07ICES-240

Testing Heuristic Tools for Life Support System Analysis

Luis F. Rodríguez, Haibei Jiang
Department of Agricultural and Biological Engineering, University of Illinois at Urbana-Champaign, IL 61801

Scott Bell, David Kortenkamp
NASA-Johnson Space Center, Houston TX 77058

ABSTRACT

BioSim is a simulation tool which captures many basic life support functions in an integrated simulation. Conventional analyses can not efficiently consider all possible life support system configurations. Heuristic approaches are a possible alternative. In an effort to demonstrate efficacy, a validating experiment is designed to compare the configurational optima discovered by heuristic approaches and an analytical approach. Thus far, it is clear that a genetic algorithm finds reasonable optima, although improved fitness function shall be required. Further, despite a tight analytical fit to data, optimization produces disparate results which will require further validation.

INTRODUCTION

NASA has recently redirected their efforts towards the development of a new vehicle, in replacement of the Space Shuttle. The recent moves are in response to the directives of the President towards a 7th Lunar landing (Bush 2004). The current suggestion is that Lunar exploration will commence by 2020 in preparation for Martian exploration. A Lunar Outpost has been described at the 2nd Exploration Conference (NASA 2006). It is likely that several consecutive missions can be strung together making total mission length on the order of 10 years, or more. Alternating crews would regularly travel to and from the outpost on a regular basis, possibly every 6 months. Consecutive trips would bring crews, resupply materials, and additional habitat resources.

As mission length increases regenerative life support systems become preferred to the resupply based systems featured on the Space Shuttle and, to a lesser extent, on the International Space Station. This suggests that the life support system shall rely more heavily on local and recycled resources than on terrestrial resources. To what extent this will be the case will depend on NASA's ability to develop cost effective regenerative technology within the current mission time frame.

In any case, as systems become more tightly constrained the necessity to understand their reliability and robustness becomes critical. Tight constraints are inherent in all NASA missions, but long-term human missions present special challenges. This is inevitable due to the extremely high mission costs, the impacts of high profile accidents upon the agency, the limitation of abort-to-Earth options available, and the unavoidable risks involved in Lunar and Martian exploration. The ongoing research described here aims to manage these risks to the extent possible.

The challenges described here involve the study of complex, multi-objective, stochastic systems. Comprehensive study of the search spaces described by these life support systems is infeasible for dynamic, non-linear simulations and therefore heuristic approaches are suggested. These have the advantage of intelligently considering the array of potential solutions based on the input of the analysts. A genetic algorithm and an analytical approach have been implemented for the purpose of optimization of a sample life support system. An ant colony optimization approach will be implemented in the near future. Once this is completed two heuristic approaches and an analytical approach will be tested to determine whether analogous results can be obtained via these alternative approaches. Preliminary results are presented here describing progress with the genetic algorithms and the analytical approach.

MISSION SCENARIO

A simplified life support system has been selected for the purpose of testing and validation. The Early Human Testing Initiative (EHTI) Phase I (later known as the Lunar-Mars Life Support Testing Project) involved a single crew member enclosed within a vacuum pressure growth chamber with a tray of hydroponically grown wheat. The baseline objective of the experiment was to demonstrate that the wheat-human system could be utilized to manage the atmospheric quality (Edeen and Barta, 1996; Lane, Sauer, and Feeback, 2002). A diagram of the basic system modeled here is included in Figure 1.

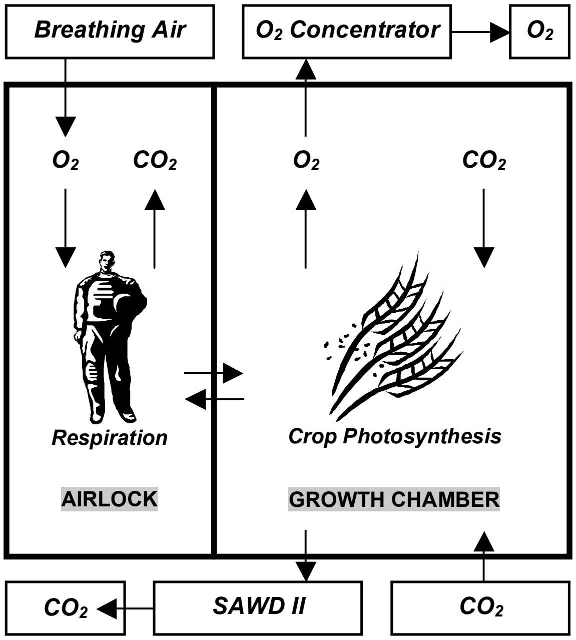

Figure 1. Schematic representation of oxygen and carbon dioxide exchange during the test.

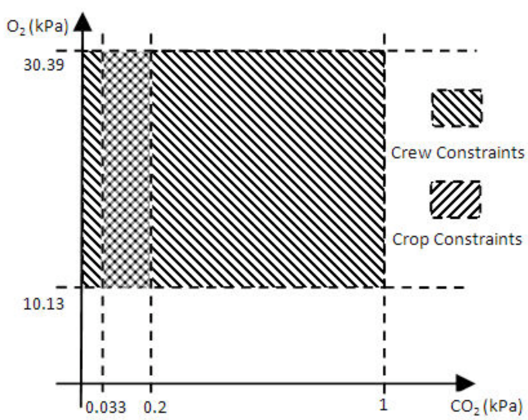

Figure 2. Atmospheric constraints on workable system states

Successful mission operations are limited to areas where the models utilized are validated and where atmospheric conditions are safe for the crewmember. The workable area is show in Figure 2. All simulations are allowed to proceed as long as the atmospheric state variables remain within the areas shown. The length of a simulated mission is the key in determining the quality of a configuration. This process is described further in the section entitled *Multi-objective Utility Function Design*.

BIOSIM - AN INTEGRATED LIFE SUPPORT SIMULATION

Over the past several years NASA has been developing an integrated life support system simulation [5]. The simulation was developed in accordance with NASA requirements and baseline assumptions for the design of an ALS [6,7]. The simulation includes detailed, stochastic models of the crew, air, water, biomass (including plant growth chambers), power, food production and solid waste recycling. Each of these components interconnects with the rest of the simulation as shown in Figure 3. The simulation can be configured to simulate a wide variety of different life support systems. This includes number, gender and ages of crew members, the size of the habitat environments, atmospheric pressure, capacities of tanks, initial levels of consumables, processing capacity of life support modules and many other variables. The simulation has sensors and actuators that connect to various controllable elements and allow for real-time control. Sensors read simulation values and can inject sensor errors. Actuators set flow rates of resources between simulation components.

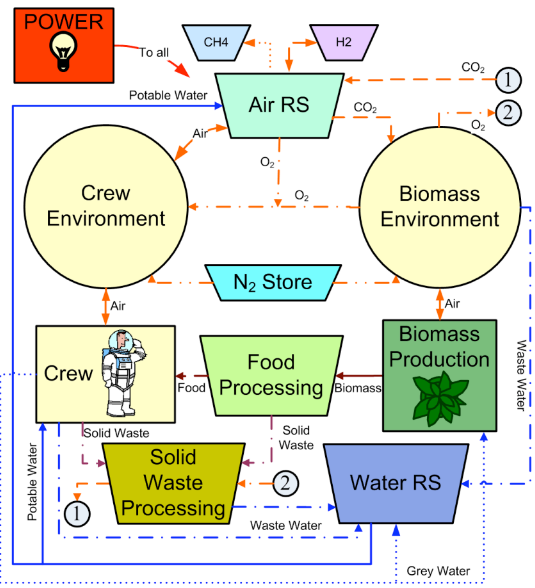

Figure 3. Typical BioSim configuration

CONFIGURATION DESCRIPTION

Five configurations, summarized in Table 1, related to the original EHTI experiment have been designed for the purpose of this study. For the most part, each configuration progressively creates a more challenging optimization problem, primarily due to the additional combinations considered with the addition of new decision variables.

In all configurations, a crop is sown on the initial simulation day. This crop is allowed to establish itself prior to the arrival of the crew member. During this time a rudimentary on/off control system is utilized to maintain carbon dioxide and oxygen partial pressures and total atmospheric pressure. Additional control variables are thus involved which becomes our interest to discover the optimal. All variables are summarized in Table 2.

Table 1. Configurations under consideration

	Decision Variables					
Config.	Area	Crop	Mix	Vol.	Control	Arrival

						Time
1	*				*	*
2	*	*			*	*
3			*		*	*
4	*		*		*	*
5	*		*	*	*	*

The crew arrival time is limited by the harvest date of the crops available, enforcing that they shall be present during harvest. The underlying assumption is that automation will be available for sowing crop, but not for harvesting. For optimization purposes, the arrival time of the crew and several set point values related to the control system are selected in each configuration studied.

Starred items in Table 1 demonstrate the decision variables included in each optimization study. For example, in the first configuration, a fixed volume of 32,000 liters is maintained and in addition to the control variables included in each configuration, the optimization techniques are asked to identify an optimal area of wheat crop.

In configuration 2, the additional complication is the selection of an optimal in addition to the area. Nine crop options are available within BioSim: dry bean, lettuce, rice, soybean, sweet potato, tomato, wheat, and white potato. Configuration 3, progresses from here selecting a mixture of crops, with a fixed area of 14 m^2. Similarly, configuration 4 considers total crop area as well as the crop mix. Finally, configuration 5 considers the system volume.

Table 2 lists the discretized ranges considered by the optimization tools. These ranges have been selected to allow the optimization tools latitude to find novel solutions, which may not have been immediately obvious to the analysts. Each range is discretized by the variable type utilized. Thus, integer variables are selected from whole numbers within the range, whereas double variables are divided up to the 16th decimal place depending on the range and the amount of digits (of 16 total) appearing before the decimal place.

EXPERIMENTAL DESIGN CONSIDERATIONS

The objective is to determine whether heuristic approaches identify optima which are as good as, better, or worse than those derived from analytical approaches. It is critical to consider the experimental factors that need to be controlled to ensure the results are comparable. These factors include things like the total number of configurations searched or the range of search spaces. The following section discusses each factor and the corresponding rationale.

Table 2. Decision variables under consideration

Attribute		Range	Variable Type
Crop Area		0-200 m^2	Double
Controller Attributes	CO_2 Level 1 Time	0-100 hr	Integer
	CO_2 Level 2 Time	100-300 hr	Integer
	CO_2 Level 3 Time	300-500 hr	Integer
	CO_2 Level 1, 2, 3 Set Point	0.01-0.1 kPa	Double
	CO_2 1, 2, 3, Total Pressure Low Rate	0-2 moles/s	Double
	CO_2 1, 2, 3, Total Pressure High Rate	0-1 moles/s	Double
	O_2 Set Point	20-30 kPa	
	O_2 Low Rate	0-10 moles/s	
	O_2 High Rate	90-100 moles/s	
Crew Arrival Time		0-504 hr	Integer
Crops		n/a	Species
Volume		20-500 m^3	Double

TOTAL NUMBER OF CONFIGURATIONS SEARCHED

The ideal condition is to provide each optimization method with an equivalent opportunity to observe BioSim performance. Based on observation of the simulation results and the consideration of computational efficiency, the number is set to be 5000. That is, unless the optimization tool successfully locates the optimal using less than 5000 configurations, the operator will terminate the search automatically after the 5000th trial.

OPERATOR CONTROLLED SEARCH LENGTH

With some heuristic techniques, the operator identifies the simulation end conditions. This is offered as opposed the maximum number of simulations considered. Operator skill at managing the search and identifying true optima is at question here.

ORIGINS OF SEARCHES

The origin of search may have great impact upon heuristic search performance especially when the path to the optimal is critical. It is also important for some regression procedures.

PSEUDO RANDOM NUMBER GENERATORS AND INITIAL SEEDS

Random number generators only approximate truly random numbers. Multiple pseudo-random number generators can be tested to determine the impact of this effect. Further, typically the seed utilized in those number generators is the system clock. The effect of the initial seed can also be considered.

DISCRETIZATION OF ATTRIBUTES

Attributes are discretized by using different variable types, such as Double and Integer. It might be possible that a search algorithm can inherently benefit from such a design, and alternative variable types shall be considered.

MULTI-OBJECTIVE UTILITY FUNCTION DESIGN

A significant challenge in the experiment design is the design of an appropriate fitness function. It needs to be well rounded and capable of balancing the defined objectives, including:

1) Maximizing crew survival length, which is regarded as a proxy for system reliability;
2) Minimizing crop area; space is at a premium within the Lunar or Martian mission and crop production is a major ESM burden;
3) Minimizing crew arrival time, which will directly affect mission length, cost, and complexity. Such a design will also reduce the challenge of designing a completely autonomous crop production system;
4) Minimizing total system volume, which consists of a fixed Airlock volume (20 m^3) [2]; and the volume of the crop growth chamber, which is decided dynamically by the product of total crop area and growth chamber's height. Based on the design by Fortson, Castillo, and Barta [8], the height is uniformly distributed between the minimal feasible height (0.8 *m*) to the maximum height (2.4 *m*) suggested by Edeen and Barta [2].

A fuzzy membership function approach has been proposed to manage these objectives, they are equally judged on a 10-point scale.

U_T, U_A, U_C, and U_V are utility contributions from mission length, crop area, crew arrival time, and system volume, respectively.

– Mission length utility contribution

$$U_T = 0 \text{, if } T_S < 0$$

$$U_T = \frac{10}{336} T_S \text{, if } 0 \le T_S \le 336$$

$$U_T = 10 \text{, if } T_S > 336$$

– Total crop area utility contribution

$$U_A = 10 - \frac{10}{200} A_T$$

– Crew arrival time utility contribution

$$U_C = 10 - \frac{10}{504} T_A$$

– Total system volume utility contribution

$$U_V = 0 \text{, if } V_T < 20$$

$$U_V = 10 \text{, if } 20 \le V_T \le 50$$

$$U_V = -\frac{10}{450} V_T + 11.11 \text{, if } 50 \le V_T \le 500$$

where T_T is the mission length, measured in hours, simulation time; T_C is the crew arrival time; T_S which equals to $T_T - T_C$ indicates the actual mission length after the crew arrives; A_T is the total crop area; V_T is the total system volume.

And the overall utility function thus becomes

$$U = \omega_T U_T + \omega_A U_A + \omega_C U_C + \omega_V U_V$$

where ω_T, ω_A, ω_C, ω_V are correspondingly the weight of each contribution factor. In the currently experiment, the weight values are all set to be one suggesting each term carries equal weight. The utility function is a unit-less measure of configuration performance and the relationship between the factors are illustrated using Figure 4, 5, 6, and 7 respectively.

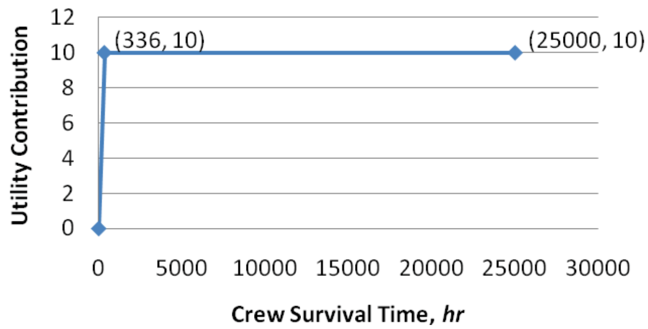

Figure 4. Crew Survival Time vs. Utility Contribution

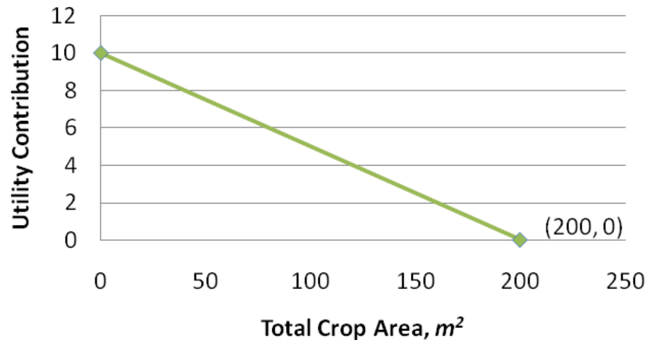

Figure 5. Total Crop Area vs. Utility Contribution

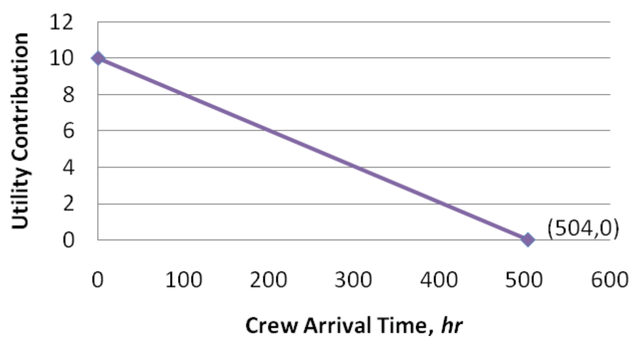

Figure 6. Crew Arrival Time vs. Utility Contribution

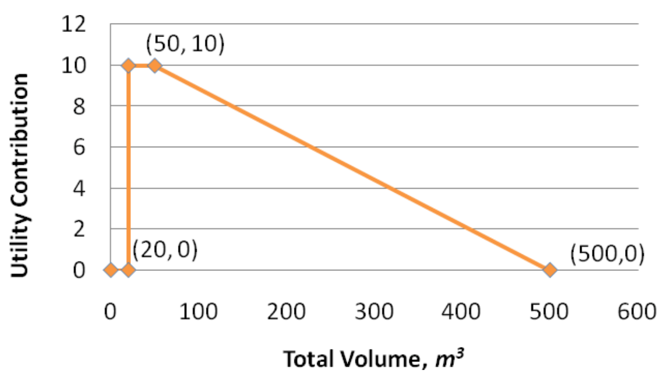

Figure 7. Total Volume vs. Utility Contribution

OPTIMIZATION APPROACHES

GENETIC ALGORITHM

Genetic algorithm is a computational heuristic algorithm based on the theory of natural selection, utilized for solving optimization problems. Genes encode inputs into the problem being optimized; in this case they represent life support configurations. The quality of each gene is judged by the selected fitness function. Natural Selection is simulated with the initialization of a population selected randomly and then selecting the best performing configurations and using these as the seeds for the next generation via simulated genetic crossover, mutation repetitively.

In our validation experiment, a genetic algorithm was integrated by using an existing open source Java genetic algorithm package JGAP[9] with BioSim simulation program. In the simplest configuration, there are 19 genes, while the most complex one will have 38. The population size has been selected to be five, with a mutation rate of fifteen.

JGAP also provides a user friendly package JGrid that allows distributed parallel computing which leverages additional computational efficiency. JGrid was implemented using four DELL workstations to simulate these configurations. Each workstation serves as a Worker, while one of them is also responsible for feeding the simulation problem as the Server, distributing work load over the network and compiling gathered results.

ANALYTICAL APPROACH

To validate the results obtained from heuristic approach, an analytical approach was developed here. BioSim performance and its relation to the utility function is converted to a nonlinear model, and thus enables classic nonlinear programming techniques to locate the optimal solution.

Randomly generated configurations are executed by the simulation tool which outputs the corresponding mission length information. Uniform random numbers are selected over for each attribute over the ranges specified in Table 2. A total number of 5000 configurations are tested as in the heuristic approaches.

SAS PROC NLIN procedure is a powerful tool which enables fitting these data into a nonlinear model. Various tests have been considered for goodness of fit for different nonlinear forms. Polynomial and exponential forms are very flexible in fitting the shape of many types of data. SAS PROC NLIN calculates the initial residual sum of squares for a predefined array of starting values and initiates the iterative regression with the best set. The objective of the iterations is the maximization of the R^2 quality of fit statistically. Variable selection is another key step towards finding a reasonable model. There are *forwardselect*, *backwardselect* and *stepselect* options available within SAS for the regression procedure, where a basic p-test is used to determine the importance of each attribute and add/eliminate the important/unimportant ones from the null/full model. The best fit is identified by comparing residual sum of squares, residual plot and the residual normal quartile plot.

A Matlab optimization function, called *fmincon*, is used here to perform the minimization of the fitted nonlinear model as an objective function. Maximization is achieved by supplying the routines with *-F(x)*. All of the constraints existing in our problems are inequalities defined by the attribute bounds. Independent M-files are coded to describe the objective function and the constraints for each configuration.

ANT COLONY OPTIMIZATION

Ant Colony Optimization is another heuristic optimization algorithm mimicking the foraging behaviour of ant colonies. In nature, ants initially walk randomly from their nest to search for food. When a path to food is found, they return to their nest leaving pheromone trails on the ground for other ants. These pheromone trails increase the odds that other ants will follow the same path to the food. As other ants follow the trail, they too drop pheromone reinforcing the path. Given time, the trail gradually evaporates reducing its attractiveness to ants. Therefore a short trail is more attractive than a longer trail, as the pheromone trail is refreshed as there is less time between the nest and the food. Furthermore, the evaporation causes the ants to explore for food rather than simply following previous paths.

Instead of looking for food, our synthetic ants are looking for the optimal configuration of BioSim. Our space is discretized into 2 dimensions, X being the attribute index, Y being the attribute value. For example, if attribute #3 is crop size and its value is constrained from 0 to 200, a possible point in space is (3,140). Each point has a pheromone value associated with it. Our synthetic ants start by picking a random value for attribute 1, then for attribute 2, etc. This creates a path through the space and a configuration for BioSim. This configuration is then simulated in BioSim. As an example of the process, let's say we'd like to optimize 4 variables in BioSim:

1. amount of water (0-500 *litres*)
2. amount of food (0-150 *kg*),
3. size of the habitat (10-1000 m^2)
4. number of crew (1-4)

An ant selects the following path: (1, 89) (2, 123) (3, 721) (4, 2). This path can be translated into a BioSim configuration by setting the water to 89 litres, the food to 123 kg, the habitat to 721 m^2, and adding 2 crew members. This configuration runs for 1200 hours. This translates into a certain amount of pheromone that the ant then deposits on each point on each point back to its starting point. A longer simulation run means the ant will drop more pheromone on the same way back to its starting point. This process continues for this ant and other ants in parallel until convergence to a near-optimal solution.

PRELIMINARY RESULTS AND DISCUSSION

Genetic algorithm analysis results for configurations one and two are currently available and are pictured in Figure 9 and Figure 10, found in the *APPENDIX*. The horizontal on both figures is the currently simulated trial, a total of 5,000 as specified previously. On the vertical is the utility for each corresponding simulation instance, depicted by the points within the chart. The thick line at the top depicts the currently simulated best configuration. The thick line within the scatter is the running average of utility. Figure 9 and Figure 10 are prototypical GA output generated dynamically as an optimization is running.

Configuration one finds an optima of 20 after roughly 1,200 trials. In configuration one, three aspects are considered within the fitness function each with a maximum value of 10: crew survival time, crew arrival date, and crop area. An output of 20 is a strong suggestion that the GA has found that sacrificing on one aspect of the fitness function, while favoring the remaining two will provide benefits. However, the actual result may not be grounded in reasonable truths. In fact this is the case: crop area selected was approximately 181 m^2, arrival time was 37 hours, and crew survival time was 25,000 hours, the maximum. These are each rewarded by the fuzzy fitness functions as 0.97, 9.2, and 10, respectively.

To combat this, a revised fitness function, with respect to crew survival has been proposed and shall be utilized in future analyses. This function is as follows and is depicted graphically in Figure 8: The originally proposed fitness function was selected to provide proxy for high reliability, but it seems the previous form was excessive. Noting the two slopes prior to $T_S = 336$ shows that rewards will be great for long simulations, but progressively worse as survival times get large for this scenario.

$$U_T = \frac{10}{336} T_S \text{ , if } 0 \le T_S \le 336$$

$$U_T = -\frac{10}{24{,}664} T_S + 10.25 \text{ , if } T_S > 336$$

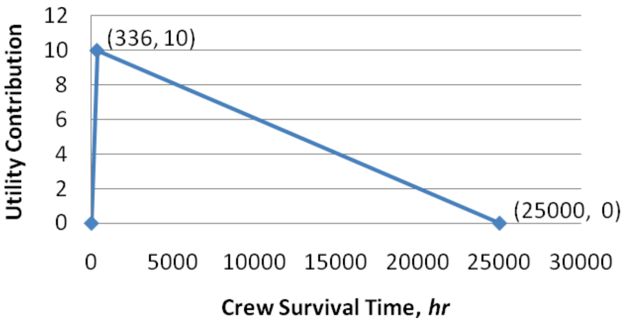

Figure 8. Newly proposed Crew Survival Time vs. Utility Contribution

Despite the challenges with the fitness function, there is no doubt that the GA is finding optima within the search space. Configuration one seems to have converged after roughly 1,200 trials. Configuration two seems to have converged after 2,500 trials onto a configuration effectively identical to that of configuration one: crop area of 189 m^2, crew arrival date of 27 hours, and total simulation length of 25,000 hours, and a crop type of wheat—the same as in configuration one. This provides some validation to the results in configuration one. It is reasonable as well that wheat was the selected crop in configuration two as this is the most photosynthetically active crop per unit area considered.

With the random scatter of simulated instances in both Figure 9 and Figure 10, it is very likely that the entire search space was effectively considered in both cases. However, the additional complication of enforcing the genetic algorithm select the optimal crop slightly more than doubles the period of time required to find an optima. This is an encouraging result as there is an effective nine-fold increase in the total volume of the search space.

The quality of fit for the analytical approach is shown in Figure 11, Figure 12, and Figure 13. A residual plot is shown in Figure 11, Figure 12 is a residual probability plot, and Figure 13 is a plot showing predicted data from the model with actual data from the BioSim simulations. Although non-linear approaches were the initial focus, a linear model has been fit with highest precision to the data and has been selected for this analysis. (For brevity, this full model has not been shown here.)

Optimization of the linear model, however, did not produce the anticipated results. An optima of 798 was found, which certainly does not fall within the range of 0-30 specified by the fuzzy fitness function. This is odd considering all data utilized for fitting had utility of 20 or less. The optimal configuration parameters are: crop area of 200 m^2, a crew arrival time of 504 hours, and dry bean as the selected crop. (Remaining attributes are not listed for brevity.) As of yet this result is not fully explained by our analyses and futher work is required. Clearly, with such large area and crew arrival time it is impossible for the fitness to be much larget than ten despite a long simulation length. Current suggestions for discrepency relate to the fact that all attributes listed in Table 2 are incorporated into the analytical model, rather than focusing principal components. However, by reducing the model to any principal components, it will be, at best, a challenging iterative problem to identify the truly optimal configuration from the remaining partial information. In any case, it is clear that further validation of any analytical model will be required for use later.

Interestingly, by simply scanning the raw data produced randomly during the data generation phase shows optima similar to those identified by the GA. This furher supports the assertion that the GA effectively considered the entire search space. However, it may suggest that the additional overhead of implementing the GA is not necessary, provided enough computaional power, and time, is available to take the brute force approach to finding an optima. Effectively each run of 5,000 simulations consumed all resources from a Dell Precision 380 machine for approximately five hours time. JGrid, though available, have not yet been utlilzed in this work, but this lab may find an effective three-fold drop in processing time, with three such machines. High performance cluster and supercomputing options have been considered for this work as well, though few a capable of implementing applications requiring the Java Virtual Machine, as in the case of BioSim.

CONCLUSIONS AND FUTURE WORK

The paper summarizes the progress made on the validation of heuristic search techniques. An experiment comparing the results from our heuristic search tool and an analytical optimization method was conducted considering the design of a simplified closed-loop life support system. Thus far, indicators suggest that with a careful design of a multi-objective utility function, a GA can find reasonable optima. However, analytical approaches require further validation, and possibly iterative solutions.

Both approaches locate a number of solutions that are reasonably workable according to mission objectives due to the ability to effectively consider the entire search space.

Current results suggest that improvements remain in the development of the fitness function. Several remaining issues, such as, the censoring of simulation length, the number of configurations tested, and the weight for each term in the utility function are being discussed for future implementation.

New research opportunities have been identified as this work continues. The genetic algorithm results and the fitted model both show that the workable results are possibly clustered at several tiers of utility. It may potentially be beneficial to mission designers if an understanding of the characteristics of these simulations which cluster around utility tiers. Regarding further development of the utility function, adding weight to each term has been proposed. This may lead to better optimization performance and improved comparability between the approaches. For example, the notion of ESM could easily be introduced to the weights. Regarding the analytical approach, other nonlinear and linear models will need to be tested using the proposed approach to determine if improving the goodness of fit is possible. And finally, the optima selected analytically should be tested for its performance in BioSim for validation.

ACKNOWLEDGMENTS

The work has been funded by NASA grant number NNJ06HA03G.

REFERENCES

1. Bush, G.W., President Bush Announces New Vision for Space Exploration Program, http://www.whitehouse.gov/news/releases/2004/01/20040114-1.html, accessed 3 March 2005, 2004.
2. Edeen, M., and Barta, D., "Early Human Testing Initiative Phase 1,"NASA Johnson Space Center Technical Rept. JSC-33636, 1996.
3. Lane, H.W., Sauer, R.L., and Feeback, D.L. (ed.), *Isolation: NASA Experiments in Closed-Environment Living,* Univelt, Inc., 2002.
4. NASA 2nd Space Exploration Conference, http://www.nasa.gov/mission_pages/exploration/main/2nd_exploration_conf.html, accessed: 25 February 2007, 2006.
5. D. Kortenkamp, S. Bell, Simulating Advanced Life Support Systems for Integrated Controls Research, Proc. Third International Conference on Environmental Systems, 2003.
6. JSC Advanced Life Support Requirements Document JSC-38571B available at http://advlifesupport.jsc.nasa.gov
7. JSC Advanced Life Support Baseline Values and Assumptions Document available at http://advlifesupport.jsc.nasa.gov
8. Fortson, Russ, and Castillo, Juan, Design of Plant Growth Chambers for NASA's Advanced Life Support Program at the Johnson Space Center, Presented at the ASAE Annual Meeting, ASAE paper no. 974026, 1997.
9. JGAP, http://jgap.sourceforge.net/

CONTACT

Questions and comments regarding this work may be addressed to Luis F. Rodríguez at lfr@uiuc.edu.

APPENDIX

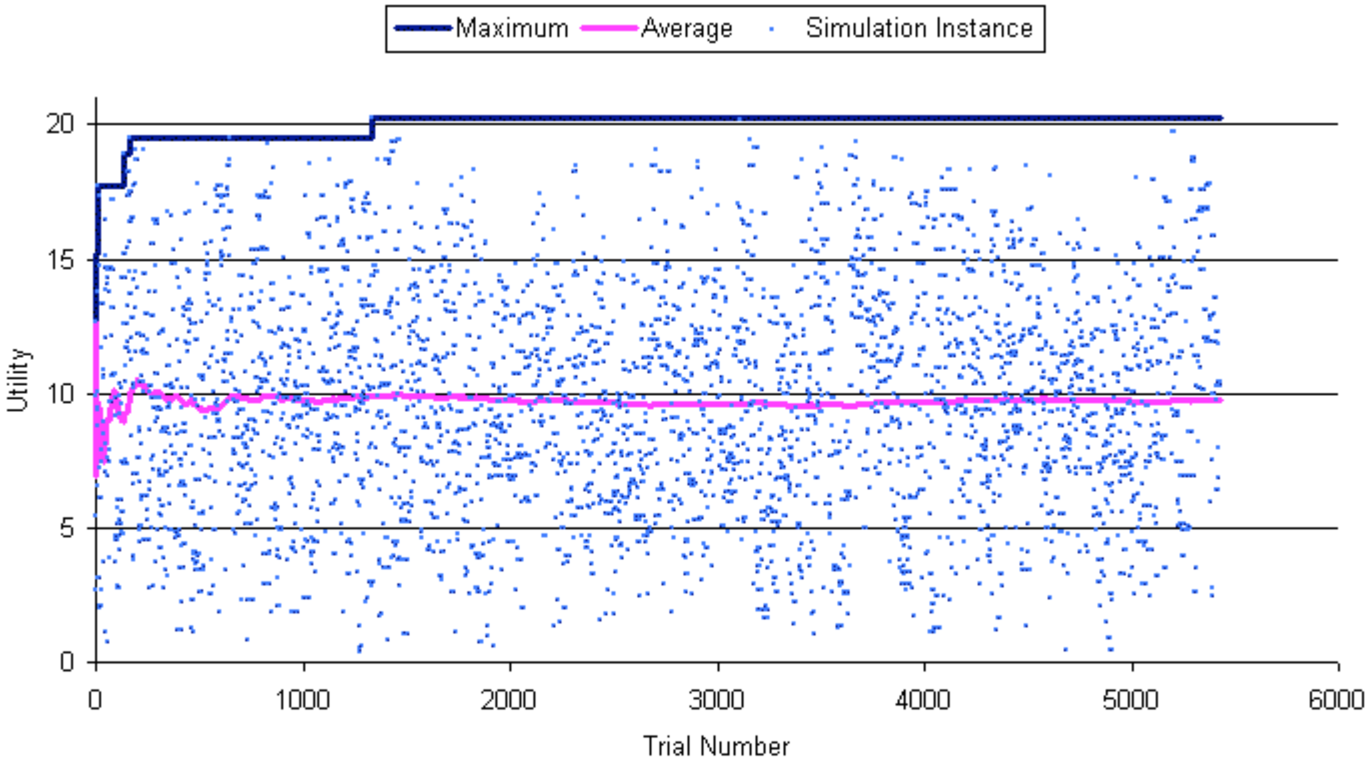

Figure 9. Genetic algorithm output for configuration one

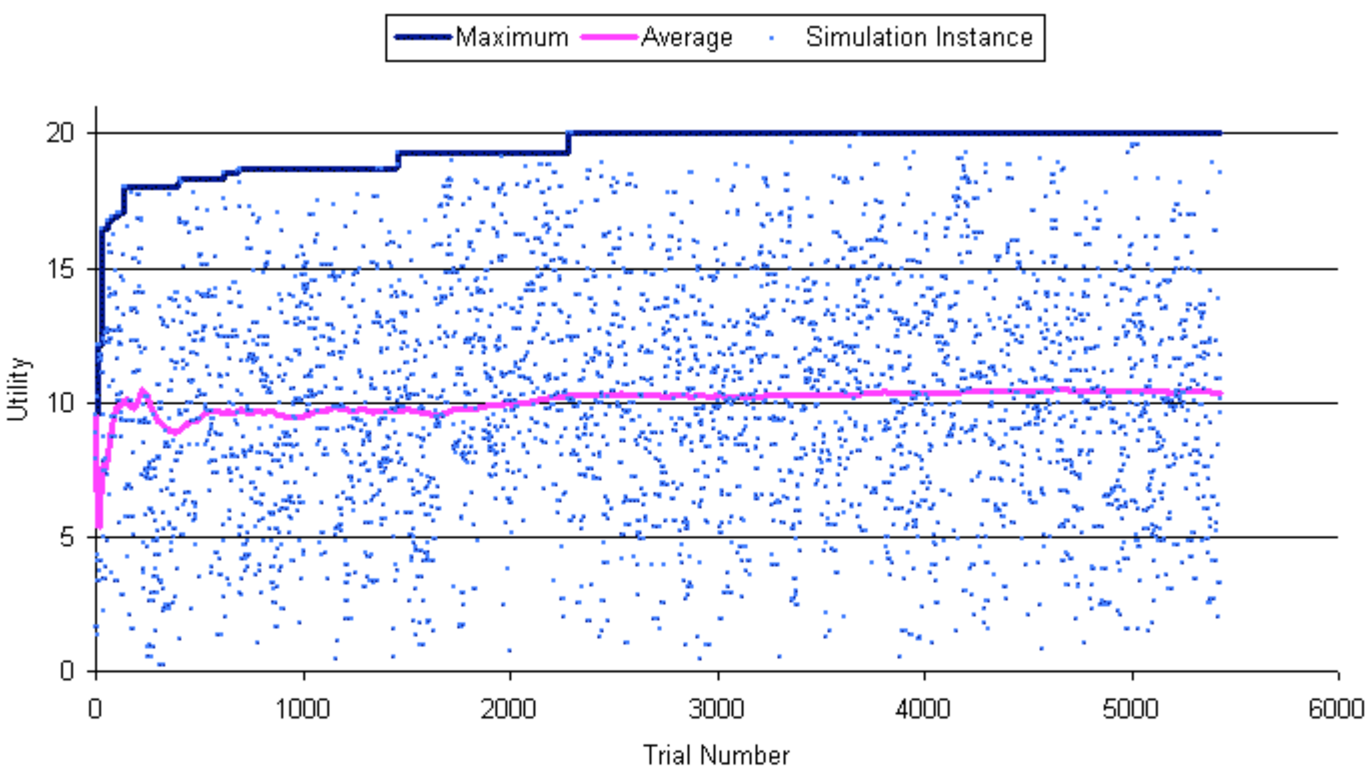

Figure 10. Genetic algorithm results for configuration two

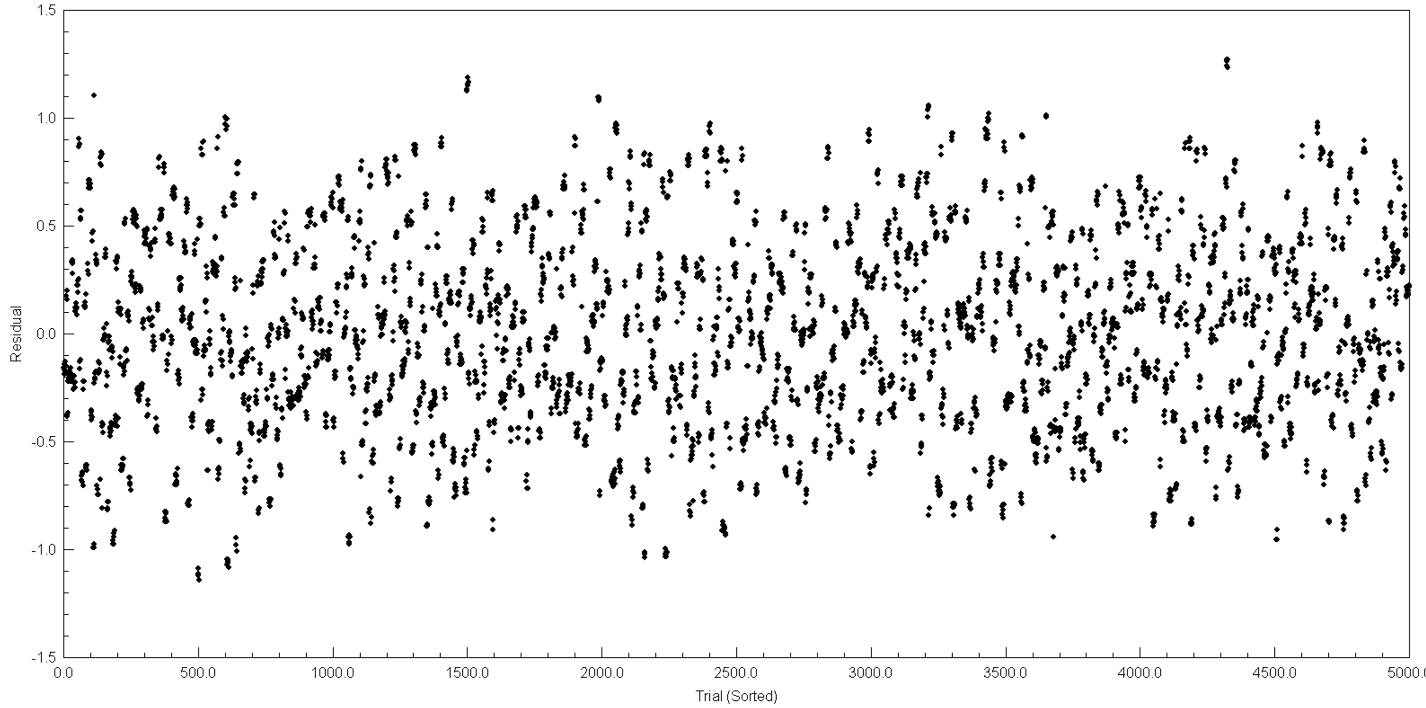

Figure 11. Residual plot of fit of analyitical model

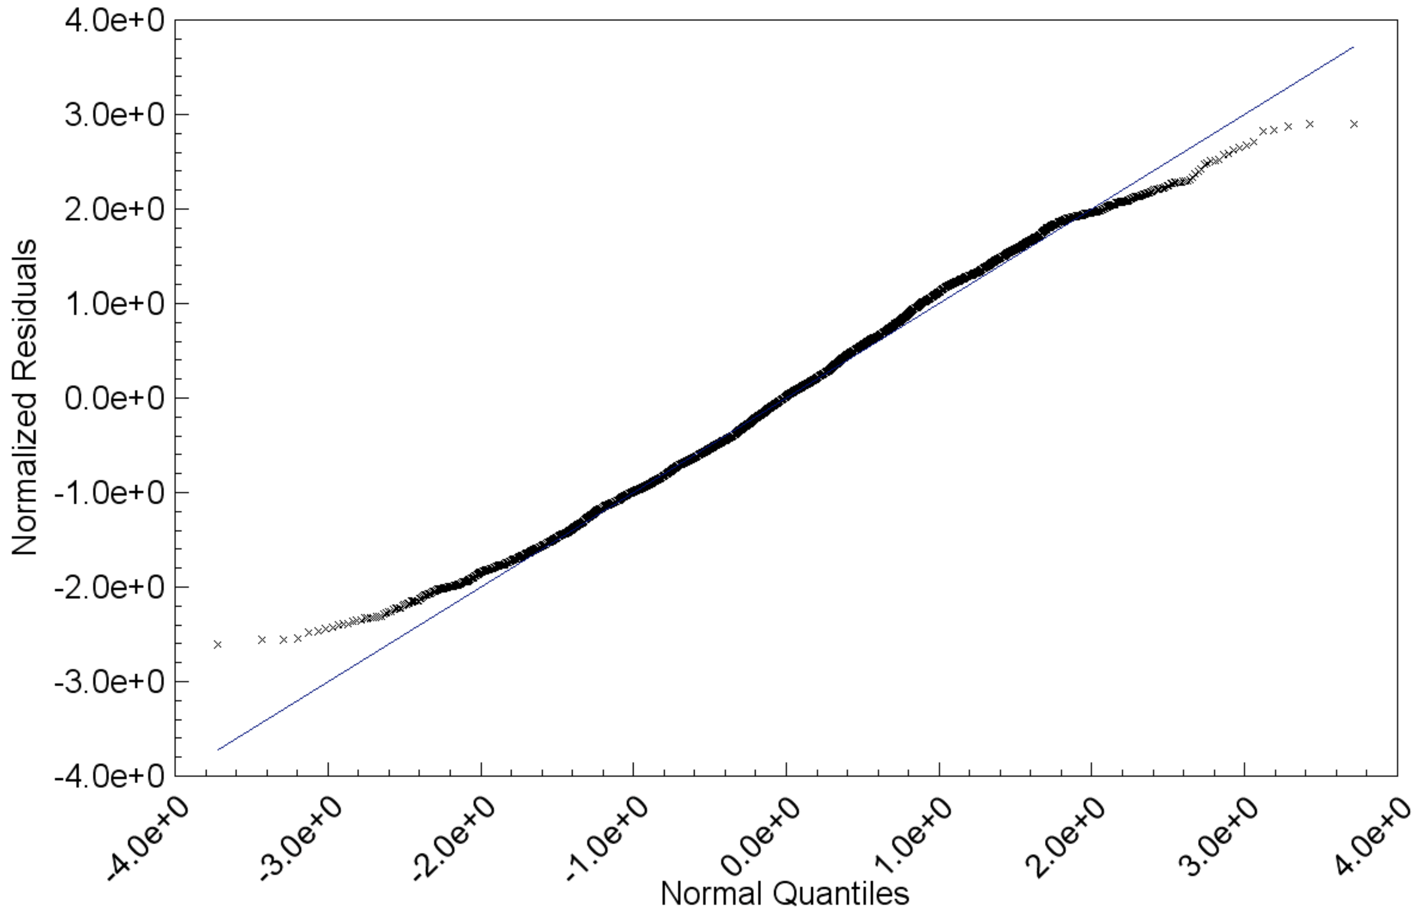

Figure 12. Residual probability plot

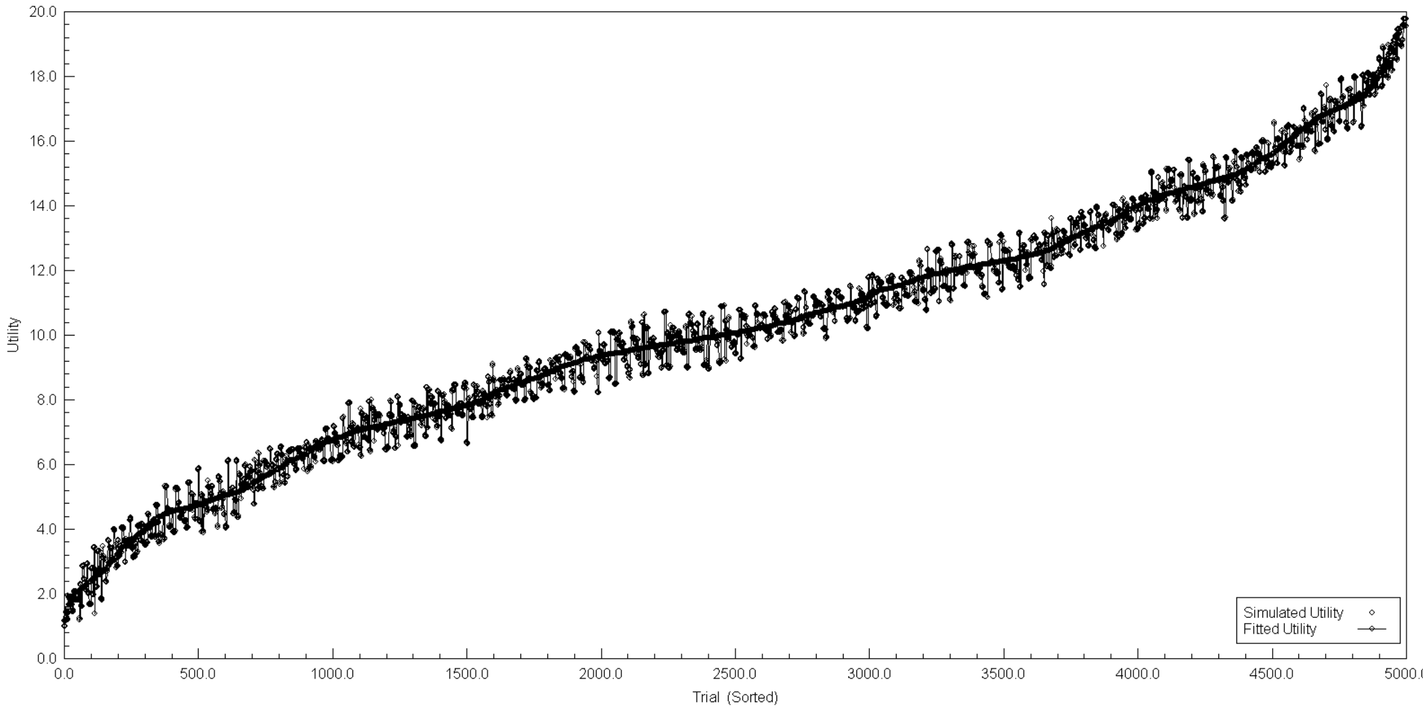

Figure 13. Analytical model output compared to actual data