Towards Force-aware Robot Collectives for On-site Construction

Nathan Melenbrink
Wyss Institute,
Institute for Computational Design and Construction

Paul Kassabian
Simpson Gumpertz & Heger

Achim Menges
Institute for Computational Design and Construction

Justin Werfel
Wyss Institute

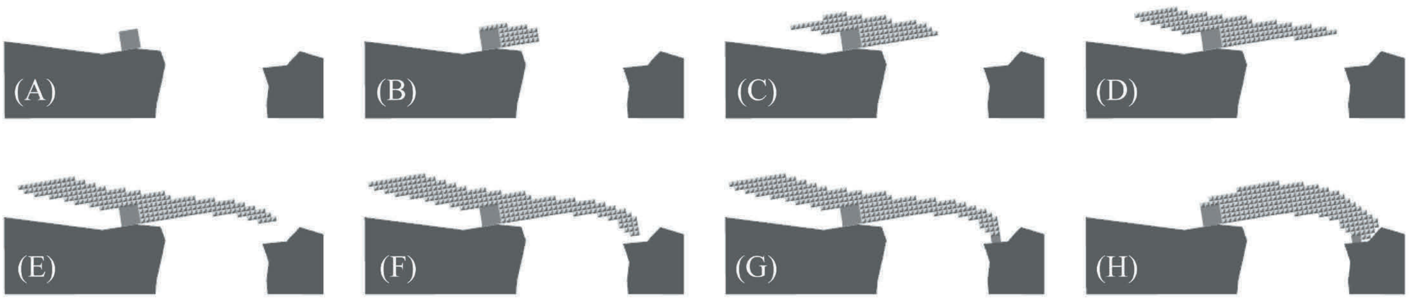

1

ABSTRACT

Due to the irregular and variable environments in which most construction projects take place, the topic of on-site automation has previously been largely neglected in favor of off-site prefabrication. While prefabrication has certain obvious economic and schedule benefits, a number of potential applications would benefit from a fully autonomous robotic construction system capable of building without human supervision or intervention; for example, building in remote environments, or building structures whose form changes over time. Previous work using a swarm approach to robotic assembly generally neglected to consider forces acting on the structure, which is necessary to guarantee against failure during construction. In this paper we report on key findings for how distributed climbing robots can use local force measurements to assess aspects of global structural state. We then chart out a broader trajectory for the affordances of distributed on-site construction in the built environment and position our contributions within this research agenda. The principles explored in simulation are demonstrated in hardware, including solutions for force-sensing as well as a climbing robot.

1 An example sequence showing how a swarm of construction robots might use forces to maintain stability in a building sequence that uses generic square blocks. In (A) a cache has been deposited on one side of a chasm. (B) Agents retrieve blocks from the cache and begin building a cantilever. (C-D) Agents read force sensors at the cache that indicate a danger of toppling into the chasm and begin counterbalancing. (E) Agents determine that the cantilever is near ground support on the far side and (F) start building downward. (G) Ground support is established on the far side (e.g., by inserting a special expanding unit to take up the remaining space between the cantilever and the ground). The force sensors now indicate that there is no longer danger of toppling, so agents accordingly begin disassembling the now-redundant counterbalancing blocks and (H) use them to strengthen the bridge. The work presented in this paper focuses on steps (B–E); the rest will be the subject of future work.

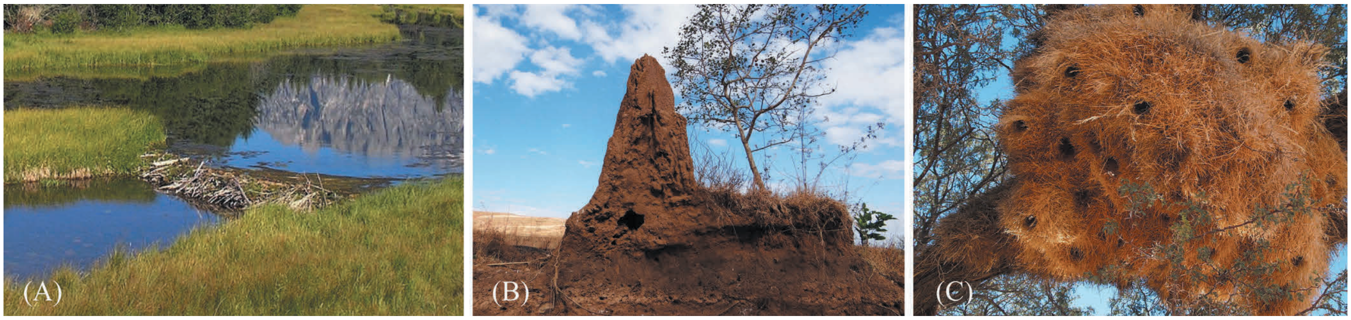

2 Examples of structures built by nature's collective builders: (A) beaver dam, (B) termite mound and (C) social weaver bird nest.

INTRODUCTION

Social weaver birds, beavers and termites are examples of nature's most accomplished cooperative builders, achieving large and resilient structures through parallel execution of simple tasks. Advances in electronics, sensing technologies, manufacturing methods and agent-based computation have only recently validated the feasibility of emulating such emergent building techniques through robotics. This affords us the opportunity to advance automation in construction in order to bring substantial advantages such as reducing cost and time on site, reducing risk related to unknown conditions, and reducing the risk of danger or injury. If full autonomy is attained, it can enable building in dangerous, remote, or otherwise challenging settings where construction is not currently feasible. When applied to architecture, it offers the possibility of conceptualizing the building process not as a binary distinction between construction and completion, but rather as an ongoing, persistently shifting response to variable high-level functional requirements.

On-Site Automation

Most evidence of robotic construction in the built environment comes in the form of off-site prefabrication with stationary robotic arms, which do not easily lend themselves to construction sites. The conventional understanding of robotics in the AEC industry tends not to consider other classes of robots that would be better suited to construction sites; for example, the Rob|Arch conference is explicitly focused on industrial arms (Rob|Arch 2014). On-site construction is typically changeable and prone to found conditions, and therefore also not typically suited to preplanned routines. Though challenging, on-site automation could enable a broad new range of building practices. The use of aerial robots in construction is starting to receive more attention; UAVs are currently being used on site for surveying, and are in development for light construction tasks (Augugliaro 2014). However, the industry tends to overlook other classes of robots, such as cable-driven robots or climbing robots, which could be more impactful than aerial robots for contemporary building practices (Sousa 2016).

Swarm Robotics

The AEC industry, being typically risk-averse, has little precedent for innovative solutions for large-scale on-site automation, so we turn to biology for inspiration (Figure 2). The two features discussed above, responsiveness to environmental conditions and agent mobility, are well-addressed by the swarm approach. Beyond merely automating existing construction practices, looking at the ways that collectives build in nature suggests a new process-model for building, inviting new control methods, construction methods, and construction machinery to challenge existing models.

While a number of classes of robots might be conducive to on-site automation, climbing robots allow for a significant expansion in the viable building height over ground-based robots, and allow for greater local accuracy and more economical power consumption than aerial or cable-driven robots (Sousa 2016). We select reactive control over predetermined building sequences, as the former is more robust to environmental perturbations. Decentralized control is more attractive than centralized control, because requiring robots to report to a centralized system would limit both the workspace and the number of robots the central control could support. Instead, a decentralized system allows for building to continue, resilient to failure of any individual robot. Robotic agents are provided with high-level user-specified requirements instead of precise blueprints; this compromises the user's design agency, but increases the likelihood that the high-level goal could be attained. Eventually, with increasing labor costs, increased safety concerns, and financial pressures, the AEC industry will almost certainly place more emphasis on distributed robots for on-site automation; one of the aims of this work is to push forward that required shift in thinking.

Irregular Environments

This research trajectory presents a theoretical framework for autonomous construction in general, though first applications may be specifically for disaster scenarios in particularly hostile environments (e.g., building a bridge across a chasm as described in Figure 1). We are interested in considering a broad range of

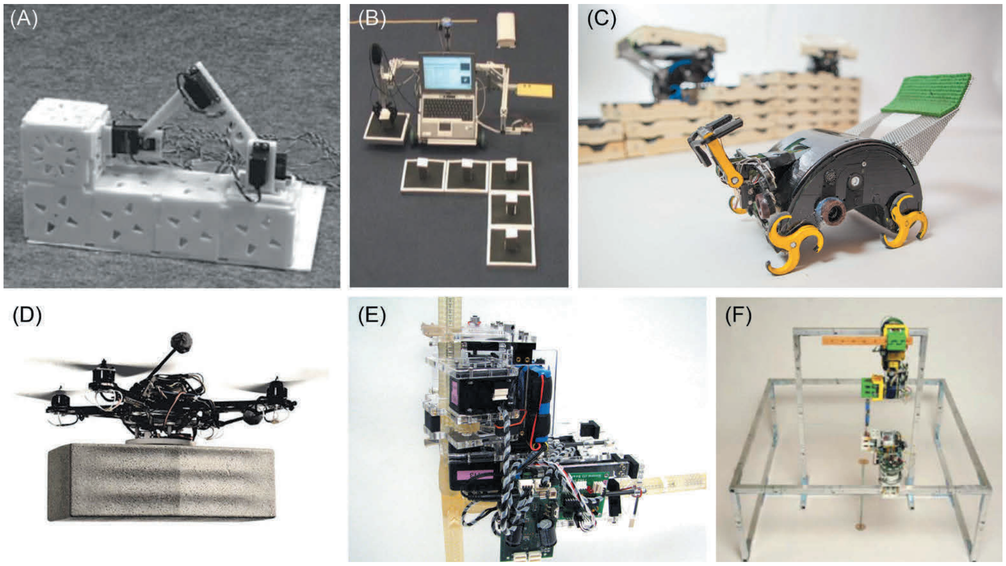

3

3 (A) AMAS robot and blocks (Terata and Murata 2004). (B) GER1LLA robot and blocks (Werfel et al. 2006). (C) TERMES robots and blocks (Werfel et al. 2014). (D) Quadrotor aircraft and block (Willman et al. 2012). (E) Truss-reconfiguring robot (Nigl et al. 2013). (F) Shady3D robots on struts (Yun and Rus 2008).

applications in the built environment, with a particular focus on the challenges of autonomous building in unstructured terrains. This requires the consideration of problems that largely have not been previously addressed, such as maintaining stability throughout an autonomous assembly sequence, which is the primary focus of this research. If the default options of scaffolding, formwork, or temporary bracing are not available, the structure itself must always be self-stable.

A fully autonomous building system would need to operate without supervision or intervention, and would need to pay attention to forces at every step of construction in order to ensure its resiliency to environmental hazards. Truss structures have been identified as an ideal entry point to on-site construction automation, because they can span long distances without requiring supplementary structures. However, if the desired structures are not trivially stable, their construction requires either anchoring or counterbalancing. Since we want to reduce the risks associated with the unknown on-site environment and the number of tasks that autonomous robots must manage, this work assumes there is no preparation of the environment before construction nor ability to anchor during construction, and instead seeks to maintain overall stability through counterbalancing. This affords the swarm the possibility of building sturdy structures that conform to challenging topographies as opposed to leveling them.

This conceptual shift towards unsupervised automation reopens exciting opportunities to rethink building life cycles or revisit, in a more serious way than previously possible, the feasibility of adapting to changing environments or functional requirements. Furthermore, it is increasingly critical to recognize the implications of impending automation, and, as architects and experts on the built environment, work with engineers from a fundamental level—or risk losing relevance in a world increasingly shaped by automation. In this paper we present an approach for maintaining stability during an autonomous building sequence generated on the fly, generalize it to other geometries and unstructured terrains, and show hardware that supports the theory.

BACKGROUND

Distributed Robotics

Significant achievements relative to this research agenda can be grouped into those pertaining to distributed robotics, responding to forces, and on-site robotics. Research in robotics has produced a variety of experimental hardware demonstrations (Figure 3), with multiple autonomous robots building three-dimensional structures, including climbing robots building with struts (Yun and Rus 2008; Nigl et al. 2013) or blocks (Terada and Murata 2004; Werfel et al. 2014), as well as flying robots building with struts (Lindsey et al. 2011) or blocks (Willmann et al. 2012). These research prototype systems are demonstrated in controlled laboratory environments, using highly specialized bespoke building materials. The challenges to refining any such systems to work reliably outside the lab, to build in natural environments using common materials and without relying on the regularities or tools available in their currently demonstrated test settings, are very substantial. Furthermore, none of these examples consider maintaining stability during construction.

Also of interest to this work is the notion of "digital materials," a discrete set of components that can connect in a finite set of ways (Popescu 2009). Digital materials facilitate automation

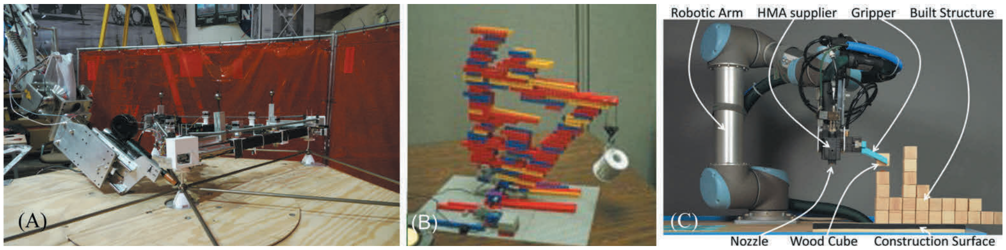

4

4 Examples of construction processes that take forces into account: (A) Intelligent Precision Jigging Robot (McEvoy et al. 2014), (B) Physical verification of computationally evolved form (Pollack et al. 1999), and (C) a robot arm builds cantilevers according to a preplanned building sequence (Brodbeck and Iida 2014).

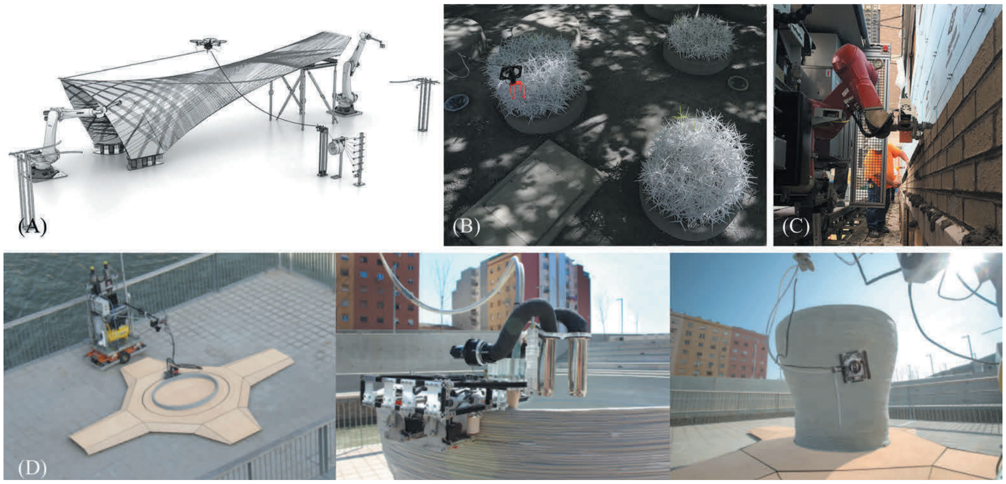

5

5 Examples of construction-site-ready robots: (A) Aerial robots used in concert with stationary arms to expand the workspace (ICD/ITKE Research Pavilion 2017), (B) Aggregate Architecture Pavilion uses a cable robot to deliver materials on site (Dierichs and Menges 2012), (C) SAM bricklaying robot (Peters and Belden 2014), (D) Minibuilders, a heterogeneous team of additive manufacturing robots (Nan 2015).

by enforcing precision while increasing structural performance, when compared to analog assemblies. The construction processes in Figure 3 all qualify as digital materials with the exception of (D), in which the geometry does not encode a finite set of possible connections.

Responding to Forces

Recent years have witnessed a considerable amount of literature on experimental modal analysis (comparing on-site impact frequency response measurements against finite element models), but it is worth noting that these techniques cannot be applied to the domain considered in this research track because they require a priori knowledge of the topology. A smaller number of studies consider responding to forces during a building sequence.

An approach shown in Figure 4A considers forces in sequence, though still relying on centralized pre-planning (McEvoy et al. 2014). Previous work describes counterbalancing features that emerge from an evolutionary algorithm that seeks to satisfy a high-level goal of spanning a gap (Pollack et al. 1999). While the authors demonstrate the stability of the resulting final configurations, they do not consider stability at each step of the building sequence or the feasibility of building each candidate structure. Researchers have extended the principles established by Pollack et al. for the purpose of a single robot arm building a cantilever (Brodbeck and Iida 2014). They consider forces to guide an evolutionary algorithm for material distribution. They discuss the use of counterbalancing for unanchored structures and propose a stability criterion of checking whether the horizontal component of the center of mass is past the edge of the building platform. This is effective in their defined case but would not generalize to an irregular environment.

Otherwise there is little research on counterbalancing beyond conventional building practices where counterbalancing is formulaic and predetermined. Self-organized construction is a relatively new field that uniquely demands this research focus.

Robotics in the Built Environment

Though challenging, there are precedents for construction-site-ready robots (see Figure 5). A number of these examples do require the environment to be prepared ahead of time, but still present interesting solutions to circumventing the workspace limitations of stationary robots (Keating et al. 2017; Nan 2015; ICD/ITKE 2017; Dierichs and Menges 2012).

SIMULATIONS

A number of different approaches were executed in simulation with the intent of evaluating our hypothesis that paying attention to local forces will allow agents to autonomously build stable structures. Simulation work was first developed with finite

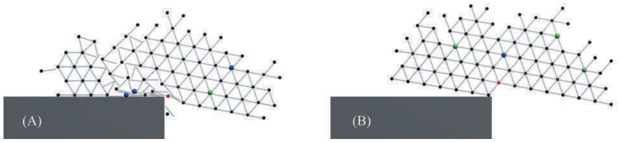

6

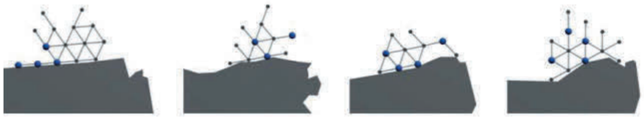

7

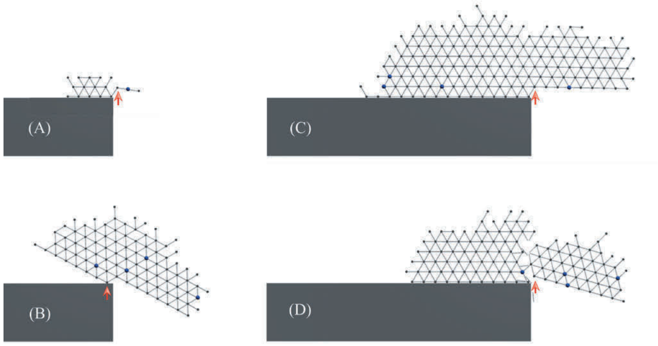

8

6 Typical failure modes of (A) breaking and (B) toppling.

7 Snapshots of trials with procedurally generated irregular terrains. The blue spheres indicate the presence of an agent at a node.

8 Typical snapshots at the moment of failure for structures under (A) the "force-unaware" variant, (B) the "force-aware" variant (no counterbalancing), (C) the "preplanned balancing" variant (force-aware), and (D) the "reactive balancing" variant (force-aware). All failures included broken joints; in (A) joints typically broke early in the trial, in (B) joints tended to break while toppling, in (C) a joint was broken but the strut did not fall away, while in (D) the joint break happened to cause the strut to fall away, leading to a cascading failure.

element models (Melenbrink et al. 2017) and further developed with a dynamic physics engine (Melenbrink and Werfel 2017). The objective of all simulations was to build a cantilever as far as possible, as if to bridge a gap, using decentralized robotic agents that respond to locally measured forces as cues. A previous consideration had been looking at vibrations to determine stability, but this technique was found to be noisy and unreliable.

For reasons described in the introduction, we focus on truss structures. Horizontal cantilevers are an apt framework for exploring stability, as they are more challenging to keep stable than fully supported structures, and would be the first step in unsupervised bridge construction, which could be particularly useful in remote or difficult terrain. Typically, bridges in such unstructured environments require scaffolding or other supplementary structural support, but truss bridges could conceivably be built without the need for additional scaffolding. For structural stability, we choose the truss geometry to be a triangular lattice, with horizontal rows. This is broadly an efficient use of material layout with both bending and shear capacity. In this work, as a first step toward developing a theoretical foundation, we considered only building in a two-dimensional vertical plane. The truss (as shown in Figure 6) is composed of individual identical rigid struts and corresponding nodes as opposed to stable cellular units (see Figure 15), though we expect that the findings from the former will generalize to the latter (see the following section on Geometry and Sensing Extension). For physical plausibility, material assumptions were made based on existing, readily available construction materials. Nodes are assumed to be capable of receiving up to six struts and measuring axial and shear forces, as was previously demonstrated (Melenbrink and Werfel 2017).

In the system we propose, agents must be able to detect and prevent structural failures. In these simulations, we consider two types of failure: toppling into the chasm and local breaking. Breaking occurs when forces at a joint between a node and strut exceed predetermined maximum thresholds. Struts are considered rigid bodies that do not buckle, as the system was designed such that this would not be a driving failure mode. Ground support under the structure is assumed to be stable. We seek to use counterbalancing (as opposed to anchoring) to eliminate the otherwise inevitable failure mode of toppling into the gap, and also use local force measurements to prevent or delay the failure mode of breaking.

In these simulations, robots do not communicate directly with each other but indirectly, inspired by the biological phenomenon of stigmergy. The principle of stigmergy is that cues left in the environment by an agent influence subsequent actions of agents that encounter that cue. In our application, robots' actions are affected by force readings they measure through the nodes of the structure, and their actions in turn affect forces throughout the structure. This approach removes the need for direct communication between robots, which therefore eliminates requirements for robots to stay in proximity to maintain connectivity, and avoids the challenges of mobile ad-hoc wireless networks in chaotic environments.

Simulation Environments

We developed two distinct simulation environments to study different aspects of system behavior. The first approximated rigid bodies with linear elements and used finite element analysis (FEA) for structural calculations (Melenbrink et al. 2017). Since FEA returns only overall static solutions, the calculations had to be reset and solved independently for every frame. This method provided significant detail at the element level. However, in order to accurately capture dynamic behavior such as friction against

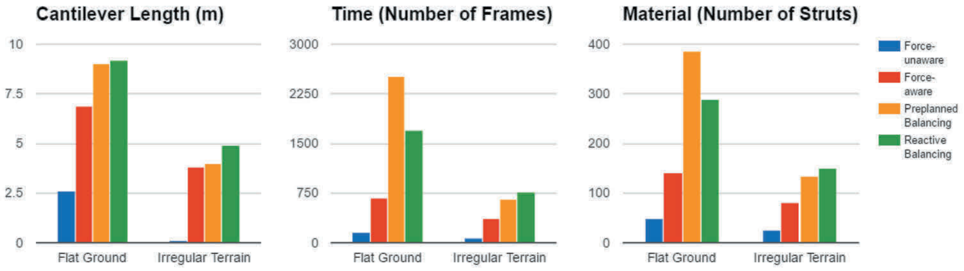

9

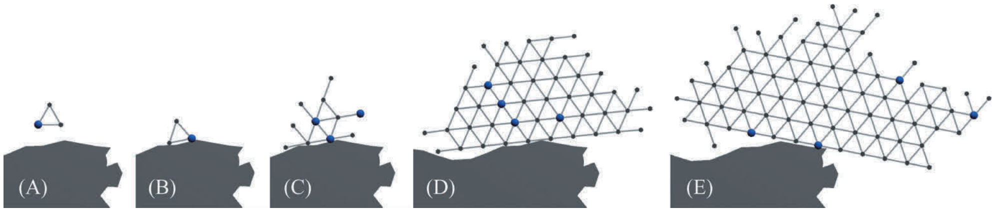

10

9 Graphs of results. There is consistently less building activity in irregular terrains than flat ground. Both "balancing" variants are able to build longer cantilevers, though at a considerable increase in time and material consumption.

10 Typical sequence for the "force-aware" (no balancing) variant in an unstructured terrain. The material distribution causes the structure to tilt, exposing newly viable positions to install struts.

the ground or collisions between elements, higher-level FEA would be required (which would be prohibitively computationally expensive for large-scale structures with changing topologies).

Accordingly, a second simulation environment was developed with Unity3D, which uses the NVIDIA PhysX 3.3 engine for rigid-body physics simulations. This simulation environment was better suited to capturing dynamic effects and interactions with an irregular terrain. Rigid-body dynamics is a branch of classical mechanics that focuses on systems of connected bodies and how they react to forces. Discrete bodies are assumed not to deform under applied loads, which simplifies calculations (thereby drastically reducing computational expense, as opposed to FEA or linear elastic models), and is more conducive to representing the discrete, rigid building elements presented here.

Rather than measuring local forces by querying the elastic deformation of node elements as needed with an FEA approach, the PhysX solver allows for direct querying of reaction forces between any two connected bodies. Irregular terrains are generated within the simulation environment for the trials that consider it. A new randomized terrain is generated with each trial, including features at various scales (Figure 7). The terrains are assigned physics properties including a collider and coefficient of friction, which the PhysX engine uses to calculate interactions between rigid bodies.

Baseline Algorithm

We first present a baseline algorithm for agent-based construction, which attempts to extend a cantilever without any regard for structural forces; it results in two distinct failure modes of breaking or toppling (see Figure 6). We then look at these two failure modes one at a time and propose algorithmic solutions to mitigate each. We propose that breaking can be forestalled by paying attention to local forces at every node encountered, and toppling can be forestalled by paying attention to the forces at the origin node. More detailed explanations of the algorithm and its variants are explained in prior work (Melenbrink and Werfel 2017).

Force Awareness

The baseline, "force-unaware" algorithm will eventually topple or break, though it tends to break before accruing enough mass to topple (see Figure 8A). Under the "force-aware" variant, agents disqualify any locomotion where the forces measured indicate that a potential structural failure might occur if the robot were to move down that strut. We ran repeated trials for both force-aware and force-unaware variants, for both anchored and resting conditions. The results of Figure 9 indicate that the proposed method of forestalling breakage was indeed successful—the length of cantilever achieved by "force-aware" agents was over twice that of the "force-unaware" agents.

Dynamic Counterbalancing

To prevent the structure from falling into the chasm when unanchored, we modify the algorithm so that agents build in a way that counterbalances the cantilever. They do this by adding material in other directions, using the weight distribution of the structure to provide stability. Two additional variants are tested, both with local force-awareness; *preplanned counterbalancing*, in which agents act to keep the structure balanced at all times, and *reactive balancing*, in which agents only add counterbalancing material when the structure is approaching overall instability.

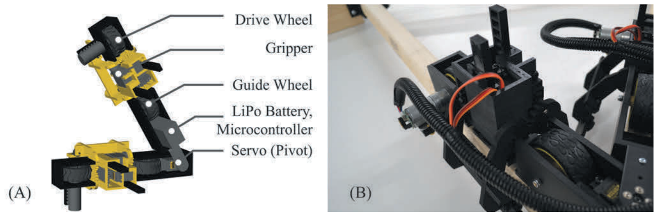

11

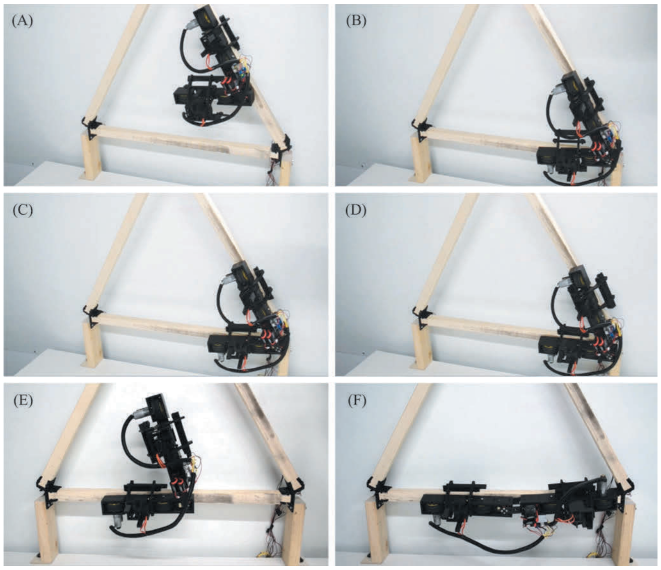

12

11 A rendering of the strut-climbing robot, consisting of a front and rear carriage, each independently capable of gripping and rolling along a strut. Not shown is a preliminary design for a strut-carrying module, which would be needed to achieve strut placement. (B) Photograph of initial prototype.

12 The locomotion sequence as the robot transitions from one strut to another. (A) The robot uses its rear carriage to traverse down the strut with its front carriage gripper open. (B) The robot has moved until it makes a hard stop at the next node. The motor encoders cease incrementing, indicating to the microprocessor to power off the motor. (C) The front carriage grips the next strut, (D) the rear carriage releases from the previous strut, (E) (taken from a different trial) the robot traverses down the new strut using its front carriage, (F) the rear carriage pivots into place and attaches to the strut. Next the front carriage will detach and pivot, and the sequence can continue. Note that the nodes are simplified to 2 vacancies as opposed to a complete 6, as a single triangle was deemed sufficient for early locomotion trials.

For the latter variant, we found that looking at shear forces can provide cues for when to begin counterbalancing, but do not reliably capture when to cease or re-initiate counterbalancing routines (Melenbrink and Werfel 2017). A more useful heuristic was found by measuring the ratio of the axial reaction forces on struts at the origin node. If this value exceeds a certain quantity, agents inherit a likelihood that they will add counterbalancing material. Otherwise, agents maintain their default bias in the direction of the cantilever.

The “no balancing” variant will predictably topple into the chasm when the cantilever reaches a critical length that is shorter

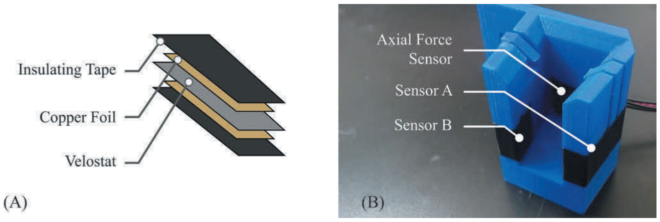

13

13 (A) The layers of the sensor package. (B) A prototype socket used to evaluate the effectiveness of this sensor in a node-and-strut assembly, as described in previous work (Melenbrink et al. 2017).

than the two other approaches. Both counterbalancing variants allow for further cantilevering by eliminating the failure mode of toppling; reactive counterbalancing performs best in terms of building furthest with the least material (see Figure 9). Future work will look at generalizing reactive counterbalancing to other geometries.

HARDWARE

To demonstrate key capabilities in a physical system, hardware prototypes were developed for an instrumented node-and-strut assembly and an autonomous strut-climbing robot. Previous work discusses these prototypes in detail (Melenbrink et al. 2017). In the following section, we generalize our force-sensing techniques to cellular geometries, which may be preferable to node-and-strut assemblies in some cases.

Robot Design

We present a design for an autonomous strut-climbing robot intended to work with the type of triangular lattice described in the Simulations section. The robot consists of two carriages capable of independently gripping struts, allowing transitions from one strut to another (Figure 11). The sequence by which a robot would use pogo pin connectors to attach to a node, measure its forces and save them to the EEPROM memory of the node’s microprocessor are described in previous work (Melenbrink et al. 2017). The sequence for installing struts and zeroing sensors is also described. The robot was able to locomote along struts and autonomously transition from one to another; however, the reliability of these operations proved a challenge. Future work will look at redesigning the morphology of the robot.

Sensor Design

A key component of this work is a custom force-sensing method, which would be required for the kinds of building tasks explored in simulation. There are a number of reasons why strain gauges were deemed impractical for this application, and were abandoned in favor of simple force sensors. Strain gauges are difficult to install and calibrate uniformly over multiple elements, and

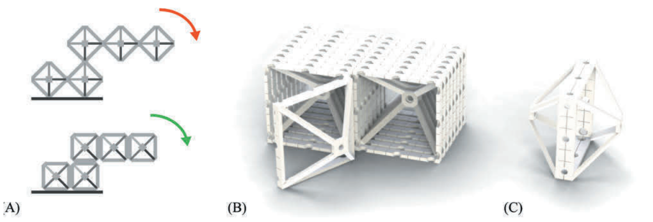

14

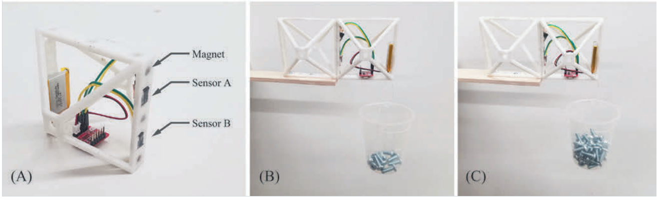

15

14 Images describing a preliminary cellular unit. (A) Previous work has determined that vertex-connected octahedra (top) can yield a desirable strength-to-weight ratio (Popescu 2006). While this arrangement is ideal in compression, it is not suitable for resisting bending, which is needed for the cantilevering described in this work. Edge-connected octahedra (below) are better at resisting bending. (B) shows how these shapes could be efficiently stacked for transport, then folded and locked into place to produce a cellular unit (C).

15 (A) Shows the instrumented cellular unit used for this evaluation. The connection to another cell is made by magnetic force, and two sensors (as described in the previous section) installed at the top (Sensor A) and the bottom (Sensor B) of the face report the force exerted on them. In a more complete system, all four edges would be equipped with sensors, and a mechanical connection would replace the magnets; this prototype tests only the accuracy of a single joint. The cell also includes a LiPo battery and a microprocessor capable of transmitting data over Bluetooth. The experimental setup is shown in (B, C); two cells are joined by magnetic force. As weights (assorted hardware) are added to the cup suspended from the cantilevering unit, the force on Sensor A decreases while the force on Sensor B increases. Once their difference crosses a preset threshold, an LED turns on (C), indicating that the joint is in danger of breaking. Known quantities of mass are then added until the joint breaks.

require additional amplification circuitry that is not needed for force sensors. Independent of amplification circuitry, strain gauges themselves range from tens to hundreds of US dollars per unit, while force sensors can be produced for a fraction of a dollar. As a structure may require hundreds of instrumented joints, this difference quickly becomes significant. Thus we instead propose a slim-package force sensor comprising a strip of force-sensitive material (e.g., Velostat) sandwiched between two copper sheets (Figure 13A). When force is exerted on the sensor, the resistance of the material decreases, yielding an increase in the voltage running through the sensor that is easily detected by a microprocessor.

The slim profile of the sensors is also advantageous because it allows the sensors to be placed directly in the load path, enabling them to report axial force, which strain gauges are unable to do.

GEOMETRY AND SENSING EXTENSION

This sensor was previously found to be effective at detecting small variations in applied load on a strut inserted into a socket, such as the one shown in Figure 13B (Melenbrink et al. 2017). This could be considered an early-stage hardware verification for the 2D system described in the Simulation section. However, due to the difficulties of a node-and-strut type system described in the Hardware section, it was deemed worthwhile to consider the possibility of construction with stable cellular units. In this section, we evaluate the same type of sensing system as applied to cellular units. The detailed design of the cellular unit itself is the subject of future work; for now we assume the use of the shape described in Figure 14. Further development will focus on this design in a more principled way—it should be structurally optimal, cost effective, able to facilitate robot locomotion, and should include a mechanical connection in addition to the magnets used for alignment.

Sensor Evaluation

The instrumented cellular unit described in Figure 14 was evaluated for its ability to reliably predict failures due to breaking by measuring two force sensors (as seen in Figure 15A). The microprocessor turns on its LED when a preset threshold in the difference between sensor readings is exceeded, indicating that the joint is in danger of breaking. We set this threshold value so that this warning occurred when the load was 2/3 of the value at which the joint would empirically break. To evaluate the reliability of both the warning signal and the joint strength, we conducted 20 trials, gradually increasing the load on the joint until failure (Figure 15). The applied load at which the warning state was triggered was 224 ± 18 g; the load at which the joint broke was 336 ± 4 g. The sensing system reliably predicted failure for all trials: the joint never broke before the LED lit, suggesting that such a system could indeed provide a method for identifying stressed joints before they break.

In future work, we will look at extending this sensor configuration to measure forces at the supply cache (similar to what is outlined in the Simulation section) in order to enable dynamic counterbalancing, which has been shown to be useful for maintaining overall stability.

FUTURE WORK

Further research will continue to consider simulation and hardware in tandem. In simulation, we will evaluate the affordances of node-and-strut vs. cellular construction systems, and systematically develop a method for determining safety factors that reliably prevent failures in changing conditions. The resiliency of the system will be rigorously studied in simulation, ensuring the system's satisfactory response to environmental hazards such as sudden gusts of wind or unstable ground. We will attempt to generalize the principle that force-awareness allows for more stable construction to three-dimensional geometries. We expect

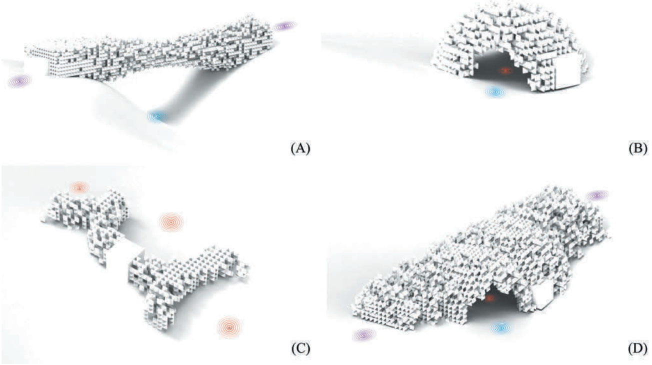

16

16 These preliminary models are snapshots in three-dimensional building sequences where force-aware robotic agents install cellular units such as described in Figure 14. Remote frequency "instruction beacons" are placed in the terrain and emit high-level distance-based instructions such as *build a wall approximately 3 m away* (red), *don't build anything within 2 m* (blue), or *target for anchoring* (purple). These instructions can be used independently or composited to form various structural typologies such as a bridge (A), canopy (B), wall (C) or hybrid (D). The resultant forms are emergent functions of the swarm's attempt to satisfy the requirements of the beacons while maintaining stability throughout the building sequence.

the results to still hold, as the prevailing failure mode is due to bending imposed by gravitational forces, which is already accounted for in the two-dimensional vertical plane. Another area for further investigation is on architectural-scale instantiations of the construction swarm in simulation. The Introduction section alluded to the possibilities of a persistently changing architecture that responds to high-level user-specified instructions, as opposed to specific blueprints. However, the formal articulation and structural affordances of such a system have yet to be explored. Figure 16 suggests some preliminary possibilities.

Future hardware development will consist of proposing a 3D construction system, either node-and-strut or cellular units. A new robotic agent will be co-designed to work in concert with the work-in-progress construction medium. Ultimately, the system should consist of a swarm of distributed robots and suitable construction materials, which would be able to reliably and autonomously assemble architectural structures that respond to high-level functional requirements while maintaining structural stability throughout the assembly.

While these are formidable challenges, in this paper we have demonstrated that one key step towards this goal, the use of force sensing and a corresponding control algorithm, is feasible, and could potentially be implemented on a variety of construction systems.

ACKNOWLEDGEMENTS

This research was supported by the Harvard University Wyss Institute for Biologically Inspired Engineering and by Autodesk Inc.

REFERENCES

Augugliaro, Frederico, Sergei Lupashin, Michael Hamer, Cason Male, Markus Hehn, Mark W. Mueller, Jan Sebastian Willmann, Fabio Gramazio, Matthias Kohler, and Raffaello D'Andrea. 2014. "The Flight Assembled Architecture Installation: Cooperative Construction with Flying Machines." *IEEE Control Systems* 34 (4): 46-64.

Brodbeck, Luzius, and Fumiya Iida. 2014. "Automatic Real-World Assembly of Machine-Designed Structures." In *International Conference on Robotics and Automation*, 1221–1226. Hong Kong: ICRA.

Dierichs, Karola, and Achim Menges. 2012. "Aggregate Structures: Material And Machine Computation Of Designed Granular Substances." *Architectural Design* 82 (2): 74–81.

Everist, Jacob, Kasra Mogharei, Harshit Suri, Nadeesha Ranasinghe, Berok Khoshnevis, Peter Will, and Wei-Min Shen. 2004. "A System For In-Space Assembly." In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2356–61. Sendai, Japan: IROS

Galloway, Kevin C., Rekha Jois, and Mark Yim. 2010. "Factory Floor: A Robotically Reconfigurable Construction Platform." In *International Conference on Robotics and Automation*, 2467–72. Anchorage, AK: ICRA.

Hjelle, David A. and Hod Lipson. 2009. "A Robotically Reconfigurable Truss." In *Proceedings of ASME/IFToMM International Conference on Reconfigurable Mechanisms and Robots*. London: ReMAR.

ICD/ITKE. 2017. "Performative Morphology: ICD/ITKE Research Pavilion." http://icd.uni-stuttgart.de/?p=19195

Keating, Steven J., Julian C. Leland, Levi Cai, and Neri Oxman. 2017. "Toward Site-Specific and Self-Sufficient Robotic Fabrication On Architectural Scales." *Science Robotics* 2 (5): 89–86.

Lindsey, Quentin, Daniel Mellinger, and Vijay Kumar. 2011. "Construction Of Cubic Structures With Quadrotor Teams." In *Robotics: Science and Systems VII*, 177–184. Cambridge, MA: MIT Press.

McEvoy, Michael, Erik Komendera, and Nikolaus Correll. 2014. "Assembly Path Planning For Stable Robotic Construction." In *International Conference on Technologies for Practical Robot Applications*. Woburn, MA: TePRA.

Melenbrink, Nathan, and Justin Werfel. 2017. "Local Force Cues For Strength And Stability In A Distributed Robotic Construction System." Forthcoming in *Journal of Swarm Intelligence: Special Issue on Self-Organised Construction*.

Melenbrink, Nathan, Paul Kassabian, Panagiotis Michalatos, and Justin Werfel. 2017. "Using Force Measurements To Guide Construction By Distributed Climbing Robots." Forthcoming in *International Conference on Intelligent Robots and Systems 2017*.

Nan, Cristina. 2015. "A New Machinecraft. Architectural Robots." In *Emerging Experience in Past, Present and Future of Digital Architecture*, *Proceedings of the 20th International Conference of the Association for Computer-Aided Architectural Design Research in Asia*, 745–54. Daegu, Korea: CAADRIA.

Nigl, Franz, Shuguang Li, Jeremy E. Blum, and Hod Lipson. 2013. "Structure-Reconfiguring Robots: Autonomous Truss Reconfiguration And Manipulation." *IEEE Robotics and Automation Magazine* 20 (3): 60–71.

Peters, Scott and Robert Belden. 2014. "SAM, The Robotic Bricklayer." *SMART: Dynamics of Masonry* 1 (4): 10–14.

Pollack, Jordan, Hod Lipson, Pablo Funes, Sevan Ficici, and Gregory Hornby. 1999. "Coevolutionary Robotics." In *Proceedings of the First NASA/DoD Workshop on Evolvable Hardware*, 208–216. Pasadena, CA: EH.

Popescu, George A., Tushar Mahale, and Neil Gershenfeld. 2006. "Digital Materials For Digital Printing." In *Proceedings of the NIP & Digital Fabrication Conference*, no. 3, 58–61. IS&T.

Rob|Arch. 2014. "General Information." http://www.robarch2014.org/conference

Sousa, José Pedro, Cristina Gassó Palop, Eduardo Moreira, Andry Maykol Pinto, José Lima, Paulo Costa, Pedro Costa, Germano Veiga, and A. Paulo Moreira. 2016. "The Spiderobot: A Cable-Robot System For On-Site Construction In Architecture." In *Robotic Fabrication in Architecture, Art and Design 2016*, 230–239. Cham, Switzerland: Springer.

Terada, Yuzuru, and Satoshi Murata. 2004. "Automatic Assembly System For A Large-Scale Modular Structure: Hardware Design Of Module And Assembler Robot." In *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2349–55. Sendai, Japan: IROS.

Wang, Zijian, and Mac Schwager. 2016. "Force-Amplifying N-Robot Transport System (Force-Ants) For Cooperative Planar Manipulation Without Communication." *The International Journal of Robotics Research* 35 (13): 1564–1586.

Werfel, Justin, Kirstin Petersen, and Radhika Nagpal. 2014. "Designing Collective Behavior In A Termite-Inspired Robot Construction Team." *Science* 343 (6172): 754–758.

Werfel, Justin, Yaneer Bar-Yam, Daniela Rus, and Radhika Nagpal. 2006. "Distributed Construction By Mobile Robots With Enhanced Building Blocks." In *Proceedings of the International Conference on Robotics and Automation*, 2787–94. Orlando, FL: ICRA.

Willmann, Jan, Federico Augugliaro, Thomas Cadalbert, Raffaello D'Andrea, Fabio Gramazio, and Matthias Kohler. 2012. "Aerial Robotic Construction Towards A New Field Of Architectural Research." *International Journal of Architectural Computing* 10 (3): 439–459.

Yun, Seung-kook, and Daniela Rus. 2008. "Self Assembly Of Modular Manipulators With Active And Passive Modules." In *Proceedings of the International Conference on Robotics and Automation*, 1477–82. Pasadena, CA: ICRA

IMAGE CREDITS

Figure 2: (A) David Mark, 2005; (B) Bishnu Sarangi, 2014; (C) Rui Ornelas, 2007

Figure 3: (A) Yuzuru Terada © 2004 IEEE. Reprinted, with permission, from (Terada and Murata 2004); (B) Justin Werfel © 2006 IEEE. Reprinted, with permission, from (Werfel et al. 2006); (C) Eliza Grinnell © 2013 Harvard School of Engineering and Applied Sciences; (D) Javier Burón, 2012; (E) Hod Lipson © 2013 IEEE. Reprinted, with permission, from (Nigl et al. 2013); (F) Seung-kook Yun © 2008 IEEE. Reprinted, with permission, from (Yun and Rus 2008)

Figure 4: (A) Michael McEvoy © 2014 IEEE. Reprinted, with permission, from (McEvoy et al. 2014); (B) Jordan Pollack © 1999 IEEE. Reprinted, with permission, from (Pollack et al. 1999); (C) Luzius Brodbeck © 2014 IEEE. Reprinted, with permission, from (Brodbeck and Iida 2014)

Figure 5: (A) Burggraf / Reichert, 2017 © ICD/ITKE; (B) Karola Dierichs, 2017 © Karola Dierichs; (C) Scott Peters 2014 © Construction Robotics; (D) Saša Jokić 2013 © IAAC

Nathan Melenbrink is a doctoral candidate at the Institute for Computational Design and Construction at the University of Stuttgart. He is currently in residency as a Research Fellow at the Wyss Institute for Biologically Inspired Engineering at Harvard University, where his research focuses on swarm robotics for construction. His previous teaching experience includes MIT, the University of Hong Kong Shanghai Study Center, the AA Shanghai Visiting School, and Virginia Tech.

Paul Kassabian is an Associate Principal at Simpson Gumpertz & Heger Inc. (SGH) and specializes in structural design of buildings, bridges, and sculptures. He develops innovative methods of design and construction using SGH's in-house materials lab, computational design approaches, and digital fabrication techniques. Paul taught at Massachusetts Institute of Technology for ten years and is currently teaching Structures at Harvard's GSD.

Professor Achim Menges is a registered architect and professor at the University of Stuttgart, where he is the founding director of the Institute for Computational Design and Construction since 2008. His practice and research focuses on the development of integral design processes at the intersection of morphogenetic design computation, biomimetic engineering and computer aided manufacturing that enables a highly articulated, performative built environment.

Dr. Justin Werfel is a Senior Research Scientist at the Wyss Institute for Biologically Inspired Engineering at Harvard University, where he leads the Designing Emergence Laboratory. His research interests are in understanding and designing complex and emergent systems, with work in areas including swarm robotics, social insect behavior, evolutionary theory, engineered molecular nanosystems, and educational technology.