Project Report
ATC-427

Review of Systems-Theoretic Process Analysis (STPA) Method and Results to Support NextGen Concept Assessment and Validation

E.P. Harkleroad
A.E. Vela
J.K. Kuchar

25 October 2013

Lincoln Laboratory
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LEXINGTON, MASSACHUSETTS

Prepared for the Federal Aviation Administration,
Washington, D.C. 20591

This document is disseminated under the sponsorship of the Department of Transportation, Federal Aviation Administration, in the interest of information exchange. The United States Government assumes no liability for its contents or use thereof.

This material is based upon work supported by the Federal Aviation Administration under Air Force Contract No. FA8721-05-C-0002. Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Federal Aviation Administration.

TECHNICAL REPORT STANDARD TITLE PAGE

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
ATC-427		

4. Title and Subtitle	5. Report Date
Review of Systems-Theoretic Process Analysis (STPA) Method and Results to Support NextGen Concept Assessment and Validation	25 October 2013
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Eric Harkleroad, Adam Vela, and James Kuchar	ATC-427

9. Performing Organization Name and Address	10. Work Unit No. (TRAIS)
MIT Lincoln Laboratory 244 Wood Street Lexington, MA 02420-9108	11. Contract or Grant No. FA8721-05-C-0002 and/or FA8702-15-D-0001

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
Department of Transportation Federal Aviation Administration 800 Independence Ave., S.W. Washington, DC 20591	Project Report
	14. Sponsoring Agency Code

15. Supplementary Notes

This report is based on studies performed at Lincoln Laboratory, a federally funded research and development center operated by Massachusetts Institute of Technology, under Air Force Contract FA8721-05-C-0002.

16. Abstract

This report provides an assessment of the applicability of Systems-Theoretic Process Analysis (STPA) to perform preliminary risk-based modeling of complex NextGen concepts, based on the observed application of STPA to Interval Management – Spacing (IM-S) as a case study. The report also considers the potential use of STPA as a formal tool for safety analysis at the Federal Aviation Administration. This report's sources include a report documenting the application of STPA performed by the MIT Systems Engineering Research Lab (SERL), previous reports, and input from other staff and aviation subject-matter experts.

17. Key Words	18. Distribution Statement
	This document is available to the public through the National Technical Information Service, Springfield, VA 22161.

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified	53	

FORM DOT F 1700.7 (8-72) Reproduction of completed page authorized

This page intentionally left blank.

EXECUTIVE SUMMARY OF FINDINGS AND RECOMMENDATIONS

This report provides an assessment of the applicability of Systems-Theoretic Process Analysis (STPA) to perform preliminary risk-based modeling of complex NextGen concepts, based on the observed application of STPA to Interval Management – Spacing (IM-S) as a case study. The report also considers the potential use of STPA as a formal tool for safety analysis at the Federal Aviation Administration (FAA). This report's sources include a report documenting the application of STPA performed by the MIT Systems Engineering Research Lab (SERL) [1], previous reports, and input from other staff and aviation subject-matter experts (SMEs).

This section summarizes the findings and recommendations in this report and provides cross-references to additional information in later sections.

Finding 1: STPA supports qualitative safety analysis through identification of hazards and the Unsafe Control Actions leading to those hazards. STPA applies primarily to the early concept development phase; however, STPA could also help to ensure adequate coverage of the hazards included in later quantitative hazard risk/likelihood analyses. The Safety Management System (SMS) process allows safety professionals and Safety Risk Management (SRM) panels freedom to apply STPA or other methods of their choice as they see fit. Multiple methods may be used in concert. Sections 2.1, 3.2.1, 3.2.2.

Finding 2: STPA, like all other risk modeling methods, cannot by itself ensure that all safety issues are captured. The breadth and depth of analysis depends on the quality of information available about the system under analysis, the expertise of those carrying out the analysis, and the resources applied. Sections 2.1.7, 2.1.8, 2.1.9.

Finding 3: The key value of STPA is in providing a structured control-system–based framework for identifying Unsafe Control Actions and causal scenarios. This framework helps direct attention to all components and information flows that may affect safety. Sections 2.1, 2.4.

Finding 4: STPA does not provide a structured framework for deriving resultant safety constraints or requirements from causal scenarios. Deriving detailed requirements from the scenarios appears to follow a more ad hoc process similar to other hazard modeling methods. Sections 2.1, 2.4.

Finding 5: Multiple duplicative requirements were derived in the observed STPA application to IM-S. STPA's method for tracing requirements to hazards does not provide a direct way to identify duplication or contradiction between related requirements. Nor does STPA provide a direct way to inform or annotate the prioritization of requirements. Sections 2.4.1, 2.4.5, 2.4.6.

Finding 6: STPA has not been demonstrated to directly identify or isolate new requirements that are introduced as a result of changes to a system (i.e., operational improvements) relative to those that may have existed prior to that change. Section 2.4.4.

Finding 7: STPA's process of organizing and numbering requirements and their traceability to Unsafe Control Actions and hazards has not been fully refined and can be confusing. A more effective and consistent hierarchical taxonomy should be developed and standardized. Section 2.4.2.

Finding 8: STPA typically uses some terms that have different meanings than those conventionally used in the FAA's SMS process. Examples include definitions of hazards and the use of the term "controller." These were addressed in MIT SERL's analysis through special notes or changes in terminology. Section 3.2.

Finding 9: SMEs should review the safety constraints or requirements generated by STPA to ensure they are consistent and specific enough to inform safety-guided design. Furthermore, additional knowledge is required to fully develop control structure diagrams and necessary process models. The level of detail required will likely necessitate the use of SMEs. Sections 2.4.5, 2.4.6, 2.4.7.

Finding 10: Automated software tools for STPA analysis are under development but are not yet complete and do not yet support all steps of the process. An automated tool for identifying unsafe control actions has not yet been fully applied to IM-S. As such, usage of the tool in an automated fashion has not been validated for NextGen operational improvements. Section 2.1.9.

Finding 11: Necessary preparation for new STPA users to perform complete studies of complex systems includes several days of training but needs to be coupled with subject matter expertise gathered through months or years in performing safety analyses. Further study is needed to more fully understand STPA training requirements. Section 3.3.1.

This report makes the following recommendations:

1. The FAA should apply STPA as an additional tool in the concept development phase to aid in developing structured inventories of hazards and Unsafe Control Actions leading to those hazards. However, STPA is not a replacement for the existing SMS process or tools.

2. Before it would be recommended for complete adoption, STPA needs additional refinement. The FAA should continue to monitor and assess the value of using STPA in future concept development activities through feedback provided by SRM and SMS experts. Specific development steps include

a. **Refine STPA automation tools:** Properly developed automated tools that standardize the process will streamline analysis and enhance quality and completeness.

b. **Standardize the safety requirements generation step:** This improvement will increase confidence in the method and permit its use in applications requiring repeatable, standard analysis by analysts with similar training.

c. **Standardize the numbering and labeling taxonomy for STPA results:** This improvement will increase traceability and help stakeholders to better understand results.

d. **Develop a means to identify and track overlapping or conflicting requirements as well as new requirements generated by changes to a system.**

e. **Validate against system behavior:** For systems whose control structure is well-defined but whose behavior is not yet fully understood, investigate the effect of STPA-generated safety constraints. Comparison of system behavior with enforcement versus without enforcement of these safety constraints will help to confirm that they do indeed prevent hazards and accidents. For concepts early in their development, this comparison might be done using simulations.

f. **Assess and standardize STPA training requirements:** Interview a broad sample of researchers who have successfully applied STPA to understand training requirements more fully. The sample should cover a variety of organizations and researchers who have completed STPA studies of complex systems. Generate a standard curriculum for new analysts to learn and practice STPA on realistic complex systems.

This page intentionally left blank.

TABLE OF CONTENTS

This page intentionally left blank.

LIST OF ILLUSTRATIONS

This page intentionally left blank.

LIST OF TABLES

This page intentionally left blank.

1. INTRODUCTION

1.1 BACKGROUND

In 2012, Lincoln Laboratory was tasked by the Federal Aviation Administration (FAA) Office of Safety and Technical Training, AJI-312, to conduct a survey of risk-based modeling and analysis techniques to support NextGen concept assessment and validation [2]. The 2012 report considered several risk modeling methods at a general survey level, but did not have the opportunity to assess modeling performance in detail. The report identified System Theoretic Process Analysis (STPA)—a safety analysis method developed at MIT—as capable of identifying hazards and analyzing risk during a system's initial development process and recommended consideration of applying STPA to NextGen concepts early in their design [3].

An operational concept early in its development process presents both great challenge and great opportunity. A less developed concept can often lack rigorous safety and risk analysis, as fewer risk modeling options are available. However, the advantage of early safety evaluation is that systems at this stage have fewer design constraints and should offer greater opportunity for safety-driven design changes; major design changes become less feasible as development progresses. There is an ongoing need to identify tools, methods, and approaches for concurrent concept development and safety evaluation of NextGen operational improvements.

In order to demonstrate the potential use of STPA as a risk modeling tool for NextGen operational improvements in early concept development, the MIT Systems Engineering Research Lab (SERL) performed an initial study in 2013 applying STPA to a specific set of Air Traffic Management (ATM) concepts. The study's focus was interval management for spacing (IM-S), including ground-based (GIM-S) and flight deck-based (FIM-S) variants; these concepts were chosen to provide examples of STPA's treatment of concepts at different stages of maturity. In parallel, Lincoln Laboratory contributed technical and program management assistance and coordination for the study. The SERL study sought to increase understanding of IM-S and to serve as a vehicle by which the potential benefits of STPA could be assessed. Though a fully detailed analysis of IM-S was not possible given the available resources, the initial results provide safety-related design insights for a substantial subset of the relevant hazards and accident types. A technical report prepared by SERL staff—with input from Lincoln Laboratory and the FAA—presents the study results [1].

More generally, the study contributes to the FAA's efforts to assess and improve the risk modeling tools applied to the National Airspace System (NAS) today and in the future.

1.2 SUMMARY OF OBJECTIVES

This report, prepared by Lincoln Laboratory, provides an assessment of the applicability of STPA to perform preliminary risk-based modeling of complex NextGen concepts, based on the observed

application of STPA to IM-S as a case study. The report also considers the potential use of STPA as a formal tool for safety analysis at the FAA. This report's sources include the September 23, 2013 SERL STPA technical report [1], previous reports, and input from other staff and aviation subject-matter experts (SMEs).

The goal of this report is to provide interpretation and guidance for FAA risk modeling decisions. Though this report considers STPA from an objective technical perspective, as determined by staff not affiliated with SERL, it should be noted that there was a contractual relationship between Lincoln Laboratory and SERL. The assessments provided here should therefore be considered relative to the context in which this work was conducted.

As a follow-up to the 2012 Lincoln Laboratory survey, this report devotes additional coverage to STPA in order to inform FAA decisions. In addition to reviewing the STPA study's technical results, this report also considers STPA in an organizational context—specifically with regards to the FAA's concept development process, SMS, and Safety Risk Management (SRM) process that assesses risk and determines necessary mitigations [4,5]. This report seeks to highlight and note benefits of STPA, required inputs and resources, and possible implementation issues. This report also considers guidance from the NASA Standard for Models and Simulations, which describes the information needed to support critical decisions based on modeling and simulation results [6,7,8].

1.3 PREVIOUS WORK

Researchers at SERL and other organizations have applied STPA to safety questions in different domains, including air traffic management, space and missile systems, robotics, water quality, food quality, and pharmaceuticals [3,9,10,11]. A common property of all of these systems is that they can be represented by control structure models that include technical, human, and social or organizational elements. Note that STPA studies cited above are not necessarily formally endorsed by safety regulators or professional groups. Furthermore, existence of these studies does not imply any particular level of acceptance of STPA but may reflect an active interest in the method among the safety community.

The 2012 Lincoln Laboratory report recommended applying a structured and repeatable hazard identification and assessment procedure. The report noted that STPA appears to provide such a framework for identifying both risks in a system and constraints needed to mitigate those risks. STPA can help expose risks because it provides a method for identifying potential hazard causes, which are known as "unsafe control actions" in the method's nomenclature. Furthermore, STPA can also help to identify mitigations because it provides a means for translating such hazard causes into safety constraints.

2. TECHNICAL ASSESSMENT

2.1 GENERAL REVIEW OF STPA

Before discussing results for a particular application, it is helpful to review the general properties of STPA. Explanations and discussions of STPA, including motivation, analysis steps, and application examples, are available in sources provided by its creators and users [3,9,10]. To summarize, STPA's essential approach is to define risk assessment as a top-down control problem rather than the more common approach of defining a bottom-up component reliability problem, like when applying Failure Mode and Effects Analysis (FMEA). Development of STPA has been motivated primarily by

- distinctions between safety and reliability,
- limitations of the probabilistic assessment of accident causal chains, and
- the roles of software and human factors [3].

Figure 1 summarizes the STPA analysis process, which includes

- defining accidents and the hazards which can lead to them,
- defining the system goals and control structure (system-level analysis),
- identifing unsafe control actions (UCAs),
- identifing causes of UCAs (causal analysis), and
- identifing necessary safety constraints and requirements [1,3].

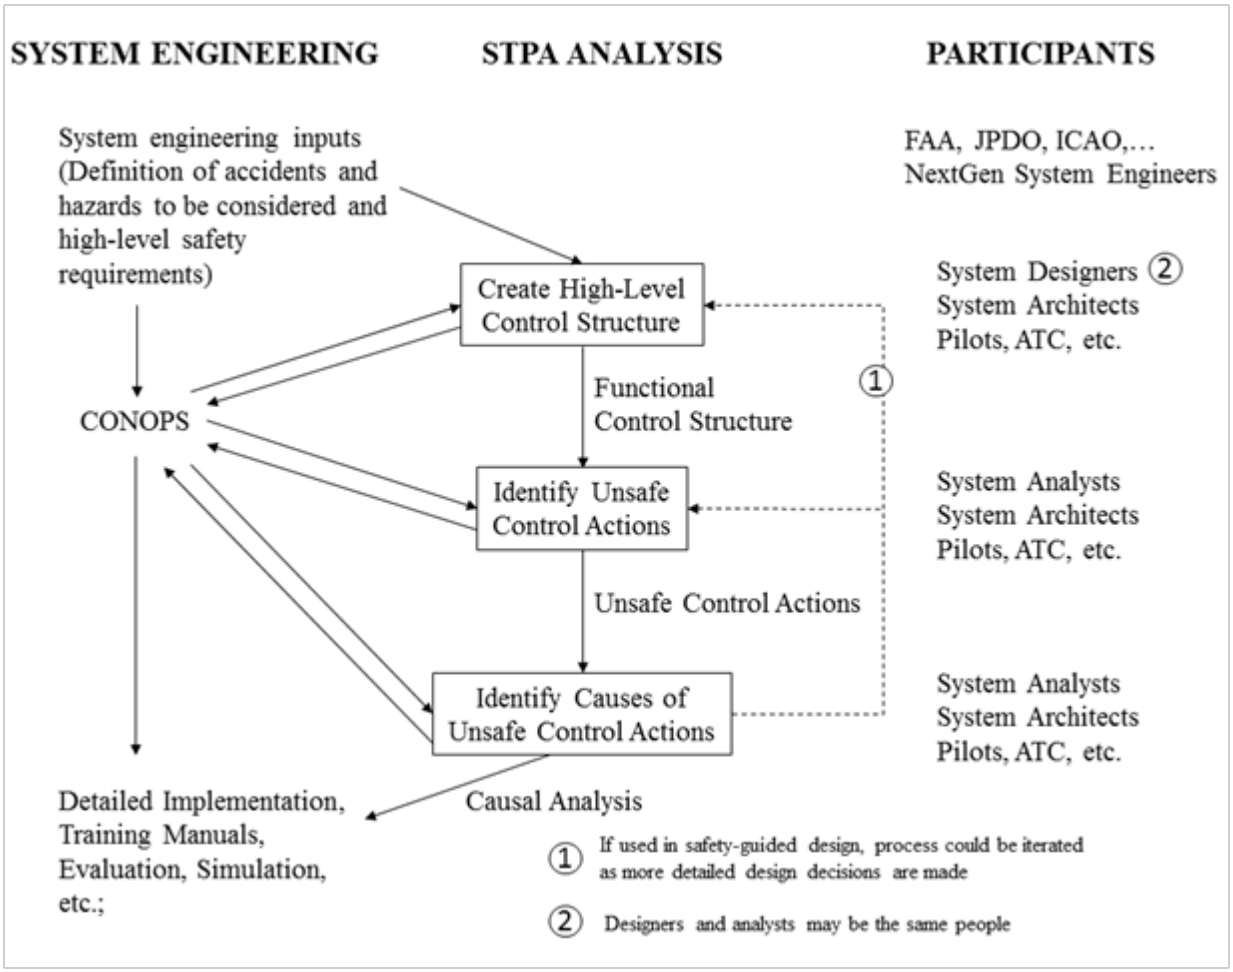

Figure 1. STPA process (from [1]).

The STPA process starts with a description of the system to be analyzed, including relevant system components and agents, their available actions, and how these actions affect other components and agents. This description could be provided by documents or other input from experts on the system; it is used to build a control structure diagram depicting the control relationships within the system [1,3].

STPA must also start with a definition of unacceptable accidents or losses and a list of situations likely to lead to them; these unsafe, pre-accident situations are defined as hazards in STPA. For example, loss of separation (LOS, considered as a hazard) could lead to a mid-air collision (considered as an unacceptable loss). Unacceptable losses could be defined broadly or narrowly depending on a study's goals; a narrower definition might include only fatal accidents, while a broader definition might also

include non-fatal (injury-producing) accidents, non-injury accidents, or less severe incidents (e.g., loss of separation). The results of STPA studies are sensitive to the assumed definition of unacceptable losses.

The next step is to identify the unsafe actions or inactions that can lead to hazards. This step identifies not only the action itself, but also the agent performing it and the relevant context—including timing, system state, and the information available to the agent. A set of four general categories of unsafe control actions helps to guide this step:

1. Not Provided when required for safety
2. Providing Causes Hazard
3. Too soon, too late, or out of sequence
4. Stopped too soon, applied too long [1,3].

Note that these categories include both actions and inactions.

Given a list of unsafe control actions, the process next determines their potential causes and associated contextual factors (causal scenarios). This step considers potential causes at any point in the control structure, including flaws in the agent's situational awareness and inadequate operation of control mechanisms or actuators. The unsafe control action, cause, and contextual factor define an unsafe scenario.

The final step determines what constraints are needed to prevent the identified unsafe actions from occurring for each unsafe scenario; these constraints can also be interpreted as requirements on system behavior. The safety constraints and requirements identified by STPA assist with safety-guided design of the system. The safety constraints and requirements can be cast as inversions of each unsafe scenario. That is to say, the requirement provides a solution to prevent a cause from ever occurring or a contextual factor from manifesting.

Figure 2 places STPA's role in the context of the overall safety analysis process during a system's lifecycle. Note that the method is most applicable in the early lifecycle stages but can help to improve safety during any stage. The FAA SRM process also covers some of the same lifecycle stages and hence there are likely opportunities for STPA to enhance that process.

In the aviation domain, STPA has previously been applied to a variety of systems and new concepts, including the In Trail Procedure (ITP), Coast Guard helicopter operations, avionics, and others discussed in recent conferences and publications [3,9,10,11]. For these applications, STPA identified necessary safety constraints and unsafe scenarios involving confusion between automation modes, conflicting data sources, and latency in operator commands and feedback; these factors may cause hazards that must be mitigated to prevent accidents. STPA studies have also identified shortcomings in organizational safety culture as important factors preventing safety progress; for example, information

sharing may be inadequate or accident prevention efforts may focus only on certain accident types, leaving others unaddressed.

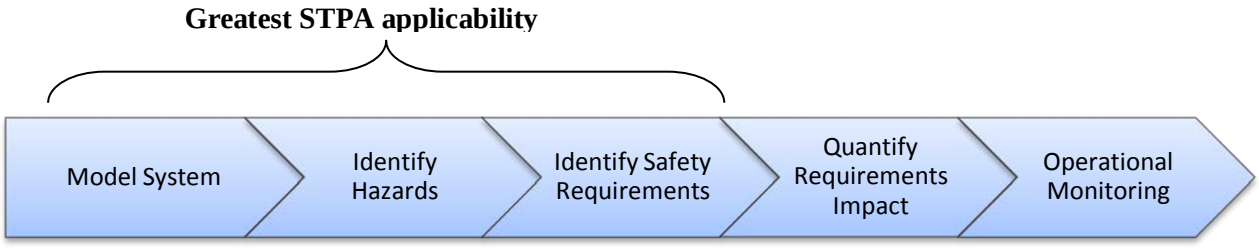

Figure 2. STPA applicability to lifecycle system safety analysis.

The remainder of this section discusses a few key characteristics of STPA—including both potential advantages and potential limitations. Characteristics of STPA are not described in full here because its creators and users have documented them elsewhere [1,2,3,9,10]. However, this section summarizes a few key examples relevant to this report. Note that some characteristics discussed in this section are not unique to STPA and also apply to other safety methods; these issues need to be addressed regardless of which methods are used.

2.1.1 Early Concept Assessment

STPA is designed to provide safety insights earlier in concept development than is appropriate for probabilistic, data-driven assessments. STPA does not require system performance data from simulations or testing; it requires only a description of a system's components and their control relationships. This ability enables analysis of systems which are not yet designed in enough detail to simulate or test.

2.1.2 Variety of Causal Mechanisms Identified

STPA is designed to identify not only direct, linear event sequences leading to accidents, but also indirect or time-delayed causal chains [3]. For example, the influence of management or organizational culture, unexpected automation behavior, and gradual system changes over time may enable accidents that would otherwise be prevented. An ability to identify these issues is important for complex systems such as Air Traffic Control (ATC) automation in which cause and effect may be widely separated in time and operator perception and control functions may be managed by software.

STPA can also consider unsafe actions that are both acts of commission, such as "providing a speed modification is hazardous when…" and acts of omission, such as "not providing a speed modification is hazardous when…" Studies that consider only acts of commission may overlook relevant causal factors. Both types of actions are considered in the IM-S STPA study.

2.1.3 Traceability

An analysis using STPA is intended to be traceable and updatable. The hazards, unsafe control actions, and safety constraints identified in STPA's top-down analysis are numbered so that the results can be updated with reasonable effort if the system under analysis changes or evolves. The analysis does not necessarily need to be completely redone after system changes [3]. This property is important for ATC systems, which tend to be long-lived, may evolve through several versions, and may need to interoperate with other, newer systems later in their service lives.

2.1.4 Human Factors

One potential limitation raised by STPA users is the method's coverage of human factors and operator behavior [10]; it is not yet clear that the method adequately addresses "cognitively complex" errors such as those related to human supervision of automated processes [1]. Because human behavior is not deterministic, constraints on it may need to be more carefully designed than constraints on hardware or software. Furthermore, there is a need to include process models that represent a human controller's mental model of the phenomena they are observing and controlling; an incomplete or inaccurate description of the controller's mental model may lead to deficiencies in any safety analysis. SERL staff has discussed the role of human operators in previous work, but further research may be needed to define flexible constraints on operator behavior [3]. These should not solely depend on checklists and should permit creative problem solving in situations not foreseen by system designers and engineers.

2.1.5 Dynamic, Real-Time Safety Constraints

STPA users have also raised concerns related to treatment of dynamic relationships between agents and the need for real-time safety constraints [10,12]. STPA is able to provide safety constraints that are conditional on the system state, but additional steps may be needed to consider systems with dynamically changing system control structures; an example would be a system where longer-term strategic and short-term tactical control are handled by different decision-makers. SERL staff is aware of this concern and plan to address it in future work [1]. Augmentation of STPA with additional tools is a possible solution, and a Boeing/NASA study discusses research on this idea [12].

2.1.6 Probabilistic Risk Assessment

STPA does not provide a probabilistic assessment of system risks, which the FAA Safety Management System requires for assessment of new operational concepts [4]. However, such a quantitative assessment is not the intended product of STPA, which seeks only to identify unsafe control actions and the constraints required to prevent them from causing accidents. STPA can contribute to safety-guided design by helping to ensure performance requirements and safety constraints are consistent, but it should not be expected to estimate accident rates or likelihoods.

STPA is complementary to probabilistic risk assessment and helps stakeholders understand safety risks in a different way. Each approach serves its own important purpose and its results should be interpreted in light of its particular strengths and limitations.

The output of probabilistic approaches has uncertainties that cannot be completely eliminated and so are typically bounded within an acceptable margin. If probability measures for each hazard are provided, then some mechanism for quantitative validation is required. While identifying probabilities in simple cases of cause-and-effect may be straightforward, the coupling of control processes brings with it conditional probabilities that are typically much more difficult to determine. For example, given several possibly interdependent inputs to a system, can the probability of an unsafe situation be easily represented? If not, then this probability must be estimated using a conservative method to avoid underestimating risk. Even with such conservative methods, the possibility of underestimating risk remains. STPA avoids this issue by not estimating probabilities at all; it instead focuses on identifying means for preventing or mitigating all possible accident scenarios regardless of likelihood.

2.1.7 Dependence on Expert Input

STPA depends on a description of a system's control structure provided by SMEs, which may be incomplete. However, such incompleteness should not prevent STPA from determining safety constraints to improve the system's design; the method is designed to find flaws such as an inadequate control structure or missing system requirements. This dependence on SMEs is not unique to STPA and applies to all methods except those that provide formal, mathematical proof of complete and correct results. As discussed in a review of validation methods, such definitive, provably valid results are typically possible only for the simplest systems [13].

2.1.8 Dependence on analyst training and experience

As with any analysis method, it is possible to apply STPA incorrectly and obtain incomplete or incorrect results. The method is not completely automated and leaves room for discretion on the part of the analyst and SMEs. For example, the analyst may fail to identify certain unsafe control actions, overlook possible causal factors, or generate incomplete safety requirements. These concerns can be mitigated by ensuring that analysts applying STPA are properly trained, that they have access to SMEs, and that they use automated analysis support tools when possible.

2.1.9 Completeness

Showing that a safety analysis comprehensively and completely covers all possible hazards may not be possible, but several properties of STPA help to address the concern. One of these is the method's top-down refinement of broad hazards (such as loss of aircraft separation); completeness is ensured if the initial, high-level list of hazards is complete and each refinement step breaks them into collectively exhaustive subsets. Verifying completeness of this top-down process is ultimately the responsibility of

SMEs, and STPA's hierarchical approach assists by providing traceability of low-level, detailed causal mechanisms back to their high-level parent causes and hazards [3].

STPA also helps to ensure completeness via a systematic approach for identifying unsafe control actions. The developers of STPA have also created an optional software-driven process that aims to automate this process; this automated process is in development and cannot yet be fully evaluated. The software-generated list of candidate unsafe control actions requires review by SMEs to eliminate those that are inapplicable, impossible, nonhazardous, or otherwise not safety concerns; however, systematic generation of a variety of unsafe control actions reduces the possibility of overlooking accident mechanisms and underestimating risk. The STPA study report provides an example of applying this process to ATC control actions related to IM-S [1]; the example shows how a discrete set of variables can describe the source, type, and context of possible control actions. The joint set of possible values of these variables provides a combinatorial expansion of the set of possible actions; this large set can sometimes be reduced to a smaller set—for example, any action leading to loss of separation is hazardous regardless of weather or terrain. Furthermore, the resulting set of actions can be represented in a compact table. This simplification and table representation help to ease expert review without sacrificing fidelity of the results [1]. An automated process is described briefly in the IM-S analysis report, but it was only applied to demonstrate STPA's systematic approach on a subset of the analysis [1]. As such, the benefit of the tool for understanding NextGen operational improvements deserves further study but cannot yet be assessed nor validated.

2.2 REVIEW OF STPA RESULTS ON GIM-S

This section reviews the observed contribution of STPA to safety analysis of the GIM-S concept. Under GIM-S, a ground-based decision-support tool provides speed advisories to air traffic controllers, who may issue these advisories to aircraft in order to achieve spacing goals [14]. Applying STPA to assess GIM-S has demonstrated that the method is able to produce safety constraints on system behavior [1]. However, additional feedback from SMEs is needed to understand the added value of the results (in comparison to previous GIM-S safety studies) and whether additional analysis via STPA should be pursued. Ultimately the value of STPA can only be understood in direct comparison to current methods.

The STPA study identifies a general, high-level principle to guide safe system design: prioritization of safety over IM-S spacing goals and operational efficiency. This principle should guide ATC personnel training and ATC display design related to IM-S. For example, personnel must be trained to prioritize resolving immediate conflicts above issuing speed advisories. Also, automated conflict alerts and other safety information should be displayed more prominently than speed advisories in ATC displays. The study makes this recommendation implicitly in its GIM-S safety constraints ([SC-1] through [SC-9]) and explicitly in several requirements [1]:

- STPA-G.1S.2.1.1 ATC must give priority to immediate conflict over any downstream demands. (Allocated to: ATC, Conflict detection tools, Interface layout)

- STPA-G.1S.3.1.1 ATC must issue conflict resolution clearance (not limited to speed modification) even if Traffic Flow Management (TFM) cannot calculate a valid speed advisory. (Allocated to: ATC, FAA Procedures)

- STPA-G.1S.3.2.1 ATC must provide appropriate action in the event of an unsafe TFM advisory. (Allocated to: ATC, FAA Procedures)

The STPA GIM-S results identify 43 unsafe control actions across the four general categories and provide nine high-level safety constraints to mitigate them [1]. The high-level constraints are each expanded to generate between 30 and 60 more specific requirements; the full results would thus include at least several hundred specific requirements. These totals may vary in future updates of the analysis. These results may also overstate the true minimally sufficient number of requirements; as noted later in Section 2.4.6, some requirements overlap because they are derived from related unsafe control actions and hence may share a common mitigation. If SMEs refine the requirements after the STPA process, the total number could increase or decrease.

The next several subsections review the GIM-S results in greater detail.

2.2.1 Common Situational Awareness Requirement

One safety constraint identified by the STPA study is that different agents involved in IM-S operations must have consistent information, including consistent surveillance and flight plans; these agents include flight crews, IM-S automation, and ATC personnel [1]. Documents describing the system's operation indicate that IM-S automation may not always receive updates such as flight plan changes, vectoring by ATC, and other clearances; such incomplete information sharing may cause misunderstandings and hazardous situations. The study identifies two methods of enforcing this safety constraint: (1) allow IM-S automation access to updated flight plans and active clearances to ensure its situational awareness is consistent with that of ATC personnel (STPA-G.1S.3.3.1, STPA-G.1S.3.4.2), and (2) provide training to allow ATC to recognize and reject inappropriate speed advisories (STPA-G.1S.1.11.1, STPA-G.1S.2.1.1). These solutions should be prioritized in this order because preventive design changes are always preferable to training or warnings that do not enforce safe behavior; Section 2.4.5 discusses this topic further.

2.2.2 Human-Automation Interactions

The study also identifies unsafe human-automation interactions as a potential issue; these may include, for example, overreliance on speed advisories or ignoring advisories. Such concerns should be investigated fully and may need attention from human factors experts. Overreliance may become a concern immediately or over time through operator adaptation to consistently reliable and valid speed advisories [1,14]. For example, en route ATC might mistakenly believe that speed advisories are verified by automation to be conflict-free or that the lack of a speed advisory implies no change in an aircraft's speed is needed; in fact, both of these statements are inconsistent with the IM-S Concept of Operations

(ConOps) document, which states that ATC retains responsibility for ensuring separation of aircraft at all times [14]. In general, the system must be correctly designed—and operators correctly trained—to ensure there are no misconceptions about GIM-S functions or outputs.

Identification of causal factors related to human interaction with IM-S automation provides an example of the human factors capability of STPA; however, the study could benefit from more in-depth consideration of human factors. For example, a human factors expert may be able to suggest additional types of unsafe control actions, causal factors, errors in mental process models, or mitigations related to operator behavior.

More generally, while human factor considerations are identified as a potential concern for STPA users (see Section 2.1.4), observation of STPA applied to IM-S suggests that it can in general identify human factor issues. STPA is a general procedure for analysis of all types of causal factors, including human and human-computer interactions. Regardless, STPA could benefit from explicit inclusion and formalization of human factors principles to help clarify its structured procedure.

2.2.3 Need for Expert Feedback

Additional feedback from SMEs would improve the study results. The study benefits from input already received from IM-S experts and program managers on details of system behavior and control relationships; this information was provided through documents and live meetings.

SERL and Lincoln Laboratory have also requested expert feedback on the completed STPA outputs—specifically whether they are reasonable and yield insights to improve safety of the system's design and operations. Though this feedback is not a formal part of the STPA process, it is a recommended best practice for validation of the results. Such feedback would also illustrate how STPA itself is received and accepted by a sample of the aviation safety community; this insight would inform any future effort to formally include STPA in FAA safety processes—an idea discussed later in this report.

2.3 REVIEW OF STPA RESULTS ON FIM-S

This section reviews the contribution of STPA to safety analysis of the FIM-S concept. Under FIM-S, air traffic controllers may issue new types of clearances related to achieving and maintaining pairwise separation; flight crews will execute these clearances using speed advisories from onboard flight management systems rather than advisories from air traffic controllers.

Applying STPA to FIM-S has yielded safety insights in the form of necessary constraints on system behavior. This analysis overlaps somewhat with the GIM-S results, but also identifies additional scenarios unique to FIM-S that should be addressed through design changes or mitigations. As with GIM-S, additional feedback from SMEs will improve the results and should be pursued.

Similar to GIM-S, a FIM-S safety constraint identified by the STPA study is that different control agents must share information to ensure consistent and common situational awareness and process models. For example, coordination between and within ATC facilities is necessary to avoid conflicting FIM-S clearances. One illustrative situation involves two different aircraft assigned spacing goals that are individually safe but later conflict when the aircraft cross paths or merge to a common stream (see Section 3.2.2 in [1]). The study provides an example of controllers in separate facilities assigning such ultimately conflicting clearances, but a similar interaction could occur between adjacent sectors in the same facility or within a single sector; in each case, such a conflict is caused by inadequate coordination and information sharing. These scenarios could be mitigated in several ways, including (1) providing cross-facility access to active FIM-S clearances in ATC displays along with appropriate training (STPA-F.15T.3.4.1, STPA-F.15T.3.4.2, STPA-F.15T.3.5.1), or (2) implementing automated checks of FIM-S clearances for conflicts (STPA-F.15T.1.6, STPA-F.15T.3.2.2). Even if automated conflict checks were included in the FIM-S concept, ATC personnel likely would still retain ultimate responsibility for issuing safe, consistent clearances.

Because FIM-S is a less mature concept than GIM-S, the STPA results are necessarily less specific. As the IM-S concept develops further, the STPA analysis results could be updated to reflect any changes in its assumptions or control structure.

2.4 STPA CHARACTERISTICS OBSERVED IN IM-S CASE STUDY

The characteristics discussed in Section 2.1 might or might not be observed in practice while applying STPA. Therefore, to provide a proper assessment regarding applicability to NextGen operational improvements, this section identifies which of them were in fact demonstrated during the IM-S analysis and which were inconsequential. Only the most important characteristics are discussed in this section, including those where new information about STPA was obtained through observation of the analysis. This section also briefly discusses STPA capabilities that were not demonstrated in the current study but that are claimed by STPA's creators or demonstrated elsewhere.

2.4.1 Hierarchical Analysis Structure

Figure 3 summarizes the hierarchical structure of the STPA analysis starting from the list of identified unsafe control actions; each additional step of the analysis expands the hierarchy to additional branches. The analysis along different branches is largely independent but not always; for example, UCAs may be produced by one causal scenario or a combination of them. This hierarchical structure assists with traceability of the process. However, due to the geometric growth in the number of elements when moving from Unsafe Control Actions to Causal Scenarios to Safety Constraints, a robust method for organizing and cataloging each item and its position in the hierarchy is needed.

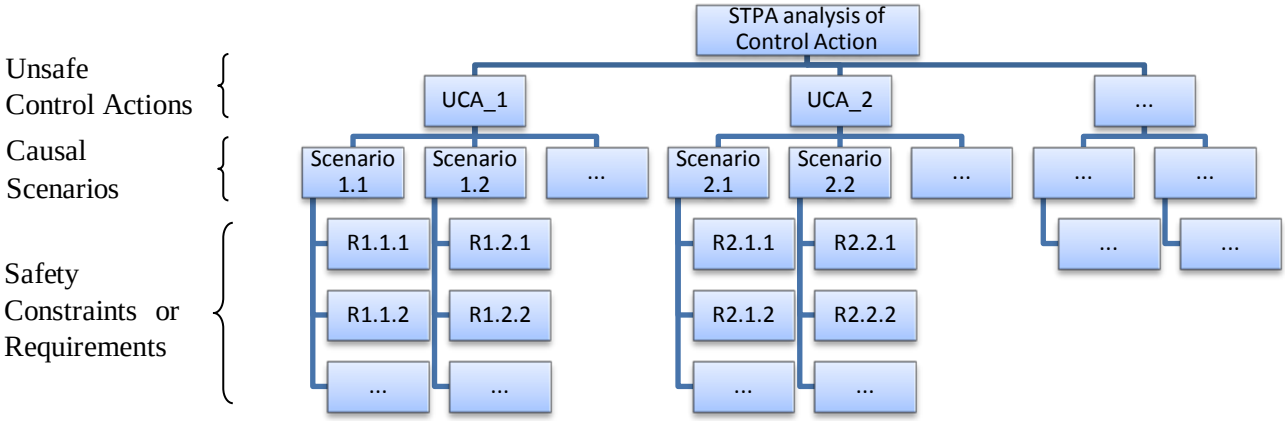

Figure 3. Hierarchical structure of STPA process results.

2.4.2 Observed Forward and Backward Traceability

To understand the usefulness of STPA in identifying hazards, causes, and safety requirements, the forward traceability of its outputs (from Unsafe Control Actions to their resultant safety constraints or requirements) must be clear and direct. That is to say, the casual scenarios should result from the unsafe control actions in a direct fashion, without confusion. Likewise, generating safety constraints from the causal scenario should be logical and not haphazard. If forward traceability can be demonstrated, it shows that, when applied correctly, the STPA method itself guides the analysis—in contrast to a gifted analyst performing an ad hoc, unstructured safety analysis (based on unique experience sets that emphasize some scenarios and neglect others). Furthermore, demonstrating forward traceability supports the argument that only a specific level of training is required to perform STPA effectively when provided with the proper documentation and information.

Forward traceability of an analysis is related to repeatability because a fully traceable process can be easily repeated and verified. It is important that the STPA process be standardized and repeatable so that different analysts with the same training on STPA would produce nearly the same results; this property is important for formal adoption of the method and repeated use across different applications.

In the case of the IM-S analysis, these properties have been largely demonstrated for the first several steps of STPA, but not for the final step of generating safety constraints and requirements. It is not clear that the complete analysis is fully traceable and could be repeated by anyone with training on STPA.

A hazard analysis method should also possess backwards traceability. Here backward traceability refers to a clear tracing from a requirement to an unsafe scenario to the associated unsafe control action.

Backward traceability allows for designers to understand the genesis of a safety requirement and why it is required to protect against a specific unsafe scenario.

Any discussion of traceability must include with it an understanding of scalability: how the number of unsafe control actions, unsafe scenarios, and safety constraints grow with increasing detail in the control structure model. As shown in Figure 3, the number of requirements grows geometrically with each unsafe control action associated with each controller identified in the control structure model. As systems become increasingly complex, traceability allows for a clear and understandable means to relate the safety constraints to contextual factors and unsafe control actions, and vice versa. Likewise, it is useful to understand how multiple safety constraints relate to each other, even when their parent unsafe control actions, unsafe scenario, or corresponding controllers are different; this issue is discussed in Sections 2.4.6, and 2.4.7.

Repeatability of Initial STPA Steps

Given the proper analyst training, earlier STPA steps appear to be repeatable and systematic; these include defining the control structure, determining unsafe control actions, and identifying certain causal mechanisms related to deterministic systems. Starting from a simple initial set of hazards and control diagrams—like that in Figure 7 or Figure 8 of the SERL technical report [1]—breakdown of individual elements can produce greater detail and a more complex system to analyze. Figure 9 of the report illustrates the result of breaking down the original control structure in Figure 8 of the report [1]. Given this information, the unsafe control actions can be generated in a clear and repeatable manner. The repeatability of the method is a product of the top-down approach and of the four generic types of unsafe control actions, which guide the process and appear to be complete.

Repeatability of Requirements Generation and Need for Expert Input on Requirements

The process of translating unsafe control actions into safety constraints and requirements does not appear to have been standardized within the STPA process. The translation appears to require specific domain knowledge of the system under analysis beyond what is included in the control structure identified early in the STPA process, and it was observed to follow a somewhat ad hoc process when applied to IM-S.

For example, the IM-S analysis produced certain requirements not directly traceable to the IM-S control structure—including a requirement that IM-S-related communications not interfere with other communications: STPA-G.1S.1.7.3 and STPA-G.2S.3.3.1 in [1]. In fact, it is not clear that these particular requirements resulted from the STPA analysis of GIM-S as there was no detailed modeling of communication systems or radar systems to support their existence. In particular, STPA-G.2S.3.3.1 was reappropriated from previously documented requirements for the Traffic Alert and Collision Avoidance System (TCAS). While these requirements may be logical, without a proper and detailed model of communication systems in STPA, they could only be specified in detail by an expert in these systems. Similarly, it is not clear where requirement STPA-G.1S.1.8.1, “Automatic Dependent Surveillance-

Broadcast (ADS-B) must provide 0.1/0.3 NM (95%) accuracy," was derived from or whether those values are truly required for GIM-S operation.

When such specialized requirements are needed, SMEs may need to assist in drafting them or to develop highly detailed functional control models and process models to support the STPA analysis. Without this knowledge, different analysts with the same training in STPA would not always generate the same set of requirements. While this conclusion may also be true for other safety analysis methods, the goal of traceability and repeatability should be pursued whenever possible.

Backward Traceability

In the final IM-S analysis results, traceability was provided by numbering the elements in each step of the analysis. Due to the hierarchical nature of the analysis, a hierarchical numbering system proved most useful and should be used to organize the results. For example, unsafe control action X might have causal mechanisms labeled X.1, X.2, X.3… and so on.

If different unsafe control actions are related or naturally grouped, their numbering or labels should be similarly grouped. Parent-child numbering should also reflect the relationships between the UCAs themselves. For example, unsafe control actions Y1 and Y2 might involve performing the same action too early versus too late; their possible causes could be labeled Y1.1, Y1.2… and Y2.1, Y2.2… with as many numbered causes as needed. If two UCAs are not related to each other, they should use distinct numbering or labels to make this clear.

The current numbering scheme used in STPA for the IM-S analysis has caused some confusion. In the STPA analysis, UCA1 through UCA11 refer to Unsafe Control Actions (UCAs) when providing (or not providing) an air traffic control clearance. For example, UCA1, UCA2, and UCA3 have been assigned to Unsafe Control Actions for not providing a clearance when one is required for safety (Table 2 in [1]). UCAs related to speed advisories (which are one form of clearance) are labeled UCA1.S through UCA11.S. Those related to vector (heading) clearances are labeled UCA1.C through UCA11.C, and so on. UCA1.C is further broken up into UCA1.C.a and UCA1.C.b, but it appears more logical to swap the order such that these would be represented by UCA.C.1.a, UCA.C.1.b. In addition, use of numbered labels is typically interpreted to imply a sequence, but there does not appear to be a sequence in this case. UCA1 through UCA11 are defined, but UCA1.S through UCA11.S do not include UCA3.S (undefined). UCA1 through UCA3 are related to not providing a clearance, and UCA4 through UCA6 are related to providing a clearance, but this distinction is not visible from the numbers alone. A better method of being able to directly interpret different types of UCAs should be developed. This change would assist stakeholders in understanding the relationships between UCAs, causal scenarios, and requirements.

Backward traceability is required to understand how causes and contextual factors are associated with an unsafe scenario. According to STPA developers, the generic control loop in Figure 4 guides the generation of unsafe scenarios [1,3]; each part of the control loop represents a possible failure point to create an unsafe scenario. As a system designer, if an unsafe scenario is presented in a safety analysis, it

would be useful to have clear understanding of where in the feedback loop the initial failure is located. STPA currently does not have a documentation procedure for clearly labeling where in the loop the failure occurs, and what are the associated components (controller, actuator, controlled process, or sensor). Documenting failed components or failed portions of the feedback loop may allow designers to find common linkages between unsafe scenarios that result in different unsafe control actions.

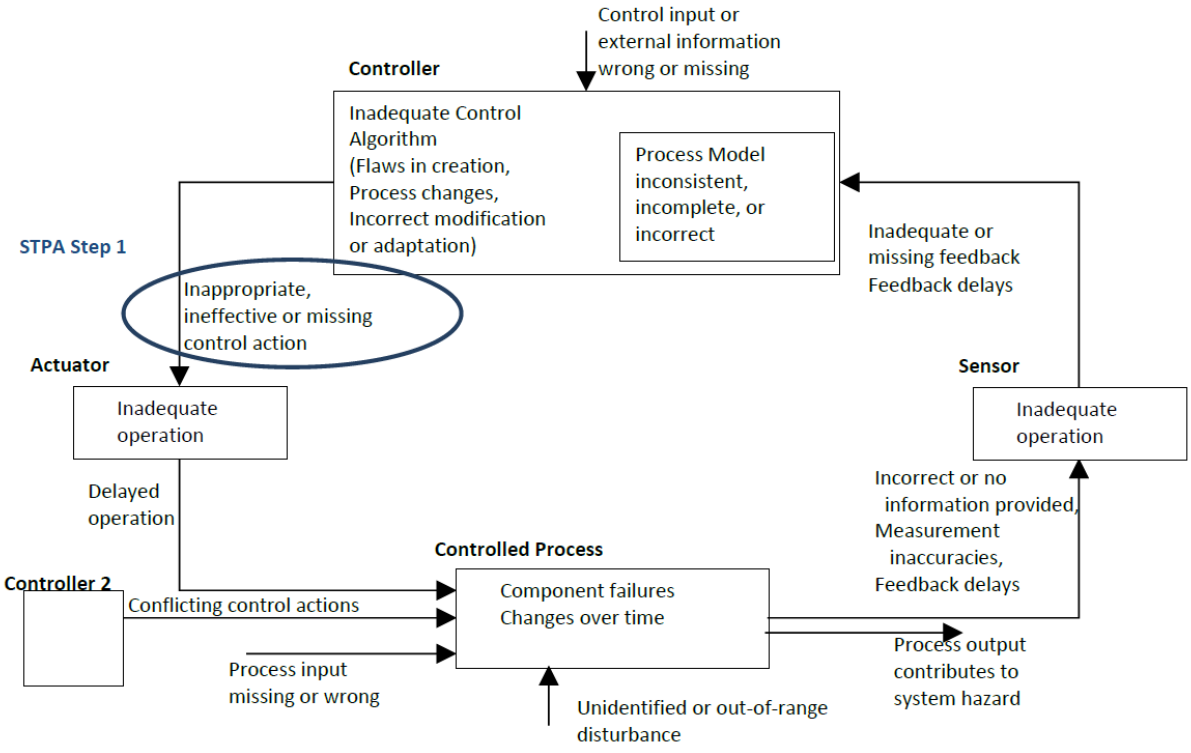

Figure 4. General control loop with causal factors [1].

2.4.3 Observed Completeness

As previewed in Section 2.1.9, the structure of STPA is observed to support development of a complete analysis of IM-S. The scope of the analysis, which was limited in 2013 by available resources, may need to be expanded in future work on FIM-S, which appears feasible given additional resources.

However, it does appear that some causal scenarios were not uncovered during the STPA analysis. One example is a case in which new information introduced by GIM-S into ATC displays obscures other information; GIM-S information could include speed advisories or stale indicator symbols. Such concerns

are implicitly included in higher-level requirements such as "STPA-G.1S.1.7.4: The IM/TFM system must not interfere with existing ATC systems or procedures," but it is not clear that such a requirement is specific enough. Additional scenarios of this type would likely be uncovered given additional SMEs and resources. Such examples underscore that STPA by itself cannot guarantee a complete safety analysis or sufficient requirements specificity.

The four generic unsafe control action types guide the analysis and appear complete at face value, but their application can be ambiguous because different unsafe control actions can overlap. As an example, consider two different unsafe control actions from the IM-S analysis [1]:

- UCA1. Not providing a clearance is hazardous if current trajectory leads to loss of separation.
- UCA7. Providing a clearance to an aircraft that was previously safe if given too late after a different clearance has been executed by same or other aircraft.

In an extreme case, suppose there is an imminent conflict requiring ATC to provide a new clearance within a few minutes; if the new clearance is delayed until after the conflict has already occurred and cleared, then ATC has essentially not provided the clearance at all. In this case, UCA1 and UCA7 appear to overlap, and it is not clear which of them actually describes the situation. When unsafe control actions overlap in this way, the result is typically a repetition or duplication of scenarios. Hence the overlap is noteworthy but does not imply a shortcoming in the analysis in terms of completeness.

Ensuring analytical completeness requires not only verifying coverage of the assumed system scope, but also validating that the scope is appropriate. One benefit of a scalable and traceable analysis is the ability to apply a step-wise approach to considering multiple interacting systems and the ability to expand the analysis scope with the addition of new agents, elements, or details. The top-down approach of STPA should permit scalability of the analysis in depth and breadth. According to the required depth of analysis, each element in a functional control structure can be decomposed. The structured process of decomposing each element ensures no key sub-elements or connections between them are missed. The reverse, bottom-up approach appears to be more likely to miss interactions between systems, especially when their level of detail differs.

The scope of the STPA analysis of GIM-S was limited to en route air traffic control and pilots and appears complete. In the case of FIM-S, there was no exploration, in this initial analysis, of aircraft flight management systems or cockpit display systems, which would require significantly more effort.

Given proper inputs and execution, STPA analysis appears to provide sufficient detail and completeness to aid designers during concept development and validation; it should be sufficient to contribute to safety-conscious NextGen operational improvements. Note that the thoroughness of any STPA analysis must ultimately be confirmed by SMEs and other applicable stakeholders.

2.4.4 Distinction between Preexisting and New Requirements

One additional issue relates to identifying UCAs, scenarios, and requirements that are introduced due to a change in a system, compared to those that may already be preexisting. For example, GIM-S introduces new information to be provided to ATC and new procedures for using that information, but no changes to communications, navigation, or surveillance (CNS) systems. A series of requirements may be derived from STPA that relate to general safety concerns with CNS but which are not specific to GIM-S. One specific example is requirement STPA-G.1S.1.8.1, "ADS-B must provide 0.1/0.3 NM (95%) accuracy" [1]. Identifying such a requirement may be appropriate for STPA in general, which seeks to identify all risks. However, given a need to expose and focus only on those new requirements levied due to a system change, STPA would not provide a direct way to uncover just those new requirements. Additional effort would be required to review the output of STPA and categorize each requirement as preexisting or new.

2.4.5 Requirements Alternatives and Safety Order of Precedence

When generating safety constraints, it would be beneficial for STPA to identify multiple alternative mitigation techniques for the same unsafe control action or causal mechanism. In some cases, safety requirements that dictate changes in a system's design are infeasible; instead, it may be necessary to provide warning devices or procedures to mitigate risk. For example, the IM-S study results in [1] suggest that potential loss of common situational awareness—discussed in Section 2.2.1—can be mitigated through some combination of design changes and operator training (STPA-G.1S.1.11.1, STPA-G.1S.2.1.1, STPA-G.1S.3.3.1, STPA-G.1S.3.4.2). If certain design changes are infeasible, designers could instead require additional operator training. Such alternatives are useful for system designers facing a variety of priorities which must be traded against each other.

In a more general context, some mitigations are able to prevent hazards entirely, while others may only be capable of reducing their risk. This idea of a mitigation hierarchy appears in the safety order of precedence shown in Table 1 as given in a Department of Defense standards document [16]. This framework prioritizes mitigation techniques. In the best case, system designs are revised to prevent hazards from occurring. When design changes cannot be made, use of safety devices is considered the next best option. The choice of mitigation techniques will depend on which parts of a system's control structure can or cannot be altered; for example, legacy systems and procedures may be less flexible than newer systems.

The linking of seemingly different safety constraints may also allow for a single mitigation to prevent multiple unsafe scenarios. That is to say, an upstream solution may provide appropriate mitigation for multiple downstream unsafe scenarios, despite the fact that the unsafe scenarios are not linked by an unsafe control action, acting controller, or unsafe scenario. As an example in today's NAS, the requirement for flight crews to read back clearances helps to prevent a variety of scenarios leading to hazards.

TABLE 1

Safety Order of Precedence (from [16])

Description	Priority	Definition
Eliminate hazards through design selection	1	Ideally, the hazard should be eliminated by selecting a design or material alternative that removes the hazard altogether.
Reduce risk through design alteration	2	If adopting an alternative design change or material to eliminate the hazard is not feasible, consider design changes that reduce the severity and/or the probability of the mishap potential caused by the hazard(s).
Incorporate engineered features or devices	3	If mitigation of the risk through design alteration is not feasible, reduce the severity or the probability of the mishap potential caused by the hazard(s) using engineered features or devices. In general, engineered features actively interrupt the mishap sequence and devices reduce the risk of a mishap.
Provide warning devices	4	If engineered features and devices are not feasible or do not adequately lower the severity or probability of the mishap potential caused by the hazard, include detection and warning systems to alert personnel to the presence of a hazardous condition or occurrence of a hazardous event.
Incorporate signage, procedures, training, and PPE	5	Where design alternatives, design changes, and engineered features and devices are not feasible and warning devices cannot adequately mitigate the severity or probability of the mishap potential caused by the hazard, incorporate signage, procedures, training, and Personal Protective Equipment (PPE). Signage includes placards, labels, signs, and other visual graphics. Procedures and training should include appropriate warnings and cautions. Procedures may prescribe the use of PPE. For hazards assigned Catastrophic or Critical mishap severity categories, the use of signage, procedures, training, and PPE as the only risk-reduction method should be avoided.

2.4.6 Necessary versus Sufficient Requirements and Requirements Optimization

Separate portions of the STPA analysis can generate the same requirements—for example, when different unsafe control actions cause the same hazard. Also, the same hazard might be prevented by a single mitigation in different causal scenarios. An example of such a hazard is loss of separation (LOS),

whose possible causes include ATC providing a conflicting speed advisory (UCA4 in [1]) and providing a speed advisory to the wrong aircraft (UCA5 in [1]); both of these example causes might be prevented by verification of IM-S algorithms and/or by adequate training. Because of such overlaps in cause and effect, a complete set of requirements can include redundancies that would need to be removed to obtain a minimal necessary set of requirements.

As discussed in Section 2.4.5, there may be different possible mitigations that are each sufficient to enforce the same safety constraint. In these cases, the STPA results can be used to identify available alternatives for design changes; this flexibility necessitates post-STPA review by system designers. If there is more than one set of requirements sufficient to ensure safety, an additional question is to consider how these requirements can be traded or optimized with different goals in mind. For example, it may be desirable to minimize the impact on cost, complexity, performance, or schedule due to additional requirements; these goals may not all be achievable together and likely require compromises. Such requirements optimization and refinement questions appear to be outside the scope of STPA analysis and likely need to be resolved by system designers and other stakeholders.

Additionally, STPA generates safety constraints with only a single priority level. It is up to SMEs to prioritize them. This need may require follow-on analysis to determine the likelihood or severity of a given scenario, but a taxonomy for how such prioritization could be included within STPA has not been defined.

2.4.7 SME Review to Provide Requirements Consistency and Refinement

Apart from requirements optimization, requirements consistency and applicability is another reason for expert review of STPA results. There may be interactions or conflicts between requirements for different causal scenarios because the safety constraints or requirements are identified independently for each causal scenario—as shown in the STPA hierarchy in Figure 3. While this property is not inherently a limitation, but rather a strength in many cases, there is a need to ensure consistency.

Additional guidance may be required to identify sets of safety constraints which, when applied to separate subsystems, may interfere with overall system performance. Alternatively, cases where constraints may provide conflicting guidance should be resolved. Certainly, some of these issues relate to all hazard analysis methods and are not unique to STPA. If traceability is improved and a formalized method of documentation is provided, it may be possible to link safety constraints to common controllers, actuators, unsafe control actions, etc. in order to better identify potential conflicts or over-constrained requirements. For example, a requirement that prevents a hazard in one scenario may cause a different hazard under a different scenario, or two requirements may be impossible to apply together. Resolution of these conflicts does not appear to be a part of the formal STPA process. Therefore, STPA-generated requirements should be reviewed by SMEs to identify and resolve any potential conflicts.

STPA-generated safety constraints and requirements should be interpreted in the appropriate context and provided to the appropriate audience. Because of their need for review by system designers

and SMEs, they should be interpreted as flexible proposals intended to increase safety—not as a final, non-negotiable set of specifications. For example, they likely cannot be provided directly to hardware manufacturers or software developers building a system; they should first be reviewed by SMEs to ensure an appropriate level of detail and feasibility.

It is important to determine what level of requirements detail is sufficient to provide insights and to guide system design decisions; the appropriate level of detail will depend on the intended uses of the requirements and their intended user types. Also, STPA results must place requirements in the proper context and carefully define any terms used. To address this concern in future studies, it may be helpful to involve experts with experience drafting system requirements.

Note that requirements-related issues discussed in this section apply not only to STPA but also to other safety analysis methods. Generating safety constraints that are consistent, and not conflicting, will require independent human review and verification of the safety analysis process as a whole.

2.4.8 Undemonstrated STPA Capabilities

Certain STPA capabilities were not demonstrated in the IM-S study but are claimed by STPA's creators or demonstrated in other studies presented in the STAMP workshop [3,10]. These capabilities were not directly addressed in the current study but should be considered in future work:

1. Analysis of social and organizational behavior related to safety—for example, the effect of organizational culture, safety processes, accident investigation procedures, and management decisions.

2. Analysis of cross-system or cross-capability interactions—for example, the interaction of different automated systems that may be designed independently and hence not explicitly designed to interoperate.

3. Analysis of system development processes, including both software and hardware. For example, does a system undergo sufficient safety testing prior to operational use?

4. Analysis of changes in a control structure and transitory effects while switching between control regimes.

The STPA application to IM-S considered only a single system or operational improvement and did not consider other systems or processes it may interact with (Figure 1 of [1]). NextGen, however, will increasingly rely on integrating multiple systems and information sources together to enable improved efficiency, safety, and reduced environmental impact. The question of cross-system interactions is therefore expected to become more important over time. Ensuring that such complex interconnected systems are developed to meet safety goals requires special care. As shown in Figure 5, for example, midterm Trajectory Based Operations (TBO) will require components and interactions spanning ground automation systems, ADS-B surveillance, cockpit flight management systems and displays, precision

navigation, datacomm, new operating procedures, and communications and collaboration tools between cockpit, facilities, and airlines—all while also supporting legacy systems and procedures as the National Airspace System transitions into NextGen.

TRAJECTORY BASED OPS	ADS-B	DataComm	NNEW	NVS	SWIM	ERAM	TFMS	ASDE-X	LAAS	RNP	RNAV	AIM	TBFM	CATMT	Procedures
MID-TERM															
Delegated Responsiblitity for Separation	X					X									X
Oceanic In-Trail Climb and Descent		X								X	X				X
Automation Support for Mixed Environment						X									
Initial Conflict Resolution Advisories		X				X									
Flexible Entry Times for Oceanic tracks			X												
Point in Space Metering						X	X						X		
Flexible Airpsace Management			X	X	X	X	X								X
Increase Capacity and Efficiency Using RNAV and RNP						X				X	X			X	X
Provide Interactive Flight Planning from Anywhere		X			X	X	X					X			

Figure 5. Example of NextGen System Integration—Trajectory Based Operations Components (from [2]).

While not yet directly observed, it should be feasible to apply STPA to systems of systems or to operational capabilities like TBO, which are built upon multiple OIs. The STPA process remains the same; however, due to greater complexity, there will likely be a greater need to include SMEs in the process early on.

In the case of multiple interacting systems, the initial step of STPA still requires a safety control structure diagram of the system—like that in Figure 8 of [1]. However, each of the subsystems need not be highly detailed. The safety control structure diagram for each of the interacting systems still includes control actions and feedback loops. In this way, unsafe control actions can still be identified.

The challenge of applying STPA to NextGen capabilities and systems arises when the unsafe scenarios and contextual factors are identified. STPA requires some level of expertise in each system component to understand how the process models, actuators, and sensors may lead to unsafe scenarios. In the case of a system of systems, or multiple OIs, expertise is required in each of the key system components. While previous safety assessments on a single operational improvement may have allowed a single practitioner to gain expertise in a system (via experience and documentation), applying this type of approach will likely not be feasible when systems grow quite large. For the case of TBO, in order to perform a complete safety analysis, a team of engineers would be required to work in concert. In this

case, because components are considered at a high level, expertise and experience will likely become increasingly vital. In particular, when working with legacy systems, the experience and knowledge of SMEs will be critical in distinguishing newly identified scenarios from preexisting unsafe scenarios, as discussed Section 2.4.4.

2.5 REVIEW OF STPA RESULTS USING NASA STANDARD FOR MODELS AND SIMULATIONS

The NASA Standard for Models and Simulations (M&S) seeks to provide a framework to evaluate evidence provided by modeling and simulation efforts in support of safety-critical decisions [6,7,8]. The standard was developed in response to a need for formal guidance in this area after the 2003 Space Shuttle Columbia accident. Note that the NASA standard and its M&S credibility assessment scale (CAS) are still in an early phase of implementation. There is little experience applying it to M&S efforts, and assessments using the standard should be only one factor among many to influence stakeholder decisions (Section VI of [6]).

This section does not provide a complete assessment of STPA using the NASA standard, but rather focused insights to inform the report's objectives. In addition to the standard's ongoing evolution, there are several reasons for this limited treatment:

- The information needed to fully apply the NASA standard to the STPA analysis of GIM and FIM has not yet been collected. For example, the standard requires M&S practitioners and decision-makers to assess the risk of using M&S results for decisions; such a risk description includes the potential safety consequences of the decision and the level of influence of the M&S results on the decision. The standard also requires a description of the limits of operation for the M&S methods, any limits which may have been exceeded, and proper use of the results (Section V of [6]). Such information could be collected during future work with STPA.

- The NASA standard's guidance depends on the context of M&S efforts—including data quality, process management, and personnel qualifications—in addition to the M&S methods themselves. Therefore, it should not be used for a comprehensive, general evaluation of a method such as STPA independent of particular application examples. STPA may be more appropriate for certain types of analyses, and the NASA standard may offer guidance on particular cases, but it likely will not offer conclusive evidence on whether to adopt the method in general. However, the standard is an appropriate framework for describing the benefits and challenges of STPA for the aviation community and is applied here in that context.

2.5.1 Credibility Assessment Scale

In order to assist decision-makers, the standard provides a credibility assessment scale (CAS) for M&S efforts. The CAS considers three assessment categories divided into eight factors. The three categories are model and simulation development, model and simulation operations, and supporting

evidence. The eight factors, with a brief explanation of each, are divided into the three categories as indicated in Table 2. The CAS is designed so that each factor is assigned a score or credibility level, then the factor scores are rolled up to a single score and complete results of the assessment provided to decision-makers. CAS developers note that the individual factor scores are more important than the rolled-up score and encourage decision-makers to consider all information provided—not just the scores themselves [7].

Modeling a system that is not yet tested experimentally or even fully designed is a challenge regardless of what methods are used. Hence, it is important to set realistic expectations for M&S results in these cases. Examples provided in documentation of the standard suggest that very few M&S efforts would achieve top scores on the CAS [7].

TABLE 2

NASA Standard for M&S: Credibility Assessment Scale (CAS) Categories and Factors (from [7,8])

Model and Simulation Development
• **Verification**: Were the models implemented correctly, and what was the numerical error/uncertainty? • **Validation**: Did the modeling and simulation results compare favorably to the referent data, and how close is the referent to the real-world system?
Model and Simulation Operations
• **Input Pedigree**: How confident are we of the current input data? • **Results Uncertainty**: What is the uncertainty in the current model and simulation results? • **Results Robustness**: How thoroughly are the sensitivities of the current model and simulation results known?
Supporting Evidence
• **Use History**: Have the current modeling and simulation been used successfully before? • **Modeling and Simulation Management**: How well-managed were the modeling and simulation process? • **People Qualifications**: How qualified were the personnel?

While the CAS is intended to score models—and STPA itself is not a model of the system under analysis—there is transference that permits evaluation of analysis methods like STPA. A reinterpretation of the CAS framework is applied in this section to better understand STPA; the reinterpretation of the three major categories is as follows: Analysis Development, Analysis Operation, and Supporting Evidence. In this section, use of the CAS is not intended to score or value STPA, but rather to understand and place in context STPA's characteristics. Therefore, this section provides information on each factor without attempting to assign CAS level scores.

2.5.2 Analysis Development

This subsection assesses arguments for STPA's benefits and quality of analysis—including the STPA analysis of GIM-S and FIM-S. The verification and validation (V&V) steps used to ensure the correctness of STPA analysis are different from those used to verify results of a numerical model or simulation. By design, STPA does not produce numerical output that can be compared to data from the real-world system. For the most part, it also does not rely on numerical input data. However, it does rely on a description of the analyzed system provided by human experts and by documents they produce. In the case of concepts in early development such as GIM-S and FIM-S, this description covers the future system's proposed behavior—not its actual, real-world behavior. Therefore, V&V of an analysis of early concepts cannot rely on real-world referent data because such data does not yet exist. However, V&V can be pursued through peer review and expert review, which were applied to the analysis of GIM-S and FIM-S and have been applied to STPA itself.

To summarize, V&V analysis development and of STPA itself have been pursued by the following means:

- Documentation of the analysis in a technical report [1]
- Stakeholder review and feedback by FAA and Lincoln Laboratory
- Subject-matter expert review and feedback on system control structures—including experts from FAA, NASA, Regulus Group, Aurora Sciences, and other organizations
- Peer-reviewed publications and conferences on STPA and results of applying it to a variety of applications [9,10,11]

Despite these peer-reviewed publications and other reports on STPA applications [1,3], it is not clear there has yet been a double-blind, side-by-side comparison between STPA and other hazard analysis methods. Where comparisons or applications of STPA are described in [1] and [3], a cursory review did not preclude the possibility that those applying STPA already had access to a prior safety analysis, which may have provided a starting point. As such, out-performing a competing analysis method only requires finding an additional hazard and unsafe scenario. In the case of a poorly performed prior safety analysis, finding additional unsafe scenarios might be quite simple. Furthermore, reports do not indicate that STPA has been verified to replicate all hazards already found by other methods.

An ideal comparison between methods would use independent results from each method as applied to the same system with the same inputs and no communication or influence between the analysis teams. It would determine where the results overlap and where they do not; this side-by-side comparison would help to identify the advantages and appropriate usage of each method, as well as any synergies between different methods. Such a comparison for STPA may already exist but was not reviewed for this report.

Once the IM-S preliminary hazard analysis is completed and published by the IM-S working group, a comparison between STPA and other methods might be possible. Such a potential comparison for IM-S would be one-way blind only—not double-blind—because STPA results have been shared with the IM-S working group for comment; their results could benefit from access to STPA results.

While STPA is an analysis method, it is built upon models from the Systems-Theoretic Accident Model and Processes (STAMP) framework. At its most basic level, the most important model underlying STPA is that of causality (i.e., cause and effect), which is an accepted principle and basis of the scientific method. While traditional safety analysis methods begin with causes (e.g., component failure in a radar's power system) to determine effects (e.g., loss of surveillance), STPA works backwards seeking to understand what causes may have led to a particular effect. It is the control loop model in Figure 4 and functional control models like that in Figure 8 of [1] that provide context to identify potential causes for any particular effect. This collection of models forms the core of STAMP. For the example of a radar system that provides surveillance, to identify potential causes for why there may be a loss of surveillance, a model of the radar along with any network and computer infrastructure is required. A greater level of detail in each model inherently provides greater opportunity to identify causes.

Some of the issues identified in Section 2.1 relate to verification and validation of underlying models. Specifically, these are human-factors, dynamic relationships, and completeness. Some of these identified limitations express the difficulty in modeling human behavior or of providing accurate descriptions of complex phenomena. Note that in many cases models are abstract simplifications of real-world systems; they are purposely created to be simplifications. Accordingly, much of the verification and validation of STPA relies on verifying and validating each of the modeling components, both individually and in composition, used in the analysis.

2.5.3 Analysis Process

This subsection assesses the quality of the input data and analysis results with respect to uncertainty and reliability.

The input pedigree for the STPA analysis refers to the quality of sources documenting the system's structure, operations, and related assumptions. These sources include documents such as Concept of Operations, RTCA standards, and other input provided by SMEs. They are used to derive the system control structure and understand the consequences of control actions. SERL cites these documents in the study's technical report, and they are considered the best available currently from the aviation community [1,14]. Several experts from different organizations were consulted during the process, and their collective opinions are also considered credible by the community. Given the level of maturity of IM-S, these inputs are likely the best available until updated versions appear.

The other two factors in this category—results uncertainty and results robustness—are intended for assessment of quantitative M&S methods and may not be directly applicable to STPA. STPA is not a well-defined mathematical algorithm operating on numerical data, but rather an analysis method which

relies on the judgment of human analysts and SMEs. Therefore, the uncertainty and reliability of the results must be considered qualitatively and is mostly covered under other CAS factors such as input pedigree and people qualifications.

Given that the STPA method has been applied correctly by qualified staff, uncertainty and robustness of the results can also be qualitatively assessed by considering the properties of the method itself. For example, does it tend to produce credible results even when there is uncertainty in the features of the system's control structure or behavior? Such uncertainty could be due to incomplete design, incomplete requirements, or disagreements between SMEs. Currently, it appears that STPA does produce such robust results because it can help to reveal shortcomings or conflicts in the system control structure. Answers to these questions will become clearer through additional discussion and could be tested through experimental comparison of STPA results with simulations and behavior of operational systems.

2.5.4 Supporting Evidence

This subsection considers evidence of the credibility of the STPA method, of this particular study, and of the researchers performing it.

Use History

As discussed in Section 2.1, the ideas behind STAMP and the STPA method have been applied to a variety of safety questions in various industries and domains [3,10,11]. These include critical aviation systems and applications such as

- In-Trail Procedure (ATSA-ITP),
- Helicopter search and rescue, and
- Friendly fire prevention.

Researchers have also applied the method to the safety of medical and automotive systems; these systems share some types of safety challenges with the NAS because their components are not always designed or evaluated for interoperability despite the requirement of interoperability.

It is not clear that the method has been formally adopted or endorsed by any safety regulators or professional associations, but it has attracted interest from the safety community [10].

Management

SERL staff performing this study were managed by and collaborated with MIT faculty, Lincoln Laboratory staff, and FAA staff.

People Qualifications

SERL staff performing this study have previous experience with STPA and are experts on the method [3,10,11]. However, they are not ATC experts or IM-S experts and did not have full-time support from such experts—only limited access through occasional discussions and written feedback. In future STPA studies, full-time support from one or more SMEs is recommended and would likely enhance the results.

2.6 ADDITIONAL SAFETY QUESTIONS ON IM-S

Potential future updates to the STPA study could investigate the additional questions listed below. Due to limited resources, these were not prioritized in the current study but may be covered at a high level in the STPA causal scenarios. Additional remaining questions are typical and expected at the conclusion of any research effort.

1. What are the possible unsafe interactions between aircraft guided by different speed control algorithms? Is it necessary to simulate or test the interactions among each possible combination of algorithms? What standards are needed for flight management systems to ensure safe IM-S operations?

2. Can aircraft guided by FIM-S and GIM-S safely operate in the same arrival stream?

3. What are the safety implications of a sudden loss or outage of IM-S/TFM automation?

4. How could changes in human operator workload affect safety during IM-S operations? This question applies to both ATC and flight crews.

5. How could unexpected weather events influence IM-S operations?

6. If IM-S clearances are not automatically shared across facilities, how should handoffs of IM-S aircraft be handled safely at facility boundaries? Would voice-based communication of clearances lead to unsafe or infeasible ATC workload levels?

7. How does the update frequency of speed advisories affect safety? What are the lower and upper bounds for safety?

3. ORGANIZATIONAL ASSESSMENT

In addition to understanding the technical merits, stakeholders performing aviation safety studies must understand whether a method such as STPA should be accepted and implemented in their organization. They must feel confident that it provides useful information to system designers, operators, management, and others. They must also understand how it can be included in the SMS framework and processes.

3.1 APPROPRIATENESS AND VALIDATION OF STPA OUTPUTS

The goal of safety analysis is the understanding and prevention of accidents. STPA contributes to this goal by providing necessary safety constraints on a system's design and operations; therefore, its outputs are generally appropriate and relevant for organizations responsible for safety. However, it is important to ensure that these constraints are specific enough to enable feasible and effective design changes or mitigations. This concern can be addressed through iterative discussion between researchers applying STPA, system designers, operations staff, requirements experts, and other SMEs. As the method generates safety constraints and associated requirements, the results can be reviewed for feasibility and effectiveness. They can then be refined as needed to reach an adequate level of specificity.

In addition to expert review, another method of validating STPA-generated safety constraints and requirements would be to assess their effect on system behavior during simulations, human-in-the-loop testing, or operational evaluation in the NAS. It is not clear that such assessment has been done to date; even though STPA has been shown to identify possible accident causes, they have not been demonstrated in controlled experiments. Comparison of system behavior with enforcement versus without enforcement of STPA safety constraints will help to confirm that they do indeed prevent hazards and accidents. For concepts early in their development, this comparison could be done using simulations if operational evaluation is not feasible.

As noted in Section 2.1.6, STPA by design does not provide a probabilistic assessment of risk, which is currently a required output of safety analysis under the FAA SRM process [4]. However, STPA can enhance and supplement such an assessment as discussed in the next section.

3.2 INTEGRATING STPA TO EXISTING SAFETY PROCESSES

This subsection discusses milestones toward acceptance of STPA and its wider use in aviation safety. It also describes a few possible approaches to applying STPA at different levels of effort—dependent on needs and preferences of users such as SRM panels.

A first step is to improve understanding and awareness of the method among the aviation community. Acceptance is not yet as common as acceptance of older methods and will take time to fully develop. This step is in progress through publications, speaking engagements by SERL staff, and

conferences attracting interest from the aviation community—including regulators and researchers from academic, industry, government, and military organizations [10].

A second step is to address concerns from the safety and aviation communities. These include STPA's treatment of dynamic interactions, human factors, and emergent behaviors discussed in Section 2.1 and in previous work—including work by STPA users [2,10]. As discussed in a Boeing/NASA study, some of these concerns are not unique to STPA and also apply to current, accepted methods used under the SMS [12].

Another possible concern is consistent definition of terms. For example, as noted in the technical report [1], STPA assumes a slightly different definition of "hazard" than is typically used by FAA. Another example is the term "controller," which can refer specifically to a human air traffic controller or generally to a software, human, or institutional agent in a system's control structure. Other terms and their usage may also differ. STPA and SMS terms may be harmonized in several possible ways:

1. agreeing to adopt a single consistent definition,

2. defining separate, distinctive terms such as "STPA hazard" versus "SMS hazard," or

3. noting clearly which definition is used in a particular context.

Options 2 and 3 are likely most feasible and would help to minimize confusion for the different communities involved. As an example, the SERL technical report uses the full wording "air traffic controller" to refer to the specific human role, and the term "control agent" to refer to a more general component in a control diagram.

Before full implementation of STPA—as with any new safety method—additional trials and demonstrations will likely be needed. These trials should help address concerns about the method and provide more evidence of its utility and importance. The SERL study of IM-S provides such a trial, but additional selective application of STPA to different systems and subsets of the NAS will answer additional questions. Comparison of STPA results with simulations and operational tests, when available, will also improve understanding and confidence in the method.

Once sufficient awareness and acceptance of STPA by the community are achieved, additional discussion will be needed to determine the best way to formally leverage the method. Feedback and guidance from the aviation community and SMS experts will be needed at this point. Experts with experience drafting SMS policies or applying the SRM process in safety studies could provide advice on how to include new methods like STPA and when they will be appropriate to use. A broad sample of opinions should be considered.

As part of this discussion, the next several subsections suggest potential approaches for integrating STPA to existing aviation safety processes; they are not intended to provide comprehensive or detailed guidance but rather a general outline of how STPA could contribute additional insights. Note that despite

its potential advantages, requiring use of STPA may not be ideal because it would restrict the flexibility in choice of methods currently permitted by the SMS.

3.2.1 Potential STPA-Assisted SRM Process

Use of STPA does not require giving up other safety modeling and analysis methods. It could be applied along with other methods to generate possible accident scenarios and improve the completeness of modeling, simulation, and analysis. For example, STPA could add structure to existing analysis processes and may reveal additional causal mechanisms not identified by other methods. In this context, STPA would supplement the existing SRM process—not replace it. SRM panels could implement it in parallel without changes to the existing SRM process—particularly given the flexibility that the SMS allows safety professionals to choose the methods they prefer for SRM studies [4,15].

STPA also need not be completely implemented in all steps in order to yield benefits. STPA users note that it can provide safety insights at each of its steps—not only at its endpoint when safety constraints are identified [10]. For example, the process of documenting a system's control structure could yield insights into what control functions may be missing or incomplete.

A more formal change would involve inclusion of STPA steps in FAA guidance on concept development so that each STPA step would occur at the appropriate time in a concept's development process. For example, identification of a concept's control structure must wait until its actors and their basic functions are documented, but it could proceed even if the relationships between the actors are not yet fully understood or finalized.

3.2.2 Dependence on Concept Maturity Level

As discussed in the 2012 Lincoln Laboratory report, STPA is intended to provide insights for systems at different concept maturity levels [2]. However, the optimal use of STPA may depend on the maturity level of the concepts under analysis. This subsection discusses potential differences in application of STPA to concepts at earlier and later maturity levels.

For less mature concepts, a full application of STPA may be most appropriate because other methods may be inapplicable or not yet applied. Less mature concepts offer less information to the analyst and their behavior is less well understood and documented. However, these concepts also have fewer design constraints and hence offer greater opportunity for uncovering feasible safety-driven design changes; the opposite is true for more mature concepts. STPA may be more readily accepted when applied to less mature concepts because there are fewer analysis options for them and they are less commonly analyzed for safety.

For more mature concepts closer to operational implementation in the NAS, full application of STPA is also worthwhile. However, a partial application may also provide insights not already provided by other methods and would require fewer resources—both cost and schedule. Partial application could

involve (1) applying only earlier steps of STPA, or (2) applying STPA only to subsystems that other methods cannot adequately treat. In case option 2 is chosen, STPA may provide the greatest benefit when applied to subsystems whose performance data is expensive or impossible to collect or whose behavior is challenging to model, simulate, or test by other methods. A disadvantage of such a partial application is that it could overlook interactions with the remaining unanalyzed portion of the system.

In all cases, a partial implementation or phased trials will likely yield greater acceptance of the method by new users, who may later accept a full implementation given initial evidence of its utility. However, partial implementation of STPA in conjunction with other methods will likely require that the complimentary method also be a top-down approach like STPA. By joining two top-down approaches, there is likely to be a beneficial overlap in portions of the analysis.

3.3 REQUIRED RESOURCES AND FEASIBILITY

This section discusses practical issues related to adoption and use of STPA.

STPA-based studies considering only risk severity and accident causes—and not considering risk likelihoods—may require fewer resources than studies that also consider risk likelihood. In addition to avoiding underestimation of risk, STPA-based studies avoid the cost and schedule constraints of estimating event likelihoods. This advantage may be important for budget-constrained stakeholders.

Despite the structure and partial automation of the STPA process, it depends on expert input, which is needed to define the system control structure and ensure appropriate unsafe control actions are considered. This report has also shown that expert input is needed to translate unsafe control actions into safety constraints and requirements. Because of the need for expert input throughout the process, STPA appears to require similar levels of expert input as other safety methods.

3.3.1 STPA User Training

A potential concern with any new analysis method is the amount of training required for analysts new to the method. The expected benefits of the new method should well exceed training and other adoption costs. Though individuals can typically achieve basic understanding of STPA with reasonable effort (days or weeks), mastery of its application likely requires longer experience and feedback from experts—potentially months or years. STPA is a structured process, but it requires discretion and judgment on the part of the analyst, which develop only with time and experience. Less experienced analysts may be able to provide adequate results, but possibly not fully developed results or a proper assessment of them. Also, because the method is based on systems thinking, the analyst needs knowledge of this field as well as assumptions and tools used to generate, format, and provide traceability of the results. Note that these training concerns are not unique to STPA and apply to any analysis method.

Because the amount and type of STPA training needed is not well-documented and may vary, this issue should be further investigated by interviewing researchers using STPA in a variety of organizations.

Understanding the amount of training required before beginning productive work with STPA would help organizations considering STPA to assess their training requirements. The STAMP workshop archive provides a list of such researchers [10]. Researchers in a variety of organizations and fields of expertise should be interviewed for a broad perspective.

Any survey on skill level and ease of learning STPA should be verified through review of STPA results generated by the same practitioners. In some cases, it has been observed that application of STPA by new practitioners was incorrect due to inconsistent use of the term “hazard” or that they applied STPA only superficially, thereby obtaining incomplete results. The reported skill level of STPA practitioners may also be inaccurate. Those that report limited benefit of STPA may lack the proper skills to apply STPA critically, and therefore do not understand its true capacity. As such, there may be a need for a formal STPA training and certification process to enhance, test, and verify the competence of practitioners.

This page intentionally left blank.

4. CONCLUSION AND RECOMMENDATIONS FOR FUTURE WORK

This report and the associated SERL technical report [1] show that STPA has been gaining more visibility among the safety and aviation communities as a potentially valuable tool for initial safety studies during the concept development phase. It should be considered as an additional tool for SRM panels, while keeping in mind that additional improvements to the STPA process may be warranted.

The SERL STPA-based study of IM-S has provided insights into the application's safety implications, but additional feedback from IM-S SMEs should be pursued—particularly regarding the completeness and coverage of STPA-generated safety constraints and requirements. Additional questions on IM-S could be investigated in future updates.

Considering STPA in the context of the SRM process shows that the method can and should be applied at the discretion of SRM panels [4]. As an additional tool for identifying accident mechanisms and mitigations, it can enhance the SRM process without formal changes to that process.

This page intentionally left blank.

GLOSSARY

ADS-B	Automatic Dependent Surveillance Broadcast
ATC	Air Traffic Control
ATM	Air Traffic Management
ATSA	Airborne Traffic Situational Awareness
CAS	Credibility Assessment Scale
CNS	Communications, Navigation, or Surveillance
CONOPS	Concept of Operations
FAA	Federal Aviation Administration
FIM-S	Flight Deck Based Interval Management–Spacing
FMEA	Failure Mode and Effects Analysis
GIM-S	Ground Based Interval Management-Spacing
ICAO	International Civil Aviation Organization
IM-S	Interval Management for Spacing
ITP	In Trail Procedure
JDPO	Joint Planning and Development Office
M&S	Models And Simulations
MIT	Massachusetts Institute of Technology
NAS	National Airspace System
PPE	Personal Protective Equipment
RTCA	Radio Technical Commission for Aeronautics
SERL	Systems Engineering Research Lab
SMEs	Subject-matter experts
SMS	Safety Management System
SRM	Safety Risk Management
STAMP	System-Theoretic Accident Model and Processes
STPA	System Theoretic Process Analysis
TBO	Trajectory-Based Operations
TCAS	Traffic Alert and Collision Avoidance System
TFM	Traffic Flow Management
UCA	Unsafe Control Action

This page intentionally left blank.

REFERENCES

[1] Cody H. Fleming, Seth M. Placke, and Nancy G. Leveson, "STPA Analysis of NextGen Interval Management Components: Ground-based Interval Management (GIM) and Flight Deck-based Interval Management (FIM)." Massachusetts Institute of Technology (September 23, 2013).

[2] E. Harkleroad, A. Vela et al., "Risk-based Modeling to Support NextGen Concept Assessment and Validation," Lincoln Laboratory, Lexington, MA, Project Report ATC-405 (2013).

[3] Nancy G. Leveson, *Engineering a Safer World: Systems Thinking Applied to Safety*. MIT Press (2012).

[4] FAA Air Traffic Organization, Safety Management System Manual, Version 2.1 (May 2008).

[5] FAA Concept Development and Validation Guidelines, Final Version 1.0 (2011).

[6] Steve R. Blattnig, James M. Luckring, et al., "NASA Standard for Models and Simulations: Philosophy and Requirements Overview," *Journal of Aircraft*, Vol. 50, No. 1 (2013), pp. 20–28. doi: 10.2514/1.C000303.

[7] M. Babula, W.J. Bertch et al., "NASA Standard for Models and Simulations: Credibility Assessment Scale," AIAA Paper 2009-1011 (January 2009).

[8] NASA Standard for Models and Simulations, NASA-STD-7009. Approval Date: July 11, 2008. standards.nasa.gov/documents/viewdoc/3315599/3315599.

[9] Nancy G. Leveson, "STPA and CAST Tutorial," First STAMP/STPA Workshop at MIT, Cambridge, MA (April 17–19, 2012). psas.scripts.mit.edu/home/1st-stampstpa-workshop-2012.

[10] 2013 STAMP Workshop, Massachusetts Institute of Technology, Cambridge, MA (March 26–28, 2013). psas.scripts.mit.edu/home/2013-workshop-presentations

[11] Cody Harrison Fleming, Melissa Spencer et al., "Safety assurance in NextGen and complex transportation systems," *Safety Science*, Volume 55 (June 2013), pp. 173–187, ISSN 0925-7535. www.sciencedirect.com/science/article/pii/S0925753512002871.

[12] Xidong Xu, Mike L. Ulrey et al., "Safety Sufficiency for NextGen: Assessment of Selected Existing Safety Methods, Tools, Processes, and Regulations," NASA/CR–2013-217801. (February 2013).

[13] R.B. Whitner and O. Balci, “Guidelines for selecting and using simulation model verification techniques,” Proceedings of the 21st Conference on Winter Simulation, editors E. MacNair, K.J. Musselman and P. Heidelberger, Washington, DC (1989), pp. 559–568.

[14] Surveillance and Broadcast Services (SBS) Concept of Operations: Arrival Interval Management – Spacing (IM-S), PMO-010, Revision 02, Final (March 1, 2013).

[15] Margaret Stringfellow, “Accident Analysis and Hazard Analysis for Human and Organizational Factors,” PhD Dissertation, Massachusetts Institute of Technology, Cambridge, MA, (2010).

[16] “Department of Defense Standard Practice, System Safety,” MIL-STD-882E (11 May 2012).