Three-Dimensional Directed Construction

Justin Werfel and Radhika Nagpal
Harvard University

We present an approach to building arbitrary solid three-dimensional structures from modular units. Our bipartite system comprises passive units, which form the structure, and active ones, which manipulate the passive units and can be removed at the end of the construction process for use in other projects. "Passive" units are not self-mobile but can communicate locally with physically attached neighbors. We present a decentralized algorithm for formation of user-specified structures, show that the required communication scales favorably with the size of the structure, and discuss how unit failures can be detected and corrected.

Self-reconfigurable modular robotics traditionally considers uniformly self-mobile modules. This universal mobility is important in a system intended to reconfigure repeatedly or frequently. However, in a system intended for construction of static structures, it may be inappropriate. Once a structure is complete and modules need not move again, the capacity for movement can be a liability: Not only is it unnecessary from that point on (and the hardware and associated complexity and expense wasted thereafter), but self-mobile units may well be less effective as structural elements than specialized passive ones. The separation into two classes of units allows passive units to be optimized for structural capabilities, and active ones to be reused elsewhere.

While structural units need not be self-mobile, a small amount of embedded processing power and a capacity for local communication can allow them to disseminate nonlocal structure information. As we show, this ability can be exploited for decentralized algorithms that avoid unwanted structural flaws and deadlock conditions, even though the timing and order of mobile agent actions is not known in advance.

In our system, we assume that mobile units can bring passive ones to a structure in progress, travel along its surface, and attach them subject to a conservative assumption on physical movement: a (cubic) unit cannot be pushed into a space one unit wide directly between two others. As a corollary, this restriction prevents more complicated situations where intended attachment sites become inaccessible or where mobile units might have to travel down difficult "tunnels".

The problem of building an arbitrary user-specified structure can be considered in two parts: (a) determining whether a passive unit may legally be attached at a given

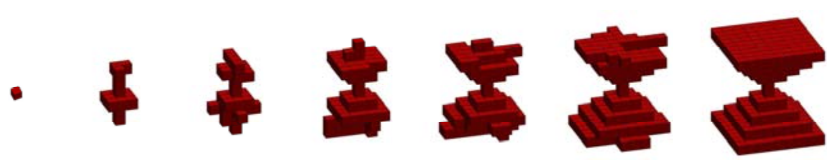

Figure 1: Successive snapshots during the construction of an hourglass-shaped structure. Only passive units attached to the structure are shown.

site, and (b) ensuring mobile units find all necessary sites.

Part (a) is addressed by local inter-unit communication. Passive units communicate minimal information about where relevant units have already been attached, and accordingly grant or deny permission to mobile units to attach further units at empty sites. A few simple rules, consonant with the physical movement constraint above, suffice to allow arbitrary solid structures to be reliably built (Fig. 1). The amount of communication necessary to achieve this scheme scales linearly with the size of the structure—that is, the number of inter-unit messages per unit remains constant regardless of structure size. Thus this approach can be feasible for very large structures without being limited by the amount of communication required.

For part (b), we compare three different approaches (each used in previous work on self-reconfigurable systems) by which mobile units find attachment sites. One is random movement along the structure surface. The second is a more systematic, directed search. The third is based on receiving directions from the passive units: those passive units with faces where another could be attached send out a numerical gradient (confined to the surface) to their neighbors; mobile units follow the gradient to the nearest available attachment site. We evaluate these approaches experimentally for a number of different structures (Fig. 2). Briefly, the first approach has the advantage of simplicity; the second eliminates some unnecessary movement and guarantees visiting all sites in finite time; the third results in an order of magnitude less movement, at the cost of an order of magnitude more communication within the structure.

Failures of passive units can be detected by their neighbors, and mobile units recruited to replace the faulty ones. We outline the approach to do so, in which the minimal number of functional units must be removed to allow access to the broken one. This process is coordinated collectively by the remaining passive units. Construction may continue elsewhere in the structure while this correction is taking place.

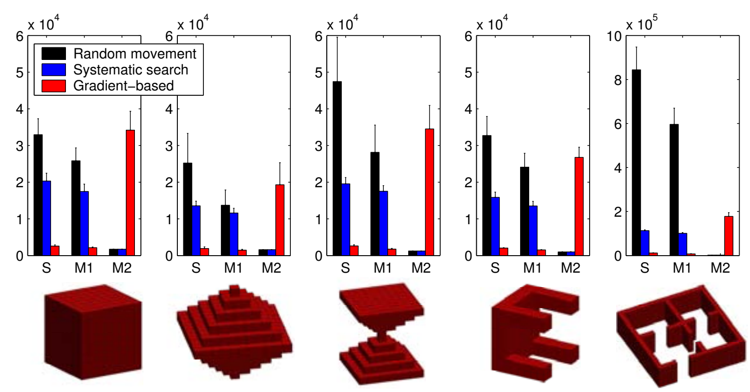

Figure 2: Experimental results for movement and communication required to build various structures. S: movement along surface; M1: messages passed from passive to mobile blocks; M2: messages passed between passive blocks.