



**DEVELOPING GAME AI FOR THE  
REAL-TIME STRATEGY GAME  
STARCRAFT**  
Post-mortem

Spring 2011

**Aleksandar Micić, Davíð Arnarsson, Vignir Jónsson**  
Computer Science B.Sc.

Leiðbeinandi: Yngvi Björnsson  
Prófdómari: Henning Úlfarsson

T-619-LOKA  
Tölvunarfræðideild

## What went well

The agent architecture turned out to work quite well in this situation. It allowed us to try out new things and swap out features with relative ease. The group spirit and morale was high throughout the entirety of the project.

## What didn't

Planning and organization during the semester could've been better. This became apparent in the last few weeks of the project when the workday got progressively longer as the deadline got closer. The original project plan assumed 3 partial man months for the AI Implementation and 2 partial man months for empirical evaluation of the AI. The implementation took longer than expected due to a number of factors.

The Brood war API did not work quite as well as we had hoped, it gave us quite a bit of difficulty when it came to having units attack one another. The problem was threefold:

Firstly, each and every unit has an attack animation, which is not provided by the API. The problem with this was that if you gave a unit any order before it finished its attack animation, it would cancel the attack, yet have its attack go on cooldown.

Secondly, the range of the unit wasn't exactly true, the dragoons were able to shoot units that stood quite a few pixels out of their supposed range given to us by the API.

Thirdly, when told to attack an enemy unit the dragoons would take several step forward before they started shooting.

We had to try and compensate for all these things, at which we mostly succeeded.

Using reinforcement learning is hard when dealing with multi-agent environments like was seen in a scenario in which many agents were required to work together in order to win. The agents showed next to no improvement over time. It all comes down to how one specifies the environment and rewards/penalizes the agents based on their behavior. A lot of time was spent fine tuning these things in the multi-agent environments to try get a more immediate improvement, but the state space just proved to massive, but in the end we just didn't have the computational power needed.

## What Surprised us

Writing the report took more effort than we expected. It was quite surprising to see the agent learn to efficiently dodge attacks after only a handful of games played. We expected it to take much longer for such behavior to emerge.

## Things we learned

We gained a deeper understanding of what it takes to write a research report. Working independantly and having to organize ourselves turned out to be a challenging endeavour.