

Conversational Policy: A case study in air traffic control

Henry Hexmoor
University of North Dakota
Grand Forks, North Dakota, 58202
hexmoor@cs.und.edu

Abstract. *We present a model of behavior that unifies behavior-based acting and conversational policies for speaking. This model is applied to the domain of air traffic control.*

Keywords. Multi-agent Systems

1. Introduction

General conversations among agents include goals and plans. We want to consider a subset of conversations that are routine and agents behave reactively with little or no reasoning. In AI and robotics, the class of systems that model routine reactive behaviors are called *behavior-based systems* [1]. These system use simple modeling techniques such as state diagrams. Routine communication among software agents have been modeled by conversational policies [3]. We believe, behavior based models of robotics that model physical actions are very similar to simple conversation models of conversation policies. In this paper we present a representational model that unifies these two models.

To demonstrate our model, we will use our simulator and tower controller software agents we have developed for air traffic control [5,6]. We will model two of the most essential functions performed by an air traffic control for (a) ensuring airplanes do

not collide, and (b) supervising the landing of airplanes.

The remainder of this paper is organized by first discussing our unified model. We then talk about the domain of air traffic control and then present application of our model in that domain.

2. A unified model for routine conversations and physical acts

We use the following concepts in our representations:

- **Speak States**— These are the meaningful mental states for our agent in keeping track of a conversation
- **Act States**— These are the meaningful mental states for our agent in keeping track of actions
- **Spoken words/phrases**— These are the words/phrases either heard or spoken by our agent.
- **Acts**— These are the physical actions that can be performed by our agent.
- **Physical conditions**— These conditions are meaningful conditions that affect the agent's decision making.

To model a conversation policy, we introduce two types of simple production-like structures we will call transitions, shown below.

physical condition * spoken word/phrase * speak state
 → speak state
 physical condition * speak state
 → spoken word/phrase

The first transition type is called *speak state transition* and accounts for changes in state. The left hand side of the transition consists of a physical condition, what the agent hears, and the current speak state. The right hand side of the transition is the agent's next speak state.

The second transition type is called *speak transition*. The left hand side of the transition is a physical condition, and the current speak state. The right hand side of this transition is the agent's next spoken word/phrase.

To model physical actions of an agent in reactive behaviors, we introduce two other types transitions, shown below.

physical condition * spoken word/phrase * act state
 → act state
 physical condition * act state
 → act

The first transition type is called *act state transition* and accounts for changes in state. The left hand side of the transition consists of a physical condition, what the agent hears, and the current act state. The right hand side of the transition is the agent's next act state. The second transition type is called *act transition*. The left hand side of the transition is a physical condition, and the current act state. The right hand side of this transition is the agent's next act. In the examples of this paper we will only demonstrate the use of transitions for conversation policy and not show uses of the

physical actions. We have used these types of rules in our other work such as [4].

3. Communication in Aviation

The FAA and other Aviation organizations have set standards for communication between pilots and air traffic controllers. In addition to these standards, individual airlines enforce their own standards for communication between their employees. If either the air traffic controller or the pilot fail to adhere to these standards, the safety of all airplanes (and their passengers) could be adversely affected. This makes the ATC a good example for study of encoding conversational policies.

The FAA outlines very specific rules for handling a potential mid-air collision. This set of rules comprises the FAA's official collision avoidance policy. This policy consists of standardized communication rules and procedures designed to ensure collision avoidance is performed in a consistent, accurate, and timely manner. It is important to note that while in flight, pilots must adhere to either Visual Flight Rules (VFR) or Instrument Flight Rules (IFR). For the purposes of this paper, we have limited our conversation policies to pilots that are using IFR. In the context of collision avoidance, pilots using IFR are dependent on the ATC to provide warnings about a potential collision, as well as recommended course corrections to avoid a collision.

The FAA's collision avoidance policy defines two levels of collision severity. The first (and lower priority) level is referred to as a *traffic call*. A traffic call is a "heads up" from the ATC to a pilot, informing him or her about surrounding traffic. When a traffic call is issued, there is no immediate threat of collision, and the distance between

airplanes is still acceptable. The second (and highest) level is referred to as a *traffic alert*. A traffic alert means that, if the airplanes receiving the traffic alert make no course correction, a mid-air collision will likely occur. At this point, the distance between two or more airplanes has become unsafe, and the collision avoidance policy must be promptly invoked to avoid a mid-air collision.

The landing conversation policy is broken down into two different types of landing, emergency and standard. Based upon which type of landing is requested by the pilot, a conversational policy is followed by taking into account the priority of the request, and the many responses possible at each stage in the conversation. An emergency landing is simply defined as a special situation that requires immediate attention and priority over other regular landings and other traffic. In most cases, the ATC will give an emergency landing conversation nearly equivalent priority to a collision avoidance conversation.

4. Collision avoidance conversation policy

Although the FAA provides a set of well-defined rules for resolving a potential mid-air collision between two airplanes, consolidating these rules into a conversation policy for collision avoidance is a difficult task.

The following are the agent *speak states*:

- 0** – Start state
- 1** – ATC has issued a traffic call
- 2** – ATC has issued a traffic alert with expedite
- 3** – Pilot has responded with “Roger”, “Affirmative”, or “Wilco”. This means pilot has committed to complying with the ATC's advice.

4 – Pilot has responded with “Unable” (this is a very rare situation). This means pilot cannot comply with the ATC's advice.

5 – Pilot ignored (did not respond to) the ATC.

6 – Pilot has completed the collision avoidance maneuver recommended by the ATC. This is one possible final state. ATC can continue to other planes.

State 0 is the start of collision avoidance policy. From either state 0 or state 1, the ATC may exit this conversation to start a new conversation, attend to another existing conversation, or simply to terminate this conversation (many collision avoidance conversations never get beyond the traffic call stage). Leaving the collision avoidance conversation is allowed at these nodes because there is not yet a critical collision that requires immediate attention. Once the policy reaches state 2 (which means a traffic alert has been issued by the ATC), ATC is not allowed to leave the conversation until the current collision avoidance conversation has been completed. The only exception to this rule is if the ATC needs to leave the current conversation to issue a traffic alert to other airplanes involved in the collision avoidance procedure. Once the ATC has issued a traffic alert, he or she expects to receive a reply from the pilot within a short period of time. The vast majority of the time the pilot will reply to the traffic alert and inform the ATC he or she will comply with the ATC's recommended course adjustment.

The following are the set of available *words/phrases*:

- S1** - [Call sign] **Traffic** [Position] [Distance] [Direction Traffic is Moving] [Traffic Type] [Altitude]
- S2** - Pilot says "Roger. Looking for traffic."
- S3** - [Call sign] **Traffic Alert** [Position] [Distance] [Direction] [Traffic Type]

[Altitude]. **Advise you to turn** (left or right) **on heading** [Heading] **and (descend or climb) to altitude** [Altitude] **Expedited**.

S4 - Pilot has responded with "Roger", "Wilco", or "Affirmative". Means pilot has committed to complying with the ATC's advice. Ex: "Roger NW318 is turning left on heading 270 and climbing to 20,000 feet."

S5 - Pilot says "Unable" to the ATC.

S6 - No words are spoken. This is considered to be "intentional silence" when the agent is listening.

S7 - Pilot contacts ATC to discuss/clarify (and probably argue about) the near collision.

The following are the *physical conditions*:

P1 - Possibility of collision is detected by ATC.

P2 - Short time is passed and the collision situation is not yet critical.

P3 - Collision situation is continued.

P4 - A severe collision threat is detected.

P5 - Collision is successfully resolved.

In the following *speak state transitions*, each agent's word/phrase is noted by a prefix of "ATC:" or "Pilot:"

P1 → 0
P1 * 0 * ATC:S1 → 1
P4 * 0 * ATC: S3 → 2
P2 * 1 * Pilot:S2 → 1
P2 * 1 * Pilot:S7 → 0
P3 * 1 * ATC:S3 → 2
2 * Pilot:S4 → 3
2 * Pilot:S5 → 4
4 * ATC:S3 → 2
2 * Pilot:S6 → 5
P4 * 5 * ATC:S3 → 2
P5 * 5 * ATC:S6 → 6
3 * ATC:S6 → 6

The following are the *speak transitions*:

P1 * 0 → ATC:S1
P2 * 0 → ATC:S1
P4 * 0 → ATC:S3
P4 * 1 → ATC:S3
P4 * 5 → ATC:S3

P5 * 3 → ATC:S6
P5 * 5 → ATC:S6
P4 * 4 → ATC:S3
P5 * 4 → ATC:S6

5. Airplane landing conversation policy

The FAA has standards for handling landing requests. We cover majority of cases but some rare conditions of argumentation and negligence are omitted. The entry point for the landing policy is either state 1 or state 2. At this point the ATC will hear the pilot request to land. The entry state may only be entered when an actual request to land has been made. Once the type of landing has been determined, conversation begins. The ATC listens to and corrects any *readback* by the pilot. *Readback* is a pilot's acknowledgement of a controller's transmission that repeats the information that the controller conveyed. Before making the transition to another state, it is important that the ATC wait for the pilot to *readback* the ATC's command and verify that it is correct. Pilot *readback* of ATC commands is where many errors can be caught and corrected, and is a big source of errors by today's air traffic controllers [2]. Unlike the collision avoidance policy, the landing policy may take several minutes to complete. During this time, the ATC will most likely enter and exit the landing conversation multiple times, especially during a regular landing.

The following are the *speak states*:

- 1 - Pilot Has Requested an Emergency Landing
- 2 - Pilot has requested a Regular Landing
- 3 - ATC has advised pilot to change Frequencies
- 4 - ATC has designated runway in which to make an Emergency Landing
- 5 - An "intentional wait" State has taken place, while the ATC "attentively" waits for *readback* from the pilot.

In 2000 Proceedings of the International Conference in Artificial Intelligence

6 - Pilot has *readback* ATC Advisement for Emergency Landing

7 - "Other Words" have been spoken by ATC back to Pilot

8 - Other Commands Issued by ATC

9 - Runway for Normal Landing Advised

10 - An "intentional wait" is taking place while the ATC "casually" waits for *readback*

11 - Pilot has been told to enter a "HOLDING pattern" by ATC

12 - ATC has advised pilot that "the Airport is Closed"

13 - Feedback has been immediately corrected

14 - Pilot has *readback* ATC Advisement

15 -Feedback has been "casually" corrected.

The following are the available *words/phrases*:

S1- Pilot Requests Emergency Landing

S2- Pilot Requests Regular Landing

S3-ATC Ignoring Pilot

S4- ATC tell to enter holding pattern

S5-ATC tells pilot that the Airport is closed

S6- ATC tells pilot to change Radio Frequency

S7- ATC advises Runway for Landing

S8- ATC waits for Readback

S9- Pilot reads back ATC command

S10- Pilot has said "Other Words" to ATC

S11- ATC says "Other Words" to Pilot

S12- ATC Corrects Pilot Readback

S13- Pilot Ignores the ATC

S14- ATC gives Pilot Other Commands

The following are the *physical conditions*:

P1 - There is a discrepancy between Pilot and ATC about runways

P2 - The emergency runway is no good and another one needs to be selected

P3 - A runway cannot be selected or something else is wrong

P4 - The airport is closed for landing

P5 - All runways are unavailable

P6 - A runway is available for normal landing

P7- There is a mis-understanding about holding command

P8 - There is a problem in communication

The following are the *speak state transitions*:

Pilot:S1 → 1
Pilot:S2 → 2
1 * ATC:S7 → 4
1 * ATC:S6 → 3
3 * Pilot:S2 → 2
4 * ATC:S8 → 5
5 * Pilot:S9 → 6
6 * ATC:S12 → 13
13 * ATC: S8 → 5
6 * ATC:S7 → 7
6 * ATC: S14 → 8
8 * ATC:S7 → 5
7 * Pilot:S1 → 1
7 * Pilot:S13 → 14
2 * ATC:S5 → 12
12 * ATC:S8 → 10
2 * ATC:S4 → 11
2 * ATC:S7 → 9
9 * ATC:S8 → 10
10 * Pilot:S9 → 14
14 * ATC:S12 → 15
15 * ATC:S8 → 10

The following are the *speak transitions*:

1 * Pilot:S1 → ATC:S7
P2 * 8 → ATC:S7
P8 * 1 → ATC:S6
4 * ATC:S7 → ATC:S8
P1 * 6 → ATC:S12
16 * ATC:S12 → ATC:S8
P3 * 6 → ATC:S7
P2 * 6 → ATC:S14
12 → ATC:S8
P4 * 2 → ATC:S5
P5 * 2 → ATC:S4
P6 * 2 * Pilot:S2 → ATC:S7
P6 * 9 * ATC:S7 → ATC:S8
P7 * 14 → ATC:S12
P7 * 15 → ATC:S8

6. Summary and Conclusion

We have developed a production-style representational framework that unifies acting and speaking. This is used to model conversational policies that are useful in

In 2000 Proceedings of the International Conference in Artificial Intelligence

routine conversations. We have demonstrated our model in the domain of air traffic control by modeling the tasks of collision avoidance and landing. We believe the similarity between models of acting and speaking allows us to build models of more complex agents that speak as well as perform physical acts.

We plan to extend our model to include shared mental model concepts that facilitate a number of agents to act as a team.

Acknowledgement

We acknowledge the diligent work of Mark Schroeder, Chris Souser, Tim Heng, and Nate Faiman for studying the FAA rules, collecting the information, and organizing them into states and spoken words/phrases for the ATC tasks we have used in this paper.

References

[1] Arkin, R. 1998. Behavior-Based Robotics, MIT Press.

[2] Cardosi, K., Falzarano, P., and Han, S., 1999. *Pilot-Controller Communication Errors: An Analysis of Aviation Safety Reporting System (ASRS) Reports*. DOT/FAA/AR-98/17 (<http://www.volpe.dot.gov/opsad/pubs.html>)

[3] Greaves, M., Holmback, H., and Bradshaw, J. 1999. *What is Conversation Policy?* In Autonomous Agents (Agents-99) Workshop titled Specifying and Implementing Conversation Policies, Seattle, WA.

[4] Hexmoor, H., 1996. Learning Routines, In Agent theories, architectures and languages, M. Wooldrige, J. Mueller, and Milind Tambe (editors), Springer Verlag.

[5] Hexmoor H., and Heng, T. 1999. ATC Tower Simulator: TACUND, Proceedings of **IASTED** on Applied Modeling and Simulation, Cairns, Australia.

[6] Hexmoor, H. and Heng, T. 2000. Air Traffic Control and Alert Agent, In Proceedings of Agents 2000, Barcelona, Spain.