

©Copyright 2015
Abdelrahman Elogeel

Selecting Robust Strategies in RTS Games via Concurrent Plan Augmentation

Abdelrahman Elogeel

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Science

University of Washington

2015

Reading Committee:

Ankur Teredesai, Chair

Matthew Alden

Martine De Cock

Andrey Kolobov

Program Authorized to Offer Degree:
Institute of Technology

University of Washington

Abstract

Selecting Robust Strategies in RTS Games
via Concurrent Plan Augmentation

Abdelrahman Elogeel

Chair of the Supervisory Committee:
Professor Ankur Teredesai
Institute of Technology

The multifaceted complexity of real-time strategy (RTS) games requires AI systems to break down policy computation into smaller subproblems such as strategic planning, tactical planning, and reactive control. To further simplify planning at the strategic and tactical levels, state-of-the-art automatic techniques such as case-based planning (CBP) produce deterministic plans. For what is inherently an uncertain environment CBP plans rely on *replanning* when the game situation digresses from the constructed plan. A major weakness of this approach is its lack of robust adaptability: repairing a failed plan is often impossible or infeasible due to real-time computational constraints resulting in a game loss. This thesis presents a technique that selects a robust RTS game strategy using ideas from contingency planning and exploiting action concurrency in strategy games. Specifically, starting with a strategy and a linear tactical plan that realizes it, our algorithm identifies a plan's failure modes using available game traces and adds concurrent branches to it so that these failure modes are mitigated. For example, approach may train an army reserve *concurrently with* an attack on the enemy, as defense against a possible counterattack. After augmenting each strategy from an available library (e.g., learned from human demonstration) our approach picks one with the most robust augmented tactical plan. Extensive evaluation on popular RTS games (StarCraft and Wargus) that share engines with other games indicates concurrent augmentation significantly improves win-rate where baseline strategy selection consistently lead to a loss.

ACKNOWLEDGMENTS

One of my life's biggest adventures, Master's studies, is drawing to an end. In this moment, I would like to thank everyone who helped me through it.

I have been privileged to work with Andrey Kolobov during my thesis work. Andrey has been passionate and keen since day one. I am grateful to Andrey's mentoring and handholding in shaping the research problem, providing guidance during experiments and elevating my writing skills. I always had fear from writing papers and getting into the world of publications but Andrey made the journey enjoyable and gave me a lot of confidence. I admit that this work was made possible thanks to Andrey's efforts.

I would like to thank the people behind the scenes who introduced me to Andrey – Professor Ankur Teredesai and Matthew Alden – as they have helped me a lot by giving guidance and feedback throughout the thesis research.

I would like to deeply thank Jeff Wilcox and Matt Gibbs, my managers at Microsoft, for giving me the flexibility in my work schedule. Thanks to their understanding and support, my master's was made possible.

It all started thanks to my friend Wael Mahmoud who introduced UWT to me, and by Professor Mohamed Ali who encouraged me to apply to the Master's program there.

Perhaps the most exciting part of my Master's thesis work has been the development of Yarmouk, an AI bot for playing StartCraft, and I must admit that without the continued help from Muhamad Lotfy it would not have happened. He has influenced a lot of my decisions related to the implementation of Yarmouk, and has given inspiration for almost every module in it. I'd like to thank Omar Tawfik and Mohammed Samir for their contributions to the Yarmouk engine as well.

Of course, none of this would have been possible without the support from my family: my wife

Hoda, my son Omar, my daughter Maryam, my parents Zakaria and Amal, my brothers Omar and Abdelkader. Neither would this effort have been worth it if they weren't around. It surely has been hard for you, my parents, to be so far away. I'd like this work to be a sign of my appreciation of your love, patience, and compassion.

Hoda, you have stood by me all the days and nights when I was buried in work,. Thank you for being with me as my friend and my life partner. I owe you a lot of time, and promise to make it all up to you! To Omar and Maryam: both of you have played a big role in my finishing the thesis by not breaking my laptop and letting me sleep peacefully at night.

TABLE OF CONTENTS

	Page
List of Figures	ii
Chapter 1: Introduction	1
Chapter 2: Background	4
2.1 Related Work	8
Chapter 3: Concurrent Plan Augmentation	11
3.1 Robust OLCBP	12
Chapter 4: Example	18
Chapter 5: Experimental Evaluation	20
Chapter 6: Conclusion	29
Bibliography	30

LIST OF FIGURES

Figure Number	Page
2.1 A screen-shot of StarCraft: Brood War. The agent is building a city focusing on infrastructure and technology buildings but does not have defenses which leads to immediate failure in case of enemy attack.	5
2.2 OLCBP cycle. Image taken from [16]	6
2.3 Example of selected path in the decision tree.	9
4.1 A tactical plan (left) with a contingent branch (right).	18
5.1 Effect of CPA on the tactical performance of the rush attack (R) and mid-game attack (M) strategies in <i>StarCraft</i> . <i>The advantage of ROLCBP is statistically significant in all experiments except for rush attack on the Bottleneck map.</i>	23
5.2 Effect of CPA on overall performance in <i>StarCraft</i> . <i>The advantage of ROLCBP is statistically significant in all experiments.</i>	24
5.3 Effect of CPA on overall performance in <i>Wargus</i> . <i>The advantage of ROLCBP is statistically significant in all experiments.</i>	25
5.4 Effect of CPA on overall performance in <i>StarCraft</i> on the Bottleneck map with the Fabian strategy. In each case, ROLCBP trains on two races and plays versus the remaining one. <i>The advantage of ROLCBP is statistically significant in all experiments.</i>	25
5.5 Training against OLCBP and playing against static AI and vice versa.	27

DEDICATION

To my dear wife Hoda,
to my children Omar and Maryam
Without them, I would have written a Ph.D. dissertation instead.

Chapter 1

INTRODUCTION

Ten years after real-time strategy (RTS) games were proposed as a challenge for Artificial Intelligence [2], computational state of the art is still well short of the human performance level [17]. Among the aspects that make RTS games so difficult for computers to play well are complexity and partial observability of the worlds they depict, their adversarial nature, and the necessity to plan in real-time with durative actions, i.e. actions are not instantaneous, but take some amount of time to complete and a high degree of concurrency. In an attempt to make RTS game policy computation more tractable, researchers have broken this problem into subproblems varying by planning timescale (strategy, tactics, reactive control) and by their function (opponent modeling, economy/resource management, infrastructure construction, and others). Some of them such as planning at micro-timescales and terrain analysis, have been tackled with relative success using reinforcement learning techniques [17], which affected at least one application area that shares characteristics with RTS games, combat environments simulation [14]. Nonetheless, AI agents in modern commercial RTS games are still far from intelligent. Despite progress in related areas, strategic and high-level tactical planning in RTS games remains a challenge making it the focus of our efforts.

Selecting game strategy and high-level tactics are tightly intertwined tasks. On one hand, we would like to choose a strategy, i.e., a sequence of subgoals to achieve to win the game, for which a tactical plan that attains these subgoals is easy to find. On the other hand, tactical plans with a good chance of winning may be available only for specific subgoal sequences, thereby implicitly constraining strategy selection. Case-based planning [7] and the On-line Case-Based Planning (OLCBP) architecture derived from it [16] are among the few automatic approaches that have

attempted to solve these planning problems jointly in RTS games.¹Taking a set of subgoal specifications and gameplay logs (e.g., generated by humans) as inputs, OLCBP builds a hierarchical task network (HTN) [19] from them. It uses this HTN to produce a deterministic strategic and tactical plan at the start of a game. It follows this plan until the game state deviates from the plan’s expectations, in which case OLCBP replans. While a significant step towards automating strategy selection, OLCBP doesn’t explicitly attempt to maximize the chance of winning the game. Moreover, its reliance on replanning causes a game loss when the existing plan fails and no new one can be found [10].

The technique we present in this thesis, which we call *concurrent plan augmentation (CPA)*, mitigates the drawbacks of OLCBP by aiming to select the strategy with a set of tactical plans that empirically will least likely require replanning when executed. We implement CPA in an architecture called *ROLCBP (Robust OLCBP)*. Like OLCBP, ROLCBP identifies promising strategies and tactical realizations for them from game logs. Importantly, the number of such distinct strategies in RTS games is usually small, allowing ROLCBP to analyze each such strategy with its tactical plans and try to address the plan’s weaknesses. For each strategy, ROLCBP first identifies the likely steps at which its tactical plan will fail. Then ROLCBP searches for plans that can be executed *concurrently with* parts of the strategy’s main tactical plan and can decrease its empirical chance of failure. For example, if the original plan calls for immediately attacking the enemy and the attack often fails in simulation, ROLCBP might propose to train an army reserve simultaneously with the attack so that the player’s base does not get overrun by the opponent’s counterattack if the player’s attack peters out. Augmenting the tactical plans for each strategy and assessing their failure probability identifies the most empirically robust strategy, while storing the augmented plans reduces the need to replan in critical situations.

We conduct an extensive empirical evaluation of ROLCBP on the popular games of *StarCraft* and *Wargus* (an open-source clone of *Warcraft II* that shares its game engine with several other real-time strategy games). It demonstrates CPA’s broad applicability and shows that, thanks to

¹See Related Work for a discussion of a hard-coded approaches.

CPA, ROLCBP not only has a significantly higher win rate against built-in AI than OLCBP, but can also win on maps where OLCBP always loses. Moreover, whenever CPA does not affect strategy choice, it still improves win rate by making tactical plans for the selected strategy much less failure-prone. Last but not least, the experiments reveal that ROLCBP easily wins against opponents very dissimilar from those encountered during training.

In summary, our work makes the following contributions:

- We present concurrent plan augmentation (CPA), a technique that helps pick robust strategic and tactical plans.
- We implement CPA in an OLCBP-based architecture called ROLCBP and extensively test it on complex commercial RTS games. The results indicate that CPA provides quantitative as well as qualitative win rate gains by improving decision-making at both strategic and tactical levels.

Chapter 2

BACKGROUND

Real-time strategy games. RTS games are a genre that typically involves managing resources, building infrastructure, training an army, and ultimately defeating other players by destroying their units and infrastructure. RTS games have several distinctive attributes that put them beyond the capabilities of most standard game tree search and reinforcement learning techniques: the size, complexity, and non-determinism of the game environment, durative and concurrent action execution, partial observability, and others as described below:

- *Complexity of the environment.* The complexity comes from the sheer state space size and the number of available action in a given decision cycle [17]
- *Durative, concurrent actions.* The actions of all players are executed in parallel in real time. The actions have durations, usually different for each action type. Moreover, each player can perform several tasks concurrently. This is in contrast to games such as chess, where players take turns making moves, move duration is immaterial, and only one piece can be moved at a time.
- *Partial observability.* States in most RTS games are partially observable, so a player cannot see what the enemy is doing right now. This is called *fog of war*.

StarCraft, developed by Blizzard EntertainmentTM, and *Wargus*, an open-source clone of *Warcraft II*, are popular RTS games that we use as experimental testbeds in this thesis. A screenshot from *StarCraft* is shown in Figure 2.1. In *StarCraft*, every player controls one of three warrior races: Zerg, Protoss, or Terran. In *Wargus*, there are two of them: humans and orcs. In both games, warriors of each race differ in their combat characteristics and in the resources required to produce

them. Accordingly, good choices of strategy and tactics vary among races, and also depend on the world configuration, called a *map*, where the game is played.



Figure 2.1: A screen-shot of StarCraft: Brood War. The agent is building a city focusing on infrastructure and technology buildings but does not have defenses which leads to immediate failure in case of enemy attack.

The influence of RTS game research goes beyond games themselves. It aids the military through battle simulations and training programs and researching possible autonomous weapon systems. Moreover, because RTS games simulate the real world it's possible to use such environments as simulation environment for robotics, construction, resource management and more.

Strategic planning in RTS games. The aforementioned characteristics of RTS games make infeasible the use of standard game tree search and reinforcement learning for strategy and high-level tactics selection. In fact, even simply finding action outcome trajectories that lead to victory under simplifying assumptions is difficult for modern planners due to durative actions and concurrency. Even if the game environment is treated as deterministic and fully observable, and the adversarial nature of the opponents is ignored (these simplifications are normally made by commercial game AI systems), modern planners have difficulty quickly identifying linear trajectories that may lead

to winning the game, due to the presence of action durations and concurrency. As a result, strategic behavior of RTS game AI is either scripted or relies on adapting plans gleaned from human gameplay. On-Line case-based planning is an example of the latter.

On-line case-based planning. OLCBP [16] is a planning architecture derived from hierarchical task networks (HTN) [19] and case-based reasoning. Its inputs, along with a game description, are a set of subgoal specifications and gameplay logs (e.g., generated by humans). OLCBP operates in a cycle with two main phases, behavior acquisition and online expansion-execution (Figure 2.2). During behavior acquisition, it heuristically matches up parts of the logged game traces with the subgoals, thereby identifying both the strategies used in the game and the tactical plans for achieving the subgoals. After some additional processing, OLCBP stores the extracted plans, called *cases*, in its *case base*.

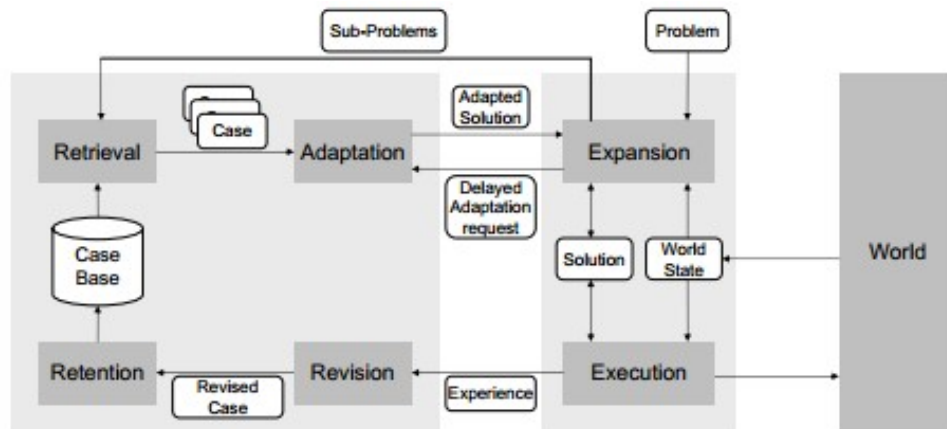


Figure 2.2: OLCBP cycle. Image taken from [16]

OLCBP assumes its opponents to be relatively static and views strategy and tactics selection as a classical planning problem with a possibility of replanning. OLCBP's expansion-execution step is responsible for dividing this planning problem into subproblems using an HTN constructed from the provided subgoals with the help of the game logs. The chosen strategy is passed to the retrieval submodule, which selects and adapts tactical plans for it from the case base. OLCBP executes these plans until the game situation deviates from the predicted one. In this case, it looks for a different tactical plan, strategy or both, given the player's current situation. Reliance on replanning is

OLCBP's significant weakness. Replanning takes time, but even a small delay in acting may mean losing the game. In addition, OLCBP's case base may not contain a back-up plan for unforeseen circumstances, also leading to a game loss.

The system retrieves a suitable plan for the current world state from a case base and sends it to the execution module, which in turn takes the HTN, evaluates it, and finds actions ready for execution. The system also provides a revision phase where a learning algorithm can be plugged in and later the revised case can be stored back into the case-base.

OLCBP handles plan failures via replanning that issues new request to the retriever to find a new plan to execute after the existing plan has failed. Empirical analysis demonstrated that relying heavily on replanning causes instability in the selected strategy which increases the probability of policy failure [5]. While replanning is certainly suitable in some cases, there are several difficulties with exclusive reliance on this approach in RTS games:

- **Computational cost** - replanning takes time for selecting a new strategy, expanding and then executing it, which is undesirable in RTS games and should be avoided as much as possible.
- **Ignorance** - in some situations the agent may not find a backup plan that is relevant to the current game state because it does not have one in its case base.
- **Resource misuse** - selecting a new strategy implies some initial groundwork that requires resources. Consuming an agent's resources on such tasks can be very expensive in critical situations where sometimes the immediate objective is just to survive.

Contingency planning. In contingency planning [18], an agent situated in a partially observable non-deterministic environment needs to devise a policy for reaching the goal from any state in which the agent may end up. Unlike in (partially observable) Markov Decision Processes (MDPs), the agent does not know the probabilities of its actions' outcomes, only the sets of possible outcomes themselves. In this thesis we show that, despite having been considered for non-adversarial settings so far, contingency planning can be potent in RTS game AI as well.

2.1 Related Work

Automatic strategy and high-level tactics selection methods have featured in several AI systems for playing RTS games; many, if not most, are derived from case-based planning. Darmok [16], I-Strategizer [4], and EISBot [21] use this technique with various additions. An agent developed for the DEFCON game employs a related technique, case-based reasoning, simulated annealing and decision tree learning. As in OLCBP, it has a pre-game learning phase that builds a case base of plans. To decide which plan to use in a new situation, the system learns a decision tree. [1].

Darmok [16] uses the OLCBP architecture described in Background section and thanks to that is capable of playing in an adaptive manner. It has a pre-game learning phase where it can observe human players and learn from their playing strategy [15]. However, it lacks real-time performance, as its plan retriever is relatively slow [11].

An extension of Darmok called I-Strategizer [4] improves the case revision phase of OLCBP and adds reinforcement learning to aid in tactical plan selection. The agent uses reinforcement learning technique, specifically Lambda Sarsa, a temporal difference approach. It also uses eligibility traces which add the concept of credibility to a plan. The agent starts in an exploration phase where it tries to use any available case and evaluate it. After learning some history about the used cases/plans the agent starts its exploitation phase where it looks for credible and successful plans. The Lambda Sarsa helps in determining case success and eligibility traces help in determining case credibility. Like Darmok, this system faces the issue of quickly handling plan failures.

EISBot [21] implements strategy selection as a combination of reactive techniques. A collection of hand-authored behaviors resides in a strategy manager for executing exact infrastructure build order. Goal-driven autonomy [12] is used to decouple strategy selection from execution. Similar to Darmok and I-Strategizer, EISBot uses a case-based planner for tactical plan selection. An agent developed for the DEFCON game [1] combines case-based reasoning, simulated annealing and decision tree learning. As in OLCBP, it has a pre-game learning phase that builds a case base of plans. The system learns a decision tree that is used to dictate what specific path to follow when the agent faces a new situation. An example of such a decision tree is shown in Figure 2.3.

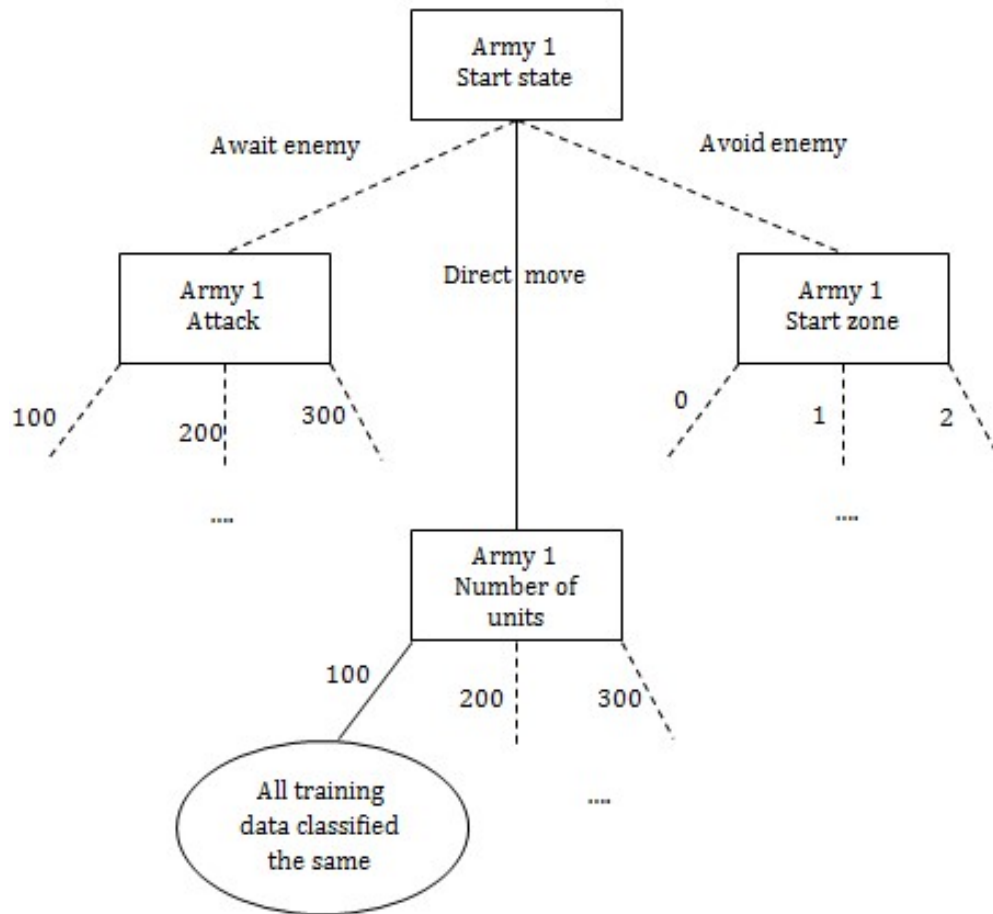


Figure 2.3: Example of selected path in the decision tree.

Hierarchical Task Networks have also been used for strategy selection. Besides OLCBP, which is partly based on HTNs, they have been used for this purpose in first-person shooters [8] and role-playing games [9]. Another potential alternative to OLCBP techniques is reinforcement learning. While we are not aware of its successful application to strategy selection in RTS games, it has been used for this purpose in first-person shooters. Hard-coded techniques for strategy and tactics selection are still strongly competitive with automatic ones, especially in complex games such as *StarCraft* [17]. Finite state machines have been particularly effective for specifying fixed patterns for strategic decision-making and other aspects of RTS games [6].

To conclude the coverage of relevant RTS game literature, we point out that the success of a

game AI system depends on many components besides strategy and tactics selection. The amount of work done on those components is enormous, and we refer the interested reader to a recent survey on the topic [17].

Several approaches related to CPA have been proposed in the area of sequential decision-making under uncertainty. Possibly the closest is *incremental contingency planning*, devised for planning Mars rover missions [3]. At a high level, it operates similarly to CPA. Like our work, this research focuses on identifying parts of the plan that need backup, but considers non-adversarial environments, doesn't face the strategy selection problem, and has different mechanisms and utility for deciding where to add contingent branches. Another important difference is that contingent branches in the Mars rover mission setting can be generated on-demand by a dedicated planner, whereas we, in the absence of such a planner, are restricted to using only those plans that can be found in OLCBP's case base.

In probabilistic planning modeled by MDPs, a related approach is *incremental plan aggregation* [20]. RFF, the planner that implements it, finds a linear plan to the goal, determines states where an agent may end up if action outcomes deviate from this plan, adds linear plans to the goal from these states, and iteratively repeats the process. Like ROLCBP, between iterations RFF estimates the "failure probability" of exiting the partially constructed policy tree. Unlike ROLCBP, RFF assumes no action concurrency, which makes choosing branching points much easier, and generates contingent plans on-demand with a dedicated planner, whereas ROLCBP is restricted to using plans from its case base. Lastly, ROLCBP's failure probability estimates are only a heuristic for adding contingency branches, because it works in an adversarial, possibly non-stationary environment.

Chapter 3

CONCURRENT PLAN AUGMENTATION

To address the drawbacks of existing strategy selection approaches such as OLCBP, we propose a technique called *concurrent plan augmentation* (CPA). CPA's high-level idea is to start with a tactical plan for each available strategy, analyze the ways in which the game situation can deviate from it, and come up with *contingent* plan branches that prevent the deviations from becoming catastrophic, to be executed *in parallel* with the main plan. Repeating these steps iteratively builds a partial *policy* for every strategy. This policy's ability to avoid failure is taken as a measure of the corresponding strategy's robustness; based on it, a strategy selection framework, e.g., OLCBP, can choose a strategy for playing the game. At first glance, CPA is reminiscent of contingency planning, but has several important differences:

- Most contingency planning algorithms assume that catastrophic failures cannot happen, i.e., that it is possible to reach the goal (win the game) starting from any state. Even the few exceptions such as [13] assume that failures can be fully avoided from any state. In contrast, failures impossible to avoid with certainty are common in RTS games, e.g., due to opponents' actions. For this reason, contingency planners can usually rely on replanning, whereas CPA attempts to pre-empt catastrophic events.
- Contingency planning mostly considers settings with no action concurrency, where actions must be executed one after another. CPA relies on concurrency for its operation.
- Unlike in contingency planning, which has no widely accepted measure of policy quality, characteristics of RTS games suggest gauging policy quality with its ability to avoid failures. CPA uses one such measure, a policy's failure probability. Although in a game this measure

is only heuristic, since it depends on the training opponent’s play, our experiments show that it works well in practice.

As described, CPA leaves many details unspecified. How to find contingent plan branches? When should a contingent branch’s execution begin? Should the policy tree be computed eagerly before the start of the game or lazily as the game progresses? Their answers depend on the RTS game and strategy selection framework to which CPA is applied. This thesis combines CPA with the OLCBP framework to produce the *Robust OLCBP* algorithm, which we describe next.

3.1 Robust OLCBP

The operation of CPA with OLCBP in an architecture we call Robust OLCBP (ROLCBP) is broken down into an offline stage (Algorithm 1) and an online stage (Algorithm 2). The offline stage is devoted to training, while the online stage exploits the learned knowledge for strategy and tactics selection. Throughout our explanations, we will be referring to ROLCBP’s pseudocode, denoting line l of Algorithm a as “line $l : a$ ”. In the following section, we illustrate the pseudocode’s operation with an example from *StarCraft*.

Offline stage. At this stage, ROLCBP relies on the OLCBP’s machinery to learn a set of strategies \mathcal{S} and a case base of tactical plans \mathcal{P} based on a human-specified set of subgoal descriptions \mathcal{G} , action descriptions \mathcal{A} , and a set of gameplay logs (line 10:1). The details of this step [16] are beyond the scope of our thesis but, importantly for the rest of ROLCBP, each plan in the resulting case base \mathcal{P} achieves a subgoal from \mathcal{G} assuming some other subset of \mathcal{G} ’s goals has already been achieved. Generally, \mathcal{P} contains several plans that achieve a given subgoal.

After tactical plan extraction, ROLCBP gathers statistics for policy augmentation. Using a simulator for the RTS game to which it is applied, ROLCBP reproduces the initial conditions of each plan $P \in \mathcal{P}$ (i.e., generates a game state in which all subgoals that P assumes to be achieved have indeed been achieved). Then, with the help of game logs, it simulates the execution of P by an OLCBP-based bot against different opponents, recording the number of times each action in P was executed and failed.

```

1 Input:  $M$  — a game map,  $\mathcal{G} = \{g_i\}_{i=1}^n$  — a set of game subgoals,  $\mathcal{A} = \{a_i\}_{i=1}^m$  — a set of
game actions,
2  $\mathcal{L}$  — a set of gameplay logs,
3 AugDegree — augmentation degree
4
5 Output: StarCraft strategy & high-level tactical plan
6
7 // Learn  $\mathcal{S} = \{S_i = (g_{j_1}, \dots, g_{j_{n_i}})\}$  — a set of strategies
8 // (subgoal sequences) &  $\mathcal{P} = \{P_i^g = (a_{j_1}, \dots, a_{j_{m_i}})\}$  —
9 // a set of high-level tactical plans, each for some subgoal  $g$ 
10  $\mathcal{S}, \mathcal{P} \leftarrow \text{OLCBP-Learn}(\mathcal{G}, \mathcal{A}, \mathcal{L})$ 
11
12 foreach  $P_i^g \in \mathcal{P}, a_j \in P_i^g$  do
13    $p_{\text{fail}}(a_j, P_i^g) \leftarrow \text{ProbOfCausingFailure}(a_j, P_i^g)$ 

```

Algorithm 1: ROLCBP: the offline stage

For the purposes of CPA, we define an action a 's failure probability $p_{\text{fail}}(a, P^g)$ as the fraction of times its execution caused a *deviation* from the corresponding tactical plan P^g and forced OLCBP to resort to replanning (lines 12:1-13:1). In reality, OLCBP's replanning is sometimes successful at attaining P_g 's goal g . Therefore, $p_{\text{fail}}(a, P^g)$ is generally a pessimistic measure of a 's quality in P^g . Nonetheless, the overestimation of a 's probability of causing an unrecoverable failure actually helps CPA produce more robust policies.

Online stage. At runtime, ROLCBP follows the MainGameLoop method (lines 3:2-10:2). It begins by selecting a game strategy with the most *robust* tactical plan (line 4:2) as explained below, using the SelectStrategy&Tactics procedure. In addition to returning the tactical plan for the chosen strategy, SelectStrategy&Tactics augments it with concurrent contingent branches. In effect, this method constructs a tactical *policy tree*, although this tree may be only partial. ROLCBP plays according to the chosen tactical plan, launching its contingent branches prescribed by the policy tree, until OLCBP's mechanisms detect a deviation from it (line 6:2).

SelectStrategy&Tactics (lines 12:2 - 25:2) operates by iterating over all strategies S in library \mathcal{S} , letting OLCBP concoct a tactical plan P_S for each strategy, augmenting this tactical plan with contingent branches using CPA, and evaluating the *failure probability* of the policy tree based on this tactical plan. The failure probability of a strategy's augmented tactical plan serves as a proxy

```

1 // See Algorithm 1 for the description of  $M, \mathcal{G}, \mathcal{S}, \mathcal{P}$ , and  $AugDegree$ 
2
3 MainGameLoop( $M, \mathcal{G}, \mathcal{S}, \mathcal{P}, AugDegree$ ) begin
4    $S, P_S \leftarrow \underline{\text{SelectStrategy\&Tactics}}(M, \mathcal{G}, \mathcal{S}, \mathcal{P}, AugDegree)$ 
5   while game has not ended do
6      $a_{fail}, P^g \leftarrow \text{PlayUntilActionFailure}(P_S)$ 
7      $P_S \leftarrow \text{ContingentBranch}[a_{fail}, P^g, P_S]$ 
8     if  $AugDegree == 1$  then  $\underline{\text{CPA}}(M, S, P_S, \mathcal{P})$ 
9     else if  $AugDegree == 0$  then  $\text{OLCBP}(M, S, g, \mathcal{P})$ 
10    else  $AugDegree \leftarrow AugDegree - 1$ 
11
12 SelectStrategy\&Tactics( $M, \mathcal{G}, \mathcal{S}, \mathcal{P}, AugDegree$ ) begin
13   foreach  $S = (g_j, \dots, g_k) \in \mathcal{S}$  do
14     // Compose a tactical plan for strategy  $S$ 
15     // by concatenating several plans from  $\mathcal{P}$ 
16      $P_S = (P_{i_j}^{g_j}, \dots, P_{i_k}^{g_k}) \leftarrow \text{OLCBP}(M, S, \emptyset, \mathcal{P})$ 
17      $LeafBranches \leftarrow \{P_S\}$ 
18     foreach  $d$  from 1 to  $AugDegree$  do
19        $NewBranches \leftarrow \emptyset$ 
20       foreach plan  $P \in LeafBranches$  do
21          $NewBranches \leftarrow NewBranches \cup$ 
22            $\cup \underline{\text{CPA}}(M, S, P, \mathcal{P})$ 
23        $LeafBranches \leftarrow NewBranches$ 
24      $p_{fail}(P_S) \leftarrow \underline{\text{EvalFailProb}}(P_S)$ 
25   return  $S, P_S$  s.t.  $p_{fail}(P_S)$  is the lowest

```

Algorithm 2: ROLCBP: the online stage - part 1

of the strategy's robustness; SelectStrategy&Tactics chooses the strategy with with lowest such probability.

The CPA logic at the heart of SelectStrategy&Tactics is captured in eponymous method (lines 1:2-15:2). CPA analyzes a tactical plan P_S by breaking it down into constituent plans P^g from case base \mathcal{P} . Recall that at the offline stage, ROLCBP estimated failure probabilities for each action in each such plan (lines 12:1-13:1). With the help of these estimates, CPA identifies actions with non-

```

1 CPA( $M, S, P_S, \mathcal{P}$ ) begin
2    $LeafBranches \leftarrow \emptyset$ 
3   foreach  $a \in P^g \in P_S$  s.t.  $p_{fail}(a, P^g) > 0$  do
4      $(g_1, \dots, g_j) \leftarrow$  achieved prefix of  $S$  if  $a$  fails
5     during  $P^g$ 's execution
6     // Compose a plan for  $S$ 's remaining subgoals
7      $P_S^+ \leftarrow OLCBP(M, S, (g_1, \dots, g_j), \mathcal{P})$ 
8     // Schedule  $P_S^+$  to start executing concurrently
9     // with  $P_S$  ASAP after  $g_j$  is achieved but
10    // before  $a$  starts executing, if possible
11    if  $ScheduleConcurrentExec(P_S, P_S^+, g_j)$  then
12       $ContingentBranch[a, P^g, P_S] \leftarrow P_S^+$ 
13       $LeafBranches \leftarrow \{P_S^+\}$ 
14    else  $ContingentBranch[a, P^g, P_S] \leftarrow \emptyset$ 
15  return  $LeafBranches$ 
16
17 EvalFailProb( $P_S$ ) begin
18   if  $P_S$  is  $\emptyset$  then return 1
19    $p_{fail}(P_S) \leftarrow 0$ 
20   foreach  $a_i \in P_S$ ,  $i$  ranging from  $|P_S|$  to 1 do
21      $p_{fail}(P_S) \leftarrow p_{fail}(a_i, P_S) \cdot$ 
22      $\underline{EvalFailProb}(ContingentBranch[a_i, P^g, P_S]) +$ 
23      $+(1 - p_{fail}(a_i, P_S)) \cdot p_{fail}(P_S)$ 
24   return  $p_{fail}(P_S)$ 

```

Algorithm 3: ROLCBP: the online stage - part 2

zero failure probabilities in P_S (line 3:2). For each of these actions a , CPA models the hypothetical situation immediately after a 's failure by determining the "latest" subgoal g_j of strategy S that will remain achieved in that situation (line 5:2). In order to be considered a valid contingency branch for a 's failure, a plan P_S^+ needs to achieve all subgoals of S after g_j . CPA delegates it to OLCBP to find such a plan (line 7:2).

When/if P_S^+ is found, it needs to be scheduled for execution. Intuitively, we would like to start P_S^+ in parallel with P_S as early as possible after g_j is achieved. (By P_S^+ 's construction, we cannot start it before achieving all of S 's subgoals up to g_j .) We would also like to schedule it before a , because otherwise in case of a 's failure, while OLCBP is looking for another plan, for some time

the player's units will not know what to do (and hence will be extremely vulnerable). Whether P_S^+ 's execution can start in this time window depends on game resource constraints. For example, it may be impossible to train several kinds of units concurrently without enough minerals. Scheduling P_S^+ is again delegated to OLCBP (line 11:2). If it turns out to be possible, P_S^+ is designated as P_S 's contingency for a 's failure (line 12:2). Otherwise, a 's failure remains unaccounted for (line 14:2). Ultimately, CPA's inability to add contingency branches to a tactical plan may point to the plan's, and hence the strategy's, inherent weaknesses.

A single CPA invocation provides back-ups for the failure of a given tactical plan, but not for failures of the back-ups themselves. The latter would require applying CPA recursively to the contingent branches. SelectStrategy&Tactics provides the option of doing that via the *augmentation degree* (*AugDegree*) parameter. The first-degree augmentation applies CPA to a strategy's main tactical plan only, the second-degree one — to that plan and its contingent branches, and so on (lines 17:2-23:2). Intuitively, higher-degree augmentations reveal complete contingency structure and provide a more thorough assessment of a strategy's robustness. Independently of augmentation degree, evaluation of a strategy's failure probability can be done in a single pass. This is due to the fact that an augmented tactical plan is always a tree. Computing its failure probability is a matter of recursively propagating information from the tree's leaves to its root (lines 17:2 -24:2).

If augmentation degree is at least 1, deviating from the main tactical plan of the chosen strategy in ROLCBP's MainGameLoop likely does *not* cause replanning, as it would in OLCBP. This is because a contingent branch for this failure has already been launched, if this was possible given the game resource constraints. This branch becomes the main execution plan (line 7:2). However, this plan may be vulnerable: recall that, depending on augmentation degree, SelectStrategy&Tactics may not have scheduled any contingencies for it. Therefore, once ROLCBP switches to this plan, if necessary it augments it with contingent branches (line 8:2), to avoid hasty replanning in case of another failure.

Practical considerations. In its computations, ROLCBP iterates over all strategies in its library \mathcal{S} , and for each of them builds a policy tree whose size is influenced by the augmentation degree parameter. Thus, the size of \mathcal{S} and augmentation degree are two major factors affecting ROLCBP's

performance. Fortunately, the following observations imply that sensible values of these quantities are small in practice:

- *RTS games typically have few viable strategies, and OLCBP's learning from human demonstration naturally identifies them.* Indeed, although in theory there are many possible goal orderings for any scenario, human players tend to use only those belonging to a small set of successful ones. Since ROLCBP's strategy library \mathcal{S} is populated by strategies extracted from human gameplay logs, it ends up containing only a few entries. This makes it possible to analyze each one of them in a short time at the beginning of the game.
- *High-degree augmentation does not pay off.* While a high-degree augmentation gives a better estimate of a strategy's robustness, and hence enables better-informed strategy choice, its computational cost grows exponentially in augmentation degree. ROLCBP can afford some deliberation time at the beginning of the game, but excessively long delay before committing to a strategy heightens the risk of the opponent's attack against the player's unprepared infrastructure. Moreover, most of the contingent branches in a highly augmented policy tree will end up not getting followed, wasting computational effort invested into them. Last but not least, high-degree contingencies are increasingly difficult to schedule because of the multiple resource conflict resolutions OLCBP needs to perform. In all our tests, *first-degree* augmentation, which applies CPA only to the main tactical plan of a policy, has provided the best balance of computational cost and strategy evaluation quality. Therefore, we set $AugDegree = 1$ in all our experiments.

Chapter 4

EXAMPLE

Suppose that after offline training (Algorithm 1), ROLCBP playing the game of *StarCraft* starts by analyzing a strategy whose subgoal ordering is depicted as a sequence of dark-gray rectangles on the *left* side of Figure 4.1. Suppose further that OLCBP, which ROLCBP uses to find a tactic for this strategy, composes a plan given by the white rectangles inside the dark-gray rectangles. Namely, the plan is $[GatherPrimaryResource, BuildBarracks, TrainMarine, AttackUnit]$, where several *TrainMarine* and several *AttackUnit* actions are executed in parallel. As Figure 4.1 shows, this tactical plan consists of primitive plans for achieving each subgoal given the previously achieved ones. For the simplicity of the example, each such subplan of the main tactical plan has only one step.

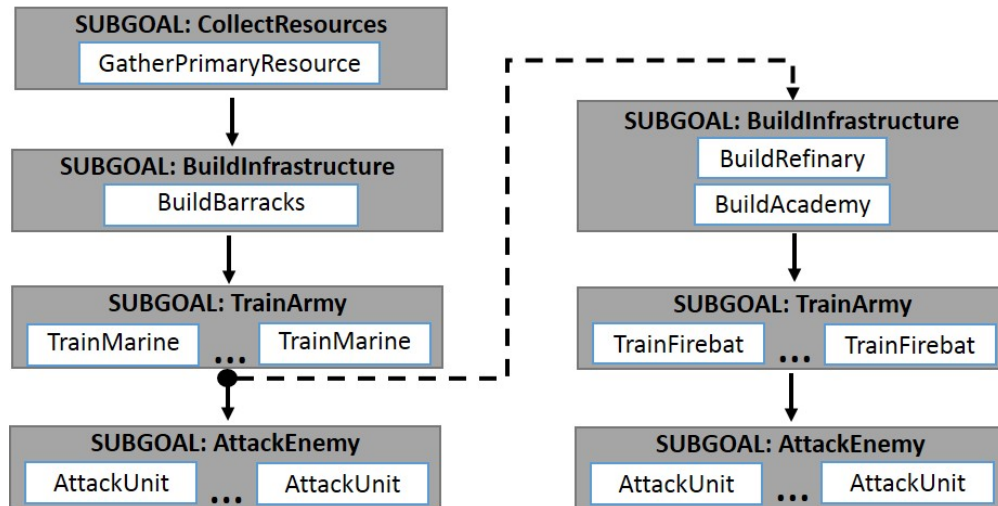


Figure 4.1: A tactical plan (left) with a contingent branch (right).

The analysis of this plan involves computing the failure probabilities of the steps that compose it. Imagine that during the offline stage (Algorithm 1), ROLCBP established that the first three

subgoals, up to and including *TrainArmy*, are always achieved successfully, e.g., because at the start of the game the opponent has too little time to reach the player's base and disrupt the plan. However, ROLCBP may have discovered that the last subgoal of the plan, *AttackEnemy* is difficult to reach: the marine units, though quick to train, are fairly weak, and their attack often fails. When it does, the *TrainArmy* subgoal also gets “unachieved”, because the army is destroyed in the attack.

Having performed this analysis, ROLCBP decides to inject a concurrent contingent plan before going for the *AttackEnemy* subgoal. This contingent branch needs to achieve all subgoals that remain unachieved if the main plan fails, i.e., *TrainArmy* and *AttackEnemy*. One such plan is shown on the right side of Figure 4.1. It trains an army of units more powerful than marines, which requires it to build additional infrastructure. This infrastructure takes a lot of resources, which are unavailable while the main army of marines is being trained, so it schedules this branch for execution right after the marines are trained and ready to attack. If the marines fail, a more potent army of firebats will be available to guard the base and attack the enemy soon afterwards.

After augmenting the main tactical plan, ROLCBP evaluates its failure probability (lines 17:3-24:3). If more strategies are available, ROLCBP analyzes them too and picks the least likely to fail.

Chapter 5

EXPERIMENTAL EVALUATION

We evaluated CPA by implementing it as part of ROLCBP in software bots for *StarCraft* and *Wargus* (an open-source clone of *Warcraft II*). Most of our experiments focus on *StarCraft*, since it is the more complicated of the two and has a better-developed publicly available experimentation infrastructure that greatly speeds up running tests on it. At the same time, *Wargus* shares its game engine *Stratagus*, the mechanism responsible for its planning capabilities, with a range of other RTS games such as *Aleona's Tales*, so our successful experiments on *Wargus* imply CPA's effectiveness on all these games as well and attest to the universality of our approach.

Our experiments address the following questions: (1) How does concurrent policy augmentation affect the bot's performance if used purely at the tactical level when a strategy is fixed, by increasing the tactical plan's robustness? (2) How does CPA affect the bot's performance if used in both at the tactical and strategic levels in ROLCBP, by guiding the strategy selection process? (3) How successfully does CPA deal with opponents who are significantly different from those it plays during the offline training stage (Algorithm 1)? Note that demonstrating the overall superiority of our bot over state-of-the-art bots for *StarCraft* or *Wargus* is *not* an objective of our experiments. Building such a system would require extensive research into components other than strategy and tactics selection modules, which is beyond the scope of this work. Instead, our evaluation uses simple algorithms in bot modules such as those responsible for low-level control, and concentrates on the influence of CPA on strategic and high-level tactical performance.

Experimental setup. We gauge the performance of ROLCBP and the baselines (built-in AI, OLCBP, or both, depending on the experiments) by running many games against an opponent and measuring the *win rate* — the percentage of won games. Win rate subsumes all other gameplay characteristics of potential interest, such as speed or memory usage: if the player's AI is unaccept-

ably slow or memory usage too high, this is reflected in a low win rate. We identify statistically significant performance differences in win rate with two-tailed z-test for equality of proportions at the 95% confidence level.

For *StarCraft*, we experiment on three maps frequently used in *StarCraft* tournaments: Blood Path, Binary Burghs, and Bottleneck. Each *StarCraft* map is characterized by the size of its grid, which is loosely correlated with the map’s difficulty; Binary Burghs Path has size 64x64, Binary Burghs – 96x96, and Bottleneck – 128x128 cells. For *Wargus*, we use maps of similar sizes: Harrow (62x62), Hills of Glory (96x96), and Dust Storm (128x128). In all figures, map size increases along the x -axis. As mentioned previously, a *StarCraft* player can control one of three races and in *Wargus* — one of two, each with its own characteristics. Our bot always plays the *Terran* race in *StarCraft* and *orcs* in *Wargus*. For offline acquisition of strategies and tactical plans, our bot uses human gameplay logs. For *StarCraft*, they are available at <http://www.teamliquid.net/replay/>.

Our implementations of ROLCBP and OLCBP are in C++. The two differ only in strategy selection, and share all other components (modules for extracting strategies and tactical plans from game logs, low-level control, etc.) The augmentation degree parameter in ROLCBP was set to 1 (see the discussion at the end of the Robust OLCBP section). The experiments were run on an Intel Core i7 2.4GHz CPU with 12GB RAM under Windows 8.1.


Effects of Tactical-Level Augmentation. In this experiment, conducted on *StarCraft*, we measure how the win rate of individual strategies changes when first-degree CPA is applied to their main tactical plans. The results are shown in Figure 5.1. We focus on two strategies in this experiment. *StarCraft* has three different races, and we measure each strategy’s win rate against each of the three races separately (Figures 5.1a, b, and c). Our bot always plays the Terran race.

To measure a strategy’s win rate against a given race (e.g., Zerg) on a given map, we first run the offline training stage of ROLCBP and OLCBP on logs of games played against *the two other* races (e.g., Protoss and Terran) controlled by *StarCraft*’s built-in AI. Then we evaluate the strategy’s performance against the target opponent race controlled by the built-in AI by comparing the win rates of OLCBP and ROLCBP when this strategy is the only one available in their strategy library. OLCBP and ROLCBP each ran every strategy 100 times against each race on each map.


The two strategies we use in this experiment are known among human *StarCraft* players as *rush attack* and *mid-game attack*. We chose these strategies because our bot judged them to be the most powerful ones it managed to extract from the game logs: in every game in the strategy selection experiments, whose results are presented later, our bot invariably choose one of these two strategies as the best one. Figure 5.1, which summarizes the results on this experiment, reveals several patterns:

- *For most combinations of strategies, maps, and train-test splits in the experiment, ROLCBP's win rate is significantly higher than OLCBP's, i.e., CPA significantly improves strategies' robustness.* This happens because OLCBP resorts to replanning when a strategy's execution goes awry. Replanning is expensive, and if it does not yield a quick solution in a critical situation, the bot loses the game. CPA reduces the need for replanning in critical situations.
- *If a strategy is inherently unsuitable to a given map, applying CPA to it does not help.* For example, all tactical plans for the rush attack strategy on the Bottleneck map have fatal flaws that prevent them from winning even after CPA.
- *CPA can improve a strategy's tactical plan qualitatively, enabling it to win games that its non-augmented version would always lose.* For example, without CPA mid-game attack on the Bottleneck map always fails, while with CPA it is able to win a percentage (albeit small) of games.


Effects of CPA on Strategy Selection. Figures 5.2 and 5.3 show the main experimental results of this thesis, which testify to the benefit of CPA for overall strategy selection in *StarCraft* and *Wargus*, respectively. For *StarCraft*, the experimental setup is the same as for the tactical-level augmentation experiments from the previous subsection, but now we allow ROLCBP and OLCBP to consider *all* strategies they extract from the game logs, instead of being restricted to any specific one. Thus, the quality of strategies' augmented tactical plans determines the strategy choice for the game. While ROLCBP makes this choice by minimizing strategies' empirical failure probability,



(a) Training vs Terran/Zerg, playing vs Protoss




(b) Training vs Terran/Protoss, playing vs Zerg




(c) Training vs Zerg/Protoss, playing vs Terran

Figure 5.1: Effect of CPA on the tactical performance of the rush attack (**R**) and mid-game attack (**M**) strategies in *StarCraft*. The advantage of *ROLCBP* is statistically significant in all experiments except for rush attack on the *Bottleneck* map.


OLCBP, our baseline, selects strategies heuristically based on the characteristics of the map and the races the players have picked, breaking ties randomly. In particular, for all *StarCraft* maps in our experiments both the rush and the mid-game attack strategies turn out to match OLCBP’s criteria, so it decides randomly between them.



(a) Training vs Terran/Zerg, playing vs Protoss



(b) Training vs Terran/Protoss, playing vs Zerg



(c) Training vs Zerg/Protoss, playing vs Terran

Figure 5.2: Effect of CPA on overall performance in *StarCraft*. *The advantage of ROLCBP is statistically significant in all experiments.*

As the plots in Figure 5.2 demonstrate, CPA gives ROLCBP a statistically significant advantage in strategy selection. At the same time, similar to the tactical-level experiments, they show that if




Figure 5.3: Effect of CPA on overall performance in *Wargus*. *The advantage of ROLCBP is statistically significant in all experiments.*

no strategy learned from the logs is very suitable for a map, then ROLCBP cannot fundamentally change AI’s performance — this is the case in the *Bottleneck* scenario. Fortunately, the fix for this issue is as simple as giving ROLCBP a more diverse set of logs. To ascertain this, we re-ran Figure 5.2’s experiments after adding logs with plays of a third strategy, known among *StarCraft* players as Fabian, in addition to rush and mid-game attack discussed earlier. This strategy is not a good option on Blood Bath and Binary Burghs, and ROLCBP indeed does not choose it there, so the results on these two maps remain as in Figure 5.2. The performance on Bottleneck, however, shown in Figure 5.4, improves by a lot across each of the three train-test splits and allows ROLCBP to beat built-in AI most of the time.




Figure 5.4: Effect of CPA on overall performance in *StarCraft* on the Bottleneck map with the Fabian strategy. In each case, ROLCBP trains on two races and plays versus the remaining one. *The advantage of ROLCBP is statistically significant in all experiments.*

The role of plan augmentation in the decision to choose the Fabian strategy for this map is vital. If failure probabilities of the main tactical plans for mid-game attack, rush attack, and Fabian are evaluated (lines 17-24 of Algorithm 2) against training opponents *before* plan augmentation, rush attack’s failure probability appears the lowest, implying that it is a better strategy. (Figure 5.1 does

not reflect this, because it shows win rates of strategies against *testing* opponents.) However, this map's dynamics make deviations from rush attack's tactical plan hard to recover from. Performing CPA and re-evaluating the augmented plans' failure probabilities reveals this and leads to the correct decision to play the Fabian strategy.

Wargus testing framework limitations has limitations like:

- The game engine is not compatible with BWAPI so the evaluation is manual not automated.
- The availability of the gameplay logs are limited.
- The game has small variations (two races) only.


Because of that we measure ROLCBP's and OLCBP's win rates against the built-in AI playing the orc race only. ROLCBP and OLCBP, which also controlled orcs, each played 20 games on each map. Results on *Wargus* (Figure 5.3) confirm our findings from *StarCraft*: CPA implemented in ROLCBP drastically improves strategy selection compared to OLCBP. On all *Wargus* maps we have tested, ROLCBP wins over half of the games against built-in AI, whereas OLCBP does not on any single map. Last but not least, since many other RTS games share the *Stratagus* game engine with *Wargus*, we can expect ROLCBP to enjoy similar advantage on them as well.

CPA's performance against unfamiliar opponents. Besides its ability to improve gameplay at strategic and tactical levels, the experiments above hint at CPA's robustness to previously unencountered opponents. As evidence of this, note that in the aforementioned setups ROLCBP's win rate was evaluated on opponent races against which it did not play during training. Nonetheless, both in training and in testing ROLCBP's opponents were controlled by the same (built-in) AI, which prompts a question: does ROLCBP perform as well when the enemy AIs in training and testing are different?


In our final set of experiments on *StarCraft*, we answer this question positively. We consider two opponent AIs: *StarCraft*'s built-in AI and OLCBP. For each *StarCraft* race, we train ROLCBP on game logs against that race controlled by the built-in AI, and measure its win rate against that race played by OLCBP. We also experiment with the case when built-in AI's and OLCBP's roles

are switched. As before, we let ROLCBP play 100 games for each race, map, and train-test AI combination.


Figure 5.5 presents ROLCBP’s win rates in this experiment. With the exception of training against OLCBP and playing against built-in AI on the Bottleneck map, OLCBP consistently and convincingly outperforms its opponents.



(a) Training vs Terran/Zerg, playing vs Protoss



(b) Training vs Terran/Protoss, playing vs Zerg



(c) Training vs Zerg/Protoss, playing vs Terran

Figure 5.5: Training against OLCBP and playing against static AI and vice versa.

Nonetheless, ROLCBP’s win rate noticeably depends on the specific combination of train and test opponent AIs; playing OLCBP after training on the built-in AI usually gives better results than vice versa. This is not surprising. OLCBP’s inferior win rates against the built-in AI in Figure

5.2c suggest that the OLCBP bot is weaker. ROLCBP's win rate suffers when it trains against this weaker AI and then faces the stronger built-in one. The converse also holds: note that the training provided by the built-in AI is so powerful that ROLCBP outmatches OLCBP on Bottleneck (Figure 5.5) despite the fact that, as already mentioned, all of the strategies in its library are fundamentally unsuitable for this map.

Chapter 6

CONCLUSION

We have presented concurrent policy augmentation (CPA), a technique for addressing shortcomings of existing methods for automatic strategy and tactics selection in real-time strategy games. These existing methods, such as OLCBP, treat strategy selection as a hierarchical deterministic planning problem with a possibility of replanning. Replanning in critical situations often fails, leading to a game loss. In contrast, CPA analyzes tactical plans for available strategies to identify their potential failure points. It attempts to mitigate the consequences of these failures by adding contingent plan branches to be executed concurrently with the main plan. We implemented CPA in an OLCBP-based architecture ROLCBP that selects strategies by evaluating the failure probabilities of their augmented tactical plans. Our experiments on popular and complex RTS games *StarCraft* and *Wargus* show that CPA gives ROLCBP significant advantages in win rate compared to OLCBP by improving both tactical and strategic decision-making. It also lets ROLCBP outperform opponents with different behavior than those ROLCBP encounters during training. In the future, we plan to improve our ROLCBP implementation and turn it into a bot competitive with entries in the *StarCraft AI* competition. We have discovered that despite superior strategic performance, our bot's low-level unit control often causes it to lose skirmishes (and hence entire games) that state-of-the-art bots would easily win. Low-level control issues account for most of the games ROLCBP lost in this thesis's experiments. Thus, despite being out of this thesis's scope, low-level control is a clear area of improvement. We believe that the usefulness of CPA extends beyond case-based planning techniques. The hypothetical advent of planners capable of efficiently generating winning trajectories for RTS games would remove CPA's dependence on plans stored in the case base and would greatly increase CPA's potential.

BIBLIOGRAPHY

- [1] Robin Baumgarten, Simon Colton, and Mark Morris. Combining AI methods for learning bots in a real-time strategy game. *Int. J. Computer Games Technology*, 2009, 2009.
- [2] Michael Buro. Call for AI research in RTS games. In *In Proceedings of the AAAI Workshop on AI in Games*, pages 139–141. AAAI Press, 2004.
- [3] Richard Dearden, Nicolas Meuleau, Sailesh Ramakrishnan., David E. Smith, and Rich Wastington. Incremental Contingency Planning. In *ICAPS'03*, 2003.
- [4] Ibrahim Fathy, Mostafa Aref, Omar Enayet, and Abdelrahman Al-Ogail. Intelligent online case-based planning agent model for real-time strategy games. In *ISDA*, pages 445–450, 2010.
- [5] Maria Fox, Alfonso Gerevini, Derek Long, and Ivan Serina. Plan stability: Replanning versus plan repair. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *ICAPS*, pages 212–221. AAAI, 2006.
- [6] Dan Fu and Ryan Houlette. The ultimate guide to FSMs in games. *AI Game Programming Wisdom 2*, 2003.
- [7] Kristian Hammond and Running Head. Case-based planning: A framework for planning from experience. *Cognitive Science*, 14:385–443, 1990.
- [8] Hai Hoang, Stephen Lee-urban, and Hector Muoz-avila. Hierarchical plan representations for encoding strategic game ai. In *In Proc. Artificial Intelligence and Interactive Digital Entertainment Conference (AIIDE-05)*. AAAI Press, 2005.
- [9] John-Paul Kelly, Adi Botea, and Sven Koenig. Offline planning with hierarchical task networks in video games. In *AIIDE'08*, 2008.
- [10] I. Little and S. Thiébaux. Probabilistic planning vs replanning. In *Workshop on International Planning Competition: Past, Present and Future (ICAPS)*, 2007.
- [11] Kinshuk Mishra, Santiago Ontañón, and Ashwin Ram. Situation assessment for plan retrieval in real-time strategy games. In Klaus-Dieter Althoff, Ralph Bergmann, Mirjam Minor, and Alexandre Hanft, editors, *ECCBR*, volume 5239 of *Lecture Notes in Computer Science*, pages 355–369. Springer, 2008.

- [12] Matthew Molineaux, Matthew Klenk, and David W. Aha. Goal-driven autonomy in a navy strategy simulation. In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2010, Atlanta, Georgia, USA, July 11-15, 2010*, 2010.
- [13] Christian Muise, Vaishak Belle, and Sheila McIlraith. Computing contingent plans via fully observable non-deterministic planning. 2014.
- [14] Jorge Muñoz, Germán Gutiérrez, and Araceli Sanchis. A human-like TORCS controller for the Simulated Car Racing Championship. In *Proceedings 2010 IEEE Conference on Computational Intelligence and Games*, pages 473–480, August 2010.
- [15] Santiago Ontañón, Kinshuk Mishra, Neha Sugandh, and Ashwin Ram. Learning from demonstration and case-based planning for real-time strategy games. In *Soft Computing Applications in Industry*, pages 293–310. 2008.
- [16] Santiago Ontañón, Kinshuk Mishra, Neha Sugandh, and Ashwin Ram. On-line case-based planning. *Computational Intelligence*, 26(1):84–119, 2010.
- [17] Santiago Ontañón, Gabriel Synnaeve, Alberto Uriarte, Florian Richoux, David Churchill, and Mike Preuss. A survey of real-time strategy game AI research and competition in starcraft. *IEEE Trans. Comput. Intellig. and AI in Games*, 5(4):293–311, 2013.
- [18] Louise Pryor and Gregg Collins. Planning for contingencies: A decision-based approach. *Journal of Artificial Intelligence Research*, 4:287–339, 1996.
- [19] Earl D. Sacerdoti. The nonlinear nature of plans. In *Proceedings of the 4th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'75*, pages 206–214, San Francisco, CA, USA, 1975. Morgan Kaufmann Publishers Inc.
- [20] Florent Teichteil-Königsbuch, Ugur Kuter, and Guillaume Infantes. Incremental plan aggregation for generating policies in MDPs. In *AAMAS'10*, pages 1231–1238, 2010.
- [21] Ben George Weber, Michael Mateas, and Arnav Jhala. Building human-level ai for real-time strategy games. In *AAAI Fall Symposium: Advances in Cognitive Systems*, volume FS-11-01 of *AAAI Technical Report*. AAAI, 2011.

VITA

Abdelrahman Elogeel is a graduate student at University of Washington Tacoma and Software Engineer at Microsoft Application Platform group working on the next generation of cloud computing platforms at Microsoft. His research interests are automated planning and decision under uncertainty applied to video game AI. Contact email is elogeel@uw.edu.