

Optimizing Propositional Networks

Chiara F. Sironi and Mark H. M. Winands

Department of Data Science and Knowledge Engineering
Maastricht University, The Netherlands
{c.sironi, m.winands}@maastrichtuniversity.nl

Abstract

General Game Playing (GGP) programs need a Game Description Language (GDL) reasoner to be able to interpret the game rules and search for the best actions to play in the game. One method for interpreting the game rules consists of translating the GDL game description into an alternative representation that the player can use to reason more efficiently on the game. The Propositional Network (PropNet) is an example of such method. The use of PropNets in GGP has become popular due to the fact that PropNets can speed up the reasoning process by several orders of magnitude compared to custom-made or Prolog-based GDL reasoners, improving the quality of the search for the best actions. This paper analyses the performance of a PropNet-based reasoner and evaluates four different optimizations for the PropNet structure that can help further increase its reasoning speed in terms of visited game states per second.

1 Introduction

The aim of General Game Playing (GGP) is to develop programs that are able to play any arbitrary game at an expert level by being only given its rules. These programs must devise a playing strategy without having any prior knowledge about the game. Moreover, the rules are given to the player just before game playing starts and usually for each game step only few seconds are available to choose a move. Thus, the player has to learn an appropriate playing strategy on-line and in a limited amount of time.

To be able to play games, a GGP program has two main components: a way to interpret the game rules, written in the Game Description Language (GDL), and a strategy to choose which actions to play.

Regarding the first component, many different approaches have been proposed to parse the game rules. Three main methods to interpret GDL can be identified: (1) Prolog-based interpreters, that translate the game rules from GDL into Prolog and then use a Prolog engine to reason on them, (2) custom-made interpreters written for the sole purpose of interpreting GDL rules, and (3) reasoners that translate the GDL description into an alternative representation that the

player can use to efficiently reason on the game. A description and performance evaluation of available GDL reasoners is given in [Schiffel and Björnsson, 2014].

Regarding the second component, most of the approaches that proved successful in addressing the challenges of GGP are based on Monte-Carlo simulation techniques and especially on Monte-Carlo Tree Search (MCTS) [Coulom, 2007; Björnsson and Finnsson, 2009]. For Monte-Carlo methods the choice of the best action to play is based on game statistics collected by sampling the state space of the game. The number of samples that Monte-Carlo methods can collect directly influences their performance. A higher number of samples will improve the quality of the chosen actions.

A faster GDL reasoner, that in a given amount of time can analyse a higher number of game states than other reasoners, can positively influence Monte-Carlo based search. Propositional Networks (PropNets) [Schkufza *et al.*, 2008; Cox *et al.*, 2009] have become popular in GGP because they can speed up the reasoning process by several orders of magnitude compared to custom-made or Prolog-based GDL reasoners. Nowadays, all the best general game players use a PropNet-based reasoner [Schreiber, 2013; Darper and Rose, 2014; Emsile, 2015].

The purpose of this paper is to analyse the performance of the implementation of the PropNet-based reasoner provided in the GGP-Base framework [Schreiber, 2013], discuss four optimizations of the structure of the PropNet and empirically evaluate their impact on the speed of the reasoning process. The performance of the custom-made GDL reasoner provided in the GGP-Base framework, called GGP-Base Prover, has been used as a reference.

The rest of the paper is structured as follows. Section 2 gives a short introduction to GDL and PropNets. Sections 3 and 4 give some details about the PropNet implementation and a description of the PropNet optimizations respectively. Section 5 presents the empirical evaluation of the PropNet and Section 6 concludes and indicates potential future work.

2 Background

2.1 The Game Description Language

The Game Description Language (GDL) is a first order logic language used in GGP to represent the rules of games [Love *et al.*, 2008]. In GDL a game state is defined by specifying

```

(role player)
(light p) (light q)
(<= (legal player (turnOn ?x)) (not (true (on ?x))) (light ?x))
(<= (next (on ?x)) (does player (turnOn ?x)))
(<= (next (on ?x)) (true (on ?x)))
(<= terminal (true (on p)) (true (on q)))
(<= (goal player 100) (true (on p)) (true (on q)))

```

Figure 1: Example of GDL game description.

which propositions are true in that state. A set of reserved keywords is used to define the characteristics of the game.

Figure 1 shows as an example the GDL description of a very simple game, where a player can independently turn on two lights (p and q). After being turned on, each light will remain on. The game ends when both lights are on and the player achieves a goal with score 100. In the figure, the GDL keywords are represented in bold.

2.2 The PropNet

A Propositional Network (PropNet) [Schkufza *et al.*, 2008; Cox *et al.*, 2009] can be seen as a graph representation of GDL. Each component in the PropNet represents either a proposition or a logic gate. Propositions can be distinguished into three types: *input* propositions, that have no input components, *base* propositions, that have one single *transition* as input, and *view* propositions, that are the remaining ones. The truth values of *base* propositions represent the state of the game. The dynamics of the game are represented by *transitions* that are identity gates that output their input value with one step delay and control the truth values of *base* propositions in the next step. The truth value of every other component is a function of the truth value of its inputs, except for *input* propositions, for which the game playing agent sets a value when choosing the action to play. Figure 2 shows as an example the PropNet that corresponds to the GDL description given in Figure 1.

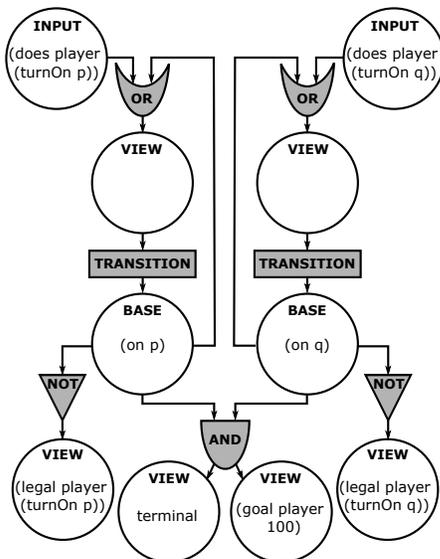


Figure 2: PropNet structure example.

3 PropNet Implementation

To create the PropNet the algorithm already provided in the GGP-Base framework is used.¹ This algorithm is implemented in the `create(List<Gdl> description)` method of the `OptimizingPropNetFactory` class and builds the PropNet according to the rules in the given GDL description.

The final product of the algorithm is a set of all the components in the PropNet, each of which has been connected to its input and output components. This set can then be used to initialize a PropNet object. The algorithm distinguishes six different types of components: *constants* (TRUE and FALSE), *propositions*, *transitions* and *gates* (AND, OR, NOT).

The GGP-Base framework also provides a PropNet class that can be initialized using the created set of components. We used this class as a starting point and implemented some changes to the initialization process to ensure that the PropNet respects certain constraints that are needed for the optimizations algorithms to work consistently. The first step of the initialization iterates over all the components in the PropNet and inserts them in different lists according to their type. While iterating over all the components, the following are the main actions that the initialization algorithm performs:

- Identify a single TRUE and a single FALSE constant, creating them if they do not exist, or removing the redundant ones.
- Identify the type of each proposition. Each proposition must be associated to one type only. A proposition that has a *transition* as input is identified as BASE type and a proposition that corresponds to a GDL relation containing the *does* keyword is identified as INPUT type. The propositions corresponding to GDL relations containing the *legal*, *goal* or *terminal* keyword are identified as LEGAL, GOAL and TERMINAL type respectively. To all other propositions the type OTHER is assigned.
- Make sure that all the INPUT and LEGAL propositions are in a 1-to-1 relation. If a proposition is detected as being an INPUT but there is no corresponding LEGAL in the PropNet, then it can be removed since we are sure that the corresponding move will never be chosen by the player. On the contrary, if there is a LEGAL proposition with no corresponding INPUT, the INPUT proposition is added to the PropNet, since the LEGAL proposition might become true at a certain point of the game and the player might choose to play the corresponding move.
- Make sure that only constants and INPUT propositions have no input components. If a different component is detected as having no inputs, set one of the two constants as its input. This action is needed because as a by-product of the PropNet creation some OR gates and non-INPUT propositions might have no inputs. The behaviour of the PropNet has been empirically tested to be consistent when such components are connected to the FALSE constant.

¹We have used a more recent and improved version than the one tested in [Schiffel and Björnsson, 2014].

Algorithm 1 Remove constant-value components

```
1: OPT0(propnet)
2:  $O_T \leftarrow \text{outputs}(\text{TRUE})$ 
3:  $O_F \leftarrow \text{outputs}(\text{FALSE})$ 
4: while  $O_T \neq \emptyset$  or  $O_F \neq \emptyset$  do
5:   REMOVEFROMTRUE( $O_T, O_F$ )
6:   REMOVEFROMFALSE( $O_T, O_F$ )
7: end while
```

4 Optimizations

The PropNets built by the algorithm given in the GGP-Base framework [Schreiber, 2013] contain usually many components that are not strictly necessary to reason on the game. This section presents four optimizations that can be performed on the PropNet structure to reduce the number of this components. Opt0 removes components that are known to have a constant truth value, Opt1 removes propositions that do not have a particular meaning, Opt2 detects more constant components and removes them, and Opt3 removes components that have no output and are not influential. All the optimization algorithms except the last one are already provided in the GGP-Base framework. The algorithms described here contain some minor modifications with respect to the original GGP-Base version in order to adapt them to the changes that were performed on the PropNet class structure.

4.1 Opt0: Remove Constant-value Components

This optimization removes from the PropNet the components that are known to be always *true* or always *false* and at the same time do not have a particular meaning for the game.

Algorithm 1 shows the main steps of this procedure. The sets O_T and O_F , at any moment, contain respectively the outputs of the *TRUE* and the outputs of the *FALSE* constant that still have to be checked for removal. At the beginning O_T contains all the outputs of the *TRUE* constant and O_F contains all the outputs of the *FALSE* constant (Lines 2 and 3). The procedures REMOVEFROMTRUE(O_T, O_F) and REMOVEFROMFALSE(O_T, O_F) (Lines 5 and 6) check the outputs of the *TRUE* and of the *FALSE* constant respectively. Algorithm 2 shows exactly which components the first procedure removes. The algorithm for the second procedure removes the outputs of the *FALSE* constant in a similar way. In the case of the *FALSE* constant, also always false GOAL and LEGAL propositions are removed since they will never be used. Moreover, whenever a LEGAL proposition is removed also the corresponding INPUT proposition is removed, since it is certain that the corresponding move will never be played.

Note that whenever a component is removed or detected as having always a constant value, it means that also its output is constant, thus its output components are connected directly to one of the two constants. In this case each output component will be added to the appropriate set (either O_T or O_F) to be checked in the next steps.

Algorithm 1 alternates between the two functions mentioned above until both sets, O_T and O_F , are empty. This repetition is needed because of the NOT gate. Whenever this gate is removed from the outputs of a constant, its outputs are

Algorithm 2 Remove true components

```
1: procedure REMOVEFROMTRUE( $O_T, O_F$ )
2:   while  $O_T \neq \emptyset$  do
3:      $c \leftarrow \text{removeComponent}(O_T)$ 
4:     switch  $cType(c)$  do
5:       case TRANSITION
6:         if  $|\text{outputs}(c)| = 0$  then
7:            $\text{propnet.remove}(c)$ 
8:         end if
9:       case NOT
10:        connect  $\text{outputs}(c)$  to FALSE
11:         $O_F \leftarrow O_F \cup \text{outputs}(c)$ 
12:         $\text{propnet.remove}(c)$ 
13:       case AND
14:        if  $|\text{inputs}(c)| = 1$  then
15:          connect  $\text{outputs}(c)$  to TRUE
16:           $O_T \leftarrow O_T \cup \text{outputs}(c)$ 
17:           $\text{propnet.remove}(c)$ 
18:        else if  $|\text{inputs}(c)| = 2$  then
19:          connect  $\text{outputs}(c)$  to other input
20:           $\text{propnet.remove}(c)$ 
21:        end if
22:       case OR
23:        connect  $\text{outputs}(c)$  to TRUE
24:         $O_T \leftarrow O_T \cup \text{outputs}(c)$ 
25:         $\text{propnet.remove}(c)$ 
26:       case PROPOSITION
27:        connect  $\text{outputs}(c)$  to TRUE
28:         $O_T \leftarrow O_T \cup \text{outputs}(c)$ 
29:        if  $pType(c) \in \{\text{OTHER, BASE}\}$  then
30:           $\text{propnet.remove}(c)$ 
31:        end if
32:     end switch
33:   end while
34: end procedure
```

connected to the other constant, thus the set of outputs to be checked for that constant will still have at least one element.

4.2 Opt1: Remove Anonymous Propositions

This optimization is trivial, nevertheless useful as it removes many useless components from the PropNet. The algorithm for this optimization (Algorithm 3) simply iterates over all the propositions in the PropNet and removes the ones with type OTHER, connecting their input directly to each of their outputs. These propositions can be safely removed as they do not have any special meaning for the game.

Algorithm 3 Remove anonymous propositions

```
1: OPT1(propnet)
2: for all  $p \in \text{propositions}(\text{propnet})$  do
3:   if  $pType(p) = \text{OTHER}$  then
4:     connect  $\text{input}(p)$  with  $\text{outputs}(p)$ 
5:      $\text{propnet.remove}(p)$ 
6:   end if
7: end for
```

Algorithm 4 Detect and remove constant-value components

```
1: OPT2(propnet)
2: Initialize all the parameters and the stack S
3: while S ≠ ∅ do
4:   [c, Pi] ← S.pop()
5:   Oc ← TOOUTPUTVALUESET(c, Pi)
6:   Pc ← Oc \ Vc
7:   if Pc ≠ N then
8:     Vc ← Vc ∪ Pc
9:     for all o ∈ outputs(c) do
10:      S.push(o, Pc)
11:     end for
12:     if cType(c) = PROPOSITION then
13:       if pType(c) = LEGAL then
14:         i ← correspondingInput(c)
15:         S.push(i, Pc)
16:       end if
17:     end if
18:   end if
19: end while
20: for all c ∈ components(propnet) do
21:   if Vc = T or Vc = F then
22:     Connect c to the appropriate constant
23:   end if
24: end for
25: OPT0(propnet)
```

4.3 Opt2: Detect and Remove Constant-value Components

This optimization can be seen as an extension of Opt0 where, before removing from the PropNet the constant value components directly connected to the *TRUE* and *FALSE* constant, the algorithm detects if there are other constant value components that have not been discovered yet.

This optimization (see Algorithm 4) associates to each component *c* in the PropNet a set *V_c* that contains all the truth values that such component can assume during the whole game. There are only four possible sets of truth values, namely:

- $N = \emptyset$: if the corresponding component can assume *neither* of the truth values.
- $T = \{true\}$: if the corresponding component can only be *true* during all the game.
- $F = \{false\}$: if the corresponding component can only be *false* during all the game.
- $B = \{true, false\}$: if the corresponding component can assume *both* values during the game.

The idea behind the algorithm is to start from the components for which the truth value that they will assume in the initial state of the game is known. It then propagates this value to each of their outputs *o* and updates the corresponding truth value set *V_o*. Whenever the truth values set of a component is updated, the algorithm propagates such changes on to its outputs components. This process will eventually end when the truth values sets of all components stop changing. Termination is guaranteed since only the truth values just added to the

truth values set of a component are propagated to its outputs and the number of possible truth values is finite.

When the algorithm starts, the set *V_c* of each component is set to *N*, since it is not known yet which values the component can assume. For each AND gate *a* the algorithm keeps track of *TI_a*, i.e. the number of inputs of *a* that can assume the *true* value. Similarly, for each OR gate *o* the algorithm keeps track of *FI_o*, i.e. the number of inputs of *o* that can assume the *false* value. These parameters are used to detect when an AND gate and an OR gate can assume respectively the *true* (if $TI_a = |inputs(a)|$) and the *false* (if $FI_o = |inputs(o)|$) value. These values are initialized to 0 for all the gates.

The algorithm exploits a stack structure *S* to keep track of the components for which the set of truth values that its input(s) can assume is changed. A pair (*c*, *P_i*) is added to the stack when the algorithm detects that an input *i* of the component *c* can also assume the values in the set $P_i \subseteq V_i$, and such values must be propagated to the component *c*. At the beginning the stack is filled with the following pairs:

- (*TRUE*, *T*), the *TRUE* constant can assume value *true*.
- (*FALSE*, *F*), the *FALSE* constant can assume value *false*.
- (*i*, *F*), for each INPUT proposition *i* in the PropNet. Each INPUT proposition can be *false* since we assume that no game exists where one player can only play a single move for the whole game.
- (*b_i*, *T*), for each BASE proposition *b_i* in the PropNet that is *true* in the initial state.
- (*b_i*, *f*), for each BASE proposition *b_i* in the PropNet that is *false* in the initial state.

During each iteration, the algorithm pops a pair (*c*, *P_i*) from the stack (Line 4) and checks if, given the new truth values *P_i* that the input *i* can assume, also the truth values *V_c* of its output *c* will change. Note that not for each type of component the set of truth values that its input can assume corresponds to the set of truth values that the component itself can output. The NOT component *n*, for example, has $V_n = T$ if its input *i* has $V_i = F$. Moreover, for an AND gate *a*, $true \in V_a \Leftrightarrow true \in V_i, \forall i \in inputs(a)$. The same holds for the *false* value for an OR gate. This means that the algorithm must first change the values in *P_i* according to the type of the component *c*, obtaining the new set of truth values *O_c* that *c* can output. This is done at Line 5 by the function TOOUTPUTVALUESET(*c*, *P_i*). Subsequently, the algorithm checks if in *O_c* there are some values *P_c* that were not in *V_c* yet (Line 6), and if so, it adds them to the set *V_c* (Line 8) and records on the stack that they have to be propagated to all the outputs *o* of *c* (Line 10). Here the algorithm treats each LEGAL propositions as if it was a direct input of the corresponding INPUT proposition, thus whenever the truth values set of a LEGAL proposition changes, the values are propagated to the corresponding INPUT proposition (Lines 12-17).

When no more changes are detected in the truth values sets (Line 3), the process terminates. At this point, the truth values set of each component is checked (Line 21) and if it equals the set *T* or *F* it is certain that the component will always be respectively *true* or *false*. It can then be disconnected from its input(s) and connected to the correct constant (Line 22).

Algorithm 5 Remove output-less components

```
1: OPT3(propnet)
2:  $Q \leftarrow \text{components}(\text{propnet})$ 
3: while  $Q \neq \emptyset$  do
4:    $c \leftarrow \text{removeElement}(Q)$ 
5:   switch  $cType(c)$  do
6:     case AND, OR, NOT, OTHER PROPOSITION
7:       if  $|\text{outputs}(c)| = 0$  then
8:          $Q \leftarrow Q \cup \text{inputs}(c)$ 
9:          $\text{propnet.remove}(c)$ 
10:      end if
11:    end switch
12: end while
```

The last step the algorithm performs consists in running the same algorithm that was proposed as Opt0 to remove all the newly detected constant components (Line 25).

4.4 Opt3: Remove Output-less Components

This optimization is also quite trivial, but helps remove some more useless components. Algorithm 5 shows this procedure: all the components in the PropNet are checked, if they are gates, or propositions of type OTHER and they have no output they are removed from the PropNet. Every time a component is removed, its inputs are added again to the set of components to be checked, since removing their outputs might have made them output-less.

5 Empirical Evaluation

5.1 Setup

To evaluate the performance of the PropNet multiple series of experiments are performed. Each of them compares the performance of the PropNet with different optimizations and combinations of optimizations. Each series of experiments poses the bases to decide which other combinations of optimizations to check. Furthermore, the optimized PropNet that shows the best overall performance is compared with the GGP-Base Prover. Both reasoners are also tested with the addition of a cache that memorizes the queries results.

The reasoners are tested using flat Monte-Carlo Search (MCS) on a set of heterogeneous games. For each reasoner the search is run from the initial state of the game with a time limit of 20s. This experiment is repeated 100 times for each of the chosen games. Such games are the following: *Amazons*, *Battle*, *Breakthrough*, *Chinese Checkers* with 1, 2, 3, 4 and 6 players, *Connect 4*, *Othello*, *Pentago*, *Skirmish* and *Tic Tac Toe*. The GDL descriptions of these games can be found on the GGP-Base repository [Schreiber, 2016].²

One of the reasons behind the choice of repeating each experiment multiple times for each game is that the number of components that the PropNet of a game has when created by the basic algorithm (i.e. without optimizations) is not always constant. This variance in the number of components could be due to the non-determinism of the order in which

²The GDL descriptions used for the experiments were downloaded from the repository on 03/02/2016.

game rules are translated into PropNet components for different runs of the algorithm. This can cause a different grounding order of the GDL description, originating more or less propositions and can also cause gates and propositions to be connected in different equivalent orders.

Another series of experiments matches two MCTS-based players, one that uses the Prover and one that uses the PropNet, against each other. We use the same 13 games that were used for the other experiments. Each player has 10s per move to perform the search. A new PropNet is built for each match in advance, before the game playing starts. For each game, if r is the number of roles in the game, there are 2^r different ways in which 2 types of players can be assigned to the roles [Sturtevant, 2008]. Two of the configurations involve only the same player type assigned to all the roles, thus are not interesting and excluded from the experiments. Each configuration is run the same amount of times until at least 100 matches have been played.

The final series of experiments further investigates the effect of the cache. It matches two MCS-based players that use the Prover, one with cache and one without, against each other, and two MCS-based players that use the best optimized PropNet, one with cache and one without, against each other. The settings are the same as in the previous experiment. At the end of each match the speed is computed by dividing the total number of nodes visited by the total time spent on the search during the whole game. Since we are only interested in the reasoning speed, for this experiment we do not consider the 10s search time per move strictly, but we allow each player to finish the current simulation when this time expires.

Before running any of the described experiments, the PropNet and all its optimized versions were tested against the Prover for consistency. For each game in the GGP-Base repository [Schreiber, 2016], for a duration of 60s, the same random simulations were performed querying both the Prover and the currently tested version of the PropNet for next states, legal moves, terminality and goals in terminal states. The results returned by the PropNet were compared with the ones returned by the Prover for consistency. All the PropNet versions passed this test on all the games in the repository, except for 12 games for which the PropNet construction could not be completed in the given time.

In all experiments, a limit of 10 minutes was given to the program to build the PropNet. The experiment that further investigates the effects of the cache on complete games was performed on an AMD Opteron 6274 2.2-GHz. All other experiments were performed on an AMD Opteron 6174 2.2-GHz.

5.2 Comparison of Single Optimizations

In the first series of experiments we compared with the basic version of the PropNet (BasicPN) the performance of each of the previously described optimizations applied singularly (Opt0, Opt1, Opt2, Opt3). Columns 2 to 6 of Table 1 show the obtained results. For each PropNet variant, for each game the first block of the table gives the average simulation speed in nodes per second, the second block gives the average number of components and the third block gives the average total initialization time (creation+optimization+state initialization) in milliseconds. The line at the bottom of each block reports

Game	BasicPN	Opt0	Opt1	Opt2	Opt3	Opt12	Opt102	Opt13	Opt1023
Avg. speed (nodes/second)									
amazons	35.1	41.4	32.7	41	40.2	38.5	41.4	32.3	41
battle	34957	49666	37877	51257	35276	59308	59697	39981	60419
breakthr.	50557	50932	65518	51357	51058	66943	66551	66833	66991
c-check.1	426374	427773	550230	444671	424516	570858	562737	541682	561634
c-check.2	125581	128623	189368	128910	127519	194442	192048	190161	193752
c-check.3	155886	157242	169352	161000	159267	175410	176162	170722	176185
c-check.4	105766	106738	127886	107153	105660	130362	130279	129194	130451
c-check.6	119650	118547	126863	113700	118783	127535	128111	127619	129000
connect4	110081	113484	105081	112920	109672	127053	126535	105978	129272
othello	290	1610	235	1604	295	1934	1894	245	1979
pentago	76336	76786	116065	76721	96782	116353	115064	117127	121108
skirmish	5887	6022	6780	6230	6151	7075	7042	7403	7600
ticTacToe	223403	228056	248769	234915	222952	259980	257285	247246	257525
Avg. rel. increase	-	40.59%	15.51%	41.44%	3.95%	70.32%	69.39%	17.38%	73.48%
Avg. number of components									
amazons	1497649	1254742	741874	1192364	1023913	623460	623460	711596	596240
battle	51197	14267	36863	14262	50721	11084	11077	36676	10902
breakthr.	10745	10678	5933	10678	10584	5900	5900	5869	5836
c-check.1	793	785	559	785	789	556	556	559	556
c-check.2	1540	1524	1179	1524	1532	1172	1172	1179	1172
c-check.3	2411	2389	1845	2236	2400	1718	1718	1845	1718
c-check.4	3159	3119	2465	2999	3133	2362	2362	2465	2362
c-check.6	4451	4411	3473	4123	4431	3238	3238	3473	3238
connect4	2164	2063	1724	1291	2114	1063	1063	1724	1056
othello	1311988	274940	1033197	274940	1305515	208510	208510	1031580	206846
pentago	3696	3706	1470	3708	2111	1464	1473	1338	1337
skirmish	126019	124267	108171	124267	78575	107296	107296	62427	61552
ticTacToe	312	291	249	291	302	239	239	249	239
Avg. rel. increase	-	-14.28%	-29.21%	-18.62%	-9.49%	-42.34%	-42.32%	-32.52%	-45.65%
Avg. total PropNet initialization time (ms)									
amazons	311335	313719	314455	417097	315637	411905	400113	312793	401559
battle	5756	6027	5897	6303	5869	6367	6233	5968	6329
breakthr.	3989	4007	4012	4358	3910	4354	4415	3982	4328
c-check.1	2699	2651	2659	2653	2707	2693	2654	2707	2652
c-check.2	2848	2773	2810	2873	2775	2848	2843	2817	2842
c-check.3	3162	3140	3159	3251	3149	3214	3186	3160	3167
c-check.4	3258	3261	3241	3473	3244	3405	3330	3275	3379
c-check.6	3225	3203	3204	3639	3205	3423	3430	3207	3395
connect4	2437	2465	2456	2698	2430	2536	2555	2417	2525
othello	35756	36486	37074	39417	36544	39170	36689	35359	37804
pentago	4249	4230	4278	4390	4232	4269	4286	4308	4325
skirmish	11887	11702	11664	12089	11824	12386	12285	11870	12577
ticTacToe	1525	1529	1523	1522	1508	1532	1535	1524	1555
Avg. rel. increase	-	0.13%	0.24%	7.69%	-0.19%	6.38%	5.18%	0.18%	5.66%

Table 1: Comparison of the optimizations and some of their combinations

the average over the 13 games of the percentage increase of the values considered in the block relative to the basic version of the PropNet (BasicPN).

The main interest is the speed increase that the optimizations induce on the PropNet, however the other two aspects are also relevant. A low number of components means less memory usage, and a shorter initialization time means more time for metagaming at the beginning of a match (or more chances to avoid timing out when the start clock time is short). From the table it seems that for most of the games, as expected, the increase in the simulation speed is related to the decrease in the number of components in the PropNet.

As can be seen, none of the optimizations outperforms the others in speed for all games. Opt0 and Opt2 seem to have the best performance in *Amazons*, *Battle*, *Othello* and *Connect 4*, while Opt1 performs best in the other games. When look-

ing at the initialization time, Opt2 is the one that increases it the most for almost all the games. Another observation is that the performance of Opt2 is overall better than the one of Opt0. This was expected because Opt2 is an extension of Opt0, thus for the same PropNet it always removes at least the same number of components as Opt0.

If we consider the speed as main choice criterion, Opt0 and Opt2 are the ones that, on average, produce the highest speed increase. However, the high average is due to the considerable relative increase that they produce in *Othello*. If we consider as starting point for the next series of experiments the optimization that produces the highest speed in most of the games, then Opt1 is the most suitable to be selected. Moreover, Opt1 is the optimization that reduces the most the number of components of the PropNet without consistently slowing down the initialization process.

Game	Prover	CachePr.	Opt1023	CacheOpt1023
Avg. speed (nodes/second)				
amazons	2.2	2.6	41	76.3
battle	44.6	48.2	60419	18989
breakthr.	220	274	66991	52602
c-check.1	2066	42629	561634	758864
c-check.2	1513	1693	193752	98255
c-check.3	1227	1309	176185	59172
c-check.4	620	689	130451	50672
c-check.6	701	758	129000	43840
connect4	209	268	129272	106262
othello	2.7	3	1979	1744
pentago	212	237	121108	49279
skirmish	22.2	23.4	7600	8034
ticTacToe	1823	243965	257525	876546

Table 2: Comparison of the PropNet with the Prover

5.3 Comparison of Combined Optimizations

In this series of experiments Opt1 is combined with other optimizations applied in sequence. In general, when we refer to OptXY we refer to the PropNet optimization obtained by applying OptX and OptY in sequence. These experiments compare the combinations of optimizations Opt13, Opt12 and Opt102 with the single optimization Opt1. The combination Opt10 has been excluded from the test since it is considered less interesting. As also previously mentioned, Opt0 always removes a subset of the components that are removed by Opt2, thus Opt10 is expected to perform less than Opt12. However, Opt0 has less negative impact than Opt2 on the total initialization time. This is why these experiments include the test of Opt102: we want to see if the application of Opt0 before Opt2 can speed up the process of Opt2 that will then run on a smaller PropNet.

The results of this series of experiments can be seen in columns 7, 8 and 9 of Table 1. Regarding the speed, Opt12 seems to be the one achieving the best overall performance. However, the performance of Opt102 is rather close, as expected, since these two combinations should remove the same number of components in each PropNet. The difference in performance is probably due to some variance in the data. Moreover, running Opt0 before Opt2 helps reducing the initialization time for large games, while it seems to have almost no effect on smaller games.

Opt13 is the one that, regarding the speed, performs worse in this series of experiments, thus it has been excluded from further tests. Among Opt12 and Opt102, it has been chosen to keep testing on top of Opt102 because of its shorter initialization time for games with large PropNets, given that its speed is still comparable with the one of Opt12.

5.4 Comparison of PropNet and Prover

In this series of experiment only one more interesting combination of optimizations is left to test: Opt1023. No further gain in performance can be obtained by repeating the same optimizations multiple times in a row, since no further change will take place in the structure of the PropNet. Thus it is not interesting to evaluate combinations of optimizations that extend Opt1023.

Game	Opt1023	Game	Opt1023
battle	100.0(± 0.0)	c-check.6	69.4(± 7.67)
breakthr.	100.0(± 0.0)	connect4	99.0(± 1.38)
c-check.2	98.0(± 2.76)	pentago	100.0(± 0.0)
c-check.3	83.3(± 7.27)	skirmish	100.0(± 0.0)
c-check.4	70.5(± 8.48)	ticTacToe	50.0(± 0.0)

Table 3: Win percentage of the PropNet-player against the Prover-player

The last column of Table 1 shows the statistics for Opt1023. For most of the games, Opt1023 seems to be the fastest. It is also the one that reduces the number of PropNet components the most. As for the initialization time, this optimization is between a few milliseconds and a bit more than 1 second slower than the basic version of the PropNet, except for *Amazons*. Optimizing the large PropNet of *Amazons* can slow down the initialization time by more than a minute.

Table 2 shows the comparison of the speed of Opt1023 with the Prover. For both of them also a cached version is tested. The GGP-Base framework [Schreiber, 2013] provides a cache structure that memorizes the results returned by the underlying reasoner and prevents it from computing the same queries multiple times. As can be seen, the cache seems to be helping only in a few games, but this is due to the simulation being run only for the first step of the game (see Subsection 5.6). For games with many states and legal moves the cache takes more time and more game steps to be filled with enough entries and really make an impact on the speed.

Looking at Table 2 it is clearly visible how the optimized PropNet achieves a much better performance than the Prover in the considered games, and Table 1 shows how it performs also better than the basic version.

As further test to evaluate the robustness of the PropNet, we compared its performance with the one of the Prover also on all the games (about 300) in the GGP-Base repository [Schreiber, 2016]. For each considered reasoner we performed once for each game the Monte-Carlo search with a time limit of 60s. On average the basic version of the PropNet was 418 times faster than the Prover and the version with all optimizations, Opt1023, was 698 times faster.

5.5 Game Playing Performance

In this series of experiments an MCTS player that uses the PropNet reasoner is matched against one that uses the Prover. Table 3 shows the win percentage of the PropNet-player against the Prover-player. In most of the games the PropNet-player achieves a win percentage close or equal to 100%. The games in which the performance of the PropNet-player seems to drop are the ones with more than 2 players. *Chinese checkers* with 4 and 6 players are the ones where the win percentage for the PropNet-player is the lowest, but it is still significantly better than the one of the Prover-player. The game of *Tic Tac Toe* is the only exception, since its state space is so small that both players can easily reach a sufficient number of simulations to play optimally and result in a tie.

No results are shown for *Amazons* and *Othello* since for both games the Prover-player could not complete even one iteration of the MCTS algorithm in the given time limit. The

Game	Prover	CachePr.	Opt1023	CacheOpt1023
Avg. speed (nodes/second)				
amazons	5.7	2316	28.1	30519
battle	45.2	2457	38656	36607
breakthr.	235	241	56275	51569
c-check.1	2273	466014	532426	862408
c-check.2	1478	93251	159935	258639
c-check.3	1105	28300	118160	133733
c-check.4	536	32684	82955	117017
c-check.6	607	5744	57008	53230
connect4	180	2455	122325	207508
othello	3.2	5502	649	80328
pentago	152	155	93185	75998
skirmish	26	4081	2997	3946
ticTacToe	1650	287380	225127	547398

Table 4: Effect of the cache on Prover and PropNet over complete games

single-player version of *Chinese checkers* is tested separately and the score is used to measure the performance of the players. This game has a relatively small search space, so both players achieved the maximum score in every match.

5.6 Cache Effect on Complete Games

Table 4 shows the results of the comparison of the Prover and the PropNet with and without cache over complete games. It is visible how the cache actually provides some benefits if used for the whole game. For the Prover it increases the speed for all the games, while for the PropNet it increases the speed for most of them. The increase is especially relevant for the games of *Amazons* and *Othello*. These results, together with the ones in Table 2, are probably an indication that usually the cache slows down the performance in the initial steps of the game, but then balances this loss towards the endgame, when the chance of finding queries results in the cache increases.

Comparing with the Prover over complete games also helps putting the PropNet into perspective with the other GDL reasoners analysed in the paper [Schiffel and Björnsson, 2014]. Even if that paper uses different experimental settings than ours, we can still make some general observations. Considering the performance of the reasoners that, like the PropNet, rely on an alternative representation of the GDL description, it seems that our implementation of the PropNet provides for most of the games a speed increase of the same order of magnitude when compared to the Prover. Moreover, for *Amazons*, *Othello* and *Chinese Checkers* with 4 and 6 players, it seems that our optimized PropNet, especially with the cache, could even achieve a better performance in similar circumstances.

6 Conclusion and Future Work

In this paper the performance of a PropNet-based reasoner has been evaluated, together with four possible optimizations of the structure of the PropNet and their impact on the performance. Even though the tested implementation of the PropNet is based on the code provided by the GGP-Base framework, the principles behind its representation and its optimizations can also be applied in general.

Experiments have shown that the use of a PropNet substantially increases the reasoning speed by, on average, at

least two orders of magnitude with respect to the GGP-Base Prover. Moreover, the addition of a combination of optimizations that reduce the size of the PropNet increases the reasoning speed further. Experiments also show that the reasoning speed increase has a positive effect on the performance of the PropNet-based player. This player achieves a win rate close to 100% in most of the games for which it is matched against an equivalent player based on the Prover.

Also the use of a cache proved to be useful in some games. For small games its effect is already visible in the first steps, while for most of the other games it helps only during later game steps. Future work could investigate a way of detecting for each game if and when the cache could be helpful.

Finally, another interesting aspect that future work could consider is the impact that the use of different strategies to propagate truth values among the components of the PropNet would have on the reasoning speed.

Acknowledgments

This work is funded by the Netherlands Organisation for Scientific Research (NWO) in the framework of the project GoGeneral, grant number 612.001.121.

References

- [Björnsson and Finnsson, 2009] Y. Björnsson and H. Finnsson. Cadiaplayer: A simulation-based general game player. *Computational Intelligence and AI in Games, IEEE Transactions on*, 1(1):4–15, 2009.
- [Coulom, 2007] R. Coulom. Efficient selectivity and backup operators in Monte-Carlo tree search. In *Computers and games*, pages 72–83. Springer, 2007.
- [Cox et al., 2009] E. Cox, E. Schkufza, R. Madsen, and M. R. Genesereth. Factoring general games using propositional automata. In *Proceedings of the IJCAI Workshop on General Intelligence in Game-Playing Agents (GIGA)*, pages 13–20, 2009.
- [Darper and Rose, 2014] S. Darper and A. Rose. Sancho GGP player. <http://sanchoggp.blogspot.nl/2014/07/sancho-is-ggp-champion-2014.html>, 2014.
- [Emsile, 2015] R. Emsile. Galvanise. https://bitbucket.org/rxe/galvanise_v2, 2015.
- [Love et al., 2008] N. Love, T. Hinrichs, D. Haley, E. Schkufza, and M. R. Genesereth. General game playing: Game description language specification. Technical report, Stanford University, Stanford, CA, USA, 2008.
- [Schiffel and Björnsson, 2014] S. Schiffel and Y. Björnsson. Efficiency of GDL reasoners. *Computational Intelligence and AI in Games, IEEE Transactions on*, 6(4):343–354, 2014.
- [Schkufza et al., 2008] E. Schkufza, N. Love, and M. R. Genesereth. Propositional automata and cell automata: Representational frameworks for discrete dynamic systems. In *AI 2008: Advances in Artificial Intelligence*, pages 56–66. Springer, 2008.
- [Schreiber, 2013] S. Schreiber. The General Game Playing base package. <https://github.com/ggp-org/ggp-base>, 2013.
- [Schreiber, 2016] S. Schreiber. Games - base repository. <http://games.ggp.org/base/>, 2016.
- [Sturtevant, 2008] N. R. Sturtevant. An analysis of UCT in multi-player games. *ICGA Journal*, 31(4):195–208, 2008.