A New Approach for Designing Safer Collision Avoidance Systems

Mykel J. Kochenderfer
James P. Chryssanthacopoulos
Roland E. Weibel

14 June 2011

1951–2011
LINCOLN LABORATORY
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

This work is sponsored by the Federal Aviation Administration under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

Aircraft Collision Avoidance

Traffic Alert and Collision Avoidance System (TCAS)

Green arc indicates required vertical rate

- TCAS uses range, altitude, and bearing measurements of local traffic
- Traffic alerts are issued to pilots to assist in visual acquisition
- Resolution advisories (RAs) advise pilots to climb or descend
- Required decades of development

Problem Statement

- **NextGen procedures will require changes to TCAS logic**
  - **Improvements to surveillance will likely be insufficient**
  - **Revising the TCAS logic is difficult and may be inadequate**
  - **A different approach to logic development may be required**
- **Challenges to collision avoidance logic development**
  - **Uncertainties due to sensor noise and aircraft behavior**
  - **Aircraft performance and operational constraints**
  - **Need to maximize safety while minimizing alerts**
- **Solution is a decision-theoretic approach**
  - **Uses explicit models of sensor and dynamic uncertainty**
  - **Optimizes logic according to performance measure**
  - **Leverages advances in computation and algorithms**

Approaches to Collision Avoidance

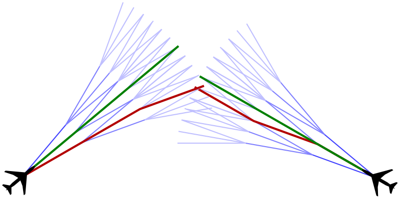

- **Nominal (green) used by TCAS**
  - **Does not consider hazardous low-probability events**
- **Worst case (red)**
  - **Potential for high rate of unnecessary course deviation**
- **Probabilistic (blue)**
  - **Improved robustness due to accounting for relative likelihood of all possible future trajectories**

Approaches to Logic Development

Legacy TCAS Development Cycle

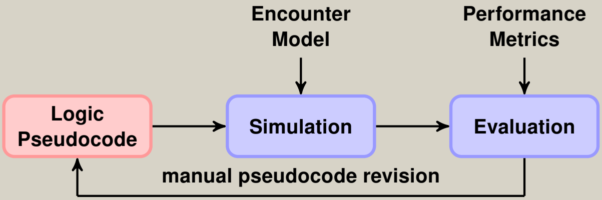

Encounter Model

Performance Metrics

Logic Pseudocode

Simulation

Evaluation

manual pseudocode revision

Logic Optimization Approach

Encounter Model

Performance Metrics

Optimization

Logic Table

Timeline

- **FY08: Initial concept for UAS (Report ATC-356)**
  - Internal seed funding
  - Applied approach to sensor systems and platforms
- **FY09: Initial exploration for TCAS (Report ATC-360)**
  - Funding from TCAS Program Office
  - Experimented with different methodologies on 2D model
- **FY10: Development and analysis (Report ATC-371)**
  - Developed 3D model, incorporated sensor noise
  - Focus primarily on single, unequipped intruder
- **FY11: Further enhancement (ongoing)**
  - Enhanced multithreat and coordination
  - Compare to alternative approaches

Logic Development and Usage

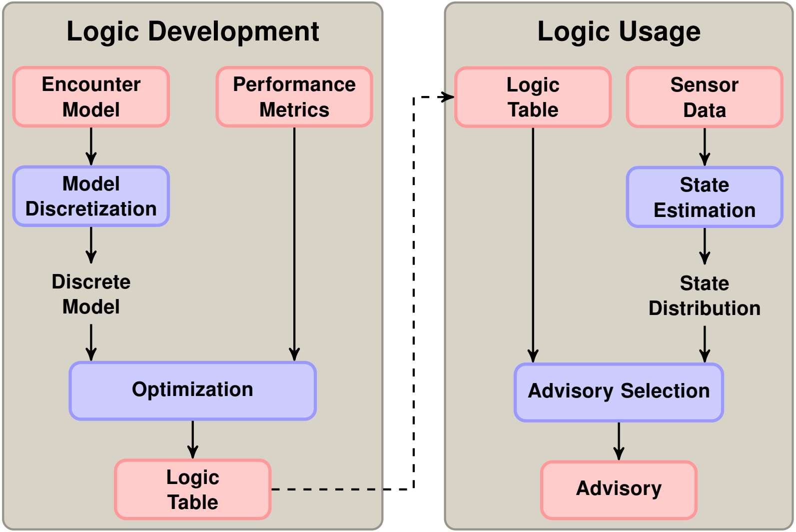

Simple 2D Encounter Model

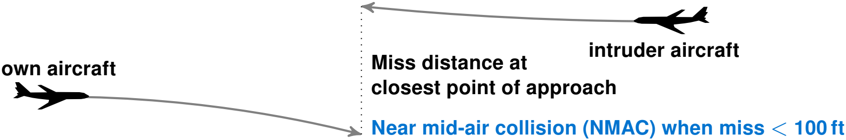

Near mid-air collision (NMAC) when miss $<$ 100 ft

Model Assumptions

- Head-on encounter
- No turning
- Constant closure rate
- Random vertical accelerations
- Advisory may be changed

State Variables

- Relative altitude
- Time to closest approach (τ)
- Own vertical rate
- Intruder vertical rate
- Advisory state (for 5 second pilot delay)

Cost Metric

- **NMAC**: 1
  - Accrued when an NMAC occurs
- **Alert**: 0.01
  - Accrued when an initial advisory is issued
- **Strengthening**: 0.009
  - Accrued when an advisory is strengthened
- **Reversal**: 0.01
  - Accrued when an advisory is reversed
- **Clear of Conflict**: -0.0001
  - Accrued every time step when no advisory is on display

These costs were chosen arbitrarily but can be adjusted to accommodate safety or operational requirements

Logic Development and Usage

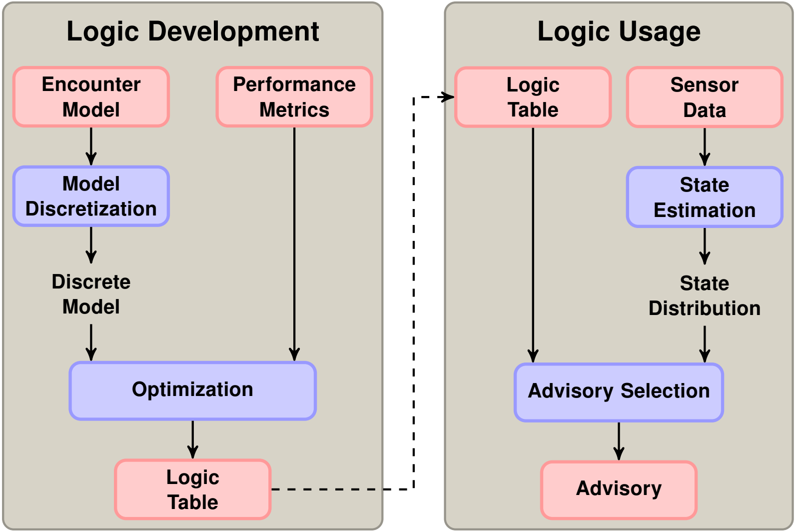

Logic Development

Encounter Model

Performance Metrics

Model Discretization

Discrete Model

Optimization

Logic Table

Logic Usage

Logic Table

Sensor Data

State Estimation

State Distribution

Advisory Selection

Advisory

Model Discretization

- **Discretize state space using a (5-D) grid**
- **Vertices correspond to discrete states (8.7 million states)**
- **Grid coarseness affects the accuracy of the model**
- **Use encounter model to determine state transition probabilities**

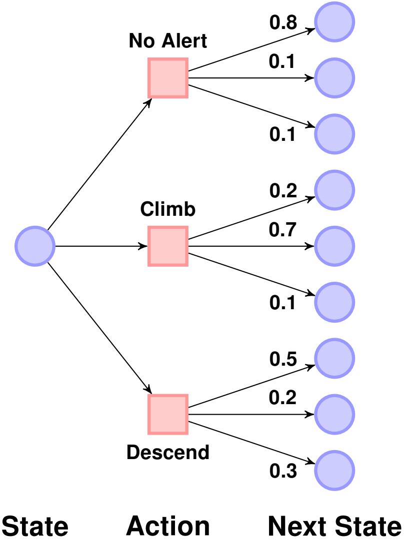

Dynamic Programming

- The expected cost table contains the expected cost for each action (and then continuing with the optimal logic) from each state
- Dynamic programming efficiently computes this table
- Table used in real time to choose actions

Notional expected cost table

State					Expected cost		
Rel. alt.	**Time to closest**	**Own vert.**	**Int. vert.**	**RA state**	**No alert**	**Descend**	**Climb**
100	19	1500	−1000	**No alert**	0.0144	0.4215	0.0190
200	20	0	0	**No alert**	0.0449	0.0339	0.4251
. . .	. . .	. . .	. . .	. . .	. . .	. . .	. . .

Logic Development and Usage

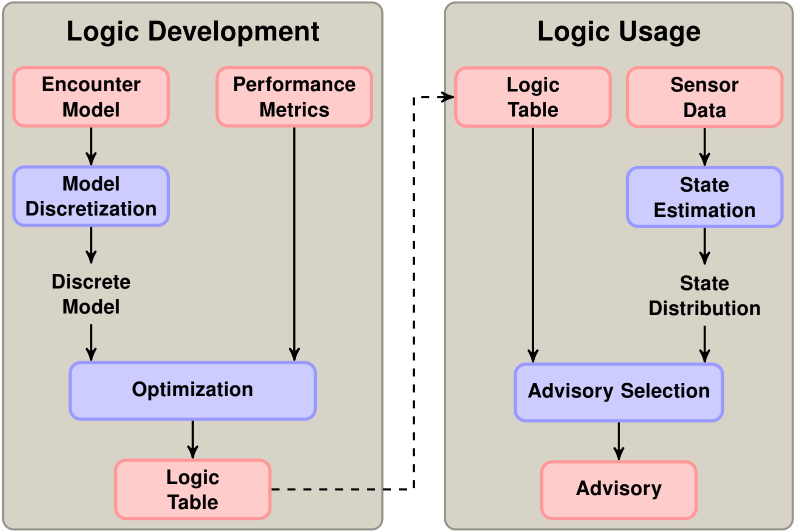

Logic Development

Encounter Model

Performance Metrics

Model Discretization

Discrete Model

Optimization

Logic Table

Logic Usage

Logic Table

Sensor Data

State Estimation

State Distribution

Advisory Selection

Advisory

Optimal Action Plot

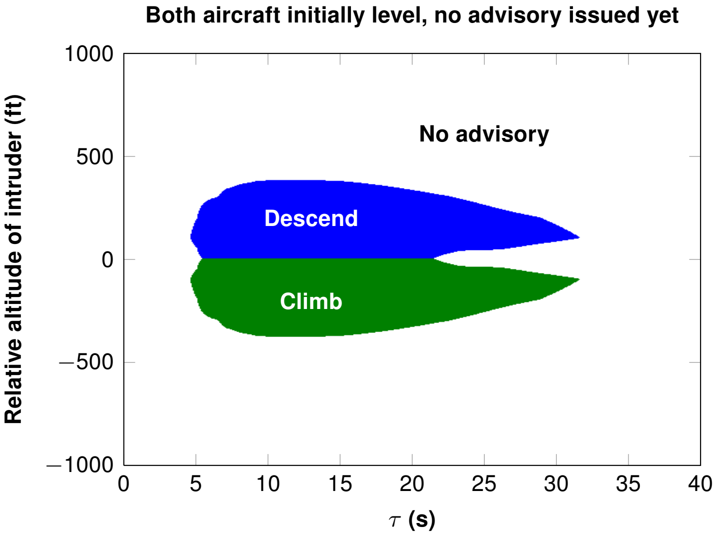

Optimal Action Plot

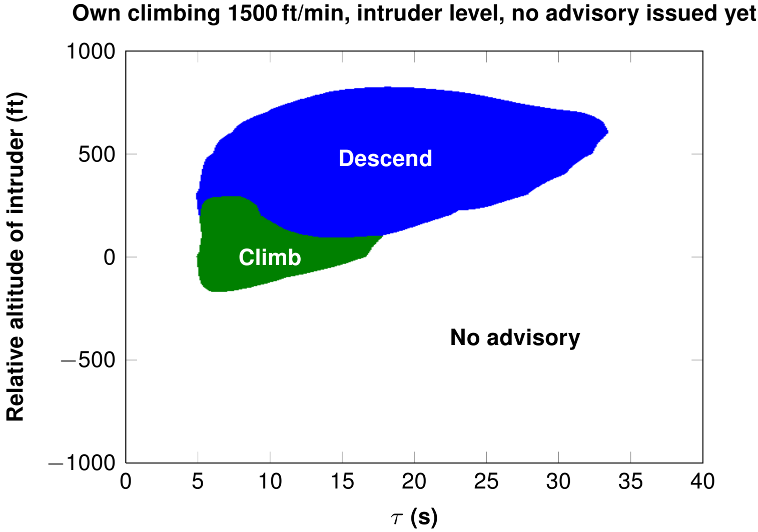

Optimal Action Plot

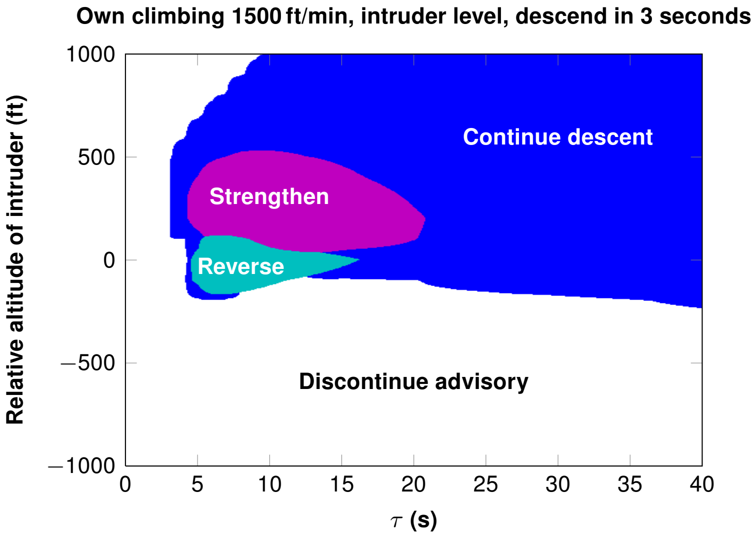

Results

Metric	DP Logic	TCAS Logic
NMACs	3	169
Alerts	690,406	994,317
Strengthenings	92,946	40,470
Reversals	9569	197,315

Performance on metrics can be traded by adjusting costs

Safety curve

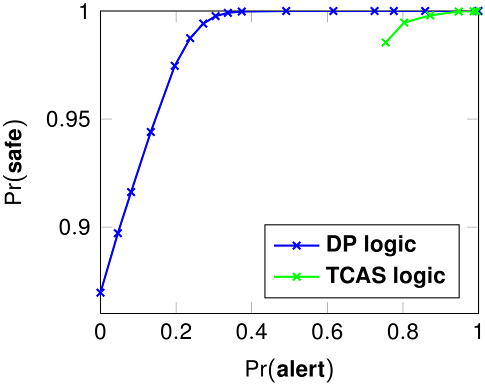

- DP curve generated by varying alert cost
- TCAS curve generated by varying sensitivity level

Example Encounter

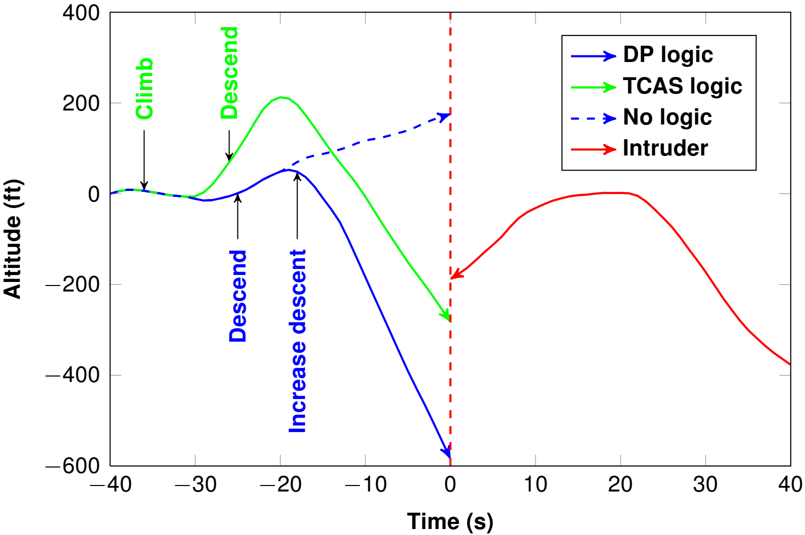

NMAC and Alert Slices

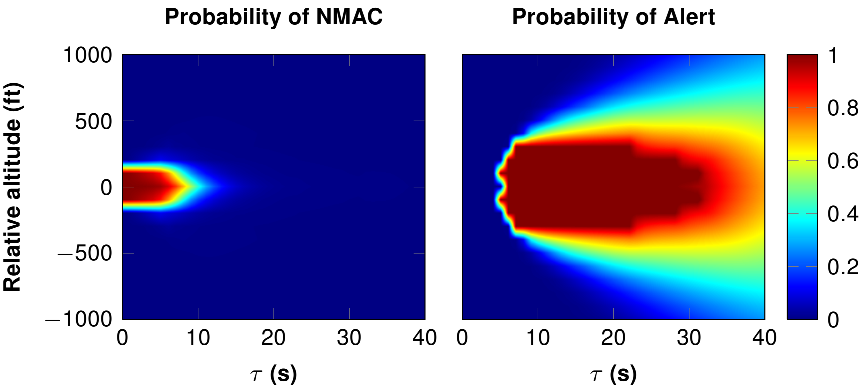

- Both aircraft level and no advisory issued
- Computed for full state space using DP algorithm

Robustness to Modeling Errors

- **Model used for optimization will not perfectly match real world**
- **Example plots vary environment model, but keeping the DP logic fixed (at 8 ft/s^2)**
- **DP outperforms TCAS even with modeling error**

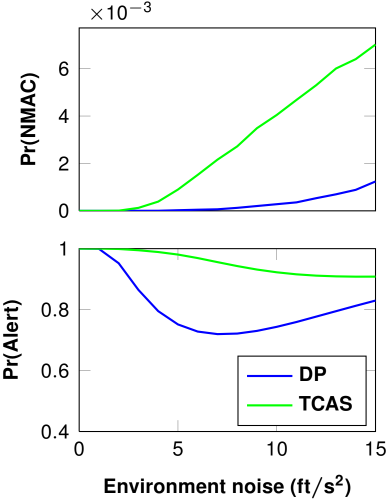

Collision Avoidance in 3-D

- When aircraft are maneuvering horizontally, τ cannot be known exactly
- Use probabilistic model to infer distribution over τ and weight cost appropriately
- Explored different methods for estimating τ

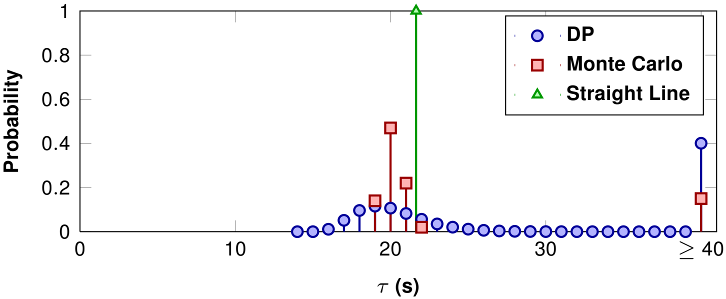

Results

	DP Logic			
Metric	**DP**	**Monte Carlo**	**Straight Line**	**TCAS Logic**
NMACs	2	11	1	101
Alerts	540,113	400,457	939,745	994,640
Strengthenings	39,549	37,975	26,485	45,969
Reversals	1242	747	129	193,582

DP logic significantly outperforms TCAS

Safety Curve

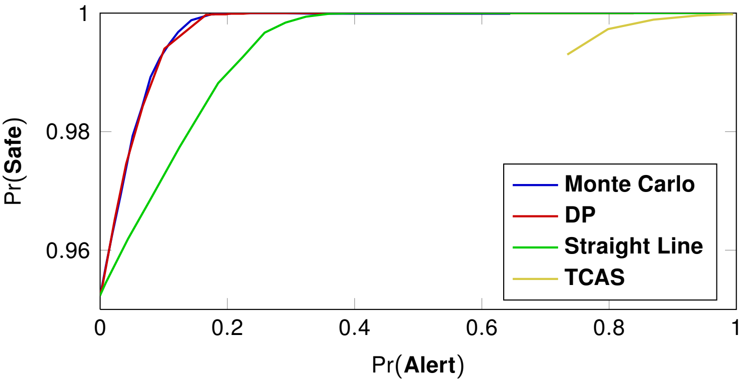

Using a distribution over τ is better than a point estimate

Probabilistic Pilot Response

- **Pilots do not respond according to the standard deterministic response model (5 second delay, 1/4 g maneuver)**
- **Modifying model to capture pilot variability leads to significantly improved performance**

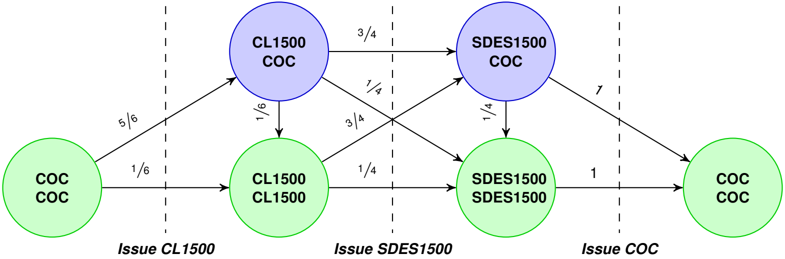

Ongoing research on more sophisticated models

Sensor Noise

- TCAS (and other systems) use most likely state (MLS) estimate
- Investigated use of full state distribution in new logic
- Involves averaging expected cost

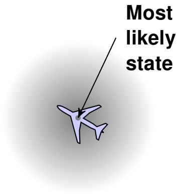

State distribution

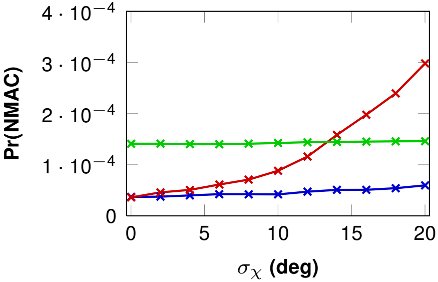

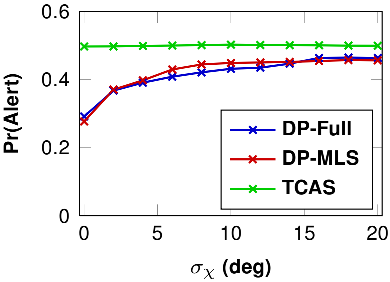

Coordination

Coordination is critical when both aircraft are equipped

- Different views of the world can lead to incompatible maneuvers (e.g., climb/climb)
- Aircraft communicate the intended sense of the maneuver
- Simply restricting action set to compatible advisories significantly improves safety
- Further work will investigate:
  - Incorporating broadcasts into model
  - Interoperability with legacy TCAS

Multiple Threat Logic

- One of the most complex pieces of TCAS logic
- Currently exploring different methods for extending approach to handle multiple threats:
  - Command arbitration: choose one of the actions associated with an intruder
  - Utility fusion: combine expected costs
  - Global strategy: find path that best avoids all intruders

	Command Arbitration		Utility Fusion			
	Closest	**TCAS-like**	**Max-sum**	**Max-min**	**Global**	**TCAS**
Pr(NMAC)	$9.662 \cdot 10^{-3}$	$4.810 \cdot 10^{-3}$	$2.206 \cdot 10^{-3}$	$1.676 \cdot 10^{-3}$	$2.714 \cdot 10^{-3}$	$2.960 \cdot 10^{-3}$
Pr(Alert)	0.640	0.682	0.604	0.682	0.581	0.764
Pr(Str.)	0.137	$8.467 \cdot 10^{-2}$	$6.991 \cdot 10^{-2}$	$8.925 \cdot 10^{-2}$	0.483	$4.430 \cdot 10^{-2}$
Pr(Rev.)	$3.644 \cdot 10^{-3}$	$5.290 \cdot 10^{-3}$	$6.632 \cdot 10^{-3}$	$9.538 \cdot 10^{-3}$	$6.996 \cdot 10^{-3}$	$6.550 \cdot 10^{-3}$

Multiple Threat Stress Test

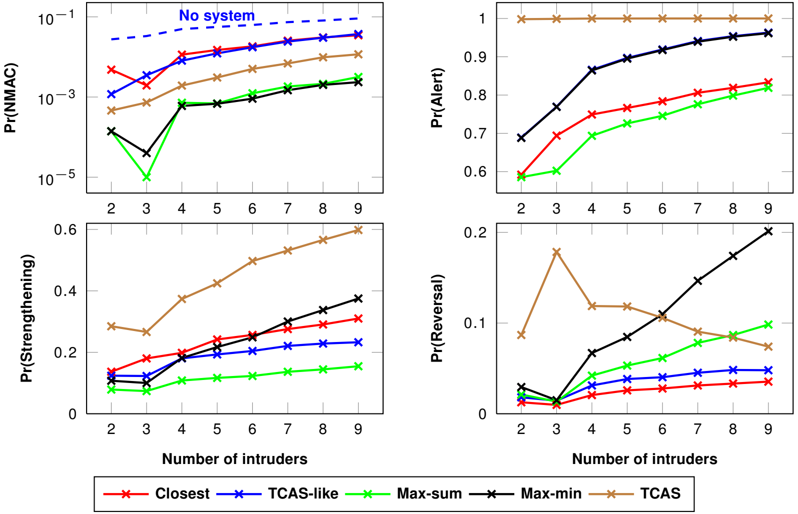

Alternative Approaches

Path planning

- Find a conflict-free path, delay alert if possible
- Uses deterministic models
- Analytic solutions
- Insensitive to dimensionality of model
- Does not account for future state information

Probability thresholding

- Computes probability of conflict for different actions
- Alert if probability exceeds some threshold
- Uses probabilistic models
- Does not account for changing action in future

Performance can be poor with increased trajectory uncertainty

Publications

- M. J. Kochenderfer and J. P. Chryssanthacopoulos, "Robust airborne collision avoidance through dynamic programming," Massachusetts Institute of Technology, Lincoln Laboratory, Project Report ATC-371, 2010.
- M. J. Kochenderfer, J. P. Chryssanthacopoulos, L. P. Kaelbling, T. Lozano-Perez, and J. K. Kuchar, "Model-based optimization of airborne collision avoidance logic," Massachusetts Institute of Technology, Lincoln Laboratory, Project Report ATC-360, 2010.
- J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Decomposition methods for optimized collision avoidance with multiple threats," in IEEE/AIAA Digital Avionics Systems Conference, Seattle, Washington, 2011.
- M. J. Kochenderfer, J. P. Chryssanthacopoulos, and R. E. Weibel, "A new approach for developing safer collision avoidance systems," in USA/Europe Air Traffic Management Research and Development Seminar, Berlin, Germany, 2011.
- J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Hazard alerting based on probabilistic models," in AIAA Guidance, Navigation, and Control Conference, Portland, Oregon, 2011.
- J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Analysis of open-loop and closed-loop planning for aircraft collision avoidance," in IEEE International Conference on Intelligent Transportation Systems, Washington, DC, 2011 (under review).
- J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Accounting for state uncertainty in collision avoidance," in Journal of Guidance, Control, and Dynamics, 2011 (in press).
- J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Collision avoidance system optimization with probabilistic pilot response models," in American Control Conference, San Francisco, Calif., 2011.
- M. J. Kochenderfer and J. P. Chryssanthacopoulos, "Partially-controlled Markov decision processes for collision avoidance systems," in International Conference on Agents and Artificial Intelligence, Rome, Italy, 2011.
- M. J. Kochenderfer and J. P. Chryssanthacopoulos, "A decision-theoretic approach to developing robust collision avoidance logic," in IEEE International Conference on Intelligent Transportation Systems, Madeira Island, Portugal, pp. 1837–1842, 2010.
- M. J. Kochenderfer, J. P. Chryssanthacopoulos, and P. Radecki, "Robustness of optimized collision avoidance logic to modeling errors," in IEEE/AIAA Digital Avionics Systems Conference, Salt Lake City, Utah, 2010.

Summary and Next Steps

- Automated optimization based on probabilistic models
  - Focuses human engineering effort on developing models and metrics
  - Improves safety while reducing alert rate
  - Is sensor agnostic
- Logic specification in terms of expected cost table
  - Reduces implementation burden of manufacturers
  - Eases verification process
  - Simplifies the update process in response to evolution of airspace
- Next steps:
  - Development: coordination and interoperability issues
  - Analysis: safety and operational acceptability
  - Certification: develop plan, engage safety community