Team 5: Force Protection with ITSim (II): Base Protection Against Ballistic Weapons

TEAM 5 MEMBERS

Pfeiffer, Volker M
IT-AmtBw, Germany

Mausberg, Niklas
Mies, Christoph
Hörcher, Sascha
Wolters, Benjamin
Fraunhofer IAIS, Germany

Sanchez, Paul PhD
Naval Postgraduate School, US

Schubert, Johan PhD,
Swedish Defence Research Agency, Sweden

Ruiz, Juan Dr
Modelling and Simulation Coordination Office, NATO

INTRODUCTION

ITSim is a general purpose simulation system for decision-support. It focuses on the simulation of coherent processes and provides additional methods for examining optimization tasks within the broader range of tasks of the German Armed Forces, the Bundeswehr. Modern warfare scenarios are dominated by asymmetric threats with complex non-linear interdependencies and interrelations that traditional techniques of analysis are insufficient to capture. For example, it is often hard to determine whether located humans are opponents (red) or just civilians (neutral). We use a base protection scenario and evaluate several active defense options against small teams firing improvised ballistic rockets at the camp. Based on the scenario introduced at the International Data Farming Workshop 18 (IDFW18) [1], the Force-Protection domain is enhanced to investigate further issues.

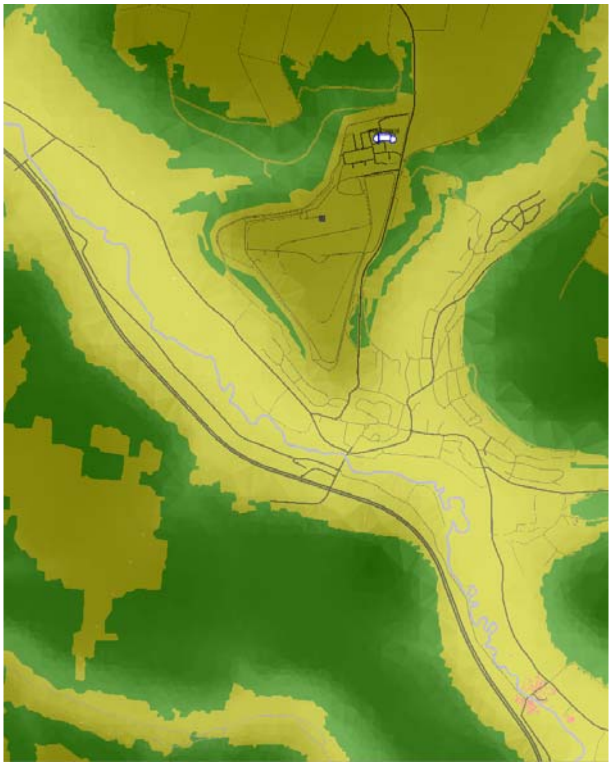

Figure 1: Base in 3D terrain with semantic information

The investigated scenario analyses exactly that aspect by using 3D terrain augmented with semantic information provided by the German Armed Forces. The data is not modeled but imported from an official data source. During this workshop, we want to answer two questions:

1. Does the consideration of semantic information (see below) result in a statistically significant change of the investigated *Measure of Effectiveness (MoE)*?
   Note that the consideration of semantic information will result in a more realistic environment model. But this more precise model raises costs in computation and modelling time. If the *MoE* is not affected by this additional effort, we can omit semantic information for this scenario.
2. Does an optimization of blue emplacements in order to increase the observed area (see below)

result in a statistically significant improvement of blue's success w.r.t. the investigated *MoE*?

Of course, we expect the answer to this question to be true. It is interesting to investigate the importance of the optimization criterion w.r.t. the investigated *MoE*. If the criterion is not important, the *MoE* will not be affected. In future, we are interested in performing several optimizations according to several criteria in order to determine the most important ones.

SCENARIO

Figure 1 depicts the investigated scenario. A blue base is located in 3D terrain with additional semantic information. Dark regions mark high terrain elevation whereas bright areas denote lower terrain. Thus, the blue base is located on a hill. The semantic information is attached directly to the terrain data. Basically, it is a classification of the terrain, including rivers, buildings, different types of wood, different types of streets and flat terrain. In figure 1, woods are visualized as green areas and plain terrain is depicted in yellow. Additionally, the dark roadmap and the blue river can be recognized. In the lower right part of figure 1, many buildings colored in red are visible. Two towers equipped with cameras are used to observe the surrounding area of the camp. They are visualized by tactical icons in the upper part of figure 1. During the course of the scenario, some Red will approach the base in order to attack it with ballistic weapons.

The key idea is that the opponents cannot be detected as Red until they start to prepare their attack. Thus, the whole approach time cannot be used to prevent the attack. After the configured preparation time, the opponents launch *n* missiles (with reload time in between) and flee afterwards.

The scenario's analysis is divided into two phases. The first one is a static classification and the second one is a simulation capturing the dynamics of the strategies.

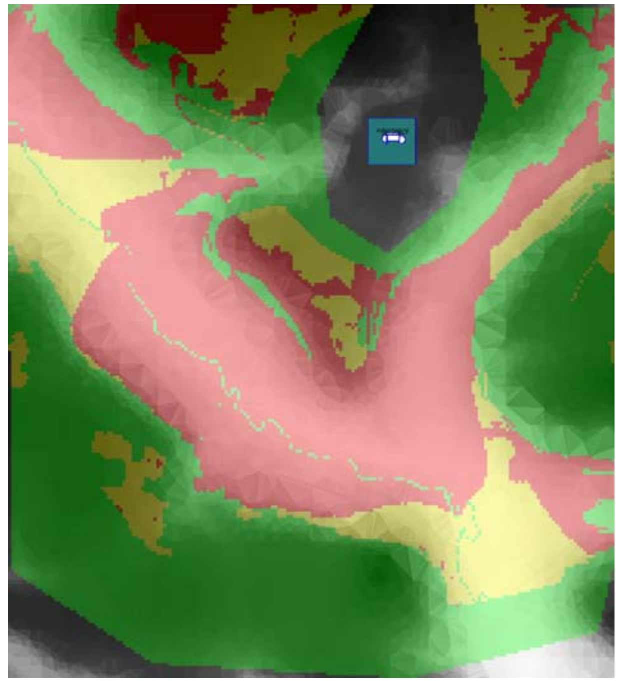

Figure 2: Result of classification considering semantics

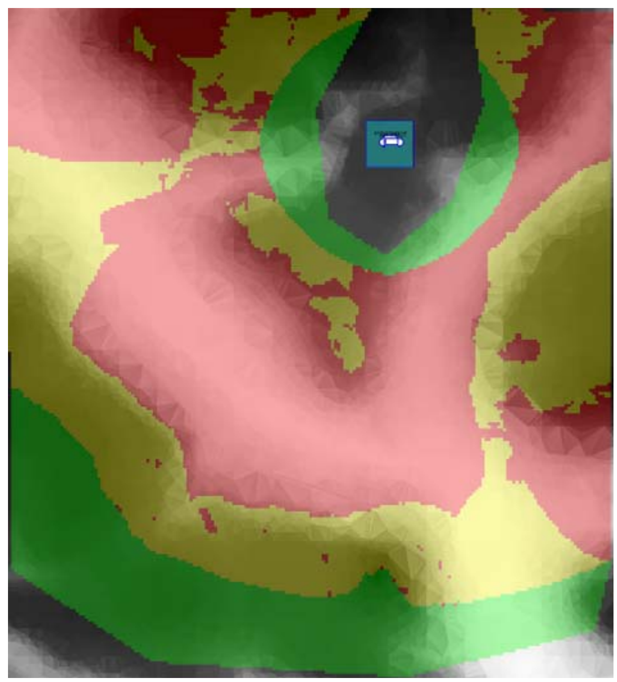

Figure 3: Classification without considering semantics

Static Classification

Before the scenario is simulated dynamically, a static classification is performed. Two important measures are vital for the strategies: Ballistic threat and line-of-sight. Areas from which the base can be attacked by ballistic weapons are called ballistically threatening. The muzzle velocity of the weapon defines its maximal distance. The terrain defines if there is an angle that results in a flight trajectory such that the base can potentially be hit. Considering the semantic information, we defined that it is impossible to attack ballistically from rivers, buildings and woods. Thus, the ballistic threat strongly depends on the given terrain and semantic augmentation. The line-of-sight denotes which areas can be observed by the cameras in the base. These cells are called observable.

In order to perform the classification, the area around the base is gridded. Afterwards, every cell, i.e. grid element, is checked if it is ballistically threatening and observable. Note that the terrain itself is not gridded but based on precise vector data. According to that classification, three cases exist:

- *Green*: A cell is not ballistically threatening, i.e. the base cannot be attacked from that cell. The Blue don't have to worry about that cell. Therefore, the cell is colored green.
- *Yellow*: A cell is ballistically threatening and observable. Thus, the base can be attacked from that cell and there is a line-of-sight to the Blue. The attackers can be detected while they prepare their attack. The cell is colored yellow.
- *Red*: A cell is ballistically threatening and not observable. Thus, the base can be attacked from that cell and there is no line-of-sight to the base. The attackers cannot be identified while they prepare their attack. This is the worst case for the blue forces and the cell is colored red.

If the semantic information is taken into account in the grid classification, the ballistic reachability is restricted by not allowing shots from woods, rivers and buildings. This reduces the opponent's area of operation. This might have an effect on the optimization, because a smaller area needs to be monitored. Thus, fewer emplacements might be needed in order to establish a certain success for Blue. The result of this classification is depicted in figures 2 and 3. Figure 2 shows the classification with considering the semantic information, figure 3 without. Considering the semantics changes the number of green, i.e. not ballistically threatening, cells from *21.8%* to *48.2%*. Thus, roughly *26.4* percent of the cells are not longer threatening for Blue. This consideration of ballistic semantics results in a limitation for Red's area of operation. By comparing figures 1, 2 and 3, we can recognize that the woods (depicted green in figure 1) are ballistically threatening in figure 3 (classification without semantics) but are not in figure 2 (classification with semantics). The same holds for the river, which can be recognized in figure 2. Note that when considering a higher grid resolution, i.e. smaller grid cells, the whole calculation gets more accurate and the river can be recognized more clearly.

Simulation of the Strategies

In order to be able to compare the semantics impact on Blue's success, several strategies have been evaluated against a given red behavior. This kind of analysis may give interesting hints to support the defending of the base. The red strategy is fixed in all experiments. It consists of the following steps:

1. *Generation*: The units are generated uniformly distributed outside the base. Their affiliation is neutral, i.e. they cannot be detected as hostile.
2. *Approach*: A yellow or red cell (i.e. a ballistic attack is possible from that cell) is selected and moved to. The unit is still not detectable as hostile.
3. *Preparation*: Two cases exist. If the attacker can detect any blue unit it gets discouraged and flees. Otherwise it starts to prepare its attack. From that point in time, it can be detected as hostile by Blue. As soon as the blue force is detected by the red unit, the Red aborts its preparation and flees. Note that the cameras' sight range is much higher than the one for regular ground troops including red attackers and blue defenders.
4. *Attack*: The Red starts to fire a previously defined number of projectiles (intended shot number) at the base. From this point in time, the attacker is detected as hostile by the blue defenders if it has not already been. Between the shots, the attacker has to reload. Afterwards, it flees.

Currently, Blue has three different strategy options to prevent ballistic bombardment at the base:

- Pursue from Base (*PfB*): A blue *Quick Reaction Force (QRF)* is located inside the base and pursues the red attackers as soon as they have been detected. The attacker can be observed by the cameras or they reveal themselves by shooting projectiles at the base.

- *Camouflaged Emplacements (CE)*: Camouflaged spotters are located outside the base. They can detect the Red but not vice versa. As soon as the red units are located, their position is reported to the base and the *QRF* starts the counterattack at the Red.
- *Show of Forces (SoF)*: Patrols move around the base. They can detect the Red and can also be detected by these. If any red force is located, the nearest patrol starts a counter attack. Note that there is no *QRF* in the base as in the other strategies. A red opponent detecting an approaching patrol is going to flee.

The semantic attributes influence the routing of the units. According to the different road types, a restriction to the top-speed along the roads is modeled. Thus, Blue's as well as Red's speed will be reduced on certain routes. One might expect less success for Blue, but this is not clear since both parties drive slower.

As *MoE* we defined the percentage of *PreventedShots* defined as the ratio of prevented shots with .respect to the number of intended shots. For example, if the red attacker intended to shots two times and has been neutralized after one shot, *PreventedShots* is *0.5*. Additionally, we measured the following *MoE*s, which are not further regarded in this report due to space limitations:

1. *PreventedShots*: The number of prevented shots at the base. This happens if the attacker is neutralized or discouraged before the attack is started.
2. *PreventedAllShots*: This binary *MoE* is true, when all intended shots have been prevented.
3. *NeutralizedAttacker*: The number of neutralized attackers.

In order to distribute the emplacements in the *CE* vignette, an optimization feature is used. The user has to define several possible positions and ITSim distributes up to *n* spotters over these possible positions. Figures 4 and 5 depict the possible emplacement position, the green icons. The visibility of the *3* active emplacements (the larger icons) is currently considered in the classification.

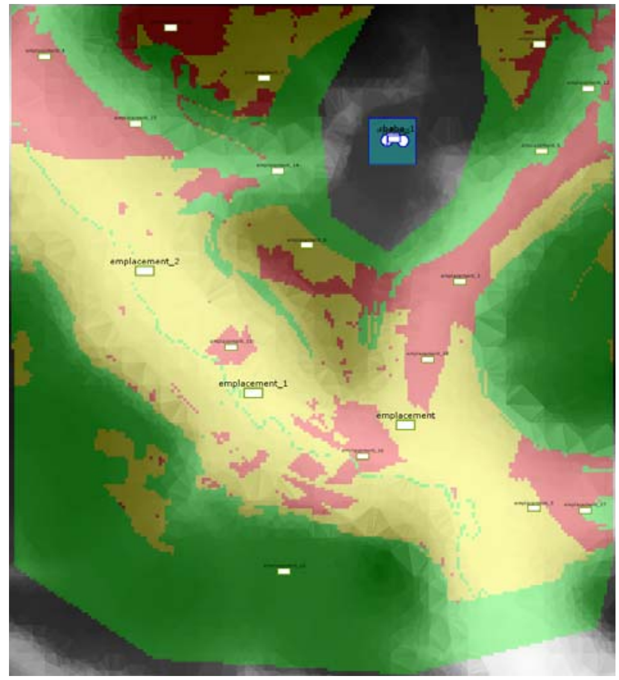

Figure 4: Classification with *3* active emplacements considering ballistic semantics

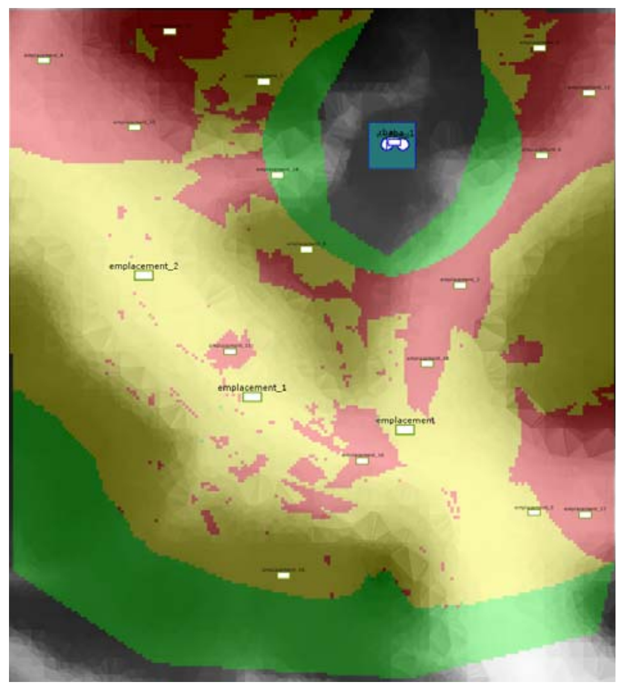

Figure 5: Classification with *3* active emplacements without considering semantics

Figure 4 shows the optimal distribution of *3* emplacements considering the ballistic semantics and figure 5 shows the same distribution on a grid without considering ballistic semantics. The green cells remain unchanged as can be seen by a comparison of figure 4 and 5 with 2 and 3, respectively. Considering the semantics, the number of green cells is *18,144* of *37,638*, i.e. *48.21* percent. When omitting the semantics, we only have *8,209* corresponding to a percentage of *21.81*. The optimization maximizes the number of yellow cells with constant green ones. Thus, the optimization criterion is the number of yellow cells divided by the sum of red and yellow cells.

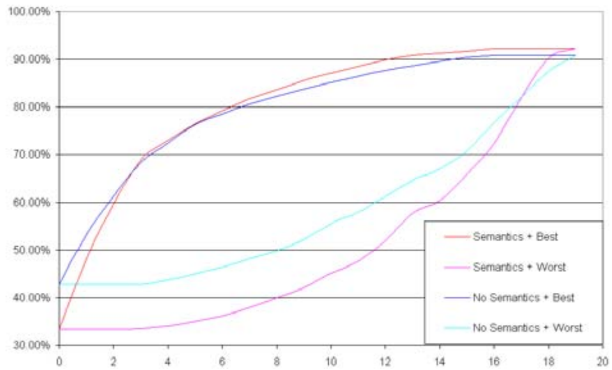

Figure 6: Percentage of yellow cells

Figure 6 shows this measure for the best and worst distribution of *0* up to *19* emplacements with and without considering semantics. The typical logarithmic shaped curve indicates that saturation is reached. The best distributions are much better than the worst ones if the number of possible decisions is sufficiently high. Note that the *0*- and *19*-distribution are exactly the same in worst and best optimization since no decision can be made. Running the optimization on the grid considering ballistic semantics gives better results if at least six emplacements are distributed than the optimization results on the non-semantic one. The advantage of considering the semantics is because of the restricted operational area for Red. Note that the non-semantic grid is superior without any spotters being distributed. The reason for this is that most cells getting green are also visible by the watch towers. In eight of nineteen distributions, the result is different.

In *SoF*, the *QRF* is not waiting in the base but patrolling in the valley as depict in figure 7. The blue line is its current route. The red attacker in the upper part of figure 7 is currently approaching its improvised fire position. In this case, a *QRF* positioned in the base would probably perform better. Note that the camera towers inside the base always support the detection of the Red. As mentioned above, the red units can only be detected after they have started preparing their attack. The *QRF* has limited time to reach the attackers before they can fire their rockets.

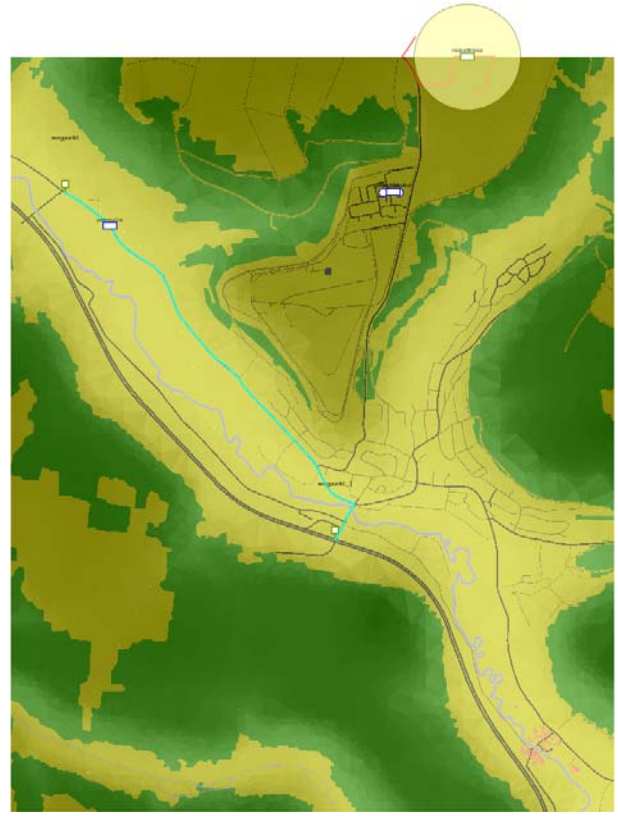

Figure 7: Running simulation of *SoF*

Although we defined three different strategy options for the blue forces, we are not interested in comparing these. As already stated in the introduction, we want to evaluate the influence of the additional semantic information as well as the optimization feature. Basically, *PfB* is a sub-strategy of *CE* (with no emplacements) as well as *SoF* with only one patrol staying in the base. In the following, we want to use all strategies to evaluate the semantics' influence on the *MoE*. With the *CE* strategies, we want to determine the impact of the optimization feature.

Factor	Min	Max	Unit
Speed *QRF*	20	75	km/h
Speed Red	20	75	km/h
Height Red	1.5	3.0	m
Setup Time Red	2	10	min
Reload Time Red	1	3	min
Mean Detection Time Camera	1	3	min
Mean Detection Time Emplacements	1	3	min
Intended Shots Red	1	3	

Table 1: *NOLH* design

RESULTS AND ANALYSIS

We used a *Nearly Orthogonal Latin Hypercube (NOLH)* [2] design with *65* design points for our experiments and crossed it with the two kinds of semantics. Thus, all experiments have been run on the grid depicted in figures 2 and 3, with and without considering the semantics for the ballistic reachability, respectively. Additionally, we run all scenarios by using the semantic routing or not. By using it, the routing algorithm is aware of the top-speed restriction on the roads. The *NOLH* design (see table 1) consists of the following parameters: Speed of blue *QRF*, meantime of detection of the cameras on the watch towers in the blue base, the meantime of detection of the blue emplacements, the speed of the red attackers, the height of Red, Red's setup time (the time needed to build up its ballistic weapon), the reload time of Red and the number of intended shots at the base. For each design point *50* replications are conducted with different seeds. In total *13,000* simulation runs for each scenario are performed. Altogether, we modeled *8* scenarios, i.e. one *PfB*, one *SoF* and six *CE* scenarios. The latter scenarios were calculated with *1*, *3* and *5* spotters, distributed according to best and worst optimization results, respectively. Thus, we performed *104,000* simulation runs.

As mentioned above, we wanted to examine the impact of considering the semantics as well as performing the optimization on blue's success, i.e. the *MoE PreventedShots*. Additionally, we compare Blue's strategies very briefly.

Impact of Optimization

The optimization result depends on the ballistic specific interpretation of the semantics. If the information is considered, different distributions are calculated as already mentioned above and can be seen in figure 6. For simplicity, we only used the results of the grid considering the semantics and evaluated this distribution with both grids, i.e. the semantic and non-semantic one. We distributed one, three and five emplacements according to the best and worst optimization result. Tables 2 and 3 show the results. Note that *0* observers correspond to the *PfB* strategy serving as baseline. The worst distribution of observers made the *MoE PreventedShots* even worse. In most cases, the best distribution outperformed the worst one significantly. Only the non-ballistic grid with one emplacement is an outlier. The reason therefore is that the optimization has been performed considering ballistic semantics but the simulation has not. Thus, the optimization criterion does not match the simulated reality. This artifact is a hint that the model of the optimization is not accurate enough and is a great indication that accurate and correct models must be used during optimization and simulation in order to get robust results. Thus, we claim that our optimization according to the visibility of the spotters improves Blue's statistically significant.

# observer	Ballistic Semantics		No Ballistic Semantics	
0	mean	0.331	mean	0.205
	std-dev	0.435	std-dev	0.362
	std-err	0.005	std-err	0.004
1	mean	0.359	Mean	0.264
	std-dev	0.442	std-dev	0.405
	std-err	0.005	std-err	0.005
3	mean	0.470	mean	0.360
	std-dev	0.468	std-dev	0.441
	std-err	0.006	std-err	0.005
5	mean	0.561	mean	0.446
	std-dev	0.463	std-dev	0.462
	std-err	0.006	std-err	0.006

Table 2: "Best" optimization of *CE*

# observer	Ballistic Semantics		No Ballistic Semantics	
0	mean	0.331	mean	0.205
	std-dev	0.435	std-dev	0.362
	std-err	0.005	std-err	0.004
1	mean	0.316	mean	0.288
	std-dev	0.428	std-dev	0.416
	std-err	0.005	std-err	0.005
3	mean	0.235	mean	0.219
	std-dev	0.389	std-dev	0.375
	std-err	0.005	std-err	0.005
5	mean	0.203	mean	0.241
	std-dev	0.374	std-dev	0.390
	std-err	0.005	std-err	0.005

Table 3: "Worst" optimization of *CE*

Impact of Semantics

"Does considering the semantics have an impact on Blue's success?" - "It depends".

The impact of the semantic information is different according to the different effects. We modeled two effects: The impact on the ballistic reachability as well as the impact on the semantic routing referred to as ballistic semantics and road semantics, respectively.

Strategy	Ballistic Semantics		No Ballistic Semantics	
PfB	mean	0.331	mean	0.205
	std-dev	0.435	std-dev	0.362
	std-err	0.005	std-err	0.004
SoF	mean	0.532	mean	0.536
	std-dev	0.463	std-dev	0.471
	std-err	0.006	std-err	0.006

Table 4: *PreventedShots* of *PfB* and *SoF*

The impact of the ballistic semantics can be read off tables 2, 3 and 4 for the different strategies. Especially with *CE*, a significant effect can be seen, since in all cases, Blue's success raises when ballistic semantics has been considered during simulation. The main reason therefore is that the semantics has been considered during the optimization step. For the baseline *PfB*, the impact is also significant. The reason therefore is that many cells that are green, i.e. not ballistically threatening, only if the semantic is considered are located far away from the base (cf. figures 2 and 3). These cells cannot be reached by the *QRF* in time with high probability. Thus, Red can launch more attacks if the ballistic semantics is not considered. Using strategy *SoF* Blue performs not significantly different in comparison to the scenario without semantics. The reason is that the patrol route (cf. figure 7) is located in an area where many cells are ballistically threatening only if the semantics are not considered. Thus, Red attackers appear near the patrol route very likely and are neutralized with high probability. Thereby note that the patrol route has not been optimized but chosen quite arbitrarily by the user. Summarizing the ballistic semantics' effects, we claim that the impact is strongly related to the blue strategies and the amount of knowledge that has been used for their optimization. It is important that this optimization must consider the same semantic information as the simulation does. The reason for *PfB* to perform better without semantics and *SoF* to perform the same is random: The *SoF* patrol route has been chosen luckily and the far cells that are very unlikely to be reached in time from the base are not ballistically threatening when considering the semantics. But it is very important to have an accurate model in order to influence the strategy by optimization and simulation.

The answer to the second question is different. In all scenarios, the road semantics had no significant effect. For brevity, we do not show any results. We think the main reason was a realistic top-speed restriction on the road. Additionally, this restriction holds for red as well as for blue forces. Since both units had quite similar speeds (cf. table 1), they were restricted similarly on the roads and no influence of the success could be determined statistically. The road semantics might get interesting if it is considered for the setup of the strategy, e.g. during an optimization of patrol routes.

Comparison of Blue Strategies

During IDFW18, we also compared the different strategy options itself. Although this comparison was not our aim at this workshop, we can compare the strategies with tables 2, 3 and 4. Note that *PfB* is a sub-strategy of *SoF* as well as *CE*. Thus, we only have to compare the latter two strategies. *CE* is supported by the optimization module and the patrol route of the modeled *SoF* scenario has been chosen arbitrarily. The baseline strategy *PfB* is the worst one for Blue if we exclude the *CE* strategies with worst distributions. If we then distribute emplacements in an intelligent, i.e. optimized, manner, Blue's success rises as can be seen by the *CE* scenarios. Following strategy *SoF* with one patrol is nearly as successful (with the defined patrol route) as distributing *5* emplacements. This result basically confirms our results from IDFW18. The reasons for *SoF*'s success are the following: On the one hand, the patrol in our scenario is in an area where attackers approach with high probability. Thus the time needed to disturb the Red is quite short. On the other hand, there exists a chance that Red senses Blue before starting to prepare its attack. Then it gets discouraged and flees without any attack. For these two reasons, *SoF* seems to be the best strategy for Blue. This superiority might even rise when the patrol routes are also optimized or more patrols are distributed.

Limitations

There are some limitations in our experiments. Concerning the optimization part, we can do a crossing of optimization and simulation runs. Although we optimized certain number of emplacements with and without considering the ballistic semantics, we only simulated the results of the former optimization. During the simulation, we then considered both ballistic grids again. In upcoming studies, we want to simulate the results of the latter optimization with both grids in order to check if a significant impact on Blue's success can be measured.

Another important limitation is that we are able to optimize the visibility of spotters, only. Additionally, we could define several optimization criterions, e.g. the reachability of possible attack points, in order to distribute *QRF*s outside the base also. One major challenge is the definition of the optimality if several criteria are considered. This is also subject to further research.

For an extensive comparison of the strategies, the following questions should be answered:

- How many emplacements/ patrols are needed to cover all cells?
- How can *n* emplacements/ patrols be distributed such that most cells are covered?
- What is a good ratio between covered cells and used emplacements/ patrols?
- How many emplacements/ patrols are needed to avoid any attack?
- How can *n* emplacements/ patrols be placed such that most attacks are avoided?
- What is a good ratio between avoided attacks and used emplacements/ patrols?

The first three questions can be answered for emplacements, but not for patrols since we currently are not able to optimize dynamic patrol routes. The last three questions could be answered by using simulation runs, but we have no hint for patrols to be defined without a corresponding optimization. Especially the third and sixth question is relevant since a notion of used resources and utility comes into account. The third question is basically answered for emplacements in figure 6.

The last limitation mentioned here is that only one specific Red behavior is modeled, which is also not evolving over time. We also always consider exactly a single attack and not multiple coordinated attacks which adapt to Blue's strategy. In order to challenge these questions at least semi-automated, we want to further extend our current approach with optimization techniques which are able to derive strategy settings automatically. Such a system could use evolutionary algorithms combined with data-farming similar to *Automated Red Teaming (ART)* [3] and *Automated Co-Evolution (ACE)* [3].

CONCLUSIONS

During this study, we evaluated the influence of considering semantic attachments to our terrain data on Blue's success in terms of the regarded *MoE PreventedShots*. We wanted to check if it is worth to invest the effort to model or import such semantic attachments and model the reaction to it. We explicitly do not want to generate results like: "Blue's *MoE* decreases when considering semantics, so we do not consider it!" We claim that such a model is more realistic and thus more accurate. The question is if we want to invest the effort to improve such a model. As can bee seen in the results of the optimization of emplacement positions, any information that is used to optimize some task must be based on an accurate model in order to generate correct results. For example, a spotter distribution generated on a ballistic grid considering semantics performs very well in a simulation considering ballistic semantics but not necessarily in a simulation without considering ballistic semantics.

We want to emphasize that the terrain data as well as the semantic attributes have not been modeled by us but are imported from a database of the German Armed Forces. We simply modeled the reaction to the attribute values, e.g. the top-speed restrictions of certain unit types on certain road types. Basically, we modeled effect of the semantics to the routing (road semantics) and to the ballistic threatening of cells (ballistic semantics).

The influence of the ballistic semantics is statistically significant if the information is used to set up the strategy. It is not significant and even arbitrary if we just evaluate a given scenario or strategy. The impact of road semantics is not significant in our study. This last result was a surprise.

Additionally, we wanted to check if our optimization according to the grid-visibility results in a significant improvement of Blue's success. During the investigation we learned that the latter question can be answered with a "yes", i.e. there is a significant improvement, if the model of the optimization is close enough to the model of the simulation.

As future work, we want to extend our approach as mentioned above. Additionally, we want to analyze the impact of the optimization, extend the system to be able to cope with several optimization criteria, develop an optimization for dynamic patrol routes and try to cope with coordinating and evolving red behavior.

REFERENCES

[1] D. Kunde, N. Mausberg, C. Mies, and B. Wolters, "Force Protection With ITSim: Base Protection Against Ballistic Weapons," in *Scythe, Proc. and Bulletin Int. Data Farming Community, Issue 6*, 2009, pp. 26-31.

[2] T. M. Cioppa, "Efficient nearly orthogonal and space-filling experimental designs for high-dimensional complex models," PhD Thesis, Naval Postgraduate School, Monterey, CA, 2002.

[3] C.S. Choo, E.C. Ng, D. Ang and C.L. Chua, "Data farming in Singapore: a brief history," *Proceedings of the 40th Conference on Winter Simulation*, 2008, pp. 1448-1455