

Tradespace Exploration of Space Settlement Architectures Using Long-term Cost and Benefit Metrics

George Lordos¹, Markus Guerster², Bruce Cameron³, Olivier de Weck⁴ and Jeffrey Hoffman⁵
Massachusetts Institute of Technology, 77 Massachusetts Ave, #33-409, Cambridge, MA 02139

¹glordos@mit.edu, 857-249-8404, ²guerster@mit.edu, ³bcameron@mit.edu,
⁴deweck@mit.edu, 617-253-0255, ⁵jhoffma1@mit.edu, 617-452-2353

Abstract— Space systems development choices made decades ago, such as decisions for the Space Shuttle and the International Space Station (ISS), still influence and constrain human spaceflight technology development in the present day. Similarly, upcoming decisions in NASA’s Artemis program and other Moon / Mars programs will likely influence the development of human spaceflight capabilities for subsequent decades. Given this, we argue that space systems architects require a new approach suitable for tradespace exploration across decades, and we have developed a new metric and tools to illustrate this approach. Typical tradespace exploration approaches trade off metrics for cost and benefit. The main families of metrics are monetary, utility and physical. Unfortunately, dollar, mass or utility metrics are fraught with weak assumptions when applied to decades-long timelines and to space settlements that will benefit from reusable rockets and in-situ resource utilization. Past work by the authors had proposed a new physics-based cost metric suitable for multi-decade architectures termed Lifetime Embodied Energy. Here, we propose a new benefit metric which may be of interest to diverse stakeholders and suitable for the comparison of alternative long-lived architectures: the Sustainable Long-term Growth Rate (SLGR). The SLGR has two components, a sustainability test to be defined by the architect and an average growth rate for key internal and external Figures of Merit (FOM). We show the application of the SLGR metric to a (future) long-term Mars base campaign using only Apollo-derived notional elements without ISRU, and to a case of a human settlement on Mars using future in-situ resource utilization technologies. The key FOM in both scenarios are habitable volume, stock of life support equipment and stock of consumables, as well as their derivative, which is carrying capacity of humans. We use a System Dynamics model to simulate the long-term growth, decline and stability behaviors of these FOM variables under each design option. For this model, we defined the SLGR’s sustainability test as avoidance of collapse, meaning that the key FOM’s per capita must remain above predefined dynamic thresholds at all times within the multi-decade model horizon. Architectures which collapse are deemed potentially fragile and discarded. For the remaining, potentially robust architectures, the growth rate of each of the key FOM across the multi-decade model horizon is its SLGR. We conclude with an assessment of SLGR compared to other metrics, a discussion of potential pros and cons, and a framing of the SLGR methodology in terms of pathfinding as opposed to traditional tradespace exploration.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. REVIEW OF COMMONLY USED TRADESPACE METRICS	2
3. DEVELOPING NEW LONG-TERM METRICS	3
4. MODELING THE SLGR.....	4
5. THE MARS SETTLEMENT CASE STUDY.....	5
6. MODEL RESULTS AND DISCUSSION.....	7
7. EPILOGUE: FROM TRADESPACE EXPLORATION TO PATHFINDING	9
APPENDIX A: MODEL DIAGRAMS	10
APPENDIX B: WHAT IS SYSTEM DYNAMICS?.....	14
REFERENCES.....	15
BIOGRAPHY.....	15

1. INTRODUCTION

Per the experience of Apollo, Space Shuttle and the International Space Station (ISS), the development of human space exploration capabilities is heavily path-dependent. Choices made decades ago still influence and constrain human spaceflight options in the present day, and thus decisions made for the Artemis program and for Mars exploration today will likely significantly influence the nature and extent of humanity’s spacefaring capabilities deep into the distant future.

One spacefaring capability of particular interest is the bootstrapping and development of persistent human settlements on other worlds. Space system architects are therefore interested in the capability of modeling the potential evolution of human space settlements through multi-decade time spans to inform the handful of early, highly impactful decisions that will largely constrain the future path of development of an off-world settlement system. In other words, space mission architects require methods, tools and metrics suitable for tradespace exploration across the very long term.

This paper reviews commonly used tradespace exploration metrics, finds that they are not optimal for the tradespace exploration of very long-lived human space settlement architectures, and proposes a suitable new metric, the

Sustainable Long-term Growth Rate metric (SLGR). Hence, the objective of this paper is to propose a new metric for very long-lived complex socio-technical systems, here specialized to human space settlements. Very long-lived socio-technical systems may require different metrics from shorter-lived systems because stakeholders, needs and system cost and performance can be much more fluid over the very long term. Using well-suited metrics for very-long-lived systems is made more important by the fact that a longer system lifetime will increase lifecycle costs and benefits, making good up-front architectural design even more important.

This paper is organized as follows: Section 2 briefly reviews the applicability of commonly used tradespace exploration metrics in the architecting of long-lived space systems. Section 3 describes a generic approach for selecting appropriate metrics for very long-lived (multi-decade) systems and shows how this approach was previously applied by the authors to develop Lifetime Embodied Energy (LEE), a physics-based long-term cost metric. Section 4 describes how the proposed long-term benefit metric (SLGR) emerges from a time-domain model. Section 5 describes the Mars settlement case study, Section 6 presents and reviews the model results, and Section 7 brings LEE and SLGR together into a proposed pathfinding paradigm for the tradespace exploration of long-lived sociotechnical architectures.

2. REVIEW OF COMMONLY USED TRADESPACE METRICS

Tradespace exploration metrics often fall into the categories of either ‘cost’ or ‘benefit’. In the tradespace paradigm, more of a benefit and less of a cost are desired, and vice versa. A third type of metric, ‘risk’, has been treated in different ways, such as a modeled variance in the expected benefits and costs per Shishko [1], or an uncertainty that can manifest as both upside and downside risk per Ross and Hastings [2]. This section reviews a selection of commonly used cost and benefit metrics for tradespace exploration, and discusses their applicability to the architecting of very long-lived complex socio-technical systems such as human settlements on other worlds. This review is intended to highlight common limitations with classical cost and benefit metrics which hinder their potential applicability to long-lived projects. The metrics can be categorized into monetary, utility, and physical basis metrics.

Monetary basis metrics

LCC — an example of a monetary basis metric is Life Cycle Cost (LCC). LCC analysis at NASA dates back to 1969 [3] and accounts for all costs, including Design, Development, Test and Evaluation (DDT&E); launch and emplacement; operations, and disposal [4]. LCC has its strength when exploring alternatives which aim to properly estimate true lifetime cost of a project. Estimates regarding potential future costs are uncertain, with uncertainty increasing as the analysis extends over longer time horizons. Uncertainty can be reduced by statistical analysis of historical costs.

Life cycle costing models sit one level above the non-recurring-cost DDT&E models and the recurring-cost Operations Cost models, incorporating their outputs as inputs and producing multi-year cost phasing as their outputs. Launch and emplacement cost is driven by system (i.e. payload) mass, destination, timeframes and orbit selection, and the operations cost component of life cycle cost analysis relies partially on DDT&E as a cost driver. This means that system mass has an indirect effect on the operations cost component of LCC via DDT&E [4].

Lordos [ibid.] notes that all model-based estimates for DDT&E, which is the first of the three major inputs into life cycle cost analysis, rely on system dry mass as the main input. Further, all of these models use historic cost data as inputs, whether directly or indirectly via Cost Estimating Relationships (CER’s), and historic cost data based on the old aerospace industry supply chain structure is declining in relevance for cost estimating purposes as newer companies have demonstrated the commercial operation of space systems at costs that vary by up to an order of magnitude.

Zapata [5] has shown that widely used cost models would have overestimated the cost of DDT&E of the Falcon 9 and Dragon system by a factor of 10, and the operating costs by a factor of 3. This large discrepancy implies that the historic cost data used by these approaches may now be out of date. Since much of the benefit of relying on historic cost data for LCC is the reduction in subjectivity of forward-looking estimates of life cycle economic costs, the value of LCC as a true life cycle metric may legitimately be questioned.

Adding to these concerns are the realities that (i) total life cycle dollar costs for a multi-year project can give rise to sticker shock which is easily politicized, and (ii) that critical new in-space activities are being contemplated, such as ISRU and long-term surface habitats, for which we have no truly comparable historic cost data at all [4].

The above underlie our motivation to search for cost and benefit metrics for in-space activities which sidestep the need to use historic monetary costs.

Utility basis metrics

MATE — Multi-Attribute Tradespace Exploration (MATE) is an approach which includes both need identification with exploration and evaluation at the architecture and the design levels [6]. The key step is the generation of a value function which aggregates and reflects the decision maker’s various preferences. This implies the need to select relative weights among these preferences, and also to map how the performance of the system translates to satisfaction, or utility, in each preference dimension. Ultimately, this enables the scoring and ranking of all the alternative system designs in accordance with the value function. The weighted utility function of the decision maker is then the benefit function, and there is also some cost function, with the two coming together into a tradespace of alternative architectures or designs.

As a result of the above, for the MATE system to be used effectively, one needs a clearly defined scope and set of stakeholders. However, with very long-lived and evolving engineering systems, such as human settlements on other worlds, the stakeholders and/or their needs may not be well defined, or might evolve and change during the useful life of the long-lived system. Even granting that ultimately, the only stakeholder is the decision maker, it is not clear that they know their needs with respect to a multi-decade project such as a city on Mars, or even that anyone could accurately compile such a list of needs. This is likely to hinder both the selection of attributes that should be included, as well as the weights to be assigned to each in the context of a multi-attribute utility function.

Physical basis metric

IMLEO — A commonly used and long-established physics-based proxy for space mission cost is Initial Mass in Low Earth Orbit (IMLEO). The use of IMLEO as a proxy for cost facilitates tradespace exploration, especially at the early stages of concept generation, selection and development where dollar cost figures would be difficult or cumbersome to estimate and assign.

IMLEO has held up well for decades as a good proxy for total space mission cost owing largely to the high launch costs: beyond their direct cost, the high launch costs underlie an intense mass-minimization and risk-retirement focus throughout the industry. Therefore, in most costing models, payload mass impacts Design, Development, Test & Evaluation costs (DDT&E). This provides a second good reason for the suitability of IMLEO as a proxy for cost.

However, changes are under way that tend to weaken IMLEO's suitability as a cost proxy for space missions, especially with respect to long term human exploration campaigns on the surfaces of other worlds: (1) As launch costs fall due to reusability of both booster and upper stage (e.g. Starship / Superheavy), other cost items which are not necessarily proportional to payload mass increasingly drive the total mission and campaign cost. For example, falling launch costs enable very large satellite constellations, motivating investment in assembly lines for the mass production of thousands of flat-packed Starlink satellites where the amortization cost of the investment depends more on the total number of satellites rather than their mass; and, (2) Campaigns of many sequential, similar missions to the same site on planetary surfaces are being planned, for example by NASA to the Mars Exploration Zones, or by private companies to envisaged outposts on the Moon and Mars. These target sites on planetary surfaces will accumulate persistent infrastructure and are likely to feature in-situ propellant production from local resources. The capability to refill propellant at the destination "resets" the rocket equation and disrupts both the reach and economics of spaceflight.

In conclusion, IMLEO has been a useful proxy for total mission cost for several decades. However, going forward,

total space mission cost is progressively being decoupled from IMLEO and will increasingly be driven by mission architecture and context instead.

3. DEVELOPING NEW LONG-TERM METRICS

Applicability to Long-lived Evolvable Sociotechnical Systems

We have reviewed one classical, well established example of a cost or benefit metric from each of the three main families of metrics – monetary, utility and physics based. In all three cases, we have identified certain issues with the potential applicability of these traditional metrics to the tradespace exploration of long-lived, evolvable socio-technical systems such as human space settlements. We now explore the possibility of the development of new long-term metrics that would address the issues identified.

Proposed Principles for New Long-term Metrics

One way to address the issues with traditional metrics identified in the previous section is to step back and try to identify principles that all proposed long-term metrics should likely have to satisfy in order to be well suited to the *ex ante* evaluation of alternative designs for long-lived, evolvable socio-technical systems. One potential principle is that for long-lived systems, a lifecycle *physics* based metric is likely to dominate a lifecycle *monetary* metric in terms of suitability. The reasoning is that this would sidestep the inherent difficulties with making assumptions regarding discount rates and inflation rates over the very long term. It would also sidestep difficulties with separating out what is true economic cost and what is an oligopolistic profit in the case where databases of historic costs are used in the cost modeling process. The operation of this principle has led Lordos [4] to propose Lifetime Embodied Energy as a candidate long-term cost and value metric for human space settlements.

A second potential principle is that a candidate 'good' benefit metric for a long-lived, evolvable sociotechnical system ought to be aligned (or at least, not misaligned) with any of the needs of the primary stakeholder(s). Of course, the challenge identified previously remains: that these needs may not be well defined or even definable. Thus, the potential principle is narrowed by incorporation of this reality to the proposition that a good long-term benefit metric is likely to be a "needs-neutral" benefit metric. This would rule out many traditional need-specific metrics such as safety, performance, value, etc.

We first review the work on Lifetime Embodied Energy, a proposed metric that is compatible with the first principle, before turning to the development of a novel metric that is compatible with the second principle.

An example: the Lifetime Embodied Energy Metric

In the search for a lifecycle, physics-based cost metric for in-space activities that is not denominated in units of mass or dollars, Lordos reviewed the literature for energy theories of value [4]. Energy-based value and cost metrics such as

embodied energy had been developed during the oil crises of the 1970's and are in use today only in the narrow sector of the energy performance of buildings. Embodied energy is calculated by summing and allocating all the past energy flows required to develop the system of interest. As these energy flows can be estimated analytically from first principles for any well-understood physical process, embodied energy as proposed by Howard Odum in his 1983 book [7] is a good candidate for a long-term, physics-based cost metric which would be compatible with the first principle described above. Odum measured embodied energy in embodied joules or emjoules, a unit that is defined below. When we add emjoules already expended to emjoules likely to be expended in the future to sustain a system in operation so that it can deliver its output, we obtain lifetime embodied energy - a generalized, physics-based lifecycle cost for any system.

Lifetime Embodied Energy Definition

Lifetime Embodied Energy (LEE) is defined as the thermodynamic sum of past, present and future work required to create, operate, maintain and decommission a system, including appropriate shares of indirect contributions from upstream systems as well as from other systems in a whole product system. The units of LEE are embodied joules; embodied joules are not conserved – they can be created and destroyed depending on losses to entropy and the method of calculation [4].

There are three main methodological approaches for embodied energy analyses: sectoral input-output models, process-based models and hybrid models. The depth and breadth of Earth-based supply chains make the application of process-based models on Earth challenging, leading to the truncation of process trees and to large error bars. These difficulties would not apply to a Mars industrial development context, thereby making application of the LEE methodology easier for space systems than it is on Earth [4].

The important methodological choice of primary energy source and system boundary in embodied energy modeling is linked to the nature of the underlying figure of merit targeted for optimization in each specific case. For example, by specifying the embodied energy of space logistics as the primary energy source and common cost denominator of all technological systems and processed resources delivered to Mars, it is possible to closely replicate the conclusions of standard IMLEO-based or Equivalent System Mass (ESM) based analyses, but also to go beyond them [4].

4. MODELING THE SLGR

Justifying the selection of SLGR as a long-term benefit metric

Given its roots in physics and thermodynamics, Lifetime Embodied Energy may have application as a long-term cost metric. However, as discussed above under the second proposed principle, a good long-term benefit metric should measure something that is potentially of value, or at least not harming, most needs of the primary stakeholder(s) over the

very long term, even when those needs are hard or even impossible to define. We recall that we have restricted the scope to only very long-lived, evolvable socio-technical architectures. Assuming that the primary stakeholder(s) / sponsor(s) of the evolvable sociotechnical system will have some future needs that will be met by this system, these stakeholders have an interest in the sustainability of the system, so that it will continue to meet the needs, and in its growth, so that it will have more embodied energy with which to carry out its functions in order to meet the various (undefined) needs. We therefore hypothesize that sustainability and growth over the long term are two attributes that, when taken together, may meet the second principle stated above, of being desirable (or not harming) most stakeholders.

One way of creating a “sustainable growth” metric which incorporates both of these attributes is by creating a dynamic model of the evolvable system at a sufficient level of fidelity to capture all potentially important / relevant modes of growth, stagnation or collapse. Within such a model, sustainability can be defined as the avoidance of collapse of key figures of merit of the evolvable socio-technical system over a very long time horizon. The choice of what constitutes collapse would be up to the architect, but the set of collapse conditions would certainly include the collapse of a key variable to zero, and might include collapse of a key variable to some fraction below its starting or peak values.

Architectures modeled in this way, and which manage to avoid collapse over the long term, will exhibit patterns of different growth rates for different internal and externally delivered variables. As we have already filtered out the architectures which exhibited collapse, these growth rates of the surviving architectures may be termed ‘sustainable’ growth rates. In view of the above, we propose a new benefit metric for use in tradespace exploration of extremely long-lived, evolvable sociotechnical systems: the Sustainable Long-term Growth Rate (SLGR).

System Dynamics models of growth, stagnation and collapse

Since SLGR requires the estimation (and likely, the simulation) of state variables across a very long time dimension, a useful first step to calculate a SLGR is the creation of a dynamic model. For this work, we have chosen to develop our model using the System Dynamics methodology [8] and the Vensim DSS software package. System Dynamics (SD) was developed by Jay Forrester at MIT, who used it in the 1960's to build the World3 model – a sociotechnical model of the Earth which was later used in the classic 1972 book *The Limits to Growth* [9]. SD is a useful modeling method for very long-lived human space settlements as it enables the time-step modeling of stocks, flow rates and information variables for both physical and social systems within the same environment. Just like *Limits to Growth*, we have chosen to abstract system variables at a very high level and to work with a small set of ‘necessary and sufficient’ key variables that are likely to constrain the main

performance metric of the settlement system, namely the carrying capacity of the human space settlement.

Having built a model which generates time series data for all key variables of interest, the next step is to convert the time series data into SLGRs for each of these key variables.

The consistent calculation of SLGR for variables of interest

An algorithm is needed to consistently calculate SLGR for variables of interest. A good algorithm would use all the time series data produced by the model for each variable and convert the entire time series to a single dimensionless growth rate, while taking into account various characteristics of the evolution of that variable. Ideally, the same algorithm should be able to convert time series data into SLGR's for variables exhibiting very different behavior modes, such as fast, slow or no growth, early, late or no stagnation, collapse or no collapse. A simple way (but not the only way) to accomplish this is to measure the time t for each variable to grow from its starting value a to the arithmetic mean μ of its entire modeled time-series, and use only these three variables to back out a compound growth rate characterizing a notional growth of the variable from its starting value to its mean time-series value. The main assumption needed to do this is the judicious selection of a long, fixed time horizon. Then, the time-series data for each variable of interest over the same time horizon yields both μ and t . Thus defined, the equation for the consistent calculation of the SLGR is:

$$SLGR = \left(\frac{\mu}{a}\right)^{\frac{1}{t}} - 1$$

5. THE MARS SETTLEMENT CASE STUDY

In the Mars settlement case study presented here, our model tracks the evolution through several decades of the following high level variables: (1) habitable, pressurized volume, measured in m^3 (2) life support system maximum capacity measured by the number of persons that can be supported, and (3) the stock of all consumables available to the settlement, measured in kg. For simplicity's sake, energy was not modeled – an energy-rich settlement was assumed instead, meaning that we assume that labor or resource bottlenecks will be felt long before energy bottlenecks. These three variables were chosen because all three are hypothesized to be required for sustainable growth; more variables may be included in the analysis as desired. All assumptions and variables in the model are encoded using stocks (boxes), flows (valves) or information links (borderless text), with a governing equation for each of these model elements. The model is decomposed into six interconnected sub-models. The first sub-model simulates the growth of habitable volume over time and is shown in Fig. 1.

As modeled, the architect of the very long-lived space settlement architecture has two main decisions to make: (1) how many tons of habitable-volume-related payload to deliver to Mars with each synod ('Hab-related payload delivered to Mars'), and what fraction of that habitable-volume-related-payload should be construction equipment for the in situ production of new habitable volume that takes advantage of in-situ resource utilization ('Kv as Fraction of Hab-related payload'). The consequences of these decisions then play out in the decades to come when the model is run, as Hab construction capital and habitable pressurized volume accumulate and depreciate with the passage of time.

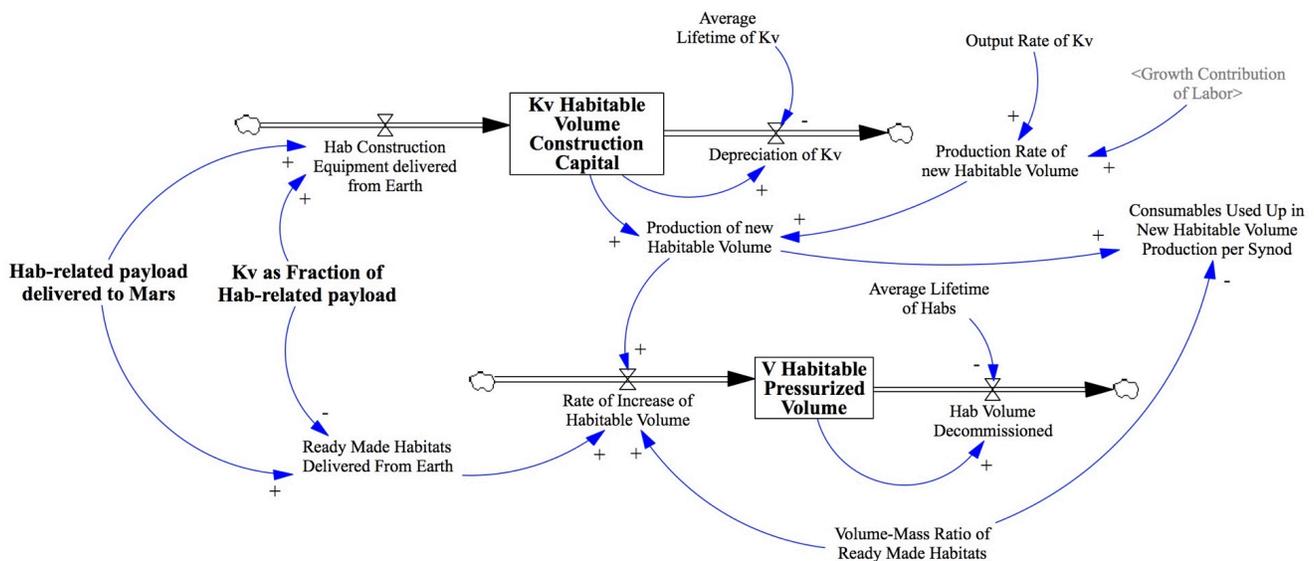


Figure 1. One of the six system dynamics sub-models, showing the payload mass allocation decision variables on the left and the two stocks: V, being the Habitable Pressurized Volume and K_v, being the Habitable Volume Construction Capital equipment. The inflows and outflows to and from these stocks are interconnected by the necessary and sufficient information variables that are required to generate their governing equations.

Another two sub-models, similar in modeling approach to Fig. 1, tackle the dilemmas of Earth-made vs. locally produced life support systems and Earth-made vs. locally sourced consumables. Tying everything together are three more sub-models¹: one models the supply and demand for human labor and the interactions of labor with the habitat, ECLS and consumables sub-models; another models quality of life in the form of ratios such as habitable volume per person or consumables per person, and links that to arrivals and departures from Mars and the availability of human labor; and finally a dashboard model control page permits what-if changes to model parameters, and also outputs results in the form of custom graphs for key variables, where long-term model behavior may be visually observed. For example, the dashboard includes a fraction slider for pre-fabricated habitats from Earth and Earth-delivered habitat construction equipment, modeling the strategic payload manifesting decision that is depicted in Fig. 1 and that the architect will have to resolve: should we take habitats with us, or should we take equipment to build them at the destination using local resources?

To implement the payload manifesting decision depicted in Fig. 1 (as well as similar decisions with respect to life support systems and consumables; for their model diagrams, see Appendix A), the model adopts an Isocost approach, whereby 180 tons of cargo and up to 100 crew are delivered to Mars with every synodic opportunity, for a modeling horizon of 60

synods (~130 years). Crew tour of duty is 5 synods (~10 years). In this model, a collapse condition is defined as any one of the key variables collapsing to less than 50% of its starting value at any point within the model horizon. Architectures which collapse can be deemed potentially unsustainable and discarded. For all architectures that are not ruled out as potentially unsustainable, the compound annual growth rate of each of the key variables is its Sustainable Long-term Growth Rate. The output vector consists of the SLGR for each of the key internal variables of pressurized volume, ECLS carrying capacity, stock of consumables, and the SLGR for the top-level figure of merit for a very long-lived human space settlement, which is the carrying capacity of humans. This carrying capacity, represented in the model as “Number of Persons on Mars”, is not a fixed number but co-determined with the desired, targeted or actual quality of life. The quality of life itself is a composite of sufficient Habitable volume, sufficient ECLS capacity and material wealth. We defined material wealth as the capability to address various needs using stockpiled consumables, and in the model this is proxied by the dimensionless ratio of the stockpile of consumables per person, divided by an 8-ton reference value, being two years of ‘nominal’ consumption per person. The reasoning is that having at least two years’ worth of stockpiled consumables per person provides material security under a variety of contingencies and Earth rescue scenarios.

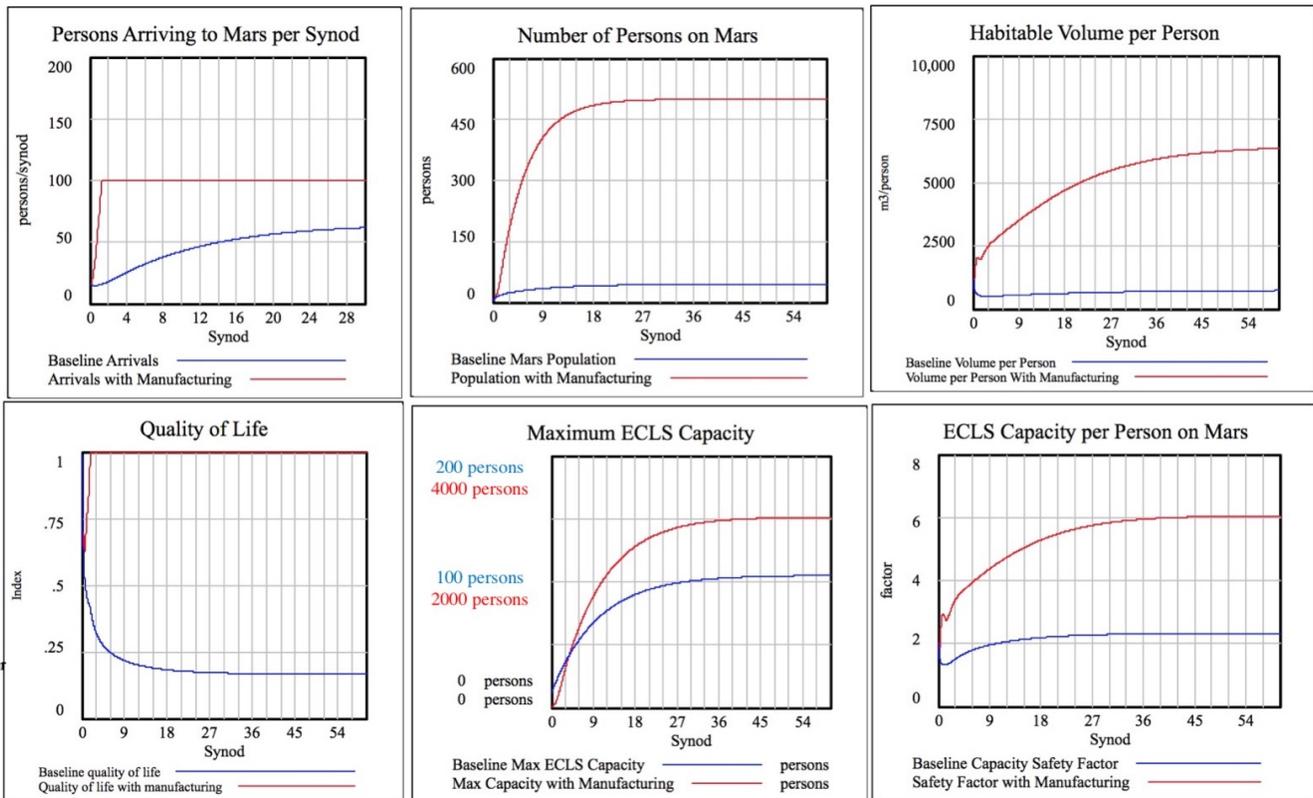


Figure 2. Model results. For the red (manufacturing on Mars) case, note the sustainable growth rate of Number of Persons on Mars, as well as the sustainable growth in Habitable Volume per Person and ECLS capacity per person.

¹ See Appendix A for the model diagrams of all sub-models

The above approach places the spotlight on the high-level tradeoff decision of building and testing the entire equipment manifest on Earth, vs. building, testing and deploying some of the equipment manifest locally on Mars. This approach abstracts away from many of the details and dynamics in the interest of maintaining model simplicity and clarity so as to communicate the idea of the SLGR. Running the model returns the evolution of the key variables across time.

6. MODEL RESULTS AND DISCUSSION

Running the model to obtain time-series data for variables

A baseline, ‘‘Apollo-like’’ case is first run with no in-situ manufacturing capabilities, i.e. with 100% of the fixed 180t cargo capacity per synod allocated to end-use systems as follows: 60t of prefabricated habitats, 60t of life support systems and 60t of consumables. These starting conditions generate the blue baseline curves shown in Fig. 2. which represent an ‘Apollo-like’ architecture: humans on Mars are logistically supported for all settlement needs from Earth. The manufacturing case is then run. One of many alternative manufacturing cases can be modeled by allocating 50% of the habitat-related payload to habitat construction equipment, 70% of the ECLS-related payload to ECLS manufacturing equipment and 50% of the ISRU-related payload to resource processing system manufacturing equipment. All three must be supplied with in situ resources and human labor to produce output, and this is represented in the model structure shown in Appendix A. The red curves in Fig. 2 represent a mission architecture which includes manufacturing of habitats, life support systems and consumables on Mars in the critical path.

SLGR comparisons for key Mars settlement variables

Continuing with the sub-model shown in Fig. 1 above, Table 1 shows the SLGR result for Habitable volume per person:

Table 1: Calculation of SLGR for the Habitable Volume per person and comparison vs. Apollo-like baseline.

Habitable Volume per person	Units	Apollo-like Baseline	Manufacturing on Mars
Starting Value	m3 per person on Mars	1200	1200
Minimum Value	m3 per person on Mars	496	1047
Mean Value	m3 per person on Mars	661.8	5118.6
Time to Attain Mean	synods	0.5	22.4
SLGR (Start to Mean)	N/A	-68.5%	6.7%
Max Decline vs StartVal	N/A	-59%	-13%

From the results in Table 1 we observe that both the Apollo-like baseline and the Mars manufacturing architectures start with the same total quantity of habitable volume per person (1200 m³). However, the Apollo-like baseline experiences a rapid (<1 synod) 59% decline in habitable volume per person from this starting value to its minimum value, and its SLGR is negative, meaning that this variable would fail the sustainability test and would flag the Apollo-like architecture to potentially be discarded. In contrast, the Mars manufacturing architecture reaches its eventual time-based mean of 5,118 m³ of habitable volume per person within about 22 synods, yielding a SLGR of 6.7%.

‘Habitable volume per person’, while interesting, is merely an internal variable relevant only for its contribution to carrying capacity: the model uses it together with other internal variables to determine the maximum number of people that can be present on Mars. This is the overall carrying capacity, which is the metric of the overall performance of the human settlement on Mars. In order to compare entire architectures against each other, we would use the SLGR of the overall carrying capacity, shown for our case study in Table 2 below:

Table 2: Calculation of SLGR for the Overall Carrying Capacity of settlement and comparison vs. Apollo-like baseline.

Overall Carrying Capacity of Settlement	Units	Apollo-like Baseline	Manufacturing on Mars
Starting Value	persons	6	6
Minimum Value	persons	6	6
Mean Value	persons	40.6	452.6
Time to Attain Mean	synods	16.9	12.5
SLGR (Start to Mean)	N/A	12.0%	41.4%
Max Decline vs StartVal	N/A	0%	0%

We note that in the Mars manufacturing case, the settlement achieved an order of magnitude higher carrying capacity (452.6 persons vs 40.6) in ~25% less time than the Apollo-like baseline case (12.5 synods vs. 16.9). The SLGR calculation yields a SLGR of 41.4% for the Mars manufacturing case, vs. a SLGR of 12% for the Apollo-like baseline. Data for the SLGR of all key variables is presented in Table 3 below.

Beyond the calculation of the SLGR, it is useful for the architect to qualitatively observe model behavior in plots of the key variables across time, such as the plots shown in Fig. 2 from which we extract the following insights:

1. Outbound crew flights gradually saturate at <60 crew in the baseline case, whereas they quickly reach their full capacity of 100 crew per flight in the manufacturing case. This is because the baseline case cannot grow its carrying capacity fast enough.
2. There is an order of magnitude difference in the long-term sustainable population on Mars, from about 50 in the baseline case to about 500 in the manufacturing case.
3. The manufacturing case, which models the construction of habitats using in-situ resources, shows a fivefold increase in pressurized habitable volume per person relative to the Apollo-like baseline. This is all the more remarkable given the order of magnitude higher population.
4. The maximum ECLS capacity and ECLS capacities per person show that the life support systems in the case with manufacturing can support ~one order of magnitude more persons than the baseline case. It can also provide a 6x safety factor relative to ~2x in the baseline, as the ECLS capacity under manufacturing grows faster than the Isocost-capped arrivals of 100 crew per synod who stay for 5 synods.
5. Not surprisingly, given the large differences in per-person margins for habitable volume, life support and consumables, the quality of life index rises to its maximum value of 1 early on in the case of manufacturing, but declines rapidly to below 0.5 and then gradually to ~0.2 in the baseline case. If the quality of life is considered a key variable by the architect, this event could also be deemed to have triggered the collapse condition for the baseline, Apollo-like architecture.

Assessment of SLGR vs. typical metrics

The above results hint at the essential relationship between SLGR and the traditional metrics reviewed in Section 2

above. Whether our (undefined) stakeholder needs were focused on having more persons on Mars, or more ECLS safety margin, or more net present value, or more quality of life, or some composite multi-attribute utility function of two or more of the above, the modeled SLGR values appear to be correlated with meeting the hidden, unobserved future need. This can be seen in Table 3, where the Mars Manufacturing architecture quantitatively dominates the Apollo-like baseline architecture in each of the seven different variables for which a SLGR was calculated. The per-person SLGRs for the internal variables are 5% - 7% for Mars manufacturing vs. zero or negative for the Apollo-like baseline. For the absolute values of the same variables, the SLGRs for manufacturing are several times higher than the baseline. If some of the key variables will turn out to be metrics for future needs, this finding – that SLGRs correlate with the likelihood of meeting future needs - is functionally similar to a capability to anticipate and foreshadow future needs.

Advantages of SLGR approach

The main advantage of the approach lies in the ability to quantitatively model and rank alternative architectures by a metric that will turn out to be correlated with a variety of future needs, including needs that have not yet been defined, and by stakeholders who have not yet been born. Another advantage is that it is relatively easy to create models at the high level of abstraction that is sufficient to study these dynamics. Such models will not be accurate, but as long as the modeled structure and behavior encodes the most consequential decisions and events, then the model should rank the alternative paths and enable order-of-magnitude comparisons of alternative architectures. This mirrors the approach in system architecture of prioritizing the handful of architectural decisions which exhibit the highest sensitivity with respect to performance outcomes and the highest connectivity (downstream impact) with the thousands of lower-level decisions.

Table 3: Calculation of SLGR for overall carrying capacity, and for all internal variables in absolute and per person terms.

Variable	Units	Dataset	Starting Value	Minimum Value	Mean Value	Synods to Attain Mean	SLGR (Start to Mean)	Max Decline vs StartVal
ECLS Carrying Capacity per person	Persons supported per person on Mars	baseline.vdfx	2	1.3	2.2	16.8	0.5%	-34%
		mars_manufacturing.vdfx	2	2	5.4	19.4	5.2%	-13%
Consumables Stockpiled per person	kg per person on Mars	Baseline	10000	1328	1605.6	11.9	-14.3%	-87%
		Manufacturing	10000	4953	66262.6	31.6	6.2%	-50%
Habitable volume per person	m3 per person on Mars	Baseline	1200	496	661.8	0.5	-68.5%	-59%
		Manufacturing	1200	1047	5118.6	22.4	6.7%	-13%
Carrying capacity of ECLS systems	Persons supported	Baseline	12	12	89.8	17.9	11.9%	0%
		Manufacturing	12	12	2522.5	17.3	36.1%	0%
Consumables Stockpiled	kg	Baseline	60000	60000	60430.2	4.6	0.2%	0%
		Manufacturing	60000	56504	32433400.0	31.1	22.4%	-6%
Habitable volume	m3	Baseline	7200	7200	27236.7	21.1	6.5%	0%
		Manufacturing	7200	7200	2431900.0	20.6	32.6%	0%
OVERALL CARRYING CAPACITY OF SETTLEMENT	persons	Apollo-like	6	6	40.6	16.9	12.0%	0%
		ISRU & manufacturing	6	6	452.6	12.5	41.4%	0%

7. EPILOGUE: FROM TRADESPACE EXPLORATION TO PATHFINDING

Given the realities of path dependence, particularly where human spaceflight is concerned, some decisions will always be far more consequential than others. The Sustainable Long-term Growth Rate metric, when used in the context of long-lived, evolvable sociotechnical systems, does not require that needs be defined or that weights be assigned when making early, consequential decisions that will affect or constrain the design of a long lived, evolvable sociotechnical system. This ability to potentially meet future needs without necessarily knowing them can be likened to pathfinding, a paradigm that would be distinct from classical tradespace exploration where the specific metrics are more clearly defined. SLGR is therefore useful in the context of long-lived, evolvable sociotechnical systems, precisely because it is hard to make predictions or to discern needs related to events that will happen in the distant future, or to know which decisions made today will influence an evolvable system in the distant future.

Hence, the uncertain future value of the undefined future needs may well end up being maximized *ex post* by the *ex ante* **pathfinding action** of always striving to maximize the growth rate of the system in a sustainable way. This is not a surprising result, as it is very similar to a ‘finding’ made by the operation of evolution by natural selection in living organisms. It is in fact yet another embodiment of the timeless Maximum Power Principle described by Odum [7]. The models presented here span times long enough to allow opportunity to the modeler to endogenize the most important interactions and, hopefully, capture the most important dynamic modes of behavior that might lead to growth, stagnation or collapse of the key variables.

“...The maximum power principle can be stated: During self-organization, system designs develop and prevail that maximize power intake, energy transformation, and those uses that reinforce production and efficiency...”

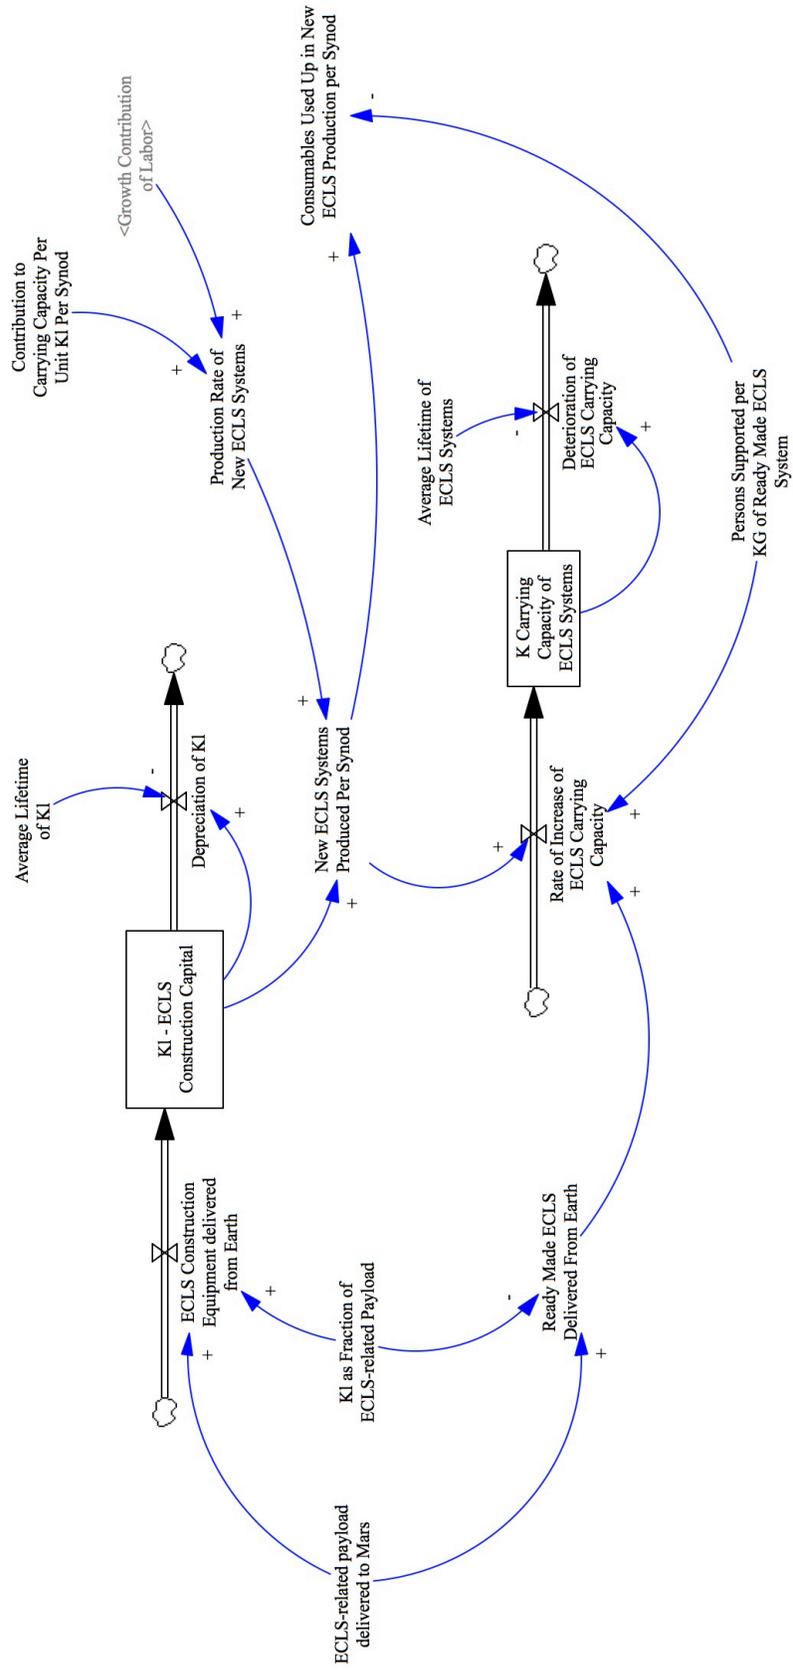
Howard Odum

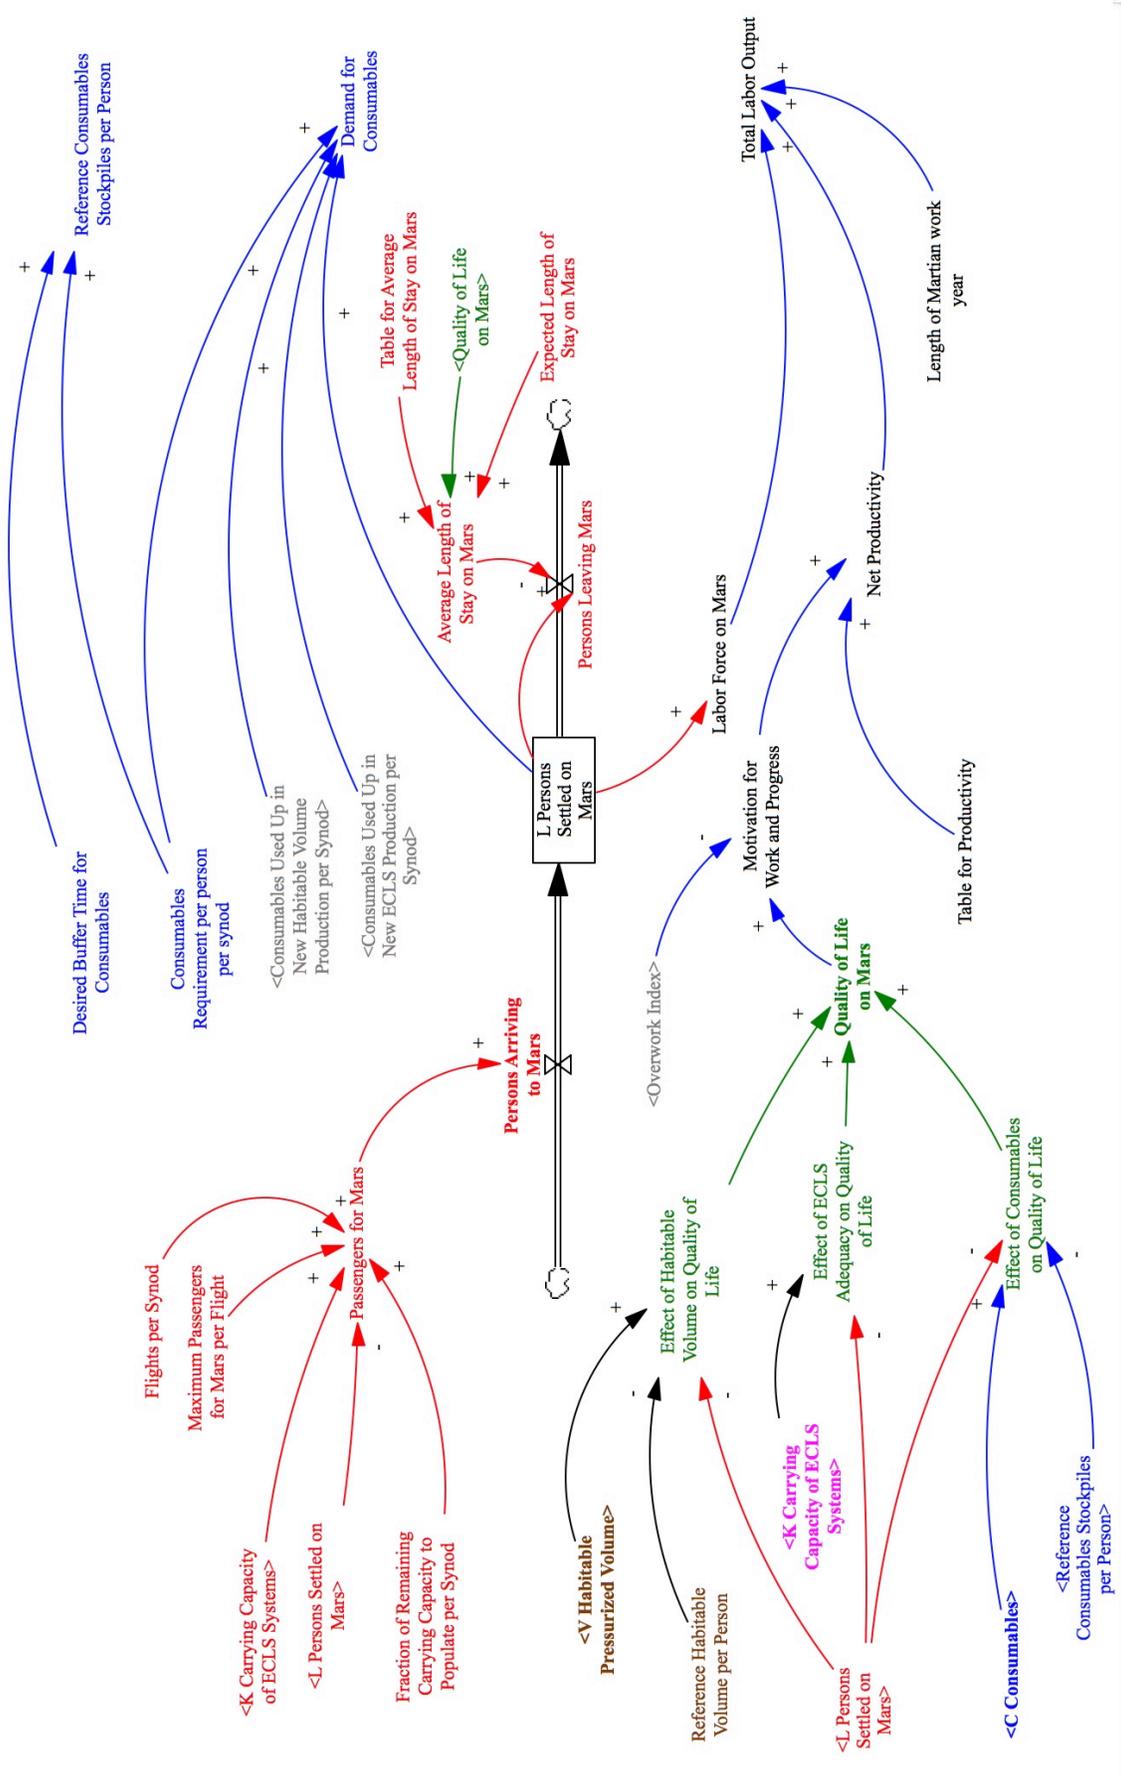
In conclusion, the previously-developed, physics-based lifetime cost of effort measured by the Lifetime Embodied Energy metric, and the beneficial potential for autocatalytic growth measured by the newly introduced Sustainable Long-term Growth Rate have been juxtaposed here as two independent metrics useful to designers of long lived, evolvable sociotechnical systems. We hope that system architects will find LEE and SLGR useful for modeling the impacts of present-day decisions on the distant future of long-lived, evolvable sociotechnical systems, such as a village on the Moon or a city on Mars.

Limitations and future work

One key limitation and weakness of the method at this stage in its development is that the modeling is deterministic and sensitive to which key variables and highest-level architectural decisions are modeled. An analysis of the sensitivity of model outcomes to certain modeling decisions may point the way to creating probability distributions of key model assumptions or elements, so that a Monte Carlo analysis of the outcomes of the system dynamics model will replace the analysis of the deterministic results. Future work will replace deterministic assumptions with probability distributions and modify the sustainability test accordingly, to test the robustness of alternative architectures against alternative ‘future histories’ and cull a larger number of architectures from the set of potentially sustainable ones.

APPENDIX A: MODEL DIAGRAMS





APPENDIX B: WHAT IS SYSTEM DYNAMICS?

The text enclosed in quote marks & cited below is the official description of System Dynamics on www.systemdynamics.org:

“Overview

System Dynamics is a computer-aided approach to policy analysis and design. It applies to dynamic problems arising in complex social, managerial, economic, or ecological systems—literally any dynamic systems characterized by interdependence, mutual interaction, information feedback, and circular causality.

The System Dynamics Approach

The approach begins with defining problems dynamically, proceeds through mapping and modeling stages, to steps for building confidence in the model and its policy implications.

Modeling and Simulation

Mathematically, the basic structure of a formal System Dynamics computer simulation model is a system of coupled, nonlinear, first-order differential (or integral) equations. Simulation of such systems is easily accomplished by partitioning simulated time into discrete intervals of length dt and stepping the system through time one dt at a time.

Feedback Thinking

Conceptually, the feedback concept is at the heart of the System Dynamics approach. Diagrams of loops of information feedback and circular causality are tools for conceptualizing the structure of a complex system and for communicating model-based insights.

Loop Dominance and Nonlinearity

The loop concept underlying feedback and circular causality by itself is not enough, however. The explanatory power and insightfulness of feedback understandings also rest on the notions of active structure and loop dominance.

The Endogenous Point of View

The concept of endogenous change is fundamental to the System Dynamics approach. It dictates aspects of model formulation: exogenous disturbances are seen at most as triggers of system behavior; the causes are contained within the structure of the system itself.

System Structure

These ideas are captured in Forrester’s (1969) organizing framework for system structure:

Closed boundary
 Feedback loops
 Levels
 Rates
 Goal
 Observed condition
 Discrepancy
 Desired action

Levels and Rates

Stocks (levels) and the flows (rates) that affect them are essential components of system structure. Stocks (accumulations, state variables) are the memory of a dynamic system and are the sources of its disequilibrium and dynamic behavior.

Behavior as a Consequence of Structure

The System Dynamics approach emphasizes a continuous view. The continuous view strives to look beyond events to see the dynamic patterns underlying them. Moreover, the continuous view focuses not on discrete decisions but on the policy structure underlying decisions. Events and decisions are seen as surface phenomena that ride on an underlying tide of system structure and behavior.”*

*Source: <https://www.systemdynamics.org/what-is-sd> (retrieved Jan 7th, 2020). Adapted from GP Richardson, System Dynamics. In Encyclopedia of Operations Research and Management Science, Saul Gass and Carl Harris, eds., Kluwer Academic Publishers, 1999/2011.

REFERENCES

- [1] R. Shishko, *Nasa Systems Engineering Handbook*. NASA, 1995.
- [2] A. M. Ross and D. E. Hastings, "The tradespace exploration paradigm," *INCOSE Int. Symp.* 2005, p. 13, 2005.
- [3] H. W. Jones, "Estimating the Life Cycle Cost of Space Systems," in *45th International Conference on Environmental Systems*, 2015, no. July, pp. 1–13.
- [4] G. C. Lordos, "Towards the Sustainable Industrial Development of Mars: Comparing Novel ISRU / ISM Architectures Using Lifetime Embodied Energy," Massachusetts Institute of Technology, 2018.
- [5] E. Zapata, "An Assessment of Cost Improvements in the NASA," in *AIAA Space 2017*, 2017, pp. 1–35.
- [6] A. M. Ross, "MULTI -ATTRIBUTE TRADESPACE EXPLORATION WITH CONCURRENT DESIGN AS A VALUE -CENTRIC FRAMEWORK FOR SPACE SYSTEM ARCHITECTURE AND DESIGN," Massachusetts Institute of Technology, 2003.
- [7] H. T. Odum, *Systems Ecology*. John Wiley & Sons, Inc, 1983.
- [8] J. Sterman, "System Dynamics Modeling: Tools for Learning in a Complex World," *Calif. Manage. Rev.*, vol. 43, no. 4, p. 25, 2001.
- [9] D. H. Meadows, D. L. Meadows, and J. Randers, *The Limits to Growth*. 1972.

BIOGRAPHY



George Lordos received a B.A. in Philosophy, Politics and Economics from the University of Oxford in 1991, a MBA from MIT's Sloan School of Management in 2000 and a Masters in Engineering and Management from MIT's System Design and Management Program in 2018. George is currently a Ph.D. Candidate with the Engineering Systems Lab at MIT's Department of Aeronautics and Astronautics, researching the industrial ecology of human settlements on the Moon and Mars. In between his education, George had a 25-year professional career as a technical project manager, strategy consultant, system architect, company director and entrepreneur.



Markus Guerster is a graduate student pursuing his PhD degree in the System Architecture Lab of the Department of Aeronautics and Astronautics at MIT. His interests include the development of tools to support the decision-making process for complex engineering systems under uncertainty. He received a M.Sc. degree (2017) with high distinction in Aerospace Engineering as well as a B.Sc. degree (2015) in Mechanical Engineering from the Technical University of Munich. He worked as a System Engineer at OHB System AG for 2 years besides his studies and is conducting a scientific study with in-tech GmbH to transfer aerospace methodologies to the railway industry.



Bruce Cameron is the Director of the System Architecture Lab at MIT and a consultant on platform strategies. At MIT, Dr. Cameron ran the MIT Commonality study, a 16 firm investigation of platforming returns. Dr. Cameron's current clients include Fortune 500 firms in high tech, aerospace, transportation, and consumer goods. Prior to MIT, Bruce worked as an engagement manager at a management consultancy and as a system engineer at MDA Space Systems, and has built hardware currently in orbit. Dr. Cameron received his undergraduate degree from the University of Toronto, and graduate degrees from MIT.



Olivier de Weck is a Professor of Aeronautics and Astronautics and Engineering Systems at MIT and a Fellow of INCOSE and Associate Fellow of AIAA. He focuses on systems engineering related to space logistics and space exploration, technology infusion analysis in human and robotic systems over time as well as multidisciplinary design optimization. He has published over 300 peer reviewed papers and 3 books and received 12 best paper awards since 2001. From 2013-2018 he served as Editor-in-Chief of the journal *Systems Engineering*, and more recently as Senior Vice President for Technology Planning and Roadmapping at Airbus.



Prof. Jeffrey A. Hoffman is a Professor of the Practice of Aeronautics and Astronautics at MIT since 2002, and co-director of MIT's Human Systems Lab since 2015. He is interested in the future of human spaceflight and in the use of the International Space Station as a testbed for future aerospace technology, especially in: human-machine interactions, EVA, flight operations and conducting laboratory research in space. From 1978-1997 he served as a NASA astronaut (five flights) and from 1997 – 2001 as NASA's European representative.