Autonomous Additive Construction on Mars

Samuel Wilkinson[1], Josef Musil[2], Jan Dierckx[3],
Irene Gallou[4] and Xavier de Kestelier[5]

Specialist Modelling Group, Foster and Partners,
Riverside, 22 Hester Road, London, SW11 4AN, UK.
[1]swilkinson@, [2]jmusil@, [3]jdierckx@,
[4]igallou@, [5]xdekeste@fosterandpartners.com

ABSTRACT

We present a conceptual construction process for an inhabitable outpost on Mars, using an autonomous multi-robot swarm approach to additively sinter layers of regolith into a protective shield over an inflatable pressurised module. The guiding design hypothesis investigated is that physically distributing risk across multiple simpler units working in parallel can improve chances of success, rather than the traditional consolidation into a single complex unit. This approach is fundamentally enabled by the decreasing size and cost of hardware. Larger numbers of more intelligent robots offers the possibility of emergent behaviour as the collective action of a complex group can be greater than that of individual independent units. In the paper we consider the general benefits of distributed redundancy, its implications for a robotic construction process using microwave power to selectively sinter layers of regolith in-situ, a subsequent habitat design with regolith shield, and a roadmap for a future technology demonstration.

INTRODUCTION

As various private and public organisations propose numerous missions to the moon, near-Earth asteroids, or Mars in the coming decades, certain trends in current solutions to surface habitation are observable. These include the use of: i) an initial robotic expeditionary mission to prepare for subsequent human arrival; ii) the combination of a delivered pressurised module for habitation and an in-situ compressive shield of regolith; and iii) large-scale additive construction with regolith. Assuming the first two points, and although there is of course dependence between them, we focus on the third for which there is specific debate about the best additive construction process. By construction process, it refers to both: the robotic system, i.e. a static vs. mobile configuration, the scale, capability, and number of units; and secondly, the material and bonding method.

Additive manufacturing offers many possibilities when applied in this context at larger scales to the construction of extra-planetary human habitats. If coupled with a robotic system, it enables the use, prior to human arrival, of local in-situ materials to work in conjunction with other delivered habitat elements. There are obvious complimentary advantages in delivering inflatable light-weight pressurised modules and a system to construct a heavy compressive regolith shield on top. The light-weight inhabitable element offers the high-tech aspects of life-support, manufactured in controlled environments on Earth, whilst the heavy-weight shield acts as protection from radiation,

micro-meteorites, and dust storms.

In the following sections we firstly describe the concept of distributed vs. consolidated redundancy which we posit will be possible in multi-robot systems in the near future. In the second section, we introduce the concept of a multi-robot swarm and its theoretical emergent behaviour. We then give a taxonometric overview of existing additive construction technologies, identify those with possible application, and introduce the process of melting regolith layers via microwaves into a proposed habitat shield with designs for three classes of construction robot. These sections come together into a proposed further work technology demonstration.

REDUNDANCY

System redundancy is often achieved through either passive or active means, where passive redundancy means an added strength capable of absorbing extra strain (this is common practice in structural engineering), and active redundancy where separate backup systems are present to come online when primary ones fail (common in electronic and computer engineering). In active redundancy, the extra backup capabilities are typically isolated or on standby, yet still integrated into the main system. In a dangerous environment this presents an issue when we consider that they are physically situated in a single fate-sharing machine; all our eggs are in one expensive basket. For example, a satellite may have passive redundancy in its physical build and active redundancy in its multiple spare computers on-board, yet its success is under risk due to consolidation of this complexity into a single device. An alternative paradigm would envisage distribution of the functionality, and risk, across multiple devices which can act together for the mission.

Many, if not all, complex systems in biology have a 'fine-grained' architecture (Mitchell, 2009), in that they consist of large numbers of relatively simple elements that work together in a highly parallel fashion. This fine-grained nature of the system not only allows many different paths to be explored, but it also allows the system to continually change its exploration paths, since only relatively simple micro-actions are taken at any time. Employing more coarse-grained actions (for instance with a single robot) would involve committing time to a particular exploration that might turn out not to be warranted. In this way, the fine-grained nature of exploration allows the system to fluidly and continuously adapt its exploration as a result of the information it obtains. Moreover, the redundancy inherent in fine-grained systems allows the system to work well even when the individual components are not perfectly reliable and the information available is incomplete or local. Redundancy allows many independent samples of information to be made, and allows fine-grained actions to be consequential only when taken by large numbers of components.

By the law of accelerating returns (Kurzweil, 1999, 2004), a generalisation extending Moore's law to technological development, the rate is observably exponential, i.e. whilst the cost and size of computation tends to zero, performance increases to infinity. Assuming this is broadly true, as it appears to be, and that it continues into the immediate future, we can take the premise that ever-increasing functionality will be possible

on ever-smaller, ever-cheaper machines. This will in turn enable progressively smaller and capable robots suitable for enabling our own presence in space. We therefore posit that it is both beneficial and possible in the near future to distribute risk across multiple such robots to construct ourselves a habitat on another planet.

SWARM ROBOTICS

In the near future it is predicted that computational intelligence (CI) and robotic technology will be sufficiently advanced to allow for a distributed system of autonomous intelligent machines. A system capable of adapting to uncertain operating environments, of self-management and system awareness, and of following high-level commands such as 'explore', 'gather materials' or 'construct habitat' (Nilsson, 2014; Truszkowski et al., 2006).

Behaviour has significance for understanding biological intelligence. It is understood as interactions between an organism and its environment where the actions of the organism affect its own perceptions, and thus its future actions and perceptions. Applying this concept to robots gives the field of autonomous robots that are behavioural machines capable of operating in partially unknown and changing environments without human intervention (Floreano and Mattiussi, 2008).

We apply these principles of behaviour and self-organization to collections of simple, autonomous robots. Simple is understood as not having sophisticated sensors, electronics or mechanics. Neither do they use global information to be centrally controlled and as a result algorithms tend to be simpler. The core idea of swarm robotics is to capitalize on simple interactions among robots in order to solve complex problems by means of emergent behaviour. This concept is inspired by a similar phenomenon found in nature, for example social insects.

Swarm robotics is mainly concerned with groups of robots that are larger than groups easily controlled with a centralized, top-down approach but within a number that is manufactured at the same or lower cost of a few complex robots. The main potential advantages of this approach is the robustness of the swarm to failure of individual robotic units. Further advantages are addition of units in real time and without changing the emergent behaviour, ability to cope with environmental noise and individual differences, and the emergent effect where the work of the swarm is greater than the sum of the work by the individual units.

Swarm robotics and uncertainty

Probabilistic robotics was developed after the 1970s when most research in robotics presupposed the availability of exact models of robots and their environments (Thrun et al., 2005). Little emphasis was placed on sensing and the intrinsic limitations of modelling complex environments. In the mid 1980s the paradigm shifted towards reactive techniques. Reactive controllers rely on capable sensors to generate robot control. Rejections of models were typical for researchers in this field. Since mid-1990s a new approach has begun to emerge: probabilistic robotics. This approach relies on statistical techniques to seamlessly integrate imperfect models and imperfect sensing.

The Martian environment is computationally difficult to simulate. The advantage of cooperative robots is that they operate locally within the immediate environment only and they do not need to understand the whole complexity, thus it is easier to simulate. Robots can use sensors to measure what has been built and update a digital model of the next printed layer to compensate for tolerances. Decentralisation provides a great robustness as no single failure can cause an overall failure (Brambilla et al., 2013).

AUTONOMOUS ADDITIVE REGOLITH CONSTRUCTION

Additive Construction (AC) refers to the application of Additive Manufacturing (AM) to an architectural scale; that is, the 3D printing of buildings. The distinctions are therefore application and scale, with AM typically for pieces less than 1m and AC for anything greater. Larger scale AC necessitates a higher deposition rate than AM, a larger construction system due to the greater size of prints, and usage of different materials. Each of these has a secondary effect of reducing the print resolution or accuracy to on average the order of 10mm, whereas AM today achieves sub-millimetre resolution.

In developing a taxonomy of additive construction techniques (Table 1), all examples can firstly be classified by one of two deposition methods: either process material internally *then* deposit (e.g. AM fused-deposition modelling (FDM)); or deposit *then* process externally (e.g. powder-based AM 3D printing). We consider the latter to be more practical in uncontrolled environments with variable material sources since it does not require internal processing which may lead to blockages.

The second classification is the binding method, generally divided into bonding (via an adhesive or chemical reaction) or by heating (sintering or melting). Sintering means heating until individual particles fuse together without melting to liquefaction (Figure 4). We propose that sintering consecutive layers of materials together is more logical as it is less energy intensive than completely melting and does not require transportation or in-situ processing of adhesives.

The third category is the material used for construction. In the examples given, this includes plastics, glasses, ceramics, metal, sand, concrete, and regolith. Without more complex in-situ material processing technologies, we are essentially left with three viable methods of layered in-situ regolith sintering: microwave, laser, or solar sintering.

In Table 1, the examples of AC highlighted in light grey have all used real or simulant lunar or Martian regolith as construction material. These have all been considered for or have potential for use in our context. Those highlighted in dark grey use regolith and layered in-situ microwave sintering, and are the closest precedent technology. All of the examples given have been demonstrated to some extent on Earth.

Site preparation

Before construction can begin however, the system must be delivered to the surface and the site must be prepared. The Entry, Descent and Landing (EDL) will happen

Table 1: Taxonometric overview of additive construction technologies

Construction Method	Binding Method	Base Material	Additives	Support Material	Comments	Example
Extrusion deposition	Chemical	Sand	Sulfur, Portland cement, plastics, Sorel cement	YES	Portland/Sorel cement are not appropriate in vacuum. Regolith + sulfur is an ISRU option.	Loughborough University, Skanska, F+P et al.
				-	Portland/Sorel cement are not appropriate in vacuum. Regolith + sulfur is an ISRU option.	WinSun
						Bruil
						WASP
						BetAbram
						CyBe
						Spetsavia
						xTreeE
						Le Roux
						TU Eindhoven
						Con3D
						Rudenko
						Contour Crafting USC
	Fused-deposition method (FDM)	Plastics	-	-	-	Dirk van der Kooij
						TechmerES, ORNL, SOM
						KamerMaker
		Glass	-	-	-	MIT Mediated Matter
		Ceramics	-	-	-	IAAC
						Robocasting SNL
	Microwave melting	Regolith	-	-	Melted in chamber and extruded.	JPL, PISCES
	Gluing	Regolith	Urethane	-	Not appropriate for large scales.	Adherent
	Powder spray	Regolith / Sand	Water, air	-	In a vacuum only possible if molten powder is sprayed.	IAAC
	Additive welding / Direct metal deposition	Metal	-	-	Requires pre-processed metal, time consuming.	Norsk Titanium Cranfield University
Layered in-situ	Chemical	Regolith / Sand	MgCl2	YES	Time consuming, energy intensive.	D-shape
						Voxeljet
	Laser sintering	Regolith	-	YES	-	Loughborough University
		Basalt	-	-	Advantageous in vacuum.	KSC
						PISCES
		Regolith / Sand	-	-	-	Aachen Uni
			-	-	-	NUS
			-	YES	-	EOS
	Solar sintering	Sand	-	YES	-	Kayser, M. (MIT)
		Regolith	-	-	-	NASA KSC
						PSI
	Microwave sintering	Regolith	-	-	Melted in-situ. Good penetration properties. Can be assisted by infrared heating. Difficult to shape.	JPL, PISCES
						University of Knoxville
	Selective inhibition sintering	Regolith	With or Without (MgO, Portland C/Water)	YES	Requires pressure. Also used for additive assembly e.g. tiles.	USC

in two phases: first the robots will be delivered to the surface for site selection and preparation; followed later by the ø4.6m habitat modules. For each, once stationary on the surface, the individual modules may be distributed within a landing radius of a few hundred metres, therefore it is necessary for them to navigate together and collect at a designated common point. A sequenced inflation-deflation of the external air bags starts a controlled roll of the modules to a shared target.

The first task for the robot system once deployed at the site is to excavate a 1.5m-deep hole for the habitat modules to sit within. The largest class of robots (RAC-D) will dig the loose regolith from the surface layer by layer, which will be moved nearby into protective berms by the medium-sized (RAC-T) robots (Figure 1a). The volume of excavated regolith is roughly equal to the amount to be printed on the shield. There may be an element of trial-and-error in the initial excavation due to unknown underground geology (i.e. hidden rocks), although this could be minimised by either selecting appropriate land formations such as a small crater, or by acoustic location.

(a) Multi-robot EDL and site preparation

(b) EDL and navigation of habitat units

(c) Stages of module deployment: (right to left) opening, inflation, and connection

Figure 1: Initial site preparation

Once the site is prepared, and the three modules are gathered together (Figure 1b), the small inflatable spheres surrounding the habitat partially deflate. The upper faces of the dodecahedron module fold up and lock in position, similarly for the lower faces to form the foundation which can fit to a rough landscape for the interior floor to be level. Subsequently, the core, shaped as a pentagonal prism, expands outwards. Each of the five vertical faces of the core is either a connection or an airlock which will move to the outer perimeter along with the inflated skin. The three deployed habitat modules are now ready for the regolith shield to be constructed on top (Figure 1c).

Regolith construction

The regolith additive construction (RAC) approach is designed for the low accuracy likely to be expected from using variable materials at an uncertain site with autonomous robots in the field. The reliability is partly in its distribution of tasks, implemented by the three classes of robot: the strategy is to 'dig, move, and melt' regolith.

The large digging robots will extract loose regolith in close proximity for the medium-sized mover robots to transport to the habitat. As the regolith is deposited, the smallest

melting robots have a microwave print head to bond one layer at a time (Barmatz et al., 2014). The regolith is positioned into rough layers of about 10mm thickness by the transporter robots, with the thickness continuously measured (Figure 3. Once a thin layer of regolith is in place, the third class of smallest robots selectively sinters patches into a hard crystalline material. The form and progressive construction of the shield is shown in Figure 2.

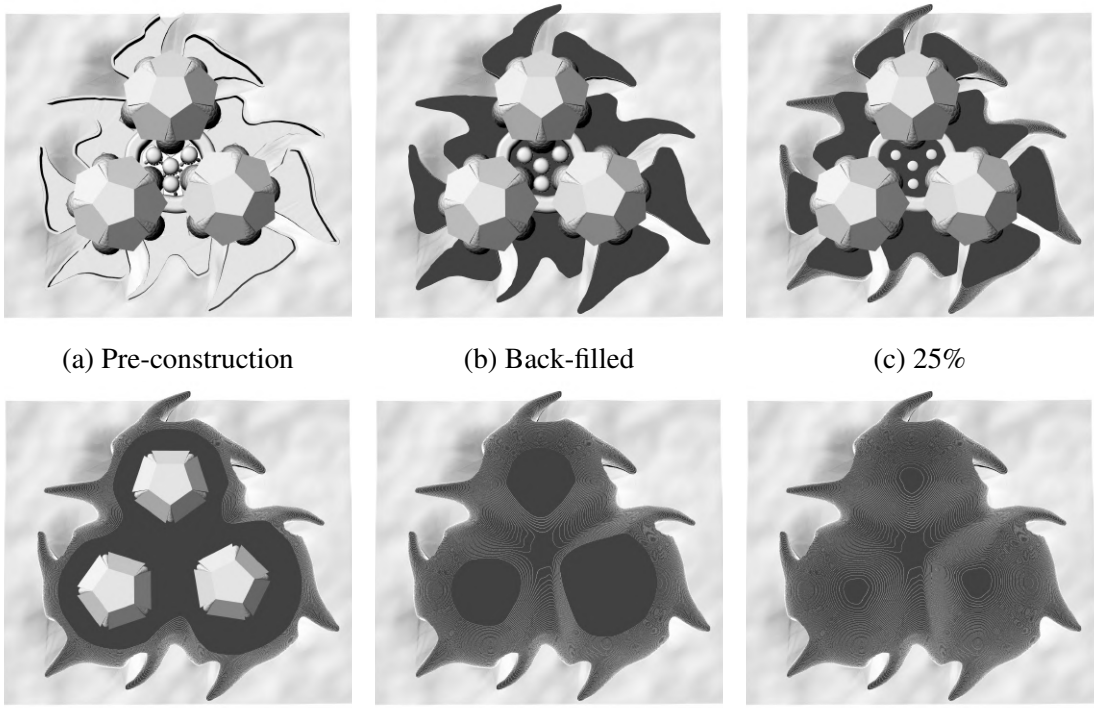

(a) Pre-construction (b) Back-filled (c) 25%

(d) 50% (e) 75% (f) Complete

Figure 2: Plan view of progressive regolith construction

Since each of the three modules has five possible connections, three of which are for entry, these areas must be left clear. As the construction increases in height, the perpendicular ramps extending radially allow for the robots to progress to the next layer.

The form of the regolith shield is driven by two key criteria. The first criteria is the minimal thickness of regolith needed to protect the inhabitants from radiation: a smaller one is applied to two of the three modules and the larger value for the third remaining module. This means that one module can serve as a safe space with extra protection in case of transient solar flares. For protection from radiation over long-term periods, rather than transporting heavy shielding from Earth, the regolith shell is a logical alternative. The largest reduction in dose equivalent (rem) occurs in the first 20g/cm^2, so assuming a regolith density of 1.5g/cm^3 the regolith depth should be at least 15cm (Simonsen and Nealy, 1991). Whilst this is a minimum depth to ensure the survival of the inhabitants, the design includes 1.5m above the work/sleep modules, and 2.5m above the communal space. This will improve the long-term health of the astronauts and provide a temporary shelter during periods of increased solar activity.

(a) Deposition and sintering of regolith layers by multiple RAC-T and RAC-M robots

(b) Operational occupied outpost with continuous repair and infrastructure construction

Figure 3: Additive construction process

The second factor of the shield's form is the ability of all construction robots to transfer themselves to the highest layer printed so far during the construction process. As a result, multiple ramp structures blended into the overall form are introduced next to every opening of each module (airlocks, windows and suitports). Because of their location, they also serve as an extra protection of these openings.

Microwave sintering

There are various possible methods available for bonding loose regolith for construction: namely chemical; with adhesive; freezing; compaction; laser or solar sintering; or microwave sintering. The last is investigated here due to its apparent robust application. Taylor and Meek (2005) and Barmatz et al. (2014) have both attempted sintering of real Lunar regolith and Martian regolith simulant (JSC-2A) respectively in the laboratory with 2.45GHz 200W microwaves. The bonding mechanism of sintering lunar regolith is shown in Figures 4a to 4d, and an example of the internal structure of a bonded sample in Figure 4e. Due to the thermal properties of the material, the sample has a heterogeneous layering of loose material, to partially bonded, through to a completely molten core.

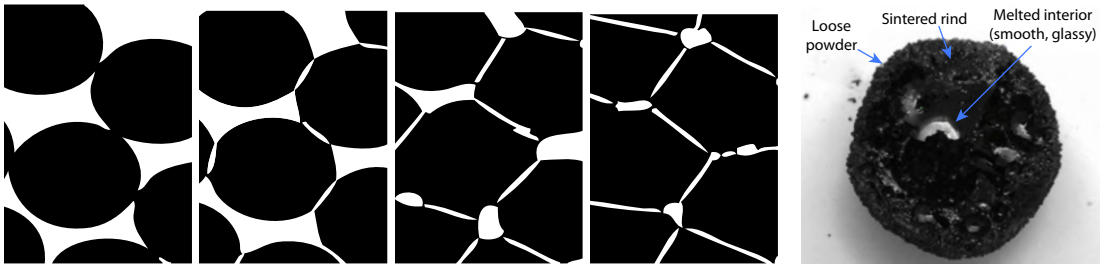

(a) (b) (c) (d) (e)

Figure 4: Regolith sintering: (a-d) Progressive sintering of lunar regolith (Hintze et al., 2009); (b) Interior of lunar JSC-2A after being heated to ˜600°C (Barmatz et al., 2014)

As classified in the earlier taxonomy, there are two methods to applying this either internally or externally. The first involves placing a layer of material and sintering it in-situ Taylor and Meek (2005). For example, the 'lunar road-paving wagon' can move back and forth with its magnetrons (microwave generators) that can be set to various frequencies and power, in order to effectively sinter the lunar regolith, thereby constructing a traffic-able road or launchpad.

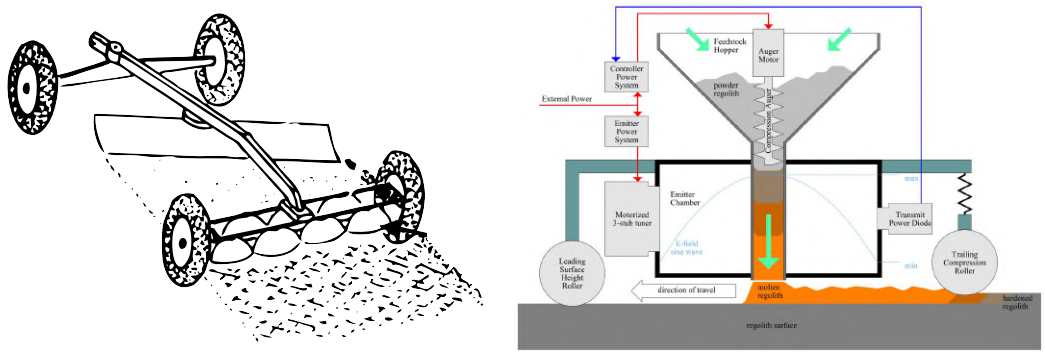

(a) (b)

Figure 5: Regolith sintering in-situ: (a) 'Lunar Road-Paving Wagon' designed by Taylor and Meek (2005); (b) Microwave printer concept by Barmatz et al. (2014)

In the second case, regolith is collected into a reservoir and fed through a microwave oven where it becomes molten and is subsequently extruded out into position (Barmatz et al., 2014).

Power beaming

The robot's batteries charge wirelessly via power beaming (Brown, 1996), a technology that is currently under development. This involves transmitting power over long distances through air, space, or optical fibres via a laser and receiving photovoltaic cells. Similar to solar power although much more intense, the transmitting station converts electricity to laser light which is transmitted to photovoltaic cells on the individual robots, where it is finally converted back to electricity.

The great advantage of wireless power is in enabling 'perpetual presence', i.e. with a constant supply of power the robots do not have to rely on finite battery lives. Tech demos with a quadcopter have achieved over 12 hours of constant flight (Nugent and Kare, 2010).

Swarm robotics behaviour

Distributed system of powering also allows for higher number of smaller robots to be controlled operate independently. A task harder to realize with larger robots, where each needs to carry its own power source able to power it over long period of time.

A programmable framework for a cooperative swarm robotic system that builds large scale architectural objects by means of additive manufacturing is here investigated. Greater precision material deposition is further enhanced by on-site measurements to

compensate for uncertainty in its simulation having stochastic parameters. Additive manufacturing is nowadays well established for small scale prototyping and manufacturing but needs to scale up in architectural context. Current research focuses on printed material research and leaves issues with architectural scale to either conventional assembly or limiting the overall size of printed object. Cooperative swarm robotics is a distributed system that allows for large scalability. Real time physical measurements are proposed in-conjunction with virtual simulation to react to under-defined and changing environmental noise.

Swarm robotics at Harvard (Werfel et al., 2014, 2011) shows a multi-agent construction system inspired by mound-building termites. First case of study tries to reverse engineer how to build a given form by a given number of independent robots. That solves the inverse problem of how low-level rules give specific outcomes, that is in general still little understood. Suggested is a further development of stochastic rules for agents for not fully pre-determined forms. This should be further developed and understood in the near future and will be suitable for distant places like Mars, where precise delivery of predefined from is not necessary, rather delivery of precise properties, like isolation and thickness of walls is necessary.

Another advantage of swarm robotics is the emergent outcome. Kilobot project at Harvard (Rubenstein et al., 2012) studies this emergent effect while keeping the price and complexity of each robot very low.

Robot design

The multi-robot system consists of three classes of robot deployed for site preparation and construction. For system diversity, robot movement is by wheels, tracks and legs respectively (Figure 6). Each robot has the ability to operate independently without any central command; therefore they need the ability for environmental sensing, local communication, and decision making. Any expected emergent behaviour from the group is dependent largely on the interactions between the individuals, either directly via communication or indirectly through the environment which they alter.

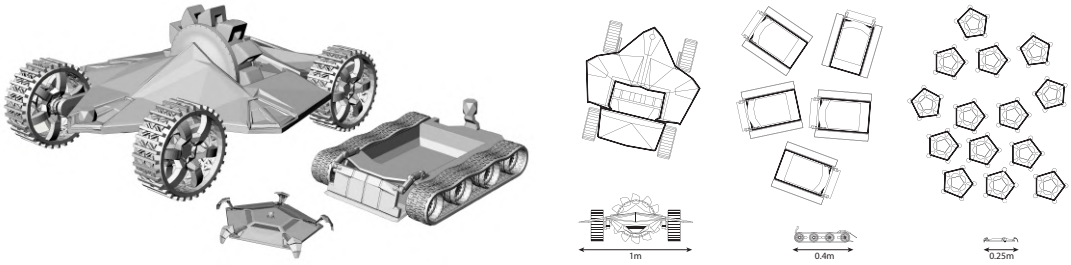

(a) (b)

Figure 6: Robot design: (large to small) RAC-D, RAC-T, RAC-M

Table 1: Construction robot specifications

Name	*Function*	*Size*	*Quantity*	*Movement*
RAC-D	'Digger', regolith excavation using a perpendicular bucket-wheel	1.00m	1	Wheels
RAC-T	'Transporter', moves and deposits regolith into thin layers	0.40m	5	Tracks
RAC-M	'Melter', microwaves patches of regolith in desired position	0.25m	10	Legs

FURTHER WORK: TECH DEMO

The next developmental steps in implementing this proposal can be broken down into four strategic development areas to be pursued. Firstly, operation of the robot system: development of individual robot mechanics, control, and communication; group test with multiple units to investigate the inverse problem of guaranteeing a satisfactory end result with input rules (physical and simulation); and power beaming from central power station to individual robots. Secondly, the additive construction process: material collection, processing, and layering; microwave sintering of in-situ regolith simulant; and combination of robotic and AC processes. Thirdly, habitat module design: EDL, navigation, inflation, deployment and connection; full-scale mock-up test with human test subjects. And finally, system integration with a fully automated field test in analogous Earth site.

CONCLUSION

The various aspects of redundancy that are explored here with the outpost delivery, construction, and design are intended to increase chances of success. A key element of this approach is allowing for flexible outcomes: from the initial navigation of the modules to find a suitable location, with the internal module layout, through to the semi-autonomous robotic additive construction.

In a largely uncharted extreme environment such as Mars, there are strict requirements on and of any such construction system. However, its specific activities and subsequent products must be less strictly controlled. Due to communication challenges and environmental unknowns, complete direct control is not possible and therefore the autonomy of, or trust placed in, the system must be high. In truth this is the case for robotic and future human missions. Relinquishing total control whilst ensuring through indirect methods the desired outcomes is the greatest challenge of this endeavour.

Acknowledgements

This work was initially entered in the AmericaMakes 3D Printed Habitat design competition, one of NASA's Centennial Challenges. We would like to thank the following companies and individuals for their assistance during the competition: Astrobotic Technology, Penelope Boston of New Mexico Tech, Malika Beggour, John Eager of the British Antarctic Survey, David A. Green of King's College London, and Rapha Clothing.

References

Barmatz, M., Steinfeld, D., Anderson, M. and Winterhalter, D. (2014), 3D microwave print head approach for processing lunar and mars regolith, *in* 'Lunar and Planetary Science Conference', Vol. 45, p. 1137.

Brambilla, M., Ferrante, E., Birattari, M. and Dorigo, M. (2013), 'Swarm robotics: a review from the swarm engineering perspective', *Swarm Intelligence* **7**(1), 1–41.

Brown, W. (1996), 'The history of wireless power transmission', *Solar energy* **56**(1).

Floreano, D. and Mattiussi, C. (2008), *Bio-inspired artificial intelligence: theories, methods, and technologies*, MIT press.

Hintze, P. E., Curran, J. and Back, T. (2009), Lunar surface stabilization via sintering or the use of heat cured polymers, *in* '47th AIAA Aerospace Science Meeting'.

Kurzweil, R. (1999), *The age of spiritual machines: how we will live, work and think in the new age of intelligent machines*, Orion.

Kurzweil, R. (2004), *The law of accelerating returns*, Springer.

Mitchell, M. (2009), *Complexity: A guided tour*, Oxford University Press.

Nilsson, N. (2014), *Principles of artificial intelligence*, Morgan Kaufmann.

Nugent, T. and Kare, J. (2010), 'Laser power for UAVs', *Laser Motive White Paper-Power Beaming for UAVs, NWEN* .

Rubenstein, M., Ahler, C. and Nagpal, R. (2012), 'Kilobot: A low cost scalable robot system for collective behaviors', *2012 IEEE International Conference on Robotics and Automation* .

Simonsen, L. and Nealy, J. (1991), Radiation protection for human missions to the moon and mars, Technical report, National Aeronautics and Space Administration, Hampton, VA (USA). Langley Research Center.

Taylor, L. and Meek, T. (2005), 'Microwave sintering of lunar soil: properties, theory, and practice', *Journal of Aerospace Engineering* .

Thrun, S., Burgard, W. and Fox, D. (2005), *Probabilistic robotics*, MIT press.

Truszkowski, W., Hinchey, M., Rash, J., Rouff, C. et al. (2006), 'Autonomous and autonomic systems: A paradigm for future space exploration missions', *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on* **36**(3), 279–291.

Werfel, J., Petersen, K. and Nagpal, R. (2011), Distributed multi-robot algorithms for the termes 3d collective construction system, Institute of Electrical and Electronics Engineers.

Werfel, J., Petersen, K. and Nagpal, R. (2014), 'Designing collective behavior in a termite-inspired robot construction team', *Science* **343**(6172), 754–758.