

Compiler for Scalable Construction by the TERMES Robot Collective

Yawen Deng¹, Yiwen Hua¹, Nils Napp², and Kirstin Petersen¹

¹ *Cornell University, Ithaca, NY 14850, USA, kirstin@cornell.edu*

² *University at Buffalo, Buffalo, NY 14260, USA*

Abstract

The TERMES system is a robot collective capable of autonomous construction of 3D user-specified structures. A key component of the framework is an off-line compiler which takes in a structure blueprint and generates a directed map, in turn permitting an arbitrary number of robots to perform decentralized construction in a provably correct manner. In past work, this compiler was limited to a non-optimized search approach which scaled poorly with the structure size. Here, we first recast the process as a constraint satisfaction problem (CSP) to apply well-known optimizations for solving CSP and present new scalable compiler schemes and the ability to quickly generate provably correct maps (or find that none exist) of structures with up to 1 million bricks. We compare the performance of the compilers on a range of structures, and show how the compilation time is related to the inter-dependencies between built locations. Second, we show how the transition probability between locations in the structure affect assembly time. While the exact solution for the expected completion time is difficult to compute, we evaluate different objective functions for the transition probabilities and show that these optimizations can drastically improve overall efficiency. This work represents an important step towards collective robotic construction of real-world structures.

Keywords: Multi-Robot Systems, Assembly, construction, Autonomy, Compiler

¹ **1. Introduction**

² Autonomous robots have the potential to revolutionize the construction
³ industry enabling rapid fabrication of inexpensive structures, novel designs,
⁴ and construction in novel settings. Researchers and industrial specialists
⁵ have proposed many solutions to these challenges, one of which involves
⁶ collectives of autonomous mobile robots which can assemble structures much
⁷ larger than the size of the individuals [1]. By focusing on distributed scalable
⁸ coordination, such systems may deploy many robots to work efficiently in
⁹ parallel and be tolerant to individual failures. Although robot collectives
¹⁰ have received a lot of attention over the past couple of decades [2], most
¹¹ demonstrations are limited to controlled laboratory settings, relatively small
¹² assemblies, and/or small collectives. Open challenges range from scalable
¹³ algorithms to capable, low-maintenance hardware. Here, we focus on the
¹⁴ former, i.e. improving the algorithmic framework in terms of how it scales
¹⁵ with the size of the structure. We present our results in the context of
¹⁶ the TERMES system presented in previous literature [3, 4, 5, 6], but our
¹⁷ approach may generalize to other distributed construction systems.

¹⁸ The TERMES hardware consists of custom bricks and simple robots capa-
¹⁹ ble of climbing on, navigating, and adding bricks to the structure (Fig. 1.A).
²⁰ Inspired by construction in social insects, the robots coordinate construction
²¹ implicitly through their environment in a scalable manner. Despite this mini-
²² malistic approach the system has been shown to assemble 3D structures with
²³ provable guarantees, by relying on a combination of an off-line compiler and
²⁴ an onboard rule set. The compiler converts the structure blueprint to a 2D-
²⁵ map with assembly locations, the desired number of bricks at each location,
²⁶ and designated travel directions between locations (Fig. 1.B). This map is
²⁷ given to an arbitrary number of robots, which follow these instructions and
²⁸ add material as determined by the onboard rule set which is dictated solely
²⁹ by the limitations of the robot platform used (Fig. 1.C-D). The scalability of
³⁰ the TERMES and similar systems is determined by several factors, including
³¹ 1) hardware cost and manufacturing complexity; 2) robot reliability and how
³² likely failures are to disrupt system progress; 3) how the coordination mech-
³³ anisms scale with the size of the collective; 4) how the compiler computation
³⁴ time scales with the size of the structure; and finally, and 5) how efficiently
³⁵ robots can reach the assembly frontier.

³⁶ Fabrication and robot reliability (points 1-2) were addressed in [4]. The
³⁷ system was designed with minimalism in mind - co-design of robots, bricks,




Figure 1: A) Photo of the TERMES system. B) Example of the map generated by the compiler (top view). The digits indicate the number of bricks at each location; arrows how robots can transition between locations. System limitations include that bricks cannot be added in between others bricks (C, dimetric view), and that robots can climb at most one brick height between neighboring locations (D, side view). The set of structures which are compilable are not necessarily intuitive. E) shows a structure which cannot be compiled, because the only way for a robot to complete the center would be an assembly move of type C.

38 and algorithms resulted in a simple robot costing $\sim \$2K$ with a 1-week assembly time. The cost of the mechanics was brought down considerably in a
 39 a subsequent paper [6]. To support reliability, the focus was not on achieving
 40 perfect behavior, but rather to enable robots to recognize and fix errors
 41 before they propagated. Scalability of the collective (point 3) was addressed
 42 implicitly by relying on the structure as a shared physical database through
 43 which the robots can coordinate [3, 5, 6]. Here, we focus instead on improving
 44 the TERMES compiler to make it feasible to compile maps of large-scale
 45 structures (point 4). The work presented in this paper includes that of the
 46 conference paper presented at the International Symposium on Distributed
 47 Autonomous Robotic Systems (DARS) 2018 [7], with an additional contribu-
 48 tion of showing how the transition probabilities between locations in the
 49 map affect structure assembly time, and how these can be optimized such
 50 that robots can complete the structure significantly faster (point 5).

51 First, we recast the compiler originally described in [3] (Sec. 3) as a back-
 52 tracking solution to a constraint satisfaction problem (CSP) with pairwise,
 53 partial, and global constraint checking. We show that the original compiler
 54 scales poorly with the size of the structure (Sec. 4). By examining the be-
 55 havior of the original search as a solution to a CSP, we are able to achieve
 56 significant improvements by formulating a new CSP that better exploits for-

58 ward checking pairwise constraints during the backtracking search (Sec. 5).
59 We then describe and prove an entirely new formulation for generating maps
60 that is not based on search, but an iterative method that builds up feasible
61 maps by considering locations in a breadth-first manner starting from the
62 exit location (Sec. 6). We show the ability of the latter to compile struc-
63 tures with up to 1 million bricks in ~ 1 min on commodity hardware. We
64 compare the performance of these compilers on different sets of structures
65 (Sec. 7), including unbuildable ones which are computationally intractable
66 for search-based compilers. Finally, we show how, after the map has been
67 compiled, construction speed may be improved simply by altering the transi-
68 tion probabilities between locations, with examples of a 2 order of magnitude
69 improvement in completion time (Sec. 8).

70 2. Related Work

71 Collective robotic construction can be achieved in a variety of ways, and
72 examples include pre-programmed robots for functional structures [8, 9],
73 template-based construction [10], centralized controllers that allow for paral-
74 lelism [11], communication-based coordination [12, 13], and compiler-based
75 systems [3, 14, 15].

76 Compilers for generating matter, which take high-level specifications and
77 generate parallel assembly steps, are used in a variety of fields, e.g. digital
78 materials [16], self-assembly, and modular robots [17]. In the construction
79 setting, compilers must take into consideration the physical constraints of
80 both building material and the robots that manipulate it. Constraints may
81 exist both in mechanisms (e.g. the ability to traverse the structure) and
82 perception/cognition (the ability to sense/remember the state of the sur-
83 rounding structure). Broadly categorized, there are two ways to approach
84 compilers [2]. The first is to define a set of sub-structures for which an assem-
85 bly plan is known, and then to decompose new structures into combinations
86 of those. The second is to compile based purely on the physical constraints
87 of the system. Although the first method makes reasoning and guarantees
88 easier, it also limits the set of structures (some structures that robots are
89 physically capable of building cannot be compiled). The second method
90 does not artificially restrict the set of buildable structures, but makes it hard
91 to reason about what is buildable. In case of the latter, it is therefore critical
92 that compilers can quickly assess whether or not a structure is buildable, or
93 potentially come up with alternative solutions [18, 5].

94 An example of the first approach include Seo et al. [15] who presented a
95 compiler for 2D assembly of simply connected structures of floating bricks
96 by boat-like robots, which decomposes structures into linear cells. Another
97 example involves that of Lindsey et al. [11, 19] who presented a compiler
98 for assembly of strut structures by teams of quadcopters. The struts could
99 be assembled into structurally stable cubes. Consequently, the compiler was
100 designed to generate assembly rules for any structure which was decompos-
101 able into such special cubic structures. Both of these systems have a concise
102 definition of the class of compilable structures.

103 The TERMES compiler is search-based and uses hardware limitations as
104 constraints. As previously mentioned, this makes it harder to infer which
105 structures are buildable. Figure 1.B and E shows structures which are build-
106 able and unbuildable, respectively, despite the fact that they differ by only
107 one location and despite the fact that it is possible for a robot to physically
108 assemble each separate location. The issue is that there is no way to consis-
109 tently order the assembly steps without violating the constraint shown in C.
110 Currently, for TERMES-like constraints, there is no good specification for
111 which structures have valid maps, other than when a map is found. This
112 is especially problematic if the compiler used is slow and has a long run-
113 time before failing. Here, we show that the compiler presented in [3] scales
114 poorly with the size and complexity of the structure, and present an alter-
115 native compilation method, such that arbitrary structures can be compiled
116 and checked quickly.

117 The second contribution of the paper concerns construction efficiency:
118 i.e. after the offline compilation, how fast can the structure be completed
119 by a given number of robots moving stochastically according to the map.
120 The randomized execution model makes global state sharing unnecessary
121 and thus makes concurrent execution between an arbitrary number of robot
122 easy. However, it also introduces inefficiencies because of 1) physical bot-
123 tlenecks which limits the number of robots that can simultaneously pass
124 through a location and 2) construction order, i.e. the need for some actions
125 to be completed before others can take place. Related work on optimizing
126 assembly plans for TERMES focus on optimizing the map structure [18].
127 Here, we leave the original map in place and instead focus on optimizing the
128 probabilities between different paths through the map. Past work on opti-
129 mizing stochastic assembly policies under such spatial- and order-constrained
130 scheduling is limited. In [20], the authors analyze stochastic assembly algo-
131 rithms constrained both by assembly orders and by raw materials through

132 chemical reaction models. In [21] the problem of optimizing transitions for
133 material transport under spatio-temporal constraints is addressed, however,
134 the transition probabilities are constrained to a relatively small parameter-
135 ized model. Efficient spatial allocation of assembly robots have been shown
136 in [22, 23, 14], with the ability to adapt to local failures and shape changes
137 through space partitioning.

138 3. Problem Formulation

139 A structure consists of a finite set of locations L that each have integer
140 x and y location, i.e. $(l_x, l_y) = l \in L$. Two locations $l, k \in L$ are said
141 to be neighbors when either the x or y differ by one, but not when both
142 are different. This type of neighbor relation corresponds to a distance of
143 1 with the Manhattan distance metric. A *path* is a sequence of locations
144 $p = (l_1, l_2, \dots, l_N)$ such that consecutive locations are neighbors. We assume
145 that all the locations for a structure are path connected, i.e. every location
146 has a path to every other location. Disconnected structures can be treated
147 as separate structures. There are two special locations, $l_{START} \in L$ and
148 $l_{EXIT} \in L$, which correspond to the start and exit locations. In a structure,
149 each location l has a target height $h_l \in \mathbb{N}$. We say that a path is *traversable*
150 if each consecutive location differs in height by at most 1, which corresponds
151 to the motion limitations of a TERMES robot.

152 In order to make a building plan for the TERMES system, we need to
153 generate a directed graph on the vertex set L . To avoid the physical assembly
154 constraint shown in Fig. 1.C the graph needs to be acyclic and a location
155 cannot have two opposing incoming edges. To ensure traversability, the graph
156 must have the additional properties that for every $l \in L$ there is a directed,
157 traversable path from l_{START} to reach l and for every $l \in L$ there is a directed,
158 traversable path to reach l_{EXIT} . l_{START} has all outgoing edges; l_{EXIT} has all
159 incoming edges.

160 In summary, the properties of a valid map are as follows:

Property 1: The map contains no cycles.

Property 2: The map contains no opposing incoming arrows.

161 *Property 3:* All locations can reach an exit on a traversable path that
is consistent with the assigned edges.

Property 4: The start can reach all locations on a traversable path
that is consistent with the assigned edges.

162 Properties 3 and 4 imply that, except for l_{START} and l_{EXIT} all locations
163 must have directed edges that point both in- and outwards. We refer to
164 this local check for Properties 3 and 4 as the sink/source-condition. We will
165 reference these properties throughout the following sections.

166 **4. Edge-CSP Compiler**

167 Past TERMES publications described a procedure for searching through
168 the space of available assignments [3]. Here, we recast this compiler as a
169 backtracking search to a CSP with pairwise, partial, and global constraint
170 checking. The CSP problem consist of variables, domains (the possible val-
171 ues for each variable), and constraints (how variable assignments affect each
172 other). The goal of backtracking search is to find an *assignment*, i.e. picking
173 from each domain one value for each variable [24, Ch6].

174 In accordance with the compiler described in [3], we make variables cor-
175 respond to edges between neighboring locations and give them a domain of
176 the two possible edge directions. We refer to this compiler as an Edge-CSP
177 compiler, further shown in Fig. 2.A. The Edge-CSP tries to pick both a good
178 variable ordering and a good domain ordering. The variable ordering is to
179 pick variables that are adjacent to already assigned edges and as close to
180 l_{START} as possible. The domains are ordered to first explore edges that point
181 “away” from l_{START} in a breadth first manner. This choice is based on the
182 observation that most edges in valid maps have this orientation.

183 We use three types of constraints. Binary constraints between edges that
184 comply with Property 2. Constraints on partial assignments which check
185 for cycles, i.e. Property 1, and checks that each location with fully assigned
186 edges other than l_{START} and l_{EXIT} complies with the sink/source-condition.
187 Constraints on the global assignment which checks Property 3-4, that every
188 location can be reached from l_{START} and that l_{EXIT} can be reached from
189 every location. The benefit of the binary checks is that constraints may be
190 propagated forward to speed up the search using forward checking [24, Ch6].
191 We use the AC3 algorithm to do this [25]. Forward checking with the binary
192 constraints enable a behavior equivalent to the “row rule” discussed in [3],
193 i.e. a behavior that causes the structure to be built from one point outwards.
194 An example of this is shown in Fig. 2.A; if v_1 is fixed, v_2 and v_3 are as well.
195 Reversely, the fixed value of v_{17} does not directly affect those around it.

196 Notice that this compiler does not take the height of the structure into
197 consideration until the final global check. The search continues until all




Figure 2: Two versions of the CSP compiler applied to a $3 \times 4 \times 1$ structure. A) In the Edge-CSP variables correspond to edges between locations. The domain for v_6 are shown as an example to the right of the structure. We can forward propagate the fixed variables, v_1 and v_4 shown in red, to fix v_2 , v_3 , and v_{11} shown in yellow according to property 2. B) In the Location-CSP variables correspond to all possible combinations of directions to and from the location. The domain for v_6 are shown as an example to the right of the structure. This scheme produces a fully connected graph in which all constraints affect each other.

198 domain combinations have been tried, or have been eliminated early by a local
 199 or partial check. The total number of possible domain combinations scales
 200 as $O(2^n)$, where n corresponds to the number of edges between locations
 201 in the structure. However, early termination of partial assignments prunes
 202 the space significantly. In general, all backtracking search may work well on
 203 structures that have many feasible solutions, but will scale poorly with large
 204 structures that have only a few or no solutions, and where bad branches in
 205 the search tree cannot be pruned early.

206 Analyzing the compiler as a CSP shows that the binary constraints for-
 207 mulated on edges limits the amount of forward checking that can be done,
 208 since each row or column results in a disconnected component of constraint
 209 arcs. Furthermore, it is not possible to use the sink/source-condition to for-
 210 ward propagate because it cannot be expressed as a binary constraint. To
 211 address these shortcomings we formulate a more efficient CSP to solve the
 212 same problem in Sec. 5.

213 5. Location-CSP Compiler

214 To speed up the backtracking search, we change the formulation of the
 215 CSP such that the variables become the locations and the domains include
 216 all combinations of travel directions on the 4 edges as illustrated in Fig. 2.B.
 217 Consequently we refer to this algorithm as a Location-CSP compiler. The
 218 benefit of this scheme is that it creates a fully connected graph, where con-

219 straints may more readily affect other variables. Note that like in the Edge-
220 CSP, cycles and structure traversability is not checked until after partial or
221 full assignment.

222 6. BFD Compiler

223 The final compiler is not based on search, but instead does an iterative
224 assignment of the edge directions in a breadth-first manner starting from
225 l_{EXIT} . Essentially, it evaluates if a location may serve as a drain (an exit-
226 like location) for the intermediate structures where locations whose travel
227 directions have been fully assigned were removed. We refer to this algorithm
228 as a Breadth-First Disassembly (BFD) compiler. The process is shown in
229 Fig. 3 and Alg. 1. Upon initialization, l_{EXIT} is added to the frontier list,
230 $Q_{frontier}$. The compiler iteratively takes a location, l_0 , from $Q_{frontier}$ and
231 checks if it can serve as a drain. To serve as a drain, l_0 must have the
232 following properties: 1) to comply with Property 2 it cannot be in between
233 two unassigned locations, 2) it needs to have a traversable path to l_{EXIT}
234 that only uses previously disassembled locations, and 3) it cannot cause a
235 disconnect in the structure which would cause a violation of Property 4. If
236 these statements are true l_0 is added to $Q_{visited}$, the edges to all neighbors are
237 assigned as ingoing, and traversable neighbors are added to $Q_{frontier}$. The
238 compiler continues to do this until $Q_{frontier}$ is empty or no solution is found.

239 The biggest overhead in the BFD compiler is the connectivity check which
240 happens each time a location is tested as a viable drain. Note that the con-
241 nectivity check takes the traversable height of the neighboring locations into
242 account. We implement two versions of this check. 1) BFD_0 : To check
243 the connectivity, the compiler conducts a breadth-first search starting from
244 l_{START} to count the number of reachable locations following unassigned edges.
245 If this count is equal to the number of unvisited locations, l_0 may serve as a
246 drain. This requires a complete check of all remaining locations ($L \setminus Q_{visited}$).
247 2) BFD : To speed up this process, we cache the connectivity computation
248 by generating a spanning tree of unvisited locations. Removing leaves in the
249 tree does not disconnect the graph, so the connectivity check can return an
250 answer without having to traverse any nodes in the spanning tree. When
251 the connectivity check is for a non-leaf node, we perform the original con-
252 nectivity check. If l_0 does not disconnect the structure we add it to $Q_{visited}$
253 and recompute the spanning tree. To create a spanning tree that is likely to
254 have leaf-nodes in $Q_{frontier}$, we add edges in breadth first manner beginning

Algorithm 1 Pseudo code for the BFD Compiler which either returns a valid map, or identifies that no such map exists. l_0 denotes the current location in question and l_i its neighboring locations. $Q_{visited}$ is the set of visited locations which have been 'disassembled', i.e. fully determined; and $Q_{frontier}$ is the frontier, i.e. locations that have traversable paths to the exit and could potentially be disassembled next.

```

1: initialize  $Q_{frontier}$  and  $Q_{visited}$  as empty
2: initialize  $map$  to be an empty graph over the vertex set  $L$ 
3: add  $L_{EXIT}$  to  $Q_{frontier}$ 
4: while  $Q_{frontier}$  is not empty do
5:   remove  $l_0$  from  $Q_{frontier}$ 
6:   if  $l_0$  is not in between two other unvisited sites (Property 2)
      and removing  $l_0$  does not disconnect the structure (Properties 3-4)
      then
        7:     Add  $l_0$  to  $Q_{visited}$ 
        8:     for each unvisited neighboring site  $l_i$  of  $l_0$  do
        9:       add edge  $(l_i, l_0)$  to  $map$ 
        10:      if  $\exists$  traversable edge from  $l_i$  to  $l_v \in Q_{visited}$  then
        11:        add  $l_i$  to  $Q_{frontier}$ 
        12: if  $|Q_{visited}| = |L|$  then
        13:   return  $map$ 
        14: else
        15:   return  $False$ 

```

255 from l_{START} following traversable edges. In Sec. 7, we show that the second
 256 method speeds up the process significantly.

257 *6.1. Proof of correctness*

258 This proof refers to the Properties 1-4 of a valid map, described in Sec. 3
 259 and Algorithm 1. The correctness proof is done by induction on the edges
 260 of visited locations for Properties 2-4. Property 1 follows from a gradient
 261 argument.

262 *Theorem 1, BFD-Compiler Correctness:* When the BFD compiler completes
 263 successfully, it produces a valid map.

264 *Proof of Theorem 1:*




Figure 3: BFD Compiler applied to a 3×3 structure. A) Consider l_{START} to be $(0,0)$ and l_{EXIT} to be $(2,2)$; B) the compiler removes $(2,1)$; C) $(1,2)$ cannot be removed because this would cause a disconnected structure; D) the compiler removes $(2,0)$; E) $(1,1)$ cannot be removed because of Property 1. The compiler continues in the same manner until l_{START} has been removed at which point it returns a valid map. Notice that the yellow arrows do not count towards the traversability check, but are needed for the robot rule set.

265 *Property 1:* The edge assignment adds directions in such a way that the
 266 newly added directions point from unvisited locations into visited locations
 267 (Lines 7–9). By following such a direction (when it is traversable) a robot is
 268 brought one step closer to l_{exit} . Each location can be labeled with the steps
 269 left to l_{exit} . Since the paths in the map move down the label gradient, they
 270 cannot contain cycles as that would require a path where the label increases.

271 *Properties 2-4:* The induction hypothesis (IH) is that the edges of visited
 272 locations have Properties 2-4, as well as the two axillary properties: (Prop-
 273 erty 5) $\forall l_q \in Q_{frontier} \exists$ a traversable path to the exit in the assigned map;
 274 and (Property 6) $L \setminus Q_{visited}$ is traversably path connected, i.e. all unvis-
 275 ited locations have traversable paths from l_{START} that only move over other
 276 unvisited locations.

277 *Base case:* $Q_{frontier}$ has only l_{EXIT} . Properties 2–4 are true for the empty
278 set, Property 5 is true because l_{EXIT} is path connected to itself, and Property
279 6 is correct because we assume that L is traversably connected.

280 *Induction step:* When adding another element l_0 to $Q_{visited}$, Property 2 is true
281 because the new element would only have two opposing incoming directions
282 if it had two unvisited neighbors. Property 3 is true, because when l_0 was
283 added to $Q_{frontier}$ one of its edges was directed to a location in $Q_{visited}$ (Line
284 9) and by Property 5 in IH there is a directed path toward the exit. Property
285 4 is true because of Property 6 in IH, l_0 can be reached from l_{START} and l_i
286 can be reached through l_0 after the new edge is added to the map (Line 9).
287 Property 5 is true because of (Line 10–11) and Property 3 in IH. Property 6
288 is true because of the second condition in Line 6. \square

289 Beyond proving that the compiler generates valid maps which work with
290 the TERMES system, we also believe that the reverse is true; i.e. that the
291 structure is unbuildable with the TERMES system if the compiler fails. The
292 intuition for this is as follows. The compiler fails when $Q_{frontier}$ is empty and
293 $|Q_{visited}| \neq |L|$. This happens when no more locations can be disassembled,
294 either because they are not traversable from visited locations (Property 3)
295 or because they are in between two other locations (Property 2). In other
296 words, the structure formed by unvisited locations could not have been built
297 because the last addition to the structure does not exist.

298 7. Comparison of Compilers

299 We next evaluate how the runtime of the compilers scale with the number
300 of locations for different types of structures (Fig. 4). These results are gener-
301 ated in a single process on a standard laptop (Intel(R) Core(TM) i7-4720HQ,
302 CPU @ 2.60GHz, quad core, 16G of RAM). Note that the compilers can han-
303 dle a wide range of structure types, however, for the purposes of analysis, we
304 focus only on square footprints in the following.

305 Fig. 4.A shows the runtime of each compiler as the number of locations
306 grow in a 1-height square structure. The Edge-CSP can compile such simple
307 structures with 10,000 bricks in around 100 s; for scale, a standard U.S. family
308 house contains around the same amount. As expected the Location-CSP does
309 slightly better because the constraints propagate more readily. Notice that
310 for small structures both BFD compilers compile about 10 times faster than
311 the CSP compilers. The BFD_0 compiler converges to quadratic growth (slope




Figure 4: Runtime of compilers versus the number of locations in different types of structures, including A) square, buildable structures of height 1, B) square, buildable structures of random height, C) square, unbuildable structures of random height, and D) unbuildable structures similar to that shown in Fig. 1.D. Insets indicate how we scale the number of locations; marks annotate mean of 10 runs (in the case of random height structures, 10 different structures of the same number of locations were tested); error bars indicate maximum and minimum runtime; and the number in the parenthesis gives the slope of the best fit line for all data in the curve.

312 2 in log-log axis), and as the structure size approaches 100,000 locations the
 313 CSPs will start to outperform it. This happens because their domain-variable
 314 ordering is especially optimized for these simple square structures so that the
 315 first tried assignment during the search is usually correct. By adding the im-
 316 proved connectivity check, the BFD outperforms all other compilers (scaling
 317 almost linearly) and can easily compile structures with up to 1 million bricks
 318 (comparable to the number of bricks in the Great Pyramid of Giza according
 319 to egyptorigins.com). Similar results can be noted when we run the compil-
 320 ers on buildable structures with randomly generated height profiles up to 7
 321 bricks tall (Fig. 4.B).

322 Fig. 4.C shows the runtime on unbuildable structures with randomly gen-

323 erated height profiles. The runtime of the Edge-CSP now varies significantly
324 because the search only terminates early if it finds a locally checkable error.
325 Such errors are more likely to be found with the Location-CSP compiler. The
326 new BFD compilers show a similar scalability as before. Fig. 4.D shows the
327 runtime for unbuildable structures, also presented in Fig. 1.D, which violate
328 Property 2 with any consistent ordering. This structure is especially slow
329 to search through, since ordering inconsistencies cannot be detected locally.
330 Each internal raft has four connectors, and each of these may, from the raft’s
331 perspective, be either a sink or a source. If it is a sink it violates property
332 2, and as a result all possible source combinations are tried first. We halted
333 compilations that exceeded 24 hours of runtime, which is why both CSP
334 compilers are only presented with a single data point. Notice again, how the
335 BFD₀ compiler scale quadratically with the size of the structure, and the
336 improved BFD compiler scales almost linearly.

337 8. Transition Probabilities

338 During the actual assembly of the structure, individual robots have no
339 knowledge of the system assembly state and can therefore not navigate di-
340 rectly towards the construction frontier. Instead they move along the directed
341 paths in the map at random, looking for open assembly locations. Once an
342 open location is encountered, the internal rule set on the robot, based on
343 restrictions shown in Fig. 1.C-D, determines whether or not material can be
344 added. To explain this rule set and how the combination of the map and rule
345 set affect the construction progress, we first introduce several terms related
346 to a location, l_i : 1) neighboring locations that lead to l_i are *parents* of l_i ; 2)
347 neighboring locations that lead from l_i are *children* of l_i ; 3) the *visit rate* of
348 l_i is the probability per unit time that a robot travels through it; and 4) the
349 *assembly time* of l_i is average time it takes for the robots to assemble l_i .

350 The TERMES rule set is discussed in detail in previous papers [3, 5]. To
351 give an intuitive overview, the rules restrict robots from adding material to
352 location l_i with height h_i , until (generally speaking) all parents and children
353 are of similar height, if such specified by the structure blueprint. Parents
354 of similar height ensures that robots never have to add a brick in between
355 two others (Fig. 1.C). Children of similar height ensures traversable paths
356 (Fig. 1.D). Because valid assembly steps are dependent on what bricks have
357 already been placed, this leads to wasted trips; i.e. cases where the robot
358 exits the structure before being able to deposit the brick it is carrying. The




Figure 5: A-B) 15×15 random height structure and its traversal map. C) Construction progress as a function of robot entries for 20 simulated runs using maps with uniform transition probabilities. The inset shows a snapshot from the simulation, robots are shown in green. D) Visit rates, λ , for each location in the structure, based on the map shown in B. E) Mean assembly time for each location in the structure, based on the 20 simulated runs. F) Sketch explaining how the exit location completion time, T_{EXIT} , depends on the completion time of the parent locations, T_A and T_B , and their transition probability, λ_7 and λ_8 .

359 combination of the rule set and map, generally speaking, makes tall structures
 360 grow in forward propagating staircases starting from l_{START} .

361 Take as an example the structure shown in Fig. 5.A, which has 406 bricks.
 362 To adhere with the map (Fig. 5.B) and rule set, this structure must grow from
 363 the upper- and left-most edge towards the exit. The construction progress is
 364 plotted in blue for 10 simulated runs in Fig. 5.C, where robots choose naively
 365 between children with equal probability. The visit rate is shown in Fig. 5.D.
 366 Note that this plot is based purely on the directed travel paths, and does not
 367 take the structure height into consideration. With such uniform transition
 368 probability, robots are unlikely to stay by the upper- or left-most edge of the
 369 structure, and are therefore unlikely to assemble locations that must be in
 370 place before downstream locations can be filled in. The overflow of robots in
 371 the center is also likely to cause bottlenecks, which further slows down the




Figure 6: The choice of transition probability may heavily affect completion time. A) Example of transition probabilities in a $3 \times 3 \times 1$ structure. B) Effect of β on structure completion time over 10,000 simulated runs, expressed in robot entries to the structure.

assembly progress. In Fig. 5.E, the plot of location assembly times clearly shows how the locations nearer the bottom- and right-most edge will require an excess of robots to file through the structure before they are completed. Analogous, the big vertical jumps in the traces in Fig. 5.C indicate times at which a robot fills in a perimeter location which is holding everything else up. Worst case, structure completion may be exceedingly slow - one run requires almost 200,000 robot entries before placing the last brick in the 406-brick structure. In the following text we reason about how transition probabilities between locations affect the construction process, and explore the potential for optimization.

8.1. Transition Model

First, we model traversal and assembly as a Poisson splitting process, i.e. robots visit a location with a rate λ , and if the location has two children with a probability of β and $1 - \beta$, the robot visits the two subsequent locations with rates $\beta\lambda$ and $(1 - \beta)\lambda$ respectively. If a location has two parents the rates add up. Our experiments show that the completion time of the structure is strongly dependent on locations that have small visit rates. In Fig. 6 we analyze a simple $3 \times 3 \times 1$ structure and the distribution in assembly time as a function of the single splitting parameter, β . The assembly times are long if the splitting parameter starves either the corner locations or the center of visit rates. The shape is asymmetric since the center receives rates from two parents, while starving either corner (small β s) affects the overall assembly time more dramatically.

Our goal is to choose splitting probabilities that minimize the expected assembly time of the last assembly location, T_{EXIT} . Since each location can

397 only be assembled after its parent locations have been assembled, completion
 398 time, T , for each location can be written as the maximum of the parent
 399 assembly times plus the additional assembly time due to the limited rate of
 400 visiting the location in question (Fig. 5.F). While a closed form expression for
 401 the assembly time due to rates is a simple exponential, closed form solutions
 402 for the maximum of two random variables requires integrating over their joint
 403 probabilities, for which we were unable to find an easy expression. Fig. 7
 404 shows a sampled probability density function (PDF) of the assembly time
 405 for each location in a $5 \times 5 \times 1$ structure. The data shows that the PDFs
 406 are heavy-tailed and that the PDF for two parent assembly times are not
 407 independent since they depend on a common ancestor, i.e. location (3,4)
 408 and (4,3) are not independently distributed, since they will both have a long
 409 tail if any of the location in the square (0,0) to (3,3) happened to have a long
 410 tail. Therefore, instead of trying to compute the actual assembly time
 411 to optimize the transition probabilities, we focus on finding visit rates λ_i for
 412 each location that produces small assembly times. We discuss this approach
 413 in the following subsection.

414 *8.2. Optimization of Transition Probabilities*

415 To formulate the optimization problem we assume that robots arrive at
 416 l_{START} with a rate of $\lambda_{START} = 1$. This means the visit rate for all other
 417 locations is between 0 and 1. We tested two different objective functions, one
 418 aiming for equally distributed visit rates ('equal-visit-rate') and one aiming
 419 to avoid visit rates below a certain threshold ('minimum-visit-rate'). We
 420 demonstrate our approach on the representative example structure shown in
 421 Fig. 5.A-B. The rate of visiting location l_i , λ_i , with parent locations l_j and
 422 visit rates λ_j , is calculated as:

$$\lambda_i = \sum_{j=1}^J \lambda_j P_{ji} \quad (1)$$

423 where P_{ji} denotes the probability of choosing location l_i from l_j and J is
 424 the total number of parent locations. The choices for P_{ji} have the additional
 425 constraint that $\sum_I P_{ji} = 1$ and $P_{ji} \in [0, 1]$. We formulate the optimization
 426 problem by defining the visit rate for l_i , λ_i , and the transition probabilities,
 427 P_{ji} , as variables, and by expressing Eq. 1 and conditions on P_{ij} as quality
 428 constraints and bounds on the variables.




Figure 7: Probability density function of assembly time for each location in a $5 \times 5 \times 1$ structure, measured in robot entries and generated through 5,000 simulated runs. Blue bars show the PDF for a uniform transition probability map; red bars for a transition probability map which was optimized according to a minimum-visit-rate. The vertical lines show their respective average.

429 In Fig. 5.D-E, we showed visit rates and assembly times for a map with
 430 uniform transition probability, which causes excess visits to the center of the
 431 structure leading to bottlenecks and wasted trips. This observation leads us
 432 to explore an equal-visit-rate objective. We partition the locations based on
 433 the number of steps it takes to reach l_{EXIT} (the distance along traversable
 434 paths in the map) and minimize the cost by:

$$Cost_{eq} = \sum_{dist} \sum_{I_{dist}} (\lambda_i - \mu_{dist})^2 \quad (2)$$

435 where I_{dist} is the set of locations that are equidistant from l_{EXIT} , and μ_{dist} is
 436 the mean visit rate of these. Using the previously formulated constraints and




Figure 8: Optimization of transition probabilities based on uniform visit rates of locations that are equidistant from l_{EXIT} on all (i) and permanent (ii) edges in the map. These results are based on the 15×15 random height structure shown in Fig. 5.A-B. A) System performance with maps of uniform transition probability (blue), and transition probabilities optimized according to (i) (red) and (ii) (yellow). B-C) Location assembly time as an average of 20 simulated runs, for (i) and (ii), respectively.

437 this objective we optimize over the transition probabilities P_{ij} and λ_i using
 438 sequential quadratic programming (SQP)¹. We test two maps: (i) one in
 439 which we include all edges in the map (Fig. 5.B), and (ii) one in which we
 440 only include edges in the map which are permanently traversible (i.e. without
 441 the blue arrows in Fig. 5.B). The second map is based on the intuition that
 442 paths which can be obstructed by locations that eventually reach a height
 443 difference above 1 restricts robot traversals later in the construction process.

444 The results are shown in Fig. 8. Optimizing transition probabilities ac-
 445 cording to equal visit rate makes the system perform significantly better than
 446 it does with a uniform transition probability map. The maximum completion
 447 time out of 20 simulated runs was $\sim 15,000$ robot entries when all edges were
 448 taken into consideration, and $\sim 20,000$ robot entries when only permanent
 449 edges were included in the optimization (Fig. 8.A). The latter generally per-
 450 forms worse than the former (Fig. 8.B-C). By analyzing the plot of assembly
 451 times, we suspect this decrease in performance occurs because individual low
 452 rate assignments have a disproportionate effect on the overall assembly time,
 453 and by enforcing equal visit rate for all locations in I_{dist} , we give equal weight
 454 to variations of λ_i that have minimal effect on the overall assembly time. For
 455 example, the visit rates for a good assignment is shown in Fig. 8.C. In it, the

¹SciPy implementation of `optimize.minimize` with the `SLSQP` option.

456 top left corner locations have vastly different rates within an I_{dist} set, but
457 still enable critical locations in the middle to have roughly equal visit rates.

458 In order to optimize assembly time while taking the structure height
459 into consideration, we instead propose the following minimum-visiting-rate
460 constraint:

$$Cost_{min} = \sum_I e^{\alpha(m - \lambda_i)} \quad (3)$$

461 where m is a minimum visit rate threshold defined for the entire structure
462 and α is scaling factor of how aggressively this minimum value is enforced
463 during optimization. When the visit rate is bigger than the minimum it adds
464 little to the overall cost, but locations that have smaller visit rates are heavily
465 penalized. For a given structure, we computed m to be the smallest visit rate
466 obtained during the equal-visit-rate optimization when all edges are present,
467 i.e. the smallest value that should be achievable by every location under
468 ideal conditions.

469 The results are shown in Fig. 9. The minimum-visit-rate constraint is
470 able to achieve significantly better performance than the equal-visit-rate con-
471 straint. By taking location heights and non-traversable edges into account,
472 it allows other locations to have unequal visit rates in order to feed locations
473 that are below the minimum visit rate m . This objective achieves equal visit
474 rates for the widest part of the structure and penalizes any other locations
475 with a visit rate less than m . With this new optimization, the 406-brick
476 structure is finished with an average of 1,200 robot entries. In other words,
477 every third robot entering the structure is able to deposit a brick. Notice
478 also, that the worst case completion time is much less than for any of the
479 other optimization schemes.

480 The tradeoff of assembly times between different locations can be seen in
481 Fig. 7. While the overall assembly time for location (5,5) of the optimized
482 transition probabilities decreases, optimization does increase the assembly
483 time of the center location (2,2). Basically, the optimized probabilities re-
484 allocate visit rate from the center to low rate locations in the corners. Re-
485 moving these low rate locations decreases the mean assembly time, and also
486 makes the assembly times more tightly clustered, reducing long outliers.

487 On a final note, it is clear that the idea of exploiting parallelism in con-
488 struction schemes that have a single point of entry heavily depends on the
489 optimization of transition probabilities. We can compare our approach with
490 one that produces a single path through the structure such that the structure




Figure 9: A) Optimization of transition probabilities based on a minimum-visit-rate constraint (green), as compared to an equal-visit-rate constraint (red), and uniform probabilities (blue). B-C) Visit rate and assembly time per location respectively, based on an average of 20 simulated runs for the minimum-visit-rate constraint optimization.

491 grows in a sequential manner and that every robot that enters can deposit
 492 a brick, in terms of the parallelism offered. With the minimum-visit-rate
 493 constraint optimization, you can achieve an increase in performance over
 494 this single-path approach if you can place more than just three robots with
 495 bricks at the construction frontier at any one point in time, which is true for
 496 most large scale structures. Another benefit of multi-path structures is that
 497 robots can take shortcuts to the frontier as opposed to having to travel over
 498 every location first.

499 9. Conclusion and Future Work

500 In summary we have presented work to address the scalability of the
 501 TERMES compiler, and demonstrated a BFD compiler which scales better
 502 than quadratic with the number of locations in the structure independent
 503 of whether or not that structure is buildable. We demonstrated this on
 504 structures with up to 1 million bricks, which were compiled on commodity
 505 hardware in minutes. We have further shown an approach by which the transi-
 506 tion probabilities between locations in the generated map can be improved
 507 for markedly faster construction speed without added hardware complex-
 508 ity. Future work will involve development of metrics by which to evaluate
 509 the compiled maps, especially in terms of the parallelism they offer, and
 510 compilers which can suggest modifications to make unbuildable structures
 511 buildable.

⁵¹² *Acknowledgements*

⁵¹³ This work was supported by the GETTYLAB, and partially supported
⁵¹⁴ by the National Science Foundation grant #1846340.

515 **References**

516 [1] K. Petersen, R. Nagpal, Complex Design by Simple Robots, Architectural Design (2017) 44–49.

517

518 [2] K. H. Petersen, N. Napp, R. Stuart-Smith, D. Rus, M. Kovac, A review of collective robotic construction, *Science Robotics* 4 (2019) eaau8479.

519

520 [3] J. Werfel, K. Petersen, R. Nagpal, Designing collective behavior in a termite-inspired robot construction team., *Science* 343 (2014) 754–8.

521

522 [4] K. Petersen, R. Nagpal, J. Werfel, TERMES: An autonomous robotic system for three-dimensional collective construction, *Robotics: Science and Systems Conference VII* (2011).

523

524

525 [5] J. Werfel, K. Petersen, R. Nagpal, Distributed multi-robot algorithms for the TERMES 3D collective construction system, in: In Modular Robotics Workshop, IEEE Intl. Conference on Robots and Systems (IROS).

526

527

528

529 [6] Y. Hua, Y. Deng, K. Petersen, Robots Building Bridges, Not Walls, in: IEEE International Workshops on Foundations and Applications of Self* Systems.

530

531

532 [7] Y. Deng, Y. Hua, N. Napp, K. Petersen, Scalable Compiler for the TERMES Distributed Assembly System, in: *Distributed Autonomous Robotic Systems*, Springer, 2019, pp. 125–138.

533

534

535 [8] M. S. D. Silva, V. Thangavelu, W. Gosrich, N. Napp, Autonomous Adaptive Modification of Unstructured Environments, *Robotics: Science and Systems* (2018).

536

537

538 [9] N. Napp, R. Nagpal, Robotic construction of arbitrary shapes with amorphous materials, *Proceedings - IEEE International Conference on Robotics and Automation* (2014) 438–444.

539

540

541 [10] T. Soleymani, V. Trianni, M. Bonani, F. Mondada, M. Dorigo, Autonomous construction with compliant building material, in: *Intelligent Autonomous Systems 13*, Springer, 2016, pp. 1371–1388.

542

543

544 [11] V. Lindsey, Q., Mellinger, D., Kumar, Construction of Cubic Structures with Quadrotor Teams, *Robotics: Science and Systems VII* (2011).

545

546 [12] C. Jones, M. J. Mataric, Toward a multi-robot coordination formalism,
547 Technical Report, DTIC Document, 2004.

548 [13] M. Rubenstein, A. Cornejo, R. Nagpal, Programmable self-assembly in
549 a thousand-robot swarm, *Science* 345 (2014) 795–799.

550 [14] D. Stein, T. R. Schoen, D. Rus, Constraint-aware coordinated construc-
551 tion of generic structures, in: 2011 IEEE/RSJ International Conference
552 on Intelligent Robots and Systems, IEEE, pp. 4803–4810.

553 [15] J. Seo, M. Yim, V. Kumar, Assembly planning for planar structures of a
554 brick wall pattern with rectangular modular robots, in: 2013 IEEE Inter-
555 national Conference on Automation Science and Engineering (CASE),
556 pp. 1016–1021.

557 [16] C. Coulais, E. Teomy, K. de Reus, Y. Shokef, M. van Hecke, Combin-
558 atorial design of textured mechanical metamaterials, *Nature* 535 (2016)
559 529.

560 [17] T. Tucci, B. Piranda, J. Bourgeois, A Distributed Self-Assembly Plan-
561 ning Algorithm for Modular Robots, in: Proceedings of the 17th Inter-
562 national Conference on Autonomous Agents and MultiAgent Systems,
563 International Foundation for Autonomous Agents and Multiagent Sys-
564 tems, pp. 550–558.

565 [18] T. S. Kumar, S. J. Jung, S. Koenig, A tree-based algorithm for con-
566 struction robots., in: ICAPS.

567 [19] V. Lindsey, Q., Mellinger, D., Kumar, Construction with quadrotor
568 teams, *Autonomous Robots* 33 (2012) 323–336.

569 [20] L. Matthey, S. Berman, V. Kumar, Stochastic strategies for a swarm
570 robotic assembly system, in: 2009 IEEE International Conference on
571 Robotics and Automation, IEEE, pp. 1953–1958.

572 [21] N. Napp, E. Klavins, Load balancing for multi-robot construction, in:
573 *Robotics and Automation (ICRA), 2011 IEEE International Conference*
574 on, IEEE, pp. 254–260.

575 [22] M. Schwager, J.-J. Slotine, D. Rus, Decentralized, adaptive control for
576 coverage with networked robots, in: *Proceedings 2007 IEEE Interna-*
577 *tional Conference on Robotics and Automation*, IEEE, pp. 3289–3294.

578 [23] M. Pavone, E. Frazzoli, F. Bullo, Distributed policies for equitable
579 partitioning: Theory and applications, in: 2008 47th IEEE Conference
580 on Decision and Control, IEEE, pp. 4191–4197.

581 [24] S. J. Russell, P. Norvig, Artificial intelligence: a modern approach,
582 Malaysia; Pearson Education Limited,, 2016.

583 [25] A. K. Mackworth, Consistency in networks of relations, Artificial Intel-
584 ligence 8 (1977) 99 – 118.