Safety and Reliability in Aerospace Cyber-Physical Systems

High Impact Applications: Future and Present

Forecasts indicate a significant increase in demand in air traffic, ranging from a factor of two to three by 2025. The ensuing shortfall could cost the U.S. billions of dollars annually in lost productivity, increased operational costs, higher fares, and lost value from flights that airlines must eliminate to keep delays to an acceptable minimum. In short, U.S. competitiveness depends upon an efficient, high capacity, flexible, safe, and environmentally compatible air transportation system.

We must develop processes and methods of assuring safety, security and reliability properties in a possibly distributed, safety-critical, real-time infrastructure. The characteristics of an air vehicle network populated by commercial, military and general aviation aircraft that can be either human-piloted or computer-controlled, need to be accurately modeled, analyzed and validated, including the interaction between components, such as autonomous and non-autonomous vehicles. Qualities such as safety, security and reliability must be maintained throughout the development, operation and evolution of such a system, in the face of control, computational, and communication challenges inherent to atmospheric flight in a shared environment. Our goal must be to develop techniques that enable the validation and eventual certification of safety, security and reliability properties via (formal) modeling and analysis as well as simulation and experimentation.

Challenges for Aerospace CPS

There is a great need to develop a methodology that enables the assessment and enforcement of a specified level of safety, security and reliability for complex software, as well as its interactions with aerospace hardware and human supervisors. One of the most difficult tasks is assuring that the system requirements specification matches the implementation. Furthermore, in order to assure that the properties are maintained throughout the lifecycle of the system, metrics must be developed to assess the safety, security and reliability qualities of the implementation, a task that is often both qualitative and quantitative in nature. These metrics, if developed early enough in the design process, can potentially be used to evaluate and perform tradeoffs of design options in terms of maintainability or even certification of the system. Finally, a rigorous framework must be established to maintain the system's safety, security and reliability throughout its evolution via fully documented property tracablility, including new software releases, and equipment upgrades. The ability to cleanly specify interaction requirements, in order to make system interactions more predictable, is critical to the composability of components and their qualities. This requires that all components have their assumptions explicit and machine checkable, and requires the development of the computer-aided synchronization of interface specifications and software code. Central to this problem will be the viability of this structure to enable certification and eventual deployment of any new technologies.

Quality Assessment with Measurement Driven Analysis and Design. There is a great need to develop a methodology that enables the assessment and enforcement of a specified level of safety, security and reliability for complex software, as well as its interactions with avionics hardware and human supervisors. One of the most difficult tasks is assuring that the system requirements specification matches the implementation. Furthermore, in order to assure that the properties are maintained throughout the lifecycle of the system, metrics must be developed to assess the safety, security and reliability qualities of the implementation. These metrics, if developed early enough in the design process, can potentially be used to evaluate and perform tradeoffs of design options in terms of maintainability or even certification of the system. Finally, a rigorous framework must be established to maintain the system's safety, security and reliability throughout its evolution via fully documented property tracablility, including new software releases, and equipment upgrades. Central to this problem will be the viability of this structure to enable certification and eventual deployment of any new technologies.

Fault Diagnosis, Recovery and Performance Degradation. Safe, secure and reliable aircraft composed with several other safe, secure and reliable aircraft do not necessarily make a safe, secure and reliable networked airspace system in a distributed environment where component interactions are not regulated by a central authority. Furthermore, software for aircraft flight management systems, collision avoidance and pilot/controller alerting functions are complex in nature. A fault protection envelope for both hardware and software is necessary so that a failure does not trigger a hazardous situation. Interface requirements and constraints must be developed which take into account all possible software, avionics and pilot and ground controller interactions under degraded conditions**.** In the near future, human-piloted vehicles will be required to share the same airspace with computer controlled, autonomous vehicle. Interactions between a human piloted plane and an autonomous vehicle must be modeled in order to design for a 'fault containment region' to contain violations of trustworthiness.

Architectures for Aerospace CPS

The architecture of aerospace CPS must focus on problems that arise uniquely from execution requirements of applications developed for the aerospace platform: a violation in liveness (progress) necessarily incurs a violation in safety. The common denominator is hardware and software that make up the overall system architecture's execution framework. That framework manages data flows, communication bandwidth, inter-process synchronization, and scheduling. It must do so in ways that meet real-time processing constraints, supports software architectures that provide trustworthiness, and meets quality of service requirements when allocating on-board communication bandwidth between computers, sensors, and actuators. This architecture could include real-time monitoring and reaction capabilities supported in hardware, and could be addressed in terms of embedded programmable hardware. In order to eliminate (or limit) error propagation, rapid detection and recovery is needed. A hardware framework for high-performance and high-dependability in which error detection and recovery firmware and programmable hardware constitute a reliability engine fully integrated with the processor (or implemented as an external FPGA-based device) is a possible solution to this issue. The application can be instrumented to instruct the processor about the desired level and type of runtime checking. Hardware support should be coupled with a software architecture to support rapid recovery.

3) High Reliability and Security Infrastructure. Rapid response to random errors and malicious attacks entail that the system must make correct decisions in an automated and autonomous manner. Integrity of the reliability and security infrastructure becomes paramount and is typically the most difficult quality to assure, as the precise conditions of field failures and security threats are difficult to anticipate or reproduce with enough realism to verify the capabilities of the infrastructure. The infrastructure must be designed to manage redundant resources across interconnected nodes, foil security threats, detect errors in both the user applications and the infrastructure components, and recover quickly from failures when they occur, in order to qualify and quantify benefits in terms of assessment, verification and maintenance of reliability, security and safety constraints.

Technical Challenges to Interface and Manipulate the Physical World

A major focus must be on the development of systems which are verifiably robust, in order to allow operation under situations with significant external environmental uncertainty combined with potentially rapidly changing conditions and objectives. Systems for optimization, scheduling, control and communication will have to interoperate efficiently and in real-time. Research must explore the organizing principles of such interactions and should investigate what are the appropriate abstractions and what is the appropriate architecture for integrating control with communication. Particular attention must be paid to developing an architecture that is amenable to future evolution, and which supports services that significantly shorten design cycle-time.

With the availability of significant computational power, there has been rapid development of tools for direct modeling of dynamical systems, leading to explicit construction of accurate models incorporating the details of both digital state-transitions and vehicle dynamics. It is imperative that important systems features be carefully and expertly identified and extracted, and that unwanted details be replaced with simple uncertainty descriptions. Designs and verification methods must be robust to model inaccuracy. The following issues dominate possible research directions: (i) Computational tractability: algorithms that scale; (ii) Predictable interaction between electromechanical (physical) and cyber (computation/communication) components; (iii) Verification of systemwide safety/reliability properties; (iv) Information/state/health management via communication over wired/wireless links.

Human decision-makers may have limited time to respond to unexpected events in which they must be the responsible agent to either carry out a decision, or veto an automated decision. In order to assess the vulnerability of an advanced system to excessively slow (or inappropriate) human decisions in emergency, two parallel approaches are required: (i) Identifying, through failure modes analysis, a catalogue of possible failures that will have an impact on human response. In doing so, we must be cognizant of the ironic tradeoff, that the less likely the failure, the longer it will take the human to respond appropriately when it does occur. (ii) Developing a computational model of the system response time for low-expectancy events so that model-predicted time required can be played off against analyst-determined time available.

Biography: Natasha Neogi is a professor in the Department of Aerospace Engineering at the University of Illinois, Urbana-Champaign. She holds joint appointments with the Departments of Computer Science and Electrical and Computer Engineering. She is an affiliate of the Institute for Aviation, as well as the Coordinated Sciences Laboratory, where her own research lab, the Autonomous Laboratory for Autonomous Embedded Systems (ALEAS), is housed. Her research focuses on the verification and validation of safety, security and reliability properties in distributed and embedded cyberphysical systems. She received her Ph.D. from MIT in 2002.