Deriving Priority Intelligence Requirements for Synthetic Command Entities

Jonathan Gratch
Stacy Marsella
Randall W. Hill, Jr.
University of Southern California, Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292
(310) 822-1510
gratch@isi.edu, marsella@isi.edu, hill@isi.edu

LTC George Stone III
JSIMS JPO
12249 Science Drive, Suite 260
Orlando, FL 32826
(407) 384-5554
george_stone@jsims.mil

ABSTRACT: *Simulation-based training is using increasingly complex synthetic forces. As more complex multi-echelon synthetic forces are employed in simulations, the need for a realistic model of their command and control behavior becomes more urgent. In this paper we discuss one key component of such a model, the autonomous generation and use of priority intelligence requirements within multi-echelon plans.*

1. Introduction

Command and control (C2) is an increasingly distributed and dynamic enterprise. In a modern campaign, distributed sensing platforms provide up-to-date situation monitoring, allowing distributed command–and control elements to plan and respond effectively to rapidly evolving situations. These innovations provide dramatic benefits, but they raise significant challenges both in training staff members to realize the system's potential, and in supporting this training with realistic models of the command and control process.

Research on simulated command entities has focused primarily on course of action development: What are the actions I can perform? How do I combine these actions to achieve my mission? This effort has begun to yield impressive results, including synthetic commanders that, in some circumstances, fully automate the plan generation process (e.g., see [1,3,4]).

Unfortunately, very little work has gone into understanding how commanders obtain the critical information they need to generate and monitor their plans. The advancement of sensor and communication technology provides a commander with unprecedented access to information, but unless the right information reaches the right person at the right time, this technology is wasted. Worse, commanders can become inundated with irrelevant or unimportant de-tails, such that access to more information actually degrades planning performance.

A key characteristic of a good commander is an ability to act proactively to determine what information is essential for the successful development and execution of his plans. In the Army, this information is formalized in terms of the Commander's Critical Information Requirements (CCIR). This is information "that must be brought to the commander's attention because of its potential impact on the decisions that he must make in order to be successful during an operation," [5]. By appropriately specifying CCIR in a operation order, a commander can manage his information flow by focusing staff and subordinates on what is essential.

Current models of C2 agents have not addressed the need for automatically deriving critical information requirements. In this article we describe a proposed approach for automating this determination within the context of synthetic Army C2 agents. We characterize the information requirements needed to support multi-echelon C2 planning. This characterization is based on an analysis of a Corps-level exercise. We show how the Commander's Critical Information Requirements (CCIR) can be derived automatically from an analysis of the information requirements that must be supported at higher echelons and

an analysis of the details of how the current echelon intends to accomplished its mission.

One area of application for this model is in training simulations that incorporate more realistic models of the command and control process. In particular, these models extend and generalize our prior work in building autonomous command and control nodes for the Synthetic Theater of War (STOW '97) and continuing under DARPA's Advanced Simulation Technology Thrust (ASTT). In these efforts, we modeled the organization of an Army aviation battalion, including the interactions between the battalion commander, his company commanders, the companies' respective vehicles, and a combat service support (logistics) unit. These interactions are supported by our Soar/CFOR agent architecture [2,3,4]. Soar/CFOR is implemented within the Soar cognitive architecture, providing substantial high-level structure and constraints that augment Soar's basic architecture. In particular, it provides data structures and reasoning capabilities informed by research in the AI planning community, and is ideally suited to command and control reasoning. The Soar/CFOR was designed as a domain independent reasoning system, and thus is not tied to any particular service or echelon. This makes the system readily configurable to other distributed planning tasks.

2. CCIR

Information is a commander's most important resource, but one must control the access to information or run the risk of becoming overwhelmed or disoriented. Any organization must address the issue of managing information flow between its members. In the business world, managers control their information by focusing on key *executive information requirements* [7]. In Army operations, commanders manage their flow of information by determining critical information for each operation and publishing these requirements through their CCIR (as described in FM 100-6, Information Operations). The CCIR provide a commander's staff and subordinate units with the guidance they need to best satisfy the commander's information needs.

CCIR is normally noted in paragraph 3d of an operations order and is broken down into three functional groups (FM 100-6):

- Priority intelligence requirements (PIR) determine what the commander wants or needs to know about the enemy, his purpose, and/or terrain
- Friendly forces information requirements (FFIR) allow the commander to determine the combat capabilities of his or adjacent friendly units
- Essential elements of friendly information (EEFI) allow the commander to determine how he must protect the force from the enemy's information-gathering systems

CCIR cannot be fixed (as in standard operating procedures). It must be precise to ensure responsiveness and, as information is perishable, it must adapt to evolving situations to ensure its relevance to the current context. A commander develops his CCIR based on the mission, his experience and higher echelon intent. In general, it is information that directly feeds the key decisions that will determine the success or failure of a mission. CCIR are typically fulfilled during execution but they also serve as important input for plan generation. For example, during the planning process a staff intelligence officer might recommend PIR to disambiguate between different possible enemy courses of action. Information requirements also drive the positions of sensors and observation posts, thereby imposing significant constraints on maneuver and logistical plan development.

Devising CCIR is no simple task; at least ten studies have attempted to analyze and classify CCIR[6]. To our knowledge, no one has attempted to model the process of devising CCIR in simulation. In our work, we have focused specifically on the problem of devising intelligence requirements related to the enemy and his purpose (PIR), and for now ignore the problem of deriving Essential Elements of Friendly Information (EEFI) and Friendly Forces Information Requirements (FFIR). The closest work in this vein is the Simulation Information Filtering Tool (SIFT) developed at STRICOM.[1] There are two aspects to handling CCIR in simulation. First, one must model the process of deriving information requirements. Second, one must model how these requirements alter the flow of information. SIFT addresses the second problem, acting as an information filtering tool. It prompts the user for inputs that define critical information requirements and then utilizes this guidance in real time to filter the flow of reports about a JANUS simulation scenario.

3. Example: Krasnovia Attacks

We illustrate the use and derivation of priority intelligence requirements by analyzing in detail how a division commander derives his intelligence requirements in the contesx of a corps-level training exercise typical of what might be performed at the National Training Center.

Let's consider an example operations order from a National Training Center (NTC) exercise in some detail. First, we paraphrase the situation which led up to the deployment of this plan, which is as follows (see Figure 1):

The Situation: Krasnovian forces, moving east, attacked across the international border, lead by the 19th Combined Arms Army (CAA) on the Inyo Front. The 19th CAA main attack was conducted by the 172d Motorized Rifle

---

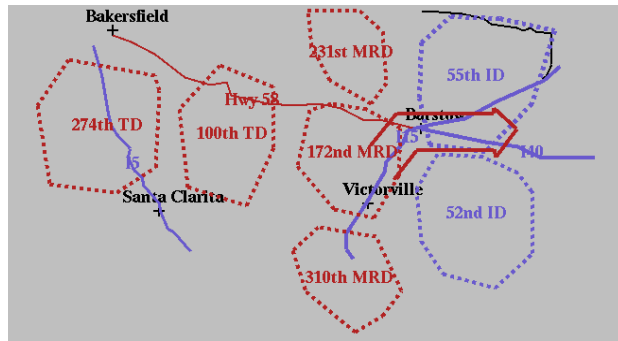[1] Information about SIFT can be found at www-mskeg.stricom.army.mil/sift/

Figure 1. II Corps situation

Division (MRD) and 100th Tank Division (TD) in the center of the sector, with support by the 231st MRD to the north and the 310th MRD to the south. Corps forces met and defeated the advance guards of the 19th CAA divisions.

The 19th's advance guard forces established defensive positions and their parent regiments have moved to the rear of these positions. No further defensive preparations have been observed. Second echelon regiments moved into position behind the lead regiments and halted. The 16th CAA, west of the 19th CAA, is moving out of its exercise areas and its forward division, the 274th TD, has paused along the International Border. Security precautions on Krasnovian convoy movements suggest possibility of introduction of chemical munitions but their use is considered unlikely.

The II Corps intent and plan is based on the assumption that the 19th CAA will resume offensive operations within the next 24 hours. The plan can be abstracted as follows:

II Corps Plan Abstract: The intent is to defeat the attacking Inyo Front forces using a mobile defense in order to shape the enemy's penetration west of Phase Line Hancock. The goal is the defeat of the 19th CAA with the Corps positioned to initiate offensive operations to secure the international border (IB).

The key component of the mobile defense is the 55th Infantry Division (Mechanized) in the north. Its plan consists of three phases, an area defense, a delay and a counterattack:

55th ID plan abstract: In Phase One, the 55th ID conducts an area defense. This includes an attack by the 55th Combined Arms Battalion (CAB) at EA (Engagement Area) Seminole to attrit the second echelon battalions of the lead regiments prior to their commitment to battle while air assets simultaneously attack the 2d echelon regiments (at Engagement Area Mohican). In Phase Two, the 55th ID on order delays to and defends at Phase Line

(PL) Hancock while holding the division's northern shoulder at PL Phoenix in order to shape the enemy's penetration. The division uses close air support (CAS) at EA CATAWBA to defeat first echelon forces. In Phase Three, the division counterattacks.

3.1 PIR

Now let's consider some of the intelligence requirements specified in the operations order. Although our main interest here is the 55th ID's PIR, let's first consider II Corps. The II Corps PIR includes, in order of priority:

1. When will the enemy first echelon regiments of the lead divisions reach the Corps forward defenses?
2. Where is the main attack of the 19 CAA?
3. Where are the main attacks of the 100 TD, the 231 MRD, the 172 MRD, and 310 MRD?
4. Where will the 274 TD, 16th CAA be committed?
5. Where are the second echelon regiments of the lead divisions?
6. Will the enemy employ NBC (Nuclear, Biological, or Chemical Weapons)? If so, when and where?

All these entries concern critical information about opposing force (OPFOR) behavior and intent, information which they will try to hide from the II Corps as part of the information warfare they will engage in. Among this information is the time and location of various elements of the CAA 19th and their attacks. This information directly affects II Corps plan; whether the assumptions about OPFOR behavior built into that plan are valid, how the execution of the plan needs to be timed and what adjustments or replanning are necessary. The information also serves as a check on the global assumption that the 19th CAA is resuming offensive operations and on what its objectives are.

Similarly, the entry concerning use of nuclear, biological or chemical weapons is here due to the OPFOR's capabilities and intelligence on the current situation. The use of these weapons would globally impact the plan's success, by altering the mobility and survivability of the Corp's forces.

Roughly, priority of the first five entries follows an implicit temporal order of when the information is needed. At the Corps level this structure tends to mirror the depth of the battlefield; first echelon regiments of lead divisions followed by the second echelon divisions and regiments. At the lowest priority is the NBC intelligence requirement which is considered unlikely and not tied exclusively to any particular phase of the operation.

This temporal structuring is more prevalent in the 55th ID (Mechanized) PIR which tends to mirror the entries in the II Corp PIR while specializing and enhancing based on its own operational concerns. Here is the 55th's PIR.

1. When will second echelon battalions of lead regiments reach EA SEMINOLE and the second echelon regiments reach EA MOHICAN?
2. When will the enemy first echelon battalions of the lead regiments reach the Division forward defenses?
3. Where are the main attacks of the regiments of the 100 TD and the 231 MRD?
4. When will enemy regiments reach EA CATAWBA?
5. Will the 16th CAA's, 274th TD be committed in the 55th Mech sector?
6. Will the enemy employ NBC? If so, when and where?

Let's consider this PIR in greater detail. Entry one follows directly from 55th's intent to simultaneously attack second echelon battalions of lead regiments at EA SEMINOLE and second echelon regiments at EA MOHICAN during Phase 1 of its plan. This early phase requires a convergence of spatial and temporal relationships that need to be explicitly affirmed. Entry two is a specialization of the Corp's PIR #1, particular to the concerns of the 55ᵗʰ ID. Furthermore, this entry pertains to the timing of the delaying action during Phase 2 of the operation and the shaping of the enemy's thrust which the delaying hopes to achieve. Similarly, entry 3 is a specialization of the II Corp's PIR #3 to the 55th's sector of operation. Here, it has special relevance to the Phase II delaying action. Entry four follows directly from the temporal/spatial constraints on the plan's action at EA CATAWBA at the end of Phase 2. Entry 5, detecting whether 274th TD will be committed to the 55th's sector, is a simple specialization of the Corp's Entry 4, detecting where 274th TD will be committed. In both cases, it has special relevance to the final phase of both plans, to counterattack. Finally, Entry 6, determining whether NBC will be employed is identical in both PIRs.

4. Sources of PIR

The sources for the 55ᵗʰ ID's PIR are illustrated in this example. For instance, the intelligence requirements often follow directly from properties of the division's plan. In particular, the attacks at SEMINOLE, MOHICAN and CATAWBA each presume that the enemy will reach a particular locale at a particular phase in the plan and that essentially will trigger the attack. These preconditions share several characteristics.

- First, they are triggering conditions for the execution of major subtasks in the division's plan. The information has a direct impact on the 55th's behavior and requires immediate response. In both cases, elements of the 55th ID must be ready to attack at that point.
- Second, achieving these preconditions is, to a large extent, outside the division's control since they depend in large part on the OPFOR's behavior. Thus

the preconditions cannot reliably be planned for but must be actively monitored.
- Third, there is considerable uncertainty associated with the preconditions. Because the information depends on OPFOR's behavior, the division needs to actively plan for the detection of these conditions. OPFOR will in turn engage in various forms of information warfare to counter such detection efforts. In the case of the attacks on SEMINOLE and MOHICAN, uncertainty is further increased because the attacks are designed to proceed simultaneously and therefore depend on the lead and second echelon Krasnovian regiments behaving consistently. In the case of CATAWBA, uncertainty is increased because it occurs further along in the battle, towards the end of Phase 2.

The division's PIR also suggests another source. In particular, the entries 5 and 6, whether the 274th will be committed in the sector and whether NBC will be used, follow directly from global concerns/presumptions about the likely Krasnovian intent and course of action. Although they impact 55th's plan, they more likely preempt or invalidate it entirely as opposed to controlling (triggering) subtask execution. In fact, the fundamental change in the assessment of the situation they represent may not be represented in the 55th's plan structure. For such reasons, the determination that they are critical intelligence requirements is more readily handled by treating them as specializations of the larger overall Corps-level PIR.

Thus, we roughly characterize the sources of the division's PIR as coming from (1) preconditions and assumptions built into their own plan about the OPFOR and (2) specializations of the II Corps PIR.

4.1. Steps required in deriving and acquiring PIR

In deriving a plan, the division and the units under it must explicitly incorporate tasks into their plans that insure timely monitoring of PIR. There may be generic approaches to achieving these monitoring tasks such as standard indicators of NBC and standard approaches to detecting those indicators. On the other hand, the monitoring tasks may also need to be adapted to the current situation due to such factors as the terrain and weather conditions. Also, the OPFOR's standard tactics will provide *key indicators* for their behavior. For instance, the operations order notes that *groups of 2-3 BMPs or 2-3 BTRs moving along battalion-sized corridors indicates combat reconnaissance patrol of marching lead regiment's advance guard battalion.*

This suggests the following key steps in automating the derivation and acquisition of PIR

- Gather PIR from own plan's major subtasks, in particular preconditions, assumptions and triggering conditions dependent on OPFOR behavior.
- Gather PIR from the higher echelon plan and specialize according to own plan and responsibilities.
- Derive subtasks for acquiring the PIR and in particular the information required by key indicators.
- Incorporate these tasks into own plan, in part using indicator rules, so that information is available when needed. This may require additional refinement or modification of the original plan.

5. Autonomous Derivation of PIR

The preceding analysis illustrates that a commander must address two sources of PIR: those that arise from higher echelon requirements and those that arise from his own planning process. It is the commander's responsibility to satisfy these requirements, either by satisfying them directly (by pre-planning sensing activities) or by suitably transforming them into guidance for subordinate units (by specifying appropriate PIR in his operations order).

We propose to model these capabilities by extending our existing model of command decision making. Under DARPA's Synthetic Theatre of War (STOW-97) and Advanced Simulation Technology thrust (ASTT) we have been developing advanced command and control simulation technology. The Soar/CFOR command agent architecture provides a domain-independent reasoning engine that supports course-of-action development, execution, and repair and has successfully modeled company and battalion-level autonomous command entities [2,3,4].

5.1. Soar/CFOR Command Agent Architecture

The Soar/CFOR agent architecture is described in greater detail in a companion paper [2]. The system is well-suited to modeling command and control nodes. It incorporates a planning algorithm that is designed to handle the dynamic and collaborative nature of command and control decision making. It builds on AI planning techniques that provide an integrated view of planning, plan execution, and plan repair.

Plans in Soar/CFOR are hierarchically organized sequences of tasks. Each task corresponds to some process and the task description includes (1) initiation or preconditions of the process, (2) completion conditions, (3) interruption conditions, and (4) the responsible entity (who performs the process). Task conditions are used to assess the validity of generated plans and enable the planner to monitor the plan's proper execution.

Three characteristics of Soar/CFOR support collaborative and organizational reasoning. First, Soar/CFOR has the ability to maintain multiple plans in memory and reason about their interactions. This allows a command agent to not only reason about his own activities, but also represent (to some level of detail) the activities of other friendly units and the projected activities of enemy units. Second, Soar/CFOR maintains explicit representations of plan management activities. These are activities that provide structure to the process of planning and implement protocols for how and when distributed planning agents should exchange information. Finally, Soar/CFOR supports the modeling of different management styles. Specifically, a domain modeler can vary the degree to which a C2 planning agent will be cooperative or antagonistic to the activities of other agents.

5.2. Soar/CFOR Capabilities related to PIR

Although the Soar/CFOR architecture does not directly model PIR, its basic capabilities address some aspects of the problem and are readily extended to this type of reasoning. The planner already reasons about the relationship between enemy activities and the planned course of action. For example, when planning a deep attack against an enemy unit in engagement area MOHICAN, a precondition of a successful attack is that the enemy enters EA MOHICAN within the appropriate time window. Since the commander does not have direct control over the enemy behavior, he must rely on intelligence to determine if the precondition will be satisfied. By the criteria described above, this precondition corresponds to a priority intelligence requirement: When will the enemy reach EA MOHICAN? Thus, by analyzing the current plan and collecting preconditions related to enemy behavior, on can automatically collect a set of PIR related to the commander's plans.

In general, PIR are too abstract to observe directly and must be inferred from more primitive information. The CFOR planner bases such inferences on a set of domain-specific indicator rules. These rules match low-level information coming from sensors and status reports and infer higher-level concepts. Some indicator rules base on Krasnovian doctrine might include:
- Groups of 2-3 BMPs and/or BRDMs moving in a stealthy manner INDICATES divisional or regimental reconnaissance.
- Groups of 2-3 BMPs or 2-3 BTRs moving along possible battalion-sized mobility corridors. INDICATES a reconnaissance patrol of a marching lead regiment's advance guard battalion.

5.2. Current Limitations

Two key capabilities must be added to fully model the use and derivation of PIR. First, the CFOR planner does not verify that its intelligence requirements will be satisfied during the course of execution, nor does it attempt to plan for sensing activities related to these requirements. Rather, the planner assumes that this information will

become available as the mission unfolds, for example, via reports from subordinates following standard operating procedures.

Second, the Soar/CFOR currently ignores any PIR it receives from its commanding unit. The only reporting the planner performs are generic reports defined in its standard operating procedures. As discussed above, this is unsatisfactory because the reports may not meet the superior unit's needs. While a filtering tool like SIFT could address part of this problem, it does not address the case where the intelligence requirements can only be met by pre-planned information gathering activities. To satisfy such requirements, we must allow the planner to reason explicitly about how to satisfy intelligence requirements.

5.3. Proposed Solution

Our proposed solution builds upon Soar/CFOR's use of indicator rules. Currently, the planner forward-chains on these rules during plan execution to infer whether any plan preconditions (intelligence requirements) are satisfied. We will extend this to backward-chain on indicator rules during plan generation time.

For each intelligence requirement (provided by higher echelons or generated as a consequence of plan generation), the system will backward-chain on indicator rules to derive a set of possible observations that could satisfy the intelligence requirements. For example, if we wish to know when a regiment's advance guard battalion is approaching EA SEMINOLE, we need to be in position to recognize something approximating a reconnaissance patrol in the vicinity of EA SEMINOLE.

Using domain-specific rules the set of possible observations will be partitioned into two sets: observations that can be satisfied by actions at the current echelon, and observations that can only be satisfied by subordinate units. The former set consists of observations that must be planed for at this level. The latter set becomes the PIR guidance for subordinate units and are added to the CCIR section of the operations order. The observations that must be planned for at the current echelon are handed off to the planner as new goals to be achieved, and they are handled much in the same way as other planning decisions.

6. Conclusion

Simulation-based training is using increasingly complex synthetic forces. As more complex multi-echelon synthetic forces are employed in simulations, the need for a realistic model of their command and control behavior becomes more urgent. In this paper we have discussed one key component of such a model, the autonomous generation and use of priority intelligence requirements within multi-echelon plans.

We have described and characterized the PIR in an example operations order. Based on that characterization, we discussed how the process of deriving PIR could be automated, allowing the derivation of PIR for the different echelons within a multi-echelon organization. Our goal is to realize this automation and use it to extent the capabilities and realism of the automated command entities.

7. Acknowledgements

8. References

[1] R. Calder, R. Carreiro, J. Panagos, G. Vrablik, B. Wise. "Architecture of a Command Forces Command Entity." Proceedings of the Sixth Conference on Computer Generated Forces and Behavioral Representation. Orlando, FL, 1996.

[2] J. Gratch and R. Hill, " Continuous Planning and Collaboration for Command and Control in Joint Synthetic Battlespaces*," Proceedings of the Sixth Conference on Computer Generated Forces and Behavioral Representation, STRICOM-DMSO*, 1999.

[3] R. Hill, J. Chen, J. Gratch, P. Rosenbloom, M. Tambe, "Soar-RWA: Planning, Teamwork, and Intelligent Behavior for Synthetic Rotary Wing Aircraft," *Proceedings of the 7th Conference on Computer Generated Forces & Behavioral Representation*, Orlando, FL., May 12-14, 1998.

[4] R. Hill, J. Chen, J. Gratch, P. Rosenbloom, M. Tambe, "Intelligent Agents for the Synthetic Battlefield: A Company of Rotary Wing Aircraft," *Proceedings of Innovative Applications of Artificial Intelligence* (IAAI-97), Providence, RI, July 1997.

[5] G. Hodge. "Commander's critical information requirements." Army Aviation, 2 21-23.

[6] M. L. McGinnis. CCIR Study, Unpublished technical paper, Operations Research Center, United States Military Academy, West Point, NY, 1996.

[7] J. C. Wetherbe. "Executive Information Requirements: Getting it Right." MIS Quarterly, March 1991.

Author Biographies

JONATHAN GRATCH is a research computer scientist at the Information Sciences Institute, University of Southern California (USC), and a research assistant professor with the computer science department at USC. He has worked for a number of years in the areas of simulation, planning, machine learning, and multi-agent systems. He received his Ph.D. from the University of Illinois.

STACY MARSELLA is a researcher at the Information Sciences Institute (ISI) University of Southern California. He is a member of ISI's Center for Advanced Research in Technology for Education. His work spans several areas of intelligent agent research, including synthetic agent teams as well as pedagogical agents for assessing human and synthetic agent teams. He received his Ph.D. from Rutgers University.

RANDALL W. HILL, JR. is a project leader at the University of Southern California Information Sciences Institute (USC-ISI) and a research assistant professor in the computer science department at USC. He received his B.S. degree from the United States Military Academy at West Point in 1978 and his M.S. and Ph.D. degrees in computer science from USC in 1987 and 1993, respectively. His research interests are in the areas of integrated intelligent systems, cognitive modeling, perception, and intelligent tutoring systems.