Finding Equilibria in Large Sequential Games of Imperfect Information[*]

Andrew Gilpin
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA, USA

gilpin@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA, USA

sandholm@cs.cmu.edu

ABSTRACT

Finding an equilibrium of an extensive form game of imperfect information is a fundamental problem in computational game theory, but current techniques do not scale to large games. To address this, we introduce the *ordered game isomorphism* and the related *ordered game isomorphic abstraction transformation*. For a multi-player sequential game of imperfect information with observable actions and an ordered signal space, we prove that any Nash equilibrium in an abstracted smaller game, obtained by one or more applications of the transformation, can be easily converted into a Nash equilibrium in the original game. We present an algorithm, *GameShrink*, for abstracting the game using our isomorphism exhaustively. Its complexity is $\tilde{O}(n^2)$, where n is the number of nodes in a structure we call the signal tree. It is no larger than the game tree, and on nontrivial games it is drastically smaller, so *GameShrink* has time and space complexity *sublinear* in the size of the game tree. Using *GameShrink*, we find an equilibrium to a poker game with 3.1 billion nodes—over four orders of magnitude more than in the largest poker game solved previously. We discuss several electronic commerce applications for *GameShrink*. To address even larger games, we introduce approximation methods that do not preserve equilibrium, but nevertheless yield (*ex post*) provably close-to-optimal strategies.

Categories and Subject Descriptors: I.2 [Artificial Intelligence], F. [Theory of Computation], J.4 [Social and Behavioral Sciences]: Economics.

General Terms: Algorithms, Economics, Theory.

Keywords: game theory, sequential games of imperfect information, automated abstraction, equilibrium finding, computer poker.

1. INTRODUCTION

In environments with more than one agent, an agent's outcome is generally affected by the actions of the other agent(s). Consequently, the optimal action of one agent can depend on the others. Game theory provides a normative framework for analyzing such strategic situations. In particular, it provides solution concepts that define what rational behavior is in such settings. The most famous and important solution concept is that of *Nash equilibrium* [36]. It is a strategy profile (one strategy for each agent) in which no agent has incentive to deviate to a different strategy. However, for the concept to be operational, we need algorithmic techniques for finding an equilibrium.

Games can be classified as either games of *perfect information* or *imperfect information*. Chess and Go are examples of the former, and, until recently, most game playing work has been on games of this type. To compute an optimal strategy in a perfect information game, an agent traverses the game tree and evaluates individual nodes. If the agent is able to traverse the entire game tree, she simply computes an optimal strategy from the bottom-up, using the principle of *backward induction*.[1] In computer science terms, this is done using *minimax search* (often in conjunction with α-β-*pruning* to reduce the search tree size and thus enhance speed). Minimax search runs in linear time in the size of the game tree.[2]

The differentiating feature of games of imperfect information, such as poker, is that they are not fully observable: when it is an agent's turn to move, she does not have access to all of the information about the world. In such games, the decision of what to do at a point in time cannot generally be optimally made without considering decisions at all other points in time (including ones on other paths of play) because those other decisions affect the probabilities of being at different states at the current point in time. Thus the algorithms for perfect information games do not solve games of imperfect information.

For sequential games with imperfect information, one could try to find an equilibrium using the normal (matrix) form, where every contingency plan of the agent is a pure strategy for the agent.[3] Unfortunately (even if equivalent strategies

---

---

[1]This actually yields a solution that satisfies not only the Nash equilibrium solution concept, but a stronger solution concept called *subgame perfect Nash equilibrium* [45].

[2]This type of algorithm still does not scale to huge trees (such as in chess or Go), but effective game-playing agents can be developed even then by evaluating intermediate nodes using a heuristic evaluation and then treating those nodes as leaves.

[3]An ϵ-equilibrium in a normal form game with any

are replaced by a single strategy [27]) this representation is generally exponential in the size of the game tree [52].

By observing that one needs to consider only sequences of moves rather than pure strategies [41, 46, 22, 52], one arrives at a more compact representation, the *sequence form*, which is linear in the size of the game tree.[4] For 2-player games, there is a polynomial-sized (in the size of the game tree) linear programming formulation (linear complementarity in the non-zero-sum case) based on the sequence form such that strategies for players 1 and 2 correspond to primal and dual variables. Thus, the equilibria of reasonable-sized 2-player games can be computed using this method [52, 24, 25].[5] However, this approach still yields enormous (unsolvable) optimization problems for many real-world games, such as poker.

1.1 Our approach

In this paper, we take a different approach to tackling the difficult problem of equilibrium computation. Instead of developing an equilibrium-finding method *per se*, we instead develop a methodology for automatically abstracting games in such a way that any equilibrium in the smaller (abstracted) game corresponds directly to an equilibrium in the original game. Thus, by computing an equilibrium in the smaller game (using any available equilibrium-finding algorithm), we are able to construct an equilibrium in the original game. The motivation is that an equilibrium for the smaller game can be computed drastically faster than for the original game.

To this end, we introduce *games with ordered signals* (Section 2), a broad class of games that has enough structure which we can exploit for abstraction purposes. Instead of operating directly on the game tree (something we found to be technically challenging), we instead introduce the use of *information filters* (Section 2.1), which coarsen the information each player receives. They are used in our analysis and abstraction algorithm. By operating only in the space of filters, we are able to keep the strategic structure of the game intact, while abstracting out details of the game in a way that is lossless from the perspective of equilibrium finding. We introduce the *ordered game isomorphism* to describe strategically symmetric situations and the *ordered game isomorphic abstraction transformation* to take advantage of such symmetries (Section 3). As our main equilibrium result we have the following:

Theorem 2 *Let Γ be a game with ordered signals, and let F be an information filter for Γ. Let F' be an information filter constructed from F by one application of the ordered game isomorphic abstraction transformation, and let σ' be a Nash equilibrium strategy profile of the induced game $\Gamma_{F'}$ (i.e., the game Γ using the filter F'). If σ is constructed by using the corresponding strategies of σ', then σ is a Nash equilibrium of Γ_F.*

The proof of the theorem uses an equivalent characterization of Nash equilibria: σ is a Nash equilibrium if and only if there exist beliefs μ (players' beliefs about unknown information) at all points of the game reachable by σ such that σ is sequentially rational (*i.e.*, a best response) given μ, where μ is updated using Bayes' rule. We can then use the fact that σ' is a Nash equilibrium to show that σ is a Nash equilibrium considering only local properties of the game.

We also give an algorithm, *GameShrink*, for abstracting the game using our isomorphism exhaustively (Section 4). Its complexity is $\tilde{O}(n^2)$, where n is the number of nodes in a structure we call the signal tree. It is no larger than the game tree, and on nontrivial games it is drastically smaller, so *GameShrink* has time and space complexity *sublinear* in the size of the game tree. We present several algorithmic and data structure related speed improvements (Section 4.1), and we demonstrate how a simple modification to our algorithm yields an approximation algorithm (Section 5).

1.2 Electronic commerce applications

Sequential games of imperfect information are ubiquitous, for example in negotiation and in auctions. Often aspects of a player's knowledge are not pertinent for deciding what action the player should take at a given point in the game. On the trivial end, some aspects of a player's knowledge are never pertinent (*e.g.*, whether it is raining or not has no bearing on the bidding strategy in an art auction), and such aspects can be completely left out of the model specification. However, some aspects can be pertinent in certain states of the game while they are not pertinent in other states, and thus cannot be left out of the model completely. Furthermore, it may be highly non-obvious which aspects are pertinent in which states of the game. Our algorithm automatically discovers which aspects are irrelevant in different states, and eliminates those aspects of the game, resulting in a more compact, equivalent game representation.

One broad application area that has this property is sequential negotiation (potentially over multiple issues). Another broad application area is sequential auctions (potentially over multiple goods). For example, in those states of a 1-object auction where bidder A can infer that his valuation is greater than that of bidder B, bidder A can ignore all his other information about B's signals, although that information would be relevant for inferring B's exact valuation. Furthermore, in some states of the auction, a bidder might not care which exact other bidders have which valuations, but cares about which valuations are held by the other bidders in aggregate (ignoring their identities). Many open-cry sequential auction and negotiation mechanisms fall within the game model studied in this paper (specified in detail later), as do certain other games in electronic commerce, such as sequences of take-it-or-leave-it offers [42].

Our techniques are in no way specific to an application. The main experiment that we present in this paper is on

---

constant number of agents can be constructed in quasi-polynomial time [31], but finding an exact equilibrium is PPAD-complete even in a 2-player game [8]. The most prevalent algorithm for finding an equilibrium in a 2-agent game is *Lemke-Howson* [30], but it takes exponentially many steps in the worst case [44]. For a survey of equilibrium computation in 2-player games, see [53]. Recently, equilibrium-finding algorithms that enumerate supports (*i.e.*, sets of pure strategies that are played with positive probability) have been shown efficient on many games [40], and efficient mixed integer programming algorithms that search in the space of supports have been developed [43]. For more than two players, many algorithms have been proposed, but they currently only scale to very small games [19, 34, 40].

[4]There were also early techniques that capitalized in different ways on the fact that in many games the vast majority of pure strategies are not played in equilibrium [54, 23].

[5]Recently this approach was extended to handle computing *sequential equilibria* [26] as well [35].

a recreational game. We chose a particular poker game as the benchmark problem because it yields an extremely complicated and enormous game tree, it is a game of imperfect information, it is fully specified as a game (and the data is available), and it has been posted as a challenge problem by others [47] (to our knowledge no such challenge problem instances have been proposed for electronic commerce applications that require solving sequential games).

1.3 Rhode Island Hold'em poker

Poker is an enormously popular card game played around the world. The 2005 World Series of Poker had over $103 million dollars in total prize money, including $56 million for the main event. Increasingly, poker players compete in online casinos, and television stations regularly broadcast poker tournaments. Poker has been identified as an important research area in AI due to the uncertainty stemming from opponents' cards, opponents' future actions, and chance moves, among other reasons [5].

Almost since the field's founding, game theory has been used to analyze different aspects of poker [28; 37; 3; 51, pp. 186–219]. However, this work was limited to tiny games that could be solved by hand. More recently, AI researchers have been applying the computational power of modern hardware to computing game theory-based strategies for larger games. Koller and Pfeffer determined solutions to poker games with up to 140,000 nodes using the sequence form and linear programming [25]. Large-scale approximations have been developed [4], but those methods do not provide any guarantees about the performance of the computed strategies. Furthermore, the approximations were designed manually by a human expert. Our approach yields an automated abstraction mechanism along with theoretical guarantees on the strategies' performance.

Rhode Island Hold'em was invented as a testbed for computational game playing [47]. It was designed so that it was similar in style to Texas Hold'em, yet not so large that devising reasonably intelligent strategies would be impossible. (The rules of Rhode Island Hold'em, as well as a discussion of how Rhode Island Hold'em can be modeled as a game with ordered signals, that is, it fits in our model, is available in an extended version of this paper [13].) We applied the techniques developed in this paper to find an exact (minimax) solution to Rhode Island Hold'em, which has a game tree exceeding 3.1 billion nodes.

Applying the sequence form to Rhode Island Hold'em directly without abstraction yields a linear program with 91,224,226 rows, and the same number of columns. This is much too large for (current) linear programming algorithms to handle. We used our *GameShrink* algorithm to reduce this with lossless abstraction, and it yielded a linear program with 1,237,238 rows and columns—with 50,428,638 non-zero coefficients. We then applied iterated elimination of dominated strategies, which further reduced this to 1,190,443 rows and 1,181,084 columns. (Applying iterated elimination of dominated strategies without *GameShrink* yielded 89,471,986 rows and 89,121,538 columns, which still would have been prohibitively large to solve.) *GameShrink* required less than one second to perform the shrinking (*i.e.*, to compute all of the ordered game isomorphic abstraction transformations). Using a 1.65GHz IBM eServer p5 570 with 64 gigabytes of RAM (the linear program solver actually needed 25 gigabytes), we solved it in 7 days and 17 hours

using the interior-point barrier method of CPLEX version 9.1.2. We recently demonstrated our optimal Rhode Island Hold'em poker player at the AAAI-05 conference [14], and it is available for play on-line at `http://www.cs.cmu.edu/~gilpin/gsi.html`.

While others have worked on computer programs for playing Rhode Island Hold'em [47], no optimal strategy has been found before. This is the largest poker game solved to date by over four orders of magnitude.

2. GAMES WITH ORDERED SIGNALS

We work with a slightly restricted class of games, as compared to the full generality of the extensive form. This class, which we call *games with ordered signals*, is highly structured, but still general enough to capture a wide range of strategic situations. A game with ordered signals consists of a finite number of rounds. Within a round, the players play a game on a directed tree (the tree can be different in different rounds). The only uncertainty players face stems from private signals the other players have received and from the unknown future signals. In other words, players observe each others' actions, but potentially not nature's actions. In each round, there can be public signals (announced to all players) and private signals (confidentially communicated to individual players). For simplicity, we assume—as is the case in most recreational games—that within each round, the number of private signals received is the same across players (this could quite likely be relaxed). We also assume that the legal actions that a player has are independent of the signals received. For example, in poker, the legal betting actions are independent of the cards received. Finally, the strongest assumption is that there is a partial ordering over sets of signals, and the payoffs are increasing (not necessarily strictly) in these signals. For example, in poker, this partial ordering corresponds exactly to the ranking of card hands.

DEFINITION 1. *A* game with ordered signals *is a tuple* $\Gamma = \langle I, G, L, \Theta, \kappa, \gamma, p, \succeq, \omega, u \rangle$ *where:*

1. $I = \{1, \ldots, n\}$ *is a finite set of players.*

2. $G = \langle G^1, \ldots, G^r \rangle$, $G^j = (V^j, E^j)$, *is a finite collection of finite directed trees with nodes V^j and edges E^j. Let Z^j denote the leaf nodes of G^j and let $N^j(v)$ denote the outgoing neighbors of $v \in V^j$. G^j is the* stage game *for round j.*

3. $L = \langle L^1, \ldots, L^r \rangle$, $L^j : V^j \setminus Z^j \to I$ *indicates which player acts (chooses an outgoing edge) at each internal node in round j.*

4. Θ *is a finite set of* signals.

5. $\kappa = \langle \kappa^1, \ldots, \kappa^r \rangle$ *and* $\gamma = \langle \gamma^1, \ldots, \gamma^r \rangle$ *are vectors of nonnegative integers, where κ^j and γ^j denote the number of public and private signals (per player), respectively, revealed in round j. Each signal $\theta \in \Theta$ may only be revealed once, and in each round every player receives the same number of private signals, so we require $\sum_{j=1}^{r} \kappa^j + n\gamma^j \leq |\Theta|$. The public information revealed in round j is $\alpha^j \in \Theta^{\kappa^j}$ and the public information revealed in all rounds up through round j is $\tilde{\alpha}^j = (\alpha^1, \ldots, \alpha^j)$. The private information revealed to player $i \in I$ in round j is $\beta_i^j \in \Theta^{\gamma^j}$ and the private information revealed to player $i \in I$ in all rounds up through round j is $\tilde{\beta}_i^j = (\beta_i^1, \ldots, \beta_i^j)$. We*

also write $\tilde{\beta}^j = \left(\tilde{\beta}_1^j, \ldots, \tilde{\beta}_n^j\right)$ to represent all private information up through round j, and $\left(\tilde{\beta}'^j_i, \tilde{\beta}^j_{-i}\right) = \left(\tilde{\beta}_1^j, \ldots, \tilde{\beta}_{i-1}^j, \tilde{\beta}'^j_i, \tilde{\beta}_{i+1}^j, \ldots, \tilde{\beta}_n^j\right)$ is $\tilde{\beta}^j$ with $\tilde{\beta}_i^j$ replaced with $\tilde{\beta}'^j_i$. The total information revealed up through round j, $\left(\tilde{\alpha}^j, \tilde{\beta}^j\right)$, is said to be legal if no signals are repeated.

6. p is a probability distribution over Θ, with $p(\theta) > 0$ for all $\theta \in \Theta$. Signals are drawn from Θ according to p without replacement, so if X is the set of signals already revealed, then

$$p(x \mid X) = \begin{cases} \frac{p(x)}{\sum_{y \notin X} p(y)} & \text{if} \quad x \notin X \\ 0 & \text{if} \quad x \in X. \end{cases}$$

7. $\succeq$ is a partial ordering of subsets of Θ and is defined for at least those pairs required by u.

8. $\omega : \bigcup_{j=1}^{r} Z^j \rightarrow \{over, continue\}$ is a mapping of terminal nodes within a stage game to one of two values: over, in which case the game ends, or continue, in which case the game continues to the next round. Clearly, we require $\omega(z) = over$ for all $z \in Z^r$. Note that ω is independent of the signals. Let $\omega_{over}^j = \{z \in Z^j \mid \omega(z) = over\}$ and $\omega_{cont}^j = \{z \in Z^j \mid \omega(z) = continue\}$.

9. $u = (u^1, \ldots, u^r)$, $u^j : \bigtimes_{k=1}^{j-1} \omega_{cont}^k \times \omega_{over}^j \times \bigtimes_{k=1}^{j} \Theta^{\kappa^k} \times \bigtimes_{i=1}^{n} \bigtimes_{k=1}^{j} \Theta^{\gamma^k} \rightarrow \mathbb{R}^n$ is a utility function such that for every j, $1 \leq j \leq r$, for every $i \in I$, and for every $\tilde{z} \in \bigtimes_{k=1}^{j-1} \omega_{cont}^k \times \omega_{over}^j$, at least one of the following two conditions holds:

(a) Utility is signal independent: $u_i^j(\tilde{z}, \vartheta) = u_i^j(\tilde{z}, \vartheta')$ for all legal $\vartheta, \vartheta' \in \bigtimes_{k=1}^{j} \Theta^{\kappa^k} \times \bigtimes_{i=1}^{n} \bigtimes_{k=1}^{j} \Theta^{\gamma^k}$.

(b) $\succeq$ is defined for all legal signals $(\tilde{\alpha}^j, \tilde{\beta}_i^j)$, $(\tilde{\alpha}^j, \tilde{\beta}'^j_i)$ through round j and a player's utility is increasing in her private signals, everything else equal:

$$\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) \succeq \left(\tilde{\alpha}^j, \tilde{\beta}'^j_i\right) \Longrightarrow$$
$$u_i\left(\tilde{z}, \tilde{\alpha}^j, \left(\tilde{\beta}_i^j, \tilde{\beta}_{-i}^j\right)\right) \geq u_i\left(\tilde{z}, \tilde{\alpha}^j, \left(\tilde{\beta}'^j_i, \tilde{\beta}_{-i}^j\right)\right).$$

We will use the term *game with ordered signals* and the term *ordered game* interchangeably.

2.1 Information filters

In this subsection, we define an *information filter* for ordered games. Instead of completely revealing a signal (either public or private) to a player, the signal first passes through this filter, which outputs a *coarsened* signal to the player. By varying the filter applied to a game, we are able to obtain a wide variety of games while keeping the underlying action space of the game intact. We will use this when designing our abstraction techniques. Formally, an information filter is as follows.

DEFINITION 2. *Let* $\Gamma = \langle I, G, L, \Theta, \kappa, \gamma, p, \succeq, \omega, u \rangle$ *be an ordered game. Let* $S^j \subseteq \bigtimes_{k=1}^{j} \Theta^{\kappa^k} \times \bigtimes_{k=1}^{j} \Theta^{\gamma^k}$ *be the set of legal signals (i.e., no repeated signals) for one player through round* j. *An* information filter *for* Γ *is a collection* $F = \langle F^1, \ldots, F^r \rangle$ *where each* F^j *is a function* $F^j : S^j \rightarrow 2^{S^j}$ *such that each of the following conditions hold:*

1. *(Truthfulness)* $(\tilde{\alpha}^j, \tilde{\beta}_i^j) \in F^j(\tilde{\alpha}^j, \tilde{\beta}_i^j)$ *for all legal* $(\tilde{\alpha}^j, \tilde{\beta}_i^j)$.
2. *(Independence) The range of* F^j *is a partition of* S^j.
3. *(Information preservation) If two values of a signal are distinguishable in round* k, *then they are distinguishable fpr each round* $j > k$. *Let* $m^j = \sum_{l=1}^{j} \kappa^l + \gamma^l$. *We require that for all legal* $(\theta_1, \ldots, \theta_{m^k}, \ldots, \theta_{m^j}) \subseteq \Theta$ *and* $(\theta'_1, \ldots, \theta'_{m^k}, \ldots, \theta'_{m^j}) \subseteq \Theta$:

$$(\theta'_1, \ldots, \theta'_{m^k}) \notin F^k(\theta_1, \ldots, \theta_{m^k}) \Longrightarrow$$
$$(\theta'_1, \ldots, \theta'_{m^k}, \ldots, \theta'_{m^j}) \notin F^j(\theta_1, \ldots, \theta_{m^k}, \ldots, \theta_{m^j}).$$

A game with ordered signals Γ and an information filter F for Γ defines a new game Γ_F. We refer to such games as *filtered ordered games*. We are left with the original game if we use the identity filter $F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) = \left\{\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right)\right\}$. We have the following simple (but important) result:

PROPOSITION 1. *A filtered ordered game is an extensive form game satisfying perfect recall.*

A simple proof proceeds by constructing an extensive form game directly from the ordered game, and showing that it satisfies perfect recall. In determining the payoffs in a game with filtered signals, we take the average over all real signals in the filtered class, weighted by the probability of each real signal occurring.

2.2 Strategies and Nash equilibrium

We are now ready to define behavior strategies in the context of filtered ordered games.

DEFINITION 3. *A behavior strategy for player* i *in round* j *of* $\Gamma = \langle I, G, L, \Theta, \kappa, \gamma, p, \succeq, \omega, u \rangle$ *with information filter* F *is a probability distribution over possible actions, and is defined for each player* i, *each round* j, *and each* $v \in V^j \setminus Z^j$ *for* $L^j(v) = i$:

$$\sigma_{i,v}^j : \bigtimes_{k=1}^{j-1} \omega_{cont}^k \times Range\left(F^j\right) \rightarrow \Delta\left\{w \in V^j \mid (v, w) \in E^j\right\}.$$

$(\Delta(X)$ *is the set of probability distributions over a finite set* X.) *A behavior strategy for player* i *in round* j *is* $\sigma_i^j = (\sigma_{i,v_1}^j, \ldots, \sigma_{i,v_m}^j)$ *for each* $v_k \in V^j \setminus Z^j$ *where* $L^j(v_k) = i$. *A behavior strategy for player* i *in* Γ *is* $\sigma_i = \left(\sigma_i^1, \ldots, \sigma_i^r\right)$. *A strategy profile is* $\sigma = (\sigma_1, \ldots, \sigma_n)$. *A strategy profile with* σ_i *replaced by* σ'_i *is* $(\sigma'_i, \sigma_{-i}) = (\sigma_1, \ldots, \sigma_{i-1}, \sigma'_i, \sigma_{i+1}, \ldots, \sigma_n)$.

By an abuse of notation, we will say player i receives an expected payoff of $u_i(\sigma)$ when all players are playing the strategy profile σ. Strategy σ_i is said to be player i's *best response* to σ_{-i} if for all other strategies σ'_i for player i we have $u_i(\sigma_i, \sigma_{-i}) \geq u_i(\sigma'_i, \sigma_{-i})$. σ is a *Nash equilibrium* if, for every player i, σ_i is a best response for σ_{-i}. A Nash equilibrium always exists in finite extensive form games [36], and one exists in behavior strategies for games with perfect recall [29]. Using these observations, we have the following corollary to Proposition 1:

COROLLARY 1. *For any filtered ordered game, a Nash equilibrium exists in behavior strategies.*

3. EQUILIBRIUM-PRESERVING ABSTRACTIONS

In this section, we present our main technique for reducing the size of games. We begin by defining a *filtered signal tree* which represents all of the chance moves in the game. The bold edges (*i.e.* the first two levels of the tree) in the game trees in Figure 1 correspond to the filtered signal trees in each game.

DEFINITION 4. *Associated with every ordered game $\Gamma = \langle I, G, L, \Theta, \kappa, \gamma, p, \succeq, \omega, u \rangle$ and information filter F is a filtered signal tree, a directed tree in which each node corresponds to some revealed (filtered) signals and edges correspond to revealing specific (filtered) signals. The nodes in the filtered signal tree represent the set of all possible revealed filtered signals (public and private) at some point in time. The filtered public signals revealed in round j correspond to the nodes in the κ^j levels beginning at level $\sum_{k=1}^{j-1} \left(\kappa^k + n\gamma^k \right)$ and the private signals revealed in round j correspond to the nodes in the $n\gamma^j$ levels beginning at level $\sum_{k=1}^{j} \kappa^k + \sum_{k=1}^{j-1} n\gamma^k$. We denote children of a node x as $N(x)$. In addition, we associate weights with the edges corresponding to the probability of the particular edge being chosen given that its parent was reached.*

In many games, there are certain situations in the game that can be thought of as being strategically equivalent to other situations in the game. By melding these situations together, it is possible to arrive at a strategically equivalent smaller game. The next two definitions formalize this notion via the introduction of the *ordered game isomorphic* relation and the *ordered game isomorphic abstraction transformation*.

DEFINITION 5. *Two subtrees beginning at internal nodes x and y of a filtered signal tree are* ordered game isomorphic *if x and y have the same parent and there is a bijection $f : N(x) \to N(y)$, such that for $w \in N(x)$ and $v \in N(y)$, $v = f(w)$ implies the weights on the edges (x, w) and (y, v) are the same and the subtrees beginning at w and v are ordered game isomorphic. Two leaves (corresponding to filtered signals ϑ and ϑ' up through round r) are ordered game isomorphic if for all $\tilde{z} \in \underset{j=1}{\overset{r-1}{\times}} \omega_{cont}^j \times \omega_{over}^r$,*

$$u^r\left(\tilde{z}, \vartheta\right) = u^r\left(\tilde{z}, \vartheta'\right).$$

DEFINITION 6. *Let $\Gamma = \langle I, G, L, \Theta, \kappa, \gamma, p, \succeq, \omega, u \rangle$ be an ordered game and let F be an information filter for Γ. Let ϑ and ϑ' be two nodes where the subtrees in the induced filtered signal tree corresponding to the nodes ϑ and ϑ' are ordered game isomorphic, and ϑ and ϑ' are at either level $\sum_{k=1}^{j-1} \left(\kappa^k + n\gamma^k \right)$ or $\sum_{k=1}^{j} \kappa^k + \sum_{k=1}^{j-1} n\gamma^k$ for some round j. The* ordered game isomorphic abstraction transformation *is given by creating a new information filter F':*

$$F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) = \begin{cases} F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) & \text{if } \left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) \notin \vartheta \cup \vartheta' \\ \vartheta \cup \vartheta' & \text{if } \left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) \in \vartheta \cup \vartheta'. \end{cases}$$

Figure 1 shows the ordered game isomorphic abstraction transformation applied twice to a tiny poker game. Theorem 2, our main equilibrium result, shows how the ordered game isomorphic abstraction transformation can be used to compute equilibria faster.

THEOREM 2. *Let $\Gamma = \langle I, G, L, \Theta, \kappa, \gamma, p, \succeq, \omega, u \rangle$ be an ordered game and F be an information filter for Γ. Let F' be an information filter constructed from F by one application of the ordered game isomorphic abstraction transformation. Let σ' be a Nash equilibrium of the induced game $\Gamma_{F'}$. If we take $\sigma_{i,v}^j\left(\tilde{z}, F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right)\right) = \sigma_{i,v}'^j\left(\tilde{z}, F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right)\right)$, σ is a Nash equilibrium of Γ_F.*

PROOF. For an extensive form game, a *belief system* μ assigns a probability to every decision node x such that $\sum_{x \in h} \mu(x) = 1$ for every information set h. A strategy profile σ is *sequentially rational at h given belief system μ* if

$$u_i(\sigma_i, \sigma_{-i} \mid h, \mu) \geq u_i(\tau_i, \sigma_{-i} \mid h, \mu)$$

for all other strategies τ_i, where i is the player who controls h. A basic result [33, Proposition 9.C.1] characterizing Nash equilibria dictates that σ is a Nash equilibrium if and only if there is a belief system μ such that for every information set h with $\Pr(h \mid \sigma) > 0$, the following two conditions hold: (C1) σ is sequentially rational at h given μ; and (C2) $\mu(x) = \frac{\Pr(x \mid \sigma)}{\Pr(h \mid \sigma)}$ for all $x \in h$. Since σ' is a Nash equilibrium of Γ', there exists such a belief system μ' for $\Gamma_{F'}$. Using μ', we will construct a belief system μ for Γ and show that conditions C1 and C2 hold, thus supporting σ as a Nash equilibrium.

Fix some player $i \in I$. Each of i's information sets in some round j corresponds to filtered signals $F^j\left(\tilde{\alpha}^{*j}, \tilde{\beta}_i^{*j}\right)$, history in the first $j - 1$ rounds $(z_1, \ldots, z_{j-1}) \in \underset{k=1}{\overset{j-1}{\times}} \omega_{cont}^k$, and history so far in round j, $v \in V^j \setminus Z^j$. Let $\tilde{z} = (z_1, \ldots, z_{j-1}, v)$ represent all of the player actions leading to this information set. Thus, we can uniquely specify this information set using the information $\left(F^j\left(\tilde{\alpha}^{*j}, \tilde{\beta}_i^{*j}\right), \tilde{z}\right)$.

Each node in an information set corresponds to the possible private signals the other players have received. Denote by $\tilde{\beta}$ some legal

$$(F^j(\tilde{\alpha}^j, \tilde{\beta}_1^j), \ldots, F^j(\tilde{\alpha}^j, \tilde{\beta}_{i-1}^j), F^j(\tilde{\alpha}^j, \tilde{\beta}_{i+1}^j), \ldots, F^j(\tilde{\alpha}^j, \tilde{\beta}_n^j)).$$

In other words, there exists $(\tilde{\alpha}^j, \tilde{\beta}_1^j, \ldots, \tilde{\beta}_n^j)$ such that

$$(\tilde{\alpha}^j, \tilde{\beta}_i^j) \in F^j(\tilde{\alpha}^{*j}, \tilde{\beta}_i^{*j}), (\tilde{\alpha}^j, \tilde{\beta}_k^j) \in F^j(\tilde{\alpha}^j, \tilde{\beta}_k^j)$$

for $k \neq i$, and no signals are repeated. Using such a set of signals $(\tilde{\alpha}^j, \tilde{\beta}_1^j, \ldots, \tilde{\beta}_n^j)$, let $\hat{\beta}'$ denote $(F'^j(\tilde{\alpha}^j, \tilde{\beta}_1^j), \ldots, F'^j(\tilde{\alpha}^j, \tilde{\beta}_{i-1}^j), F'^j(\tilde{\alpha}^j, \tilde{\beta}_{i+1}^j), \ldots, F'^j(\tilde{\alpha}^j, \tilde{\beta}_n^j))$. (We will abuse notation and write $F_{-i}'^j\left(\hat{\beta}\right) = \hat{\beta}'$.) We can now compute μ directly from μ':

$$\mu\left(\hat{\beta} \mid F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right), \tilde{z}\right) =$$
$$\begin{cases} \mu'\left(\hat{\beta}' \mid F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right), \tilde{z}\right) \\ \quad \text{if } F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) \neq F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) \text{ or } \hat{\beta} = \hat{\beta}' \\ p^* \mu'\left(\hat{\beta}' \mid F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right), \tilde{z}\right) \\ \quad \text{if } F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) = F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right) \text{ and } \hat{\beta} \neq \hat{\beta}' \end{cases}$$
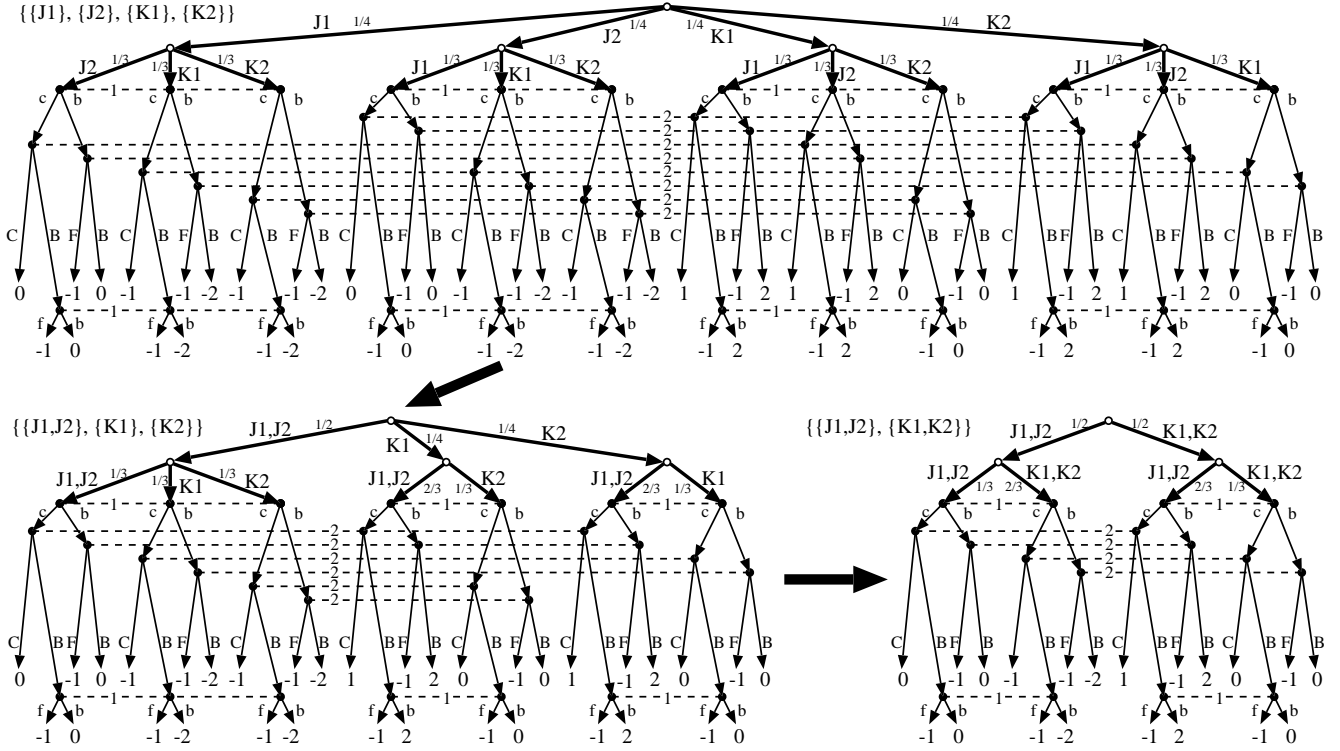
Figure 1: *GameShrink* applied to a tiny two-person four-card (two Jacks and two Kings) poker game. Next to each game tree is the range of the information filter F. Dotted lines denote information sets, which are labeled by the controlling player. Open circles are chance nodes with the indicated transition probabilities. The root node is the chance node for player 1's card, and the next level is for player 2's card. The payment from player 2 to player 1 is given below each leaf. In this example, the algorithm reduces the game tree from 53 nodes to 19 nodes.

where $p^* = \frac{\Pr\left(\hat{\beta} \mid F^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right)\right)}{\Pr\left(\hat{\beta}' \mid F'^j\left(\tilde{\alpha}^j, \tilde{\beta}_i^j\right)\right)}$. The following three claims show that μ as calculated above supports σ as a Nash equilibrium.

CLAIM 1. *μ is a valid belief system for Γ_F.*

CLAIM 2. *For all information sets h with $\Pr(h \mid \sigma) > 0$, $\mu(x) = \frac{\Pr(x \mid \sigma)}{\Pr(h \mid \sigma)}$ for all $x \in h$.*

CLAIM 3. *For all information sets h with $\Pr(h \mid \sigma) > 0$, σ is sequentially rational at h given μ.*

The proofs of Claims 1-3 are in an extended version of this paper [13]. By Claims 1 and 2, we know that condition C2 holds. By Claim 3, we know that condition C1 holds. Thus, σ is a Nash equilibrium. □

3.1 Nontriviality of generalizing beyond this model

Our model does not capture general sequential games of imperfect information because it is restricted in two ways (as discussed above): 1) there is a special structure connecting the player actions and the chance actions (for one, the players are assumed to observe each others' actions, but nature's actions might not be publicly observable), and 2) there is a common ordering of signals. In this subsection we show that removing either of these conditions can make our technique invalid.

First, we demonstrate a failure when removing the first assumption. Consider the game in Figure 2.[6] Nodes a and b are in the same information set, have the same parent (chance) node, have isomorphic subtrees with the same payoffs, and nodes c and d also have similar structural properties. By merging the subtrees beginning at a and b, we get the game on the right in Figure 2. In this game, player 1's only Nash equilibrium strategy is to play left. But in the original game, player 1 knows that node c will never be reached, and so should play right in that information set.


Figure 2: **Example illustrating difficulty in developing a theory of equilibrium-preserving abstractions for general extensive form games.**

Removing the second assumption (that the utility functions are based on a common ordering of signals) can also cause failure. Consider a simple three-card game with a deck containing two Jacks (J1 and J2) and a King (K), where player 1's utility function is based on the ordering

---
[6]We thank Albert Xin Jiang for providing this example.

K $\succeq$ J1 $\sim$ J2 but player 2's utility function is based on the ordering J2 $\succeq$ K $\succeq$ J1. It is easy to check that in the abstracted game (where Player 1 treats J1 and J2 as being "equivalent") the Nash equilibrium does not correspond to a Nash equilibrium in the original game.[7]

4. GAMESHRINK: AN EFFICIENT ALGORITHM FOR COMPUTING ORDERED GAME ISOMORPHIC ABSTRACTION TRANSFORMATIONS

This section presents an algorithm, *GameShrink*, for conducting the abstractions. It only needs to analyze the signal tree discussed above, rather than the entire game tree.

We first present a subroutine that *GameShrink* uses. It is a dynamic program for computing the ordered game isomorphic relation. Again, it operates on the signal tree.

ALGORITHM 1. *OrderedGameIsomorphic?* $(\Gamma, \vartheta, \vartheta')$

1. If ϑ and ϑ' have different parents, then return false.

2. If ϑ and ϑ' are both leaves of the signal tree:

    (a) If $u^r(\vartheta \mid \tilde{z}) = u^r(\vartheta' \mid \tilde{z})$ for all $\tilde{z} \in \underset{j=1}{\overset{r-1}{\times}} \omega^j_{cont} \times \omega^r_{over}$, then return true.

    (b) Otherwise, return false.

3. Create a bipartite graph $G_{\vartheta,\vartheta'} = (V_1, V_2, E)$ with $V_1 = N(\vartheta)$ and $V_2 = N(\vartheta')$.

4. For each $v_1 \in V_1$ and $v_2 \in V_2$:

    If OrderedGameIsomorphic? (Γ, v_1, v_2)
        Create edge (v_1, v_2)

5. Return true if $G_{\vartheta,\vartheta'}$ has a perfect matching; otherwise, return false.

By evaluating this dynamic program from bottom to top, Algorithm 1 determines, in time polynomial in the size of the signal tree, whether or not any pair of equal depth nodes x and y are ordered game isomorphic. We can further speed up this computation by only examining nodes with the same parent, since we know (from step 1) that no nodes with different parents are ordered game isomorphic. The test in step 2(a) can be computed in $O(1)$ time by consulting the $\succeq$ relation from the specification of the game. Each call to *OrderedGameIsomorphic?* performs at most one perfect matching computation on a bipartite graph with $O(|\Theta|)$ nodes and $O(|\Theta|^2)$ edges (recall that Θ is the set of signals). Using the Ford-Fulkerson algorithm [12] for finding a maximal matching, this takes $O(|\Theta|^3)$ time. Let S be the maximum number of signals possibly revealed in the game (*e.g.*, in Rhode Island Hold'em, $S = 4$ because each of the two players has one card in the hand plus there are two cards on the table). The number of nodes, n, in the signal tree is $O(|\Theta|^S)$. The dynamic program visits each node in the signal tree, with each visit requiring $O(|\Theta|^2)$ calls to the *OrderedGameIsomorphic?* routine. So, it takes $O(|\Theta|^S|\Theta|^3|\Theta|^2) = O(|\Theta|^{S+5})$ time to compute the entire ordered game isomorphic relation.

While this is exponential in the number of revealed signals, we now show that it is polynomial in the size of the signal tree—and thus polynomial in the size of the game tree

---

[7]We thank an anonymous person for this example.

---

because the signal tree is smaller than the game tree. The number of nodes in the signal tree is

$$n = 1 + \sum_{i=1}^{S} \prod_{j=1}^{i} (|\Theta| - j + 1)$$

(Each term in the summation corresponds to the number of nodes at a specific depth of the tree.) The number of leaves is

$$\prod_{j=1}^{S} (|\Theta| - j + 1) = \binom{|\Theta|}{S} S!$$

which is a lower bound on the number of nodes. For large $|\Theta|$ we can use the relation $\binom{n}{k} \sim \frac{n^k}{k!}$ to get

$$\binom{|\Theta|}{S} S! \sim \left(\frac{|\Theta|^S}{S!}\right) S! = |\Theta|^S$$

and thus the number of leaves in the signal tree is $\Omega(|\Theta|^S)$. Thus, $O(|\Theta|^{S+5}) = O(n|\Theta|^5)$, which proves that we can indeed compute the ordered game isomorphic relation in time polynomial in the number of nodes, n, of the signal tree.

The algorithm often runs in *sublinear* time (and space) in the size of the game tree because the signal tree is significantly smaller than the game tree in most nontrivial games. (Note that the input to the algorithm is not an explicit game tree, but a specification of the rules, so the algorithm does not need to read in the game tree.) See Figure 1. In general, if an ordered game has r rounds, and each round's stage game has at least b nonterminal leaves, then the size of the signal tree is at most $\frac{1}{b^r}$ of the size of the game tree. For example, in Rhode Island Hold'em, the game tree has 3.1 billion nodes while the signal tree only has 6,632,705.

Given the *OrderedGameIsomorphic?* routine for determining ordered game isomorphisms in an ordered game, we are ready to present the main algorithm, *GameShrink*.

ALGORITHM 2. *GameShrink* (Γ)

1. Initialize F to be the identity filter for Γ.

2. For j from 1 to r:

    For each pair of sibling nodes ϑ, ϑ' at either level $\sum_{k=1}^{j-1} (\kappa^k + n\gamma^k)$ or $\sum_{k=1}^{j} \kappa^k + \sum_{k=1}^{j-1} n\gamma^k$ in the filtered (according to F) signal tree:

        If OrderedGameIsomorphic?$(\Gamma, \vartheta, \vartheta')$, then $F^j(\vartheta) \leftarrow F^j(\vartheta') \leftarrow F^j(\vartheta) \cup F^j(\vartheta')$.

3. Output F.

Given as input an ordered game Γ, *GameShrink* applies the shrinking ideas presented above as aggressively as possible. Once it finishes, there are no contractible nodes (since it compares every pair of nodes at each level of the signal tree), and it outputs the corresponding information filter F. The correctness of *GameShrink* follows by a repeated application of Theorem 2. Thus, we have the following result:

THEOREM 3. GameShrink *finds all ordered game isomorphisms and applies the associated ordered game isomorphic abstraction transformations. Furthermore, for any Nash equilibrium, σ', of the abstracted game, the strategy profile constructed for the original game from σ' is a Nash equilibrium.*

The dominating factor in the run time of *GameShrink* is in the r^{th} iteration of the main for-loop. There are at most

$\binom{|\Theta|}{S}S!$ nodes at this level, where we again take S to be the maximum number of signals possibly revealed in the game. Thus, the inner for-loop executes $O\left(\left(\binom{|\Theta|}{S}S!\right)^2\right)$ times. As discussed in the next subsection, we use a union-find data structure to represent the information filter F. Each iteration of the inner for-loop possibly performs a union operation on the data structure; performing M operations on a union-find data structure containing N elements takes $O(\alpha(M,N))$ amortized time per operation, where $\alpha(M,N)$ is the inverse Ackermann's function [1, 49] (which grows extremely slowly). Thus, the total time for *GameShrink* is $O\left(\left(\binom{|\Theta|}{S}S!\right)^2 \alpha\left(\left(\binom{|\Theta|}{S}S!\right)^2, |\Theta|^S\right)\right)$. By the inequality $\binom{n}{k} \leq \frac{n^k}{k!}$, this is $O\left((|\Theta|^S)^2 \alpha\left((|\Theta|^S)^2, |\Theta|^S\right)\right)$. Again, although this is exponential in S, it is $\tilde{O}(n^2)$, where n is the number of nodes in the signal tree. Furthermore, *GameShrink* tends to actually run in *sublinear* time and space in the size of the game tree because the signal tree is significantly smaller than the game tree in most nontrivial games, as discussed above.

4.1 Efficiency enhancements

We designed several speed enhancement techniques for *GameShrink*, and all of them are incorporated into our implementation. One technique is the use of the union-find data structure for storing the information filter F. This data structure uses time almost linear in the number of operations [49]. Initially each node in the signalling tree is its own set (this corresponds to the identity information filter); when two nodes are contracted they are joined into a new set. Upon termination, the filtered signals for the abstracted game correspond exactly to the disjoint sets in the data structure. This is an efficient method of recording contractions within the game tree, and the memory requirements are only linear in the size of the signal tree.

Determining whether two nodes are ordered game isomorphic requires us to determine if a bipartite graph has a perfect matching. We can eliminate some of these computations by using easy-to-check necessary conditions for the ordered game isomorphic relation to hold. One such condition is to check that the nodes have the same number of chances as being ranked (according to $\succeq$) higher than, lower than, and the same as the opponents. We can precompute these frequencies for every game tree node. This substantially speeds up *GameShrink*, and we can leverage this database across multiple runs of the algorithm (for example, when trying different abstraction levels; see next section). The indices for this database depend on the private and public signals, but not the *order* in which they were revealed, and thus two nodes may have the same corresponding database entry. This makes the database significantly more compact. (For example in Texas Hold'em, the database is reduced by a factor $\binom{50}{3}\binom{47}{1}\binom{46}{1}/\binom{50}{5} = 20$.) We store the histograms in a 2-dimensional database. The first dimension is indexed by the private signals, the second by the public signals. The problem of computing the index in (either) one of the dimensions is exactly the problem of computing a bijection between all subsets of size r from a set of size n and integers in $\left[0,\ldots,\binom{n}{r}-1\right]$. We efficiently compute this using the subsets' *colexicographical ordering* [6]. Let $\{c_1,\ldots,c_r\}$, $c_i \in \{0,\ldots,n-1\}$, denote the r signals and assume that

$c_i < c_{i+1}$. We compute a unique index for this set of signals as follows: $index(c_1,\ldots,c_r) = \sum_{i=1}^{r} \binom{c_i}{i}$.

5. APPROXIMATION METHODS

Some games are too large to compute an exact equilibrium, even after using the presented abstraction technique. This section discusses general techniques for computing approximately optimal strategy profiles. For a two-player game, we can always evaluate the worst-case performance of a strategy, thus providing some objective evaluation of the strength of the strategy. To illustrate this, suppose we know player 2's planned strategy for some game. We can then fix the probabilities of player 2's actions in the game tree as if they were chance moves. Then player 1 is faced with a single-agent decision problem, which can be solved bottom-up, maximizing expected payoff at every node. Thus, we can objectively determine the expected worst-case performance of player 2's strategy. This will be most useful when we want to evaluate how well a given strategy performs when we know that it is not an equilibrium strategy. (A variation of this technique may also be applied in n-person games where only one player's strategies are held fixed.) This technique provides *ex post* guarantees about the worst-case performance of a strategy, and can be used independently of the method that is used to compute the strategies.

5.1 State-space approximations

By slightly modifying *GameShrink*, we can obtain an algorithm that yields even smaller game trees, at the expense of losing the equilibrium guarantees of Theorem 2. Instead of requiring the payoffs at terminal nodes to match exactly, we can instead compute a penalty that increases as the difference in utility between two nodes increases.

There are many ways in which the penalty function could be defined and implemented. One possibility is to create edge weights in the bipartite graphs used in Algorithm 1, and then instead of requiring perfect matchings in the unweighted graph we would instead require perfect matchings with low cost (*i.e.*, only consider two nodes to be ordered game isomorphic if the corresponding bipartite graph has a perfect matching with cost below some threshold). Thus, with this threshold as a parameter, we have a knob to turn that in one extreme (threshold $= 0$) yields an optimal abstraction and in the other extreme (threshold $= \infty$) yields a highly abstracted game (this would in effect restrict players to ignoring all signals, but still observing actions). This knob also begets an *anytime* algorithm. One can solve increasingly less abstracted versions of the game, and evaluate the quality of the solution at every iteration using the *ex post* method discussed above.

5.2 Algorithmic approximations

In the case of two-player zero-sum games, the equilibrium computation can be modeled as a linear program (LP), which can in turn be solved using the simplex method. This approach has inherent features which we can leverage into desirable properties in the context of solving games.

In the LP, primal solutions correspond to strategies of player 2, and dual solutions correspond to strategies of player 1. There are two versions of the simplex method: the primal simplex and the dual simplex. The primal simplex maintains primal feasibility and proceeds by finding better and better primal solutions until the dual solution vector is feasible,

at which point optimality has been reached. Analogously, the dual simplex maintains dual feasibility and proceeds by finding increasingly better dual solutions until the primal solution vector is feasible. (The dual simplex method can be thought of as running the primal simplex method on the dual problem.) Thus, the primal and dual simplex methods serve as *anytime* algorithms (for a given abstraction) for players 2 and 1, respectively. At any point in time, they can output the best strategies found so far.

Also, for any feasible solution to the LP, we can get bounds on the quality of the strategies by examining the primal and dual solutions. (When using the primal simplex method, dual solutions may be read off of the LP tableau.) Every feasible solution of the dual yields an upper bound on the optimal value of the primal, and vice versa [9, p. 57]. Thus, without requiring further computation, we get lower bounds on the expected utility of each agent's strategy against that agent's worst-case opponent.

One problem with the simplex method is that it is not a primal-dual algorithm, that is, it does not maintain both primal and dual feasibility throughout its execution. (In fact, it only obtains primal and dual feasibility at the very end of execution.) In contrast, there are interior-point methods for linear programming that maintain primal and dual feasibility throughout the execution. For example, many interior-point path-following algorithms have this property [55, Ch. 5]. We observe that running such a linear programming method yields a method for finding ϵ-equilibria (*i.e.*, strategy profiles in which no agent can increase her expected utility by more than ϵ by deviating). A threshold on ϵ can also be used as a termination criterion for using the method as an anytime algorithm. Furthermore, interior-point methods in this class have polynomial-time worst-case run time, as opposed to the simplex algorithm, which takes exponentially many steps in the worst case.

6. RELATED RESEARCH

Functions that transform extensive form games have been introduced [50, 11]. In contrast to our work, those approaches were not for making the game smaller and easier to solve. The main result is that a game can be derived from another by a sequence of those transformations if and only if the games have the same *pure reduced normal form*. The pure reduced normal form is the extensive form game represented as a game in normal form where duplicates of pure strategies (i.e., ones with identical payoffs) are removed and players essentially select equivalence classes of strategies [27]. An extension to that work shows a similar result, but for slightly different transformations and *mixed reduced normal form* games [21]. Modern treatments of this prior work on game transformations exist [38, Ch. 6], [10].

The recent notion of *weak isomorphism* in extensive form games [7] is related to our notion of restricted game isomorphism. The motivation of that work was to justify solution concepts by arguing that they are invariant with respect to isomorphic transformations. Indeed, the author shows, among other things, that many solution concepts, including Nash, perfect, subgame perfect, and sequential equilibrium, are invariant with respect to weak isomorphisms. However, that definition requires that the games to be tested for weak isomorphism are of the same size. Our focus is totally different: we find strategically equivalent *smaller* games. Also, their paper does not provide algorithms.

Abstraction techniques have been used in artificial intelligence research before. In contrast to our work, most (but not all) research involving abstraction has been for single-agent problems (*e.g.* [20, 32]). Furthermore, the use of abstraction typically leads to sub-optimal solutions, unlike the techniques presented in this paper, which yield optimal solutions. A notable exception is the use of abstraction to compute optimal strategies for the game of Sprouts [2]. However, a significant difference to our work is that Sprouts is a game of perfect information.

One of the first pieces of research to use abstraction in multi-agent settings was the development of *partition search*, which is the algorithm behind *GIB*, the world's first expert-level computer bridge player [17, 18]. In contrast to other game tree search algorithms which store a particular game position at each node of the search tree, partition search stores *groups* of positions that are similar. (Typically, the similarity of two game positions is computed by ignoring the less important components of each game position and then checking whether the abstracted positions are similar—in some domain-specific expert-defined sense—to each other.) Partition search can lead to substantial speed improvements over α-β-search. However, it is not game theory-based (it does not consider information sets in the game tree), and thus does not solve for the equilibrium of a game of imperfect information, such as poker.[8] Another difference is that the abstraction is defined by an expert human while our abstractions are determined automatically.

There has been some research on the use of abstraction for imperfect information games. Most notably, Billings *et al* [4] describe a manually constructed abstraction for Texas Hold'em poker, and include promising results against expert players. However, this approach has significant drawbacks. First, it is highly specialized for Texas Hold'em. Second, a large amount of expert knowledge and effort was used in constructing the abstraction. Third, the abstraction does not preserve equilibrium: even if applied to a smaller game, it might not yield a game-theoretic equilibrium. Promising ideas for abstraction in the context of general extensive form games have been described in an extended abstract [39], but to our knowledge, have not been fully developed.

7. CONCLUSIONS AND DISCUSSION

We introduced the ordered game isomorphic abstraction transformation and gave an algorithm, *GameShrink*, for abstracting the game using the isomorphism exhaustively. We proved that in games with ordered signals, any Nash equilibrium in the smaller abstracted game maps directly to a Nash equilibrium in the original game.

The complexity of *GameShrink* is $\tilde{O}(n^2)$, where n is the number of nodes in the signal tree. It is no larger than the game tree, and on nontrivial games it is drastically smaller, so *GameShrink* has time and space complexity *sublinear* in

---

[8]Bridge is also a game of imperfect information, and partition search does not find the equilibrium for that game either. Instead, partition search is used in conjunction with statistical sampling to simulate the uncertainty in bridge. There are also other bridge programs that use search techniques for perfect information games in conjunction with statistical sampling and expert-defined abstraction [48]. Such (non-game-theoretic) techniques are unlikely to be competitive in poker because of the greater importance of information hiding and bluffing.

the size of the game tree. Using *GameShrink*, we found a minimax equilibrium to Rhode Island Hold'em, a poker game with 3.1 billion nodes in the game tree—over four orders of magnitude more than in the largest poker game solved previously.

To further improve scalability, we introduced an approximation variant of *GameShrink*, which can be used as an anytime algorithm by varying a parameter that controls the coarseness of abstraction. We also discussed how (in a two-player zero-sum game), linear programming can be used in an anytime manner to generate approximately optimal strategies of increasing quality. The method also yields bounds on the suboptimality of the resulting strategies. We are currently working on using these techniques for full-scale 2-player limit Texas Hold'em poker, a highly popular card game whose game tree has about 10^{18} nodes. That game tree size has required us to use the approximation version of *GameShrink* (as well as round-based abstraction) [16, 15].

8. REFERENCES

[1] W. Ackermann. Zum Hilbertschen Aufbau der reellen Zahlen. *Math. Annalen*, 99:118–133, 1928.
[2] D. Applegate, G. Jacobson, and D. Sleator. Computer analysis of sprouts. Technical Report CMU-CS-91-144, 1991.
[3] R. Bellman and D. Blackwell. Some two-person games involving bluffing. *PNAS*, 35:600–605, 1949.
[4] D. Billings, N. Burch, A. Davidson, R. Holte, J. Schaeffer, T. Schauenberg, and D. Szafron. Approximating game-theoretic optimal strategies for full-scale poker. In *IJCAI*, 2003.
[5] D. Billings, A. Davidson, J. Schaeffer, and D. Szafron. The challenge of poker. *Artificial Intelligence*, 134:201–240, 2002.
[6] B. Bollobás. *Combinatorics*. Cambridge University Press, 1986.
[7] A. Casajus. Weak isomorphism of extensive games. *Mathematical Social Sciences*, 46:267–290, 2003.
[8] X. Chen and X. Deng. Settling the complexity of 2-player Nash equilibrium. *ECCC*, Report No. 150, 2005.
[9] V. Chvátal. *Linear Programming*. W. H. Freeman & Co., 1983.
[10] B. P. de Bruin. Game transformations and game equivalence. Technical note x-1999-01, University of Amsterdam, Institute for Logic, Language, and Computation, 1999.
[11] S. Elmes and P. J. Reny. On the strategic equivalence of extensive form games. *J. of Economic Theory*, 62:1–23, 1994.
[12] L. R. Ford, Jr. and D. R. Fulkerson. *Flows in Networks*. Princeton University Press, 1962.
[13] A. Gilpin and T. Sandholm. Finding equilibria in large sequential games of imperfect information. Technical Report CMU-CS-05-158, Carnegie Mellon University, 2005.
[14] A. Gilpin and T. Sandholm. Optimal Rhode Island Hold'em poker. In *AAAI*, pages 1684–1685, Pittsburgh, PA, USA, 2005.
[15] A. Gilpin and T. Sandholm. A competitive Texas Hold'em poker player via automated abstraction and real-time equilibrium computation. Mimeo, 2006.
[16] A. Gilpin and T. Sandholm. A Texas Hold'em poker player based on automated abstraction and real-time equilibrium computation. In *AAMAS*, Hakodate, Japan, 2006.
[17] M. L. Ginsberg. Partition search. In *AAAI*, pages 228–233, Portland, OR, 1996.
[18] M. L. Ginsberg. GIB: Steps toward an expert-level bridge-playing program. In *IJCAI*, Stockholm, Sweden, 1999.
[19] S. Govindan and R. Wilson. A global Newton method to compute Nash equilibria. *J. of Econ. Theory*, 110:65–86, 2003.
[20] C. A. Knoblock. Automatically generating abstractions for planning. *Artificial Intelligence*, 68(2):243–302, 1994.
[21] E. Kohlberg and J.-F. Mertens. On the strategic stability of equilibria. *Econometrica*, 54:1003–1037, 1986.
[22] D. Koller and N. Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and Economic Behavior*, 4(4):528–552, Oct. 1992.
[23] D. Koller and N. Megiddo. Finding mixed strategies with small supports in extensive form games. *International Journal of Game Theory*, 25:73–92, 1996.
[24] D. Koller, N. Megiddo, and B. von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2):247–259, 1996.
[25] D. Koller and A. Pfeffer. Representations and solutions for game-theoretic problems. *Artificial Intelligence*, 94(1):167–215, July 1997.
[26] D. M. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50(4):863–894, 1982.
[27] H. W. Kuhn. Extensive games. *PNAS*, 36:570–576, 1950.
[28] H. W. Kuhn. A simplified two-person poker. In *Contributions to the Theory of Games*, volume 1 of *Annals of Mathematics Studies, 24*, pages 97–103. Princeton University Press, 1950.
[29] H. W. Kuhn. Extensive games and the problem of information. In *Contributions to the Theory of Games*, volume 2 of *Annals of Mathematics Studies, 28*, pages 193–216. Princeton University Press, 1953.
[30] C. Lemke and J. Howson. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12:413–423, 1964.
[31] R. Lipton, E. Markakis, and A. Mehta. Playing large games using simple strategies. In *ACM-EC*, pages 36–41, 2003.
[32] C.-L. Liu and M. Wellman. On state-space abstraction for anytime evaluation of Bayesian networks. *SIGART Bulletin*, 7(2):50–57, 1996.
[33] A. Mas-Colell, M. Whinston, and J. R. Green. *Microeconomic Theory*. Oxford University Press, 1995.
[34] R. D. McKelvey and A. McLennan. Computation of equilibria in finite games. In *Handbook of Computational Economics*, volume 1, pages 87–142. Elsevier, 1996.
[35] P. B. Miltersen and T. B. Sørensen. Computing sequential equilibria for two-player games. In *SODA*, pages 107–116, 2006.
[36] J. Nash. Equilibrium points in n-person games. *Proc. of the National Academy of Sciences*, 36:48–49, 1950.
[37] J. F. Nash and L. S. Shapley. A simple three-person poker game. In *Contributions to the Theory of Games*, volume 1, pages 105–116. Princeton University Press, 1950.
[38] A. Perea. *Rationality in extensive form games*. Kluwer Academic Publishers, 2001.
[39] A. Pfeffer, D. Koller, and K. Takusagawa. State-space approximations for extensive form games, July 2000. Talk given at the First International Congress of the Game Theory Society, Bilbao, Spain.
[40] R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a Nash equilibrium. In *AAAI*, pages 664–669, San Jose, CA, USA, 2004.
[41] I. Romanovskii. Reduction of a game with complete memory to a matrix game. *Soviet Mathematics*, 3:678–681, 1962.
[42] T. Sandholm and A. Gilpin. Sequences of take-it-or-leave-it offers: Near-optimal auctions without full valuation revelation. In *AAMAS*, Hakodate, Japan, 2006.
[43] T. Sandholm, A. Gilpin, and V. Conitzer. Mixed-integer programming methods for finding Nash equilibria. In *AAAI*, pages 495–501, Pittsburgh, PA, USA, 2005.
[44] R. Savani and B. von Stengel. Exponentially many steps for finding a Nash equilibrium in a bimatrix game. In *FOCS*, pages 258–267, 2004.
[45] R. Selten. Spieltheoretische behandlung eines oligopolmodells mit nachfrageträgheit. *Zeitschrift für die gesamte Staatswissenschaft*, 12:301–324, 1965.
[46] R. Selten. Evolutionary stability in extensive two-person games – correction and further development. *Mathematical Social Sciences*, 16:223–266, 1988.
[47] J. Shi and M. Littman. Abstraction methods for game theoretic poker. In *Computers and Games*, pages 333–345. Springer-Verlag, 2001.
[48] S. J. J. Smith, D. S. Nau, and T. Throop. Computer bridge: A big win for AI planning. *AI Magazine*, 19(2):93–105, 1998.
[49] R. E. Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of the ACM*, 22(2):215–225, 1975.
[50] F. Thompson. Equivalence of games in extensive form. RAND Memo RM-759, The RAND Corporation, Jan. 1952.
[51] J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1947.
[52] B. von Stengel. Efficient computation of behavior strategies. *Games and Economic Behavior*, 14(2):220–246, 1996.
[53] B. von Stengel. Computing equilibria for two-person games. In *Handbook of Game Theory*, volume 3. North Holland, Amsterdam, 2002.
[54] R. Wilson. Computing equilibria of two-person games from the extensive form. *Management Science*, 18(7):448–460, 1972.
[55] S. J. Wright. *Primal-Dual Interior-Point Methods*. SIAM, 1997.