An Untethered Mobile Limb for Modular In-Space Assembly

Sawyer Brooks, Peter Godart, Paul Backes, Brendan Chamberlain-Simon, Russell Smith, Sisir Karumanchi
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
Sawyer.Brooks@jpl.nasa.gov

Abstract— **In-space assembly can enable new types of spacecraft and structures which are too large or fragile to be carried on a rocket in an assembled form, and robotic systems can make in-space assembly feasible and cost-effective. Such systems should be able to assemble large and complex structures while imposing minimal launch mass and mission risk. We propose an autonomous robotic limb, henceforth referred to as "Limbi," which is self-mobile and symmetric. Two identical electromechanical docking mechanisms serve as end-effectors. With either end-effector anchored to a base structure, the other can grab modular elements and attach them to the growing structure. Power and computing are provided by the spacecraft through these docks, enabling Limbi to walk end-over-end across the structure without a battery or tether. We have constructed and tested a prototype system in a planar workspace that demonstrates the mobility and assembly capabilities of the proposed limb. We also introduce the concept of "Limboids," consisting of multiple Limbi robots temporarily attached to each other to form more complex kinematic chains. The resulting configurations are application-specific and can be tailored to the degrees of freedom, range of motion, and general dexterity required by a particular task. Because Limbi and Limboids can assemble large and complex structures with minimal robotic complexity, the development of this class of robots is a critical step forward in low-risk and lightweight assembly.**

TABLE OF CONTENTS

978-1-4673-7676-1/16/$31.00 ©2016 IEEE

1. INTRODUCTION

In this paper, we define an architecture for in-space assembly known as "Limbi." The key features of the proposed architecture include end-over-end mobility with a single limb and semi-autonomous assembly of modular structures. A Limbi-style robot is powered and controlled through the structure it assembles, via electromechanical docks at either end of the limb, so that it does not require a battery or tether and can build indefinitely large modular structures.

Physically, Limbi is a multiple degree-of-freedom serial kinematic chain. Both ends of Limbi are electromechanical docks which provide a rigid mechanical connection to structural elements while also offering an electrical interface to provide power to the limb via the structure. In the context of this paper, "end-over-end mobility" will refer to the ability for the limb to dock to a structure at both ends, then release itself at one end while still receiving power through the other end, and finally maneuver the free end to another docking point and repeat the process. A "module" is an independent and disconnected element of a structure, which can attach to other modules and serves as an electrical pass-through to provide power to the limb. "Modular assembly" is the ability to manipulate and attach these modules to each other. The module-to-module interface does not need to be identical to the limb-to-module interface, but it does need to provide a mechanical and electrical connection between modules.

The Limbi architecture is minimal in that it is fully capable of these mobility and assembly tasks, as a single unit, while using no more actuators than the necessary degrees of freedom defined by the structure's geometry. This minimizes mass, complexity, and risk. Because limb power is provided through the structure, the limb can be untethered and does not require a battery. This allows it to work in an indefinitely large workspace, so the size of structures it can assemble is not limited by the robot.

We also describe the mechanical and electrical design of a testbed prototype of a Limbi robot, which is used in our laboratory to test new docking mechanisms, control methods, and autonomy algorithms. We explain some of the lessons learned through development of this testbed prototype, which help inform the practical requirements for Limbi-style robots, and we outline a plan for adding new features and functionality to the prototype.

Finally, we introduce a novel concept of "Limboids," which are two or more Limbi robots working in serial and parallel configurations to achieve tasks that would otherwise only be possible with specialized robots. We specifically focus on Limbi and Limboids in the context of in-space assembly, referring primarily to systems which are assumed to be in orbit around Earth or other bodies. However, the concepts may also be applicable to terrestrial assembly applications, which would introduce the additional constraint of gravity, or even aquatic assembly and inspection which would introduce buoyancy and waterproofing challenges.

2. Prior Work

In many respects, Canadarm2 on the International Space Station meets the same qualifications as a Limbi-style robot. It is capable of end-over-end mobility between a few "Power/Data Grapple Fixtures" (PDGFs) around the exterior of the station. It has also been directly used to attach new pieces to the station, such as airlocks and solar panels. However, it is designed to work with direct control from astronauts or ground controllers. An end-goal for the Limbi architecture is to support autonomy at all levels of control and all but high levels of planning and error-handling, so that human operators need only to specify a high level task plan. This will enable much more rapid assembly processes while reducing the demand for communication bandwidth. Still, Canadarm2 is a flight-tested example of one possible use case for the Limbi architecture.

Similarly, a mobile limb to assemble cubic modules was proposed by the FZI Research Center in Germany as part of their "Intelligent Building Blocks for On-Orbit Satellite Servicing" ([1] and [2]). The FZI research briefly proposes using a limb with end-over-end mobility to build satellites from cube-shaped elements, but does not detail the design requirements for such a limb or other applications of the technology.

Finally, another related concept was explored with a robot named Skyworker at Carnegie Mellon University [4]. Unlike Limbi, Skyworker has two "feet" which walk on a scaffolding structure, and a separate manipulator for carrying cargo, making it slightly more complex but also more capable than a single Limbi robot. To carry cargo with a single Limbi-style robot would require either a passive cargo-mounting point at the center of the limb, or for Limbi to incrementally move the cargo, then itself. As described here under "Limboid Concept," the capabilities of Skyworker can be replicated if multiple Limbi robots work together. Another key distinction between Skyworker and the Limbi architecture is the end-effectors on both robots. On Skyworker, the feet are grippers which mechanically connect to a scaffolding but do not include an electrical connection, and the robot is battery-powered. Under the Limbi architecture, power is transferred to the limb via electromechanical docks. This allows for indefinite range and duration without recharging, but also limits flexibility in terms of attachment points to the structure.

3. Testbed Prototype

To demonstrate the applicability of this research for in-space assembly, and to serve as a testbed for future research, we built a laboratory prototype of a Limbi-style robot. The actuators powering the limb are Schunk PowerCube rotary actuators, which were available freely in our lab. The electromechanical docks at the ends of the limb are made from 3D printed custom parts and off-the-shelf electromagnets. A set of representative modules for the limb to assemble was built from aluminum framing. To reduce complexity and permit the use of low-power actuators, it was decided early in the project to demonstrate Limbi in a planar workspace. The limb and modules are supported by casters on a smooth, flat, aluminum surface. By imposing this constraint, the effect of gravity can be neglected, and fewer, less powerful, actuators are needed. The limb and modules resting on the testbed floor are shown in Figure 2.

Limb Geometry and Kinematics

With one side of the limb docked, the other end has a maximum of three degrees-of-freedom: x, y, and θ. However, the current prototype uses four joints, giving a redundant degree-of-freedom which can be used to maneuver the center of the limb around obstacles. The geometry of the limb is shown in Figure 1. Note that because the limb is typically connected to a module on one or both sides, the geometry is considered with modules attached to both sides of the limb, and the first and last link lengths are in fact the distances to the center of the attached modules.

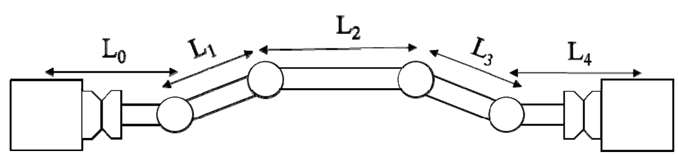

Figure 1: Limbi Prototype Geometry

To optimize the dimensions of the limb, with the constraint that the limb must be symmetric, a MATLAB script was used to iteratively test and score many combinations of link lengths with a fixed total limb length. To prioritize the ability to dock two modules next to each other, as well as the ability to walk end-over-end between neighboring modules, configurations which offered high dexterity with the two ends of the limb in close proximity were desirable. Each combination of link lengths was itself iteratively tested for a large set of joint angles, and the resulting geometry was scored according to the inverse of the distance between the two ends. The total score for a combination of link lengths was the sum of all tested joint-space configurations.

The results of this analysis suggest that the links at both ends of the limb (L_0 and L_4) should be as short as mechanically possible, and that L_1 and L_3 should each be approximately half the length of the middle link (L_2). A more refined version of this approach might be used to optimize dimensions of future iterations of Limbi, including

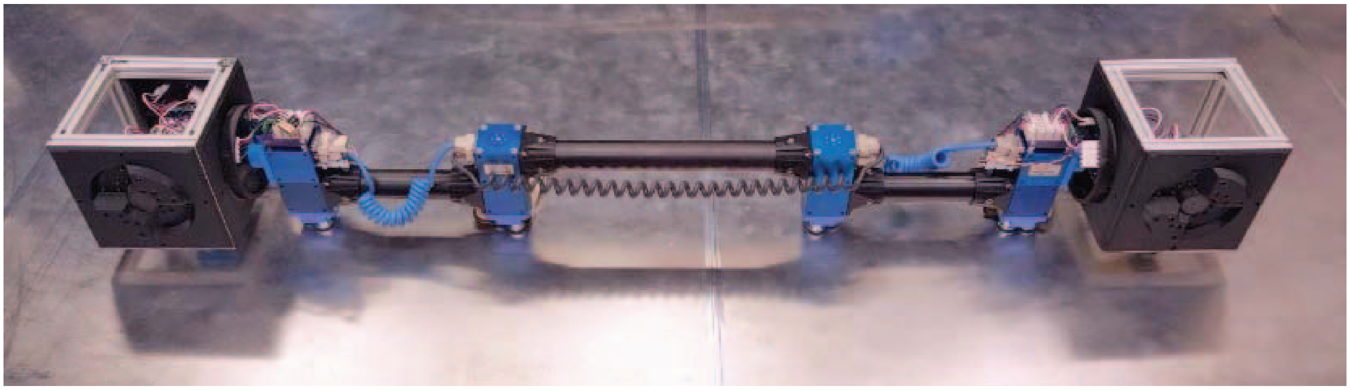

Figure 2: Testbed prototype of Limbi, with a module docked at both ends.

those operating in three-dimensional space. Including mechanical considerations, the final dimensions chosen for the link lengths are as follows: $L_0 = L_4 = 0.20$ meters, $L_1 = L_3 = 0.33$ m, and $L_2 = 0.54$ m, for a total limb length of 1.60 m including the modules.

Electromechanical Docking Mechanisms

To simplify the prototype, we designed a dock that could be used both for connecting the limb to the modules, and for connecting the modules to each other. This mechanism was designed with three primary constraints: rigidity, androgyny, and requiring minimal mating force.

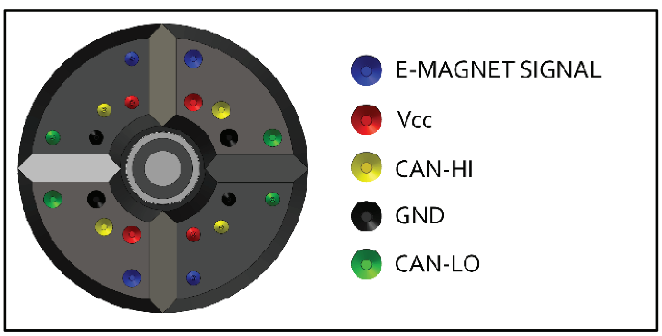

Figure 3: Front view of the electromechanical docks

The holding force is provided by an electromagnet at the center of the dock, which can be controlled by a digital signal from the limb. Two identical docks can mate if only one of the electromagnets is powered, because the other is a large block of ferrous metal. This supports the requirement for an androgynous dock. Additionally, the mechanism does not require the limb to exert significant force to successfully dock, because the dock is simply brought into position and the electromagnet is turned on.

A set of opposing chamfers along the vertical and horizontal axes ensures consistent alignment when two docks are held together by the electromagnet. The geometry of the chamfers also prevents any non-mating pin/pad pair from making contact (for example, power and ground).

Power and data are transferred through the dock by means of spring pins and exposed copper pads on printed circuit-boards (PCBs). Each dock has two sets of spring pins (in the top-left and bottom-right quadrants in the image above), and two mirrored sets of copper pads. The layout is symmetrical around the vertical and horizontal axes so that the docks are androgynous and can be rotated by 180 degrees.

Unfortunately, the 80-lb holding force electromagnets used in this prototype did not provide sufficient rigidity. The pull strength of an electromagnet drops rapidly with distance. As a result, the docks were not resistant against high moments, because as soon as the electromagnets separated slightly, the holding force would plummet and the limb would come undocked. Limb motions had to be planned to avoid exerting significant moments. Additionally, a slight amount of flexibility in each dock caused significant oscillations at the end of the limb, though it is unclear how much of this could be attributed to the materials used for the docks (3D-printed ABS plastic).

The spring pins and PCB pads were observed to suffer some corrosion after repeated use. This corrosion was most noticeable on the power and ground pads, and can likely be attributed to sparks crossing an air-gap immediately before and after docking. This could be mitigated by using better materials and surface finishes for the pads (the pads in the prototype were simply exposed copper traces on a PCB). Additionally, if Limbi was working in the vacuum of space, the sparks across the air gap would not occur. Regardless, this reveals a need for careful material selection for future versions of the docking mechanism. An internal switch could also be used to prevent power from flowing until the mechanical connection is verified.

Module Design

The prototype modules are made of 80-20 aluminum framing and are roughly 8” on a side. Docks are attached to three sides of each cube. The spring pads and copper pins on the docks are connected to a common PCB mounted at the center of the cube, so that power and data can be transferred between any two docks through the PCB. The circuit board also holds a small microcontroller which toggles an electromagnet in response to a digital signal from the limb.

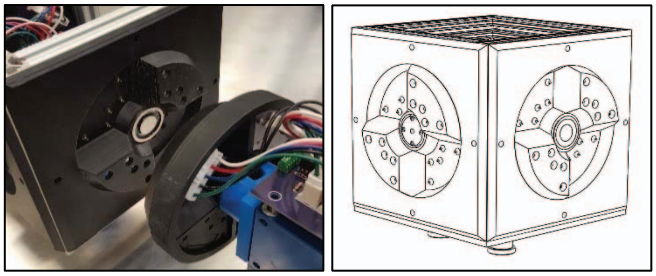

Figure 4: Left, the limb prepares to dock with a module. Right, the design for the representative modules.

Although the prototype modules are cube-shaped, the architecture can support many different shapes and sizes of structural elements. The docks on the modules are identical to the end effectors on the limb, with a small exception for the scope of the demo. With a pre-planned assembly pattern, certain module cube faces did not need to feature a dock that could actively engage and disengage. As such, steel plates were used in place of some electromagnets to reduce the cost of the prototype. If this docking architecture is used again, the steel plates will be replaced with electromagnets so that the docks are truly androgynous. This would optimize re-configurability.

System-level tests revealed unwanted compliance in the electromechanical docks when modules are assembled in long chains. This implies several important considerations for future hardware solutions. In order to reasonable model a set of assembled modules as a single body, it is necessary for the docks to react with sufficiently high forces and moments once engaged. Because structural elements can range in mass, shape, and size, a standardized docking system must be designed to be scalable, both in size and in number. Autonomous mating typically incorporates hardware compliance for error stack-up, which is not amenable to rigid docking. One possible solution is to add force/torque sensors to both ends of every limb so that active force control can be used to correct misaligned docks. The alternative is to utilize a two-part docking system, with a "soft dock" and a "hard dock" state. The soft dock would account for system misalignments, and the hard dock would increase the rigidity of the soft dock.

Electrical Architecture

The Schunk PowerCube actuators communicate over a Controller Area Network (CAN) Bus and draw power from common power and ground lines connecting all of the actuators. A 120-Ohm resistor on either end of the limb between the CAN-Hi and CAN-Lo wires satisfies the CAN protocol, which requires a net 60-Ohm resistance between CAN-Hi and CAN-Lo. A simplified electrical diagram showing the limb attached to two modules is shown in Figure 5.

This bus architecture conveniently allows the limb to be connected to power and data at either end, or at both ends simultaneously. This supports Limbi's end-over-end mobility. Initial concerns about small voltage differences on

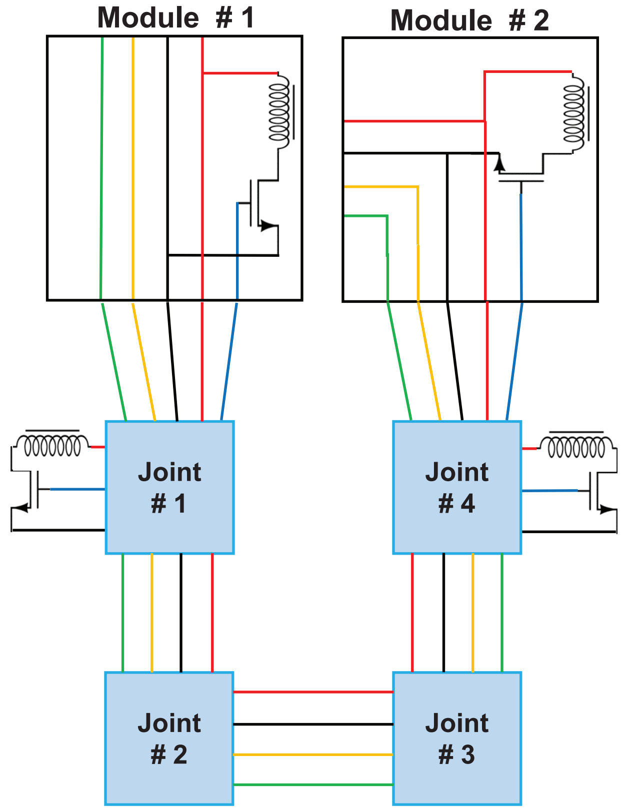

Figure 5: A simplified electrical diagram of the limb, as attached to two modules. The color-coding is the same as in Figure 3.

the power and ground lines, and about CAN timing differences when connected to both ends simultaneously, were alleviated through simple validation testing. If a future version of Limbi is longer by an order of magnitude, the timing differences may begin to cause communication problems and a different bus architecture will need to be used.

The electromagnets on the ends of the limb itself are activated directly using digital output pins on the first and last PowerCubes. A HI signal opens a 60V N-type MOSFET allowing approximately 0.25 amps at 24V to power the electromagnet. Toggling the electromagnets inside the modules, for module-to-module docking, is slightly more complex because the modules themselves do not communicate over CAN (however, they can draw from the power supply). As a result, the limb must be able to turn on the electromagnet inside of a module, leave that electromagnet on even when the limb isn't directly connected to the module, and disable it again if necessary for disassembly. This is accomplished by using an ATTINY 8-pin microcontroller on a custom PCB inside the module to listen for a short HI pulse on the "signal" pin of the module. When the pulse is received, the ATTINY toggles the state of the electromagnet, again through a 60V N-type MOSFET. To turn the electromagnet inside a module on or off, the limb uses a second digital output pin on the first and last PowerCubes to send this short HI pulse to the dock's signal pin while connected to the module.

Control Software

At this time, the sequences used to demonstrate the limb are pre-programmed commands for joint-space moves and Cartesian moves. To perform a joint-space move, the software generates a trapezoidal velocity trajectory for each joint and scales the trajectories so that all joints reach their targets simultaneously. To move the end-effector in Cartesian space, a set of waypoints are recursively generated between the current position and the desired position. Since the limb has a redundant degree of freedom, joint angles for each waypoint are calculated with a simple inverse kinematics solver, with the constraint that the second and third joints maintain the same angle. Further work must be done to better utilize the redundant degree-of-freedom for efficient trajectories or obstacle avoidance.

The main control loop runs at 200 Hz, obtains the desired joint position and velocity from the trajectory, applies proportional feedback to the velocity, and sends the command over the CAN Bus to each of the PowerCube actuators. The PowerCubes have an internal controller to interpolate between sequential waypoints at a higher rate.

Online Visualization and Offline Simulation

A visualization and simulation tool that was developed for other projects was used as part of the Limbi prototype software. The tool allows commands and command sequences to be simulated and visualized in an offline mode before sending them to hardware. Additionally, it shows the status of the hardware when running in online mode. A screenshot of the environment is shown below – the green limb is the current simulated/real position, and the gray limb is a target position. Moving forward, the environment will need to be improved to support visualization of the structure and other objects as they are manipulated by the limb.

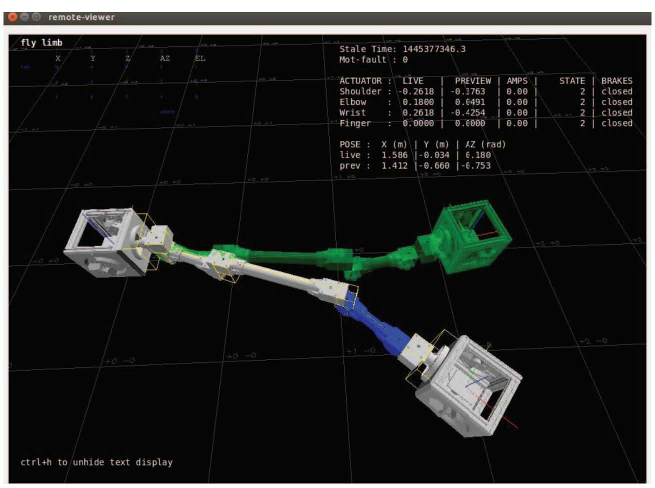

Figure 6: The visualization and simulation environment.

4. Technology Demonstration

The key capabilities of modular assembly and untethered end-over-end mobility were demonstrated by building a structure from three modular elements. The limb started attached to a single module, which was fixed in place. It then grabbed a second module and attached it to the first module, then moved itself so that the limb was connected to the second module instead of the first. Next, it repeated this process by connecting a third module to the second module. The demonstrated sequence of events is shown below.

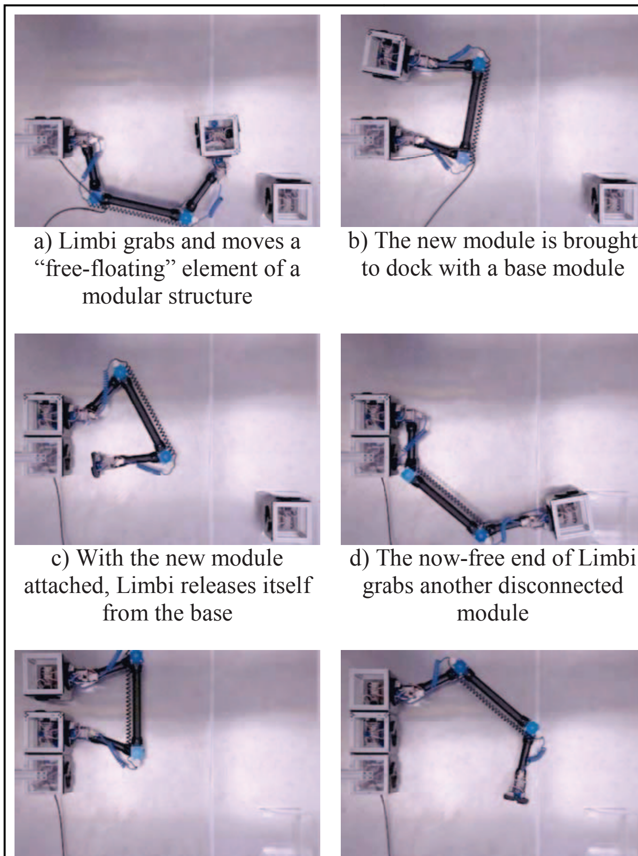

a) Limbi grabs and moves a "free-floating" element of a modular structure

b) The new module is brought to dock with a base module

c) With the new module attached, Limbi releases itself from the base

d) The now-free end of Limbi grabs another disconnected module

e) Limbi docks this module to the growing structure

f) Limbi can move between docks on the assembled structure as needed.

Figure 7: A demonstration of the Limbi prototype assembling three modules.

5. Application Areas

Spacecraft Assembly

Spacecraft assembly is the original use case which motivated Limbi. For this application, Limbi would be launched along with a shipment of modular spacecraft components, and would assemble these components as shown in the above demonstration. Since commencing the project, we have also recognized a wide variety of additional areas where Limbi could enable new breakthroughs.

3D Printed Assembly

Many new space structures can be enabled by 3D-printing parts in space, so that the parts do not need to withstand high launch forces. Once these parts are printed, they must be assembled, which is an application for which Limbi would be well-suited as a lightweight mobile robot.

Reconfigurable Rovers

JPL specializes in robotic rovers for planetary exploration, and the Limbi concept would be particularly well-suited for a reconfigurable rover with a mobile arm. A set of docking points on the chassis would allow the arm to be positioned in a variety of positions and orientations, which could significantly extend the rover's reachable workspace. For example, Limbi could be mounted on the front of the rover for surface sampling tasks, and then reposition itself to the top surface to distribute samples between difference instruments and tools. The limb could therefore be smaller and lighter than a fixed-base limb required to reach the entire workspace.

Distributed Sensing Network

High mission value could be gained by using Limbi robots to redistribute sensors on a spacecraft, or by using sensors already built into Limbi such as cameras to monitor spacecraft health. For example, a small Limbi robot could crawl around the outside of a re-entry vehicle to inspect ceramic tiles for damage, without requiring a much larger fixed-position arm.

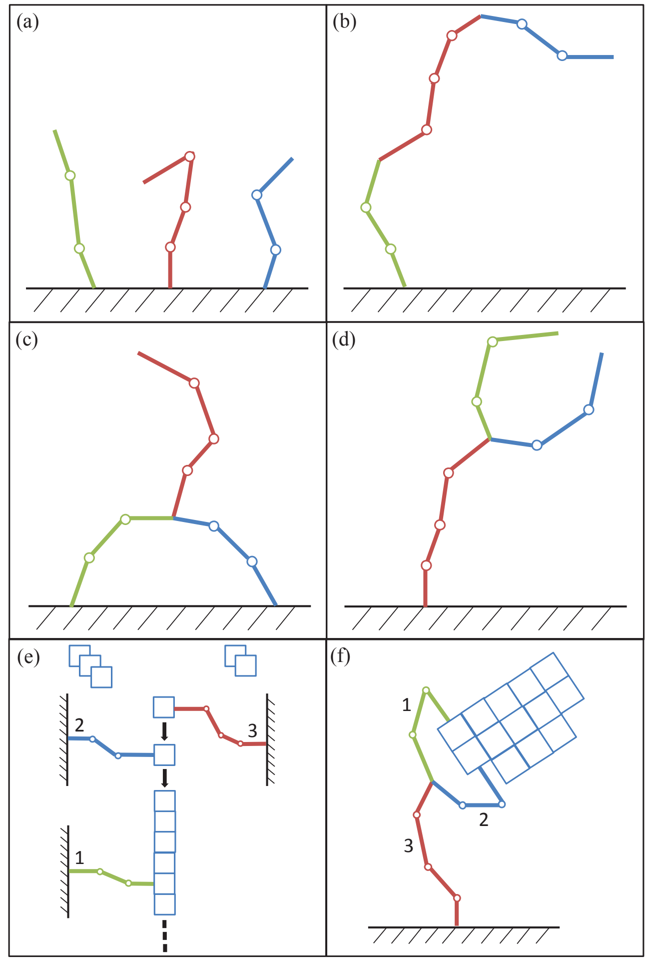

Figure 6: Limboid configurations and example uses.

6. LIMBOID CONCEPT

Limboids, which we define as a self-mobile robotic system that is fully reconfigurable at the limb level, are the next logical extension of Limbi. This system would not require passive elements (e.g. chassis, hubs, and scaffolding) and would be capable of reconfiguring itself, and thus it can adapt to changes in workspace requirements to remain effective at every scale without redundant components. In a typical use-case for Limbi, the singular limb is designed to be the optimal solution for assembling and maintaining a single spacecraft of preordained size and geometry. While this solution applies to a large subset of current in-space assembly and servicing tasks, there is still a need for a lightweight reconfigurable robotic system that can handle more general tasks on a variety of spacecraft and platforms. Limboids fill this necessity, as they are capable of both gross and dexterous manipulation at arbitrary scales and can thus be embedded in a factory setting with the assurance that they will continue to meet these demands over numerous spacecraft design cycles. Limboids can be decomposed into four primitive configurations, each with dedicated functionality. As is illustrated in the top four panels of Figure 8, these configurations consist of (a) multiple Limbi-type robots working in parallel to perform independent or cooperative operations, (b) limbs combined into an arbitrarily long serial chain to adapt to a changing workspace, (c) a walker-manipulator that can manipulate objects while moving, and (d) a torso with multiple arms for gross manipulation of large objects. Figures 8e and 8f show two examples of how a Limboid system would enable the same set of three independent limbs to perform two different tasks that would otherwise require multiple single-purpose robots. In first example, the primitive configuration illustrated in Figure 8a is used to assemble an indefinitely long chain of modules with Limb 1 holding and repositioning the chain while Limbs 2 and 3 grab and attach more modules. In the second example, the Limboid system reconfigures itself with Limbs 1 and 2 attaching themselves to the end of Limb 3, and proceeds to perform gross manipulation of a large object. The results of this paper have demonstrated two features of a Limboid system necessary to facilitate these processes. First, limbs can reliably navigate to docks in known locations, and second, the limb can draw power and communicate through the structure to which it attaches itself, thereby eliminating the need for an external tether anywhere in the system. Future work will expand these features to include a docking architecture that allows an indefinite number of limbs to attach at a single node and flexible control algorithms that can adapt to the ever-changing kinematic models of Limboids.

Example use case: In-space factory

Due to the high degree of re-configurability that is unmatched by existing robots, Limboids would be ideal for staffing in-space factories for spacecraft assembly, asteroid processing, or satellite maintenance.

Automated factories on earth, such as those used for automobile manufacturing, typically rely on large workforces of robotic arms that are fixed to the floor or mobile on rails. For most space applications, it would be infeasible to launch such a large robotic workforce, and it may not be practical to maintain a human workforce to reconfigure the robots for different tasks. It will be preferable to use a smaller number of robots that can reconfigure and move themselves around the factory, and a Limboid workforce with access to a tool crib would meet this requirement for many types of factories. A small set of Limbi robots could perform each of the following tasks:

- Unpacking and assembling a modular "factory floor" with a set of docking points where Limboids could attach.
- As a long serial kinematic chain, reaching out from the factory floor to dock with an incoming shipment.
- Unpacking spacecraft components or raw materials from incoming packages and delivering those pieces to different workstations.
- Assembling many varied modular components to build an entire working spacecraft or satellite.
- Wielding tools like grippers or drills to perform small-scale dexterous manipulation.
- Safely moving a large spacecraft around the factory by breaking into separate Limbi modules and grabbing the part in many locations.
- As a large multi-fingered gripper (where each finger is a single Limbi), Limboids could grab or transport boulders or chunks of asteroids.

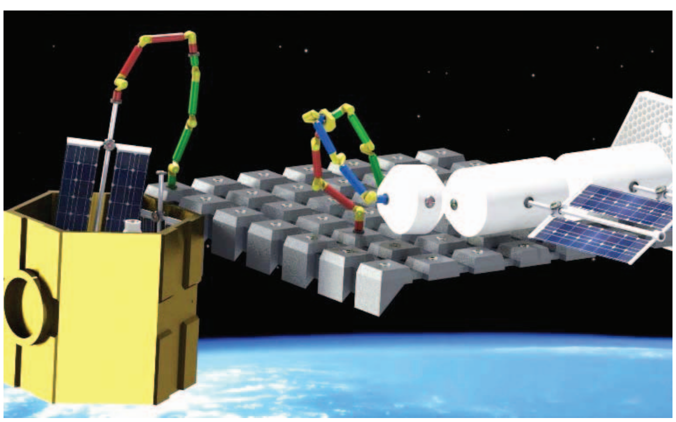

Figure 7: A rendering of Limboids being used to unpack a shipment of components on an in-space factory and build a small satellite. Separable limbs are shown in different colors.

7. FUTURE WORK

Many of the qualifications that are part of the Limbi architecture were left out of the prototype and demonstration due to time and budget constraints. In particular, these important qualifications include high-level autonomy and highly-rigid electromechanical docks. Other elements of future work listed here are not essential components of the Limbi architecture as we have defined it, but will be useful for practical implementations.

Autonomy

A key component of the proposed Limbi architecture is high-level autonomy, but it was not developed for the testbed demonstration. Current operations with existing limbs in space, such as Canadarm2 or those on JPL's Mars rovers, require extensive human-in-the-loop planning and verification. This is feasible with small tasks to perform (such as moving a single structural element, or drilling a hole), but would quickly become a bottleneck for assembling complex structures requiring many operations and arm moves. An autonomous Limbi system should require minimal high-level specifications from ground controllers (for example, the geometry of the final structure and the initial locations of all parts), and then autonomously generate, simulate, and follow a detailed plan. While following the generated plan, the system must be fault-tolerant to the extent that it can detect an anomaly, attempt to automatically correct the condition (for example, a misaligned dock), and notify ground controllers otherwise. This becomes particularly critical for Limboid systems with multiple limbs working on different tasks, so that one limb may detect a fault and wait for further instruction, while other limbs carry on with their assigned plans. Developing this autonomy will be a major focus of Limbi and Limboids moving forward.

Docking Mechanism Refinement

The electromechanical docks in our testbed are powered by 80-lb electromagnets, which are suitable for prototyping purposes but would be woefully inadequate for flight applications. We have found that they lack the desired rigidity and strength, and require very precise alignment before they will mate. Additionally, they draw continuous power. The next generation of docks should offer greater rigidity and strength, while simultaneously being robust against slight misalignment.

Tools and Tool Cribs

The capabilities of Limbi, and particularly the proposed Limboid system, can be enhanced further by allowing the robots to pick up compatible tooling (such as grippers, drills, or wrenches) from a tool crib or depot and drop the tools back off at the same location. The same docking interface that is used between Limbi and structural components could be used between Limbi and these tools.

Minimum Power Path Planning

Power is a precious resource for space operations and must be conserved where possible. Closely related to the need for autonomy, a deployed Limbi system should be able to generate power-saving plans at both a high level (planning the optimal sequence of moves to assemble or walk around a structure) and a low level (searching for the shortest paths through joint-space to move from one position to the next).

Integrated Vision System

Cameras can be added to either or both ends of Limbi, and could be used for inspection of assembled components as well as navigation and control (i.e. ensuring precise alignment with visual targets on a dock).

Wireless Communication

The current testbed limb uses a CAN Bus for communication, which has performed well but requires two additional wires running through the entire structure and limb, as well as additional pins on the electromechanical docks. Other wired communication protocols, such as EtherCAT, require even more data lines. It would be preferable for the limb(s) building a structure to communicate wirelessly with a single base computer. Wireless power transmission could be explored, although it would likely come at a high cost to power efficiency.

8. SUMMARY

Limbi builds on a successful architecture first proven in space by Canadarm2. The testbed prototype, working in a planar space, has already been shown to be capable of autonomous pre-determined tasks. The hardware will be invaluable for the team's future research work focused on integrating computer vision and higher levels of autonomy to enable rapid unaided in-space assembly. Additionally, the same hardware design will be used to build and demonstrate a first-of-a-kind Limboid system. Simultaneously, the team will use lessons learned from assembling and working with the hardware to build a refined prototype with cameras, tools, more robust docking mechanisms, and wireless communication. The long-term project goal is to prepare Limbi and Limboids for flight projects and eventual planetary applications.

ACKNOWLEDGEMENTS

Funding to develop the Limbi prototype was generously provided by JPL's National Space Technology Applications (NSTA) program office, with advocacy from Edward Chow and Joseph Sauvageau.

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Rudra Mukherjee of JPL's Robotics section helped bring together the Limbi team and pushed us to successfully develop the prototype with a small budget and limited time.

REFERENCES

[1] M. Goeller, J. Oberlaender, K. Uhl, A. Roennau, and R. Dillman, "Modular Robots for On-Orbit Satellite Servicing," in Proc. IEEE Int. Conf. Robotics and Biomimetics. Dec. 2012.

[2] J. Oberlaender et al, "Management and Manipulation of Modular and Reconfigurable Satellites", in Proc. Robotik, 2012.

[3] M. Hancher and G. Hornby, "A Modular Robotic System with Applications to Space Exploration," in Proc. IEEE Int. Conf. Space Mission Challenges for Information Technology, 2006.

[4] P. Staritz, S. Skaff, C. Urmson, W. Whittaker, "Skyworker: A Robot for Assembly, Inspection and Maintenance of Large Scale Orbital Facilities." IEEE International Conference on Robotics and Automation, June 2001.

[5] David Akin, Brian Roberts, Stephen Roderick, Walter Smith, and Jean-Marc Henriette, "MORPHbots: Lightweight Modular Self-Reconfigurable Robotics for Space Assembly, Inspection, and Servicing" AIAA-2006-7408, AIAA Space 2006 Conference and Exhibit, San Jose, California, Sep. 19-21, 2006

BIOGRAPHY

Sawyer Brooks *joined JPL as a Robotics Software Engineer in 2015. Prior to that, he studied Robotics (M.S.E) and Mechanical Engineering (B.S.E) at the University of Pennsylvania. As a member of the Manipulation and Sampling group, his work involves designing new types of robotics limbs and software support on robotics arms for Mars missions. His research interests include bringing greater autonomy to robotics in aerospace, robots for planetary habitat assembly, and multi-robot cooperation.*

Peter Godart *received Bachelor of Science Degrees in Mechanical Engineering and Electrical Engineering from MIT in 2015. Currently Peter is a technologist in the Robotic Manipulation and Sampling Group within JPL's Mobility and Robotic Systems Section. There his research focus involves the development of adaptive control and path-planning algorithms for reconfigurable robotic systems.*

Paul Backes *is group supervisor of the Manipulation and Sampling group in the Mobility and Robotic Systems Section at Jet Propulsion Laboratory, Pasadena, CA, where he has been since 1987. His current activities focus on planetary manipulation, sample acquisition, and sample handling. He received the BSME degree from U.C. Berkeley in 1982, MSME in 1984 and Ph.D. in 1987 in Mechanical Engineering from Purdue University. Dr. Backes received the 1993 NASA Exceptional Engineering Achievement Medal for his contributions to space telerobotics, 1993 Space Station Award of Merit, Best Paper Award at the 1994 World Automation Congress, 1995 JPL Technology and Applications Program Exceptional Service Award, 1998 JPL Award for Excellence, 1998 NASA Software of the Year Award Sole Runner-up, and 2004 NASA Software of the Year Award. He has served as an Associate Editor of the IEEE Robotics and Automation Society Magazine.*

Brendan Chamberlain-Simon *began working as a Technologist at JPL in June of 2015 after graduating from Columbia University with a BS in Mechanical Engineering. Brendan's work involves synthesizing theoretical and empirical models for high-fidelity simulation of robotic hardware. He also works on novel robotic systems for in-space assembly, as well as mobility and manipulation on planetary bodies. Prior to his work at JPL, Brendan worked at Swamp Works at NASA's Kennedy Space Center and in Professor Peter Allen's Robotics Laboratory at Columbia University.*

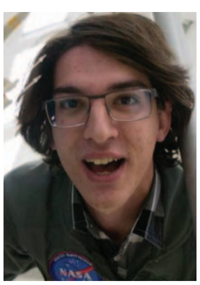

Russell Smith *received a B.S. in Mechanical Engineering from the California Institute of Technology in 2014. He has been employed as a Roboticist at JPL since 2014. His work has focused on electrostatic dry adhesives for vertical surface mobility platforms, and electromechanical systems for in-space assembly.*

Sisir Karumanchi *is a Robotics Technologist at the Manipulation and Sampling Group under the Mobility and Robotic Systems Section at Jet Propulsion Lab, NASA. He received a Bachelors degree in Mechatronic Engineering from the University of Sydney in 2005, and completed his Ph.D. with the Australian Centre for Field Robotics at the University of Sydney in 2010. Between 2011-2014, he was a postdoc at the Massachusetts Institute of Technology in the Robotic Mobility Group. His research interests include mobile manipulation, semi-autonomous control, guaranteed motion safety, vehicle-terrain interaction, and Bayesian nonparametric inference.*