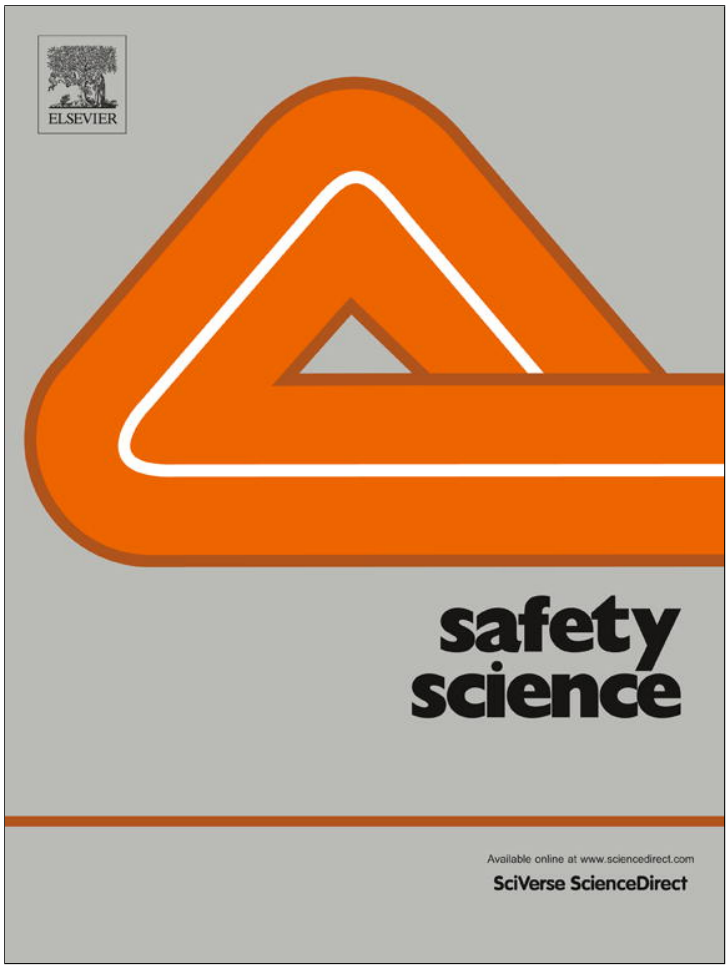

(This is a sample cover image for this issue. The actual cover is not yet available at this time.)

Safety Science 55 (2013) 173–187

Contents lists available at SciVerse ScienceDirect

Safety Science

journal homepage: www.elsevier.com/locate/ssci

Safety assurance in NextGen and complex transportation systems ☆

Cody Harrison Fleming [a], Melissa Spencer [a], John Thomas [a], Nancy Leveson [a,*], Chris Wilkinson [b]

[a] MIT, United States
[b] Honeywell Aerospace, United States

ARTICLE INFO

Article history:
Received 8 May 2012
Received in revised form 17 December 2012
Accepted 19 December 2012

Keywords:
Air transportation
System safety
Hazard analysis

ABSTRACT

The methods currently used to assure the safety of planned changes to our air transportation systems were developed 50 years ago for systems composed primarily of hardware components and of much less complexity than the systems we are building today. These methods are not powerful enough to handle the complex, human and software intensive systems being planned and introduced today. This paper describes an alternative and demonstrates it on a new NextGen procedure to allow more flight level changes over oceanic and other regions with limited radar coverage. The new approach and results are compared with the results obtained by the more traditional methods being used for NextGen.

1. Introduction

The plan to transform the National Airspace System (NAS), called NextGen, changes the system through an evolution from a ground-based air traffic control system to a satellite-based system of air traffic management (FAA, 2011). The overarching goals of NextGen are to (1) reduce flight delays by improving airport operations; (2) improve aviation's impact on the environment through reduced CO_2 emissions and fuel use; and (3) make the airspace safer via more precise tracking, improved information-sharing, and implementing a Safety Management System (EUROCAE ED-78A/RTCA DO-264, 2002).

As changes are designed and implemented to realize the NextGen goals, assurance is necessary that the current high level of safety will not be degraded. The complexity of the current system and the changes envisioned makes this process challenging. Powerful tools will be required to assure aircraft and airspace safety.

Traditional approaches to safety analysis assume that accidents are caused by component failures (Leveson, 1995; Roland and Moriarty, 1983). They therefore focus on reliability analysis techniques, particularly fault tree or event tree analysis. The goal of these traditional approaches is to determine scenarios of component failures that together will lead to an accident or loss event. Failures may be single or multiple and are usually assumed to be random. After the component failure scenarios are identified, engineers use fault tolerance or fail-safe techniques to protect against hazards caused by the identified failures and to increase individual component integrity. A fly-fix-fly approach augments the design techniques with investigation of accidents and potentially serious incidents in great depth and recommendations made from the results to prevent reoccurrences.

This approach has been very effective in the past because there have been relatively few changes in the basic aircraft or air traffic control design; the systems are relatively simple; technology has changed slowly; engineers have been able to use very conservative design approaches; and the system components can be effectively decoupled so that interactions can be anticipated, simplified, and guarded against. This approach, by itself, is becoming less effective, however, as these assumptions start to be violated in our new or enhanced system designs.

Software is increasingly an important part of systems and allows enormously more complex and tightly coupled systems to be constructed. The potential for accidents arising from unsafe interactions among non-failed components, i.e., unplanned system and software behavior, is increasing. NextGen components, for example, may involve more than just one aircraft and one onboard system and include multiple aircraft, ground controllers, space-based systems, and communication links between aircraft. The traditional hardware-oriented safety engineering techniques focusing on failures do not adequately handle these types of new accident causes.

In addition, human roles are changing from direct control to supervision of automation, which requires more cognitively complex human decision-making. Like software, the changing roles of pilots and ground controllers introduces the potential for new causes of accidents that are not well handled by today's failure-oriented and hardware-oriented approaches.

☆ This research was partially supported by the NASA Aviation Safety Program (Contract NNL10AA13C).

* Corresponding author. Address: Massachusetts Institute of Technology, Room 33-334, 77 Massachusetts Ave., Cambridge, MA 02142, United States. Tel.: +1 617 258 0505; fax: +1 617 253 7397.
E-mail address: leveson@mit.edu (N. Leveson).

0925-7535/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ssci.2012.12.005

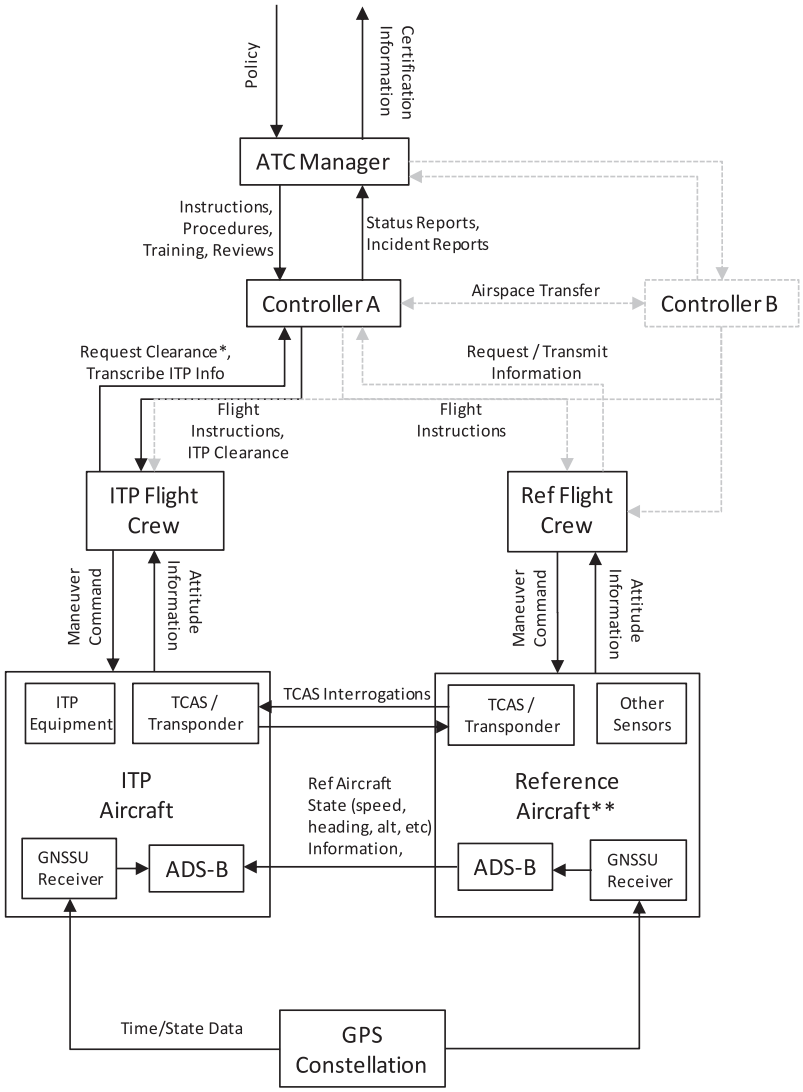

Fig. 1. Safety control structure for ATSA-ITP.

To deal with these new accident causes, more powerful tools are needed. This paper describes and demonstrates a new approach to safety analysis based on systems and control theory rather than reliability theory. Safety is treated as a control problem rather than a "prevent failures" problem, allowing not only consideration of the causes of the component failure accidents that were predominant in the past but also the new causality factors that are increasingly important today. This approach can be applied to NextGen and to other upgrades to complex transportation systems.

To demonstrate and evaluate the approach, we use a new Air Traffic Control (ATC) procedure, called Airborne Traffic Situational Awareness In-Trail Procedure (ATSA-ITP). ATSA-ITP, or simply ITP, was chosen because the safety analysis and underlying methodology has already been documented in DO-312 (RTCA, 2008). ITP provides a real-world case study with which to compare our safety assurance philosophy and analytical techniques being proposed to those being used by the Federal Aviation Administration (FAA), European Organisation for the Safety of Air Navigation (EUROCONTROL), and their associated organizations.

We first describe our new approach and the results of applying it to ATSA-ITP. We then compare the results to those of the ITP safety analysis documented in DO-312, particularly the difference philosophical underpinnings of these different approaches.

2. Using STAMP and STPA for safety assurance

The significant technical changes envisioned for NextGen creates a necessity for a new, more powerful model of accident causality that better represents today's complex, socio-technical systems. The new model used in our analysis, called Systems Theoretic Accident Model and Processes (STAMP) (Leveson, 2012), extends the types of accidents and causes that can be considered by including non-linear, indirect, and feedback relationships among events. In this way, the traditional causality model is extended to consider new types of accident causes arising from component interactions (rather than just component failures), cognitively complex human mistakes, management and organizational errors, software errors (particularly requirements errors), etc. Accidents or unacceptable losses can result not only from system component failures but also from interactions among system components—both physical and social—that violate system safety constraints.

2.1. Systems Theoretic Accident Model and Process (STAMP)

In systems theory, emergent properties (like safety) associated with a set of components are related to constraints upon the degree

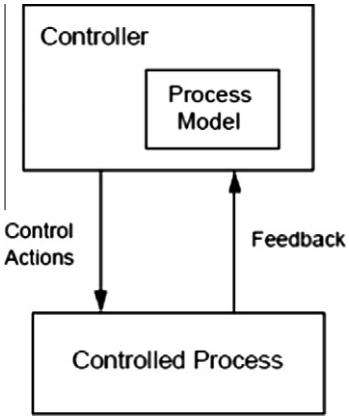

Fig. 2. A simple control loop and process model.

of freedom of those components' behavior (Checkland, 1981). System safety, then, can be reformulated as a system control problem rather than a component reliability problem: accidents or losses occur when component failures, external disturbances, and/or dysfunctional interactions among system components are not handled adequately or controlled—where controls may be managerial, organizational, physical, operational, or manufacturing.

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

Section 2.2 briefly describes the hazard analysis procedure, called STPA (System Theoretic Process Analysis), used to identify the system constraints necessary to ensure safe development and operation of complex socio-technical systems. The results are shown using a model-based specification method, called Intent Specifications, that was initially developed for the certification of TCAS II (Leveson et al., 1994) and later extended (Leveson et al., 2000). Intent Specifications capture the results of the hazard analysis in a readable, reviewable way by people from multiple disciplines.

2.2. System Theoretic Process Analysis (STPA)

STPA is a hazard analysis technique built on STAMP. As described above, accidents are viewed in STAMP as resulting from inadequate enforcement of constraints on system behavior. Fig. 1 shows a safety control structure to enforce safety constraints (for the example air traffic control procedure in this paper). Each hierarchical level of the control structure represents a control process and control loop with actions and feedback. Each component of the control structure has different responsibilities with respect to enforcing safe system behavior. The reason behind the inadequate enforcement may involve classic component failures, but it may also result from unsafe interactions among components operating as designed or from erroneous control actions by software or humans.

STPA consists of the following general framework, which will be described throughout the rest of this subsection.

- STPA Step 0 – Identify System Level Properties.
  - Accidents and Hazards (Section 2.5).
  - Hierarchical Control Structure (Section 2.2).
- STPA Step 1 – Identify potentially Unsafe Control Actions (Sections 2.2 and 2.5).
- STPA Step 2 – Identify potential causal factors for Unsafe Control Actions (Section 2.5).

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

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

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

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

Before providing details about STPA, the demonstration system used in this research procedure is defined and the approach currently being used to assure safety in NextGen described.

2.3. ATSA-ITP

Because of the limited radar coverage in oceanic and other remote airspace, air traffic controllers have historically relied on procedural separation rules to ensure safe traffic flow and minimum separation between aircraft. Aircraft in these sectors generally fly long, pre-defined flight paths (tracks), and air traffic controllers have limited options for knowing the positions of all aircraft in a sector at the same time or the ability to directly communicate with aircraft. To compensate for the limitations, these sections of the airspace often have much larger separation requirements than those applied to airspace with greater surveillance and communication coverage. The large separation minima, however, place severe limits on the capacity of a given track in a remote airspace.

Changing flight levels allows for greater fuel efficiency due to changes in aircraft weight during the course of a flight (as fuel is burned). Because of the large separation requirements, a desired flight level might often not be available due to the presence of

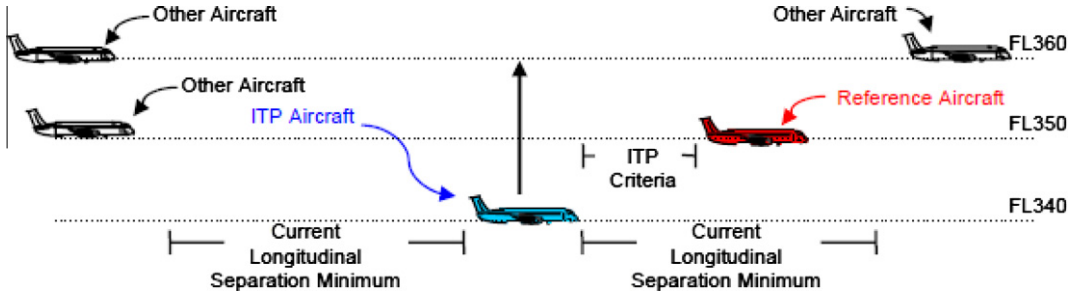

Fig. 3. ITP following climb (RTCA, 2008).

"blocking" aircraft in intervening flight levels that fall within the minimum longitudinal separation distance.

The new Airborne Traffic Situational Awareness In-Trail Procedure (ATSA-ITP, referred to here as just ITP) will allow many of these previously blocked flight level changes to occur. ITP enables either leading or following Same Track aircraft to perform a climb or descent to a requested flight level through intervening flight levels. The crew will use information derived on the aircraft to determine if the criteria for applying the ITP procedures are met with respect to one or two Reference Aircraft at intervening flight levels. Note that standard separation minima between aircraft may not hold at times during the ITP maneuver. The ITP equipment must ensure that the reduced separation minima, as defined for ITP, are observed.

In the ITP procedure, the flight crew determines if the criteria for an ITP request are met. If the criteria are satisfied, the flight crew may request an ITP, identifying the Reference Aircraft in the request. ATC verifies that the ITP and Reference Aircraft are Same Track and that the maximum Closing Mach Differential will not be exceeded. If the controller determines that separation minima will be met with all Other Aircraft, the climb or descent request may be granted. The controller does not determine or verify the separation distance from the Reference Aircraft (RTCA, 2008).

An example of one the six potential ITP maneuver geometries is shown in Fig. 3. In this diagram, the ITP aircraft wants to pass the Reference Aircraft.

ITP consists of four phases: the initiation phase, instruction phase, execution phase, and termination phase (RTCA, 2008).

1. ITP Initiation phase: The preparation for performing the application consists of realizing the desire and assessing the appropriateness for requesting an ITP maneuver by the flight crew. This includes the identification of the Reference Aircraft in the procedure and transmission of the ITP request to the ground controller.
2. ITP Instruction phase: The ITP clearance is issued by the controller, and reevaluated by the flight crew.
3. ITP Execution phase: The cleared ITP Aircraft performs the ITP maneuver, maintaining the required rate of climb/descent and speed as directed by the ITP clearance. Conducting an ITP maneuver is similar operationally to standard climbing/descending maneuvers.
4. ITP Termination phase: The procedure is terminated once the ITP Aircraft has achieved the requested Flight Level or an abnormal event results in premature termination of the ITP maneuver.

2.4. Current process used to assure safety in ITP

Hazard analysis is usually performed on a completed design to assess the safety of the designed system. For upgrades to the air transportation system, a different strategy is needed. The high-level system design is being produced and procedures are being defined but the components will be designed and produced by different companies and the designs of each will differ from the others. Each company producing components used in the ITP procedure cannot be held responsible for assuring the safety of the overall system.

The assurance of the safety of ITP, therefore, involves ensuring that the overall architectural design and the defined ITP procedures will together not lead to a hazardous system state and then generating specific safety requirements for each of the component manufacturers to ensure that their components will integrate into the system in a safe manner. Some components already exist, such as ADS-B, and they need to be evaluated to ensure that their behavior does not violate the ITP safety requirements. For ITP, then, the goal of the safety analysis is to evaluate what is required for safe behavior of the system as a whole and then to generate specific safety requirements for the behavior of the individual components that will ensure safe operations.

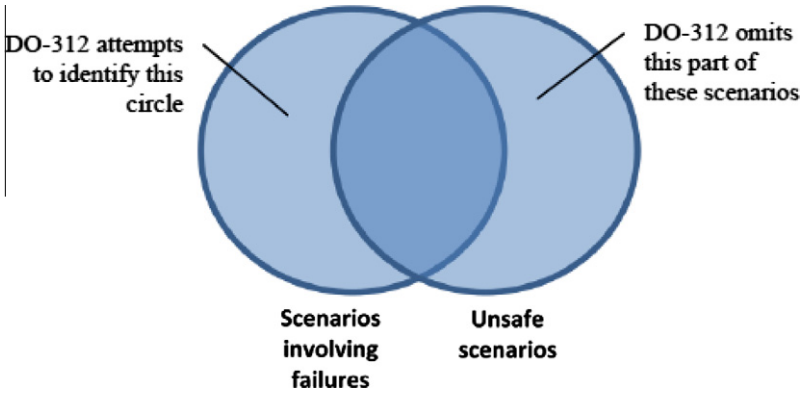

Fig. 4. The consequences of equating safety and reliability.

The safety analysis process adopted for NextGen and used for the ITP procedure uses traditional modeling and hazard analysis techniques. A Collision Risk Model was created to identify the aircraft state vectors needed to avoid a collision. The state vectors provide the information to determine the parameters that must be satisfied before the ITP is performed. We assume here that this model is correct and that the associated Operational Performance Assessment appropriately correlates with the risk model.

This paper concentrates on the hazard analysis, which for ITP is called the Operational Hazard Analysis and is documented in DO-312 (RTCA, 2008). There are four steps in this Operational Hazard Analysis: (1) identify hazards, (2) allocate severity classes, (3) determine probability of occurrence (Pe), and (4) assign a safety objective.

2.4.1. Identifying hazards

An operational hazard is defined as an event that may "arise when the system is in a faulted mode." The failure events are identified by applying three failure modes to each action in the ITP procedure description: (1) action is not available or not executed, (2) action is performed incorrectly or is performed using incorrect information, and (3) action is executed in non-suitable conditions or executed out of sequence. The set of operational hazards is then determined by grouping the failure events that lead to a similar consequence.

An operational hazard appears to simply be an event that follows some failure. There does not seem to be any tie of an operational hazard to an accident or incident, which is the usual definition of a hazard. For example, a hazard identified for ITP is that ATC incorrectly rejects an ITP clearance (and, therefore, the ITP is not executed). Although such an event is not desirable, it is not unsafe. This definition of hazard leads to defining all failures as hazards, not just those that can lead to a loss, and essentially equating reliability and safety. "Non-hazardous" hazards are later ignored in the analysis so this limitation is not serious but it does involve extra work and may be confusing to those familiar with the usual definition of hazard.

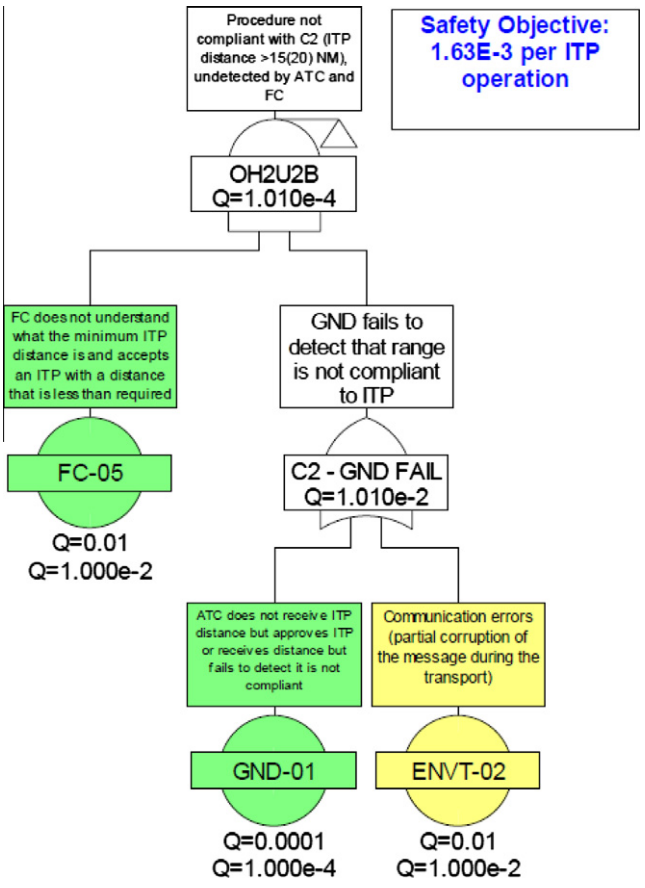

Fig. 5. Example fault tree from DO-312.

More important, however, is that the equating of a hazard and a failure leads to omitting some unsafe states (see Fig. 4). While some failure scenarios are unsafe, some unsafe scenarios lie outside the realm of a failure and are not included in the safety analysis when only failures are considered.

2.4.2. Allocating severity classes

Considering the characteristics of the environment in which the ITP is to be used and any mitigation means (primarily procedures) that help to reduce the hazard effects, the hazards are classified from 1 to 5 in terms of their effect on operations, occupants, air crew, air traffic service, etc. The "non-hazardous hazards" identified in the first step are eliminated from further consideration here. A maximum tolerable probability called a Safety Target is assigned to each outcome so that more severe outcomes have smaller tolerable probabilities.

2.4.3. Determining probability of occurrence

Event trees are first used to identify the different possible chains of events that can result from each hazard given the barriers (controls) in place and to quantify the probabilities of each adverse outcome. The use of events trees seems odd as they were developed for much simpler systems and cannot possibly identify all the events leading to a hazard in a system as complex as ITP. The use of event trees also requires assigning probabilistic values to mitigating effects of the barriers (controls). For example, the aircraft crew detecting an aircraft's proximity during an interrupted ITP maneuver through visual means and taking appropriate action to avoid an NMAC (Near Mid-Air Collision) is assigned a probability of success of 0.80 and by means other than unaided visual acquisition and responding properly is assigned the probability of success of 0.90. These numbers appear arbitrary and not justified in the ITP safety analysis except to say that they were calculated based on computer simulations and expert feedback.

Because human actions are involved in this analysis and in the chain of events, probabilities are needed for these human actions. Operational safety workshops were held with pilots, controllers, and operations experts. The participants were provided with various types of human error that could occur in ITP and were asked to qualitatively assess the likelihood of occurrence as *very often*, *often*, *rare*, or *very rare*. These qualitative rankings were later assigned a quantitative value: An error that may happen *very often* is assigned a probability of occurring between 1% and 10%, while a *very rare* human action is assigned a probability of occurring of less than 0.01%. For example, "the probability that the ITP aircraft flight crew will level off at an intermediate (incorrect) flight level is assumed to occur no more than Very Rare," or less than 0.01% of the time.

Beginning from the identified operational hazards, fault trees are used to identify causal factors. The analysis stops when an identified cause (called a Basic Cause) describes a single failure, a human error, or an environmental factor. Fig. 5 shows one of the fault trees. The top level event in the figure is that the procedure is performed when one of the criteria for safe ITP is not satisfied and neither the flight crew nor ATC detects this non-compliance.

According to the analysis in the fault tree, failure to comply can occur either because the flight crew does not understand the minimum distance required or the ATC does not receive the required data or fails to detect non-compliance due to a communication error involving corruption of data during transport. There are, of course, many other reasons for communication errors but these are ignored in the analysis. The leaf nodes in the tree are assigned

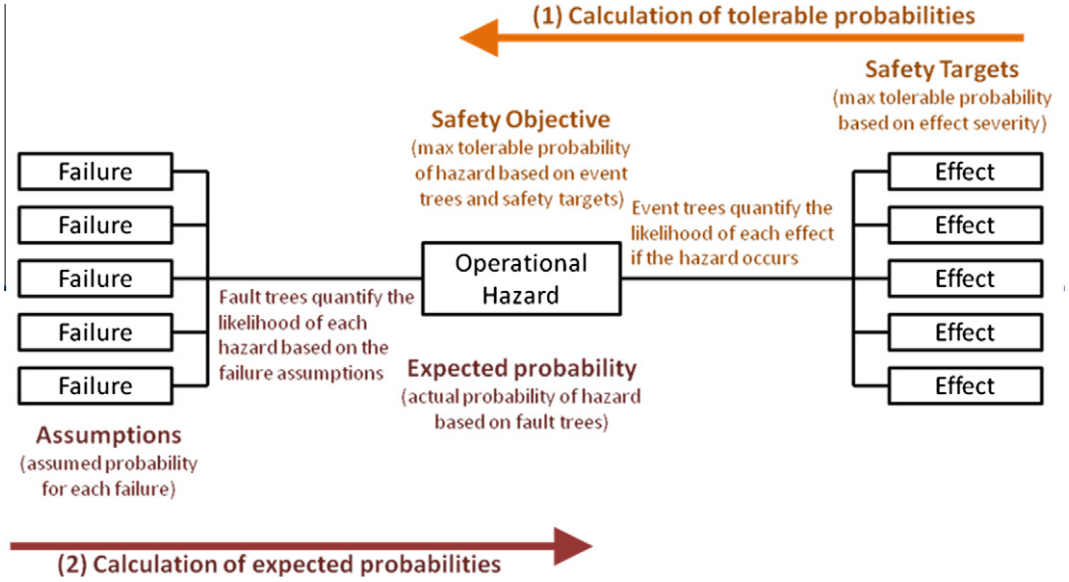

Fig. 6. Summary of the official process to ensure safety in ITP.

a probability and combined using the AND/OR logic of the tree to calculate an overall probability of occurrence for the hazardous event at the top of the tree.

Note that the fault tree assumes independent behavior, however the interaction and behavior of the flight crew and ATC may be coupled, with the parties exerting influence on each other or both being influenced by high-level system conditions.

2.4.4. Assigning safety objectives

A safety objective is allocated to each hazard by using the probabilistic event tree to compute a maximum tolerable likelihood for the hazard such that on average the Safety Targets will be met. The safety objectives are then compared to the expected probability of occurrence for each hazard (as determined by the fault tree for that hazard) to ensure that each safety objective is met. The overall process is summarized in Fig. 6.

2.4.5. Using the results

The results of the analysis are used to create two types of safety requirements:

- *Safety operational requirements*: For example, the flight crew shall maintain the required mach number during the maneuver, and if during an ITP maneuver the ITP flight crew detects that the climb/descent rate is not compliant, the crew shall attempt to rectify the deficiency and follow regional contingency procedures.
- *Safety probabilistic requirements*: For example, the likelihood that the ITP equipment provides an undetected erroneous relative track angle to the flight crew shall be less than 1E-3 per flight hour.

In addition, a number of safety probabilistic assumptions are identified during the analysis. For example, the probability that ATC does not receive ITP Distance (as part of the ITP climb/descent request) but approves the ITP procedure or fails to detect that the ITP Distance received in the request is not compliant, is assumed to occur no more frequently than Very Rare.

2.5. Applying STPA to ITP

The new hazard analysis technique based on STAMP is called STPA (System Theoretic Process Analysis). The process involved and results are very different than the more traditional approach described in the previous section. Although, as will be seen in Section 3, STPA is more powerful than the approach currently being used in terms of it identifying more paths to hazards, STPA is surprisingly also easier to use.

As in the traditional safety analysis, the process starts by identifying hazards, although hazards are not equated to failures. Instead, as in system safety engineering for defense and space systems, a hazard is defined as a system state that under worst-case environmental conditions will lead to a loss or accident. This definition encompasses more than simply the states following component failures but undesired states that can result from any cause. The general hazards for aircraft include:

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

Because ITP is at first only going to be used on long oceanic tracks, hazard H-1 is the most relevant at this time and was the focus of our analysis. H-1 leads to the high-level system safety requirement/constraint: "The ITP must not cause a pair of controlled aircraft to violate minimum separation standards." The STPA hazard analysis identifies ITP system and component requirements necessary to enforce this constraint.

STPA is performed on a functional control diagram of the system. An example of a functional control diagram for the ITP-related parts of the system is shown in Fig. 1. Levels of control above the ATC manager are not shown.

Using this functional control diagram, STPA hazard analysis has two goals. The first goal is to identify potentially hazardous control actions, i.e., specific instances of the control actions shown on the downward arrows for each component that could lead to a system-level hazard by violating a system safety constraint. These identified unsafe control actions are used to refine the high-level safety constraints/requirements into more detailed safety requirements.

The second goal of STPA is to analyze the system to determine the potential scenarios that could lead to providing one of the unsafe control actions. These scenarios can then be used to derive detailed safety requirements/constraints for the system components to ensure that all the components operating together cannot create a system hazard, in this case H-1 (violation of minimum separation standards). The system and component-level requirements are used to design controls and to certify the safety of the system and of its components.

2.5.1. STPA step one

The first step of STPA identifies control actions for each component that can lead to one or more of the defined system hazards. The four general types of unsafe control actions were shown above, i.e., providing a control action that leads to a hazard, not providing a control action that is needed to prevent a hazard, incorrect timing or sequencing, and applying a control action too long or stopping it too soon. A table can be used to document the hazardous control actions identified, as in Tables 1 and 2. The hazardous control actions can then be translated into high-level system and component safety requirements and constraints.

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

Once the tables are created, the identified unsafe control actions are rewritten as high-level system safety constraints. The constraints are refined further, in a top-down system engineering process, during STPA Step Two. Although some (or most) of these may seem obvious, they are used not only in the refinement to more detailed requirements but also to provide a chain back to the unsafe control actions from the more detailed requirements and constraints in order to document where these requirements came from and therefore why they are needed. For example, the high-level safety constraints on the ITP flight crew generated from Table 1 are:

SC-FC.1: The flight crew must not execute the ITP when it has not been approved by ATC.
SC-FC.2: The flight crew must not execute an ITP when the ITP criteria are not satisfied.
SC-FC.3: The flight crew must execute the ITP with correct climb rate, flight levels, Mach number, and other associated performance criteria.
SC-FC.4: The flight crew must not continue the ITP maneuver when it would be dangerous to do so.
SC-FC.5: The flight crew must not abort the ITP unnecessarily. (Rationale: An abort may violate separation minima).
SC-FC.6: When performing an abort, the flight crew must follow regional contingency procedures.
SC-FC.7: The flight crew must not execute the ITP before approval by ATC.
SC-FC.8: The flight crew must execute the ITP immediately when approved unless it would be dangerous to do so.
SC-FC.9: The crew shall be given positive notification of arrival at the requested FL.

Similar safety constraints on ATC can be generated from Table 2:

2.5.2. STPA step two

The second step of STPA examines each control loop in the safety control structure to identify potential causal factors for each hazardous control action, i.e., the scenarios for causing a hazard. This process will basically refine the high-level safety requirements identified in Step One. Fig. 7 shows a generic control loop that can be used to guide this step. While STPA Step One focuses on the provided control actions (the upper left corner of Fig. 7), STPA Step Two expands the analysis to consider causal factors along the rest of the control loop.

One reason a safety constraint might be violated is that the process model of the controller is incorrect, for example, the flight crew thinks it is safe to execute the ITP when it is not. The incorrect process model, in turn, may be the result of inadequate feedback provided by a failed or erroneous sensor or the feedback may be delayed or corrupted. Alternatively, the designers may have omitted a feedback signal or the flight crew may have received incorrect input from ATC or from other input devices (such as ADS-B).

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

Fig. 8 shows part of the Step Two analysis for unsafe flight crew behavior in executing the ITP. A more complete analysis can be found in Fleming et al. (2012). The required process model and responsibilities of the flight crew are shown in the Flight Crew Controller Box. The unsafe control actions are written on the top left connection from the Controller to the actuator while some of the possible causes of unsafe flight crew behavior are shown on the other connections in the control loop.

Compare the results with the fault tree shown in Fig. 5. The STPA causal factors include the basic communication errors included in the Fault Tree Analysis (FTA), but also include additional reasons for communication errors as well as guidance for understanding human error within the context of the system. Communication errors may result, for example, because there is confusion about multiple sources of information (for either the flight crew or ATC), confusion about heritage or newly implemented communication protocols, or simple transcription or speaking errors. There is no way to quantify or verify the probabilities of any of these sources of error for many reasons, particularly because the errors are dependent on context and the operator environments are highly dynamic and, in fact, not necessarily designed yet. Instead of assuming that humans will rarely "fail," our analysis assumes they will make mistakes. The identified causal factors can be restated as requirements and then used to improve the system design to try to eliminate or mitigate these human errors, such as ensuring that the air traffic controllers and pilots have access to the information they need to make safe decisions and create safe control commands and that this information is available when they need it. The information produced by the Step 2 analysis may also be used in developing ATC and flight crew operational procedures and training.

2.5.3. Specification and analysis of the hazard analysis results

For our specification of the safety requirements for ITP, we used an intent specification. An intent specification is a specification and model-based development framework supporting system design and other system engineering activities, intended to assist humans at all organizational levels in dealing with complexity by providing more readable and reviewable specifications and providing support for change management. Intent specifications are based on psychological research in human problem solving and on basic principles of system theory and system engineering (Leveson et al., 2000).

Table 1
Potentially hazardous control actions (CA) for the flight crew (FC).

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

[a] The unsafe control actions "ITP aircraft levels off above [and below] requested flight level" are not included in the following as separate unsafe control actions as they are equivalent to the unsafe control action "ITP executed incorrectly.".

Table 2
Potentially hazardous control actions for ATC.

Control action	Not providing CA causes hazard	Providing CA causes hazard	Wrong timing/order of CA causes hazard	CA stopped too soon or applied too long
Approve ITP request		Approval given when criteria are not met	Approval given too early	
		Approval given to incorrect aircraft	Approval given too late	
Deny ITP request				
Abnormal Termination Instruction	Aircraft should abort but instruction not given	Abort instruction given when abort is not necessary	Abort instruction given too late	

Any specification environment could be used that has the following features: (1) facilitates the tracing of system-level requirements and design constraints down into detailed design and implementation and the documentation of design rationale, (2) assists in the assurance of various system properties (such as safety) in the initial design and implementation, and (3) reduces the costs of implementing changes and re-analysis when the system is changed, as it inevitably will be.

Fig. 9 shows an example of the traceability that is required during the original design phase and, even more important, for the re-analysis of the safety of ITP if changes are made in the future. Starting the analysis from scratch for every planned change is not usually feasible. Safety analysis of the changes in any system requires appropriate documentation of the original safety analysis. Traceability is a critical part of that documentation.

Fig. 9 shows two levels of the intent specification, the system level analysis and requirements and the system-level architectural design decisions that result from these requirements. Traceability is provided by hyperlinks and rationale is embedded in the document. The documentation of the design rationale and of the underlying assumptions about the operational environment in which ITP will exist is especially important in operational safety analyses. When conditions change such that the assumptions are no longer true, then a new safety analysis should be triggered. In the traditional system engineering specification approach, these assumptions may be included in a safety analysis document (or at least should be), but are not usually traced to the parts of the implementation they affect. Therefore, even if the system safety engineer knows that a safety analysis assumption has been changed, it is very difficult and resource-intensive process to figure out which parts of the design used that assumption.

Intent specifications also contain a model-based specification of the behavioral requirements for each system component, again traced to the system-level architectural design decisions and thus the hazards and hazard analysis. Black-box requirements, free of any component design decisions, are specified using an executable and formally analyzable language that was developed for the FAA to specify TCAS II requirements during its certification activities. Fig. 10 shows part of the tabular notation used to specify the ITP requirements for the case study. The complete specification can be found in Fleming et al. (2012). This tabular notion represents a formal specification of the system (or component) states required to generate a necessary transition or prevent an unwanted transition. In the ITP Display example shown in Fig. 10, each column on the right side of these "AND/OR" tables represent the combination of states required to change the ITP display mode. For example, if Display Mode Input is Display ITP **AND** Reference Aircraft Data is Correct, then the equipment should display the ITP data.

3. Comparing the results of the two analyses

The current approach to safety analysis for NextGen can be compared to our new approach both in terms of the actual safety requirements generated and in the underlying philosophical differences.

3.1. Safety requirements identified

In our comparison of the official analyses with the STPA results, we included with the DO-312 results a set of additional requirements provided by the FAA that are being used for ITP (FAA, 2010). Our STPA analysis identified nineteen high-level safety requirements that were not in either of the two official NextGen documents. Table 3 shows some of the omitted requirements identified using STPA.

One other example of an important omitted requirement involves the reference aircraft, which (reasonably) has no requirements levied on it. DO-312 assumes that the reference aircraft will not deviate from its flight plan during ITP execution. There should be a contingency or protocol in the event that the reference aircraft does not maintain its expected speed and trajectory, for example, because of an emergency requiring immediate action.

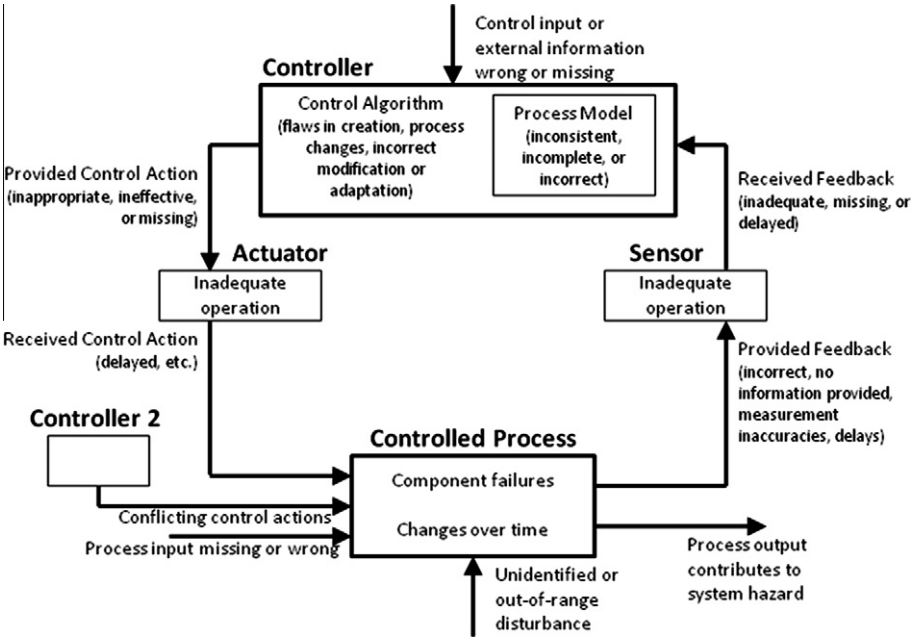

Fig. 7. General control loop with causal factors.

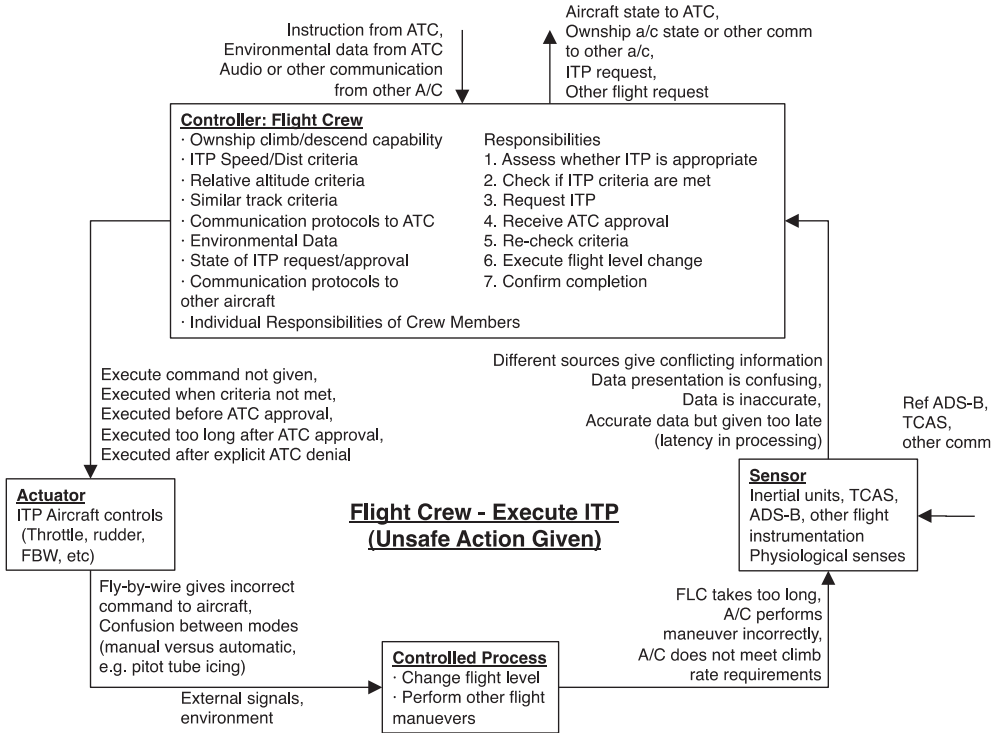

Fig. 8. Partial causal analysis for the flight crew executing the ITP unsafely.

3.2. Philosophical differences

The larger and more important comparison is not with the results of the specific techniques provided in DO-312—they could be changed or improved—but rather the basic philosophical differences between the traditional approach to hazard analysis represented by the methods used in DO-312 and that of the new systems-theoretic approach described in this paper. The results of any hazard analysis are inextricably tied to the overall philosophy and viewpoint of the analysis approach, the causal factors identified, and the certification method implied by the safety analysis approach. Table 4 summarizes the comparison.

3.2.1. Analysis differences

The DO-312 safety assessment is based on the assumption that the system operates nominally and that accidents result due to deviation from nominal behavior at the component or subsystem level. The Collision Risk Model created for ITP calculates probabilities based on nominal system behavior, where the probability of longitudinal overlap—a potential crash scenario—is the aggregation

System Level Analysis and Requirements

Hazard
[H-1] A pair of controlled aircraft violate minimum separation standards (←[A-1],→[1.2])

Causal Analysis
[FC.1] FC believes aircraft climb/descent capability is greater than it is (process model inconsistency)
[FC.2] FC does not receive communication from ATC (inadequate/missing feedback)

Safety Requirements
[1.2] ITP shall provide the flight crews of aircraft operating in procedural airspace the ability to determine if an ITP maneuver is appropriate (←[H-1], ←[FC.1], ↓[2.1])
[1.3.1] ITP training shall include communication protocols (both channels and appropriate syntax) between flight crew and air traffic control (←[H-1], ←[FC.2], ↓[2.5])

…

System-Level Architectural Design Decisions

Design Decisions
[2.1] The ITP flight crew must check that the following criteria are fulfilled before requesting an ITP clearance. This requirement does not imply that an individual assessment of each criterion is carried out by the flight crew, but rather that each criterion is assessed by either the flight crew, or by automation (see section 3.5.1). Although not required for all criteria, it is recognized that automation may provide a more predictable solution. (↑[1.2], [1.9],[1.26],[1.27])

[2.1.1] The Ownship climb/descend capability criteria are considered passed if and only if the ITP Aircraft can climb/descend in the desired direction at a rate of 300 fpm or more.

Design Rational: Initiation Distance Criteria and other geometric values ([2.1.2], [2.1.3], [2.1.4]) were selected such that when a Flight Level change at 300 fpm is performed with the related 20 or 30 kts Closing Ground Speed Differential, the distance between the aircraft does not become less than the ITP Separation Minimum (i.e., 10 NM).

[2.5] ATC shall include a minimum data set in its clearance, in order to minimize the risk of confusion during communication between the Flight Crew and ATC. This information includes the Reference Aircraft ID and the cleared-to Flight Level in the ITP clearance.(↑[1.3.1])

Fig. 9. ITP equipment specification partial example.

of errors in aircraft attitude and environmental assumptions. While the underlying mathematics is fine, the problem lies in the modeling assumptions, i.e., that the ITP and Reference Aircraft will always maintain minimum separation and that error propagation is due solely to instrumentation error and environmental conditions. Safe behavior is imposed by providing contingencies for the identified off-nominal behaviors or conditions.

In contrast, instead of assuming that accidents are caused by expected incorrect behavior (called a "design basis accident in the nuclear power community), STAMP/STPA assumes a worst-case scenario and identifies potential scenarios that could lead to that worst case. The use of worst-case system behavior has the potential for identifying a greater set of contingencies for "off-nominal" behavior.

A second difference is the underlying model of accident causality used. Most all current hazard analysis and safety assessment techniques are based on a chain-of-events model of causality, where the events represent component failures. Each failure event leads to the next one in the chain with a direct relationship between the two. Traditional safety engineering techniques then focus on preventing or reducing the probability of component failure to prevent accidents. Fault tree analysis (FTA) (at least as commonly used and as used in DO-312) also assumes that most of the component failure modes are independent. Human operators and software are treated as if they fail like mechanical hardware, and probabilities of failure are assigned to them. Given the critical role that software and human decision making play in modern complex systems, these assumptions are unrealistic (discussed further below).

Accidents often arise due to unanticipated failures or due to unsafe interactions among non-failed components. Starting a hazard analysis from failures puts the analysis at risk of identifying only a subset of the possible causes, as opposed to beginning with hazards and identifying the interactions that could possibly lead to hazardous states, including those not involving component failure (see Fig. 4).

STAMP and STPA assume the worst case in identifying accident causality. That is, STPA starts with an accident, identifies hazards that may lead to an accident, and then identifies causal factors, including interactions among system components that could possibly lead to hazardous states. Table 5 shows the alternative definitions and the different identified hazards that result from these definitions.

DO-312, starting from events arising from a "faulted mode," identifies six operational hazards. To identify the operational hazards, abnormal events were identified by looking at failure modes related to each of the expected actions throughout the ITP phases. Each abnormal event was then traced forward in time to create a chain of events that leads to an outcome such as an accident or inconvenience. An event from each chain is then selected and labeled as an operational hazard. Note that the same operational hazard may appear in more than one chain of events.

This process identified Operational Hazards OH-3, OH4, and OH-5 (shown in Table 5) as having "no effect on safety" and therefore are not analyzed further. OH-2 and OH-6 are identified as

Control Mode

ITP Display Mode

Description: ITP Control is the primary Control Mode of the ITP Equipment. The equipment can be starting up, operational, or experiencing an internal fault.

Comment: System automatically defaults to display ITP Data over displaying Other Data and sleeps after 10 minutes of idle time.

References: Reference AC Data State, ITP AC Data State, Display Mode Input

Appears In: ITP Distance Display, Ground Speed Differential Display, Data Quality Passes Display, Data Quality Fails Display, Relative Track Angle Similar Track Status Display, Reference AC ID Display, Data Input Error Display

DEFINITION

= Off

System Start	T

= Data Fault Detected

System Start	F	F	F	F	F	F
Reference AC Data State in state Obsolete Reference AC Data	T	*	*	*	*	*
Reference AC Data State in state Inaccurate Reference AC Data	*	T	*	*	*	*
Reference AC Data State in state Low Integrity Reference AC Data	*	*	T	*	*	*
ITP AC Data State in state Obsolete ITP AC Data	*	*	*	T	*	*
ITP AC Data State in state Inaccurate ITP AC Data	*	*	*	*	T	*
ITP AC Data State in state Low Integrity ITP AC Data	*	*	*	*	*	T

= Display ITP Data

System Start	F	F
Display Mode Input is Display ITP Data	T	*
Display Mode Input is Obsolete	*	T
Reference AC Data State in state Correct Reference AC Data	T	T

= Display Other Data

System Start	F
Display Mode Input is Display Other Data	T
Display Mode Input is Obsolete	F

= Sleep

System Start	F	F
Time Since Display Mode Input was Last Received > 600 seconds	T	*
Display Mode Input is Sleep	*	T

Fig. 10. Formal, tabular specification partial example.

having the potential for gravest impact and were the focus of the safety analysis. OH-1, interruption of an ITP maneuver (which in STAMP would be considered a cause of a hazard and not a hazard itself) was deemed to have less significant impact on safety. This decision reflects the overarching philosophy of assuming nominal or best-case behavior. Our STPA analysis, in fact, identified interruption of the ITP maneuver in a later step as an important potential cause of a hazard (violation of minimum separation requirements) and we derived requirements to constrain and protect against this behavior.

The DO-312 safety analysis procedure leads to overlooking important scenarios, such as the case where the ITP maneuver is abandoned because it is discovered after starting the ITP that the ITP criteria are not met. In some cases, the DO-312 analysis even overlooks the abnormal events that were used to derive the hazard in the first place. For example, OH-1 was identified in part by the possibility of an ACAS (airborne collision avoidance system, or TCAS) resolution advisory causing the crew to interrupt the maneuver, but the analysis of OH-1 and the resulting fault tree completely omit that scenario.[1]

The list of hazards used for STPA (in the far right column of Table 5) includes general airspace hazards. Because ITP is at first only going to be used on long oceanic tracks, H1 is the most relevant hazard and was the focus of our analysis. In the future, if ITP is to be allowed in other locations, the analysis must be augmented to include any of the other hazards that become relevant.

DO-312, starting from events arising from a "faulted mode," identifies six operational hazards. To identify the operational hazards, abnormal events were identified by looking at failure modes related to each of the expected actions throughout the ITP phases. Each abnormal event was then traced forward in time to create a chain of events that leads to an outcome such as an accident or

[1] In fact, the fault tree analysis of OH-1 only identifies two basic causes for OH-1: an equipment failure (e.g., engine failure) or a misuse of traffic information by the flight crew.

Table 3
Additional requirements and constraints.

Safety-related constraint	Rationale
If ATC is using traffic data from multiple sources to monitor traffic surrounding an ITP maneuver, they must have a clearly defined hierarchy of which data to use	With a new procedure, ATC may be getting traffic data from multiple sources (e.g. the ITP aircraft, other ADS-B equipped aircraft, traditional fix point communications). In order to issue the proper clearance and have a clear picture of current traffic conditions, ATC must always give precedence to the various sources in the same hierarchy
ITP equipment must be used by the flight crew only to determine if a FL change is feasible and to collect the necessary data to transmit to local ATC	The requirements on the ITP equipment as stated in DO-312 do not address the equipment's ability to do anything more than determine the feasibility of an ITP maneuver prior to execution of the maneuver. If the equipment is used for more than this (e.g. to monitor the Reference Aircraft during execution), the other intended uses must be analyzed for safety as well
ATC must be provided with a mechanism for knowing how much time has elapsed since the ITP request was made	Humans are not very good at estimating the time that has elapsed, particularly under stressful or distracting circumstances. ATC must be able to accurately determine the elapsed time since a request has been made during evaluation of request in order to avoid issuing clearance too late (i.e. when conditions have changed)
The window of time between ITP reassessment and execution must be less than TBD (to be determined) minutes	In order to ensure that traffic conditions remain as they were when criteria were reassessed, a precise upper limit on the time between reassessment and ITP execution must be defined
ATC must grant clearance for ITP within TBD min of request	The phrase "without delay" is used in current documentation. The ITP procedures must include the definition of what constitutes waiting too long to issue clearance. A precise definition of what it means for approval to be given too late is needed to avoid confusion between the ATC and the ITP aircraft flight crew
If ATC notices a potentially hazardous traffic scenario, they must assess if an abnormal termination of the ITP maneuver is necessary and initiate it	ATC must assume that an abnormal termination is always possible and communicate with ITP flight crew accordingly. While it is possible that emergency situations will mean that this communication comes too late, ATC should always assume that an abnormal termination is possible unless they know otherwise. This would include any unexpected (i.e. emergency) maneuvers by the reference (non-ITP) aircraft
The ITP flight crew must initiate an abnormal termination of ITP without delay when they believe they will enter a hazardous situation	Flight crew must be trained to initiate an abnormal termination when traffic conditions are dubious, and not assume that ATC will initiate any needed termination
ITP equipment must not be used during the procedure, i.e., once the ITP flight crew begins the flight level change, ITP equipment should not be used to assess nearby traffic	FAA memo says (Section 3.1.4) that the flight crew must monitor ITP conditions throughout maneuver but DO-312 says the FC is not required to do so (A.2.3.1.2). We created this requirement under the assumption that ITP equipment would be implemented to display criteria prior to ITP execution but not necessarily during execution and could cause an unnecessary abnormal termination if used during execution for monitoring traffic

inconvenience. An event from each chain is then selected and labeled as an operational hazard. Note that the same operational hazard may appear in more than one chain of events.

This process identified OH-3, OH4, and OH-5 (shown in Table 5) as having "no effect on safety" and therefore are not analyzed further. OH-2 and OH-6 are identified as having the potential for gravest impact and were the focus of the safety analysis. OH-1, interruption of an ITP maneuver (which in STAMP would be considered a cause of a hazard and not a hazard itself) was deemed to have less significant impact on safety. This decision reflects the overarching philosophy of assuming nominal or best-case behavior. Our STPA analysis, in fact, identified interruption of the ITP maneuver in a later step as an important potential cause of a hazard (violation of minimum separation requirements) and we derived requirements to constrain and protect against this behavior.

The DO-312 safety analysis procedure leads to overlooking important scenarios, such as the case where the ITP maneuver is abandoned because it is discovered after starting the ITP that the ITP criteria are not met. In some cases, the DO-312 analysis even overlooks the abnormal events that were used to derive the hazard in the first place. For example, OH-1 was identified in part by the possibility of an ACAS (TCAS) resolution advisory causing the crew to interrupt the maneuver, but the analysis of OH-1 and the resulting fault tree completely omit that scenario.[1]

The list of hazards used for STPA (in the far right column of Table 5) includes general airspace hazards. Because ITP is at first only going to be used on long oceanic tracks, H1 is the most relevant hazard and was the focus of our analysis. In the future, if ITP is to be allowed in other locations, the analysis must be augmented to include any of the other hazards that become relevant.

3.2.2. Differences in causal factors identified

An important distinction between STPA and the FTA performed in DO-312 is the assumption about basic causal factors. DO-312 assumes or prescribes independent probabilities for off-nominal behavior (which is common in analyses using fault trees or event trees). STPA, in contrast, accounts for sub-system interaction component interaction accidents (where no components may have failed) and, in fact, assumes that not only can causal factors be dependent but also that the behavior of one component might be highly influential on other aspects of the system.

STPA also recognizes that software does not "fail," but merely performs the way it was designed: it can therefore be hazardous due to flawed requirements (or implementation) or unsafe interactions with the rest of the system. Most software-related accidents arise due to flaws in the software requirements (Leveson, 1995) so getting a complete and correct set of safety-related requirements/constraints on software behavior is key to preventing software-related accidents.

Human operators also do not fail in the sense that hardware does and most of their failures are not random. Instead, humans are influenced by the design and operation of the overall system and the operational context and can thus make unsafe decisions due to the factors in Fig. 7, such as incorrect mental models of the process they are controlling, possibly due to missing or incorrect feedback.

The human error identification process used for ITP is very incomplete, but the problems involve more than just incompleteness. Human error is treated in exactly the same way as a physical failure, that is, as a deviation from a predefined behavior or procedure. Unfortunately, this treatment of human error oversimplifies

Table 4
General comparison of approaches.

	DO-312	STAMP/STPA
Analysis Philosophy	Success oriented, i.e. it assumes nominal case then tries to predict probability of deviation Provides set of contingencies for off-nominal behavior Emphasis on preventing or reducing failures Assumes most failure modes are independent	Assumes worst-case scenario, i.e. it starts with accident, then hazards, then causal factors and assumes that any of the causal factors can happen Emphasis on enforcing constraints on system (and thus component) behavior Accounts for sub-system interactions and how these influence safety-related behavior
Causal Factors	Considers only hardware failures, or treats operators and software as if they are hardware (e.g. leaves on a fault tree with assigned probabilities of failure)	Assumes accidents are caused by lack of enforcement of safety constraints on the behavior of the system and its components Assumes that software does not "fail" but can still be hazardous due to flawed requirements or unsafe interactions with rest of system Human operators perform within the context of a larger system design and, like software, do not necessarily "fail" but can make unsafe decisions
Certification Method	Assign performance goals or necessary probabilities of failure, then manufacturer attempts to assure compliance	Specify safety constraints derived from STPA, based on safety-related control actions and required component behavior, which manufacturer implements.

Table 5
Hazard analysis comparison.

	DO-312	STPA
Hazard Definition	An event that may arise when the system is in a faulted mode; events leading to an OH are called its Basic Causes and Abnormal Events, and can either be system failures, human errors, procedures dysfunctions or failures and conditions external to the application itself Or, any condition, event, or circumstance that could induce an operational effect	A system state or set of conditions that together with a particular set of worst-case environmental conditions, will lead to an accident (loss)
Hazard Identification	OH-1: Interruption of an ITP maneuver (flight crew abandons the maneuver) OH-2: Execution of an ITP clearance not compliant with ITP Criteria OH-3: ITP request not accepted by ATC OH-4: Rejection by the flight crew of an ITP clearance not compliant with the ITP Criteria OH-5: Rejection by the flight crew of an ITP clearance compliant with the ITP Criteria OH-6: Incorrect execution of an ITP maneuver	H1 – a pair of controlled aircraft violate minimum separation standards H2 – aircraft enters unsafe atmospheric region H3 – aircraft enters uncontrolled state H4 – aircraft enters unsafe attitude (excessive turbulence or pitch/roll/yaw that causes passenger injury but not necessarily aircraft loss) H5 – aircraft enters a prohibited area

it as a binary decision between right and wrong. Many of the most important situations involved in accidents are overlooked because they are difficult or impossible to model in this way, including when:

- The correct behavior is not predefined or not clear.
- The prescribed behavior is thought to be incorrect by the person responsible for following it.
- Procedures conflict with each other, or it is not clear which procedure applies.
- The person has multiple responsibilities or goals that may conflict.
- The information necessary to carry out a procedure is not available or is incorrect.
- Past experiences and current knowledge conflict with a procedure.
- The procedure is misunderstood or the responsibility for the procedure is unclear.
- The procedure is incorrect.

The DO-312 analysis produced a short list of human errors such as "flight crew fails to detect inadequate climb/descent rate." Because basic causes are the "lowest level of failure," human errors are not analyzed in further detail. No attempt is made to understand why the human errors may arise or to prevent them. Instead, all errors are assumed to occur randomly at a given probability rate. If necessary, mitigation attempts are made to reduce the chance that human errors will lead to a hazard. Mitigation is done by adding barriers called "mitigation measures". Interestingly, every mitigation measure is simply a new procedure that is imposed on the humans.

Because human behavior was treated as random and no attempt was made to explain or understand the potential human errors, the possibility of eliminating or reducing errors was precluded. Human behavior is usually not random but is influenced by the current context, the information observed (rather than simply the information available (Dekker, 2004; Dekker, 2006), and constructed beliefs about the operation of the system. Human behavior is also heavily dependent on interactions with other system components and past experiences. The treatment of human error as independent, random events precludes the possibility of eliminating or reducing errors in the first place. Eliminating errors through system design is typically more effective than managing hazards through mitigation alone. There is also no guarantee that humans will perform better when additional (mitigation) procedures are added, and they may actually perform worse because of the added workload.

Table 6
Comparison of Specifications.

	DO-312	STPA
Requirements and Assumptions	Assumption AS.40 The probability that ATC does not receive ITP Distance (as part of the ITP climb/descent request) but approves ITP procedure or fails to detect that ITP Distance received in the request is not compliant, is assumed to occur no more frequently than Very Rare AS.12 The corruption of information because of human factors (HF) occurs no more than Often	Requirement [1.1.2] ITP shall provide the flight crews of aircraft operating in procedural airspace the ability to determine a clear procedure for communicating data about the desired flight level change and necessary state data to the local air traffic controller [1.2.1.1] Once ITP request has been made, all communication between ATC and the FC must occur on the same communication channel [1.2.1.2] All communication protocols must include definitions of when a communication is complete [1.10]–[1.17] (see appendix B) [1.18] ATC must have access to current∗ knowledge of the velocity, heading, and location of all aircraft involved in ITP request Assumption: ATC will have this knowledge as part of their overall ability to maintain separation, regardless of ITP clearances

Instead, a safety analysis should not only identify *what* humans can do wrong, but also *why* and *how it can be avoided*. This is the goal of STPA. Assuming that the flight crew and ATCO are not intentionally malicious, this identification requires understanding the conditions under which each erroneous decision can make sense to them and modifying or adding system design requirements to help make the correct decisions obvious.

3.2.3. Differences in certification approach

Due to the differences outlined above, the certification method that results is necessarily dissimilar. Table 6 shows examples from the specifications used for each of these approaches related to the communication problems identified in the fault tree shown in Fig. 5. The certification process for DO-312 compliance is performance based: The document specifies performance goals or necessary minimum probabilities of component and system failure in order to meet the assumptions used in the fault or event trees.

For the performance goal approach, manufacturers must assure compliance with the performance metrics. Although many electro-mechanical components have sufficient heritage to yield an accurate probabilistic assessment, such individual physical component statistics may not hold in a complex system and are not useful for new components or old components operating in new environments, as is anticipated for NextGen. Even greater problems arise in defining probabilities for software and human errors. A good case can be made that such probabilities do not exist or cannot be determined even if they do. In practice, manufacturers either ignore the parts of their systems for which probabilistic data cannot be obtained, they estimate them, or they may even assign arbitrary numbers such as 10^{-4} per flight hour for all software errors.

In contrast, the analysis produced by STPA results in a specification of behavioral constraints (requirements) that must be satisfied by the manufacturers to have their products certified for use in the system. This approach, in fact, was the one used for TCAS where behavioral requirements (using AND/OR tables as shown earlier in Fig. 10) were provided and the manufacturers of TCAS boxes used standard assurance methods, such as DO-178B, to show that the requirements were satisfied.

4. Conclusions

The method being used in the hazard analysis for NextGen incorporates traditional hazard analysis methods (fault tree and event trees). The basic problem is that these methods were developed almost 50 years ago for systems composed primarily of hardware and of much less complexity than the transportation systems we are building today. They are not powerful enough to handle accidents involving interactions of components (system design errors) that do not involve failure of individual system components, software, and the cognitively complex tasks being assigned to human operators today. More powerful methods are needed to assure and certify safety in upgrades and changes to the NAS and other transportation systems. This paper describes a possible alternative and applies it to a new ATC procedure called ITP. The alternative approach identifies a larger class of causes and provides a more comprehensive set of requirements and potential design solutions to prevent accidents.

But the goal of this paper was not simply to compare two approaches to safety analysis. Most important is to understand the important philosophical differences that underlie them and the implications of these differences for selecting appropriate tools for safely integrating changes into complex transportation infrastructures. Different assumptions underlie the analysis methods and thus the causal factors that can be identified by each and the certification method implied for introducing new or altered procedures into the air transportation system.

5. Acronyms

ACAS	Airborne Collision Avoidance System
ADS-B	Automatic Dependent Surveillance-Broadcast (satellite-based navigation)
ATC	Air Traffic Control
ATSA–ITP	Airborne Traffic Situational Awareness–In-Trail Procedure
CA	Control Action
EUROCONTROL	European Organisation for the Safety of Air Navigation
FAA	Federal Aviation Administration
FC	Flight Crew
FTA	Fault Tree Analysis
ITP	In-Trail Procedure
NAS	National Airspace System
NextGen	Operational and technological enhancements envisioned for NAS
NMAC	Near Mid-Air Collision
OH-x	Operational Hazard, number "x", defined in DO-312 (RTCA, 2008)
STPA	Systems Theoretic Process Analysis (Hazard Analysis)
STAMP	Systems Theoretic Accident Model and Processes
TCAS II	Traffic Collision Avoidance System

References

Checkland, Peter, 1981. Systems Thinking, Systems Practice. John Wiley & Sons, New York.

Dekker, Sidney, 2004. Ten Questions about Human Error. CRC.

Dekker, Sidney, 2006. A Field Guide to Understanding Human Error. Ashgate Publishing.

EUROCAE ED-78A/RTCA DO-264, 2002. Guidelines for Approval of the Provision and Use of Air Traffic Services Supported by Data Communications, March 2002.

FAA, 2011. What is NextGen. <http://www.faa.gov/nextgen/why_nextgen_matters/what/>.

FAA, 2010. Interim Policy and Guidance for Automatic Dependent Surveillance Broadcast (ADS-B) Aircraft Surveillance Applications Systems Supporting Oceanic In-Trail Procedures (ITP).

Fleming, Cody H., 2012. Spencer, Melissa, Leveson, Nancy, Wilkinson, Chris. Safety Assurance in NextGen, NASA Technical. Report NASA/CR-2012-217553.

Leveson, N.G., Mats, P.E., Heimdahl, H., Hildreth, H., Reese, Jon D., 1994. Requirements specification for process-control systems. IEEE Transactions on Software Engineering, SE-20, no. 9, September, 1994.

Leveson, N.G., 2000. Intent specifications: an approach to building human-centered specifications. IEEE Transactions on Software Engineering, SE-26, no. 1, January 2000.

Leveson, Nancy G., 1995. Safeware. Addison-Wesley.

Leveson, Nancy G., 2012. Engineering a Safer World. MIT Press.

Roland, Harold E., Moriarty, Brian, 1983. System Safety Engineering and Management. John Wiley & Sons, New York.

RTCA, 2008. Safety, Performance and Interoperability Requirements Document for the In-Trail Procedure in the Oceanic Airspace (ATSA-ITP) Application, DO-312, Washington DC, June 19, 2008.