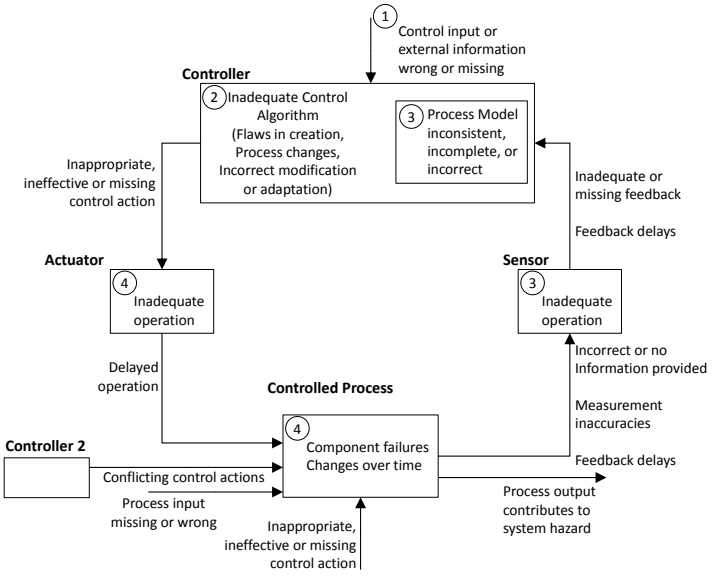

Figure 4 - General Control Loop with Causal Factors

As an example, Figure 5 shows the process model components for the flight crew, the basic algorithm used by the flight crew (responsibilities) and some of the causes for the flight crew to provide an unsafe command.

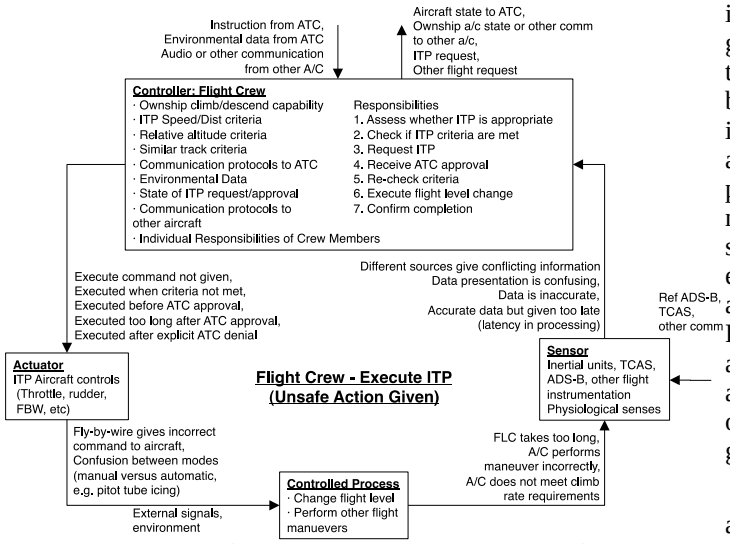

Figure 5 - Some reasons for the FC to execute ITP when it is unsafe to do so

Once the second step of STPA has been applied to determine potential causes for each hazardous control action identified in STPA Step 1, the causes should be eliminated or controlled in the design at the system level or detailed behavior requirements must be levied on the system components.

E. Comparison to Traditional Methods

STPA found more potential causes of the hazard considered (violation of separation requirements) than the traditional hazard analysis performed on ITP [9]. In this comparison, we included with the DO-312 analysis a set of additional requirements provided by the FAA for ITP [10]. Our analysis identified nineteen safety requirements that were not in either of the two official NextGen documents. One example of an important omitted requirement involves the reference aircraft. DO-312 assumes that the reference aircraft will not deviate from its flight plan during ITP execution. There needs to be a contingency or protocol in the event that the reference aircraft does not maintain its expected speed and trajectory, for example, because of an emergency requiring immediate action.

At a more detailed comparison level, consider the unsafe control action involved in the ITP procedure being performed by the flight crew when the minimum distance criterion for safe ITP is not satisfied and neither the flight crew nor ATC detects this non-compliance. The official fault trees used in DO-312 assume that failure to comply can occur either because the flight crew does not understand the minimum distance required or the ATC does not receive the required data or fails to detect non-compliance due to a communication error involving corruption of data during transport. There are, of course, many other reasons for communication errors but these are ignored in the original fault tree analysis of ITP. A fault tree also assumes independent behavior, but the interaction and behavior of the flight crew and ATC may be coupled, with the parties exerting influence on each other or both being influenced by high-level system conditions.

In contrast, the STPA causal factors include the basic communication errors included in the fault tree, but also include additional reasons for communication errors as well as guidance for understanding human error within the context of the system. Communication errors may result, for example, because there is confusion about multiple sources of information (for either the flight crew, or ATC), confusion about heritage or newly implemented communication protocols, or simple transcription or speaking errors. There is no way to quantify or verify the probabilities of any of these sources of error for many reasons, particularly because the errors are dependent on context and the operator environments are highly dynamic and, in fact, not necessarily designed yet. Perhaps that is why they were omitted from the fault trees. In any case, the rigorous process used by STPA, in our experience and the experience by others who have used it on many types of systems, leads to a more complete analysis with more guidance on what to include compared to other methods.

Despite being more powerful than traditional hazard analysis techniques, in all instances where STPA and the traditional analysis techniques were applied to the same real system, STPA took less effort and time [11].

V. USING STPA IN THE CONCEPT FORMATION STAGE

The previous section described the use of STPA when basic operations have already been defined, but much of NextGen is still being developed and is at the preliminary hazard analysis (PHA) phase. PHA usually involves identifying the high-level system hazards and then determining their risk in terms of severity and likelihood. A PHA has been done by the JPDO on TBO using traditional hazard analysis techniques [8].

The JPDO PHA identified a useful list of hazardous hardware failures, such as ADS-B or Data-Link Communication outage that can be used to make design decisions to reduce the likelihood of these failures leading to a hazard. The existing PHA focus on failures, however, made the results with respect to software and human errors less helpful. Identified hazards such as ‘controller confusion’ or ‘software anomaly’ have proposed mitigations that include ‘comprehensive system testing,’ ‘training’ and ‘phased implementation.’ While these human error and software anomaly mitigations may be helpful, they do not address the underlying, system-wide complexity issues that may result in system degradation and accidents and, most important, they do not address how to make specific TBO design decisions in order to eliminate or mitigate the potentially unsafe software or human behavior.

In this section we describe how STPA could be used to do a preliminary hazard analysis on TBO that provides more information for the conceptual design process.

A. Trajectory Based Operations

TBO is a shift from the current ATM strategy of Clearance-Based Operations that operates with very little automation to a system where aircraft will follow four dimensional paths (latitude, longitude, altitude, time) called trajectories [5]. When fully realized these trajectories will represent the gate-to-gate movement of aircraft and be the basis for ATM that seeks to resolve conflicts by altering trajectories and coordinating

responses to unplanned events such as weather and dynamic airspace needs.

The implementation of TBO will bring drastic changes to the roles of ATC and the tools they use, particularly how they load-share with automated systems. New location technologies such as ADS-B and communication capabilities such as net-centric data links used to enable TBO will change the way information is shared and increase the visibility of the airspace state to all users. Figure 5 is a diagram of the proposed information flow under TBO.

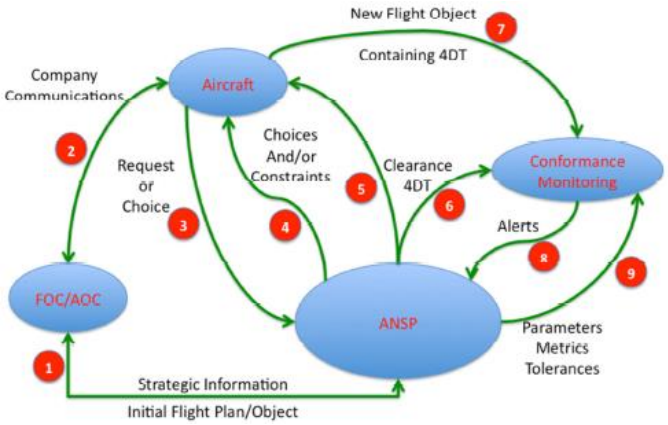

Figure 6 - TBO Information Flows [5]

A central aspect of TBO is trajectory negotiation, i.e., the process by which trajectories will be proposed, altered and executed. In the fully realized system, trajectory negotiation will be used for all aspects of flight management with safety critical clearances being an occasional exception. This concept of trajectory negotiation is an example of a design problem where STPA could assist in system architecture design and definition.

Reference [5] divides negotiations into 'Strategic' and 'Tactical,' which are loosely defined as resolutions beyond and up to a 20 minute time horizon respectively. Negotiation will have several phases, each of which will likely require different safety control structures as the hierarchy of control will change depending on the time horizon. The tradeoffs and design of these control structures can be informed by using STPA to analyze their different safety characteristics. Use of STPA can help further clarify differences between Strategic and Tactical timeframes by evaluating differences in control hierarchy and the needs of decision makers in each context. STPA analysis could show the need for additional time horizons and/or aid in the creation of decision rules regarding when and how to switch between them.

For the sample analysis shown here, an en-route tactical time horizon control structure is considered. The term Air National Service Provider (ANSP) is used to represent ATC to reflect both the official TBO literature and the idea that ATM will have to adapt from the current ATC model.

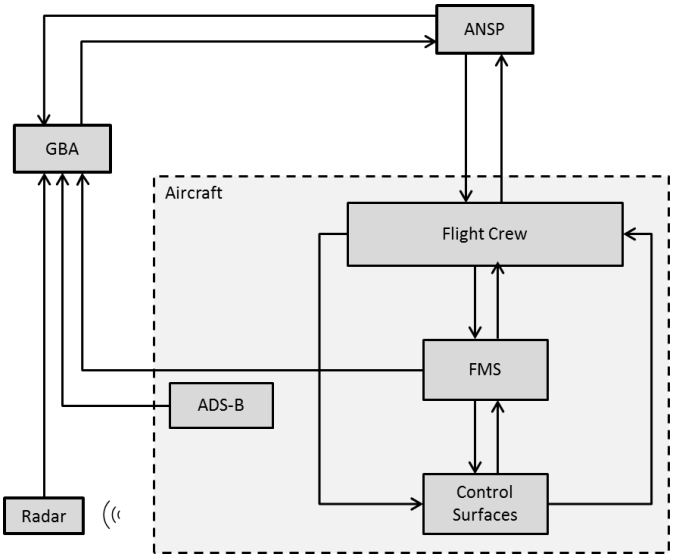

Figure 7 – Possible TBO Control Structure

B. *STPA Step 0*

The hazards listed in 4.2 are still applicable. As with ITP, hazard H-1 is the most relevant and leads to the high-level system safety constraint: "TBO control must not cause a pair of controlled aircraft to violate minimum separation standards." The hazard analysis identifies TBO system and component requirements necessary to enforce this constraint. Once again, STPA is performed in two steps, shown below.

C. *STPA Step 1*

Step 1 of STPA has been performed as described in section 4.3 to TBO resulting in the following matrix of hazardous control actions for the ANSP.

TABLE II. POTENTIALLY UNSAFE CONTROL ACTIONS FOR THE ANSP

Controller: ANSP	**Not Providing Causes Hazard**	**Providing Causes Hazard**	**Wrong Timing/Order Causes Hazard**	**Stopped Too Soon/Applied Too Long**
Approve 4DT	Not hazardous (however, delay leads aircraft to an imminent collision, or no provision but FC executes trajectory anyway)	Approved 4DT leads to LOS Approved 4DT is different than proposed 4DT provided to wrong aircraft	Approved 4DT is 'out of date' 4DT approved after atmospheric conditions (or other) change	N/A
Deny 4DT	4DT will lead to LOS Current 4DT is unsafe and new constraints are not provided	Denial pushes conflict resolution into imminent time horizon	Denial comes too late and leads aircraft to an imminent collision	N/A

Controller: ANSP	Not Providing Causes Hazard	Providing Causes Hazard	Wrong Timing/Order Causes Hazard	Stopped Too Soon/Applied Too Long
Command Exceptional 4DT/Vector	Nominal 4DT maintained during imminent emergency (and ANSP has separation responsibility)	Commanded path leads to MAC Aircraft is not capable of executing commanded trajectory Command conflicts with onboard RA No emergency exists and execution overloads crew or injures crew/passengers Command provided to wrong aircraft	Out of sequence with onboard warnings or given too late	N/A

These unsafe control actions can be translated into high-level TBO safety constraints. Step 2 analysis provides detailed information about how these constraints might be violated in a particular system design concept. Each unique architecture proposal (physical and procedural) will yield different results as components are removed, changed, added or shifted within the system. The analysis results may be compared to further understand the impacts of design decisions and will yield safety constraints to be followed during design and implementation of both technical subsystems and procedures. For example, as various trajectory monitoring systems are evaluated and developed (implementations of radar, ADS-B, GPS and others), STPA Step 2 analysis will assist designers in considering the way varying data sources are used.

VI. Conclusions

This paper has described a new hazard analysis method that is based on system engineering and control systems theory principles. It provides a more complete analysis method for the complex and software-intensive systems than the traditional methods currently being used. In addition to hardware failure, this method captures component interaction behavior, including software requirements, human-in-the-loop enforcement of safety constraints, and timing. STPA may be used throughout the entire design process from early concept design to certification upon system completion, all the while promoting safety-guided design.

References

[1] N. G. Leveson, Engineering a Safer World, MIT Press, 2012.

[2] D. Fowler, G. Le Galo, E. Perrin, and S. Thomas, So it's Reliable But is it Safe? A More Balanced Approach to ATM Safety Assessment, *Proceedings of the 7th US / Europe Seminar on ATM Research & Development, Barcelona*, July 2007.

[3] ICAO, Doc 9854 Global Air Traffic Management Operation Concept, 2005.

[4] JPDO, Concept of Operations for the NextGen Air Transportation System v3.2, September 2010.

[5] JPDO, TBO Study Team Report, December 2011.

[6] C. H. Fleming, M. Spencer, N. G. Leveson, C. Wilkinson, Safety Assurance in NextGen, NASA Technical report NASA/CR-2012-217553.

[7] Peter Checkland, *Systems Thinking, Systems Practice*, John Wiley & Sons, New York, 1981.

[8] JPDO, Capability Safety Assessment of TBO, February 2012.

[9] RTCA/DO-312. Safety, Performance, and Interoperability Requirements Document for the In-Trail Procedure in Oceanic Airspace (ATSA-ITP) Application, RTCA Incorporate, Washington DC.

[10] FAA, Interim Policy and Guidance for ADS-B Surveillance Applications Systems Supporting Oceanic In-Trail Procedures, 2010.

[11] N. Leveson, A New Approach to System Safety, *Safety-Critical Systems Symposium*, Bristol, U.K, February 2013

Author Biography

Cody Fleming is pursuing a doctoral degree in Aeronautics and Astronautics at the Massachusetts Institute of technology. He holds a BS degree in Mechanical Engineering from Hope College and masters in Civil Engineering from MIT. Prior to returning to MIT, he spent 5 years working in space system development for various government projects.

Nancy Leveson is Professor of Aeronautics and Astronautics and also Engineering Systems at the Massachusetts Institute of Technology. She is an elected member of the National Academy of Engineering. Professor Leveson conducts research on the topics of system safety, software and system engineering, and human-computer interaction. In 1999 she received the ACM Allen Newell Award for outstanding computer science research and in 1995 the AIAA Information Systems awarded for "developing the field of software safety and for promoting responsible software and system engineering practices where life and property are at stake." In 2005 she received the ACM Sigsoft Outstanding Research Award. She has published over 200 research papers and is author of two books, *Safeware: System Safety and Computers* published by Addison-Wesley and *Engineering a Safer World* published by MIT Press.

Seth Placke is pursuing a master's degree in Engineering Systems at the Massachusetts Institute of Technology. He holds a BS degree in Mechanical Engineering from North Carolina State University.