Architecture Modeling of In-Situ Oxygen Production and its Impacts on Lunar Campaigns

Ariane B. Chepko[1] and Olivier de Weck[2]
Massachusetts Institute of Technology, Cambridge, MA, 02139

Diane Linne[3]
NASA Glenn Research Center, Cleveland, OH, 44135

Edgardo Santiago-Maldonado[4]
NASA Kennedy Space Center, FL, 32899

and

William A. Crossley[5]
Purdue University, West Lafayette, IN, 47906

In-situ lunar oxygen production has the potential to reduce the cargo mass launched from Earth necessary to sustain a lunar base. As research and development in lunar oxygen production continue, modeling tools are being used to help characterize the many possible system architectures and guide decisions for future plant designs. Using the previously built NASA In-Situ Resource Utilization (ISRU) System Model, an optimization tool was developed to facilitate exploration of the design space of the different system architectures represented in the model. For each architecture, an optimization of the continuous design space is performed using a gradient-based search. In instances when the gradient-based search cannot converge, the tool changes to simulated annealing, a heuristic method. Nine primary lunar oxygen production system architectures were optimized to minimize system mass for oxygen production levels from 500 kg/yr to 6000 kg/yr. Good designs minimized mass and maximized produced oxygen with system masses in the range of 100 kg to 700 kg. Preliminary results show that two particular architectures populate the Pareto-optimal front of best designs for most production levels, making them attractive for further investigation. An economy of scale of .837 was found using a power law regression, indicating that some economy of scale exists (values less than one have economy of scale) and that launching fewer, higher-capacity plants will be less massive overall than many small-capacity plants to achieve the same total production level. A simplified comparison of lunar-produced oxygen for crew breathing supply and ECLSS (environmental control and life support systems) technologies was performed with a space logistics planning tool, SpaceNet. For all but the most advanced ECLSS technologies, use of in-situ oxygen over a 10-year campaign resulted in more than 12,000 kg of consumables cargo launch mass savings.

Nomenclature

α = economy of scale
C = investment cost

---

[1] Graduate Research Assistant, Aeronautics and Astronautics, MIT/37-344a, AIAA Student Member.

[2] Associate Professor, Aeronautics and Astronautics, MIT/33-410, AIAA Associate Fellow.

[3] Senior Research Engineer, Propulsion and Propellants Branch, NASA/GRC MS 301-3, AIAA Senior Member

[4] Research Engineer, Applied Science Division, NASA/KSC KT-D-3

[5] Associate Professor, School of Aeronautics and Astronautics, 701 West Stadium Ave., West Lafayette, IN 47907-2045, AIAA Associate Fellow

f	=	objective function
g_i	=	inequality constraint
k	=	scaling coefficient
N	=	production quantity per unit time
x_i	=	design variable
x_{LB}	=	design variable lower bounds
x_{UB}	=	design variable upper bounds
ECLSS	=	Environmental Control and Life Support Systems
ISRU	=	In-Situ Resource Utilization
PEM	=	Proton Exchange Membrane (electrolysis system)
SA	=	Simulated Annealing
SQP	=	Sequential Quadratic Programming

I. Introduction

As NASA prepares for a return to the Moon, an opportunity is unfolding to extend the frontier of space exploration beyond the quick forays of the past and create a new paradigm in our approach to space mission planning. Almost every aspect of space mission sustenance relies entirely upon materials and supplies carried from Earth including propellants, pressurants, crew consumables, and tools. Shifting this Earth-based dependence to consider utilizing in-situ resources may enable more sustainable exploration, higher scientific returns, and lower mission costs[1]. The concept of In-Situ Resource Utilization (ISRU) is simply using resources that can be found at a present location. Because of NASA's lunar plans, lunar-derived resources have the most relevant interest in the short to medium term (2010-2030).

The work presented in this paper addresses oxygen extracted from the lunar regolith. Oxygen is a valuable resource that can be used as crew air, in fuel cells, or oxidizer for propellant. Oxygen extraction technologies are under development in a variety of organizations, but full-scale demonstrations and characterization of system scaling effects of the recent technologies have not occurred or are still at an immature stage. Because lunar oxygen extraction is a new technology with little historical background and experience to draw from in choosing a system design, the motivation arises to develop a modeling tool to analyze the numerous technologies and inform future decisions for lunar oxygen production system architectures.

This research, in collaboration with the NASA ISRU team, is developing a system modeling approach and design tool that captures the varying ISRU technology alternatives, enabling direct comparison of oxygen production system performance, overall mass, power, and sizing[2,3]. The primary goals for the tool are:

1. to enable easy reconfiguration of the system model between technology options to represent different system architectures,

2. to easily allow for future expansion of the model and addition of new technology options, and

3. to provide design space exploration and optimization methods that examine both the parameters specific to single system or subsystem designs and the effects of different combinations of technologies.

Generally, system models involve tying together models of individual subsystems or disciplines to capture system-level trade-offs and interactions. A system model for an aircraft may involve linking the analysis of the structures to an aerodynamics model and a propulsion model. The resulting tool allows a designer to optimize the system as a whole and understand the effects of one subsystem on another. This approach is known as Multidisciplinary Design Optimization (MDO), and has been employed in engineering fields for the last twenty years, with its usage increasing as computational power developed. Numerous examples of system modeling and optimization exist in the literature, including design of a blended-wing-body aircraft[4], communications satellites[5], and diesel engine exhaust treatment systems[6]. Some system models focus only on sizing parameters, such as in Ref.4, where a particular system architecture is chosen, and the size and shape of the wing, angle of attack, etc are the design variables. In Ref. 5, technology alternatives are considered as part of the optimization problem formulation. That study developed a satellite sizing tool that included choices for types of communications hardware, solar cells, batteries, and launch vehicles, however; the most of the analysis models were kept to first-order approximations and empirical fits to historical data.

Studies of ISRU and its mission benefits have also been performed for decades, but until recently, most work has been confined to paper studies and a few bench-test experiments of individual ISRU components, such as chemical reactors, electrolyzers, or cryocoolers[7,8]. The Eagle Engineering report of 1988[9] is one of the most comprehensive studies performed, reviewing thirteen chemical processes for lunar oxygen production and developing conceptual designs of pilot oxygen production plants. They focused on very large production rates (144-1500 Mt O_2/yr) that

would support oxygen use as spacecraft propellant oxidizer. This study developed detailed engineering models of an oxygen extraction process called hydrogen reduction and performed trade studies between a few subsystem technologies. Other ISRU modeling to date has incorporated previously reported data and analyses of ISRU chemical processes with economic models to provide a higher-level ISRU scenario analysis tool.[10]

NASA's ISRU system model discussed in this paper integrates detailed engineering analyses of the technology components that comprise a variety of oxygen production systems[2,3]. It captures the system-level interactions of specific subsystems and also allows for testing different types of subsystems (i.e. in the airplane model, evaluating a jet-powered propulsion system in place of a turbo-prop, etc…). Thus, the ISRU system model can be used to explore the design space of a specific type of oxygen extraction plant and to compare types of oxygen plants to each other.

This paper will provide a brief overview of the primary oxygen production processes included in the system model in Section A. The description of the ISRU system model framework and approaches to optimization of oxygen plant size, performance, and architectures will be presented in Section II. An analysis of space logistics for supplying crew breathing oxygen that compares varying capabilities of life support systems with ISRU technology is shown in Section III.

A. Lunar Oxygen Extraction

Lunar oxygen extraction can be accomplished via several different chemical processes. Three of these are receiving primary focus for development: hydrogen reduction, carbothermal reduction, and electrowinning. Of these three, only hydrogen and carbothermal reduction have currently been modeled for inclusion in the ISRU system tool. An oxygen plant system design consists of a reactor that uses one of these chemical processes and the supporting hardware necessary to operate the reactor and condition the reaction products such that the result is stored liquid oxygen of a specified purity.

Lunar regolith is comprised of about 42% oxygen by mass that exists in the form of glasses (mostly SiO_2) and metal oxides. The hydrogen reduction processes targets the iron oxide (FeO) in the glasses and metal oxides such as ilmenite, $FeOTiO_2$, Eq. (1). Hydrogen reduction is a reaction capable of extracting between 1% and 5% by mass of oxygen-to-regolith[9,11]. Regolith is heated to between 700°C and 1000°C and exposed to hydrogen that reacts with weakly bonded oxygen to form water. Water vapor is removed from the reactor and electrolyzed to form oxygen that is sent to storage and hydrogen that is recycled to the reactor. Hydrogen reduction systems require high-temperature reaction chambers, a hydrogen gas flow system, water separation and electrolysis systems, along with mechanisms to handle regolith into and out of the system and the separate gas streams. A schematic of hydrogen reduction is shown in Fig. 1.
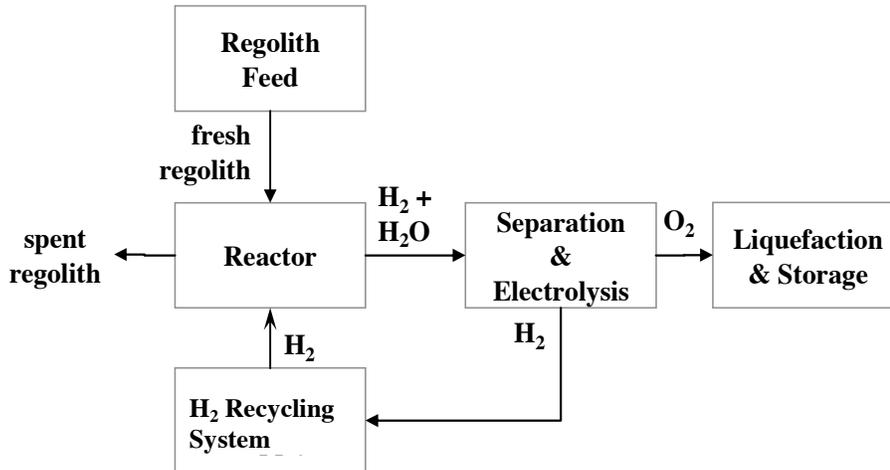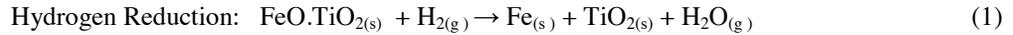
$$\text{Hydrogen Reduction:} \quad FeO.TiO_{2(s)} + H_{2(g)} \rightarrow Fe_{(s)} + TiO_{2(s)} + H_2O_{(g)} \tag{1}$$

$$\text{Electrolysis:} \quad H_2O_{(g)} + \text{electricity} \rightarrow H_{2(g)} + 1/2\ O_{2(g)} \tag{2}$$

Figure 1. Hydrogen reduction schematic.

In the carbothermal process, the regolith is heated to a higher temperature of about 1625°C and methane gas is flowed through the chamber[12]. At this temperature, methane cracks into carbon and hydrogen, and carbon is deposited into the melt. The deposited carbon then reduces the metal oxide forming carbon monoxide. Hydrogen and carbon monoxide are passed into a separate reactor that converts the gas streams to methane and water. The water is electrolyzed and methane and hydrogen are re-circulated. This process can extract approximately 15-20% by mass of oxygen-to-regolith[9]. Carbothermal systems require two high-temperature reactors, a water separation and electrolysis system, hydrogen and methane gas flow systems, and regolith handling. Example reactions are shown below (Eq. (3-5)) with a schematic (Fig. 2). MOx represents a generic metal oxide in the notation.
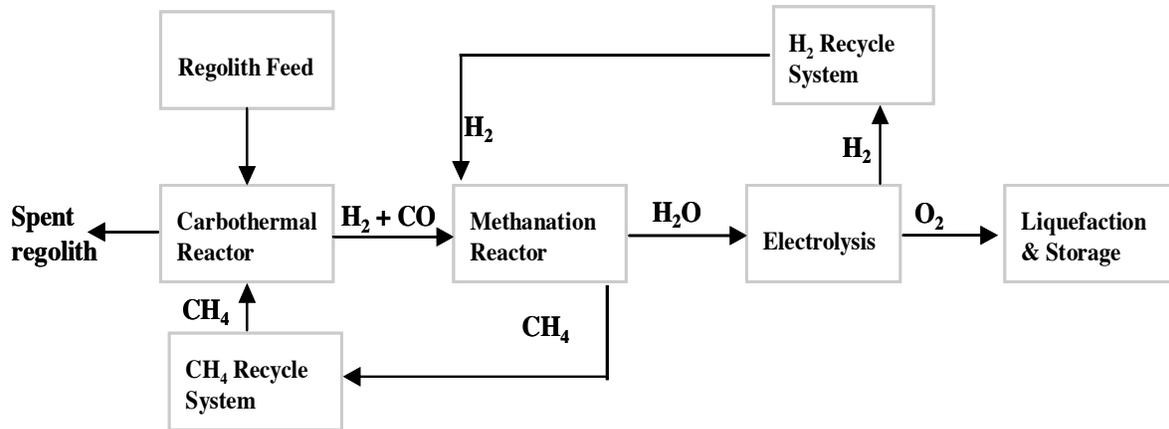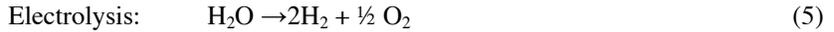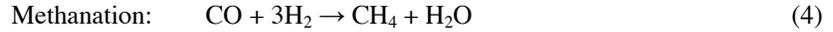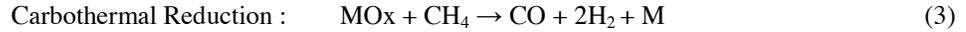
$$\text{Carbothermal Reduction :} \qquad MOx + CH_4 \rightarrow CO + 2H_2 + M \qquad (3)$$

$$\text{Methanation:} \qquad CO + 3H_2 \rightarrow CH_4 + H_2O \qquad (4)$$

$$\text{Electrolysis:} \qquad H_2O \rightarrow 2H_2 + \tfrac{1}{2} O_2 \qquad (5)$$

Figure 2. Carbothermal reduction schematic.

The baseline approach is for the above processes to occur in batch fashion, i.e. discrete quantities of regolith are deposited to and removed from the reactor, rather than allowing for continuous flow of regolith and reactants. Aside from overall reaction efficiency, the overall production rate of liquid O_2 therefore depends mainly on the number of reactors, the size of each reactor and the speed of reaction. The role of the system model is to quantify these relationships, while satisfying all physical and chemical constraints.

II. ISRU System Model

A. Model Framework

The ISRU System Model being developed by NASA is an evolvable tool that can expand as more ISRU technologies are developed and system designs tested. It is intended to work in conjunction with experimental test beds, using test data to develop more accurate models, and using system model outputs to guide future designs and experiments. To be as general as possible, and thus able to incorporate as much of the potential design space as possible, the ISRU System Model is structured according to the *functional decomposition* of the system. There are three primary functions: oxygen production, power supply, and liquefaction with storage. Each of these functions can be accomplished by one or more technology options or processes, which in turn can be broken down into subsystems and basic components (Fig. 3). By analyzing each function in the system, a set of function definitions and stable interfaces can be defined such that any model can be added to the system to accomplish a specific function as long as it has the required interfaces. For example, any model of a hydrogen reduction reactor must have interfaces defining gas feed into the system, gas flow out of the system, and regolith flow in and out of the system.

The internal details of the reactor design may vary widely, but to accomplish the function of a hydrogen reduction reactor, it must specify those fundamental interfaces.
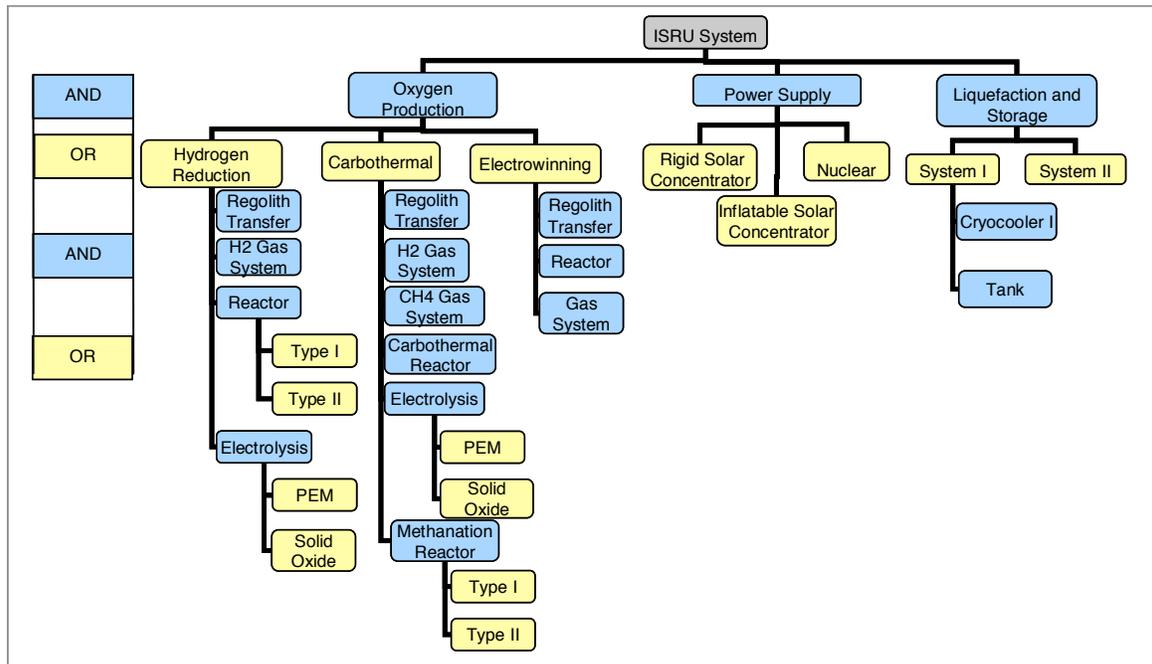

Figure 3. Functional breakdown of ISRU system partially populated with technology options.

Decomposing the system into its constituent elements reveals a pattern in the levels of the hierarchy. Each alternating level consists of either subsystems whose functions must be *combined together* to create their parent function ("AND" levels), or a set of choices from which only one is needed to achieve its parent function ("OR" levels). At each "OR" level, there is a repository for the different model alternatives (since the OR is an exclusive or: XOR) that perform the function described by each block, i.e. all models for a PEM electrolyzer are contained in the repository for that block. By defining the static interfaces throughout the entire system, a "black box" approach to the analysis models can be adopted, creating a "plug and play" model architecture at multiple levels. A new hydrogen reduction reactor model can be plugged in and tested in the system framework, or, on a higher level, a complete hydrogen reduction system model can be added to the framework. With these repositories, the system model contains all potential system configurations. The inputs and outputs of the models are linked together, choosing one option at each "OR" level, to form an end-to-end system model configuration.

The ISRU System Model is integrated using Phoenix Integration's ModelCenter, a tool that enables data flow between separate models. ModelCenter has a graphical interface that shows variable linkages and allows control and execution of all models tied into the environment. The ISRU System Model is populated with subsystem and component models developed by the NASA ISRU team[11,12]. A linking tool was created to enable simple, programmatic re-linking between "OR" level options based on user selections. This allows a search or optimization tool to automatically reconfigure the system model to different system architectures.

B. System Model Optimization
1. System Sizing Optimization

In considering optimization of the ISRU System Model, there are two levels at which optimization is of interest: a high-level comparison between system architectures, and, to allow fair comparison of architectures, a detailed optimization of system sizing parameters and performance. On the detailed level, for an individual architecture (comprised of a set of given technology choices), there are several key design variables that dictate the overall mass and oxygen production quantity. These include the number of batches of regolith processed per day, the size of the

reactor, the number of reactors, mass of regolith per batch, reactant gas flow rates, and pre-heat time to raise the regolith to operating temperature. Each "inner-level" optimization problem's objective function is to minimize system mass subject to constraints on annual oxygen production and reaction batch times. Evaluation of the objective requires executing the string of component models for a particular architecture. This results in a highly non-linear design space.

Two approaches to accomplish the inner-level optimization have been tested: a sequential quadratic programming (SQP) algorithm based in ModelCenter that is a gradient-based search using finite difference approximations for gradient calculations, and simulated annealing. The problem formulation is stated as follows:

$$\text{Minimize:} \qquad f(x_1, x_2, x_3, x_4, x_5, x_6) \qquad (6)$$

Where: f = system mass scaled by a constant
x_1 = number of batches per day
x_2 = mass of regolith per batch [kg]
x_3 = warm-up time [hr]
x_4 = reactor diameter [m]
x_5 = H_2 flow rate [mol/s]
x_6 = CH_4 flow rate [mol/s]

Subject to:

$$g_1 = 1 - \frac{O_2 \, produced}{O_2 \, desired} \leq 0 \qquad (7)$$

$$g_2 = \frac{\left(Time\_reaction + Time\_warmup + Time\_regolith\_transfer\right)}{24/NumberBatchPerDay} - 1 \leq 0 \qquad (8)$$

$$x_{LB\,i} \leq x_i \leq x_{ub_i} \qquad (9)$$

The first constraint, g_1, enforces a minimum desired O2 production level, otherwise the system would minimize reactor mass and lead to a lightweight but useless system. The number of time constraints, g_2, present in the optimization problem depends on the number of reactors- with three reactors, three equations for g_2 are formulated, allowing for simultaneous operation of reactors: one reactor dumps spent regolith and takes un-reacted regolith in while the second is preheating and the third is reacting.

The SQP solver is the fastest, but it can settle into local minima and sometimes cannot converge on a feasible solution. Tests with the SQP solver with the initial problem formulation above proved to be erratic and very dependant on the initial starting point (Fig. 4). Table 1 demonstrates the variation in final x vector, x^*, and $f(x^*)$ for different initial values. Tests 1 and 2 have very similar initial conditions, but result in significantly different final values. Test 3 had quite different initial conditions, but converged to a result close to Test 1. In all of these cases, the optimizer terminated with a suboptimal, feasible point, unable to reach an optimum condition. This is unacceptable when considering that in order to perform a system architecture optimization, the inner optimizer will be required to run consecutively many times. It must be robust enough to allow a wide-spread design space search.
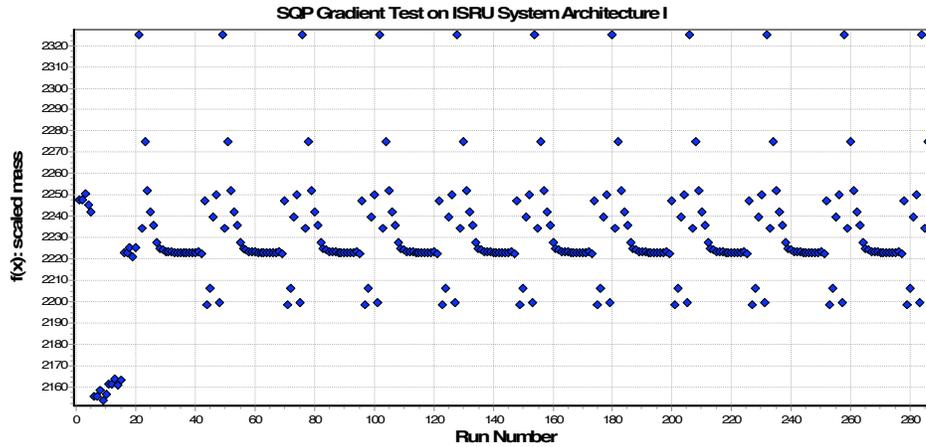
Figure 4. Unsatisfactory Convergence history of SQP test run.

Table 1. Comparison of initial and final points in gradient optimizer runs

	Test 1		Test 2		Test 3	
	Initial Guess (x^o)	Final Point (x^*)	Initial Guess (x^o)	Final Point (x^*)	Initial Guess (x^o)	Final Point (x^*)
Batches per Day	5	22	5.00	7	20	22
Mass Regolith per Batch (kg)	10	32.5	10.00	109.5	20.00	32.0
Warm Up Time (hr)	1	1.1	2.00	2.91	2.00	1.08
Reactor Diameter (m)	0.1	.24	0.10	.41	0.30	0.24
O_2 Produced (kg/yr)	143.52	1780.0	141.26	1930.1	1004.0	1800.0
$f(x^*)$	354.1	1129.49	236.1	2001.0	827.6	1129.43

To address the erratic behavior of the gradient solver, the problem was re-formulated by scaling the objective function, constraints, and design variables. Scaling can be critical to achieving stable and efficient algorithms, especially when using real design models[13,14]. Poor scaling occurs when there are large differences between the order of magnitude of the design variables and the function, causing the Jacobian and Hessian matrices used by the SQP algorithm to be ill-conditioned. Often scaling involves reformulating variables and objective function to be of a magnitude of unity, but this can be somewhat problem dependant, and experimentation is necessary to achieve consistent results. As noted in Ref. 13, "scaling is the single most important, but simplest, reason that can make the difference between success and failure of a design optimization algorithm".

The scaled utility metric chosen for this problem was to maximize the ratio of oxygen produced to system mass (overall mass efficiency). The design variables were scaled by their upper bounds to put all changing parameters on the scale between zero and one.

$$\text{Minimize:} \quad \text{Mass}/O_{2\,produced} : \quad f(x_1, x_2, x_3, x_4, x_5, x_6) \qquad (10)$$

7

Where: x_1 = batches per day ratio
x_2 = mass of regolith per batch ratio
x_3 = warm-up time ratio
x_4 = reactor diameter ratio
x_5 = H_2 flow rate ratio
x_6 = CH_4 flow rate ratio

Subject to:

$$g_1 = 1 - \frac{O_2\,produced}{O_2\,desired} \leq 0 \tag{11}$$

$$g_2 = 1 - \frac{O_2\,produced}{O_2\,desired} \geq -.1 \tag{12}$$

$$g_3 = \frac{\left(Time\,\mathrm{Re}\,action + TimeWarmup + Time\,\mathrm{Re}\,golithTransfer\right)}{24/NumberBatchPerDay} - 1 \leq 0 \tag{13}$$

$$\frac{x_{LB\,i}}{x_{UB_i}} \leq \frac{x_i}{x_{UB_i}} \leq 1 \tag{14}$$

The new constraint, g2, sets an upper limit on the size of the system. The scaled problem significantly improved the SQP performance, enabling it to converge for most cases on a minimum point (Fig. 5). Convergence may still be in local minima, but the problem is smoother and better behaved. Starting at a few different initial points can help avoid accepting local minima.
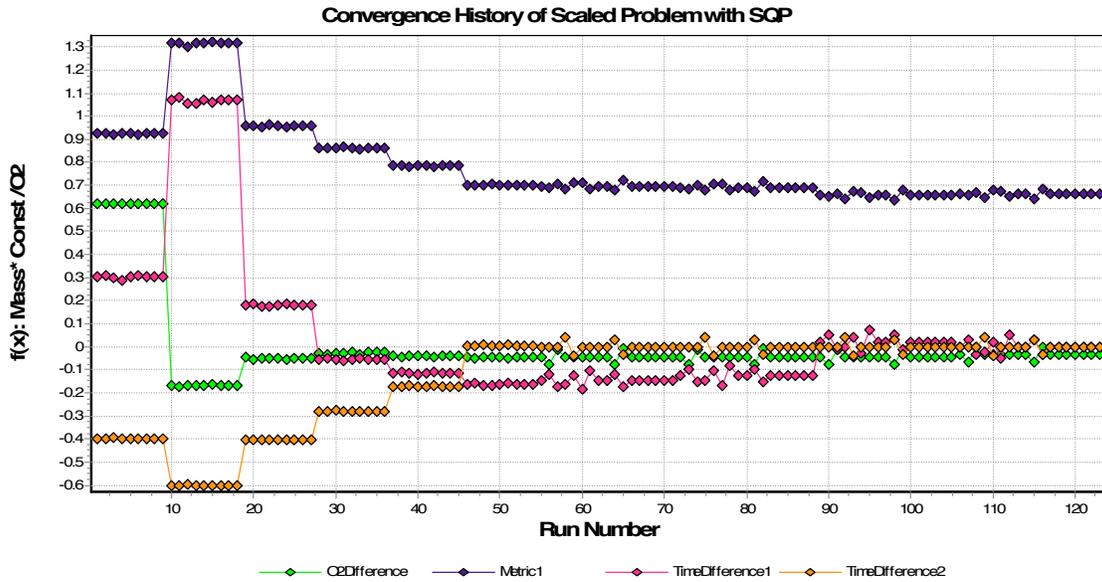

Figure 5.   Convergence history of SQP test run with scaled problem formulation.

8

To handle the cases when the SQP algorithm cannot converge on a minimum, the optimization tool switches to a heuristic search method, simulated annealing. Heuristic search algorithms such as simulated annealing do not use gradient information to guide the search, but instead use probabilistic operators that direct a random search through the design space. Simulated annealing is a process modeled after the behavior of molecular energy states as hot metal cools and freezes to a minimum-energy crystalline state[15]. A random starting point is chosen, a perturbing function changes the design variables, and the subsequent function evaluation is accepted if it is better than the last. If the objective value is not better than the previous objective, it may still be accepted with some probability, based on a "temperature" value. This probability decreases as the algorithm progresses, mimicking the temperature cooling in the metal annealing process. Simulated annealing is a powerful tool to apply to highly non-linear systems because it can overcome being trapped in local minima.

A MATLAB-based simulated annealing algorithm was modified for the ISRU problem. Constraints were handled using a quadratic exterior penalty method, thus the objective function will have a penalty added to it when constraints are violated. The magnitude of the penalty changes with the square of the constraint violation. Simulated annealing is computationally expensive and can take hundreds to thousands of function evaluations to converge on a global optimum. For the ISRU System Model, to balance between time and fidelity of the results, an additional termination criteria was set to end the search after 100 evaluations without finding a new, better point. This will usually result in termination without an optimum design, but it still provides a good answer for system comparisons and can be used as a starting point for finer-tuned further optimization if needed.

From trial runs with varying architecture options, simulated annealing took between 200 and 700 iterations to terminate, compared to SQP which was generally between 100-300 evaluations. Each function evaluation took from 1 to 5 minutes, depending on the system architecture configuration. Because of the computational time required for simulated annealing, it was only used in the ISRU System Model optimization tool after trying two or more initial starting points with the SQP gradient method. If a feasible point could not be found, the optimization tool switched to simulated annealing. A drawback to simulated annealing aside from computation time is that constraint handling via the penalty method can still result in infeasible designs. Simulated annealing results need to be checked for feasibility after termination.

Figures 6 and 7 show the convergence history of the same architecture case shown in the SQP figures (Fig.4-5). In Fig. 6, the progression of the simulated annealing run can be seen as more "inferior" points are accepted earlier in the history when the temperature parameter is higher, and as the system cools down, fewer jumps with smaller perturbation levels are accepted. The impact of the function penalty is demonstrated in Fig. 7, causing a sharp increase in the objective for combined constraint violations.
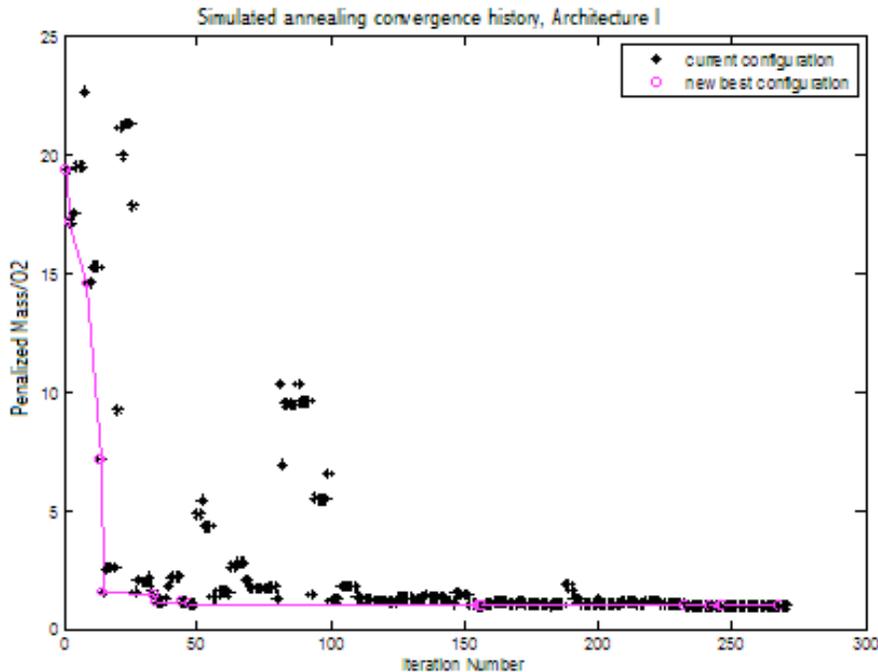

Figure 6. Convergence history for simulated annealing, desired O_2 of 2 Mt. Note: for this run, the upper limit in Eq. (12) on O_2 Production was not enforced.
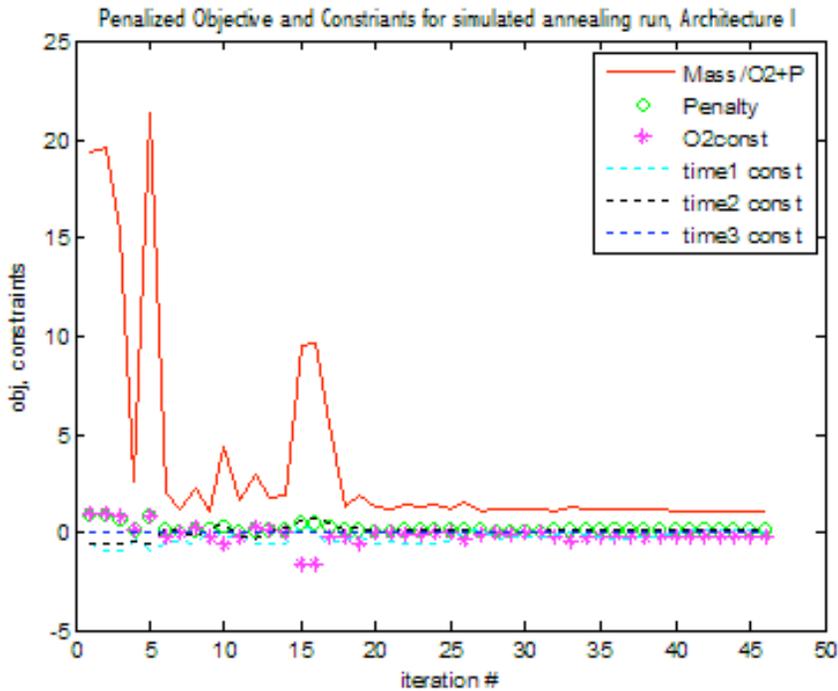
Figure 7. Objective and Constraint history for simulated annealing run.

2. System Architecture Exploration and Economy of Scale

With the optimization tool that can size individual architectures, the outer-level comparison between different system architectures can be performed. This exploration requires considering the technology options at the "OR" levels throughout the model as the design variables, and the goal is to search the trade space of combinations of technologies to understand how they perform for different levels of oxygen production (see tree in Fig. 3). The approach to explore the discrete design space of the ISRU System Model depends on the number of combinations of decisions that are available. The current population of technology options in the system model that have a significant impact on system performance result in only nine primary architectures for consideration (including secondary, lower impact choices in the model results in 120-800 combinations). These discrete variables include reactor types, electrolysis types, oxygen production methods, condenser types, and number of reactors. Because nine combinations is small enough to fully enumerate, the outer-level architecture comparison is executed by running an inner-level optimization on each architecture and directly comparing the results. As the ISRU System Model develops, the possible technology combinations will grow to be in the hundreds of thousands. At that point, full enumeration will no longer be feasible, and another search method will be required. This is the subject of ongoing work.

A MATLAB-based tool was written to control the full enumeration search and inner-level optimizations. Each architecture was executed for varying oxygen production levels to explore the overall relationship of system mass to oxygen production level. This relationship is of ultimate interest to mission planners who have to decide between deploying an ISRU plant on the Moon versus shipping an equivalent amount of consumables from Earth. Mass of the electrical power system is not included in this analysis. Annual oxygen production was varied from 500 kg/year to 6000 kg/year. The higher end of the range caused convergence issues in the models and resulted in mostly infeasible designs. Further work refining the models and constraints between component models is necessary to achieve a more robust system model.

Figure 8 shows the preliminary results of the full system optimization. A thorough validation and verification process is still required to assess the adequacy of the ISRU System Model and optimization tool. Nine architectures were tested at increasing production levels to find the Pareto front of non-dominated designs and to approximate an optimal production-cost function. While there is a mix of architectures on the Pareto front at the lower production levels, System II and System IV consistently populate the set of optimal designs for all production levels. This
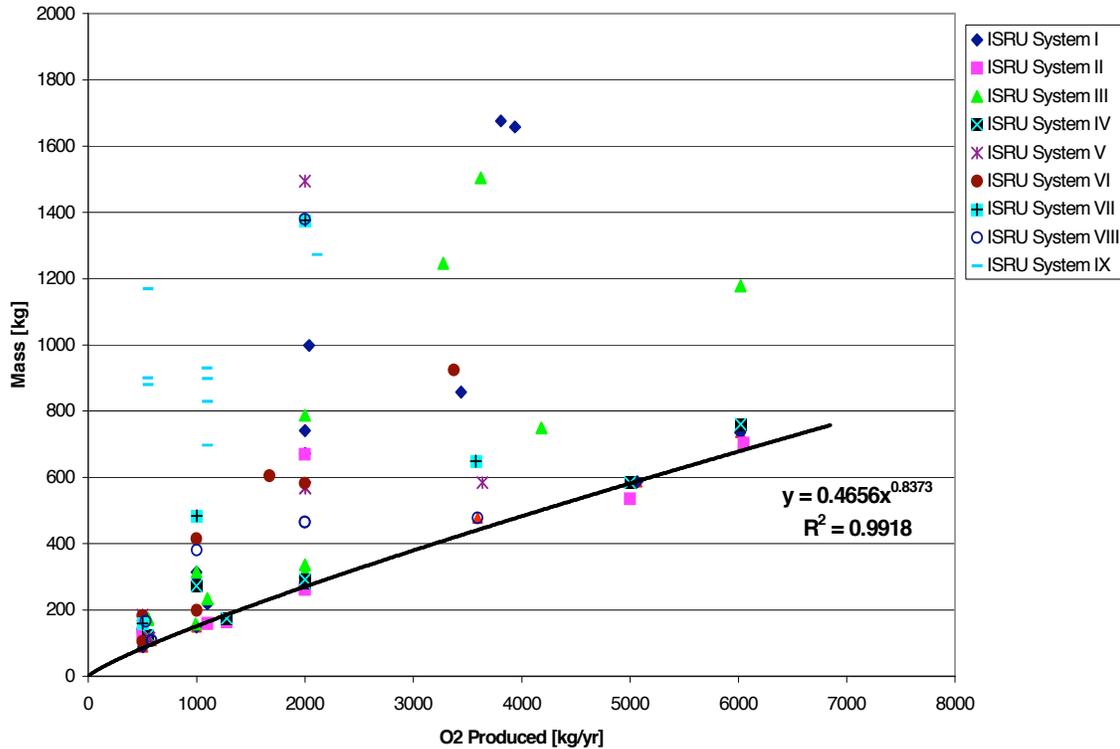
Figure 8. Preliminary results of system mass of ISRU oxygen plants for varying production quantities.

indicates that these particular system configurations should be investigated in further and may serve as a starting point for more detailed design optimization.

The set of best designs were fit with a power-law equation to determine if an economy of scale is present. Economy of scale is a concept generally considered when planning the rate and quantity of capacity addition to existing systems such as power plants and chemical production plants[16]. It exists when the per-unit-cost of a plant decreases as output or capacity increases. A commonly used model to quantify this is shown in Eq.(15)[16], with the empirically fit coefficients from Fig. 8 shown in Eq. (16).

$$C = k * N^{\alpha} \qquad (15)$$

Where:

C = investment cost
k = scaling constant
N = production quantity per unit time
α = scale coefficient

Pareto front production-cost function: $\qquad C = .466 * N^{.837} \qquad (16)$

Values of α with $0 \leq \alpha < 1$ exhibit economy of scale. It should be noted that this model allows investment cost to approach zero as capacity approaches zero. In practice, there is an operational range for applying economy of scale, as there is always some fixed-cost of producing a plant. It is the value of α that is important, as it indicates how the cost-increase will change as capacity is increased. The empirical fit of ISRU System Model optimization in Fig. 8 resulted in an α value of .837, which does exhibit a modest economy of scale. This indicates that the cost-per-plant (system mass) of oxygen production plants decreases as production capacity increases. Based on these results,

to achieve a 10 Mt/yr oxygen production rate, two 5 Mt/yr plants would cost approximately 1.1 Mt of launch mass, while five 2 Mt/yr plants would cost 1.3 Mt of launch mass. While these numbers are based on preliminary results of the ISRU System Model and require further testing, this may be an important consideration if large-scale lunar oxygen production becomes a future need.

III. ISRU Impacts on Space Logistics

One of the primary applications for in-situ produced oxygen on the Moon is to provide crew breathing supplies. A space logistics planning and simulation tool, SpaceNet[17,18], was used to assess the trade-off between Earth-based and lunar-derived oxygen. The baseline mission scenario for this study had a mission duration of 10 years with four crew members living on the lunar surface. A 4 person crew consumes approximately 1500 kg of oxygen per year in this SpaceNet scenario. In-situ oxygen supply was compared to the re-supply from Earth needed for four options of ECLSS (environmental control and life support systems) technology, ranging from minimal recycling capabilities to advanced life support technologies (see Table 2). All but the most basic option of ECLSS carry water from Earth to electrolyze for oxygen. The basic options involve lower levels of technology risk and lower levels of system closure. The more advanced options involve high closure for water recovery and additional systems for re-capturing and recycling oxygen. The primary difference between the most advanced options is the additional systems to aid in recovery of carbon dioxide and oxygen recycling. Cases with ISRU incorporated assume that an ISRU plant handles all of the oxygen demand, thus some of these additional oxygen recycling systems are not used, and the ISRU plant provides the full 1550 kg/yr demand (which serves as an upper-limit estimate in this study).

Table 3 presents the demand differences between the four ECLSS scenarios and the cases with ISRU included. The values shown include cargo tank mass. The ISRU plant mass for a 1550 kg/yr demand was determined using Eq. (16) from the preliminary ISRU System Model data and was found to be approximately 220 kg. The last column of Table 3 shows the mass savings of launched cargo over the 10-year mission. This analysis was limited to

Table 2: ECLSS Technology Option Capabilities

ECLSS Technology	Capability
Minimum	Some water recycling, small gas leak rate, no CO2 capture, no water electrolysis
Basic	Some water recycling, CO2 capture and water electrolysis for oxygen recycling
Intermeadiate	Higher level of water recycling, CO2 capture and water electrolysis for oxygen recycling
Advanced	Highest level of water recycling, CO2 capture, Methane reformer and electrolysis for oxygen recycling

Table 3.   Cargo Mass of Water and Oxygen Supplies for ECLSS vs ISRU scenarios

	Cargo Demand w/out ISRU Oxygen Production		Cargo Demand w/ ISRU Oxygen Production		
ECLSS Technology	Water Launched from Earth [kg/yr]	Oxygen Launched from Earth [kg/yr]	Water Launched from Earth [kg/yr]	Oxygen Launched from Earth [kg/yr]	10 year Cargo Mass Savings with ISRU [kg]
Minimum	849	3174	849	1554	31,740
Basic	2077	0.0	849	0.0	12,280
Intermediate	1384	0.0	156	0.0	12,284
Advanced	96	0.0	81	0.0	150
*With ISRU, Basic - Advanced options no longer using CO2 capture and water electrolysis for oxygen recycling					

comparing the quantities of consumables launched from Earth. It does not include the comparison in mass between ISRU plant mass and ECLSS technology system mass. However, for all but the most advanced ECLSS scenarios, subtracting ISRU plant mass from the cargo savings still results in a positive mass savings, without considering the effect of ECLSS systems that would no longer be needed when using in-situ oxygen.

Therefore we can make two main observations with respect to the use of lunar ISRU for breathable oxygen: a) there is a tradeoff between investing in advanced ECLSS systems and oxygen ISRU production. These investments are not mutually exclusive but should be co-optimized in the context of an overall campaign. b) Given the relationship between oxygen produced per year and ISRU system mass, it appears that an ISRU plant may break even within the first year of operations on the lunar surface. This requires, however, a high level of reliability (100% uptime, no spares required) and sufficient available power.

IV. Conclusions and Future Work

An optimization tool for NASA's ISRU System Model was developed for lunar oxygen production plants that provides a means for analyzing and comparing different system configurations and technology options.   It enables users to apply optimization to varying system designs and understand the effect of sizing parameters on system performance. Optimization tests performed using SQP and simulated annealing found that between these two methods, a feasible, minimum point can usually be found. However, higher oxygen production levels did cause more erratic behavior in the models, resulting in few converged solutions. Further tests with these methods and refinement of the ISRU component models is necessary to yield a fully robust system optimization.

Initial tests of nine potential system architectures for increasing quantities of oxygen production provided a non-dominated set of optimized designs. Two particular system architectures populated the Pareto-optimal set of designs for most production levels tested. These are attractive architectures to investigate further. This set of designs demonstrated that ISRU oxygen production plants may exhibit an economy of scale, indicating that investment in larger plants would result in less total launched mass in a long-term build-up of oxygen production capability. This conclusion will be tested as more models are developed and experimental data provides validation and verification.

An initial analysis of ISRU impacts to lunar-base supply chain logistics indicated that if ISRU plant masses scale as per Figure 8, using in-situ derived oxygen for crew breathing air may be beneficial for low-to-medium capability ECLSS systems. A more detailed trade study is necessary to compare ISRU to the advanced ECLSS systems to determine how the system masses compare. ISRU and advanced ECLSS have some overlap in subsystem technology, so it is conceivable that they may have similar system masses or may share in common parts.   Future work in lunar campaign analysis will include full integration of ISRU analysis into SpaceNet, assessment of supply chain impacts of oxygen production for life support and propellant over different lunar architecture scenarios, and investigation of Constraint Satisfaction Problem (CSP) techniques for architecture optimization.

References

[1]Sanders, G., Duke, M. (eds), "NASA In-Situ Resource Utilization Capability Roadmap Final Report." Documents Archive, Lunar and Planetary Institute, Houston, TX, May 2005 (unpublished).

[2]Steffen, C., Freeh, J., Linne, D., Faykus, E., Gallo, C., Green, R., "System Modeling of Lunar Oxygen Production: Mass and Power Requirements", *Proceedings of Space Nuclear Conference 2007*, Paper 2049, Boston, June 2007.

[3]Santiago-Maldonado, E., Linne, D., "ISRU System Model Tool: From Excavation to Oxygen Production", *Space Resources Roundtable IX,* LPI Contribution No 1375, Lunar and Planetary Institute, Houston, 2007.

[4]Wakayama, S., "Blended-Wing-Body Optimization Problem Set-up", *8th AIAA/USAF/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, AIAA-2000-4740, Sept 2000.

[5]Hassan, R., Crossley, W., "Multi-Objective Optimization of Communication Satellites with Two-Branch Tournament Genetic Algorithm", *Journal of Spacecraft and Rockets*, AIAA, Vol 40 No2, March 2003.

[6]Graff C., de Weck O., "A Modular State-Vector Based Modeling Architecture for Diesel Exhaust System Design, Analysis and Optimization", AIAA-2006-7068, 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Portsmouth, Virginia, 6 - 8 Sep 2006.

[7]Sanders, G., "ISRU: An Overview of NASA's Current Development Activities and Log-Term Goals", 38^{th} *Aerospace Sciences Meeting & Exhibit*, AIAA 2000-1062, Jan 2000.

[8]Rice, E., Gustafson, R., "Review of Current Indigenous Space Resource Utilization (ISRU) Research and Development", 38^{th} *Aerospace Sciences Meeting & Exhibit*, AIAA 2000-1057, Jan 2000.

[9]Eagle Engineering, Inc. "Conceptual Design of a Lunar Oxygen Pilot Plant: Lunar Base Systems Study", NASA-CR-172082, July, 1988.

[10] Belachgar, H., Labib, M., et all, "FERTILE Moon: Feasibility of Extraction of Resources and Toolkit for In-Situ Lunar Exploration", *International Space University Final Report,* International Space University, Strasbourg, FR, 2006.

[11]Hegde, U., Balasubramaniam, R., Gokoglu, S., "Analysis of Thermal and Reaction Times for Hydrogen Reduction of Lunar Regolith", *Space Technology and Applications International Forum—STAIF 2008,* American Institute of Physics, pp 195-202, 2008.

[12]Balasubramaniam, R., Hegde, U., Gokoglu, S., "Carbothermal Processing of Lunar Regolith Using Methane", *Space Technology and Applications International Forum—STAIF 2008*, American Institute of Physics, pp 157-161, 2008.

[13]Papalambros, P., Wilde, D., *Principals of Optimal Design Modeling and Computation*, 2^{nd} ed., Cambridge University, New York, 2000, pp 349-352.

[14]Wilcox, K., Wakayama, S., "Simultaneous Optimization of a Multiple-Aircraft Family", Journal of Aircraft, AIAA, Vol 40, No 4, July 2003.

[15] S. Kirkpatrick, C. G. Jr., and M. Vecchi, "Optimization by Simulated Annealing", *Science,* Vol. 220, No 4598, pp. 671–680, 1983.

[16]Lieberman, M., "Market Growth, Economies of Scale, and Plant Size in the Chemical Processing Industries", *Journal of Industrial Economics*, Vol 36, No. 2, pp 175-191, 1987.

[17] de Weck O.L., Simchi-Levi D., Shishko R., Ahn J., Gralla E., Klabjan D., Mellein J., Shull A., Siddiqi A., Bairstow B, Lee G., "SpaceNet v1.3 User's Guide", NASA/TP-2007-214725, January 2007.

[18]Armaer, N., Siddiqi, A., de Weck, O., Lee, G., Jordan, E., Shishko, R., "Design of Experiments in Campaign Logistics Analysis", *AIAA Space 2008 Conference & Exposition*, San Diego, California, 9-11 Sep. 2008.

14