Hybrid Qualitative Simulation of Military Operations

Thomas Hinrichs[†], Kenneth Forbus[†], Johan deKleer[††], Sungwook Yoon[††], Eric Jones[‡], Robert Hyland[‡], Jason Wilson[‡]

[†]Department of EECS,
Northwestern University,
2133 Sheridan Road, Evanston IL 60208
{t-hinrichs, forbus}@northwestern.edu

[††]PARC, (Palo Alto Research Center)
3333 Coyote Hill Road
Palo Alto, CA 94304 USA
{dekleer, sungwook.yoon}@parc.com

[‡]BAE Systems AIT
Intelligent Systems Division
6 New England Executive Park
Burlington, MA 01803
{eric.k.jones, robert.hyland,
jason.r.wilson}@baesystems.com

Abstract

Our goal is to enable military planners to rapidly critique alternative battle plans by simulating multiple outcomes of adversarial plans. We describe a novel simulator, SimPath, that combines qualitative reasoning, a geographic information system (GIS), and targeted probabilistic calculations to envision how adversarial battle plans can play out. We outline the problem and describe the overall operation of the simulator. We then explain how qualitative process theory is extended with actions to model military tasks, how envisioning is factored to reduce combinatorial explosions, and how probabilities are computed for transitions and used to filter possibilities. Empirical results, including an experiment conducted by an independent evaluator, are summarized. The results show that it is possible to identify dozens of possible outcomes on each of 9 combinations of adversarial plans (COAs) in under two minutes. We close with a discussion of future work.

Introduction

In battle planning, military personnel consider multiple alternative courses of action (COAs), as well as possible adversarial plans. The uncertainty in an enemy's plans and of battlefield outcomes is part of what makes this task so difficult. Traditional numerical simulations provide little help, since they require extensive setup, many *ad hoc* numerical assumptions, and significant computational power to produce even one possible outcome for one scenario. Monte Carlo simulation can be used to generate samples from the space of possible outcomes, but at the cost of even more computational power and without any guarantee that the space of behaviors is sampled adequately. In principle, qualitative simulation should be perfect for this application. In qualitative simulation, the infinite space of possibilities is characterized by a finite (but possibly large) set of states, called an *envisionment*. Envisionments could be analyzed to detect possible opportunities and blind alleys (Price, 2000) and used to track execution as a battle unfolds. However, military battles are far more complex than previous QR domains. First, military tasks combine aspects of continuous change with *intent* (e.g., deciding to break off an attack if one's losses are too heavy). Second, reasoning about COAs requires complex geospatial reasoning. Third, battlefield envisionments can be very large, due to the number of entities involved. Fourth, probabilities must be estimated for state transitions, both to prune truly unlikely possibilities and to help track execution. These tough challenges are part of what DARPA's Deep Green program (Surdu and Kittka, 2008) is tackling.

This paper describes SimPath, our qualitative simulator for military operations. We begin by outlining the problem and discussing the overall structure of the simulator. Next we describe the modeling of military tasks, extensions to qualitative process theory (Forbus, 1984), and role of qualitative spatial representations. The computation of probabilities is described, including their use in filtering. Empirical evaluation results are discussed next, closing with future work.

This work was sponsored by The Defense Advanced Research Agency (DARPA) Information Processing Technology Office (IPTO) under the Deep Green program. The views and conclusions in this document are those of the authors and should not be interpreted as representing the official policies, either expressly or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

A brief guide to Courses of Action

Figure 1 illustrates a simple pair of COAs. The red diamonds and blue rectangles represent enemy and friendly Mechanized Infantry Battalions, respectively. Much of the COA is specified graphically, including arrows (here indicating Avenues of Approach) and symbols denoting particular tasks. Military tasks, such as Destroy and Fix (i.e., prevent the other unit(s) from moving), often involve multiple sequential actions, and being adversarial, have highly uncertain outcomes. Military organizations are hierarchical, with tasks at one echelon being implemented by tasks at the echelon below it, specified by COAs at their level, and so on. This is a simple COA: many involve a dozen or more units even at the highest-level echelon, which can expand to many more lower-echelon units.

Hinrichs, T., Forbus, K., de Kleer, J., Yoon, S., Jones, E., Hyland, R., and Wilson, J. (2010). Hybrid Qualitative Simulation of Military Operations. In 24th International Workshop on Qualitative Reasoning, August 8-10, Portland, Oregon, USA.

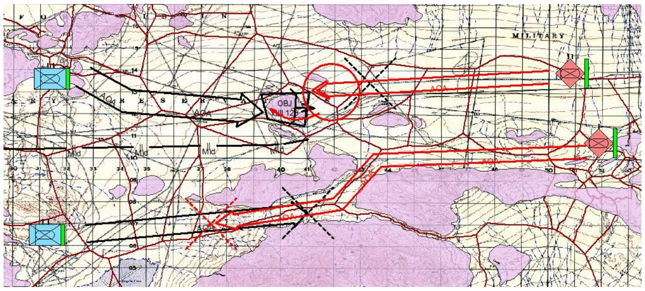

Figure 1: Example Course of Action

COAs are highly spatial and are specified with graphical symbols, including different kinds of regions (e.g., *engagement areas*, where a commander intends to fight an enemy force) and paths (e.g. *avenues of approach*, the route an attacking force is taking). These spatial symbols decompose terrain into functionally meaningful units relative to the COA, and hence form a kind of qualitative spatial representation (Forbus, Usher, & Chapman, 2003).

Military planners operate under a harsh information environment: Opponents hide their intent and capabilities, and information about one's own units can be incomplete. Consequently, planners must consider multiple possible opponent COAs, and often consider multiple COAs for themselves. War-gaming, i.e., playing out the possible outcomes of these plans, is largely a manual process today. There are three reasons for this. (1) Many parameters are unknown. (2) Quantitative models are not close to the accuracy of models in science and engineering, due to the variety of factors involved, some of them psychological. (3) Multi-trajectory numerical simulation methods, such as Monte Carlo simulation, require massive amounts of computation to achieve reasonable accuracy.

This problem is a natural match for qualitative reasoning. Many of the qualitative spatial representations have already been articulated by domain experts. The dynamics of battles can be represented at an appropriate level of precision by qualitative states and transitions. However, it is not enough to know that a unit may or may not survive an engagement; planners also want to know the probabilities of different outcomes, which requires a quantitative component. Quantitative calculations can also help constrain qualitative envisionment, thus the system we describe is a hybrid qualitative/quantitative simulator.

SimPath Overview

The input to SimPath consists of three Blue COAs and three Red COAs. Each COA uses its own terrain and task graphics. The output is a *futures graph*, an envisionment with probabilities for each state transition, representing how all nine COA pairs might play out.

Given the properties of COAs, it might first appear that envisioning would be completely intractable. After all, envisioning is exponential for most domains, and the addition of human intent to the vast number of physically relevant quantities (e.g., vehicle/weapon properties, attrition, ranges) means that the potential number of states is astronomical. Fortunately, the domain also imposes significant constraints over purely physical domains: Units are constrained by the COAs to do what they are tasked to do, as opposed to considering all possible actions. Moreover, a COA constrains units' spatial locations, and movement rates constrain when they can be in particular locations. This allows us to factor the envisionment into much smaller and more tractable problems.

Here is the core idea underlying SimPath: (1) Spatial reasoning is performed to find the possible spatial intersections of units, based on where they might be during the COA. A geographic information system is used to ground the qualitative spatial reasoning (Donlon and Forbus, 1999); while there are interesting innovations in SimPath's spatial reasoning, we do not consider them further in this paper for brevity. (2) Temporal reasoning is performed to identify which of the spatial intersections are possible, using interval calculations over possible speeds for the units involved. (3) Each of these spatio-temporal intersections becomes an *engagement*, which typically represents a battle, where the dynamics of the interactions between units must be simulated. Envisionments are constructed for each engagement separately, thereby factoring the futures graph and drastically reducing its size. Engagements are envisioned sequentially, so that results of earlier engagements constrain subsequent engagements (e.g., a unit that is destroyed is no longer a participant). (4) When an envisionment for an engagement is constructed, a *stochastic simulator* (described below) calculates probabilities for each state transition. Transitions that are extremely unlikely are pruned, further reducing the graph.
The rest of this paper focuses on qualitatively modeling military tasks so they can be effectively envisioned, and on computing probabilities, since these are some of the most innovative components of SimPath.

Qualitative modeling of military tasks

Over time, military doctrine has codified tasks in semi-formal representations, consisting of natural language plus graphical descriptions. Qualitative representations are useful for formalizing military tasks because they unfold over time and often have continuous effects. Moreover, they do not have fixed outcomes, since their results depend on context, including what the opponent does. The representation must ultimately be *compositional* so that models can be constructed independently and combined algorithmically. Many aspects of them can be cleanly expressed in QP theory (Forbus, 1984). However, other aspects require significant extensions, notably discrete actions such as deciding to break off an engagement.

Compositionality

We achieve compositionality by first reducing military tasks to sequences and combinations of a smaller set of about twenty types of primitive behaviors, such as

`ConductTacticalManeuver` and `AssaultEnemyPosition`. These behaviors may still interact with enemy behaviors in complex ways, so they are further reduced to an even smaller set of *process instances*. The vocabulary of these processes is quite small and directly mirrors the types of capabilities that units have, such as mobility (`MovementProcess`), combat (`DirectFireProcess`), obstacle crossing or preparing defenses. As per QP theory, these processes have conditions that determine when they are active, and consequences that represent influences between quantities (see Table 1). The set of influences that hold in a situation constitute a qualitative differential equation, including in addition causal information. The vocabulary of influences consists of *direct* influences (i+, i-) which are essentially positive and negative constituents of qualitative derivatives, and *indirect* influences (qprop+, qprop-) which express monotonic functional dependencies. By reducing 156 types of military tasks down to four types of qualitative influences, the envisionment algorithm is able to compose their effects in a domain-independent way.

The engagement envisionment algorithm

Envisioning engagements is implemented very similarly to the original QPT algorithm. In order to envision possible next states, we begin with an initial qualitative state. This contains a set of inequalities on continuous quantities, and the set of all possible process instances and model fragments, some of which are initially active.

The algorithm has 3 main steps. Given an initial state, (1) resolve active influences on the quantities to determine whether they are increasing or decreasing, (2) identify the possible thresholds or *limit hypotheses* that the quantities could next encounter, (3) for each consistent combination of those quantity limits, identify which processes and model fragments would be active, and construct possible next states and explicit transitions. This process is repeated until no more states can be produced.

Specifically, SimPath resolves influences on a quantity by identifying all the direct influences on it and if there are competing influences, determining which direction of influence will dominate. This may involve numerically comparing rates of processes or looking for explicit inequalities in the current state. If there are no direct influences, it looks for qualitative proportionalities. Recall that these may be nonlinear and cannot be arithmetically combined to determine dominance. If there are competing indirect influences on a quantity, its direction of change cannot be resolved.

Process	Quantities	Conditions	Consequences
Movement	D *Distance to destination* F *Unit force strength* S *Speed*	D > 0 F > 0	(i- D, S)
Direct Fire	D *Target range* W *Weapon range* F *Unit force strength* E *Enemy force strength* R *Rate of fire* H *Hit rate* P *Defensive preparations*	D < W F > 0 E > 0	(i- E, H) (qprop+ H, R) (qprop- H, P)

Table 1: Example process types

Limit analysis consists of finding possible thresholds for each changing quantity. The inequality conditions on the instantiated processes and model fragments define the universe of limits to consider. The resolved direction in which a quantity is changing determines which limits could be encountered next. For unresolved quantities, both possible directions are hypothesized.

The limit hypotheses correspond to possible changes from the current state. The possible next states are defined by the consistent combinations of quantity changes taken from among the cross-product of limit hypotheses. I.e., all consistent combinations of thresholds being reached or not reached are considered, as determined by a depth-first search. Pruning by consistency is important for maintaining accuracy, with the side-effect of greatly reducing the size of the envisionment.

To determine which processes will be active in a new state, the envisioner queries the conditions on all process instances and determine if they are entailed by the states's inequalities. Once the consistent relationships and active processes are determined, a next situation object is constructed and compared to existing situations in the futures graph. If such a situation already exists, a transition is added from the current situation to that situation, otherwise the new situation is added to the futures graph.

Sequencing processes

Because military tasks are intentional rather than causal, they often involve executing sequences of behaviors, and the envisionment need not branch through all physically possible behaviors simultaneously. For example, in most contexts a unit will engage targets of opportunity on their way to their objective and will complete that engagement before proceeding on. Sequencing ensures that the unit will not be moving towards two destinations at once and end up in inconsistent locations. To implement sequencing, we implemented a Boolean condition in the form of an *after* predicate and a *started* predicate. The after predicate prevents one process from becoming active until after the other has started. When a process is history-dependent, it causes the envisioner to record the precondition processes that have started and to consult this record to enforce sequential process activation.

Discrete actions & decisions

Although the majority of military tasks have non-negligible durations, there are some that are effectively instantaneous. A *decision* to withdraw, or to change formation, or reinforce another unit is essentially an action. We treat actions as atomic, guaranteed transformations of the current state that are run to exhaustion at the beginning of each envisionment cycle. Actions are represented as quantity-conditioned model fragments (like processes) but instead of influences, their consequences are add-lists and delete-lists of Boolean conditions (Forbus, 1989).

Actions are primarily used to manipulate intent, the will to pursue some coherent activity. When that activity can change based on quantitative thresholds, we represent intent explicitly as a Boolean condition on the processes that support the activity. Actions to start or terminate an activity simply add or retract the intent condition from the current state. This makes it possible to model activities such as *Follow-and-Assume*, in which one unit follows another and takes over its task when the lead unit is no longer capable.

Modeling actions and intent this way also enables us to represent simple decision processes that may branch directly on the initial situation rather than be influenced over time. For example, when faced with an overwhelmingly superior force, a commander may withdraw before suffering losses. Instantaneous actions allow us to model that behavior.

Modeling space and terrain

While spatial reasoning helps determine what engagements to envision, spatial considerations also contribute to the envisionment itself and to the calculation of probabilities afterward. Within an engagement envisionment, the key spatial representation is *relative distance*. That is, space is projected down to one dimensional distance quantities.

Working with distance quantities introduces some complexities, since projection inevitably loses information. First, the 'as the crow flies' distances may not reflect actual travel distances. A route may take units closer together and then farther apart (think switchbacks on mountain roads). The envisionment algorithm assumes routes approach or depart monotonically. Nonmonotonic routes must be broken up into segments, each corresponding to a distinct movement process.

The second complexity is consistency checking. Given only a set of distance inequalities, how do we know if an additional distance inequality is consistent? Unlike other linear quantities, distances don't combine transitively because angular information has been projected away. So for example, if the distance from location A to B is 1000 meters, and the distance from B to C is 500 meters, then the distance from A to C could be anywhere from 500 to 1500 meters.

The distance consistency algorithm checks a proposed distance inequality by computing maximum and minimum bounds on that distance given a qualitative state. It first uses a breadth-first search to identify a traversal of the distance graph equivalent to the query distance (e.g., AB + BC in the previous example). It then finds the biggest leg in this path and adds and subtracts the cumulative remainder of the traversal (truncating at zero) to determine max and min distances. The algorithm uses interval arithmetic because each distance may itself be bounded by inequalities. It also keeps track of open and closed ranges entailed by different predicates (e.g., $>$ vs. $\geq$).

Integration with Stochastic Simulator

In addition to envisioning possible futures, SimPath also assesses the likelihood of those futures. This is important to help commanders assess COAs during planning, to support tracking and situational awareness during execution, and to improve the efficiency of the simulation. Transitions with sufficiently low probabilities are pruned from the futures graph, removing entire trajectories from further consideration. The probabilities are also used to help track execution during battle and maintain situational awareness, a topic beyond the scope of this paper.

SimPath computes probabilities of state transitions after each round of qualitative envisionment using a *stochastic simulator*. For each qualitative state, the simulator creates a distribution of sample outcomes that can be classified into buckets corresponding to the different qualitative transitions. A sample outcome is constructed by stepping through time and simulating each influenced quantity. Distances, for example, are updated as a function of preferred vehicle speeds, terrain and route. Force strength, or more accurately, platform attrition, is simulated using a random variable as input to Lanchester equations (Lanchester, 1916) to probabilistically determine when a unit has lost one of its platforms. The (Lanchester) *lethality* of a platform is determined by combat power of the platform and modified by situational factors including terrain features, visibility, and postures and therefore is constant over a situation. The ability of a unit to attrit an enemy unit is determined by the product of platform lethality and the number of attacking platforms in the friendly unit. The simulator stops and produces a sample outcome when one of the quantities reaches a threshold value that defines a qualitative state transition. The simulator produces 1000 samples for each situation.

Although the stochastic simulator is slow compared to qualitative envisionment, the integration of the two is practical because they mutually constrain each other. The factored envisionment constrains the number of situations that must be simulated, and conversely, the stochastic simulator rules out numerically impossible trajectories. Whereas the qualitative envisionment may indicate two possible outcomes to an engagement, the stochastic simulator may show one outcome to be vanishingly unlikely, and no subsequent futures will be envisioned for the losing unit.

Putting it all together

The more quantities and influences represented, the more possible thresholds and futures states are generated. This quickly adds up. The main strategy for controlling this is to divide the simulation into smaller, non-interacting pieces and envision them separately. We call this *factoring*.

The battlefield is factored by identifying the collections of tasks in the opposing COAs that can interact. In maneuver warfare, tasks can interact if the trajectories of their assigned units intersect in space and time. The *geospatial reasoner* computes the spatial intersections based on the ranges of weapons systems, rather than merely the unit footprints. It also considers ground control measures such as boundary lines and phase lines. Next, the *task grouper* computes possible temporal intersections of these tasks using a Simple Temporal Network (STN) (Dechter *et al*., 1991). The resulting set of interacting tasks is translated into a qualitative situation containing processes that influence a common set of quantities. For example, the COA pair in Figure 1 shows two parallel engagements between blue (friendly) and red (enemy) units. The engagement in the north can be envisioned independently of the engagement in the south, significantly reducing the number of possible states produced.

Experience & Empirical Evaluation

In qualitative modeling, it is critical to omit needless detail. Each process and quantity has the potential to significantly increase the branching factor of the envisionment. We quickly learned to avoid modeling quantities such as fuel or ammunition, because this led the envisioner to propose that every unit might run out of those resources at every step along the way. Unit boundaries are another example. COAs can be very explicit about constraining maneuvers to stay within boundaries. Such invariants should be left out of a qualitative model because the envisioner will turn them into limits and model units crossing them.

Compositionalithy can be both a blessing and a curse. It has the benefit that a model that is valid for a 1x1 interaction is also valid for a 1x2 or a 2x2 interaction. However, when multiple units interact, the complexity and size of the envisionment grows rapidly. We found factoring to be essential in simulating large COAs. In fact, for a typical COA, factoring yielded an average 17-fold improvement just in the number of terminal situations. This result led us to define additional rules for factoring based on different ground control measures, and quickly became the primary means for controlling state explosions.

The stochastic simulator took the majority of the execution time. It typically succeeded in pruning between a third and a half of the situations in an envisionment. The downstream effect of this can be enormous.

External Evaluation

In addition to our internal evaluations, an external team of researchers and subject matter experts (SMEs) conducted a formal evaluation in 2009. This evaluation tested fidelity, the faithful execution of tactical tasks, and scalability, the time required to simulate 9 COA pairs.

To assess fidelity, two SMEs subjectively evaluated limited, task-focused "vignettes." These were designed to measure how well the behavioral models simulated the range of expected outcomes. Each SME scored 28 vignettes, covering 16 types of tasks and decisions, and assigned each a numeric score, which was grouped into one of three categories: *acceptable* (appropriate answers); *partially acceptable* (capability needs some work); or *unacceptable* (capability is not adequate or does not exist). As shown in Figure 2, SimPath simulated 12 vignettes with acceptable fidelity and 14 vignettes with partially acceptable fidelity for a wide range of tasks.

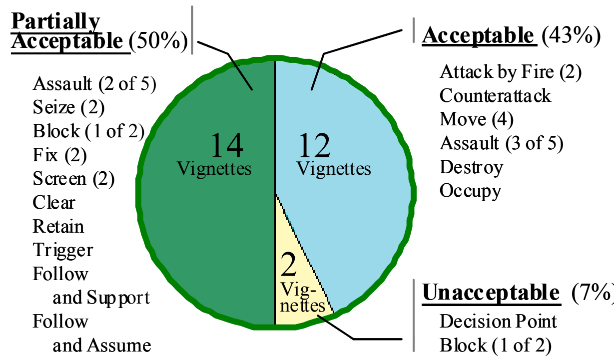

Figure 2: SimPath simulated 93% of vignette tests at a minimum fidelity or greater, covering a range of tactical tasks.

For each vignette, the SMEs answered 17 questions that assessed the realism of combat results, duration, execution, movement, use of terrain, and other modeling details to better characterize SimPath's strengths and weakness.

SimPath scored well with movement, use of terrain, use of engagement ranges, and multiple believable outcomes with realistic probabilities. More work is needed to model task synchronization and Decision Points, which are explicit contingent branches in the COA. Tasks that required decomposing units to subunits also did poorly. Some FollowAndAssume tasks were factored incorrectly, leaving the leader and follower in separate engagements.

Scalability tests assessed performance on more realistic COAs, which are bigger than vignettes, and have more "moving parts". They involve chaining together multiple engagements and exercise factoring of the battlespace. There were four COA scenarios (defined by terrain and force structures). For each scenario, three Blue and three Red COAs were produced and simulated against each other in nine pair-wise combinations.

In the scalability tests, Simpath's speed was better than expected. The goal was to simulate nine COA pair combinations in a half hour. By simulating each COA pair in parallel on a multi-core server and offloading GIS code to laptops, execution time was under two minutes. We take this as vindication for our hybrid qualitative approach.

On the other hand, there were some fidelity issues in the COA tests. Because each COA contained many subtasks, any incompletely modeled behavior tended to bring down the overall score. For instance, because the Decision Points were not implemented, some units did not move. Also, units in different engagements were not synchronized as intended. Since the simulator works at the level of individual engagements, coordination across engagements will require more sophisticated bookkeeping, such as that provided by an ATMS mechanism (de Kleer, 1986).

Related Work

Most military simulations operate at a fixed resolution and produce a single future trajectory. This is a reasonable approach for training simulations such as OneSAF (Parsons and Surdu, 2005), where humans may play a role alongside simulated agents in a massive multiplayer simulation. For the purpose of assessing plans and determining how they may succeed or fail in execution, it is important to record multiple outcomes and the audit trail showing how they came about. Doing this efficiently drove the decision to use qualitative simulation. In fact, extrapolating from the execution time of a single future trajectory in SimPath compared to the real-time simulation in OneSAF suggests a speedup factor of over a million.

An alternate approach is to use Monte Carlo simulation to produce an approximation of an envisionment (Atkin, Westbrook, & Cohen, 1999). This can be used to gain insights on small scenarios, but the exponential nature of Monte Carlo simulation makes it difficult to scale.

Future Work and Conclusions

SimPath is under active development. As part of Deep Green, its scope is expanding to cover intelligence, surveillance and reconnaissance (ISR), indirect fires, engineering tasks, aviation units, and ultimately full-spectrum warfare. Because these transcend simple move-and-shoot scenarios, they present new challenges. Since these tasks are not as local, detecting interactions (or non-interactions) is more difficult. Finally, counter-insurgency (COIN) modeling breaks assumptions about well-defined tasks, allegiances, and short, uniform time-scales.

Although qualitative reasoning is a mature technology, its application to military simulation has driven a number of innovations in this project:

- *Factored envisioning* – taking apart the situation, envisioning sub-graphs and merging them back together greatly reduces the size of the simulation.
- The highly *spatial* nature of the envisionment. Correctly capturing consistency and entailment of situations when 2D space is projected down to 1D distance quantities required new algorithms.
- Supporting intent and discrete *actions* (such as breaking fire, assuming another unit's mission) is also unusual in a process envisioner, though the idea has been explored previously (Forbus, 1989).
- Integration with a constrained *numerical simulator* permits transition pruning and probability estimation.

In addition to their importance for military simulation, we believe that these innovations may open up the use of envisioning for planning other complex human organization operations, such as sensor resource planning, disaster response and financial market analysis.

References

Atkin, M., Westbrook, D. and Cohen, P., 1999. Capture the Flag: Military Simulation Meets Computer Games. *AAAI Spring Symposium on AI and Computer Games.*

Dechter, R., Meiri, I., and Pearl, J., 1991. Temporal Constraint Networks. *Artificial Intelligence*, 49:61-95

de Kleer, J., 1986. An Assumption based TMS, *Artificial Intelligence*, 28:127-162.

deKleer, J., Forbus, K., Hinrichs, T., Sungwook, Y., and Jones, E.K., 2009. Factored Envisioning. In *Proceedings of the 23rd Annual Workshop on Qualitative Reasoning,* Ljubljana, Slovenia.

Forbus, K., 1984. Qualitative Process Theory. *Artificial Intelligence,* 24:85-168.

Forbus, K., 1989. Introducing actions into qualitative simulation. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1273-1278.

Forbus, K., Usher, J. and Chapman, V., 2003. Qualitative spatial reasoning about sketch maps. *Proceedings of the Fifteenth Annual Conference on Innovative Applications of Artificial Intelligence*, Acapulco, Mexico.

Gilmer, J.B. and Sullivan, F.J., 1998. Alternative Implementations of Multitrajectory Simulation. In *Proceedings of the 1998 Winter Simulation Conference.*

Lanchester, F.W., *Aircraft in Warfare: The Dawn of the Fourth Arm, London.* Constable & Co., 1916.

Parsons, D. and Surdu, J.R., The U.S. Army's Next Generation Simulation Modeling the Response to the World's Future Threat. *Proceedings, 2005 NATO Modeling and Simulation Group Conference.* Warsaw, Poland, October 2005.

Price, C.J., 2000. AUTOSTEVE: Automated Electrical Design Analysis. In *Proceedings ECAI-2000*, 721-725. Amsterdam, The Netherlands: IOS.

Surdu, J.R. and Kittka, K., 2008. "The Deep Green Concept," In *Spring Simulation Multiconference 2008 (SpringSim '08), Military Modeling and Simulation Symposium (MMS)*, 14-18 April 2008, Ottaway, Canada.