05ICES-72

Using dynamic simulations and automated decision tools to design lunar habitats

Scott Bell
S&K Technologies Inc. at NASA Johnson Space Center ER2

Luis Rodriguez
USRA at NASA Johnson Space Center EC2

David Kortenkamp
Metrica Inc. at NASA Johnson Space Center ER2

ABSTRACT

This paper describes the role of transient simulations, heuristic techniques, and real-time integrated control in designing and sizing habitat life support systems. The integration of these three elements allows for more accurate requirements to be derived in advance of hardware choices. As a test case, we used a typical lunar surface habitat. Large numbers of habitat configurations were rapidly tested and evaluated using automated decision support tools. Through this process, preliminary sizing for habitat life support systems were derived. Our preliminary results show that by using transient simulations and real-time control, we substantially reduced the system mass required to meet mission goals. This has greater implications for general systems analyses and for life support systems. It is likely that transient models, real-time integrated control, and other analyses capable of capturing the uncertainties of systems can be useful for systems analyses much earlier in the system development life cycle than has previously been considered.

INTRODUCTION

Sizing and optimization of habitat life support systems is an ongoing challenge for the life support community. Technology choices, buffer sizes, power plant sizing, crop planting area and subsystem flow rates all need to be determined in order to design an appropriate and optimal system that meets mission requirements. Often these decisions are made using steady state simulations and labor intensive searches [1]. Our approach uses transient simulations and automated searches to test large numbers of habitat configurations, including those that may not be obvious. By using an integrated simulation the interaction between subsystems can also be evaluated which can provide very different design requirements than when subsystems are considered independently. Furthermore by using a transient model, control can be leveraged to provide a more precise (and often smaller) sizing approximation. The study of malfunctions, crew schedules, and reliability also become possible with transient modeling, all of which can be rapidly tested with heuristic search techniques. In previous work [2] we examined how an heuristic tool called a genetic algorithm [3] can be used with a dynamic simulation to find optimal life support configurations. In this paper, we extend our work to include a real-time PID system controller [4] and again used a genetic algorithm and a transient life support system model called BioSim [5] to find an optimal habitat configuration for a 90 day mission to the moon.

DECISION TOOLS

TRANSIENT MODELS

Transient models vary their output based on previous inputs and allow the course of the simulation to be altered while it is running. For integrated life support, this allows for buffer sizing, control, malfunctions, and other properties to be examined [5]. Furthermore, system dynamics and nonlinearities can cause damaging instabilities in ALS systems [7,8]. By introducing even simple transient modeling early in the design process, we can estimate more appropriate sizes of ALS systems.

Our dynamic habitat simulation is based on an existing simulation called BioSim [5]. BioSim has been developed at NASA JSC over the past three years and is still under active development. It is a generic habitat simulation, so our first task was to create a specific instance of BioSim for a lunar habitat. Because instances of BioSim are stored in an eXtensible Markup Language (XML) file, creating a new instance of BioSim does not require changing any computer code. Instead, a GUI (Graphical User Interface) tool is being developed to connect modules with the resource productions and consumptions. The configuration can then be saved and read directly into BioSim, which loads and instantiates the simulation.

HEURISTIC TOOLS

Optimization of any system requires searching through a large solution space to find optima. For relatively small solution spaces this can be done exhaustively, i.e., try every possible solution to guarantee an optimal one. For linear problems, techniques such as linear programming can identify optimal solutions, if they exist. As solution spaces get larger and models are no longer linear, this becomes increasingly challenging. The problems described in this paper have solution spaces that can be greater than 2^{24}. Further, many non-linearities are considered within BioSim, including crew schedules, crop production, malfunctions, and general stochastic processes. Thus, there is great benefit to utilizing heuristic tools capable of intelligently searching through the search space, identifying promising zones where optima may lie.

There are many tools for searching large solution spaces. For example, Monte Carlo techniques[8,9] probabilistically probe the solution space to hone in on a satisfactory solution. Reinforcement learning [10] uses feedback from the environment to search through the solution space. Hill climbing [12] is a simple technique to move towards a locally optimal solution in a large space. Finally, genetic algorithms[3] use biology-based insights to find solutions. It is important to remember that most of these approaches cannot guarantee an optimal solution since they do not exhaustively search the entire solution space. One exception is reinforcement learning which can produce provably optimal solutions in very constrained situations. Instead, these approaches exploit the underlying topology of the solution space to move in the most promising direction and find locally optimal solutions. In this paper we use genetic algorithms as our automated search tool-–however it should be noted that our methodology can be analogously pursued with any number of such tools.

REAL-TIME CONTROL

Real-time control means using computer programs to read sensor values and set actuators to establish or maintain an equilibrium. Real-time control is also referred to as closed-loop control. Effective real-time control can reduce buffer sizes by reducing oscillations in resource usage over time. This is similar to how a very well controlled car moves straight down the road while a poorly controlled car weaves from side to side. In the former case the road can be narrower (i.e., the buffer smaller) and thus less costly for the same performance. By examining real-time control in conjunction with buffer and component sizing in designing life support systems, we can more accurately predict the system sizes and performances. The real-time controller used in our experiments is quite simple. In the future we will evaluate more sophisticated controller such as those described in [13,14].

GENETIC ALGORITHM

Genetic algorithms are based on the paradigm of natural selection. Successful genes breed with other successful genes to create offspring. If the offspring themselves are successful they breed. After each generation the entire population of genes should be more and more successful. The keys to using a genetic algorithm approach are to encode the problem as a gene and to develop a fitness function that measures a gene's success.

We implemented a genetic algorithm where the gene is a description of the initial configuration of the transient models (e.g., crop size, storage capacities, process flow rates, etc.). The genetic algorithm program directs BioSim to utilize this initial configuration and create an instance of the transient model accordingly. The model is simulated until consumable resources are exhausted and the mission can no longer continue. A fitness function is utilized to compare the quality of each configuration tested by the genetic algorithm. The fitness function is how the genetic algorithm evaluates genes to determine which genes to save and breed and which genes to discard. The intention of this fitness function is to identify the optimal design for the 90 day lunar mission. An optimal design will have minimum mass while meeting mission objectives. The fitness function combines the length of the mission and the Equivalent System Mass (ESM) of the configuration to make these assessments as described in the Metrics Section, below. ESM is a measure of the predicted launch pad mass of the configuration, where a lower mass is preferred for exploration missions [15]. Thus, in our experiments “preferred” genes are those configurations that resulted in the longest running time with the least mass. We fixed a maximum mission length at 90 days as suggested by the Lunar reference mission [16]. Thus after 90 days configurations were only distinguished by their ESM.

GENES

Each gene in our genetic algorithm represents a single life support systems configuration. It consists of between 18 and 21 attributes, each an integer. At first these integer values are randomly assigned to each attribute. Over time, the genetic algorithm adjusts each attribute's value to search for an optimal solution. We looked at three different habitat scenarios: 1) no biomass; 2) biomass and crew in the same environment; and 3) biomass and crew in separate environments. (More detail on the simulated scenario is provided below in the section entitled Example ALS Mission.) The attributes and the range of values for each scenario are below.

Gene Attribute Name	Gene Attribute Range
O_2 Injector flowrate	0 - 30 mol/hr
OGS power consumption	0 - 2500 watts
VCCR power consumption	0 - 100000 watts
PowerPS power production	0 - 7000000 watts
CrewEnvironment volume	0 - 20000000 L
O_2 Store capacity	0 - 2000 mol
O_2 Store initial level	0 - capacity mol
PowerStore capacity	0 - 5000000 watts
PowerStore initial level	0 - capacity watts
FoodStore capacity/initial level	0 - 2000 kg
FoodStore initial level	0 - capacity kg
WaterRS power consumption	0 - 3000 watts
PotableWaterStore capacity	0 - 8000 L
PotableWaterStore initial level	0 - capacity L
GreyWaterStore capacity	0 - 4000 L
GreyWaterStore initial level	0 - capacity L
DirtyWaterStore capacity	0 - 5000 L
DirtyWaterStore initial level	0 - capacity L

Table 1: Gene attributes for all three scenarios

Gene Attribute Name	Gene Attribute Range
CO_2 Injector flowrate	0 - 30 mol/hr
Tomato Shelf Size	0 – 1.5 m^2
Lettuce Shelf Size	0 - 20 m^2

Table 2: Gene attributes for scenario 2

Gene Attribute Name	Gene Attribute Range
PlantEnvironment volume	0 - 20000000 L
CO_2 Injector flowrate	0 - 30 mol/hr
Tomato Shelf Size	0 – 1.5 m^2
Lettuce Shelf Size	0 - 20 m^2

Table 3: Gene attributes for scenario 4

Please note, the odd numbering (1, 2, 4) of the scenarios is an artifact of precursor studies performed by this research team with similar scenarios in the past [2].

METRICS

During technology development it is necessary to be able differentiate potential components, subsystems, and systems when making funding, scenario, architecture, and design decisions. Metrics may be defined to highlight favorable aspects of an individual technology or suite of technologies to assist these decisions. The challenge lies in defining useful metrics that may be determined objectively. Often objective metrics cannot be defined to consider all critical aspects of the system and subjective metrics are utilized instead. Generally, it is beneficial to have several metrics considering various aspects of the system providing a complete perspective of critical system issues.

Equivalent system mass (ESM) is currently the predominant metric utilized in system analyses within Advanced Life Support. Typical use of ESM involves trade studies of life support technology currently under development. ESM is a cost related metric acting as a proxy for launch costs by predicting launch pad mass. NASA endeavors to minimize launch pad mass as it is predicted that as much as $30,000 could be spent to impulse one kilogram of mass into Martian orbit. In general, we determine ESM by simulating components within an assumed scenario and solving a steady state mass balance for average throughput rates [15]. Subsequently, components are sized for providing the determined throughput rates. This provides the mass, power, volume, and cooling requirements for each component within the system. Power, volume, and cooling requirements are converted to mass using cost equivalecy factors (Eq. 1). The cost equivalency factors are derived to reflect the anticipated support hardware that will provide the necessary infrastructure of the system [17].

$$ESM = \Sigma_{\forall i} \begin{bmatrix} M_i + (V_i \cdot V_{eq}) + (P_i \cdot P_{eq}) \\ + (C_i \cdot C_{eq}) + (CT_i \cdot D \cdot CT_{eq}) \end{bmatrix}$$

Eq. 1: Formula to find ESM

The most prominent use of ESM is the determination of the ALS metric (Eq. 2). The ALS metric is utilized as a progress report and is reported to the US Congress. It compares the potential of utilizing Advanced Life Support Technologies with the current state of the art in the International Space Station.

$$ALS\ Metric = \frac{ESM_{ISS}}{ESM_{ALS}}$$

Eq. 2: ALS Metric

Reliability analysis enables the consideration of several metrics that are not currently considered by ESM, but is proposed to be added as a tool for systems analysis in Advanced Life Support. Reliability has several implications critical for technology selection including the determination of maintenance and redundancy requirements, contingency planning, and the availability of life support functions. Flight ready hardware is not currently available for the consideration of rigorous reliability analyses, thus integrated system modeling can be utilized as an alternative. However, steady state mass balances, as in ESM, are not adequate. Simulation of variable performance and malfunctions are both necessary to effectively consider the reliable performance of a system. The modeling and analysis described here is a precursor for analyses that is planned for studying reliability.

For the purpose of this study we developed a metric for differentiating between the configurations tested by the genetic algorithm (Eq. 3). In our experiment we are searching for configurations which can successfully execute a 90 day lunar mission. We quickly discovered that several configurations are capable of achieving this goal. Thus, a second term was added to the fitness function designed to reduce the consumption of resources.

$$f = w_t t + \frac{\Sigma_{\forall i} (w_i a_{i,max} - w_i a_i)}{\Sigma_{\forall i} (w_i a_{i,max})}$$

Eq. 3: Fitness Function with ESM proxy

In Eq. 3, f defines the fitness of a configuration, t is the length of the mission in hours, w_i is the weight of attribute i, a_i is the value of attribute i and $a_{i,max}$ is the maximum value of attribute i (the list of attributes can be found in Table 1). The function f is a unitless measure of fitness. Thus, the weight w_t has unit hr^{-1} and takes a value equal to one.

The first term, $w_t t$, measures the length of time that the simulation runs without failure in integer values, as BioSim is a discrete event simulation. The second term will never be greater than one, which allows the first term to dominate. That is, a missions lasting 2160 hours will always be rated as more fit than any mission lasting fewer than 2159 hours, no matter the amount of resources consumed. In this analysis, we are most concerned with 90 day Lunar missions (2160 hours), so we have instructed BioSim to terminate any simulations that complete 90 days. This renders the first term equal to 2160 for all 90-day missions, whether or not a configuration may have been capable of surviving for longer than that time. Since all 90-day missions will be more fit than any mission lasting less than 90 days, the second term in the fitness function then differentiates configurations, selecting those that minimize resources. The second term considers the sizing of the various subsystem attributes, their relative weights, and rewards configurations with components sized smaller than their maximum values. The attributes utilized and their weights are shown in Table 1.

Weights in Table 4 are derived from the theory behind the determination of ESM. Each attribute being sized by the genetic algorithm imparts some mass load upon the system. The weights are chosen to predict this load based upon the anticipated mass of the item, or the mass equivalency suggested for aspects such as power or volume. Thus, a portion of the numerator of the second term, $\Sigma_{\forall i} (w_i a_i)$, in the fitness fuction (Eq. 3) can be utilized as a proxy for ESM. It will not be ESM itself, but rather it will be a relative measure of the ESM components traded by the GA. Its units will be kg, as in ESM. Results of proxy analysis is included with the analysis of genetic algorithm results.

Attribute	***Attribute Unit***	***Weight***	***Weight Unit***
O_2 Injector	mol/hr	2.1606E-04	kg-hr/mol
CO_2 Injector	mol/hr	2.1606E-04	kg/W
VCRR power req	W	3.3279E-01	kg/W
OGS power req	W	1.3785E-01	kg/W
Power production	W	6.2000E-02	kg/W
Crew/crop volume	L	1.3310E-01	kg/L
O_2 Store Capacity	mol	1.0877E-03	kg/mol
O_2 Store initial level	mol	3.2000E-02	kg/mol
Power Store Capacity	W	6.8700E-01	kg/W
Power Store initial level	W	0.0000E+00	kg/W
Food Store Capacity	kg	2.3600E+00	kg/kg
Food Store initial level	kg	0.0000E+00	kg/kg
Tomato/Lettuce shelf size	m^2	3.7010E+01	kg/m^2
WaterRS power req	W	2.5625E-01	kg/W
Water (Potable, Grey, Dirty) Store Capacity	L	6.8373E-02	kg/L
Water (Potable, Grey, Dirty) Store initial level	L	1.0000E+00	kg/L

Table 4: Each configurable attribute and its contribution to the utility function.

EXAMPLE ALS MISSION

For the experiments in this paper we implemented a specific instance of the simulation to reflect a lunar habitat. The instance was designed with information from an internal JSC memo describing a lunar reference mission [16]. The reference mission assumes a four person crew with equal numbers of men and women. Mission length is 90 days with the habitat initiated and operating nominally upon crew arrival. The landing site is the lunar south pole with the sun above the horizon 80% of the time and surface temperatures between 210K and 230K during the day. The habitat atmosphere is composed of 29% oxygen at an overall pressure of 65.5 kPa and a leakage rate of 0.00224 kg/day. Food is shipped in most circumstances (although we looked at the addition of small salad crops) and is 0.257 kg/crewmember-day moist food and 0.665 kg/crewmember-day of dry food. Air, water, and waste

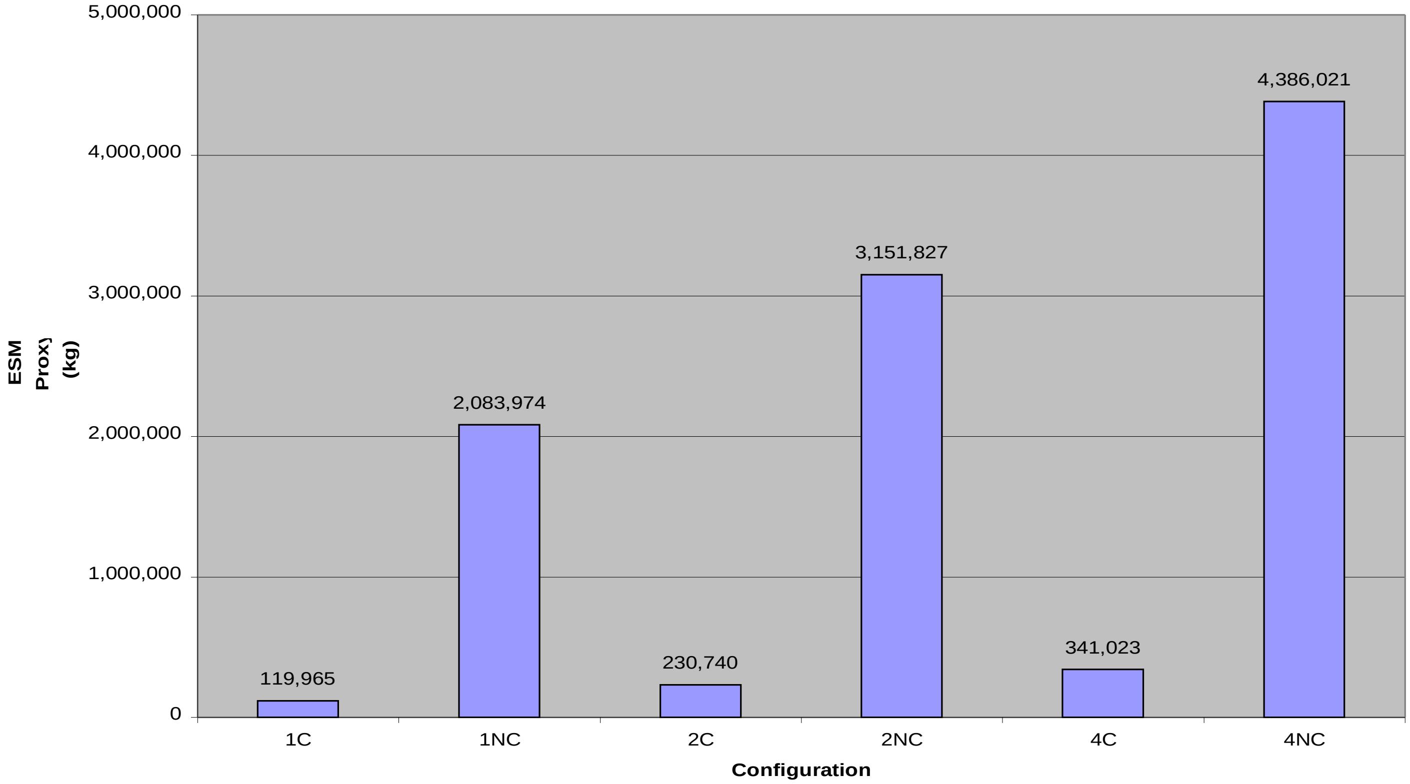

Figure 1. Chart showing ESM proxy among different ALS configurations.

recovery systems are part of the habitat. One four-hour EVA by one crew member was performed each day of the mission. The EVA takes place through an airlock that is 3.7 m^3 in size and 10% of the airlock atmosphere is lost each time the airlock is used.

Parameters such as the size of the habitat (specifically air volume), the size of the recovery systems, the amount of salad crops and the size of the power subsystem were not fixed and were determined through analysis of simulation results.

EXPERIMENTS

Our genetic algorithm ran through 6 different lunar mission scenarios described in Table 5.

Scenario Name	***Description***
1NC	No crops without control.
1C	No crops with control.
2NC	Crops with integrated with crew cabin without control.
2C	Crops with integrated with crew cabin with control.
4NC	Crops with separate cabin from crew without control.
4C	Crops with separate cabin from crew with control.

Table 5. Lunar mission scenarios.

Each of these scenarios were then run 4 times, which we termed A, B, C, and D. Each experiment was stopped when convergence was determined in the genetic algorithm. Convergence was determined when the genetic algorithm settled upon a configuration that maximized the fitness function.

The controller was very basic. Oxygen and carbon dioxide were regulated to sea level earth compositions. Oxygen was generated using the OGS when the O_2 store was determined low, namely when the levels dropped below 30%. Likewise when the potable water store dropped below 30%, the WRS was turned on to produce more potable water.

The ALS mission ended when resources for either the crew or the crops ran too low, or the gas composition of the cabin became dangerous, or if the 90 day limit had been reached.

RESULTS & DISCUSSION

Components were successfully sized in each of the size scenarios simulated by the genetic algorithm. Each run of the genetic algorithm produced a most optimal configuration. The output from each of the four simulation runs was averaged to find a predicted mean value sizing for each component. Further, 95% confidence intervals for the predicted size of each component was also determined. Based on the average sizing, the proxy of ESM was also determined.

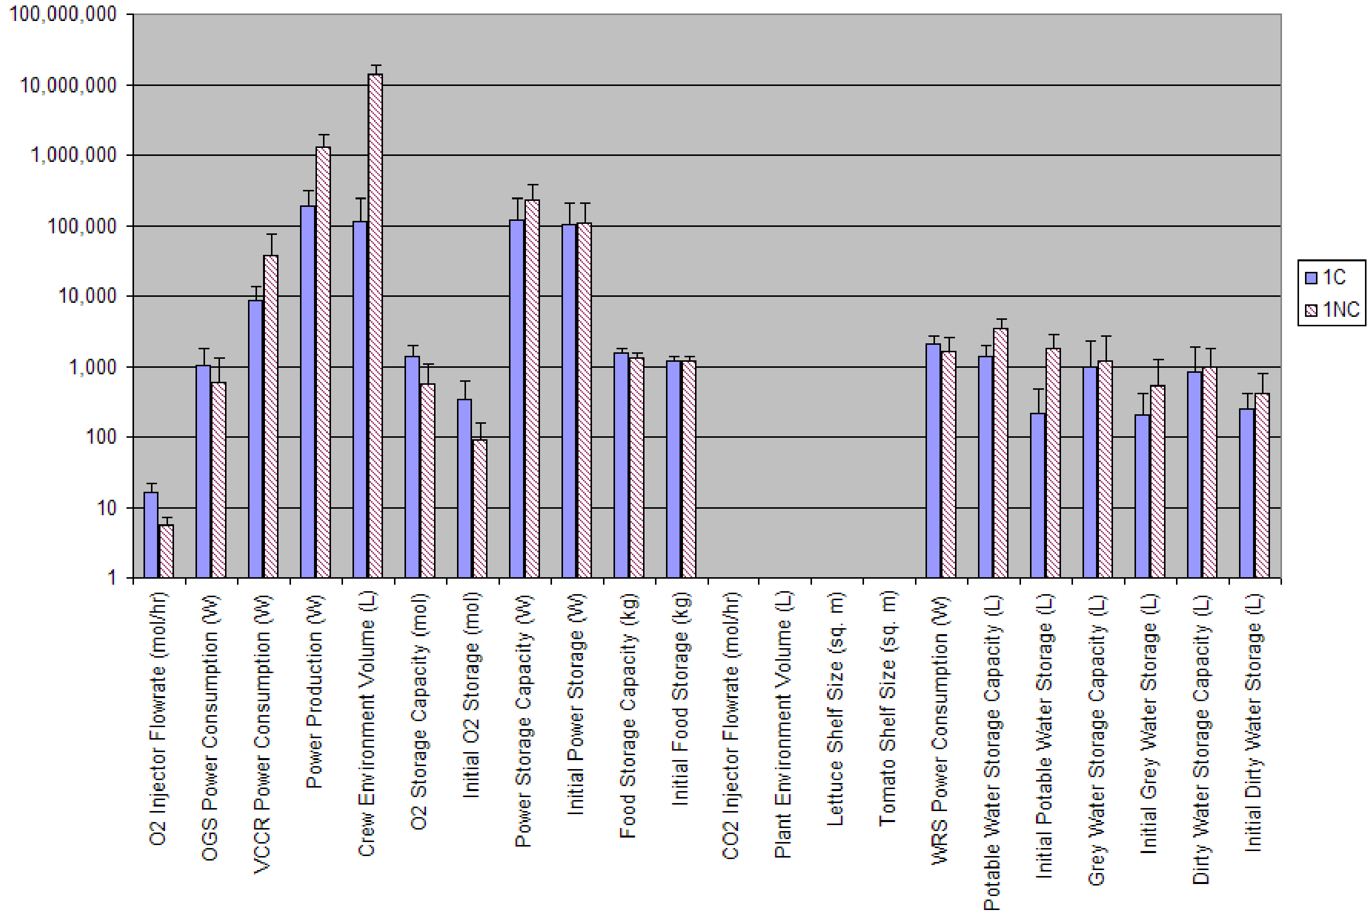

Figure 2. Logarithmic chart of ESM proxy for configuration 1.

Figure 1 depicts the relationship between ESM proxy and each of the configurations simulated. Configuration 1C has the lowest proxy value, although 2C and 4C are relatively similar in size. The uncontrolled cases, 1NC, 2NC, and 4NC each were effectively over an order of magnitude larger. The differences between the controlled simulations are not currently considered to be significant, as will be demonstrated in the discussion of the subsequent figures.

It is not unexpected that the controlled case should outperform the uncontrolled case with respect to the proxy. The uncontrolled case effectively simulates steady state operations of the system as components are operated continuously, although the they may lead the system outside the operable envelope of the system. Alternatively, the controller monitors several system state variables and determines which components within the system may be turned on or off during the simulation. With this basic level of decision making utilized to operate the components within the system, it becomes possible to alter the sizing of some components and reduce the overall proxy. The significance of these results lies in the simplicity of the controller utilized. Many more sophisticated integrated control regimes exist including: reinforcement learning, market based control, and neural networks. Each can be utilized similarly to optimize the system and improved results should be anticipated. We plan to demonstrate this in the future.

Figure 2 compares the output from the controlled and uncontrolled simulation of configuration one. Please note the use of a logarithmic scale on the vertical axis, which enables display of all the components on one chart. Due to this, only the positive error bars are shown, representing the 95% confidence interval, as the negative bars are distorted in this view. Seven of the 18 components are differently sized by an order of magnitude, explaining the disparity found in the ESM proxy (Figure 1), although not all can be shown to be significantly different. The most notable results which are significant include power production and the volume of the environment. Since the fitness function is based upon ESM, power and volume apparently are the most attractive areas for the genetic algorithm to consider for reduction of resources due to their relatively high cost equivalencies. This is the focus of the second term in the fitness function. It should also be noted that in some cases the genetic algorithm sized components such as oxygen storage and the power consumption of the OGS larger in the controlled case than in the uncontrolled case. This is the benefit of effective control. As it turns out, there is a proxy benefit to size certain components larger, and operate them in a batch format, rather than

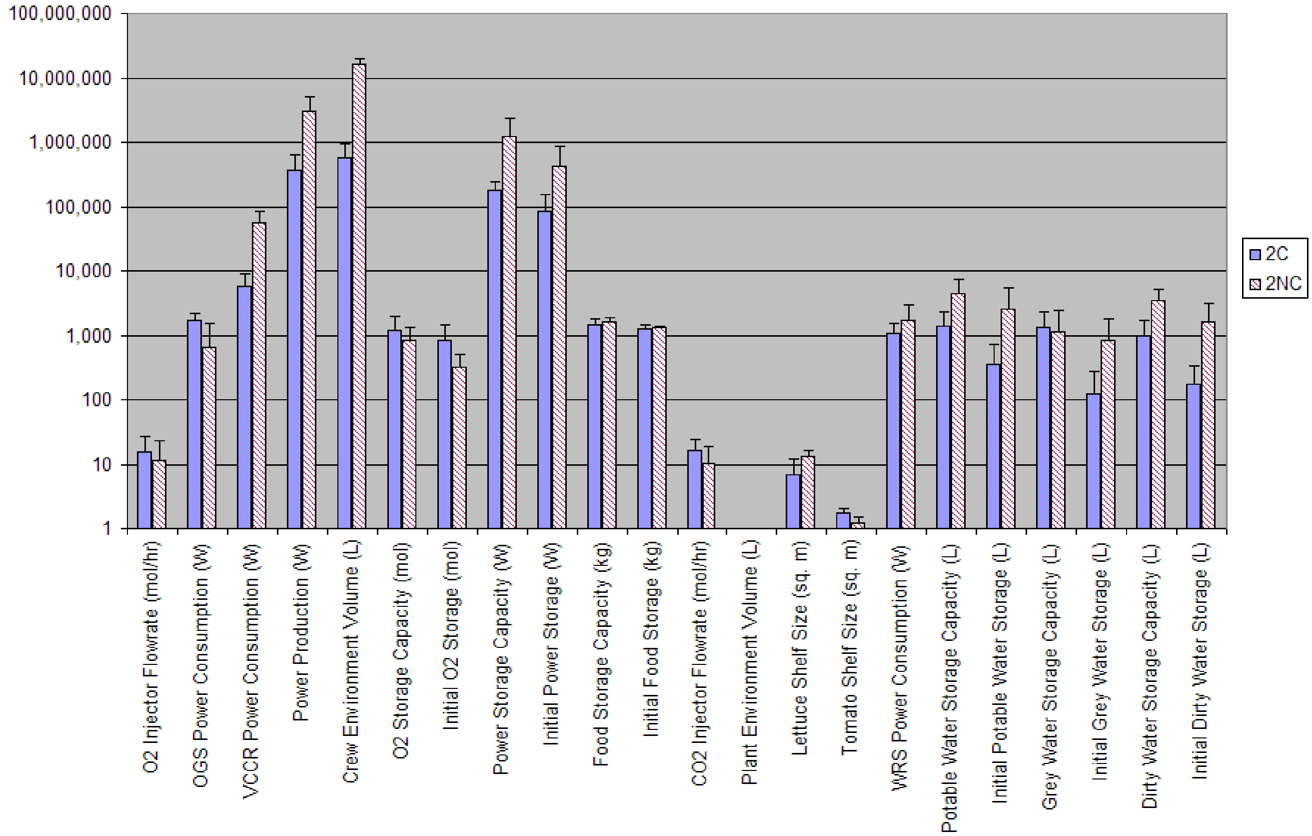

Figure 3. Logarithmic chart of ESM proxy for configuration 2.

continuously as in the uncontrolled case, as dictated by the controller.

Figures 3 and 4 show analogous results as in Figure 2 for configurations 2 and 4, respectively. In each case power and volume were distinctly smaller in the controlled case versus the uncontrolled case. In addition, a similar disparity was identified with respect to power storage. Configuration 4 sized a dedicated crop volume, which similarly was sized smaller in the controlled simulations.

Figure 5 compares the sizing chosen by the genetic algorithm for the 3 controlled scenarios. Across the board, it is apparent that only 3 components are sized so differently that an order of magnitude separates the output: VCCR, lettuce area, and initial potable water. However, even with these differences, it is likely that these components are not significantly different from a statistical perspective. Thus, it is concluded that the differences in the proxy of ESM depicted in Figure 1, are anticipated to not be significant among the controlled scenarios. The only differences that remain beyond the components which are sized in each scenario are those that are only sized in certain scenarios. Thus, the additional cost incurred by configurations 2 and 4 are due to the CO_2 injector, the crop environment volume, and the lettuce and tomato growth chambers. This is particularly evident in configuration 4, where the genetic algorithm sized a very large crop chamber, larger in fact than the crew volume. It is likely that other configurations may exist with smaller crop volumes, but the genetic algorithm failed to identify them. Future analyses are planned utilizing a more sophisticated genetic algorithm to increase the confidence in results.

IMPLICATIONS

A significant implication of this study is that steady state analyses may not be adequate for the complete study of life support systems. Certainly there is no disputing that certain research requiring transient models, such as integrated controls research or reliability, is impossible with steady state models. To support such research the necessary models will be developed. However, the question remains whether transient models should be utilized in areas where traditionally steady state models are utilized.

In the case of ESM analyses, typically steady state analyses are utilized. For components known to operate with time variant performance, average values are assumed to be representative. In a transient model, this can be simulated by assuming the components operate continuously at an average rate, as was done in the uncontrolled cases here. ESM proxy for the uncontrolled cases was determined based on these average values. Interestingly, in the controlled cases, the ESM proxy is determined based upon peak rates which define the operating state of components when they turn on. Thus, some integration of the actual usage of the components would be necessary to determine the average usage

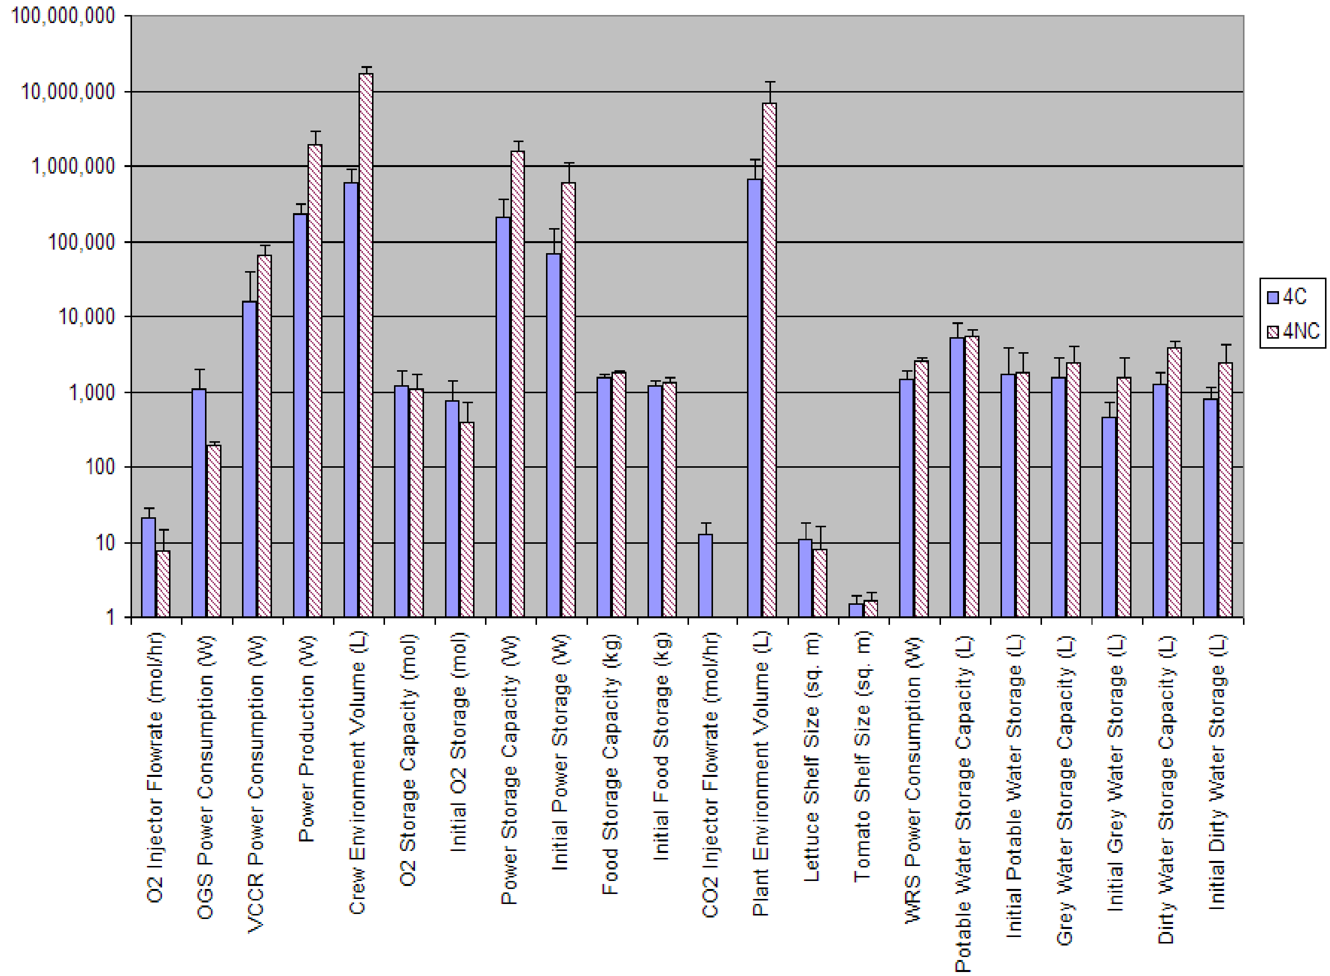

Figure 4. Logrithmic chart of ESM proxy for configuration 4.

rates that would be most comparable to the uncontrolled ESM proxy. Nonetheless, by determining ESM proxy based on the peak rates, results in a conservative prediction and an averaged prediction would produce a lower result.

The challenge remains in the verification and validation of such analyses. Future work proposes to consider the results produced by traditional steady state analyses and comparing it with that of the transient model utilized here to provide some preliminary verification. In addition, the development of the ability to increase the amount of analyses possible, including the study of reliability, is a major area for future investigation.

CONCLUSION

This paper shows that transient simulations coupled with an automated search tool and a real-time controller can configure and size lunar habitats. Dynamic simulations allow both control and insertion of malfunctions to be considered during the design process. By using an automated search tool, we were able to search through large search spaces of lunar configurations including more than 2^{24} possible combinations. Though our conclusions with respect to design requirements are preliminary, the transient simulation and automated search can aid human designers find configurations that might never have been tried otherwise. The large differences between the controlled and uncontrolled simulations reinforce this point. Furthermore, different malfunctions and control schemes can be tested rapidly for viability allowing greater flexibility in the design process. Ultimately, it may be worthwhile to consider the use of transient models not only in the study of control systems and reliability analyses, but also in areas traditionally dominated by steady state analyses, such as determining ESM.

ACKNOWLEDGMENTS

The authors thank Theresa Klein for implementing the original controller for this scenario. This work is supported by NASA Grant No. NNJ04HG90G.

REFERENCES

1. Yeh, H.Y., C. Brown, F. Jeng, C. Lin and M. Ewert, "ALSSAT Development Status and its Applications in Trade Studies,", *Proceedings of the 34th*

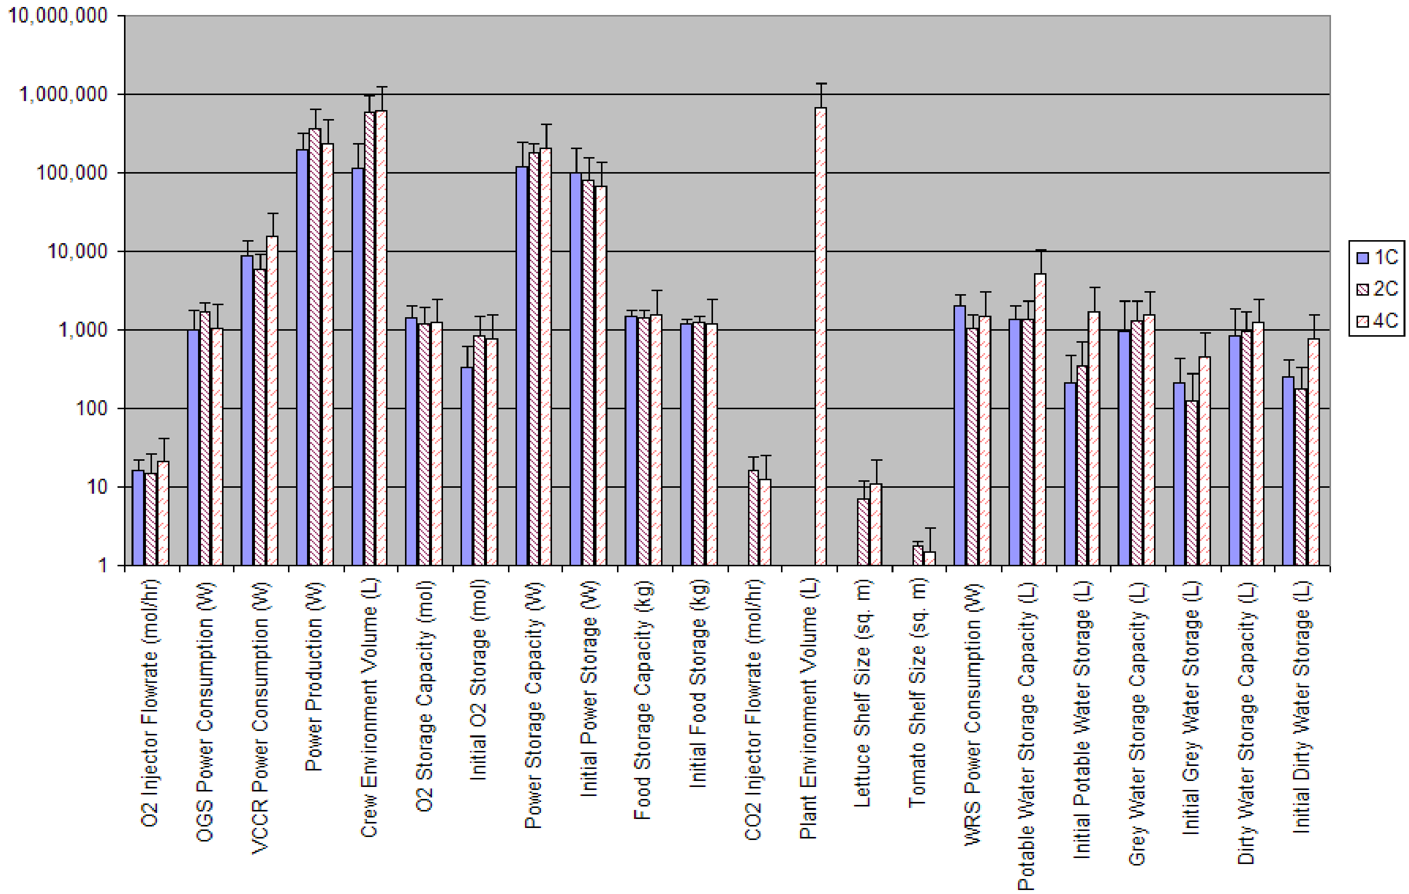

Figure 5. Logarithmic chart comparing ESM proxies of all configurations.

International Conference on Environmental Systems, 04ICES-60, 2004.

2. Kortenkamp, D., S. Bell and L. Rodriguez, "Simulating Lunar Habitats and Activities to Derive System Requirements," in Proceedings of the First AIAA Space Exploration Conference, 2005.
3. Holland, J., *Adaptation in Natural and Artificial Systems*, University of Michigan Press, 1975, Ann Arbor MI.
4. Klien, T., D. Subramanian, D. Kortenkamp, and S. Bell, "Using Reinforcement Learning to Control Life Support Systems," *Proceedings of the 34th International Conference on Environmental Systems*, 2004.
5. Kortenkamp, D and S. Bell, "Simulating advanced life support systems for integrated controls research," *Proceedings of the 33rd International Conference on Environmental Systems*, 2003.
6. Jones H. W., "Transient Modeling of ALS Systems," *Proceedings of the 33rd International Conference on Environmental Systems*, 2003.
7. Jones H. W., "Survey of Modeling Failures, Transients, and Nonlinearities in Advanced Life Support Systems," *Proceedings of the 31st International Conference on Environmental Systems*, 2001.
8. Jones H. W., "Nonlinear transient Models in Life Support," *Proceedings of the 32nd International Conference on Environmental Systems*, 2002.
9. Haldar A. and S. Mahadevan, *Probability and Statistical Methods in Engineering Design*, John Wiley & Sons, Inc., New York, 2000.
10. Ang A. H-S. And W. H. Tang, *Probability Concepts in Engineering Planning and Design Volume II Decision, Risk, and Reliability*, John Wiley & Sons, Inc., New York, 1984.
11. Sutton, R. and A. Barto, *Reinforcement Learning: An Introduction*, MIT Press, Cambridge, MA, 1988.
12. Nilsson, N., *Principles of Artificial Intelligence*, Tioga Publishing Company, Palo Alto, CA, 1980.
13. Abdelwahed, S., J. Wu, G. Biswas, J. Ramirez and E. Manders, "Online Fault Adaptive Control for Efficient Resource Mangement in Advanced Life Support Systems," *Habitation: International Journal for Human Support Research*, Vol. 10, No. 2, pp. 105-116, 2005.
14. Bonasso, R.P., D. Kortenkamp and C. Thronesbery, "Intelligent Control of a Water-Recovery System," AI Magazine, Vol. 24, No. 1, pp. 19-44, 2003.
15. Levri et al., "Advanced Life Support Equivalent System Mass Guidelines Document" NASA ARC NASA/TM—2003-212278, September 2003.

16. Hanford, T., "Transient Modeling Challenge: A Lunar Reference Mission for a 90-Day Habitat," *NASA JSC Draft Document*, 2004.
17. Hanford A. J., “Advanced Life Support Baseline Values and Assumptions Document,” NASA JSC CTSD-ADV-484 A, JSC 47804, NASA CR-2004-208941, August, 2004.

CONTACT

BioSim is available at http://www.traclabs.com/biosim. Questions and comments can be directed to Scott Bell at scott.e.bell1@jsc.nasa.gov, David Kortenkamp at kortenkamp@jsc.nasa.gov, or Luis Rodriguez at luis.f.rodriguez1@jsc.nasa.gov.

ACRONYMS

XML – eXtensible Markup Language

GUI – Graphical User Interface

ALS – Advanced Life Support

ESM – Equivalent System Mass

WRS – Water Recovery System

VCCR – Variable Configuration CO_2 Removal System

OGS – Oxygen Generation System

PowerPS – Power Production System

PID – Proportional Integral Derivative