CSCI-8940: An Intelligent Decision Aid for Battlefield Communications Network Configuration

W.D. Potter
Department of Computer Science & Artificial Intelligence Center
University of Georgia

Abstract

The Intelligent Decision Aiding Representation System (IDARS) is a frame-based object-oriented platform that is intended to support the intelligent tutoring of U.S. Army planning personnel in battlefield communications network configuration. U.S. Army communications networks are configured using Mobile Subscriber Equipment (MSE) which provides communications support for a typical five-division corps in an area of up to 15,000 square miles. The major goals of IDARS are the development of an Intelligent Tutoring System (ITS) and an Intelligent Decision Aid (IDA) for network configuration using Mobile Subscriber Equipment. Within an ITS or IDA, the module that provides expertise in the target domain is the Expert Module. In the IDARS Network Expert Module, called IDA-NET, we propose to use a Darwinian "natural selection" approach to provide intelligent MSE network configuration.

Introduction

The need for reliable communications support is of paramount importance as evidenced by the military deployment to South Florida in support of hurricane relief efforts there. Mission planners spend many hours developing various scenarios that incorporate all aspects of the mission. The communications aspect of a mission, that is, the development of a battlefield communications network that will support both wire and radio telephone communications requirements, is a particularly difficult task because of the numerous dimensions involved such as what types of components are needed, and how many of each component type is needed.

Network planners achieve expert level capabilities only after extensive training and experience with communications equipment and mission planning. The Intelligent Decision Aiding Representation System (IDARS), a frame-based object-oriented platform, is intended to support the intelligent tutoring of U.S. Army planning personnel in battlefield communications network configuration. Combined with instruction provided by experienced network planners, network planning personnel acquire knowledge of the formal network aspects as well as useful rules of thumb for network planning. The major goals of IDARS are the development of an Intelligent Tutoring System (ITS) and an Intelligent Decision Aid (IDA) for network configuration using Mobile Subscriber Equipment (MSE).

In order to achieve these goals, the single most important module of an ITS/IDA, the Expert Module, must be developed. The Expert Module provides expertise in the target domain (other modules include the Student Module, the Instruction Module, and the Interface Module). The Expert Module provides the primary foundation for student instruction because it manages (solves or knows the solution to) the problem the student is learning to solve. Student remediation as well as remediation strategies are based on the variance between the solution provided by the expert module and the solution developed by the student.

In this report, we discuss the development of a research prototype expert module called the Intelligent Decision Aid Network Expert Module (IDA-NET). In IDA-NET, we use a natural selection approach to provide intelligent MSE network configuration. Based on the pre-defined mission requirements, IDA-NET determines which components are needed to support the mission, and how many of each component type is needed in order to support the communications customers (subscribers). [In a later extension to IDA-NET, the locations (hilltops) where backbone components should be placed in the battlefield will be determined.] Here we use a model of natural selection which has an integer based representation.

Background

During a training session, a student attempts to configure a battlefield communication network using a set of available MSE components (such as Node Centers, which form the backbone of the network, various capacity Extension Nodes, Control Centers, NATO interface nodes, and Radio Access Units). The network must function properly, support the current mission, and satisfy other configuration constraints (such as constraints imposed by U.S. Army doctrine, the terrain, and certain limitations of the components, such as proper connectivity). The instructor may provide guidance to the student during the various stages of network design. This guidance ranges from simple hints (for example, about component compatibility) to detailed analysis (for example, identification of the major misconceptions leading to a faulty network configuration).

From a computational perspective where an intelligent training system assumes the instructor's role, "intelligent" student guidance consists of two primary parts. Part one deals with evaluating intermediate network configurations (i.e., incomplete configurations) in order to provide the student with information on whether or not the partial network could lead to a satisfactory network. In addition, part one deals with determining whether or not a completed network satisfactorily meets all the requirements and constraints defined in the training session. To support part one, the system must be able to solve the network configuration problem. In the Intelligent Tutoring System (ITS) literature, this function is contained in the Expert Module, the primary focus of the project reported here.

Part two involves determining the level of instructor/student interaction appropriate for the student's level of expertise. Interaction levels range from the most basic form of guidance, such as providing basic definitions, to more sophisticated guidance, such as a very abstract hint that would allow the student to independently resolve the current problem. To support part two, the system must be able to: 1) determine (as accurately as possible) the student's level of expertise, and 2) determine the appropriate instructor actions. These functions are contained in what are called the student diagnosis module and the instructional module, respectively. These functions are outside the scope of this project but are being addressed within the overall IDARS effort.

Problem Complexity

Configuring an optimal network from various differing components, component quantities, and locations is the type of problem that is affected by combinatorial explosion. That is, as the number of network components increases, the number of possible network configurations increases drastically. Using a familiar example, *N* network components may be configured into *N!* (N-factorial) networks (i.e., having *N* components in each network). This situation worsens, for example, if we allow varying network sizes, and have numerous location options about where to place each component.

Because of the combinatorial explosion in the number of possible networks, configuration algorithms that are guaranteed to find the optimal network structure via an exhaustive or near exhaustive search are of little value when dealing with networks with over, say, 8 or 10 components because of their extensive runtimes. Since a typical MSE five-division corps contains over 40 Node Centers (not to mention the other types of components), it is clear that we need a good heuristic approach to solve the MSE configuration problem.

Previous Work

In the development of a good heuristic approach, two methods or knowledge-based system approaches are available to us. These are the rule-based (experiential) approach using typical if-then rules, and the functional (deep or associative) approach based on knowledge about the structure and behavior of the network and its components. The functional approach follows the "reasoning from first principles" paradigm. This is quite different from the rule-based approach which is driven by "rules of thumb". Regardless of the approach, the major points of emphasis are the heuristic synthesis of a satisfactory network. Our approach, however, could be considered a combination of the rule-based and functional paradigms (although we do not have a typical collection of if-then rules that are manipulated in some fashion, expert rules-of-thumb are incorporated into IDA-NET). For a network configuration application, synthesis means incorporating a given set of network components and a set of constraints associated with the components (and network) into a network that satisfies the mission goals and constraints. Synthesis can be thought of as the design or planning of a network.

Probably the most famous expert system to be developed for design applications is R1 (XCON) which is used to configure computer systems from customer specifications. An early example of a network design system using heuristics is DESIGNET, developed by Bolt, Beranek, and Newman in the early 1980's. DESIGNET is a rule based design aid that focuses on an iterative user interface approach to network configuration based on the process a designer goes through during the design process.

Examples of expert systems/design aids for network configuration that have been developed more recently include MAPCon, ELAND, and a system for packet adio network design developed by Ruston. MAPCon, developed at the Industrial Technology Institute in Ann Arbor, Michigan, configures communications networks that use the Manufacturing Automation Protocol. It has two main similarities to IDA-NET which include a frame-based representation (as in IDARS), and a combined experiential and functional knowledge approach. However, MAPCon focuses on accurate component parameter settings rather than components, and quantities.

ELAND, developed in Europe, configures distributed information systems operating on a local area network (ELAND stands for Expert Local Area Network Designer). ELAND is slightly different from the other network design systems mentioned here in that it focuses on the specific type of LAN to install, the computer systems needed, and the software needed to support a user's set of networking requirements.

The design system developed by Ruston focuses on the design of high-performance packet radio networks. It proposes changes to the network's parameters that, hopefully, will improve the performance of the network. This assumes that the components and their locations have been pre-specified, and deals solely with network node connectivity issues which is a slightly different problem than that addressed by IDA-NET.

Mobile Subscriber Equipment

U.S. Army communications networks are configured using Mobile Subscriber Equipment which provides communications support for a typical five-division corps in an area of up to 15,000 square miles. The communications system is characterized as an area-switched system which provides both voice and data communications. The MSE network supports radio subscribers using VHF (Very High Frequency) radio links, wire subscribers using hardwired or cabled links, connectivity to other NATO forces, and connectivity to commercial networks.

The backbone of an MSE network is made up of over 40 Node Centers (see Appendix A). Attached to the backbone are various types of Extension Nodes, Control Centers, NATO Interfaces, and Remote Access Units. The differences between the Extension Nodes are directly related to the number of wire subscribers that can be supported and the actual fashion in which they are supported. A Large Extension Node can support nearly 200 wire subscribers. Version 1 of a Small Extension Node can support just under 30 wire subscribers and a Version 2 can support just over 40 wire subscribers. Radio subscribers are supported by Remote Access Units with a capacity of up to 50.

Our Natural Selection Approach

We use a heuristic search routine that is guided by a model of Darwin's theory of natural selection or the survival of the fittest. Here the fittest means the most highly ranked solution in a large solution space. The basic idea behind our search strategy is to generate solutions that converge on the global maximum (i.e., the best solution in the search space) regardless of the "terrain" of the search space. One characteristic of our approach (we hope) is that it is relatively unaffected by hill-climbing or being misled by some local maximum. The key to finding the global maximum lies in the ability to evaluate and compare possible optimal solutions.

The basic operations involved in our natural selection approach (NSA) are: 1) mate selection, 2) crossover, and 3) mutation. The major data structure is an integer string representing the possible solutions. In NSA terms, an integer string corresponds to an individual, and a set of individuals is called a population. The fitness or strength of an individual is computed using some objective or fitness function, and is used to compare an individual with other individuals in the same population. During mate selection, parent strings are stochastically selected, according to their fitness, from the current population. Then, parent strings are "mated" via crossover to produce offspring for the next generation. Fitter parents contribute more offspring to the next generation than weaker parents because they have a higher probability of being selected for mating. This is the step that models the process of natural selection in nature.

Crossover, the second operation, determines the characteristics of a child or next generation individual. In nature, children inherit good as well as bad features of their parents in varying degrees of dominance. Crossover performs this same function in our NSA. One of the simplest crossover approaches is to split each parent string at the same randomly chosen location and swap their tail sections. This insures a certain amount of inheritance and ideally, the good/strong features will dominate the children. The inheritance of features that produce stronger children throughout the generations is the source of the ability to converge on the global maximum in a relatively short time. However, our engineers have developed a alternate technique called uniform crossover which they want us to use.

The last basic operation is called mutation. Mutation is that extremely rare "glitch" in the inheritance mechanism that introduces or modifies some feature with unpredictable consequences. Mutation occurs immediately after the creation of a next generation individual yet before the next generation has become static. Once the new generation becomes static, we move forward in order for it to become the new current generation. Ideally, mutants would contain some useful features that may have been inadvertantly lost in earlier generations.

Each position in the sequence corresponds to a particular component type, and the values of the integers range between the minimum and maximum Corps complement for the component type. Because of the special nature of the NSA individuals, we will use a uniform crossover technique instead of a more typical crossover technique such as one- or two-point crossover. Briefly, uniform crossover creates new generation individuals by randomly selecting sequence elements from either one of the two parent sequences. These random selections are positional so that the new individual's elements are within the proper range for the corresponding component. Another scheme, called average crossover, takes the average of the parent sequence elements for use as the new generation individual sequence elements. This method also insures positional range constraint satisfaction. Although we will not be using average crossover, we should plan for just in case the engineers want us to fit it in.

One item that we want to experiment with is the use of an "elitist" policy during the evolutionary process from one generation to another. The elitist policy that we follow conditionally seeds a new generation if the fitness of the weakest next generation individual is less than the fitness of the strongest current generation individual. In this case, we simply replace the new weakest individual with the current strongest individual. However, we have not observed any noticeable improvement in experiments so far.

Strength Of An Individual

As we have seen, our approach uses a model of natural selection to converge on an optimal or near optimal solution. As with any heuristic approach, the NSA is not guaranteed to find the optimal solution. However, due to its functional nature, we hope it finds the optimal or near-optimal solution. The model mechanics involve the crossover, mutation, and mate selection operators. The primary driving force behind the GA is the mechanism which determines the strength of each individual, called a fitness or objective function.

The NSA objective function takes into consideration several important issues in determining the "fitness" of an individual (solution). These issues include the mission requirements, the known support capacities of all the various components, the Division and Corps equipment complement, and the known minimum component requirements (such as, any network must have at least one Control Center) into consideration while determining the objective function value. In addition, issues such as the combined total number of components, and connectivity doctrine are included to help distinguish competing individual solutions. Therefore, we have several goals for the objective function:

1) minimize the total number of components,
2) avoid violating any of the minimum or maximum constraints,
3) avoid violating connectivity constraints,
4) minimize any penalty incurred due to not satisfying the mission requirements,
5) minimize any penalty incurred due to providing too much support, and
6) optimize the mix of abundant components with the more sparse types of components.

This last goal deals with the actual complement of the derived component list. The issue here is to attempt to use "cheaper" or "more plentiful" components instead of unnecessarily deploying the more "expensive" types of components. For example, a version two Small Extension Node supports more wire subscribers than a version one (a version two also requires more personnel to operate and deploy than a version one). However, the Corps complement of version twos is much smaller than that of version ones. This leads to heuristics developed for the judicious use of version twos (the more "expensive" type of component). Goal 6 is an attempt to balance the mix of version ones with version twos, for example.

Penalties are incurred for the above issues in a graduated fashion; small violations incur small penalties and large violations incur large penalties. Then, they are combined to determine the fitness, a product of several terms based on the above goals, of an individual solution.

Based on Node Center quantities, connectivity capacities are determined by the MSE planning strategy. For example, within a Division, Node Centers should have complete connectivity. However, connectivity between adjacent Divisions is not complete, as is the case with connectivity between Corps Node Centers and Divisions. Also, Node Centers are limited in the number of connections that can be made to other Node Centers while still providing connectivity to other MSE components. Again, a graduated penalty is incurred.

Appendix A

MSE Components

Component	Qty	Wire	Radio	Antenna
Node Centers	42	24	0	12-UHF (V3)
Large Extension Node	9	176	0	2-UHF (V4)
Small Extension Node 1	168	26	0	1-UHF (V1)
Small Extension Node 2	56	41	0	1-UHF (V1)
System Control Center	7	0	0	0 (cable)
Remote Access Unit	92	0	25	1-UHF (V1)
NATO Interface	4	0	0	1-UHF (V2)
UHF Links	0	0	0	0

Component Structure:

NC LEN SEN-1 SEN-2 SCC RAU NAI