Marine Corps Efforts to Harness Massive Computing Power

Capt John P. Kresho, USMC
Marine Corps Combat Development Command

Today's advances in agent-based, or complex adaptive systems, models (that we often refer to as distillations) have given us the potential to capture some of the adaptability and other key factors inherent in conflict. However, to explore even a single question of interest, a couple of data points are often insufficient given the landscape of possible outcomes within a given scenario. Thus, in order to gain insight into the question at hand, it is necessary to produce thousands, hundreds of thousands, or even millions of data points.

Initially, a single user would create a scenario with a distillation on a personal computer. The user would then run the model many times to explore how appropriately the distillation world seemed to represent the real world he/she was envisioning. The amount of data points depended on the patience of the user because his personal computer could NOT be used to do other tasks while the model was executing. The user manually changed input parameters to explore how well the scenario was targeting the question.

Given the aforementioned situation where the user is changing 5 parameters, each over a range of 2 values, the number of combinations is 32, as illustrated in Figure 1. Multiplying this by 2 random seeds for each parameter set, the user must execute the model 64 times on his personal computer. If everything works as planned and the scenario is suitable, the user would then like to execute the distillation over the same 5 parameters, but each with a

range of 10 different values. Additionally, each parameter set could be run with 5 different random seeds. This turns out to be 500,000 runs, too many for a single computer.

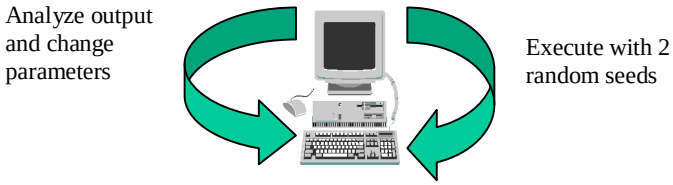


```
For Parameter_1 = 1 To
2
 For Parameter_2 = 1 To 2
  For Parameter_3 = 1 To 2
   For Parameter_4 = 1 To 2
    For Parameter_5 = 1 To 2
     EXECUTE MODEL
```

Figure 1

In order to produce the required amount of data points, the Marine Corps attained support from the Maui High Performance Computing Center (MHPCC). MHPCC's resources include IBM mainframe computers that are capable of running distillation models millions of times without human intervention. Experienced programmers within MHPCC are able to write computer code that will execute a given model on the numerous computer processors over a wide range of input parameters. The user can thus send a particular scenario, along with the range of input parameters, to MHPCC, as illustrated in Figure 2.

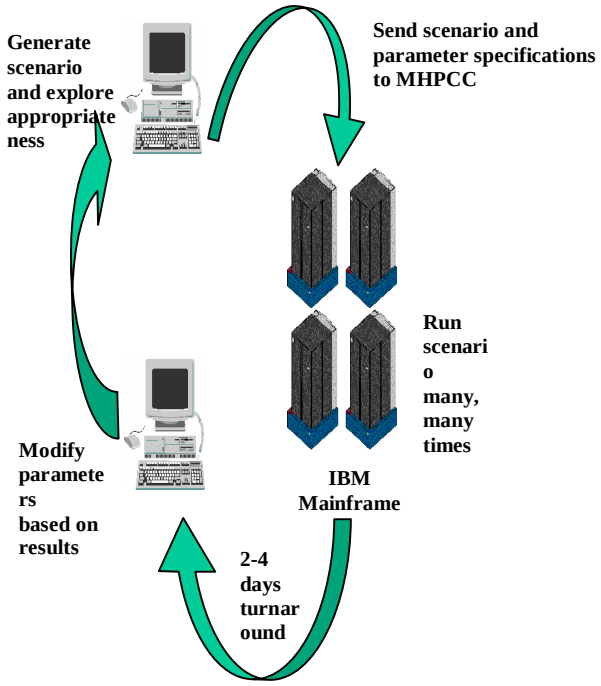

Figure 2

MHPCC's IBM Mainframe resources provide excellent results for very large problems, but may not be required for scenarios of lesser magnitude. In order to provide the Marine Corps with a self-contained "supercomputer" type platform, they turned to a PC based network of computers. The Marine Corps already possesses a large network of Microsoft Windows NT™ Servers and Workstations. In conjunction with MHPCC, a solution was sought to take advantage of the Marine Corps' available hardware assets and its existing knowledge of the Microsoft Windows NT™ Operating System.

As a result, MHPCC and the Marine Corps fielded a cluster of 8 Windows NT workstations, each with 2 Intel Pentium III

processors. This cluster of workstations was termed “Gilgamesh” and provides the Marine Corps with an in-house “supercomputer”. Housed at Studies and Analysis, Marine Corps Combat Development Command, Quantico, VA, Gilgamesh provides the Marine Corps analysts the ability to develop distillation scenarios and produce over 100,000 data points within 24 hours.

Wishing to build on the success of Gilgamesh, the Marine Corps continued its quest to increase its computing power. Within the Marine Combat Development Command, there are over 250 computers that have several hours of unused processor cycles each day. If only these unused processor cycles could be used to execute agent-based models, the Marine Corps could increase its “supercomputing” power tenfold.

From further research, MHPCC recommended the use of Condor, a product developed by the Computer Science Department of the University of Wisconsin. Condor allows an organization to take advantage of its collection of computer resources.

As shown in Figure 3, a single Microsoft Windows NT™ workstation serves as the Condor Cluster Controller. After registering with the Condor Controller, any computer within the Marine Corps Combat Development Center can be used to execute an agent-based model. The Condor Controller continually tracks the computers that have available processor cycles, including those in the Gilgamesh cluster. As work is submitted, Condor queues each model run and distributes it to an available computer. Condor also stores any output created in a central location for later analysis. At the time this article was written, the Marine Corps’ Condor cluster included over 20 workstations. By the middle of 2001, 50-100 workstations are planned to be part of the Condor cluster.

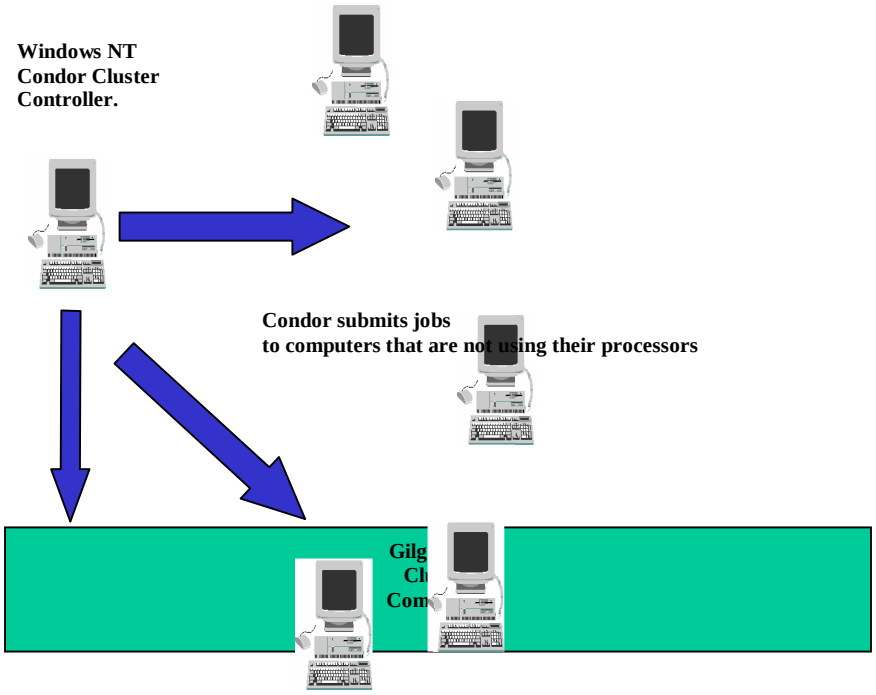

Figure 3

With the increasing speed of processors and the declining prices of hardware, the Marine Corps has been able to purchase a massive amount of computer power for each staff member within the Marine Corps Combat Development Command. The combination of this new computing power and advances in software has allowed the Marine Corps' analytical community to harness and focus this technology toward meeting the need of producing enormous amounts of data. Through this combination as well as Gilgamesh and supercomputers, the computing power is now available to allow us to explore vast possibility spaces in our efforts to answer questions.