

The Structural Power of Reconfigurable Circuits in the Amoebot Model

Andreas Padalkin ✉ 
Universität Paderborn, Germany

Christian Scheideler ✉ 
Universität Paderborn, Germany

Daniel Warner ✉ 
Universität Paderborn, Germany

Abstract

The *amoebot model* [Derakhshandeh et al., SPAA 2014] has been proposed as a model for programmable matter consisting of tiny, robotic elements called *amoebots*. We consider the *reconfigurable circuit extension* [Feldmann et al., JCB 2022] of the geometric (variant of the) amoebot model that allows the amoebot structure to interconnect amoebots by so-called *circuits*. A circuit permits the instantaneous transmission of signals between the connected amoebots. In this paper, we examine the structural power of the reconfigurable circuits.

We start with some fundamental problems like the *stripe computation problem* where, given any connected amoebot structure S , an amoebot u in S , and some axis X , all amoebots belonging to axis X through u have to be identified. Second, we consider the *global maximum problem*, which identifies an amoebot at the highest possible position with respect to some direction in some given amoebot (sub)structure. A solution to this problem can then be used to solve the *skeleton problem*, where a (not necessarily simple) cycle of amoebots has to be found in the given amoebot structure which contains all boundary amoebots. A canonical solution to that problem can then be used to come up with a canonical path, which provides a unique characterization of the shape of the given amoebot structure. Constructing canonical paths for different directions will then allow the amoebots to set up a spanning tree and to check symmetry properties of the given amoebot structure.

The problems are important for a number of applications like rapid shape transformation, energy dissemination, and structural monitoring. Interestingly, the reconfigurable circuit extension allows polylogarithmic-time solutions to all of these problems.

2012 ACM Subject Classification Theory of computation → Distributed computing models; Theory of computation → Computational geometry

Keywords and phrases programmable matter, amoebot model, reconfigurable circuits, spanning tree, symmetry detection

Digital Object Identifier 10.4230/LIPIcs.DNA.2022.8

Funding This work has been supported by the DFG Project SCHE 1592/6-1 (PROGMATTER).

1 Introduction

The *amoebot model* [6, 4] is a well-studied model for programmable matter [15] – a substance that can be programmed to change its physical properties, like its shape and density. In the geometric variant of this model, the substance (called the *amoebot structure*) consists of simple particles (called *amoebots*) that are placed on the infinite triangular grid graph and are capable of local movements through *expansions* and *contractions*.

Inspired by the *nervous* and *muscular system*, Feldmann et al. [8] introduced a *reconfigurable circuit extension* to the amoebot model with the goal of significantly accelerating fundamental problems like leader election and shape transformation. As a first step, they



© Andreas Padalkin, Christian Scheideler, and Daniel Warner;
licensed under Creative Commons License CC-BY 4.0


28th International Conference on DNA Computing and Molecular Programming (DNA 28).

Editors: Thomas E. Ouldridge and Shelley F. J. Wickham; Article No. 8; pp. 8:1–8:22

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Feldmann et al. [8] have proposed joint movements to the amoebot model where an expanding amoebot is capable of pushing other amoebots away from it, and a contracting amoebot is capable of pulling other amoebots towards it. (a) and (b) show a joint expansion (from left to right) resp. a joint contraction (from right to left) of the yellow amoebots. The left side of (b) shows stripe $A(S, u, N)$ (see Section 1.3). The stripe can expand without causing any conflicts. (c) and (d) show exemplary conflicts. (c) If the yellow amoebots contract, the red amoebots will collide. In order to avoid the collision, the red amoebots have to contract as well. (d) If the yellow amoebots expand, the amoebot structure will tear apart. In order to maintain all connections, the red amoebots have to expand as well.

showed that leader election, consensus, compass alignment, chirality agreement, and various shape recognition problems can be solved in at most $O(\log n)$ time. This paper continues this line of work by considering a number of additional problems:

First, we consider the *stripe computation problem* where, given any connected amoebot structure S , an amoebot u in S , and some axis X , all amoebots belonging to axis X through u have to be identified. Second, we consider the *global maximum problem*, which identifies an amoebot at the highest possible position with respect to some direction in some given amoebot (sub)structure. A solution to this problem can then be used to solve the *skeleton problem*, where a (not necessarily simple) cycle of amoebots has to be found in the given amoebot structure which contains all boundary amoebots. A canonical solution to that problem can then be used to come up with a canonical path, which provides a unique characterization of the shape of the given amoebot structure. Constructing canonical paths for different directions will then allow the amoebots to set up a spanning tree and to check symmetry properties of the given amoebot structure.

The problems have a number of important applications. The stripe computation problem is important to avoid conflicts in joint amoebot contractions and expansions (see Figure 1), which is critical for rapid shape transformation. A spanning tree is an important step towards energy distribution from amoebots with access to energy to amoebots without such access [5], and canonical skeleton paths as well as symmetry checks are important for structural monitoring and repair.

1.1 Geometric Amoebot Model

In the *geometric amoebot model* [4], a set of n amoebots is placed on the infinite regular triangular grid graph $G_\Delta = (V, E)$ (see Figure 2a). An amoebot is an anonymous, randomized finite state machine that either occupies one or two adjacent nodes of G_Δ , and every node of G_Δ is occupied by at most one amoebot. If an amoebot occupies just one node, it is called *contracted* and otherwise *expanded*. Two amoebots that occupy adjacent nodes in G_Δ are called *neighbors*. Amoebots are able to move through *contractions* and *expansions*. However, since our algorithms do not make use of movements, we omit further details and refer to [4] for more information.

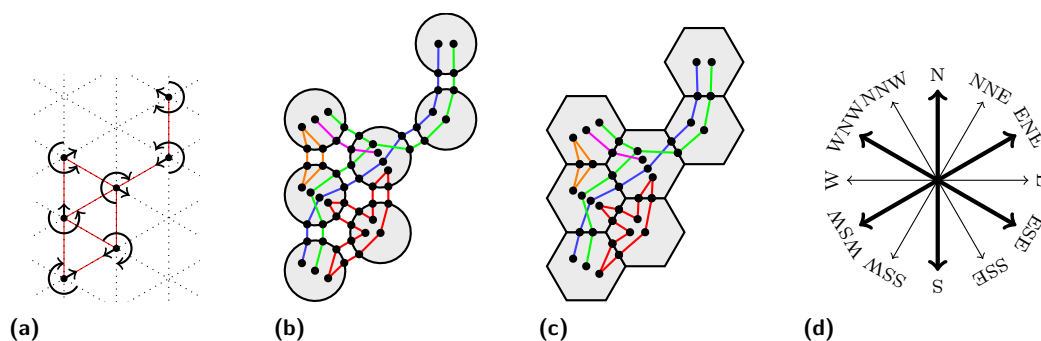


Figure 2 (a) shows an amoebot structure S . The dotted lines indicate the triangular grid G_Δ . The nodes indicate the amoebots. The arrows show their chirality and compass orientation. The red edges indicate the graph G_S . (b) and (c) show an amoebot structure with $k = 2$ external links between neighboring amoebots. The amoebots are shown in gray. The nodes on the boundary are the pins, and the ones within the amoebots the partition sets. An edge between a partition set Q and a pin p implies $p \in Q$. Each color indicates another circuit. (a) and (b) are taken from [8]. (d) shows the cardinal directions. The thick arrows indicate the cardinal directions along the main axes, and the thin ones the cardinal directions perpendicular to the main axes.

Each amoebot has a compass orientation (it defines one of its incident edges as the northern direction) and a chirality (a sense of clockwise or counterclockwise rotation) that it can maintain as it moves, but initially the amoebots might not agree on their compass orientation and chirality. In this paper, we assume that all amoebots share a common compass orientation and chirality. This is reasonable since Feldmann et al. [8] showed that all amoebots are able to come to an agreement within the considered extension (see Section 1.4).

Let the *amoebot structure* $S \subseteq V$ be the set of nodes occupied by the amoebots. By abuse of notation, we identify amoebots with their nodes. We say that S is *connected* iff G_S is connected, where $G_S = G_\Delta|_S$ is the graph induced by S . In this paper, we assume that initially, S is connected and all amoebots are contracted. Also, we assume the fully synchronous activation model, i.e., the time is divided into synchronous rounds, and every amoebot is active in each round. On activation, each amoebot may perform a movement and update its state as a function of its previous state. However, if an amoebot fails to perform its movement, it remains in its previous state. The time complexity of an algorithm is measured by the number of synchronized rounds required by it.

1.2 Reconfigurable Circuit Extension

In the *reconfigurable circuit extension* [8], each edge between two neighboring amoebots u and v is replaced by k edges called *external links* with endpoints called *pins*, for some constant $k \geq 1$ that is the same for all amoebots. For each of these links, one pin is owned by u while the other pin is owned by v . In this paper, we assume that neighboring amoebots have a common labeling of their incident external links.

Each amoebot u *partitions* its *pin set* $P(u)$ into a collection $\mathcal{Q}(u)$ of pairwise disjoint subsets such that the union equals the pin set, i.e., $P(u) = \bigcup_{Q \in \mathcal{Q}(u)} Q$. We call $\mathcal{Q}(u)$ the *pin configuration* of u and $Q \in \mathcal{Q}(u)$ a *partition set* of u . Let $\mathcal{Q} = \bigcup_{u \in S} \mathcal{Q}(u)$ be the collection of all partition sets in the system. Two partition sets are *connected* iff there is at least one external link between those sets. Let L be the set of all connections between the partition sets in the system. Then, we call $H = (\mathcal{Q}, L)$ the *pin configuration* of the system and any connected component C of H a *circuit* (see Figure 2b). Note that if each partition set of \mathcal{Q}

is a *singleton*, i.e., a set with exactly one element, then every circuit of H just connects two neighboring amoebots. However, an external link between the neighboring amoebots u and v can only be maintained as long as both u and v occupy the incident nodes. Whenever two amoebots disconnect, the corresponding external links and their pins are removed from the system. An amoebot is part of a circuit iff the circuit contains at least one of its partition sets. A priori, an amoebot u may not know whether two of its partition sets belong to the same circuit or not since initially it only knows $\mathcal{Q}(u)$.

Each amoebot u can send a primitive signal (a *beep*) via any of its partition sets $Q \in \mathcal{Q}(u)$ that is received by all partition sets of the circuit containing Q at the beginning of the next round. The amoebots are able to distinguish between beeps arriving at different partition sets. More specifically, an amoebot receives a beep at partition set Q if at least one amoebot sends a beep on the circuit belonging to Q , but the amoebots neither know the origin of the signal nor the number of origins. Note that beeps are enough to send whole messages over time, especially between adjacent amoebots. We modify an activation of an amoebot as follows. As a function of its previous state and the beeps received in the previous round, each amoebot may perform a movement, update its state, reconfigure its pin configuration, and activate an arbitrary number of its partition sets. The beeps are propagated on the updated pin configurations. If an amoebot fails to perform its movement, it remains in its previous state and pin configuration, and does not beep on any of its partition sets.

In this paper, we will utilize the dual graph of the triangular grid graph, i.e., a hexagonal tessellation, to visualize amoebot structures (see Figure 2c). Thereby, we reduce each external link to a single pin. Furthermore, in order to improve the comparability of circuit configurations, we add pins to each side of the hexagon.

1.3 Problem Statement and Our Contribution

Let $D_m = \{N, ENE, ESE, S, WSW, WNW\}$ be the set of all cardinal directions along the axes of G_Δ , and $D_p = \{E, SSE, SSW, W, NNW, NNE\}$ the set of all cardinal directions perpendicular to the axes of G_Δ (see Figure 2d). In the following, we state the considered problems.

First, we consider the *stripe computation problem*. Let $X(v, d) \subseteq V$ denote the nodes of G_Δ that lie on the axis through the node $v \in V$ into the cardinal direction $d \in D_m \cup D_p$. For $R \subseteq V$, we call the set $A(R, v, d) = R \cap X(v, d)$ a *stripe* of R (see Figure 1b). Note that a stripe is not necessarily connected. Let an amoebot $u \in S$ and a cardinal direction $d \in D_m \cup D_p$ be given, i.e., each amoebot $v \in S$ knows the cardinal direction d and whether $v = u$. The goal of each amoebot $v \in S$ is to determine whether $v \in A(S, u, d)$. Our *stripe algorithm* solves the stripe computation problem after $O(\log n)$ rounds.

Second, we consider the *global maxima problem*. Let a cardinal direction $d \in D_m \cup D_p$ and a non-empty set $R \subseteq S$ be given, i.e., each amoebot $v \in S$ knows the direction d and whether $v \in R$. The goal of each amoebot $v \in S$ is to determine whether $v \in \operatorname{argmin}_{w \in R} f_d(R, w)$ where $f_d(R, w)$ denotes the number of amoebots in R that lie in direction d from amoebot w . We call $\operatorname{argmin}_{w \in R} f_d(R, w)$ the set of global maxima of R with respect to d . Our *global maxima algorithm* solves the global maxima problem after $O(\log^2 n)$ rounds w.h.p.¹

Third, we consider the (*canonical*) *skeleton problem*. An amoebot u is a *boundary amoebot* iff it is adjacent to an unoccupied node in $V \setminus S$. Otherwise, we call u an *inner amoebot*. A (potentially non-simple) cycle of amoebots is a *skeleton* iff the cycle contains all boundary

¹ An event holds *with high probability (w.h.p.)* if it holds with probability at least $1 - 1/n^c$ where the constant c can be made arbitrarily large.

amoebots in S . Note that the skeleton may contain inner amoebots. An amoebot structure computes a skeleton C iff each amoebot knows its predecessor and successor for each of its occurrences in C . The goal of the skeleton problem is to compute an arbitrary skeleton.

Since skeletons are not unique, we define a *canonical skeleton* with respect to a cardinal direction $d \in D_m \cup D_p$ and a sign $s \in \{+, -\}$ (abbreviated as (d, s) -skeleton). We defer the definition to Section 4.1. Let a cardinal direction $d \in D_m \cup D_p$ and a sign s be given, i.e., each amoebot $v \in S$ knows the cardinal direction d and the sign s . The goal of the canonical skeleton problem is to compute the canonical skeleton. Our *canonical skeleton algorithm* solves the (canonical) skeleton problem after $O(\log^2 n)$ rounds w.h.p.

Our algorithms for the remaining problems are based on skeletons. However, they split the skeletons into paths that we call *skeleton paths*. For the canonical skeletons, we define a *canonical skeleton path* by specifying a splitting point. Our canonical skeleton algorithm determines this point in parallel to the computation of the canonical skeleton.

Fourth, we consider the *spanning tree problem*. A *tree* is a cycle-free and connected graph. A *spanning tree* of an amoebot structure S is a tree $T = (S, E_T)$ with $E_T \subseteq E$. An amoebot structure computes a spanning tree T if each amoebot $u \in S$ knows whether $\{u, v\} \in E_T$ for each neighbor $v \in N(u)$. The goal of the spanning tree problem is to compute an arbitrary spanning tree. Our *spanning tree algorithm* solves the spanning tree problem after $O(\log^2 n)$ rounds w.h.p.

Finally, we consider the *symmetry detection problem*. The goal of that problem is to determine whether the amoebot structure features rotational or reflection symmetries. Our *symmetry detection algorithm* solves the the problem after $O(\log^5 n)$ rounds w.h.p.

1.4 Related Work

The reconfigurable circuit extension was introduced by Feldmann et al. [8]. Note that the amoebot model [1] and its circuit extension [11, 12, 13, 14] can in principle be realized.

Feldmann et al. [8] have proposed solutions for *leader election* (see Section 2), *consensus*, *compass alignment*, *chirality agreement*, and various *shape recognition* problems. Both, the alignment of the compasses and the agreement on a chirality requires $O(\log n)$ w.h.p. This makes our assumption of a common compass orientation and chirality reasonable.

To our knowledge the stripe computation, global maxima, (canonical) skeleton, and symmetry detection problem have not been considered within the standard amoebot model. However, regarding the global maxima problem, Daymude et al. [2] have considered the related problem of determining the dimensions of an object (a finite, connected, static set of nodes) in order to solve various *convex hull problems*. Their approach can be easily adjusted to compute the global maxima of the amoebot structure. However, it requires $O(n)$ rounds.

The spanning tree problem is widely studied in the distributed algorithms community, e.g., [10] (also see the cited papers within the reference). The *spanning tree primitive* is one of the most used techniques to move amoebots, e.g., [3, 2, 7, 9]. Beyond that, spanning trees were applied to distribute energy [5]. However, the construction requires $\Omega(D)$ rounds where D is the diameter of the amoebot structure (e.g., see [3]). Since $D = \Omega(\sqrt{n})$, our solution is a significant improvement.

2 Preliminaries

In this section, we enumerate important primitives given in previous papers.

Global Circuit. If each amoebot partitions its pins into one partition set, we obtain a single circuit that interconnects all amoebots. We call this circuit the *global circuit* [8].

Leader Election. We make use of a generalized version of the leader election algorithm proposed by Feldmann et al. [8]:

► **Theorem 1.** *Let C_1, \dots, C_m be sets of candidates. For each $i \in \{1, \dots, m\}$, let C_i be the circuit that connects all candidates of set C_i . Let $C_i \cap C_j = \emptyset$ hold for all $i \neq j$. An amoebot structure elects a leader from each set of candidates after $\Theta(\log n)$ rounds w.h.p.*

In the classical leader election problem, the amoebot structure S has to elect a leader only from the set $C_1 = S$.

Chains. We call an ordered sequence $C = (u_0, \dots, u_{m-1})$ of m amoebots a *chain* iff (i) all subsequent amoebots u_i, u_{i+1} are neighbors, (ii) each amoebot in C except u_0 knows its predecessor, and (iii) each amoebot in C except u_{m-1} knows its successor.


Boundary Sets. We adopt the definition for the *boundary sets* from [7]: Consider $G_{V \setminus S} = G_{\Delta} |_{V \setminus S}$. The connected components of $G_{V \setminus S}$ are called empty regions. The number of empty regions is finite since S is finite. Let R_1, \dots, R_m denote the empty regions. For $i \in \{1, \dots, m\}$, the boundary set B_i is the neighborhood of R_i in S , i.e., $B_i = \{u \in S \mid \exists v \in R_i : \{u, v\} \in E\}$. There is exactly one infinite empty region since S is finite. We call the corresponding boundary set the *outer boundary set*, and the others *inner boundary sets*. Naturally, each boundary set can be represented as a (potentially non-simple) cycle.

In order to distinguish inner and outer boundaries, we apply the *inner outer boundary test* by Derakhshandeh et al. [7]. They accumulate the angles of the turns while traversing the cycle of the boundary set once. An outer boundary set results in a sum of 360° , and an inner boundary set results in a value of -360° . Note that it is sufficient to count the turns by 60° modulo 5. This allows us to accumulate the sum with constant memory. However, the traversing requires $O(n)$ rounds. We accelerate the accumulation by the following result by Feldmann et al. [8]:

► **Theorem 2.** *Let $C = (v_0, \dots, v_{m-1})$ be a chain within the amoebot structure where $m \in \mathbb{N}$ denotes the length of the chain. Let $k \in \mathbb{N}$ be constant. Let $x_i \in \{0, \dots, k-1\}$ for all $i \in \{0, \dots, m-1\}$. Suppose that for each $i \in \{0, \dots, m-1\}$, amoebot v_i knows the value x_i . Then, the chain computes $x = \sum_{i \in \{0, \dots, m-1\}} x_i \pmod k$ after $O(\log m)$ rounds.*

► **Corollary 3.** *A boundary set can determine whether it is an inner boundary set or the outer boundary set within $O(\log n)$ rounds w.h.p.*

Synchronization of Procedures. Sometimes, we want to execute a procedure on different subsets in parallel, e.g., we apply the inner outer boundary test on all boundary sets at once. The execution of the the same procedure may take different amounts of rounds for each subset. The amoebots of a single subset are not able to decide when all subsets have terminated. In order to synchronize the amoebot structure with respect to the procedures, the amoebots periodically establish the global circuit. The amoebots of each subset that has not yet terminated beep on that circuit. The amoebot structure may proceed to the next procedure once the global circuit was not activated.



■ **Figure 3** PASC algorithm. Each figure shows the chain at the beginning of the i -th iteration. The circles indicate the primary (P) and secondary (S) partition sets. The blue bordered amoebot denotes the reference amoebot u_r . Yellow amoebots are active, and gray amoebots are passive. All primary and secondary partition sets that are part of the primary resp. secondary circuit of u_r are depicted in red resp. cyan.

3 Computing Identifiers

In this section, we assign *identifiers* in \mathbb{Z} to the amoebots. Let (x_{k-1}, \dots, x_0) denote the *two's complement representation* of $-x_{k-1} \cdot 2^{k-1} + \sum_{i=0}^{k-2} x_i \cdot 2^i$. By abuse of notation, we identify the identifiers with their two's complement representation. Each amoebot computes a two's complement representation of its identifier. In the first subsection, we compute successive identifiers along chains. In the second subsection, we compute spatial identifiers with respect to a cardinal direction. In the third subsection, we show two applications for the identifiers.

3.1 Successive Identifiers along the Chain

In this section, we compute successive identifiers along a chain with respect to an amoebot that we call the *reference amoebot*. Let $C = (u_0, \dots, u_{m-1})$ be a chain of amoebots. Let u_r be an arbitrary reference amoebot within the chain, e.g., chosen by position (for example $r = 0$), or by a leader election. We assign identifiers id_{C, u_r} according to the following rules:

$$\text{id}_{C, u_r}(u_r) = 0 \tag{1}$$

$$\text{id}_{C, u_r}(u_{i+1}) = \text{id}_{C, u_r}(u_i) + 1 \text{ for } 1 \leq i < m \tag{2}$$

Note that $\text{id}_{C, u_r}(u_i) = i - r$.

In order to compute the identifiers, we utilize a procedure on the chain of amoebots proposed by Feldmann et al. [8] that we henceforth refer to as the *primary and secondary circuit algorithm* (PASC algorithm). Originally, the algorithm has been used as a subroutine for Theorem 2.

In the following, we explain how the PASC algorithm works (see Figure 3). So, let $C = (u_0, \dots, u_{m-1})$ be a chain of m amoebots. If an amoebot occurs multiple (but constantly many) times, it operates each position independently of each other which is possible by using sufficiently many pins. Each amoebot is either *active* or *passive*. Initially, each amoebot is active. The algorithm iteratively transforms active amoebots into passive amoebots while keeping amoebot u_r active.

At the beginning of an iteration, the amoebots establish two circuits as follows. Each amoebot has two partition sets that we call the *primary and secondary partition set*. We connect the primary and secondary partition sets by the following two rules² (see Figure 3): (i) If an amoebot is active, we connect its primary partition set to the secondary partition set of its predecessor, and its secondary partition set to the primary partition set of its predecessor. (ii) If an amoebot is passive, we connect its primary partition set to the primary partition set of its predecessor, and its secondary partition set to the secondary partition set of its predecessor.

We obtain two circuits through all amoebots (see Figure 3). Each amoebot defines the circuit containing its primary partition set as its *primary circuit*, and the circuit containing its secondary partition set as its *secondary circuit*. Clearly, we obtain two disjoint circuits along the chain. We refer to [8] for a detailed construction.

After establishing these circuits, the iteration utilizes two rounds. In the first round, amoebot u_r activates its primary circuit. Each active amoebot that has received the beep on its secondary circuit beeps in the second round on its secondary circuit. These amoebots become passive amoebots in the next iteration. The algorithm terminates when the second round is silent. At this point, amoebot u_r is the remaining active amoebot. Feldmann et al. [8] have proven the following lemma:

► **Lemma 4.** *The PASC algorithm terminates in $\lceil \log m \rceil$ iterations, resp. $O(\log m)$ rounds.*

We obtain the identifiers from the PASC algorithm as follows. Let $k = \lceil \log m \rceil$ be the number of iterations. For $0 \leq i < k$, let r_i be the first round of the $(i + 1)$ -st iteration. Note that in each of these rounds, each amoebot of the chain receives a beep either on its primary circuit or its secondary circuit. Hence, in round r_i , amoebot x interprets a beep on the primary circuit as $x_i = 0$ and a beep on the secondary circuit as $x_i = 1$.

► **Lemma 5.** *The PASC algorithm computes id_{C,u_r} .*

3.2 Spatial Identifiers

In this section, we compute identifiers relative to the spatial positions of the amoebots with respect to a cardinal direction $d \in D_m \cup D_p$ and a reference amoebot $u_r \in S$. Let d' denote the direction obtained if we rotate d by 90° counterclockwise, e.g., $d' = N$ for $d = E$. First, consider $d \in D_p$ (see Figure 4a). Let $\mathcal{A}_d = \{A(S, v, d') \mid v \in S\}$ denote a set of stripes. We first assign identifiers to \mathcal{A}_d . Afterwards, we extend these identifiers to the nodes in S .

Observe that \mathcal{A}_d partitions S into disjoint stripes. These stripes form a chain \mathcal{C}_d if we think of the stripes as nodes such that two nodes are adjacent if the corresponding stripes are neighbors (see Figure 4a). The order of the chain is given by the cardinal direction d : The successor of a stripe is the neighboring stripe in direction d . Let $\text{succ}(A)$ denote the succeeding stripe of stripe A . Let $A_r = A(S, u_r, d')$. We assign identifiers $\text{id}_{\mathcal{C}_d, A_r}$ according to the following two rules:

$$\begin{aligned} \text{id}_{\mathcal{C}_d, A_r}(A_r) &= 0 \\ \text{id}_{\mathcal{C}_d, A_r}(\text{succ}(A)) &= \text{id}_{\mathcal{C}_d, A_r}(A) + 1 \text{ for all } A \in \mathcal{A}_d \end{aligned}$$

Finally, we define $\text{id}_{d, u_r}(v) = \text{id}_{\mathcal{C}_d, A_r}(A(S, v, d'))$ for all nodes $v \in S$.

² In comparison to [8], we have adjusted the rules by mirroring the connections.




Figure 4 Spatial identifiers. (a) show the spatial identifiers with respect to $d = E$, and (b) with respect to $d = N$. Each color indicates a stripe (in S). The white hexagons indicate the neighborhood of S . The thick boundary indicates the reference amoebot u_r . (c) If we do not include the neighborhood of S , then amoebot u is unable to determine whether amoebot v belongs to the preceding stripe (orange) or the stripe preceding its preceding stripe (yellow).


In order to compute the identifiers, we transfer the concept of primary and secondary circuits from a chain of amoebots to a chain of stripes (compare with Section 3.1): From a global perspective, each stripe knows its predecessor and its successor, is either active or passive, and has a primary and secondary partition set. The primary and secondary partition sets are connected by the following rules: If a stripe is active, its primary partition set is connected to the secondary partition set of its predecessor, and its secondary partition set is connected to the primary partition set of its predecessor. If a stripe is passive, its primary partition set is connected to the primary partition set of its predecessor, and its secondary partition set is connected to the secondary partition set of its predecessor. There are no further connections. In order to keep the amoebots within a stripe synchronized, we additionally require that each amoebot has access to the primary and secondary partition set of its stripe. In the following, we explain how to construct the circuits that satisfy the aforementioned properties.

Note that the neighborhood of each amoebot only contains amoebots of the same stripe, the preceding stripe, and the succeeding stripe. Since we assume common compass orientation and chirality, each amoebot is able to determine to which stripe each neighbor belongs. We now define pin configurations that satisfy the aforementioned properties. Each pin configuration has two partition sets that we call the primary and secondary partition set. We connect the primary and secondary partition sets of an amoebot to the primary and secondary partition sets of adjacent amoebots such that the aforementioned properties are reflected locally. For example, consider two adjacent amoebots u and v such that u belongs to an active stripe and v belongs to u 's preceding stripe. We connect u 's primary partition set to v 's secondary partition set, and u 's secondary partition set to v 's primary partition set. Since we assign the same identifiers to amoebots of the same stripe, we connect their primary and secondary partition sets, respectively. We obtain two pin configurations: one for amoebots of active stripes and one for amoebots of passive stripes (see Figure 5).

► **Lemma 6.** *The construction satisfies the aforementioned properties.*

Now, we simply apply the PASC algorithm on the chain of stripes to compute $\text{id}_{C_{d,A_r}}$, i.e., each amoebot $v \in S$ computes $\text{id}_{C_{d,A_r}}(A(S, v, d'))$ that equals $\text{id}_{d,u_r}(v)$ by definition.

► **Lemma 7.** *The PASC algorithm computes id_{d,u_r} for $d \in D_p$.*



■ **Figure 5** Utilized pin configurations. We utilize the former two pin configurations for $d = E$. We utilize the latter four pin configurations for $d = N$. The first argument denotes the state of the amoebot’s stripe and the second argument denotes the state of the preceding stripe, respectively.

Next, consider $d \in D_m$. We discuss the necessary modifications in comparison to $d \in D_p$. Note that an amoebot is unable to locally determine to which stripe its neighbors belong (see Figure 4c). We therefore perform the PASC algorithm on the union of the amoebot structure and its neighborhood (see Figure 4b). Note that each neighbor of an amoebot $v \in S$ belongs either to one of the two preceding stripes or to one of the two succeeding stripe.

Each amoebot tracks when the stripes of its neighbors become passive as follows. Recall that all stripes are initially active. Suppose that each amoebot $v \in S$ knows whether (its stripe and) the stripes of its neighbors are active or passive at the beginning of an iteration of the PASC algorithm. Hence, v also knows how these stripes are interconnected. Thus, v can conclude from the signal it receives on the signals received by the stripes of its neighbors and with that whether these become passive.

There are two reasons for the tracking. First, some of the stripes are not occupied by any amoebots. The tracking allows each amoebot $v \in S$ to activate the correct circuits for each of its neighbors. Second, the connections between an amoebot and its neighbor of the stripe preceding its preceding stripe depends on the states of its stripe and its preceding stripe. We obtain four pin configurations (see Figure 5).

► **Lemma 8.** *The PASC algorithm computes id_{d,u_r} for $d \in D_m$.*

Note that we can generalize this technique to arbitrary directions as long as the identifiers within a neighborhood only differ by some constant.

3.3 Applications

We now consider two applications for the identifiers, namely the stripe problem and the global maxima problem.

First, consider the stripe problem. Our stripe algorithm simply executes the PASC algorithm with u as the reference amoebot, i.e., $u_r = u$. Note that this sets the identifier of u to 0, i.e., $\text{id}_{d,u}(u) = 0$. By construction, $\text{id}_{d,u}(v) = \text{id}_{d,u}(u) = 0$ holds for all $v \in A(S, u, d)$. We obtain the following theorem.

► **Theorem 9.** *The stripe algorithm solves the stripe computation problem in $O(\log n)$ rounds.*

Next, consider the global maxima problem. Recall that we compute the global maxima of a set $R \subseteq S$ (see Section 1.3). By construction, $\text{argmin}_{w \in R} f_d(R, w) = \text{argmax}_{w \in R} \text{id}_{d,u_r}(w)$ holds for any reference amoebot u_r . The idea of our global maxima algorithm is therefore to execute the PASC algorithm and to determine the highest identifier. By choosing an $u \in R$ as the reference amoebot, i.e., $u_r = u$, we ensure that the maximal identifier is non-negative. In order to determine the maximum of non-negative numbers, we apply the consensus algorithm by Feldmann et al. [8] that agrees on the highest input value. First, the amoebot structure

establishes the global circuit (see Section 2). Each amoebot $v \in R$ with $\text{id}(v) \geq 0$ transmits its identifier starting from the most significant bit. If a transmitting amoebot observes a beep in a round it does not beep, it stops its transmission. Only an amoebot with the highest identifier is able to transmit its identifier until the end.

However, since each amoebot can only store a constant section of its identifier, and since the PASC algorithm provides the identifiers from the least significant bit to the most significant bit, we have to partially recompute the identifiers after each bit. In order to identify the correct bit, we simply use two binary counters where the first counter indicates the current bit, and the second counter the current iteration of the PASC algorithm. These can be realized along a chain such that with the help of circuits, the incrementation, the decrementation, and the comparison only require $O(1)$ rounds. Using a chain along the outer boundary set ensures a sufficient length of the counters. We obtain the following theorem.

► **Theorem 10.** *The global maxima algorithm computes the global maxima in $O(\log^2 n)$ rounds w.h.p.*

4 Skeletons

This section deals with (canonical) skeletons. In the first subsection, we canonicalize and construct the skeletons. In the second and third subsection, we show two applications for skeletons by showing how to construct spanning trees and how to detect symmetries.


4.1 Canonicalized Construction

The general idea to construct a skeleton is to start with the cycles given by the boundary sets and to fuse these into a single cycle, i.e., a skeleton (see Figures 6 and 14). In order to fuse two boundary cycles, we have to determine a path between them. The boundary cycles are split at the respective endpoints of the path and connected along the path. Note that the fused cycle uses the path twice. The difficulty lies in finding paths between the boundary cycles and in avoiding the creation of new cycles. The construction is correct, i.e., we obtain a single cycle, iff the boundary sets and paths form a tree with the boundary sets as the nodes and the paths as the edges. In order to obtain a skeleton path, the skeleton is split at an arbitrary point, e.g., chosen by a leader election.

We now canonicalize the construction of a skeleton and a skeleton path. Recall that we define the canonical skeleton (path) with respect to a cardinal direction $d \in D_m \cup D_p$ and a sign $s \in \{+, -\}$. For the canonical skeleton, we have to define how the paths are constructed, and how the boundary cycles and paths are exactly connected. For the canonical skeleton path, we have to define the point to split the canonical skeleton. Afterwards, we show how the amoebot structure computes the canonical skeleton in a distributed fashion.

We start with the construction of the canonical skeleton. Let $\rho_s(d, x)$ denote the direction obtained if we rotate direction d by x degrees counterclockwise if the sign s is positive, and clockwise if the sign s is negative. Let $d_p = d$ if $d \in D_m$ and $d_p = \rho_s(d, 30)$ if $d \in D_p$. For each inner boundary set B , we construct a path as follows. First, we compute the global maxima of B with respect to direction d . Let B_d denote these global maxima. Then, we compute the global maximum of B_d with respect to direction $\rho_s(d, 90)$. Let u_B denote the global maximum.

► **Lemma 11.** *Amoebot u_B is adjacent to exactly one node in R , namely the one in direction $\rho_s(d_p, 180)$. Further, no amoebot in direction d_p of u_B is adjacent to a node in R .*



■ **Figure 6** (a) shows the initial situation. The red lines indicate the boundary cycles. (b) shows the $(N, +)$ -skeleton. The yellow and orange amoebots indicate the global maxima of the boundary sets. The orange amoebots indicate the starting points of the paths. The blue lines indicate the paths between the boundary cycles. The node indicates the location where the cycle is split. (c) shows the spanning tree obtained from the $(N, +)$ -skeleton. The node indicates the root.

The path starts at u_B and goes straight in direction d_p until it reaches an amoebot v_B of another boundary set. The existence of v_B is guaranteed by the outer boundary set. Clearly, all nodes of the path are occupied by amoebots. Note that the path may be trivial, i.e., $u_B = v_B$. There is only a single case where v_B is part of two boundary sets unequal to B . In this case, we take the boundary of v_B in direction $\rho_s(d_p, 60)$.

▶ **Lemma 12.** *The boundary sets and paths form a tree.*


It remains to define how the cycles and paths are exactly connected. We define that the cycle runs along the tree without crossing itself.

Next, consider the construction of the canonical skeleton path. We determine the splitting point u_{B_O} by applying the same procedure as for the starting points of the paths on the outer boundary set B_O . That is, we first compute the global maxima of B_O with respect to direction d , and then compute the global maximum of these global maxima with respect to direction $\rho_s(d, 90)$. If the canonical skeleton visits u_{B_O} multiple times, we pick a predefined position with respect to d_p (see Figure 11).

▶ **Lemma 13.** *The canonical skeleton (path) visits each bond at most twice. Thus, the canonical skeleton has linear complexity.*

In the following, we present our canonical skeleton algorithm that computes the canonical skeleton and the splitting point for the canonical skeleton path in parallel. Some instructions are performed on different subsets in parallel. In order to keep the amoebot structure synchronized, we apply the synchronization primitive (see Section 2). In a preprocessing step, each boundary set determines whether it is an inner or outer boundary set (see Corollary 3).

The canonical skeleton algorithm follows our construction of the canonical skeleton. In the first step, we compute the starting points of the paths and the splitting point by performing the global maxima algorithm on each boundary set B with respect to direction d , and on each resulting set B_d with respect to direction $\rho_s(d, 90)$. However, the computation of the boundary sets may interfere with each other since the boundary sets may intersect. In order to circumvent that problem, we add two additional external links and let each boundary set use the two external links closer to the corresponding empty region (see Figures 7a to 7c). Note that an amoebot is not able to determine whether two adjacent nodes belong to the



■ **Figure 7** Computation of the global maxima of each boundary set with respect to $d = N$. (a) and (b) show the original construction for the outer and inner boundary set, respectively. The gray amoebot does not participate in the computation for the inner boundary set. (c) shows the construction for the computation in parallel. Amoebot u is unaware that two of its adjacent nodes belong to the same empty region. (d) shows the circuit utilized to identify the paths. The yellow amoebots are boundary amoebots. The gray amoebots are inner amoebots.

same empty region (see Figure 7c). Hence, it can only construct the primary and secondary circuits along the cycle. Subsequently, it handles each of its occurrences within the cycle separately. Nonetheless, the adjusted construction still satisfies the necessary properties given in Section 3.2.

The second step is the computation of the paths from the starting points straight into direction d_p . The canonical skeleton algorithm proceeds as follows. Each inner amoebot connects all pins of its neighbors in directions d_p and $\rho_s(d_p, 180)$, and each boundary amoebot connects all pins to its neighbors in direction d_p and $\rho_s(d_p, 180)$, respectively (see Figure 7d). Each starting point without a second boundary activates the circuit to its neighbor in direction d_p . Each amoebot that receives a beep is part of a path from the starting point straight into direction d_p . Finally, we obtain the following theorem:

► **Theorem 14.** *The canonical skeleton algorithm computes a (canonical) skeleton (path) in $O(\log^2 n)$ rounds w.h.p.*

4.2 Spanning Tree

We now show how a skeleton can be utilized to construct a spanning tree. We assume that we have already computed a (not necessarily canonical) skeleton (see Section 4.1). Our spanning tree algorithm consists of two phases. We first outline the goal of each phase. In the first phase, we construct a tree spanning all amoebots of the skeleton path. In the second phase, we add the remaining amoebots to the tree. Now, consider the first phase. We make use of the following lemma.

► **Lemma 15.** *Let $G = (V, E)$ be a connected graph. Let $\pi = (v_1, \dots, v_m)$ be a path in G . Let $V' \subseteq V$ be the set of all amoebots on the path π . Let $\pi(v)$ denote the first edge in π incident to v . Then, $T = (V', E')$ with $E' = \bigcup_{v \in V' \setminus \{v_1\}} \{\pi(v)\}$ is a tree.*

In order to determine the first occurrence of each amoebot, we apply the PASC algorithm on the path with v_0 as the reference amoebot (see Section 3.1). Each amoebot is able to determine its first occurrence by simply comparing the identifiers of all its occurrences. Each amoebot notifies the predecessor of its first occurrence.

Next, consider the second phase. For each amoebot v not included in the skeleton $S \setminus V'$, we add an edge from it to its northern neighbor w to the spanning tree. Note that v is an inner amoebot such that w has to exist. Otherwise, v would be included in the skeleton. Each amoebot $v \in S \setminus V'$ notifies its northern neighbor. We obtain the following theorem:

► **Theorem 16.** *Given a skeleton, the spanning tree algorithm computes a spanning tree after $O(\log n)$ rounds. Altogether, it requires $O(\log^2 n)$ rounds w.h.p.*

4.3 Symmetry Detection

We now show how to detect rotational symmetries and reflection symmetries. Due to the underlying infinite regular triangular grid graph G_Δ , there is only a limited number of possible symmetries. More precisely, an amoebot structure can only be 2-fold, 3-fold or 6-fold rotationally symmetric, and reflection symmetric to axes in a direction of $D_m \cup D_p$. Moreover, the problem is complicated by the facts that the symmetry point may be an unoccupied node of G_Δ or not a node of G_Δ at all, and that the symmetry axis may not be occupied by any amoebots.

Recall that we define a canonical skeleton by two parameters: the direction d and the sign s . Note that rotating the direction results in a rotated construction, and inverting the sign results in a reflected construction. Hence, a symmetric amoebot structure implies a symmetric construction of canonical skeletons. The idea of our symmetry detection algorithm is therefore to compare the canonical skeletons. We compare the skeletons according to the following observations (compare Figure 8).


► **Observation 17.** *An amoebot structure is 2-fold rotationally symmetric if the $(N, +)$ -skeleton and the $(S, +)$ -skeleton are symmetric. An amoebot structure is 3-fold rotationally symmetric if the $(N, +)$ -skeleton and the $(ESE, +)$ -skeleton are symmetric. An amoebot structure is 6-fold rotationally symmetric if it is 2-fold and 3-fold rotationally symmetric.*

Let $d \in D_m \cup D_p$ and let d' denote the direction obtained if we rotate d by 90° counterclockwise. An amoebot structure is reflection symmetric to an axis in direction d if the $(d', +)$ -skeleton and the $(d', -)$ -skeleton are symmetric. Note that due to symmetry, it is enough to only check half of $D_m \cup D_p$.

In order to compare two canonical skeletons, we map each canonical (d, s) -skeleton path to a unique bit string by having each amoebot on the path store a partial bit string of constant length encoding the direction of its successor relative to direction d and sign s . Consequently, the comparison of two canonical skeletons is reduced to the comparison of the corresponding bit strings of the two skeletons. In the following we show how such a comparison of two bit strings is possible in polylogarithmic time.

To this end, we consider the *string equality problem*: Let $A = (A_0, \dots, A_{m-1})$ and $B = (B_0, \dots, B_{m'-1})$ be two chains of amoebots with reference amoebots A_0 and B_0 , holding bit strings $a = (a_0, \dots, a_{m-1})$ and $b = (b_0, \dots, b_{m'-1})$. We show how a and b can be checked for equality in time $O(\log^5 m)$ w.h.p. using probabilistic polynomial identity testing.

We first give a high-level overview of our solution: Since we can compare the length of A and B by comparing the identifiers $\text{id}_{A, A_0}(A_{m-1}) = m - 1$ and $\text{id}_{B, B_0}(B_{m'-1}) = m' - 1$ of the last amoebots of the chains bit by bit with the PASC algorithm (see Section 3.1) in time $O(\log m)$, we can assume $m = m'$ in the following. Let $c \in \mathbb{N}$. Chain A generates a prime $p \geq 2m$ and repeats the following procedure: A samples r uniformly at random from $[p]$ and sends the pair (p, r) to chain B . Chain A computes $f_a(r) = \sum_{i=0}^{m-1} a_i r^i \pmod{p}$, and chain B computes $f_b(r) = \sum_{i=0}^{m-1} b_i r^i \pmod{p}$ and sends the result to chain A which outputs “ $a \neq b$ ”



■ **Figure 8** Symmetries of an amoebot structure. The amoebot structure is 3-fold rotational symmetric since the $(N, +)$ - and the $(ESE, +)$ -skeleton are symmetric, but not 2-fold or 6-fold rotational symmetric since the $(N, +)$ - and the $(S, +)$ -skeleton are not symmetric. Further, the amoebot structure is reflection symmetric to an axis into northern direction since the $(N, +)$ - and the $(N, -)$ -skeleton are symmetric, but not reflection symmetric to an axis into eastern direction since the $(E, +)$ - and the $(E, -)$ -skeleton are not symmetric.

if $f_a(r) \neq f_b(r)$ and repeats the procedure otherwise. After $c \lceil \log m \rceil$ repetitions, A outputs “ $a = b$ ”. Note that $a = b$ implies $f_a(r) = f_b(r)$. From the Schwartz-Zippel lemma follows that the one-sided error probability for a single repetition is $\Pr[f_a(r) = f_b(r) \mid a \neq b] \leq m/p \leq 1/2$. It follows $\Pr[A \text{ outputs “} a = b \text{”} \mid a \neq b] \leq 1/m^c$.

We now describe the algorithm in more detail: First, we describe a *block primitive* that we use to divide the chain A into blocks of length $k = O(\log m)$ where $k = 2^{\lceil \log \lambda \rceil}$, $\lambda = 2l$ and $l = \lceil \log m \rceil + 2$. Note that $k \geq \lambda$. We have $k \leq m$ for $m \geq 44 =: \eta$. From here on we assume $m \geq \eta$ (in case $m < \eta$, the chains A and B can simply compare their bit strings deterministically). Since the PASC algorithm terminates after $\lceil \log m \rceil$ iterations, we can easily determine amoebot A_λ by using the PASC algorithm 2 times and forwarding a marker after every iteration. Then we use the PASC algorithm again, with the following addition (compare Figure 13): For an amoebot let Q be the partition set on which it received a beep (either its primary or secondary partition set). An active amoebot (except A_0) splits Q into singletons. We obtain a circuit between each pair of consecutive active amoebots. Then, A_0 beeps on Q . If A_λ receives a beep, the procedure terminates, otherwise we continue with the next iteration. After termination, exactly the amoebots A_{ik} are active. Since we can directly compare the bits a_i of the amoebots between the last active amoebot and A_{m-1} with the corresponding bits b_i of chain B in time $O(\log m)$, we assume in the following w.l.o.g. that $k \mid m$ holds. This enables us to divide the amoebots of A into m/k chains $C_i = (A_{ik}, \dots, A_{(i+1)k-1})$ of length k with reference amoebot A_{ik} .

Now we describe how A generates a prime $p \geq 2m$. Chain A samples an l -bit integer $p = (p_0, \dots, p_{l-2}, 1)$ uniformly at random such that A_i stores p_i . Note that the most significant bit is fixed to 1 and therefore $p \in [2^{l-1}, 2^l]$, in particular $2m \leq p < 4m$. We check deterministically whether p is a prime by checking in parallel for all $2 \leq t < m$ whether

$t \mid p$ (note that $\lfloor \sqrt{p} \rfloor < m$): First, p and $t = \text{id}_{A, A_0}(A_{ik}) = ik$ (using the PASC algorithm) are stored in the first l amoebots of every chain C_i in time $O(\log m)$. Then, all chains C_i repeat the following procedure in parallel for at most k times: If $t \geq 2$, check whether $t \mid p$ in time $O(\log^2 m)$ using binary long division with remainder. Abort the prime testing, if $t \mid p$, otherwise increment t .

We repeat the entire procedure at most $3cl^2$ times or until we have successfully sampled a prime p . The runtime for the prime generation is $O(\log^5 m)$. We now analyse the probability for the event E_{fail} that no prime is generated. Using non-asymptotic bounds on the prime-counting function, one can show that the fraction of integers in $[2^{l-1}, 2^l[$ that are prime is at least $1/(3l)$. It follows: $\Pr[E_{\text{fail}}] \leq (1 - 1/(3l))^{3cl^2} \leq 1/e^{cl} \leq 1/m^c$

We now address the probabilistic polynomial identity testing, focusing on the computation of $f_a(r)$ for $r \in [p]$. We use the previously determined division of the chain A into blocks of length $k = O(\log m)$. Assume that p , r and $e = \text{id}_{A, A_0}(A_{ik}) = ik$ are stored in the first l amoebots of every chain C_i . All chains C_i repeat the following procedure in parallel for k times: Compute $s^{(i)} = a_e r^e \pmod{p}$ using modular exponentiation via the right-to-left binary method. Using binary long multiplication and division with remainder, this step is possible in time $O(\log^3 m)$. Then, increment e . Once all chains C_i have completed the j -th repetition, we compute the sum of the $s^{(i)}$ modulo p and store it in the first l amoebots of chain A using a generalization of Theorem 2. The summation is possible in time $O(\log^2 m)$. The computed sum is then added to a running total modulo p .

Finally, after k repetitions, the result $f_a(r)$ is stored in the amoebots A_0, \dots, A_{l-1} . The runtime for the polynomial identity testing is $O(\log^4 m)$.

Note that the size of the outer boundary set is $\Omega(\sqrt{n})$, which is also a lower bound for the size of a canonical skeleton. Hence, we get the following result:

► **Theorem 18.** *The string equality problem on chains of length $O(m)$ can be solved in $O(\log^5 m)$ rounds w.h.p. Therefore, the symmetry detection problem can be solved in $O(\log^5 n)$ rounds w.h.p.*

Additionally, we can compute the amoebot occupying the symmetry point and amoebots on the symmetry axis, but due to the similarity to the applications in Section 3.2, we just sketch the algorithm. The idea is to identify some symmetric amoebots, e.g., global maxima, and to compute symmetric identifiers with these as reference amoebots (see Figure 12). We output all amoebots that receive the same identifier for each reference amoebot.

5 Conclusion and Future Work

In this paper, we have proposed polylogarithmic-time solutions for a range of problems. First, we have computed spatial identifiers in order to compute a stripe through a given amoebot and direction, and the global maxima of the given amoebot structure with respect to a direction. Using these results, we have constructed a canonical skeleton path, which provides a unique characterization of the shape of the given amoebot structure. Constructing canonical skeleton paths for different directions will then allow the amoebots to set up a spanning tree and to check symmetry properties of the given amoebot structure.

Our solutions could be useful for various applications like rapid shape transformation, energy dissemination, and structural monitoring. The details have to be worked out in future work. Beyond that, we think that exploring further applications for the spatial identifiers and the skeleton would be interesting.

References

- 1 John Calvin Alumbaugh, Joshua J. Daymude, Erik D. Demaine, Matthew J. Patitz, and Andréa W. Richa. Simulation of programmable matter systems using active tile-based self-assembly. In *DNA*, volume 11648 of *Lecture Notes in Computer Science*, pages 140–158. Springer, 2019.
- 2 Joshua J. Daymude, Robert Gmyr, Kristian Hinneenthal, Irina Kostitsyna, Christian Scheideler, and Andréa W. Richa. Convex hull formation for programmable matter. In *ICDCN*, pages 2:1–2:10. ACM, 2020.
- 3 Joshua J. Daymude, Kristian Hinneenthal, Andréa W. Richa, and Christian Scheideler. Computing by programmable particles. In *Distributed Computing by Mobile Entities*, volume 11340 of *Lecture Notes in Computer Science*, pages 615–681. Springer, 2019.
- 4 Joshua J. Daymude, Andréa W. Richa, and Christian Scheideler. The canonical amoebot model: Algorithms and concurrency control. In *DISC*, volume 209 of *LIPICs*, pages 20:1–20:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- 5 Joshua J. Daymude, Andréa W. Richa, and Jamison W. Weber. Bio-inspired energy distribution for programmable matter. In *ICDCN*, pages 86–95. ACM, 2021.
- 6 Zahra Derakhshandeh, Shlomi Dolev, Robert Gmyr, Andréa W. Richa, Christian Scheideler, and Thim Strothmann. Brief announcement: amoebot - a new model for programmable matter. In *SPAA*, pages 220–222. ACM, 2014.
- 7 Zahra Derakhshandeh, Robert Gmyr, Thim Strothmann, Rida A. Bazzi, Andréa W. Richa, and Christian Scheideler. Leader election and shape formation with self-organizing programmable matter. In *DNA*, volume 9211 of *Lecture Notes in Computer Science*, pages 117–132. Springer, 2015.
- 8 Michael Feldmann, Andreas Padalkin, Christian Scheideler, and Shlomi Dolev. Coordinating amoebots via reconfigurable circuits. *J. Comput. Biol.*, 29(4):317–343, 2022.
- 9 Giuseppe Antonio Di Luna, Paola Flocchini, Nicola Santoro, Giovanni Viglietta, and Yukiko Yamauchi. Shape formation by programmable particles. *Distributed Comput.*, 33(1):69–101, 2020.
- 10 Ali Mashreghi and Valerie King. Broadcast and minimum spanning tree with $o(m)$ messages in the asynchronous CONGEST model. *Distributed Comput.*, 34(4):283–299, 2021.
- 11 Jennifer E. Padilla, Ruojie Sha, Martin Kristiansen, Junghuei Chen, Natasha Jonoska, and Nadrian C. Seeman. A signal-passing dna-strand-exchange mechanism for active self-assembly of dna nanostructures. *Angewandte Chemie International Edition*, 54(20):5939–5942, 2015.
- 12 Dominic Scalise and Rebecca Schulman. Controlling matter at the molecular scale with dna circuits. *Annual Review of Biomedical Engineering*, 21(1):469–493, 2019. PMID: 31167101. doi:10.1146/annurev-bioeng-060418-052357.
- 13 Shalin Shah, Jasmine Wee, Tianqi Song, Luis Ceze, Karin Strauss, Yuan-Jyue Chen, and John Reif. Using strand displacing polymerase to program chemical reaction networks. *Journal of the American Chemical Society*, 142(21):9587–9593, 2020. PMID: 32364723. doi:10.1021/jacs.0c02240.
- 14 Tianqi Song, Abeer Eshra, Shalin Shah, Hieu Bui, Daniel Fu, Ming Yang, Reem Mokhtar, and John Reif. Fast and compact dna logic circuits based on single-stranded gates using strand-displacing polymerase. *Nature Nanotechnology*, 14(11):1075–1081, November 2019. doi:10.1038/s41565-019-0544-5.
- 15 Tommaso Toffoli and Norman Margolus. Programmable matter: Concepts and realization. *Int. J. High Speed Comput.*, 5(2):155–170, 1993.

A Overview

An overview of our results is given by Table 1.

■ **Table 1** An overview of our algorithmic results.

Problem	Required pins	Runtime	Section	Theorem
Stripe	2	$O(\log n)$	Section 3.3	Theorem 9
Global maxima	2	$O(\log^2 n)$ w.h.p.	Section 3.3	Theorem 10
Canonical skeleton	4	$O(\log^2 n)$ w.h.p.	Section 4.1	Theorem 14
Spanning tree	4	$O(\log^2 n)$ w.h.p.	Section 4.2	Theorem 16
Symmetry detection	4	$O(\log^5 n)$ w.h.p.	Section 4.3	Theorem 18

B Omitted Proofs

Proof of Corollary 3. In order to accelerate the inner outer boundary test, we apply Theorem 2. Recall that each boundary set can organize itself into a cycle. We apply Theorem 1 to split the cycle into a chain. Each amoebot knows its predecessor and successor on the cycle and therefore also within the chain. Let x_i denote the angle at amoebot v_i and let $k = 5$. Finally, note that each boundary has $O(n)$ amoebots since each amoebot has at most three local boundaries. Thus, the sum of the angles can be accumulated after $O(\log n)$ rounds. ◀

Proof of Lemma 5. Since the reference amoebot u_r activates its primary circuit, it always receives a beep on its primary circuit. This implies $(u_r)_i = 0$ for $0 \leq i < k$. Thus, Equation 1 holds.

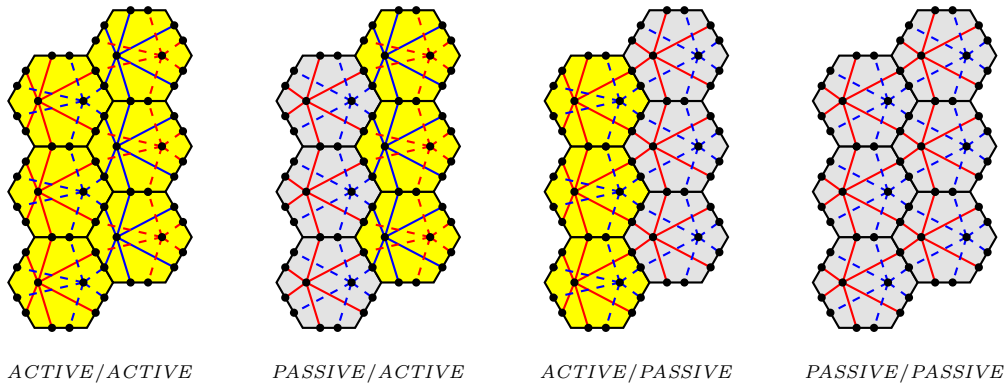
Let y be the successor of x . We have to show that Equation 2 holds, i.e., $\text{id}_{C,u_r}(y) = \text{id}_{C,u_r}(x) + 1$. We have two cases: (i) y never becomes passive, and (ii) y becomes passive. The first case directly implies $\text{id}_{C,u_r}(x) = (1, \dots, 1) = -1$ and $\text{id}_{C,u_r}(y) = (0, \dots, 0) = 0$. For the second case, let l denote the iteration where y becomes passive. Hence, y has received a beep on its primary circuit in the first $l - 2$ iterations and a beep on its secondary circuit in the $(l - 1)$ -st iteration, i.e., $y_{l-1} = 1$ and $y_i = 0$ for $0 \leq i < l - 1$. Since y is active, x has received a beep on its secondary circuit in the first $l - 2$ iterations and a beep on its primary circuit in the $(l - 1)$ -st iteration, i.e., $x_{l-1} = 0$ and $x_i = 1$ for $0 \leq i < l - 1$. Since y is passive from the l -th iteration, $x_i = y_i$ holds for $l \leq i < k$. We obtain

$$\begin{aligned} \text{id}_{C,u_r}(x) &= (x_{k-1}, \dots, x_l, 0, 1, \dots, 1) \\ \text{id}_{C,u_r}(y) &= (y_{k-1}, \dots, y_l, 1, 0, \dots, 0) = (x_{k-1}, \dots, x_l, 0, 1, \dots, 1) + 1 \end{aligned}$$

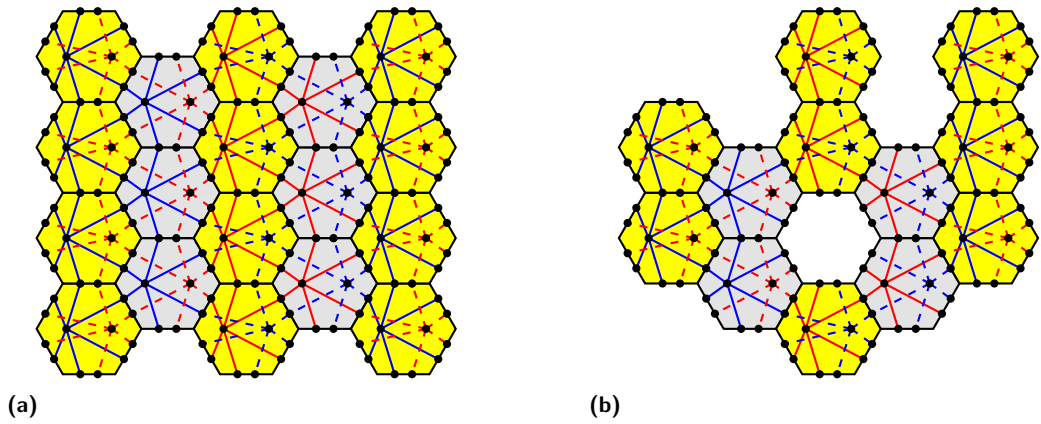
since $2^{l-1} = \sum_{i=0}^{l-2} 2^i + 1$. ◀

Proof of Lemma 6. For the sake of analysis, we first consider an infinite amoebot structure where $S = V$. Afterwards, we transfer the results to arbitrary connected amoebot structures.

Consider a single stripe. Let *ACTIVE* denote the pin configuration used for amoebots of active stripes, and *PASSIVE* denote the pin configuration used for amoebots of passive stripes. Both pin configurations define a primary and secondary partition set. All amoebots within the stripe connect their primary and secondary partition sets, respectively (see Figure 9). We define the union of all primary resp. secondary partition sets within a stripe as the primary resp. secondary partition set of the stripe. Note that each amoebot has access to both partition sets and is able to distinguish between them.



■ **Figure 9** Sections of various stripes of an infinite amoebot structure.



■ **Figure 10** (a) A section of the infinite amoebot structure. (b) An arbitrary amoebot structure. The connectivity of the remaining amoebots is preserved.

Next, consider the connections to the preceding stripe. The connections within the pin configuration *ACTIVE* are selected in such a way that an active stripe connects its primary partition set exclusively to the secondary partition set of the preceding stripe, and its secondary partition set exclusively to the primary partition set of the preceding stripe (see Figure 9). Similar, the connections within the pin configuration *PASSIVE* are selected in such a way that an active stripe connects its primary partition set exclusively to the primary partition set of the preceding stripe, and its secondary partition set exclusively to the secondary partition set of the preceding stripe (see Figure 9). Note that there are no further connections.

The crucial property of our construction is that any two amoebots are connected by any arbitrary path of amoebots in the same fashion, i.e., either both primary partition sets are connected to the secondary partition set of the other amoebot, or their primary and secondary partition sets are connected, respectively (see Figure 10a). This allows us to remove amoebots without separating the circuits as long as the amoebot structure stays connected (see Figure 10b). ◀

Proof of Lemma 7. By Lemma 6, we can apply the primitive of primary and secondary circuits to the chain of stripes as long as each amoebot knows whether its stripe is active or passive. This implies that we can perform the PASC algorithm that by Lemma 5, computes id_{C_d, A_r} .

It remains to show that each amoebot knows whether its stripe is active or passive throughout the execution of the algorithm. Initially, this is trivially true since each stripe is active. A stripe becomes passive once it receives a beep on its secondary circuit. Each amoebot of the stripe can observe this beep since by construction, each amoebot has access to the secondary partition set of its stripe. Thereafter, the stripe stays passive. ◀

Proof of Lemma 8. The proof works analogously to the one of Lemma 7. ◀

Proof of Lemma 12. Let $\text{rank}(B) = \min_{w \in B} f_d(S, w)$ be the rank of an inner boundary set B , and let $\text{rank}(B_O) = -1$ be the rank of the outer boundary set B_O (compare to the definition of the global maxima in Section 1.3). The rank of an inner boundary set is lower than the rank of another inner boundary set if its global maxima are further in direction d than the global maxima of the other inner boundary set. Furthermore, the outer boundary set has a lower rank than all ranks of the inner boundary sets.

We claim that for each inner boundary set B , we construct a path from B to another boundary set B' such that $\text{rank}(B) > \text{rank}(B')$ holds. Clearly, this relationship cannot be cyclic. The lemma immediately follows since we construct a path for each inner boundary set. We prove the claim in the following.

Lemma 11 excludes the possibility of a self-loop, i.e., $B \neq B'$ holds. The claim holds by definition if B' is the outer boundary set. Suppose that B' is an inner boundary set. The claim also holds if the path from u_B to v_B is not trivial since $\text{rank}(B) = f_d(S, u_B) > f_d(S, v_B) > \text{rank}(B')$ holds.

Suppose that the path is trivial, i.e., $u_B \in B$ and $u_B \in B'$. Let $w \in R'$ be a node adjacent to u_B . Note that $f_d(S, u_B) \geq f_d(S, w)$ holds since otherwise, $R = R'$ and with that $B = B'$ would hold. We go from w into direction d_p until we reach an amoebot $x \in B'$. Note that $f_d(S, w) > f_d(S, x)$ holds since for $V \setminus R_O$, f_d is strictly monotonically decreasing if we go into direction d_p . This amoebot exists since B' is an inner boundary set. The claim holds since $\text{rank}(B) = f_d(S, u_B) \geq f_d(S, w) > f_d(S, x) > \text{rank}(B')$ holds. ◀

Proof of Lemma 11. Amoebot u_B has to be adjacent to at least one node in R . Otherwise, $u_B \notin B$ would hold. It is easy to see that if that node would lie in another direction than $\rho_s(d_p, 180)$, either u_B would not be a global maximum of B with respect to d , or u_B would not be the global maximum of B_d with respect to direction $\rho_s(d, 90)$. By the same reasoning, u_B would not be a global maximum of B with respect to d if any amoebot in direction d_p of u_B would be adjacent to a node in R . ◀

Proof of Lemma 13. The canonical skeleton (path) visits a bond either due to a boundary cycle or due to one of the paths. Each local boundary (a common unoccupied adjacent node of the endpoints) adds one visit. Each bond has at most two local boundaries. Each path adds two visits. Due to Lemma 11, a bond cannot be part of more than one path. A bond cannot be part of a boundary cycle and a path at the same time since the path would stop at either endpoints due to the unoccupied adjacent node. Altogether, each bond is visited at most twice. ◀


Proof of Theorem 14. The preprocessing step requires $O(\log n)$ rounds w.h.p. (see Section 2). The first step requires $O(\log^2 n)$ rounds for the computation of global maxima (see Section 3.3). The second requires $O(1)$ rounds. Altogether, the canonical skeleton algorithm requires $O(\log^2 n)$ rounds w.h.p. ◀

Proof of Lemma 15. In order to prove that T is a tree, we show that T is cycle-free and connected. Each edge $\pi(v) = (v', v)$ implies that v' appears before v in π . Clearly, this relationship cannot be cyclic.

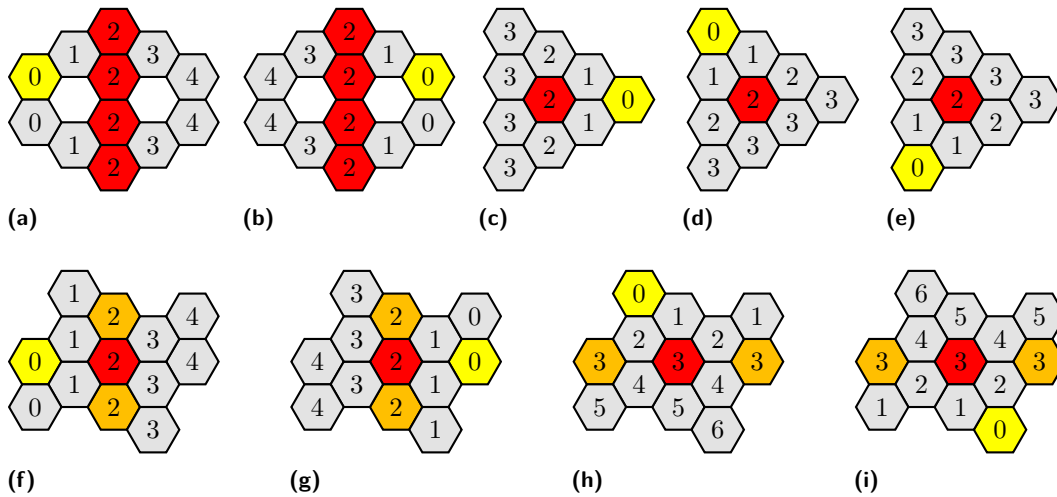
We prove that T is connected by induction on the path π . The induction base holds trivially for v_1 . Suppose that all nodes up to node v_i are connected within T . Consider node v_{i+1} . If it is not the first occurrence of v_{i+1} on the path, then v_{i+1} is already connected by induction hypothesis. Otherwise $\pi(v_{i+1}) = \{v_i, v_{i+1}\} \in E'$. This edge connects v_{i+1} to all nodes up to node v_i since these are connected by induction hypothesis. ◀

Proof of Theorem 16. The correctness follows from Lemma 15. The first phase requires $O(\log n)$ rounds (see Section 3.1). The second phase requires $O(1)$ rounds. Altogether, the spanning tree algorithm requires $O(\log n)$ rounds. ◀


C Omitted Figures




■ **Figure 11** Predefined splitting point with respect to d_p . For $d_p = N$, the figure shows all cases where the canonical skeleton visits u_{B_O} multiple times. By similar arguments as for Lemma 11, there are no other cases. The red lines indicate the canonical skeleton. The node indicates the splitting point.



■ **Figure 12** Symmetry point and symmetry axis. (a) and (b) show the computation of a symmetry axis. (c) to (e) show the computation of a symmetry point for a 3-fold or 6-fold rotationally symmetric amoebot structure. (f) to (i) show the computation of a symmetry point for a 2-fold rotationally symmetric amoebot structure. In this case, we have to compute two axes so that the symmetry point lies on the intersection of these. The yellow amoebot indicates the reference amoebot. The red (and orange) amoebots receive the same identifier for each reference amoebot.



■ **Figure 13** Block primitive. The figure shows how a chain A with reference amoebot A_0 and a marked amoebot A_λ can be divided into blocks of length $k \geq \lambda$ where $k = 2^{\lceil \log \lambda \rceil}$ using the PASC algorithm (compare Figure 3) with an additional step after each iteration. Here, we have $\lambda = 3$ and $k = 4$. The first line shows the chain at the beginning of the i -th iteration of the PASC algorithm. The circles indicate the partition sets. The blue bordered amoebot is A_0 , and the green bordered amoebot is A_λ . Yellow amoebots are active, and gray amoebots are passive. The second line shows the configuration resulting from the following additional step: For an amoebot let Q be the partition set on which it received a beep (depicted in red; either its primary or secondary partition set). An active amoebot (except A_0) splits Q into singletons. We obtain a circuit between each pair of consecutive active amoebots. Then, A_0 beeps on Q . If A_λ receives a beep, the procedure terminates, otherwise continue with the next iteration. After termination, exactly the amoebots A_{ik} are active.



■ **Figure 14** The figure shows $(NNW, +)$ -skeleton. The yellow and orange amoebots indicate the global maxima of the boundary sets. The orange amoebots indicate the starting points of the paths. The blue lines indicate the paths between the boundary cycles. The node indicates the location where the cycle is split.