REPRESENTATIONS OF AND STRATEGIES FOR STATIC INFORMATION, NONCOOPERATIVE GAMES WITH IMPERFECT INFORMATION

by

Bradley Bart

B.Math, University of Waterloo, 1996

A THESIS SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
MASTERS OF SCIENCE
in the School
of
Computing Science

SIMON FRASER UNIVERSITY
August 8, 2000

APPROVAL

Name: Bradley Bart

Degree: Masters of Science

Title of thesis: Representations of and Strategies for Static Information, Non-cooperative Games with Imperfect Information

Examining Committee: Dr. Mark Drew
Chair

Dr. James P. Delgrande, Senior Supervisor

Dr. Fred Popowich, Supervisor

Dr. Ramesh Krishnamurti, SFU Examiner

Date Approved:

Abstract

In this thesis, we will examine a card game called MYST, a variant of the Parker Brothers' classic board game Clue. In MYST, a set of cards is divided uniformly among a set of players, and the remaining cards form a hidden pile. The goal of each player is to be the first to determine the contents of the hidden pile. On their turn, a player asks a question about the holdings of the other players, and, through a process of elimination, a player can determine the contents of the hidden pile.

MYST is one of few static information games, wherein the position does not change during the course of the game. To do well, players need to reason about their opponents' holdings over the course of multiple turns, and therefore a sound representation of knowledge is required. MYST is an interesting game for AI because it ties elements of knowledge representation to game theory and game strategy.

After informally introducing the essential elements of the game, we will offer a formal specification of the game in terms of first-order logic and the situation calculus developed by Levesque et al.. Strategies will be discussed including: existence of a winning strategy, randomized strategies, and bluffing. Implementation of some strategies will be discussed.

Acknowledgments

I would like to acknowledge the support of NSERC for money, Simon Fraser University for accepting me into the program, Jim Delgrande for money and for being helpfully skeptic, Oliver Schulte for insights on game theory, and Peter Gvozdjak and Petr Kubon for their LaTeX help. Their support was instrumental to the existence of this thesis.

Contents

List of Figures

Chapter 1

Introduction

Through the annals of history, games have been played as a source of amusement, mental exercise and competition. *Difficult* games have traditionally been the most interesting, usually because a best strategy is neither easily found nor readily apparent. Difficult usually means that there is a finite, but large space of variations, for which a human could not possibly memorize the best move. Games with many variations become the subject of study over time. Chess and bridge are two examples of popular games with a finite, but very large number of variations.

Grand masters of popular games often publish works describing their suggested strategies, or what they feel are some of the best plays. In chess, many have published "opening books" describing the best possible variations for the first 10 to 15 moves by each side. The authors are not mathematicians, however, and so the books do not guarantee that such and such a move is the best possible first move for White, for instance. Such books are a summary of experience and so they represent heuristic strategies.

Mathematical provability about games with large numbers of variations has remained insurmountable even with the help of supercomputers. Since their creation, humans have postulated that computers would be able to out play humans at games. In fact game playing was one focus of initial Artificial Intelligence (AI) research, but the computers of the era were not fast enough to compute the existing heuristics in reasonable time. However, with the advancing speed of supercomputers, we have witnessed computer programs triumph at the world championship level in checkers with Chinook [17] [18] and in chess with Deep Blue [12] [19]. These recent triumphs have restimulated research in game playing.

Checkers and chess are representatives from the class of games with *perfect information.*

Informally, perfect information means that at any point during the game, each player knows the position with absolute clarity. Furthermore, each player can hypothesize the next move and derive the next position, again with absolute clarity. Given enough resources, a program could search the entire space of positions and determine the best moves for each player, using backwards induction. However, this was impractical for live play and so heuristics and search were coupled together to achieve a locally best result. GO is another example of a game with perfect information, but the search space is still too large for a computer to do an efficient heuristic search.

We can contrast games with perfect information against those with *imperfect information*. Informally, imperfect information means that at some point during the game some players are not aware of the entire position. Most card games are examples of games with imperfect information, the most popular of which are bridge and draw poker. Bridge and draw poker are similar to checkers and chess because the available information (the positions of the cards) may change during the course of the game. They are called games with *nonstatic information*.

In this thesis we will examine a card game called MYST, a variant of the game of Clue, which is a game with static and imperfect information. In MYST, a set of cards is divided uniformly among a set of players, and the remaining cards form a hidden pile. The goal of each player is to be the first to determine the contents of the hidden pile. On their turn, a player asks a question about the holdings of the other players, and, through a process of elimination, they can determine the contents of the hidden pile.

Clue is an interesting game to study for several reasons.

- Clue has not been formally studied.
- There is enough variation in Clue for it to be an interesting, replayable game.
- Most games focus on the competitive dissemination of resources, but Clue focuses on the competitive dissemination of knowledge.
- Since there is only one type of information about the position, i.e. what cards each player holds, then the dissemination of knowledge about positions is easy to specify.
- Once a card's position has been determined in Clue, it remains in play, which is usually not the case in most card games, like bridge.

This raises some interesting questions about similar games. How do we specify knowledge? Since each player is trying to glean as much knowledge as possible whilst giving away as little as possible, how could we quantify knowledge to determine who is winning? Could players use the actions of other players to gain knowledge for themselves? How should we base strategies around knowledge? Fortunately, Clue is simple enough to analyze, and contains enough variation to make the conclusions profound.

There are many aspects to Clue, which touch various areas of AI. We will give a formal specification of the game, which will involve first-order logic and the situation calculus, and we will develop strategies for Clue, which will involve elements of algorithms and game theory. But we will need a model for knowledge, one which handles knowledge in a world with static information.

In this chapter, we will informally describe the essential elements of the game of Clue and its variant MYST. Next we will describe the relevant aspects of game theory, what it means for a player to have knowledge, and finally the framework known as the situation calculus. Before we proceed, we should distinguish between the terms *information* and *knowledge*.

Information is a statement which is true, but may or may not be known by a player to be true; *knowledge* is the information that a player can prove.

1.1 Informal Specification of Clue and MYST

The classic game of Clue is a board game originally released by Parker Brothers and is now owned by Hasbro. What is the premise of the game? Six guests are invited to Boddy Mansion for a dinner party with Mr. Boddy. After dinner, Mr. Boddy is found dead on the main staircase, apparently the victim of foul play. The details of the crime are unknown and are to be determined by the guests (the players). The objective of Clue is to be the first player to correctly determine the murderer, the murder weapon and the room in which Mr. Boddy was killed.

There are six choices for the suspect: Miss Scarlet, Colonel Mustard, Professor Plum, Mrs. Peacock, Mr. Green, Mrs. White; six choices for the murder weapon: knife, revolver, rope, lead pipe, wrench, candlestick; and nine choices for the room: kitchen, study, ballroom, dining room, billiard room, lounge, library, conservatory, hall. Each of the listed choices is represented exactly once in a deck of 21 playing cards. Before the game begins, the cards are partitioned into their three sorts and one card from each sort is randomly selected and

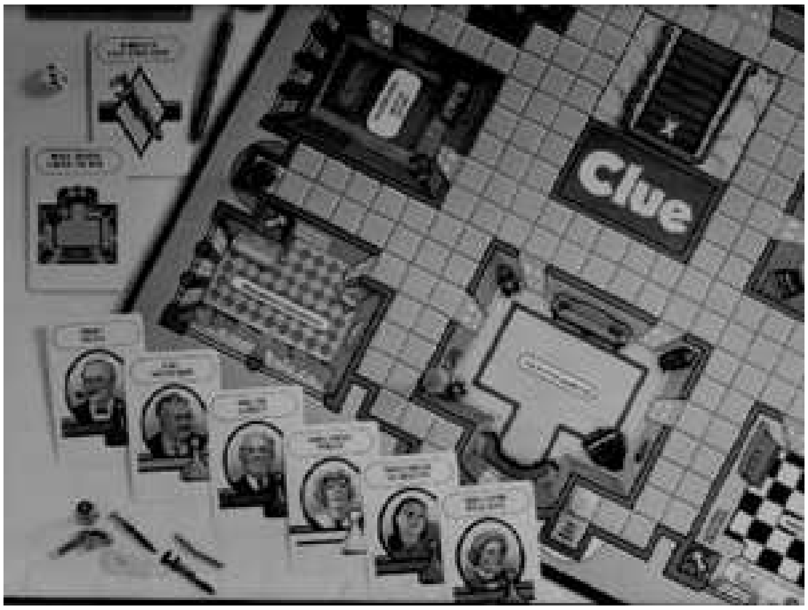

Figure 1.1: The classic game of Clue

hidden inside an envelope. These cards describe the details of the crime, and if a player correctly guesses the cards contained within the envelope, then they win the game. The remaining 18 cards are shuffled and dealt to each player in turn until they are exhausted. These cards are not shown to any other player, except by means of a *suggestion*.

The players move around Boddy Mansion on a game board and every time they enter a room, they are permitted to make a suggestion about the details of the crime. In other words, they simultaneously suggest a suspect, a murder weapon and a room. The first player in a clockwise progression who is able to contradict the suggestion, must privately show them a single card which contradicts it. Players take turns making suggestions until eventually someone has enough evidence to deduce the contents of the envelope and win the game. A player may guess the contents of the envelope at the end of their turn, and if they are wrong they may not make any other suggestions or guesses.

For example, suppose player 1 makes the suggestion that, “It was Mrs. White with the candlestick in the conservatory.” Player 2 is the first to the left of player 1, and she checks if they have any of the appropriate cards to contradict the suggestion. Player 2 does not have any of the Mrs. White, candlestick or conservatory cards and says, “Pass.” Player 3 checks for the same cards and so on until a player can contradict the suggestion or all opponents

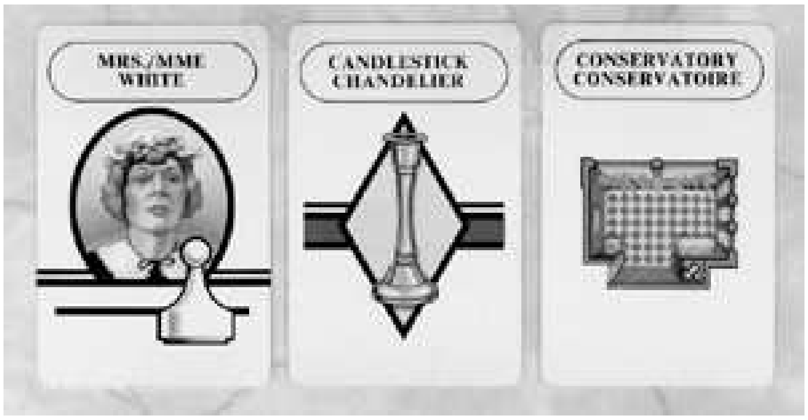

Figure 1.2: The suggestion "It was Mrs. White with the candlestick in the conservatory."

pass. Suppose that in this example, player 3 holds both the candlestick and the conservatory cards. She is not required to show both of them to player 1, so she picks the candlestick card and places it face down and slides it to player 1 for her inspection. After verifying the identity of the card, player 1 slides the card back to player 3. Notice that player 3 did not announce to all the players which card she showed player 1. It was a private show.

There are other rules in the classic game of Clue.

1. A player may only make a suggestion about the room they currently occupy.

2. A suggestion about a player automatically moves that player to the room about which the suggestion was made.

3. When moving between rooms, no player may walk through a square occupied by another player.

We have reformulated the original game of Clue into a simpler, more general game we call MYST[1], short for *mystery*. We can informally describe MYST as follows.

Consider the game MYST where there is a finite set of cards S from which $m \geq 1$ are selected randomly to be the unknown mystery pile and the remainder are dealt uniformly to each of $p \geq 2$ players in turn. To keep things simple and fair, we consider only the case in which each player holds an equal number of cards, say $n \geq 1$ of them. Therefore, $|S| = m + np$ and we define the constant $c = m + np$ for the sake of convenience.

[1]This is not to be confused with the puzzle adventure game MYST produced by the computer game company Cyan, Inc..

Each player, starting with player 1, gets a turn in rotation. On their turn, they *must* ask a question about $q \geq 1$ different cards of the form, "Do you have at least one of the following cards: $s_{a_1}, \ldots, s_{a_q}$?" This player is referred to as the *poser*. If the next player in rotation can answer *yes* to this question, then they will prove it to the poser by privately showing them the card and player 1's turn is over. However, if the answer is *no*, then the next player will either show a card or say *no* and so on. The player whose turn it is to answer the poser's question is called the *responder*. The turn order proceeds in a clockwise fashion for both posers and responders.

After asking their question, a player may guess the contents of the mystery pile. If they are right, then they will win the game; if wrong, they cannot ask another question or guess again. The game ends if a winner has been found or if all players have been eliminated from unsuccessful guesses.

Most of MYST can be easily expressed in terms of finite sets. Let $P = \{1, 2, \ldots, p\}$ be the set of p players, who take their turns in numerical order. Let $S_0 \subset S$ be the set of cards dealt to the mystery pile, and let $S_i \subset S$ be the set of cards dealt to player i. So, $|S_0| = m$ and $|S_i| = n$, when $i \in P$, and all of the subsets S_j are a partition of S. i.e.

$$\bigcup_{j=0}^{p} S_j = S \quad \text{and} \quad \forall 0 \leq i < j \leq p \ \ S_i \cap S_j = \emptyset$$

We can also define a question in terms of sets. Let Q be a *question* such that $Q \subset S$ and $|Q| = q$, and define $Q_i = Q \cap S_i$. Q_i represents the cards in the question held by player i, or, in the special case of Q_0, Q_0 represents the cards in the question which are in the mystery pile S_0. It follows that since $S_0, S_1, \ldots, S_p$ is a partition for S, then $Q_0, Q_1, \ldots, Q_p$ is a partition of Q. We adopt this notation of subscripted Qs to represent such partitions.

To be formal and complete, the constants m, n, p, q are positive integers and $p \geq 2$.

Clue is a special case of MYST, where $c = 21$, $m = q = 3$ and $2 \leq p \leq 6$. However, p does not divide 18 in all cases for p so some players might get extra cards. Furthermore, the construction of S_0 and any question is subject to the partitioning of S into the suspect, weapon and room subsets. Therefore there are more constraints in Clue than we permit in MYST, and so it is not conclusive that by solving MYST that we will solve Clue.

The generalization into MYST was done in order to make Clue easier to analyze. These are our motivations.

1. The suspect, weapon and room sets are clumsy to work with. Pooling the three

partitions into one makes the questions easier to specify. However, since the space of possible questions and possible mystery piles is larger for MYST than for Clue, MYST is no easier a game than Clue.

2. The suspect, weapon and room sets are not of equal size.
3. There is no guarantee that each player will start with the same number of cards, which might give one player a clear advantage.
4. Instead of moving from room to room to make suggestions, which involves unpredictable die rolling, we permit each player to make a suggestion every turn.
5. It is possible to force a draw via multi-player collusion.

In fact, all draw situations are the result of the "other rules" of Clue, stated earlier. Here are two examples of how collusion could force a draw.

1. Suppose that you had already proven it was Mr. Green with the candlestick and you had eliminated every room except the billiard room and the library. Since you already know the identity of the killer and the murder weapon, all that remains is the room and if you could make it to either room and make a suggestion, you would be assured a victory. However, two other (mean) players who are in rooms far away from the billiard room and the library each take turns making suggestions which include your player's name in the suggestion! These suggestions constantly teleport you to their rooms and can prevent you from ever reaching your destination. This is similar to the situation in chess where a draw is forced by perpetually placing the king in check.
2. Because no player may walk through a square occupied by another player, a large group of game tokens may blockade the advancement of another player, or perhaps block off all the entrances to a specified room.

In the next chapter, we will give a formal specification of MYST in terms of formal logic and the situation calculus.

1.2 Game Theory

No discussion of game strategy would be complete without referring to game theory. Game theory formalizes all games and seeks optimal strategies in competitive situations. In this

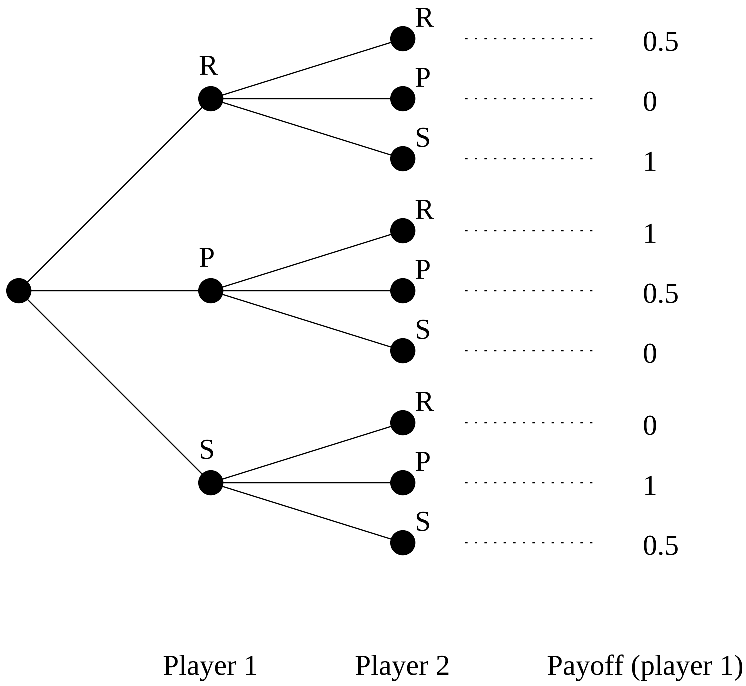

Figure 1.3: The game tree for rock-paper-scissors if player 1 goes first.

section, we describe the essential elements of game theory and highlight those relevant to this research. We introduce the fundamentals by characterizing the classic game of rock-paper-scissors.

The game of rock-paper-scissors is a two-player game. Each player must make one *decision*: choose one of *rock*, *paper* or *scissors*. If one player chooses rock and the other scissors, then the player who chose rock wins and earns a *payoff* of one point to their opponent's nil. Similarly, scissors win over paper, and paper wins over rock, each combination earning a payoff of 1-0. If both players make the same choice, it is a tie and each player earns $\frac{1}{2}$. If player 1 goes first and player 2 goes second, then all possible games can be described by the *game tree* shown in Figure 1.3, where each path from the root to a leaf represents a complete two-decision game.

Since rock-paper-scissors is a noncooperative game, each player will make the decision which maximizes their payoff. Player 1 goes first, so player 2 can earn a payoff of 1-0 every game by employing a *strategy*. A strategy is a function which makes a decision at each node in the game tree. Player 2's strategy is this. If player 1 picks scissors, player 2 picks rock; if player 1 picks paper, player 2 picks scissors; and if player 1 picks rock, player 2 picks

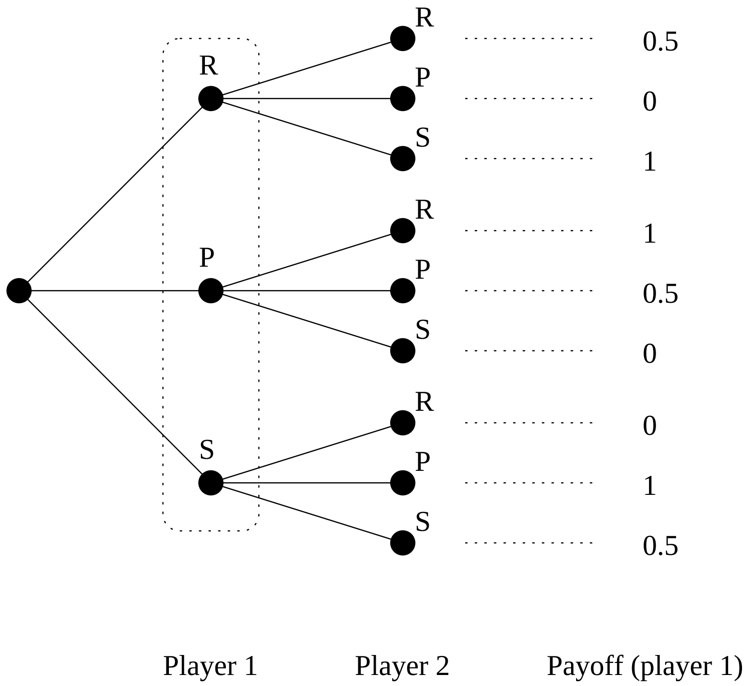

Figure 1.4: The game tree for rock-paper-scissors if both players decide simultaneously.

paper. This strategy could be determined by using backwards induction or by finding the path(s) associated with evaluating $\min_{\text{Player 1}} \max_{\text{Player 2}} payoff(\text{Player 2})$. Player 1 can guarantee no better than 0-1 either. Since rock-paper-scissors has *perfect information* player 2 can always view player 1's action and respond to it.

The more standard version of rock-paper-scissors is one in which player 2 is not privy to such information. In actual play, the two decisions happen simultaneously, and so we need *information sets*. Figure 1.4 shows the modified game tree. The two figures are identical except for the dotted box representing the information set. Player 2 cannot distinguish between the three nodes in the box, that represent player 1's decision. The second game is a game of *imperfect information* and the optimal strategy cannot easily be found by using backwards induction.

The goal of game theory is to find the best strategies for each player. Here we introduce *equilibrium*, originally defined by von Neumann and later refined by Nash and others [5]. Loosely speaking, an equilibrium point is a local minimum in the payoff function. For two players, a pair of strategies is in equilibrium if one player's opponent cannot do better by switching to another strategy. In rock-paper-scissors, if player 1 decides that it is best to choose *rock* every time, then player 2 could earn a payoff of 1-0 by choosing *paper* every time. Therefore, consistently choosing rock is not an equilibrium strategy.

Sometimes it is best to probabilistically choose a strategy, which is called a *mixed strategy*. In rock-paper-scissors, a mixed strategy offers the only equilibrium. If player 1 should uniformly pick one of rock, paper or scissors with equal probabilities, then their expected payoff is $\frac{1}{2}$ no matter what strategy player 2 chooses, and furthermore, $\frac{1}{2}$ is the best possible expected payoff. This concludes the example.

There are many works on equilibria in game theory some of which focus on the two-person aspects [13] [14], and some of which focus on the n-person aspects [5] [15]. The most important result is the theorem of von Neumann that claims the existence of a mixed equilibrium in a finite game with imperfect information.

Theorem 1.1 (von Neumann) *Every finite game of imperfect information has a mixed strategy equilibrium.*

We have already introduced the concept of perfect information and static information games. Although it is not central to the core of this thesis, we should contrast imperfect information against incomplete information.

A game has complete information if all players can construct the entire game tree. Informally stated, a player has complete information if they know everything that could possibly happen during the course of the game. Incomplete information means that at least one player either does not know all of the rules or does not know if they know all the rules. Note that most commonly played games have complete information. Human players who do not have complete information usually feel the game is unfair when an unknown rule is sprung upon them. Therefore games of incomplete information are rare, unpopular and asymmetric.

The card game Mao is the best example of a card game with incomplete information. In Mao, players are assessed penalties whenever they break a rule, and players try to learn the rules from the penalties assessed. Players may also develop new rules as play progresses. Some might also suggest that baseball is a game of incomplete information because it is impossible to assimilate all of the rules. However, the rules of baseball are well-defined and, although complicated and detailed, failure to assimilate the rules is considered a deficiency of the player, not the game. To maintain symmetry, it is assumed that: *All players are* *perfect reasoners.* Fagin et al. [4] refers to this as *logical omniscience* and discusses it at length. In this thesis, we will call this Axiom 0 and discuss it in Section 1.3.

Another key assumption is that all players are aware of and remember what actions have taken place in the course of the game. This is called *perfect recall* and is also incorporated in Axiom 0. Once again, human players are likely to forget the game history in a long game, but symmetry again motivates us to make this assumption.

We conclude this section with a table of some popular complete information games, classified by their types of information. Static, perfect information games are not interesting because they completely lack any variation.

Clue	static	imperfect
Stud Poker	static	imperfect
Bridge	nonstatic	imperfect
Draw Poker	nonstatic	imperfect
Gin Rummy	nonstatic	imperfect
Checkers	nonstatic	perfect
Chess	nonstatic	perfect
Monopoly	nonstatic	perfect

1.3 Reasoning About Knowledge

For this thesis, all reasoning about knowledge utilizes the possible worlds model summarized from Fagin et al. [4]. If player p considers the truth of some piece of information ϕ, p is said to *know* ϕ if and only if ϕ is true in all worlds that p considers possible, and we write $K(p, \phi)$[2]. The possible worlds model is best illustrated by an example.

Consider a standard deck of cards uniformly distributed among 4 players, each player receiving 13 cards. The goal of this example is to determine what player 1 knows after the deal of the cards. Before looking at his cards, player 1 considers all of the worlds which are possible from his perspective. Player 1 can conceive of a world in which the ace of spades is in player 2's hand, and also of a world in which it is in his own hand. The location of the ace of spades, or any card for that matter, could be in any player's hand (with equal

[2] We use a function notation instead of the familiar expression $K_p\phi$ because we intend to augment K with a situation argument.

probability) and therefore its location is not known by player 1. In fact, before looking at his own cards, player 1 only knows that a partitioning of the cards has occurred. In other words, he knows that the ace of spades is held once and only once and that each player holds exactly 13 cards.

After player 1 has examined his cards, the number of possible worlds is reduced. If he, for instance, holds the ace of spades, then he will eliminate all worlds except where he holds the ace of spades and the other players do not hold the ace of spades. In similar fashion, he knows that he holds every other card in his hand, but again he does not know the location of the cards not in his hand. Therefore, the worlds that player 1 considers possible are represented exactly by the possible distributions of the cards he does not hold. Furthermore, since the initial distribution of the cards was uniform, he cannot decree that one possible world is more likely than another. The example is thus concluded.

We will now introduce *common knowledge*. Let G be a subset of players and ϕ be a piece of information. $C(G, \phi)$ means that "it is common knowledge among the players in G that ϕ." Informally, $C(G, \phi)$ means that everyone in G knows ϕ, and everyone in G knows that everyone in G knows ϕ, and everyone knows that everyone knows that everyone knows ϕ, and so on. Algebraically, if we let $E(G, \phi) \equiv \bigwedge_{p \in G} K(p, \phi)$ represent "everyone in G knows ϕ", then we define $E^k(G, \phi)$ by the recurrence

$$E^1(G, \phi) \equiv E(G, \phi) \qquad E^{k+1}(G, \phi) \equiv E(G, E^k(G, \phi))$$

and then write

$$C(G, \phi) \equiv \bigwedge_{k=1}^{\infty} E^k(G, \phi).$$

Although first-order logic prevents the use of an infinite conjunction, there are ways to axiomatize $C(G, \phi)$ using a finite list of axioms, which capture the meaning understood in the definition above. Such a characterization can be found in [4].

We now present the S5 Axioms which describe the most fundamental manipulations of $K(p, \phi)$.

- *Distribution Axiom*

$$(K(p, \phi) \wedge K(p, (\phi \supset \psi))) \supset K(p, \psi)$$

- *Knowledge Generalization Rule*

$$\text{If, in all possible worlds } \phi \text{ holds true, then } K(p, \phi).$$

- *Truth Axiom*

$$K(p, \phi) \supset \phi$$

- *Positive Introspection Axiom*

$$K(p, \phi) \supset K(p, K(p, \phi))$$

- *Negative Introspection Axiom*

$$\neg K(p, \phi) \supset K(p, \neg K(p, \phi))$$

We might also hypothesize how to describe information that players do not know. "Player p does not know ϕ" is written as $\neg K(p, \phi)$. This can allow for some very complicated expressions with nested Ks. For example, $\neg K(1, K(2, \neg K(1, \phi)))$ means "Player 1 does not know that player 2 knows that player 1 does not know ϕ." Meta-knowledge about what players do not know has limited usefulness in the study of Clue and will not be considered any further.

There are some other assumptions about our belief system that must be considered. The first is *logical omniscience*, which is the assumption that all players are perfect reasoners. Every player knows all tautologies and is able to derive all consequent tautologies, and every player knows all rules and axioms, and can always understand the full consequences of any action. Logical omniscience can be summarized by the following closure property of knowledge: an agent is *logically omniscient* if, whenever he knows all of the formulae in a set Ψ, and Ψ logically implies the formula ϕ in all possible worlds, then the agent also knows ϕ. A more complete characterization can be found in [4].

The second is *perfect recall*, which is the assumption that all players can remember the history of the game. Humans usually do not have perfect recall because they usually cannot accurately remember what happened more than (say) 20 turns ago. The notions of recall and history brings another knowledge issue to light. Games like Clue take several turns and so a player's knowledge about another player's holdings will be increased incrementally over time. Since Clue is a game with static information, the information does not change over time. Therefore, if a player knows the position of a card, that knowledge will never go away. We call this *immutable knowledge*. It follows that knowledge of immutable knowledge is also immutable, but assertions like "Player 1 does not know that player 2 holds card s_x," i.e. $\neg K(1, \phi)$, are not immutable because player 1 may (on a future turn) discover the true location of card s_x.

We will refer to logical omniscience and perfect recall collectively as *Axiom 0*. Although Axiom 0 does not include the reality of imperfect human opponents, it makes the analysis of strategies mathematically easier.

Maintaining knowledge over the course of many turns requires a mathematical quantification of history. Here, we employ the situation calculus whose focus is the situation structure: a complete list of previous actions. In this work, we will augment the knowledge predicates K, E and C with a third, situation argument s.

1.4 Situation Calculus

In this section we describe the essential elements of the situation calculus, described in the survey paper of Levesque et al. [10] and originally developed by Hayes and McCarthy [11]. We employ the situation calculus because it neatly describes the notions of history and time in a framework that is suitable to formalize Clue.

The situation calculus is a second order language with equality. It has three disjoint sorts: *action* for actions, *situation* for situations and *object* for everything else. First, we will define the notion of a situation. Loosely speaking, a situation is a list of actions in the order in which they occurred. In fact, a situation s could be viewed as a LISP list structure, with functions $car(s)$, $cdr(s)$, $cons$ and so forth. $car(s)$ would represent the most recent action and the $cdr(s)$ would represent the history previous to this action. Levesque et al. define a function $do(a, s)$ which is similar to $cons$ in LISP. $do(a, s)$ returns a situation with a as its most recent action and s as the rest of its history. The special symbol $Init_0$[3] is used as the initial situation, which is analogous to a starting state or a NIL list.

Formally, we define the symbols $do : action \times situation \rightarrow situation$, $Init_0$, $\sqsubset$ and $\sqsubseteq$

[3] Levesque et al. uses S_0, but that symbol is used extensively in this work for some other purpose.

according to the following four axioms.

$$do(a_1, s_1) = do(a_2, s_2) \supset (a_1 = a_2) \wedge (s_1 = s_2) \tag{1.1}$$

$$\neg(s \sqsubset Init_0) \tag{1.2}$$

$$s \sqsubset do(a, s') \equiv s \sqsubseteq s' \tag{1.3}$$

$$(\forall P).P(Init_0) \wedge (\forall a, s)[P(s) \supset P(do(a, s))] \supset (\forall s)P(s). \tag{1.4}$$

Axiom (1.1) is a unique names axiom for situations. Two situations are the same if and only if they are the same sequence of actions. Axioms (1.2) and (1.3) axiomatize the $\sqsubset$ relation. Informally, $s \sqsubset s'$ if s is a proper subhistory of s'. Note also that $s \sqsubseteq s'$ is a shorthand for $s \sqsubset s' \vee s = s'$. Axiom (1.4) is a second order induction for situations.

The situation structure is important for Clue because it formalizes gaining knowledge over time. If a player discovers the location of an opponent's card, he has immutable knowledge in that and all subsequent situations. However that same player did not know its location in any proper subhistory. Therefore, the predicate $K(i, \phi)$ depends on the situation, and, where relevant, s will be augmented as a third argument.

Situations also encapsulate the history of actions and the order in which they occurred. In most games, the opponent's history can be used to deduce their game plan.

The language of the situation calculus may also include:

- Countably infinitely many *variable symbols* of each sort.
- The binary predicate symbol $Poss : action \times situation$. The intended interpretation of $Poss(a, s)$ is that it is possible to perform the action a in situation s. $Poss(a, s)$ will be used to define what actions each player is allowed on their turn.
- Countably infinitely many *predicate symbols* to denote situation independent relations.
- Countably infinitely many situation independent *functions* whose output are of sort object.

- Countably infinitely many situation independent functions whose output are of sort action. These are called *action functions*, and they are used to denote actions with input variables such as

$$play(\text{ace of spades}), asks(i, Q), \text{etc.}$$

- Countably infinitely many predicate symbols, which are situation dependent. These are called *relational fluents*, and can be used to keep track of things which change during the course of the game. For example, $K(i, \phi, s)$ is a relational fluent.

- Countably infinitely many functions, which are situation dependent. These are called *functional fluents*, but they will not be employed in this work.

We will define the axioms for the special predicate $Poss(a, s)$, which describe all possible actions, and the Successor State Axioms which describe the relational fluents. We are also allowed some uniform formulae. Informally, a formula in the situation calculus is uniform if and only if it does not depend on two situations with different histories. A more formal description of uniform formulae can be found in [10]. For example, the S5 Axioms, the tautologies of first order logic and the axioms describing the situation $Init_0$ are all uniform.

We will use the situation calculus to formalize MYST in Chapter 2.

1.5 Thesis Overview

In this thesis, we will first give a formal description of MYST in terms of the situation calculus and all game axioms, described in Chapter 2. Chapter 3 will describe some general results and properties of MYST and Chapter 4 will describe some strategies and strategic principles of MYST. The body of the research concluded, Chapter 5 will summarize the work and suggest avenues of future work.

Chapter 2

Formal Specification of MYST

In this chapter, we will formalize the game of MYST in terms of logical constructs and the situation calculus developed by Hayes and McCarthy [11] and described in the survey paper of Levesque et al. [10]. This will give a solid foundation for the informal description given in Section 1.1.

The organization of this chapter is modular, but similar to the way Levesque breaks up his portions of the situation calculus. We first describe the constants which are fundamental to MYST, which include the integers m, n, p, q, the cards and the players. Next we describe all possible actions and fluents, whose values change over the course of the game. We then describe the initial game conditions for the case of the initial situation $Init_0$ and then describe which actions are possible in which situations. Finally, we will describe the Successor State Axioms which will axiomatize how the values of the fluents are affected by actions.

2.1 Constants

In the game of MYST, we will need 4 integers which are represented by the symbols m, n, p, q. All of these must be positive integers and $p \geq 2$. We must also have the concept of a card and the concept of a player. Informally, there will be $c = m + np$ cards, denoted by the symbols $s_1, \ldots, s_c$, and p players, denoted by the integers $1, \ldots, p$. We will define the predicates $card(x)$ and $player(i)$ to encapsulate these symbols into their two sorts. The cards will not have any order to them, but the players will have a cyclical turn order, determined by the function $incp(i)$, which will return the player who is next to act after player i. Without loss of generality, $incp(i)$ will add 1 to (i *modulo* p).

In summary, we define symbols m, n, p, q of type integer, $card(x)$ and $player(i)$ of type predicate, and $incp(i)$ of type function.

$$m \geq 1, \;\; n \geq 1, \;\; p \geq 2, \;\; q \geq 1. \tag{2.1}$$

For simplicity sake, we define the constant c:

$$c \;=\; m + np, \tag{2.2}$$

and then the predicates $card(x)$ and $player(i)$:

$$\begin{aligned} card(x) \;&\equiv\; (x = s_1) \vee \cdots \vee (x = s_c) \\ &\equiv\; \bigvee_{i=1}^{c} (x = s_i) \end{aligned} \tag{2.3}$$

$$\begin{aligned} &(s_1 \neq s_2) \wedge (s_1 \neq s_3) \wedge \cdots \wedge (s_{c-1} \neq s_c) \\ \equiv\; &\forall 1 \leq i < j \leq c \;\; (s_i \neq s_j) \end{aligned} \tag{2.4}$$

$$\begin{aligned} player(i) \;&\equiv\; (i = 1) \vee \cdots \vee (i = p) \\ &\equiv\; \bigvee_{j=1}^{p} (i = j), \end{aligned} \tag{2.5}$$

and finally, $incp(i)$:

$$\begin{aligned} incp(1) &= 2 \\ incp(2) &= 3 \\ &\vdots \\ incp(p) &= 1 \end{aligned} \tag{2.6}$$

or, more succinctly:

$$(player(i) \wedge player(j) \wedge j = i \bmod p + 1) \;\supset\; incp(i) = j. \tag{2.7}$$

We will need one more constant: the special initial situation constant, denoted by $Init_0$[1]. $Init_0$ will represent the initial game situation for MYST, before any players have had an opportunity to act.

[1] Recall that $Init_0$ is usually denoted by S_0 in Levesque et al., but in this thesis S_0 represents the mystery pile.

2.2 Turns and Phases

In MYST, every player takes their turn in sequence, starting with player 1. In principle, the turn sequence is determined by the function $incp(i)$, and in practice it is the integers $1, \ldots, p$ in numerical sequence. Each player's turn is divided into a sequence of phases. The phase sequence is illustrated by the flowchart shown in Figure 2.1.

To expound on the flowchart, once a player i's turn begins they *must* ask a question (*askphase*), unless they have been eliminated due to a wrong guess (*endphase*). Players in order of sequence by $incp(i)$ either answer *yes* or *no* until a *yes* has been answered or all players have been exhausted by *no*'s (*ansphase*). The player to answer *yes* shows player i a card in the question (*showphase*). After that, player i may guess or not guess (*guessphase*). If they guess correctly, they win and the game is over; if they guess incorrectly, they are eliminated and they enter the *endphase*. A wrong guess will permit player i to enter only the *endphase* for the remainder of the game. If player i does not guess, they will enter the *endphase*. Player i's turn ends and player $incp(i)$'s turn begins.

This flowchart and the previous informal description is a pictorial representation of the situation calculus predicate $Poss(a, s)$. Whether or not an action is possible will be determined mostly by the current phase, but other fluents which keep track of the current player and which players have been eliminated will also be involved. The formal axiomatization of $Poss(a, s)$ will be given in Section 2.7.

The finite state diagram in Figure 2.2 shows the relationship between the phases. Only one phase may be active at a time, and only the described actions on the transitions can change the phase.

From these two diagrams, we should observe a strong relationship between actions and phases. A game can only progress from one phase to another in a specific order and only after certain actions have been performed. On the other hand, each action can only be performed in a certain phase. The lock-step between phases and actions will be a firm base for defining the axioms for $Poss(a, s)$. But first we will need to define the necessary actions and fluents, which have only been mentioned in the figures so far.

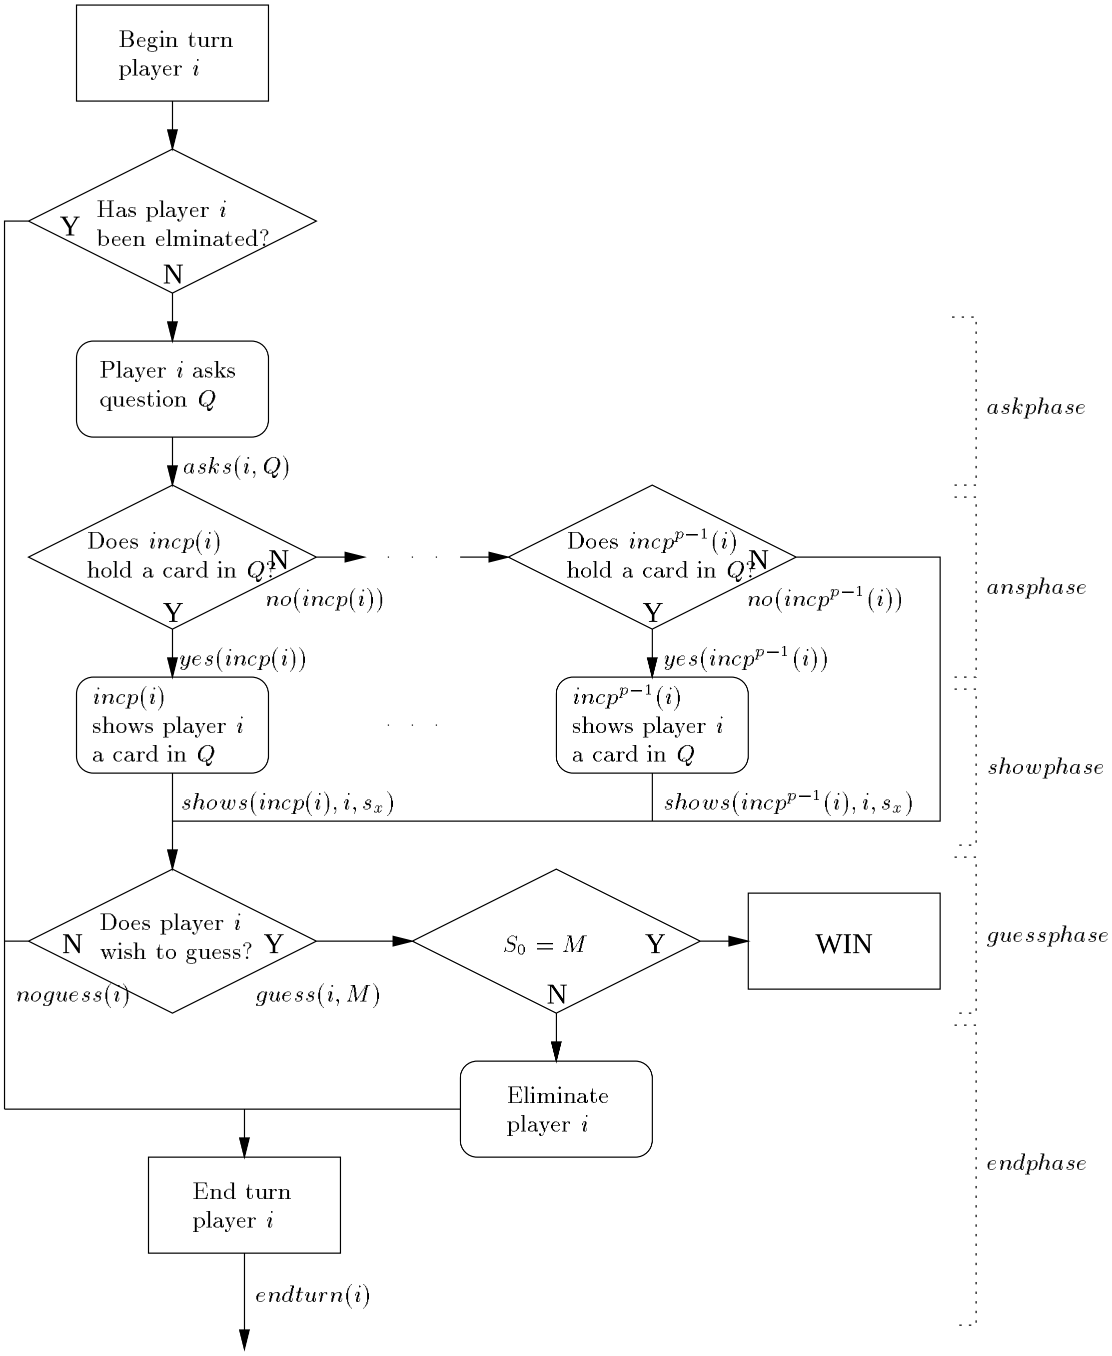

Figure 2.1: Flowchart for a single turn of MYST

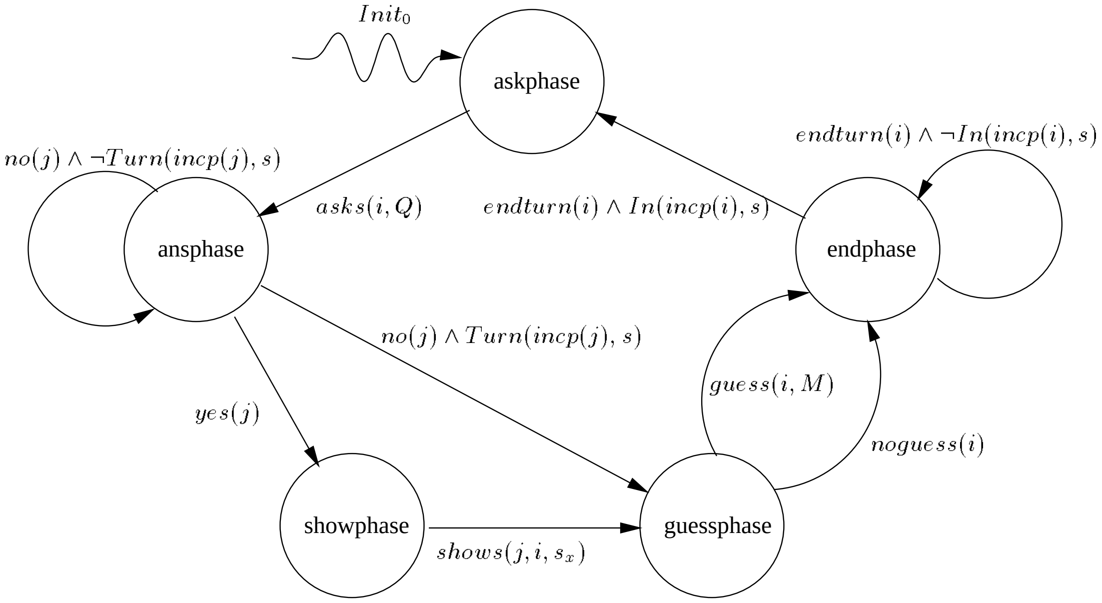

Figure 2.2: Finite state diagram illustrating the progression of phases

2.3 Notation Abbreviations

One may have already remarked about the definitions of the predicates $card(x)$ and $player(i)$, in that they are equivalent to the member predicates of the well-defined finite sets S and P which were introduced in Section 1.1. Since many of the axioms can be neatly specified in terms of S, P and the standard set theoretic operators, we will embed our set theoretic notions in first-order logic and then adopt this more succinct notation. For instance, we could assign a different name to each of the possible subsets of S (there are 2^c of them, but a finite number nonetheless) and then construct the standard operators $\cap$ and $\cup$, the relation $\subset$, and the function $sizeof(A)$, which would (albeit exhaustingly) tie these subsets together. However, we feel that this would obfuscate the meaning of the axioms we are presenting, and henceforth we will let

$$\begin{aligned} S &\equiv \{s_1, \ldots, s_c\} \quad \text{and} \\ P &\equiv \{1, \ldots, p\}. \end{aligned} \tag{2.8}$$

This exhaustive technique is but only a single implementation of a finite set in terms of first-order logic. Other implementations are possible, but need not be discussed here.

In order to simplify our equivalences even further, we will adopt a notation for variables.

We will tend to use the variables:

- i, j as indices for players, where typically i is the poser and j is the responder,
- s_x to represent a single card,
- G to represent a subset of players, i.e. $G \subseteq P$,
- Q to represent a set of cards in a question,
- M to represent a set of cards in a guess of the mystery pile,
- S_j to represent the set of cards held by player j,
- S_0 to represent the set of cards contained in the mystery pile,
- a to represent an action, and
- s to represent a situation.

Henceforth, we will no longer explicitly list "type-checking" predicates to ensure that s_x is a card (for instance). It will be assumed that Q is verified to be a subset of S and of size q, and similarly that $M \subset S$ is of size m, even though not explicitly stated.

2.4 Actions

The following is the list of action functions and their informal descriptions. The sequence in which actions may occur will be defined by the predicate $Poss(a, s)$ in Section 2.7.

Note that the player arguments i and j could be omitted in all action functions, because of the use of the fluents which maintain the current player's turn, however, we have retained this argument in the interest of readability.

Note also that the first argument always represents the player performing the action.

$asks(i, Q)$: Player i asks question Q.

$no(j)$: Player j says *no* to question Q, asked by player i.

$yes(j)$: Player j says *yes* to question Q, asked by player i.

$shows(j, i, s_x)$: Player j shows card $s_x \in Q \cap S_j$ to player i who asked Q.

$guess(i, M)$: Player i guesses that $S_0 = M$.

$noguess(i)$: Player i makes no guess.

$endturn(i)$: Player i ends their turn.

2.5 Fluents

The following is a list of fluents and their informal descriptions. The evaluation of the fluents will depend on the situation s. The value of the fluents in $Init_0$ will be defined in Section 2.6. Their values will change as situations develop according to the Successor State Axioms given in Section 2.9.

Fluents Describing the Location of Cards

$h(i, s_x, s)$: Player i holds card s_x.

$h(0, s_x, s)$: The mystery pile holds card s_x.

Fluents Describing Knowledge

$K(i, \phi, s)$: Player i knows ϕ.

$C(G, \phi, s)$: ϕ is common knowledge for all players in $G \subseteq P$.

Fluents Describing the State of the Game

$In(i, s)$: Player i has not yet been defeated due to a wrong guess.

$Question(Q, s)$: Question Q was the most recently asked question.

$Gameover(s)$: The game is over.

Fluents Describing the Turn Order and Phases

$Turn(i, s)$: It is player i's turn.

$AnsTurn(j, s)$: It is player j's turn to answer the question.

Fluents Describing the Phases

$askphase(s)$: It is the ask phase.

$ansphase(s)$: It is the answer phase.

$showphase(s)$: It is the show phase.

$guessphase(s)$: It is the guess phase.

$endphase(s)$: It is the end phase.

2.6 The Initial Situation $Init_0$

Upon initially dealing the cards, each player will hold exactly n cards and the mystery pile will hold exactly m cards. The cards of S are partitioned among the players and the mystery pile. The following three axioms describe all possible partitions.

Disjoint Axiom If player i holds card s_x, then a second player j (or the mystery pile) does not hold s_x. Likewise, if the mystery pile holds s_x, s_x is not held by any player.

$$0 \leq k \leq p \wedge h(k, s_x, Init_0) \supset \left[\bigwedge_{0 \leq k \neq k' \leq p} \neg h(k', s_x, Init_0) \right]. \tag{2.9}$$

Existence Axiom Every card is held by at least one player (or the mystery pile). This is the converse of the Disjoint Axiom.

$$\bigvee_{j=0}^{p} h(j, s_x, Init_0). \tag{2.10}$$

Set Size Axioms Player i holds exactly n cards and the mystery pile holds exactly m cards.

$$\forall \Sigma \; [\Sigma \subset S \wedge |\Sigma| = n] \supset \left[\bigwedge_{s_x \in \Sigma} h(i, s_x, Init_0) \equiv \bigwedge_{s_x \in S - \Sigma} \neg h(i, s_x, Init_0) \right] \tag{2.11}$$

$$\forall \Sigma \; [\Sigma \subset S \wedge |\Sigma| = m] \supset \left[\bigwedge_{s_x \in \Sigma} h(0, s_x, Init_0) \equiv \bigwedge_{s_x \in S - \Sigma} \neg h(0, s_x, Init_0) \right] \tag{2.12}$$

Collectively, we will refer to axioms (2.9), (2.10), (2.11) and (2.12) as the *Partition Axioms.*

Alternatively, we could have specified the Partition Axioms in terms of set theory. Their set equivalents would be

$$\begin{aligned} \forall 0 \leq i < j \leq p \;\; S_i \cap S_j = \emptyset & \qquad (\textit{Disjoint Axiom}) \\ \bigcup_{i=0}^{p} S_i = S & \qquad (\textit{Existence Axiom}) \\ |S_0| = m, \quad \forall i \;\, player(i) \supset |S_i| = n & \qquad (\textit{Set Size Axioms}) \end{aligned} \tag{2.13}$$

Once the cards are held by the players in $Init_0$, they know they hold their cards.

Knowledge Initialization Axiom If player i holds card s_x, then it is common knowledge for player i that he holds card s_x. i.e.

$$C(\{i\}, h(i, s_x, Init_0), Init_0) \tag{2.14}$$

Initialization of Other Fluents Here we define the value of the rest of the fluents in the initial situation $Init_0$. Informally, they assert that when the game begins

- No one has been eliminated;

$$\forall i \;\; In(i, Init_0)$$

- No one has asked a question;

$$\forall Q \;\; \neg Question(Q, Init_0)$$

- Player 1 is in the *askphase* of their turn; and

$$askphase(Init_0)$$
$$Turn(1, Init_0)$$

- The game is not over.

$$\neg Gameover(Init_0)$$

2.7 *Poss*ible Actions

The transitions of the flowchart shown in Figure 2.1 provide a pictorial representation of the situation calculus predicate $Poss(a, s)$. We are now ready for its formal axiomatization.

It should be well noted that no action is possible if the game is over. Therefore the fluent $\neg Gameover(s)$ is implicitly conjoined with all of the following axioms, but omitted for the sake of readability.

- $asks(i, Q)$

  Player i *must* ask a question *iff*

  1. it is their turn and
  2. it is the *askphase*.

  Note that player 1 starts in the *askphase* and therefore $Poss(asks(1, Q), Init_0)$ is true, implied by the consequences of the fluents defined for $Init_0$. Player i may choose any *valid* question Q which satisfies the necessarily implied type-checking: $Q \subset S$ and $|Q| = q$.

  $$Poss(asks(i, Q), s) \equiv Turn(i, s) \wedge askphase(s) \tag{2.15}$$

- $no(j)$

  Player j must answer *no iff*

  1. $Q \cap S_j = \emptyset$, and
  2. It is player j's turn to answer and
  3. It is the *ansphase*.

  $$\begin{aligned} Poss(no(j), s) \equiv\ & ansphase(s) \wedge AnsTurn(j, s) \wedge \\ & Question(Q, s) \wedge (Q \cap S_j = \emptyset) \end{aligned} \tag{2.16}$$

- $yes(j)$

  Player j must answer *yes iff*

  1. $Q \cap S_j \neq \emptyset$ and

2. It is player j's turn to answer <u>and</u>
3. It is the *ansphase*.

$$\begin{aligned} Poss(yes(j), s) \equiv\ & ansphase(s) \wedge AnsTurn(j, s) \wedge \\ & Question(Q, s) \wedge (Q \cap S_j \neq \emptyset) \end{aligned} \tag{2.17}$$

- $shows(j, i, s_x)$

  Player j must show player i a card *iff* they just answered *yes*.

$$\begin{aligned} Poss(shows(j, i, s_x), s) \equiv\ & Turn(i, s) \wedge AnsTurn(j, s) \wedge (s_x \in Q \cap S_j) \wedge \\ & showphase(s) \wedge Question(Q, s) \end{aligned} \tag{2.18}$$

- $guess(i, M)$ and $noguess(i)$

  Player i may *guess* (or *noguess*) only in the *guessphase* of their turn.

$$Poss(guess(i, M), s) \equiv Turn(i, s) \wedge guessphase(s) \tag{2.19}$$

$$Poss(noguess(i), s) \equiv Turn(i, s) \wedge guessphase(s) \tag{2.20}$$

- $endturn(i)$

  Player i must end their turn *iff* they are in the *endphase*.

$$Poss(endturn(i), s) \equiv Turn(i, s) \wedge endphase(s) \tag{2.21}$$

2.8 Uniform Formulae

The definitions of the constants, predicates and the fluents in the situation $Init_0$ are all uniform in MYST. Furthermore, all tautologies in first-order logic are uniform as well as all of the knowledge axioms in our belief system such as Axiom 0 and the S5 Axioms stated in Section 1.3. The only other uniform formulae unique to MYST are those pertaining to the Turn, Question and Phase Fluents.

The implications in (2.22) state that only one player at a time may have a turn; only one player at a time may answer a question; and only one question at a time may be asked.

$$\begin{aligned} \forall i, j\ (Turn(i, s) \wedge i \neq j) \ &\supset \neg Turn(j, s) \\ \forall i, j\ (AnsTurn(i, s) \wedge i \neq j) \ &\supset \neg AnsTurn(j, s) \\ \forall Q, Q'\ (Question(Q, s) \wedge Q \neq Q') \ &\supset \neg Question(Q', s) \end{aligned} \tag{2.22}$$

The five implications below exhaustively state that only one phase may be active in any situation.

$$\begin{aligned}
askphase(s) &\supset \neg ansphase(s) \wedge \neg showphase(s) \wedge \neg guessphase(s) \wedge \neg endphase(s) \\
ansphase(s) &\supset \neg askphase(s) \wedge \neg showphase(s) \wedge \neg guessphase(s) \wedge \neg endphase(s) \\
showphase(s) &\supset \neg askphase(s) \wedge \neg ansphase(s) \wedge \neg guessphase(s) \wedge \neg endphase(s) \\
guessphase(s) &\supset \neg askphase(s) \wedge \neg ansphase(s) \wedge \neg showphase(s) \wedge \neg endphase(s) \\
endphase(s) &\supset \neg askphase(s) \wedge \neg ansphase(s) \wedge \neg showphase(s) \wedge \neg guessphase(s)
\end{aligned}$$

2.9 Successor State Axioms

In this section, we list the Successor State Axioms for all of the fluents, that is, the axioms which associate actions with the values of the relational fluents.

Card Fluent Axiom

Once the cards are initially held by the players in situation $Init_0$, they are fixed and their locations do not change for the entire game.

$$\begin{aligned}
h(i, s_x, do(a, s)) &\equiv h(i, s_x, s) \\
h(0, s_x, do(a, s)) &\equiv h(0, s_x, s)
\end{aligned} \tag{2.23}$$

Therefore h is independent of the situation argument, and so we will abbreviate $h(i, s_x, s)$ by $h_{i,x}$ from this point forward.

Knowledge Fluent Axioms

Immutable Knowledge The fluent $K(i, \phi, s)$ is supposed to be analogous to the concept of modal knowledge, but in a situation dependent setting. In the game of MYST, we need to reason about how the players perceive their holdings. Recall that the objective of the game is to come up with the correct guess of the contents of the mystery pile S_0. As we will later prove, there are only two ways to do such a thing. Both of them rely on determining the whereabouts of the outstanding cards, the placement of which is determined by the Partition Axioms and is fixed throughout the game by the Card Fluent Axiom (2.23). Therefore, the primary knowledge and reasoning will be about the fluent $h_{i,x}$.

Since $h_{i,x}$ is independent of the situation, then so is $\neg h_{i,x}$ by the Partition Axioms. Any knowledge about h or $\neg h$ is an immutable, objective truth. Therefore, immutable knowledge in any history of s is known in s as well.

$$\begin{aligned} \forall i, j, s_x, s, s' \; (s' \sqsubseteq s \wedge K(i, h_{j,x}, s')) \; \supset \; K(i, h_{j,x}, s) \\ \forall i, j, s_x, s, s' \; (s' \sqsubseteq s \wedge K(i, \neg h_{j,x}, s')) \; \supset \; K(i, \neg h_{j,x}, s) \end{aligned} \tag{2.24}$$

Inductively, it can be seen that knowing any of the knowledge of the form (2.24) is also immutable, and so on. Therefore Axiom (2.24) holds for the common knowledge fluent C as well as K.

Reasoning about what the players do not know (i.e. knowledge which is not immutable) will not be considered in this work, as its usefulness is very limited.

Shows Axiom To describe the knowledge consequence of $shows(j, i, s_x)$, we need the following axiom.

$$C(\{i, j\}, h_{j,x}, do(shows(j, i, s_x), s)) \tag{2.25}$$

Game State Axioms

We define the Successor State Axioms for the fluents $In(i, s)$, $Question(Q, s)$ and $Gameover(s)$ as follows.

- Player i is still in *iff* they did not guess wrong.

$$\begin{aligned} In(i, do(a, s)) \; \equiv \; & In(i, s) \wedge \\ & \neg(a = guess(i, M) \wedge M \neq S_0) \end{aligned} \tag{2.26}$$

- The last question asked does not change until a new question is asked.

$$\begin{aligned} \forall Q, Q' \; Question(Q, do(a, s)) \; \equiv \; & a = asks(i, Q) \; \vee \\ & (Question(Q, s) \wedge \neg(a = asks(i, Q'))) \end{aligned} \tag{2.27}$$

- The game is over when any player has had a successful guess or all players have been eliminated.

$$\begin{aligned} Gameover(do(a, s)) \; \equiv \; & (a = guess(i, M) \wedge M = S_0) \; \vee \\ & \bigwedge_{i=1}^{p} \neg In(i, do(a, s)) \end{aligned} \tag{2.28}$$

Turn Order Axioms

A player's turn does not change until the $endturn(i)$ action has occurred. The next player is determined by $incp(i)$.

$$\begin{aligned} Turn(incp(i), do(a,s)) \equiv\ & (Turn(i,s) \wedge a = endturn(i)) \vee \\ & (Turn(incp(i), s) \wedge \neg(a = endturn(incp(i)))) \end{aligned} \tag{2.29}$$

Similarly, for $AnsTurn(j,s)$, players answer in their turn order. It is player j's turn to answer if the previous player asked the question or the previous player said no. If player j said yes, then $AnsTurn(j,s)$ is true in the resulting situation.

$$\begin{aligned} AnsTurn(incp(i), do(a,s)) \equiv\ & a = asks(i,Q) \vee \\ & a = no(i) \vee a = yes(incp(i)) \end{aligned} \tag{2.30}$$

Phase Axioms

The Successor State Axioms for the phase fluents are in direct correlation with the finite state diagram shown in Figure 2.2.

$$\begin{aligned} askphase(do(a,s)) &\equiv a = endturn(i) \wedge In(incp(i)) \\ ansphase(do(a,s)) &\equiv a = asks(i,Q) \vee (a = no(j) \wedge \neg Turn(incp(j), s) \\ showphase(do(a,s)) &\equiv a = yes(i) \\ guessphase(do(a,s)) &\equiv a = shows(j,i,s_x) \vee (a = no(j) \wedge Turn(incp(j), s) \\ endphase(do(a,s)) &\equiv (a = endturn(i) \wedge \neg In(incp(i))) \vee \\ & \quad a = noguess(i) \vee a = guess(i,M) \end{aligned} \tag{2.31}$$

2.10 Players Gaining Knowledge

In this section, we will explore the methods by which knowledge is transferred between players. Axiom 0 asserts that all players are perfect reasoners that know all of the axioms and have perfect recall. At each stage of MYST, each player is trying to determine the location of all the cards, i.e. they are trying to gain knowledge about the fluent $h_{i,x}$. For which actions can a player gain such knowledge?

Reexamining the function $Poss(a,s)$, we find four actions which exchange such knowledge. They are: $no(j)$, $yes(j)$, $shows(j,i,s_x)$ and $guess(i,M)$. The actions no, yes and *shows* are more frequent actions than *guess*, and they are much more instructive to examine: a correct *guess* will end the game, and so only an incorrect *guess* stands to glean any information. We list it anyway for sake of completeness.

Knowledge Gained from $no(j)$

From (2.16), if $no(j)$ is a legal action, then we can conclude that

$$\forall Q \;\; Question(Q, s) \;\supset\; C\left(P, \bigwedge_{s_x \in Q} \neg h_{j,x}, do(no(j), s)\right). \tag{2.32}$$

Stated in words, if player j answers *no* to question Q, then it is common knowledge for all players that player j does not hold any of the cards in Q.

Knowledge Gained from $yes(j)$

From (2.17), if $yes(j)$ is a legal action, then we can conclude that

$$\forall Q \;\; Question(Q, s) \;\supset\; C\left(P, \bigvee_{s_x \in Q} h_{j,x}, do(yes(j), s)\right). \tag{2.33}$$

Stated in words, if player j answers *yes* to question Q, then it is common knowledge for all players that player j holds at least one of the cards in Q.

Note that the knowledge gained by player i from $shows(j, i, s_x)$ (2.25) implies the common knowledge gained by all players P from the previous answer $yes(j)$ in (2.33).

Knowledge Gained from $shows(j, i, s_x)$

This was described by Axiom (2.25).

Knowledge Gained from $guess(i, M)$

If $guess(i, M)$ is a legal action and the game continues then, player i guessed wrong and therefore $S_0 \neq M$. All players can conclude that one of the cards was incorrect.

$$a = guess(i, M) \wedge \neg Gameover(do(a, s)) \;\supset\; C\left(P, \bigvee_{s_x \in M} \neg h_{0,x}, do(a, s)\right) \tag{2.34}$$

Chapter 3

Properties of MYST

3.1 A Crucial Theorem

It is important to recall at this point that MYST is a non-cooperative game, with specific, clearly defined goals. For any player to have a certain win, they must be able to determine the entire contents of the mystery pile with certainty. Mathematically speaking, player i has a certain win in situation s *iff* for some set Σ of size m,

$$K\left(i, \bigwedge_{s_x \in \Sigma} h_{0,x}, s\right) \tag{3.1}$$

The winning condition in (3.1) could be obtained by slowly acquiring several $K(i, h_{0,x}, s)$ over the course of a game. Alternatively, (3.1) might be attainable in one fell swoop as a consequence of the Set Size Axiom (2.12). In either case, any positive knowledge regarding the holdings of the mystery pile cannot be determined without first determining the card(s) are not held in the remaining places. This is the subject of our first theorem.

Theorem 3.1 *If, in situation s, player i has knowledge $K(i, h_{0,x}, s)$, then it is both sufficient and necessary that he used the Existence Axiom (2.10) or the Set Size Axiom (2.12) to prove it. i.e.*

- $K\left(i, \bigwedge_{1 \le j \le p} \neg h_{j,x}, s\right)$ *or*
- $K\left(i, \bigwedge_{s_y \notin S_0} \neg h_{0,y}, s\right)$

Proof: It is easily seen that these conditions are sufficient, by use of the appropriate axiom. The proof of necessity follows directly from the axioms in the formal specification. There are only two axioms which can conclude $h_{0,x}$ and they are the Existence Axiom and the Set Size Axiom. So any conclusions about $h_{0,x}$ must use one of the premises of these two axioms. □

We might try to extend Theorem 3.1 to include reasoning about the cards not held by the mystery pile. We would have to make a similar conclusion about the Disjoint Axiom, and this is not possible. Here is a counterexample wherein a player will know a card is not held in the mystery pile without first knowing which player holds it.

Suppose that in a 5 player game of MYST all players ask two-card questions. Player 1 asks $\{s_1, s_2\}$ and player 2 says *yes*, and then on a later turn, player 3 asks $\{s_1, s_2\}$ and player 4 says *yes*. All of players $\{1, 2, 3, 4\}$ have enough knowledge to deduce the location of both cards $\{s_1, s_2\}$, but what about player 5?

Intuition suggests that the cards $\{s_1, s_2\}$ can be only in one of two configurations: $h_{2,1} \wedge h_{4,2}$ or $h_{4,1} \wedge h_{2,2}$ (see equation $(*)$ below), and therefore neither card can be in anyone else's hand, including the mystery pile. Player 5 knows $(h_{2,1} \vee h_{2,2}) \wedge (h_{4,1} \vee h_{4,2})$ and can resolve this using the axioms as follows:

$$
\begin{aligned}
& (h_{2,1} \vee h_{2,2}) \wedge (h_{4,1} \vee h_{4,2}) \\
\equiv\ & (h_{2,1} \wedge h_{4,1}) \vee (h_{2,1} \wedge h_{4,2}) \vee (h_{2,2} \wedge h_{4,1}) \vee (h_{2,2} \wedge h_{4,2}) \\
\equiv\ & F \vee (h_{2,1} \wedge h_{4,2}) \vee (h_{2,2} \wedge h_{4,1}) \vee F && \text{(Disjoint Axiom)} \\
\equiv\ & (h_{2,1} \wedge h_{4,2}) \vee (h_{2,2} \wedge h_{4,1}) && (*) \\
\equiv\ & \left(\bigwedge_{i \neq 2} \neg h_{i,1} \wedge \bigwedge_{i \neq 4} \neg h_{i,2} \right) \vee \left(\bigwedge_{i \neq 2} \neg h_{i,2} \wedge \bigwedge_{i \neq 4} \neg h_{i,1} \right) && \text{(Disjoint Axiom)} \\
\equiv\ & \left(\bigwedge_{i \neq 2,4} \neg h_{i,1} \wedge \bigwedge_{i \neq 2,4} \neg h_{i,2} \wedge \neg h_{4,1} \wedge \neg h_{2,2} \right) \vee \\
& \quad \left(\bigwedge_{i \neq 2,4} \neg h_{i,1} \wedge \bigwedge_{i \neq 2,4} \neg h_{i,2} \wedge \neg h_{2,1} \wedge \neg h_{4,2} \right) \\
\equiv\ & \bigwedge_{i \neq 2,4} (\neg h_{i,1} \wedge \neg h_{i,2}) \wedge [(\neg h_{2,1} \wedge \neg h_{4,2}) \vee (\neg h_{4,1} \wedge \neg h_{2,2})] \\
\supset\ & \bigwedge_{i \neq 2,4} (\neg h_{i,1} \wedge \neg h_{i,2})
\end{aligned}
$$

Therefore player 5 knows $\neg h_{0,1}$ and $\neg h_{0,2}$, but might not know for certain the location of either $\{s_1, s_2\}$. In fact, player 5 does not need to know the exact locations of these cards to win, which was asserted by Theorem 3.1.

To conclude the discussion of what is and what is not possible, we will present some elementary knowledge consequences of asking questions. Disjunctions are generated only by questions for which some opponent asks the question and some other opponent answers *yes*. In practice, these end up comprising most of the questions, and so a disjunction is usually generated every turn. As we will see in Section 3.2, deductions using disjunctions could be computationally expensive.

Here are the knowledge consequences.

A response of *no*. By Equation (2.32), if a player responds *no* to a question, the information gleaned is the same for every player. The poser does not glean any more information than any other player who witnesses this action. Note also that all information is expressed as a conjunction of primitive predicates.

A response of *yes* and a *shows* action. Examine and compare equations (2.25) and (2.33). A response of *yes* yields a disjunction of $h_{i,x}$ predicates, except for the poser and the responder who said *yes*. Because of the rules of the game, the responder must follow their *yes* response by a *shows* action, and therefore, the poser will for certain know the location of a card in the responder's hand. Therefore, from the poser's perspective, the information gleaned by (2.33) is an implication of the information gleaned by (2.25) and so the poser gleaned no less than any of his opponents.

Moreover, the result of the *yes* response is of no direct use to the responder, because the Knowledge Initialization Axiom (2.14) implies the information gleaned by (2.25). One could make a similar self-knowledge observation for the *no* response.

Everyone responded *no*. Every player desires this result on their turn. If every opponent says *no* to a question Q, then for every card $s_x \in Q$ and opponent j, the poser can conclude $\neg h_{j,x}$. However, from the Knowledge Initialization Axiom (2.14) and the Partition Axioms, the poser knows which cards they do not hold. For every card

$s_x \in Q$ which is also not held by the poser, we obtain

$$\begin{aligned} & \forall i \ player(i) \supset \neg h_{i,x} \\ \equiv & \bigwedge_{player(i)} \neg h_{i,x} \\ \supset & \ h_{0,x} \qquad \text{(Existence Axiom)} \end{aligned}$$

Therefore, a question for which all opponents respond *no* determines a portion of the cards in the mystery pile.

This is the only result of a question which guarantees that *a portion* of the mystery pile is determined. Theorem 3.1 confirms this because it guarantees only two ways of determining positive truths about the holdings of the mystery pile. If the entire pile has not been determined by deductive elimination via the Set Size Axiom, then the card must have been determined by deductive elimination via the Existence Axiom.

In conclusion, there are two principle methods to determine the holdings of the mystery pile. The first method is by determining portions of the pile, one at a time, until the whole mystery pile has been determined. We call this the *direct method*, and we call the questions that yield a *no* response from all opponents thereby determining a portion of the mystery pile, *direct questions*. The second method is by determining the $c - m$ cards not contained in the mystery pile and then, by the Set Size Axiom, conclude the m cards in the mystery pile. We call this the *indirect method*.

3.2 Implementation of the Set Size Axiom (SSA)

How can we write an algorithm so that a player can take advantage of the Set Size Axiom (2.11) and (2.12)? First, the SSA asserts that if a player knows all the cards in S_0, then they have won. By Theorem 3.1, only direct questions can determine a portion of S_0. Algorithmically, we can keep track of S_0 by maintaining a ternary array and a counter. Each cell of the array will represent a different card and its entry will be one of $\{yes, no, ?\}$. It will be *yes* if the player knows the card is in S_0, *no* if the player knows the card is not in S_0, and ? if the player does not know. The counter will total the *yes* entries, until there are m of them.

Second, the SSA asserts that if player i knows that player j holds n cards, they can deduce he does not hold the remaining cards. If the n cards are known to be held for certain, then this is no more interesting to implement than for S_0 because the same ternary array and counter will suffice. However, it is possible to use the SSA to make conclusions even if there is some uncertainty about player j's cards. To illustrate this, let us consider the following example game of MYST.

Consider a 3 player game where $n = q = 2$ and $S = \{s_1, s_2, \ldots, s_{15}\}$. Player 1 has asked two questions so far, and each time, player 2 answered *yes* and showed player 1 a card. Player 1's questions were $Q^1 = \{s_1, s_2\}, Q^2 = \{s_3, s_4\}$. Of course, player 1 has completely determined player 2's hand, but player 3 may still be in doubt. Player 3 uses (2.32) on Q^1, Q^2, and concludes $[(h_{2,1} \vee h_{2,2}) \wedge (h_{2,3} \vee h_{2,4})]$, and therefore deduces that player 2 must hold two of these cards and therefore cannot hold any of $\{s_5, s_6, \ldots, s_{15}\}$, because $|S_2| = 2$.

What deduction takes place for player 3 to make such a conclusion?

If player i is suggested to hold two cards, say $S_i = \{s_a, s_b\}$, then they are suggested not to hold $S - S_i$. Using the SSA, we get

$$h_{i,a} \wedge h_{i,b} \equiv \bigwedge_{x \neq a,b} \neg h_{i,x}. \tag{3.2}$$

Therefore for player 3 in this example we get

$$\begin{aligned} & K(3, [(h_{2,1} \vee h_{2,2}) \wedge (h_{2,3} \vee h_{2,4})], s) \\ \equiv\ & K(3, [(h_{2,1} \wedge h_{2,3}) \vee (h_{2,1} \wedge h_{2,4}) \vee (h_{2,2} \wedge h_{2,3}) \vee (h_{2,2} \wedge h_{2,4})], s) \end{aligned}$$

Since each conjunction describes player 2 holding at least two cards, player 3 can conclude from the disjunction that player 2 holds at least two cards in one of these configurations by using the substitution in (3.2) as follows.

$$\begin{aligned} & K(3, [(h_{2,1} \wedge h_{2,3}) \vee (h_{2,1} \wedge h_{2,4}) \vee (h_{2,2} \wedge h_{2,3}) \vee (h_{2,2} \wedge h_{2,4})], s) \\ \equiv\ & K\left(3, \left[\left[\bigwedge_{x \neq 1,3} \neg h_{2,x}\right] \vee \left[\bigwedge_{x \neq 1,4} \neg h_{2,x}\right] \vee \left[\bigwedge_{x \neq 2,3} \neg h_{2,x}\right] \vee \left[\bigwedge_{x \neq 2,4} \neg h_{2,x}\right]\right], s\right) \\ \equiv\ & K\left(3, \left[\bigwedge_{x \neq 1,2,3,4} \neg h_{2,x}\right] \wedge [(\neg h_{2,2} \wedge \neg h_{2,4}) \vee (\neg h_{2,2} \wedge \neg h_{2,3}) \vee (\neg h_{2,1} \wedge \neg h_{2,4}) \vee (\neg h_{2,1} \wedge \neg h_{2,3})], s\right) \\ \supset\ & K\left(3, \left[\bigwedge_{x \neq 1,2,3,4} \neg h_{2,x}\right], s\right) \end{aligned}$$

And therefore any card not in any conjunction is known to be not in player 2's hand. We can now conclude $K\left(3, \bigwedge_{5 \leq x \leq 15} \neg h_{2,x}, s\right)$.

This concludes the example. What can we glean from similar examples, but in which the search space is large? If, by witness, player 3 observes player 2 answer *yes* to five questions to get the expression $K(3, \phi, s)$, where

$$\phi = (h_{2,1} \vee h_{2,2}) \wedge (h_{2,2} \vee h_{2,3}) \wedge (h_{2,3} \vee h_{2,4}) \wedge (h_{2,4} \vee h_{2,5}) \wedge (h_{2,5} \vee h_{2,1}),$$

then all five of $s_1, s_2, \ldots, s_5$ might be in player 2's hand. But really, we are interested in the minimum required number of cards which make ϕ true. If the minimum equals n, then we can apply the SSA. In this example ϕ, there must be at least 3 cards in player 2's hand and we therefore need $n = 3$ to use the SSA.

This is algorithmically equivalent to expanding the expression ϕ, but that would lead to a likely 2^γ conjunctions stemming from γ original disjunctions. What is the underlying problem and can it be solved efficiently?

Problem: Let ϕ be a monotonic logical expression in conjunctive normal form, where each clause has two literals, and ϕ may contain up to c different literals.

Question: Are there β different literals, $l_1, \ldots, l_\beta$ such that $\{l_1, \ldots, l_\beta\} \vdash \phi$?

This is the (monotonic) MIN 2-SAT problem which is known to be reducible to VERTEX COVER [6], and therefore is NP-hard. We conclude that there is no known efficient way of practically applying the SSA algorithmically for arbitrary c. It should be noted, however, that for most practical instances of game-playing, c will be no larger than a few dozen and expanding ϕ can be easily handled by today's processors.

3.3 Solvable and Unsolvable Instances of MYST

In what instances of MYST does a solution exist? Since we can define any instance of MYST in terms of the constants m, n, p, q, we might reformulate the above question and ask: for which 4-tuples of constants (m, n, p, q) can a player determine the cards in the mystery pile with certainty?

Let us assume that this player lives in a "vacuum" wherein the opponents will not divulge any information about their hands by the questions they ask. Alternately stated, this player can only glean information from the questions they ask. For readability, we will we assume the role of this player and speak in the first person. Our goal is to determine if there always exists a sequence of questions which will determine S_0 with certainty. Such a sequence will be referred to as a *winning sequence of questions.* We show that such a sequence exists *iff* $1 \leq q \leq m+n+1$.

Without any loss of generality, in each of the four lemmas and the discussion which follows, it is assumed that we are acting as player 1 and that our opponents are acting in the most rational manner.

Lemma 3.1 *A winning sequence of questions exists when $1 \leq q \leq m$.*

Proof: Since $|S_0| = m$ is finite, we can enumerate its contents, say by the sequence $s_0, \ldots, s_{m-1}$.

We define the winning sequence of $\lceil \frac{m}{q} \rceil$ questions $Q^1, \ldots, Q^{\lceil \frac{m}{q} \rceil}$ as follows.

$$Q^i = \{s_x : x = ((i-1) * q + a) \bmod m;\ 0 \leq a \leq q-1;\ a \in \mathcal{Z}\}.$$

It is clear that all of $Q^1, \ldots, Q^{\lceil \frac{m}{q} \rceil}$ are direct questions. Since $\bigcup_{j=1}^{\lceil \frac{m}{q} \rceil} Q_0^j = S_0$, then by the Partition Axioms, the result follows. □

Lemma 3.2 *A winning sequence of questions exists when $m < q \leq m+n$.*

Proof: Let $T \subseteq S_1$, where $|T| = q - m$. It is possible to find such a subset since $m < q \leq m+n$, and therefore $0 < q - m \leq n = |S_1|$.

Now let $Q = T \cup S_0$. Note that $|Q| = (q-m) + m = q$ because of the Disjoint Axiom.

The single direct question Q is sufficient to win MYST. By the Existence Axiom, $Q_0 = Q \subseteq S_0$ and since $|Q_0| = |S_0|$, then the mystery pile is completely determined. □

Lemma 3.3 *A winning sequence of questions exists when $q = m+n+1$.*

Proof: Unlike in the above lemmata, no direct questions are possible when $q = m+n+1$. This is easily seen by applying the Pigeonhole Principle: for every question of size q, there is at least one card which is in an opponent's hand.

By Theorem 3.1, we must use the indirect method to win the game, i.e. determine the location of all the cards which are not in the mystery pile. We claim that for $q = m+n+1$, there exists a question which will locate an opponent's card whose location we did not already know.

To prove the claim, suppose that at a certain stage of the game, player i has shown us α cards as the result of some previous questions. Now, $0 \leq \alpha < n$, otherwise if $\alpha = n$ we know all there is to know about player i's hand. Our strategy is: do not ask about the α cards we already know.

So, let's ask about the n cards in our own hand, plus S_0, plus a card in the opponent's hand that we have not yet seen. This question will only be answered *yes* by player i, and they must prove it to us by showing us the card we have not yet seen! □

Lemma 3.3 is interesting because it makes comments about how hard a game MYST could be for large q. When $q = m + n + 1$, it is possible for any single player to construct at most $n(p-1)$ questions and determine S_0 with absolute certainty, and this certainly continues to hold as an upper bound when q is smaller. In a future section we will describe an algorithm which achieves the upper bound independently of the initial distribution of the cards. This algorithm is also the optimal strategy for the case when $p = 2$ and $q = m+n+1$.

For a lower bound, it is possible (but improbable) for us to succeed in only n questions, simply by asking questions constructed using the claim and the cards held by player p. While on the subject of lower bounds, for $1 \leq q \leq m$, the lower bound of questions required is $\min\{n(p-1), \lceil \frac{m}{q} \rceil\}$ and for $m < q \leq m + n$, the lower bound is 1 question.

Lemma 3.4 *No winning sequence of questions exists when $q > m + n + 1$.*

Proof: It is sufficient to prove the lemma for $q = m + n + 2$, because the argument is extendible for any larger q. We can claim that for $q = m + n + 2$, our right hand opponent (player p) can prevent us from knowing at least one card in his hand. If the claim is true, then by using Theorem 3.1, neither the direct method nor the indirect method is possible and a guess cannot be made with certainty of a win.

Let $\mathcal{S}_p \subset S_p$ be the set of cards that player p has shown us to date and suppose $|\mathcal{S}_p| = n - 1$. If we can show that no question can force player p to disclose his last card, we have proven the claim.

Consider the existence of such a question Q. If Q makes reference to any card held by players $2, \ldots, p-1$, then that opponent will answer *yes* before player p has to respond, and Q will gain us nothing about player p's holdings. Therefore, we are left with $2n + m$ cards which could possibly be in Q.

- n cards in our hand. (S_1)
- m cards in the mystery pile. (S_0)
- 1 card in player p's hand which is unknown. $(S_p - \mathcal{S}_p)$
- $n-1$ cards in player p's hand which are known. $(\mathcal{S}_p)$

By the pigeonhole principle, any question with $n+m+2$ cards will contain at least one card, $s_x \in \mathcal{S}_p$. Since it is common knowledge between players 1 and p that player p holds all the cards in $\mathcal{S}_p$, and since player p is acting in the most rational manner, player p will show us card s_x again. $\square$

Theorem 3.2 *A winning sequence of questions exists* iff $1 \leq q \leq m+n+1$.

Proof: This follows directly from the proofs of the previous four lemmas. $\square$

Chapter 4

The Two-Player Version of MYST

The only analyzable game which is free from the aspects of collusion (but not luck) is the two player version of MYST. For instance, in a game with three (or more) players, a player may intentionally ask questions whose *yes* response from one opponent will guarantee to supply extra information to a third party. What alliances each player chooses to make (or break) with other players and the strategy involved in picking alliances requires the analysis of n-person game theory. For the instance of MYST, this analysis is infeasible.

Assuming $p = 2$ makes MYST axiomatically simpler. For instance, the knowledge gleaned from a response of *yes* (2.33) can never be witnessed by a third party because there are only two players. Therefore, this knowledge is always implied by the knowledge gleaned from the subsequent *shows* action (2.25) or by the Knowledge Initialization Axiom (2.14). (Refer to the discussion in Section 3.1.) Also, the *shows* action, which implies common knowledge shared between the poser and the responder, reduces to common knowledge among all players in a two player game.

In fact, there can only be two results from a question in a two player game.

Lemma 4.1 *In a two-player game of MYST, there are only two possible outcomes to a question Q:*

- *proving a single card in the opponent's hand,* or
- *proving $Q_0 \subseteq S_0$.*

Proof: Without loss of generality, suppose player 1 asks Q. Player 2 can respond either *yes* or *no*. A *yes* response will be followed by a *shows* action which will prove a single card in the opponent's hand, and a *no* response makes Q a direct question thereby proving $Q_0 \subseteq S_0$. $\square$

Knowledge gleaned from questions with a response of *yes* is common knowledge between both players. But what if a player responds *no*? We will later prove that if both players separately glean the same information from a *no* response, then that information is also common knowledge.

In the next section, we will present a randomized algorithm which is optimal in the case of maximum q. The ideas presented here will carry over into strategies for smaller q.

4.1 An Optimal Strategy for $q = m + n + 1; p = 2$

In this section, we present a summary of an optimal strategy for when $q = m + n + 1$. The maximum q case reduces the axiom set even further than suggested by Lemma 4.1, because, by the Pigeonhole Principle, no matter what question is asked by the poser, there exists a card in the question held by the opponent and therefore, they must respond *yes*.

Here is a summary of the strategy.

Suppose player 1 goes first and player 2 second. Let $\mathcal{S}_i$ be the set of cards that player i has shown player $incp(i)$. At any turn, player 1 (and symmetrically, player 2) will ask a question constructed as follows. Pick a *random* subset of q cards from $S - \mathcal{S}_1 - \mathcal{S}_2$. If $|S - \mathcal{S}_1 - \mathcal{S}_2| < q$, then pick all of them plus some cards from $\mathcal{S}_1$ to make up the difference.

The construction of this algorithm and proof of its optimality are the result of the following two claims.

Claim 4.1 *For $q = m+n+1$, the* most *information that a player can glean from a question is the location of one of their opponent's cards.*

Claim 4.2 *For $q = m + n + 1$, a player can ask a question which gives no information to their opponent.*

Proof of Claim 4.1: Using the Pigeonhole Principle, for every question of size $q = m+n+1$, there exists a card held by their opponent. Therefore, the only possible outcome (listed in Lemma 4.1) is to prove a single card in the opponent's hand. $\square$

Proof of Claim 4.2: Without loss of generality, assume that player 1 is the poser and player 2 is the responder. Assuming a uniform distribution of the cards in $Init_0$, player 2 considers all possible worlds equally likely.

Now suppose player 1 asks a random question Q, as prescribed in the strategy summary. Player 2 will know the locations of some cards in Q and at the very least he will know Q_2. The cards about which player 2 is unaware might be in player 1's hand or in the mystery pile. But since the question was randomly chosen with equal probability, player 2 cannot probabilistically partition Q. Therefore, all possible worlds for player 2 remain equally likely. □

Claim 4.3 *The algorithm is optimal.*

Proof: Since each question constructed by the algorithm gleans the maximum possible amount of information, while simultaneously preventing the opponent from gleaning any information, then the algorithm is optimal. □

Randomizing questions in this way could be viewed as bluffing, because one might ask questions pertaining to cards in one's own hand. However, since this sort of randomization will still glean information from the opponent, it carries no risk. We might better describe the randomization process as *camouflage* for one's own questions. We will discuss bluffing and camouflage in more detail in Section 4.2.

This algorithm guarantees to find a card in the opponent's hand on every turn and therefore will establish the contents of the opponent's hand with certainty in n turns. This algorithm can be extended for more players, however it may not remain optimal. The extended algorithm will continue to guarantee that the location of a card will be gleaned on any given question, and therefore, the mystery pile will be determined in no more than $n(p-1)$ questions, which was the upper bound suggested in Section 3.3.

This algorithm can also be extended for smaller q. But since it is possible to construct direct questions when $q < m+n+1$, it might require up to $n(p-1)+m-1$ questions to determine the mystery pile. No matter which extension, the principle of the algorithm remains the same: a question is chosen randomly from a set useful questions.

Note that the algorithm presented is optimal for choosing a question, and we have ignored the possibility that either player might try to guess before the entire contents of the mystery pile become known. The following theorem states when each player should guess.

Lemma 4.2 *After the initial deal, there are* $\begin{pmatrix} m+n \\ m \end{pmatrix}$ *different possible worlds for each player. After asking* α *questions, there are* $\begin{pmatrix} m+n-\alpha \\ m \end{pmatrix}$.

Proof: There are $m+2n$ cards in total and, in $Init_0$, a player will know the location of only their n personally held cards, leaving $m+n$ cards whose location is uncertain. By the Set Size Axiom (2.11), there are $\begin{pmatrix} m+n \\ m \end{pmatrix}$ possible distributions of the cards, all of them equally likely. After α questions, the poser will know α cards in their opponent's hand and therefore, there remains $m+n-\alpha$ cards and $\begin{pmatrix} m+n-\alpha \\ m \end{pmatrix}$ possible worlds. □

Lemma 4.3 *After player 1's* n^{th} *question, they will know* <u>*with certainty*</u> *the set* S_0. *Furthermore, player 1 and 2 can do this analysis and conclude this.*

Proof: We have already argued that n questions will determine n cards in the opponent's hand with certainty. Using the Set Size Axiom (2.11), player 1 can deduce the mystery pile. Since both players are logically omniscient (Axiom 0), they are able to conclude this lemma. □

Theorem 4.1 *Player 2's best chance to win is to guess the contents of* S_0 *after their* $n-1^{st}$ *turn. Their probability of winning is at most* $\frac{1}{m+1}$. *Player 1's best chance to win is to ask their* $n-1^{st}$ *question and not guess.*

Proof: If player 2 does not guess after their $n-1^{\mathrm{st}}$ turn, player 2 has no chance by Lemma 4.3. By Lemma 4.2, the number of possible worlds is $\begin{pmatrix} m+n-(n-1) \\ m \end{pmatrix} = m+1$, and therefore their probability of success is $\frac{1}{m+1}$. Game theory suggests that player 2 cannot help but gain by guessing at this stage of the game.

How should player 1 react to player 2's early guess strategy? Should player 1 prevent player 2's guessing strategy by guessing after their $n-1^{\mathrm{st}}$ turn, or should they let player 2 guess?

In a case by case analysis,

1. Player 1 guesses at turn $n-1$
   (a) Player 1 was right: $P(\text{Player 1 wins}) = \frac{1}{m+1}$.
   (b) Player 1 was wrong: $P(\text{Player 1 loses}) = \frac{m}{m+1}$.
2. Player 1 does not guess at turn $n-1 \Longrightarrow$ player 2 guesses at turn $n-1$.
   (a) Player 2 was right: $P(\text{Player 1 loses}) = \frac{1}{m+1}$.
   (b) Player 2 was wrong: $P(\text{Player 1 wins}) = \frac{m}{m+1}$.

It is clear that decision 2 is better than decision 1 for player 1. □

4.2 Camouflage and Bluffing

The algorithm presented for $q = m + n + 1$ uniformly selects a question from a set of useful questions, with the intention of camouflaging the poser's holdings. Nearly all of the useful questions name a card held by the poser with the intention of avoiding predictability. In fact, a player must not choose their questions in a predictable fashion, because the questions will divulge more information than they will stand to gain. Here is an example.

Suppose player 1 guesses α cards from their hand mixed with $q - \alpha$ cards from the rest, and that player 2 can predict player 1 will use this strategy. We should note that $\alpha < q$ or else player 1 is not asking a useful question. Player 2's counter-strategy is to ask the same question at least α times, but replacing already shown cards one at a time with other (randomly chosen) cards. After γ questions, where $\gamma \geq \alpha$, player 2 will have determined the α cards used in player 1's original question. Since player 1 is acting predictably, player 2 will consequently know the locations of the remaining $(q - \alpha)$ cards. We should view this last consequence as "bonus knowledge" for player 2.

Can player 1 recoup the loss? After γ questions, player 1 will determine the locations of γ cards, but after the same γ questions, player 2 will have located γ cards as well, plus the bonus $q - \alpha$ cards. Therefore player 2 has the advantage.

It is clear from this analysis, that any good strategy must vary the constant α in some random fashion.

Although some might view the subterfuge of camouflage as an act of bluffing, it is not so because the poser does not risk anything by asking such a question. A true bluff would be to ask a nonuseful question with the intention of misleading the responder.

When $q > 1$, bluffing is of no use. For maximum q, a question could be asked that gleans some fixed amount of information, while simultaneously divulging no information to the opponent. The question could be constructed by uniformly picking from a set of questions which were guaranteed to glean some information, i.e. useful questions. Since all questions are equally likely, the opponent cannot probabilistically reduce their number of possible worlds. This technique is still valid for smaller q, even as small as $q = 2$.

Can it ever be advantageous to bluff completely? In other words, could it be to a player's advantage to intentionally ask a question whose cards are all in their own hand? It is clear that this is really a bluff and not the random camouflage discussed earlier: such a question gleans no information for the poser, and can only serve to mislead their opponent. Our claim is that bluffing is not an optimal strategy when $q > 1$.

Claim 4.4 *Any strategy which employs bluffing is not optimal when $q > 1$.*

Proof: Recall that MYST is a game of gleaning information to conclude the contents of the mystery pile. Since information does not change and immutable knowledge does not go away, then if a player has more information than his opponent, they are more likely to win. Since player 1 goes first, then they have a decided advantage over player 2, because he can pick a question which is guaranteed to glean himself some information. In terms of probabilities, player 1's chances of winning MYST using this "slow and steady" approach are at least 0.5.

Now assume that a bluffing strategy is optimal for player 1. Player 2's counter-strategy is to use the "slow and steady" approach. Player 2 simply ignores what player 1 is doing, and the turn immediately after player 1 has bluffed, player 2 will ask a question which will (on average) glean themselves more information than player 1. Or, in other words, player 2 treats all of player 1's turns as bluffs, and therefore can never be misled. When player 1 bluffs, they have wasted a turn, and it is now as if player 2 had gone first. Since the "slow and steady" approach yields a probability of at least 0.5 for player 2, then player 1 can do no better than a probability of 0.5.

Since, player 1 could have used the "slow and steady" approach and done better, the bluffing strategy is not optimal by contradiction. □

That concludes our analysis for when $q > 1$. But what if $q = 1$? Single card questions do not offer the luxury of extra cards to camouflage our intentions. In fact, in order to glean any information from a question, a player must ask a question about a card not in their own hand. Their opponent will respond *yes* or *no*: the *yes* response means the card in question is in his opponent's hand; the *no* response means the card is in the mystery pile. If it is assumed that each opponent never bluffs, then a *no* response will tell *both* players the card in question is in the mystery pile.

The main consequence of no bluffing is that players can guess the mystery pile using information gleaned by their opponents about the mystery pile. The prevention of early guessing seems to be the payoff for the bluffer. If a player bluffs, they lose the information they would normally have gleaned, however, they will prevent their opponent from guessing early, because their opponent may not be able to trust that every question was not a bluff. However, habitual bluffing is also not recommended, because the opponent may begin to employ an "echo" strategy wherein they repeat all of their opponent's questions, hoping to catch a bluff and win the game.

The above analysis suggests that a pure bluffing or a pure non-bluffing strategy will not yield the best expected result. Game theory suggests that a *mixed strategy* may be a solution. A mixed strategy is a set of pure strategies and a probability distribution that describes the strategy to employ. Such a strategy will not yield a repeatable result, but will yield a fixed expected value no matter which counter-strategy the opponent might employ.

Correct analysis of the mixed strategies requires the complete game tree (all possible strategies and their outcomes) and some matrix analysis. Unfortunately, even for the case of $n = q = 1$ and $m = 3$, the game tree made the analysis difficult. However, a game tree exists, it is finite and all of its outcomes are computable. Even though MYST is a game of imperfect information, we can still create the game tree where the payoffs represent the probabilities of winning. Now we can use the Theorem of von Neumann [13] about mixed strategies to conclude the following theorem.

Theorem 4.2 *There exists a mixed strategy for MYST which is optimal.*

4.3 Strategies for $1 < q \leq m + n$

According to Lemma 4.1, there are two possible outcomes to a question, a response of *yes* and a response of *no*. A *yes* response reveals the position of a card in the responder's hand

and makes it common knowledge. A *no* response reveals a portion of the mystery pile to the poser, and this is not common knowledge.

But what if both players (on separate turns) determine the same portion of the mystery pile? We claim that the knowledge in common about the mystery pile is common knowledge. To establish the claim, let's define some new sets.

Let $\mathcal{S}_i$ be the set of cards that player i has shown player j. By Lemma 4.1, both players have common knowledge about $\mathcal{S}_1 \cup \mathcal{S}_2$. Let $\mathcal{M}_i$ be the set of cards that player i knows are in the mystery pile S_0. The contents of $\mathcal{M}_i$ are not necessarily common knowledge (usually not). But we claim that for every $s_x \in \mathcal{M}_1 \cap \mathcal{M}_2$, $h_{0,x}$ is common knowledge.

Theorem 4.3 *$\mathcal{M}_1 \cap \mathcal{M}_2$ is common knowledge.*

Proof: Take any card $s_x \in \mathcal{M}_1 \cap \mathcal{M}_2$. By Theorem 3.1, there exists questions Q^1 and Q^2 which satisfied Lemma 4.1 for each respective player. i.e.

$$s_x \in Q^1, \quad Q^1_2 = \emptyset, \quad s_x \in Q^2, \quad Q^2_1 = \emptyset.$$

Without loss of generality, we will complete the argument from player 1's perspective. Player 1 has proven using Q^1 that $s_x \in S_0$. Player 1 knows the above lemmata, so he knows that player 2 knows Q^2_0. Since player 1 knows $s_x \in S_0$, he knows that $s_x \in Q^2_0$ and therefore $K(1, K(2, h_{0,x}, s), s)$. By following this argument by induction, $h_{0,x}$ is common knowledge. $\square$

Why is Theorem 4.3 important? Since the positions of all cards in $\mathcal{S}_1 \cup \mathcal{S}_2$ are common knowledge, these cards can be treated like they have been eliminated from play entirely; the original game has been reduced to a game of MYST with a smaller value of n. Similarly, Theorem 4.3 is a reduction result, because $\mathcal{M}_1 \cap \mathcal{M}_2$ can be eliminated from play. Therefore, we could mathematically quantify a two-player game of MYST in terms of the number of cards in each other's hand which remain undiscovered, and the number of cards each player has discovered in the mystery pile which are not common knowledge. i.e. the four-tuple:

$$(|S_1 - \mathcal{S}_1|, |S_2 - \mathcal{S}_2|, |\mathcal{M}_1 - \mathcal{M}_2|, |\mathcal{M}_2 - \mathcal{M}_1|).$$

Unfortunately, for the sake of analysis, it is difficult to quantify and compare the two different types of information gained by $\mathcal{S}_i$ and $\mathcal{M}_i$ and an optimal algorithm was not found. However, since the game is finite, an optimal strategy exists, and for $q > 1$ we claim:

Claim 4.5 *For a two-player game of MYST where $1 < q \leq m + n + 1$, the probability that player 1 wins is at least 0.5.*

Proof: To prove the claim we follow a similar argument used by Nash to prove the existence of a winning strategy for HEX [3], known as "strategy stealing." Suppose, by means of contradiction, that player 2 has an optimal strategy that guarantees the probability that player 2 wins is more than 0.5. In other words, this strategy means that the probability player 1 wins is less than 0.5. Since player 1 goes first, he can ask a useful random question at no cost and then employ player 2's strategy himself to get a probability of more than 0.5. $\square$

Note that we proved player 1 has a clear advantage, without exhibiting an optimal strategy. However, some general principles must be followed in order to have a successful strategy. As discussed in Section 4.2, players should always camouflage their questions and should never bluff when $q > 1$. In the next section, we will use backwards induction to analyze when $q = 1$ and bluffing is not allowed.

4.4 Backwards Induction for $q = 1$

In this section, we will examine the probability of winning if $q = 1$, which, as established in Section 4.2 is the only case where bluffing might be useful and camouflage is impossible. Initially we will assume that bluffing is not permitted, i.e. that players may only ask questions about cards that are not in their own hand.

A question about a card in the mystery pile makes that card's position known to both players (under the no bluffing assumption), which effectively reduces m by 1. A question about a card in the opponent's hand effectively reduces the opponent's n by 1. With these two results in mind, it becomes natural to describe any game in progress by a state vector: (m, n_1, n_2), where

- m is the current number of unknown cards in the mystery pile;
- n_1 is the current number of unknown cards in player 1's hand; and
- n_2 is the current number of unknown cards in player 2's hand.

All values for m, n_1 and n_2 must be positive for the game to continue, because a value of 0 means that one player has completely determined the mystery pile. Since, on each player's turn, a question will reduce one of these numbers by 1, we can conclude that the game is finite and that the state vectors are related by a partial order.

After a player asks a question, which places them in state (m', n_1', n_2'), they have two options.

1. Try to guess the mystery set. Note that in this case, the game ends.

$$P(\text{Guess right}, (m', n_1', n_2')) = \begin{pmatrix} m' + n_2' \\ m' \end{pmatrix}^{-1}$$

2. Let his opponent play.

$$P(\text{No guess}, (m', n_1', n_2')) = 1 - P(\text{win}, (m', n_2', n_1'))$$

Since they are behaving rationally, then they will take the maximum payoff of these two options. This description suggests a recursive relationship, described below. Note that we drop the "win" argument from all probability expressions.

Basis:

$$P(m, n_1, n_2) = 1 \qquad \text{if} \qquad m = 0 \quad \text{or} \quad n_2 = 0.$$

Recursive Relation:

$$\begin{aligned} P(m, n_1, n_2) = {} & \frac{m}{m+n_2} \max\left\{ \begin{pmatrix} m + n_2 - 1 \\ m - 1 \end{pmatrix}^{-1}, 1 - P(m-1, n_2, n_1) \right\} + \\ & \frac{n_2}{m+n_2} \max\left\{ \begin{pmatrix} m + n_2 - 1 \\ m \end{pmatrix}^{-1}, 1 - P(m, n_2 - 1, n_1) \right\} \end{aligned} \tag{4.1}$$

Because the state vectors are related by a partial order, any probability can be computed in polynomial time using dynamic programming. Here are some observations.

- If $n_1 < n_2$ then $P(m, n_1, n_2) \le 0.5$.
- If $P(m, n_1, n_2) < 0.5$ then $n_1 < n_2$.

- For $n_1 = n_2 = n$, $\lim_{m,n\to\infty} P(m, n_1, n_2) = 0.5$.

Although analyzing the large (m, n) cases might prove laborious, there appear to be some patterns for $n = 1$ or $m = 1$. They are summarized by the following two theorems, which tend to support the validity of the limit.

Theorem 4.4 *Let $n_1 = n_2 = 1$. For every integer $\alpha \geq 1$,*

$$\begin{aligned} P(m = 1, n_1, n_2) &= 1 \\ P(m = 2\alpha, n_1, n_2) &= \frac{\alpha+1}{2\alpha+1} \\ P(m = 2\alpha+1, n_1, n_2) &= \frac{1}{2} \end{aligned}$$

Theorem 4.5 *Let $m = 1$ and $n_1 = n_2 = n$. For every integer $n \geq 2$,*

$$P(m, n_1, n_2) = \frac{n+2}{2n+2}$$

Proof: We will prove Theorem 4.4 by a straightforward induction on m. Theorem 4.5 could be proved in a similar fashion.

The basis cases for $P(m = 1)$, $P(m = 2)$ and $P(m = 3)$ are verified by simple substitution into the recurrence (4.1).

Assume true for some $\alpha \geq 1$ that $m = 2\alpha$ and $m = 2\alpha + 1$, i.e.

$$\begin{aligned} P(m = 2\alpha) &= \frac{\alpha+1}{2\alpha+1} \\ P(m = 2\alpha+1) &= \frac{1}{2} \end{aligned}$$

and show true for $m = 2\alpha + 2$ and $m = 2\alpha + 3$. i.e.

$$\begin{aligned} P(m = 2\alpha+2) &= \frac{\alpha+2}{2\alpha+3} \\ P(m = 2\alpha+3) &= \frac{1}{2} \end{aligned}$$

$$
\begin{aligned}
P(m = 2\alpha + 2) &= \frac{1}{2\alpha+3} + \frac{2\alpha+2}{2\alpha+3}\left[\max\left\{\frac{1}{2\alpha+2}, 1 - P(2\alpha+1)\right\}\right] \\
&= \frac{1}{2\alpha+3} + \frac{2\alpha+2}{2\alpha+3}\left[\max\left\{\frac{1}{2\alpha+2}, \frac{1}{2}\right\}\right] \\
&= \frac{1}{2\alpha+3} + \frac{2\alpha+2}{2\alpha+3}\left[\max\left\{\frac{1}{4}, \frac{1}{2}\right\}\right] \qquad (\alpha \geq 1.) \\
&= \frac{1}{2\alpha+3} + \frac{2\alpha+2}{2\alpha+3} \cdot \frac{1}{2} \\
&= \frac{\alpha+2}{2\alpha+3}
\end{aligned}
$$

$$
\begin{aligned}
P(m = 2\alpha + 3) &= \frac{1}{2\alpha+4} + \frac{2\alpha+3}{2\alpha+4}\left[\max\left\{\frac{1}{2\alpha+3}, 1 - P(2\alpha+2)\right\}\right] \\
&= \frac{1}{2\alpha+4} + \frac{2\alpha+3}{2\alpha+4}\left[\max\left\{\frac{1}{2\alpha+3}, 1 - \frac{\alpha+2}{2\alpha+3}\right\}\right] \\
&= \frac{1}{2\alpha+4} + \frac{2\alpha+3}{2\alpha+4}\left[\max\left\{\frac{1}{2\alpha+3}, \frac{\alpha+1}{2\alpha+3}\right\}\right] \\
&= \frac{1}{2\alpha+4} + \frac{2\alpha+3}{2\alpha+4} \cdot \frac{\alpha+1}{2\alpha+3} \\
&= \frac{\alpha+2}{2\alpha+4} \\
&= \frac{1}{2}
\end{aligned}
$$

□

Chapter 5

Conclusions and Future Work

MYST is a game of static information which, because of its question asking actions, has strong ties to first-order logic and reasoning. To quantify the nature of MYST, we have used elements of the situation calculus and analyzed MYST using elements of game theory.

Although lengthy but modular, the specification of MYST using the framework of the situation calculus has proven useful. All of the essential elements of MYST were axiomatized, including:

- turn order;
- possible actions;
- augmenting of knowledge by situations, and implementation of immutable knowledge;
- the Partition Axioms and Knowledge Initialization Axiom;
- gleaning knowledge from actions *no*, *yes*, *shows* and *guess*; and
- game ending conditions.

In examining MYST ourselves, we proved in Theorem 3.1 that there are only two possible ways that a player can conclude a card, say s_x, is in the mystery pile. The two ways are by first concluding:

- that s_x is not in any player's hand; or
- that $c - m$ different cards are not in the mystery pile.

This theorem is important because it describes what the players must accomplish in order to win. The theorem also suggests two different types of questions: *direct questions* for which all players say *no* and *indirect questions* for which some player says *yes*.

Using Theorem 3.1, we were able to prove Theorem 3.2 which asserted that a winning sequence of questions exists *iff* $1 \leq q \leq m + n + 1$. This is an interesting result for two reasons. First, it is independent of the number of players. Second, the proof of the theorem suggested an optimal algorithm for the largest possible q: $q = m + n + 1$.

We were also able to show that an algorithmic implementation of the Set Size Axiom is NP-hard, and therefore might be computationally expensive. This analysis also suggests that resolving disjunctions may be equally as expensive.

An optimal question constructing algorithm was presented for $q = m+n+1$ which utilized the structure of the proof of Theorem 3.2. This led to the optimal strategy presented in Theorem 4.1, which asserted that player 2's best strategy is to guess after they have asked their $n - 1^{\text{st}}$ question. The question construction algorithm uniformly picks one from a set of questions which will glean new information. The random choice served as *camouflage* for the poser's own holdings, therefore gaining the maximum amount of information, while giving none away. Game theory concludes that camouflage is required when $q > 1$, or else the opponents will take advantage.

In Claim 4.4, we showed that any strategy which employs bluffing is not optimal when $q > 1$, however, bluffing might be useful when $q = 1$.

In searching for optimal strategies for smaller q, we proved the reduction Theorem 4.3, which asserts that all knowledge in common is common knowledge. Therefore, if a card's position is known by both players, then we can eliminate it from play entirely, thereby reducing the number of cards in play. Therefore it was possible to develop a recurrence relation (4.1) for $q = 1$. Using dynamic programming, we computed some probabilities for this stochastic process, and we were able to prove Theorems 4.4 and 4.5 for the small cases.

In conclusion, we have presented some useful results that touch many areas of AI, including first-order logic, knowledge representation, the situation calculus and game theory, the center point of which is the game of MYST, our variant of the game of Clue. In formalizing the problem, we used the framework of the situation calculus to characterize the notion of immutable knowledge. The presented axiomatization of MYST will allow other research about this problem within our framework.

The axiomatization describing the gleaning of knowledge from actions is the basis for all strategy development. In the other half of this work, we presented strategies for playing MYST, the conclusions for which were based on the earlier axiomatization. Therefore all the presented strategies and theorems about playing MYST are useful, since MYST (and Clue) have never been formally studied.

Although we have made much progress, there is always future work to be done. In closing, here are some problems that would build on this work or would have made this work more complete.

- Incorporate the partitioning of S, just like in real Clue.

  Let S be a set of m sorts (in Clue, $m = 3$). Restrict S_0 to contain exactly one card of each sort and restrict all questions Q to contain exactly one card of each sort. Note that $m = q$.

  Theorem 3.2 still applies, so every instance is solvable since $q = m < m + n + 2$. Theorem 3.1 applies also as does the requirement of mixed strategies. Does partitioning S make the game any easier or any harder?

  This is an important problem because this property of Clue was not carried over to the variant MYST. Computer programs which play Clue will have to address this issue.

- Implement a computer program which plays Clue.

- For $q = 1$, show the recurrence relation (4.1) implies the limit:

  $n_1 = n_2 = n, \lim_{m,n \to \infty} P(\text{win}, (m, n_1, n_2)) = 0.5.$

  Theorems 4.4 and 4.5 surely suggest this result is true. This would be a great stochastic process research problem.

- Examine the strategies in multi-player games of MYST.

  The two-player assumption made the knowledge axioms lot easier to handle. Could observations be made in a three-player game that could not be made in a two-player game? Also, how should a player pick their questions in a multi-player game?

Bibliography

[1] Nesmith C. Ankeny. *Poker Strategy: Winning with Game Theory.* Basic Books, Inc., 1981.

[2] Elliott M. Avedon and Brian Sutton-Smith. *The Study of Games.* John Wiley & Sons, 1971.

[3] John D. Beasley. *The Mathematics of Games.* Oxford University Press, 1989.

[4] Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning About Knowledge.* The MIT Press, 1995.

[5] Ference Forgó, Jenö Szép, and Ferenc Szidarovszky. *Introduction to the Theory of Games: concepts, methods, applications.* Kluwer Academic Publishers, 1999.

[6] Dan Gusfield and Leonard Pitt. A bounded approximation for the minimum 2-sat problem. *Algorithmica*, 8(2):103–117, 1992.

[7] Patrick J. Hayes. The second naive physics manifesto. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal Theories of the Commonsense World*, chapter 1. Ablex Publishing Corporation, 1985.

[8] D. Marc Kilgour. Struggles for survival: A study of duels and related games. In *SIGNUM*, volume II, pages 13–21, 1975.

[9] Maurice Lancaster. Games mathematicians play. In *SIGNUM*, volume II, pages 6–12, 1975.

[10] Hector Levesque, Fiora Pirri, and Ray Reiter. Foundations for a calculus of situations. *Electronic Transactions on Artificial Intelligence*, 2:159–178, 1998.

[11] J. McCarthy and P.J. Hayes. Some philosophical problems from the standpoint of artificial intelligence. In D. Michie and B. Meltzer, editors, *Machine Intelligence 4*, pages 463–502. Edinburgh University Press, 1969.

[12] M. Newborn. *Kasparov versus Deep Blue: Computer Chess Comes of Age.* New York: Springer, 1997.

[13] T. Parthasarathy and T.E.S. Raghavan. *Some Topics in Two-Person Games.* American Elsevier, Inc., 1971.

[14] Anatol Rapoport. *Two-Person Game Theory: The Essential Ideas.* University of Michigan Press, 1966.

[15] Anatol Rapoport. *N-Person Game Theory: Concepts and Applications.* University of Michigan Press, 1970.

[16] Eric Rasmusen. *Games and Information: An Introduction to Game Theory.* Basil Blackwell Inc., 1989.

[17] Jonathan Schaeffer. Marion tinsley: human perfection at checkers? *Math. Sci. Res. Inst. Publ.*, 29:115–118, 1996.

[18] Jonathan Schaeffer and Robert Lake. Solving the game of checkers. *Math. Sci. Res. Inst. Publ.*, 29:119–133, 1996.

[19] J. Uiterwijk. The kasparov-deep blue match. *International Computer Chess Association Journal*, 19(1):38–41, 1996.