The Canonical Amoebot Model: Algorithms and Concurrency Control

Joshua J. Daymude and Andréa W. Richa (Arizona State University) Christian Scheideler (Universität Paderborn)

DISC 2021, Freiburg, Germany (Hybrid Event) — October 5–7, 2021







Programmable Matter

Programmable matter is a substance that can change its physical properties autonomously based on user input or environmental stimuli.



RGR 2013



"Kilobots" RCN 2014





"Particle Robots" LBBCRHRL 2019

The Geometric Amoebot Model (Before This Paper)

The (geometric) **amoebot model** is an abstraction of programmable matter.

- Space: triangular lattice G_{Δ} .
- Amoebots can be contracted (one node) or expanded (two adjacent nodes).
- Amoebots are anonymous, have only constant-size memories, communicate with immediate neighbors, and have no global compass.
- Self-actuated movements via expansions, contractions, and handovers.
- Sequential, weakly fair adversary: one amoebot acts per time, every amoebot acts infinitely often.



A Short History of the Amoebot Model

~30 papers: Models, shape formation, leader election, object coating, Markov chains, and more!



Related Models

The amoebot model considers <u>active</u> entities in <u>space</u>, unlike <u>non-spatial</u>, <u>passive</u> models like population protocols, network constructors, and tile-based self-assembly.

The most closely related model is **autonomous mobile robots** (see [FPS 2019] for a recent overview), specifically those using discrete, graph-based models of space.

Similarities

- Anonymous individuals.
- Active movements w/o global orientation.
- Limited compute/sensing capabilities.
- Persistent memory (*F*-state). [BDS 2008]
- Limited communication (luminous robots). [DFPSY 2012, DFPSY 2016, DFCPSV 2017]

Differences

- Amoebots respect physical exclusion.
- Mobile robots use look-compute-move cycles that are difficult to reconcile with traditional message-passing.
- Mobile robots have a hierarchy of adversarial schedulers; amoebots usually assume sequential activations.

Overview of Results

1) The Canonical Amoebot Model

- Provides a standardized hierarchy of assumption variants to better facilitate comparisons.
- Models communication via message passing to address concurrency using standard adversarial activation models.

2) Sufficient Conditions for Concurrent Correctness

- Identifies sufficient conditions for amoebot algorithm correctness under any adversary.
- Shows that an algorithm for hexagon formation exists that satisfies these conditions.

3) A General Framework for Concurrency Control

• Converts algorithms that terminate under a sequential adversary and satisfy certain conventions into algorithms that exhibit equivalent behavior under an asynchronous adversary.

The standardized hierarchy of assumption variants:

	Variant	Description
Space	General [*] Geometric ^{*,†}	G is any infinite, undirected graph. $G = G_{\Delta}$, the triangular lattice.
Orientation	Assorted ^{*,†} Common Chirality [*] Common Direction Common	Assorted direction and chirality. Assorted direction but common chirality. Common direction but assorted chirality. Common direction and chirality.
Memory	Oblivious Constant-Size ^{*,†} Finite Unbounded	No persistent memory. Memory size is $\mathcal{O}(1)$. Memory size is $\mathcal{O}(f(n))$, some function of the system size. Memory size is unbounded.
Concurrency	Asynchronous [†] Synchronous [*] <i>k</i> -Isolated Sequential [*]	Any amoebots can be simultaneously active. Any amoebots can simultaneously execute a single action per discrete round. Each round has an evaluation phase and an execution phase. No amoebots within distance k can be simultaneously active. At most one amoebot is active per time.
Fairness	Unfair [†] Weakly Fair [*] Strongly Fair	Some enabled amoebot is eventually activated. Every continuously enabled amoebot is eventually activated. Every amoebot enabled infinitely often is activated infinitely often.

* have been considered in previous work; † are the focus of algorithmic results in this work.

The Canonical Amoebot Model

In the canonical amoebot model, amoebot functionality is partitioned into:

- A higher-level **application layer** where algorithms are defined in terms of operations.
- A lower-level system layer that executes an amoebot's operations via message passing.

Operation	Return Value on Success	
Connected (p)	TRUE iff a neighboring amoebot is connected via port p	$p = \bot$ Return
CONNECTED()	$[c_0, \ldots, c_{k-1}] \in \{N_1, \ldots, N_8, \text{FALSE}\}^k$ where $c_p = N_i$ if N_i is the locally identified neighbor connected via port p and $c_p = \text{FALSE}$ otherwise	$start \rightarrow (A.x) \rightarrow success \rightarrow (Xetum)$
$\operatorname{Read}(p,x)$	The value of x in the public memory of this amoebot if $p = \bot$ or of the neighbor incident to port p otherwise	$r = \mathbf{read_ack}(x, x_{val})$
$WRITE(p, x, x_{val})$	Confirmation that the value of x was updated to x_{val} in the public memory of this amoebot if $p = \bot$ or of the neighbor incident to port p otherwise	Check for connection via port p Yes Tead_request(x) in port p Wait for response Wait for
$\operatorname{Contract}(v)$	Confirmation of the contraction out of node $v \in \{\text{HEAD}, \text{TAIL}\}$	
$\operatorname{Expand}(p)$	Confirmation of the expansion into the node incident to port p	Disconnection
$\operatorname{Pull}(p)$	Confirmation of the pull handover with the neighbor incident to port \boldsymbol{p}	
$\operatorname{Push}(p)$	Confirmation of the push handover with the neighbor incident to port \boldsymbol{p}	FAILURE
Lock()	Local identifiers of this amoebot and the neighbors that were locked	(a) $\operatorname{READ}(p, x)$
$\mathrm{Unlock}(\mathcal{L})$	Confirmation that the amoebots of \mathcal{L} were unlocked	

The Canonical Amoebot Model

Algorithms in the canonical amoebot model are specified in terms of **actions**:

 $\langle label \rangle : \langle guard \rangle \rightarrow \langle operations \rangle$

- *label* specifies the action's name.
- *guard* is a Boolean predicate determining whether this action is currently **enabled**.
- operations specifies the computation and sequence of operations to perform if enacted.

Example from Hexagon-Formation:

 $\alpha_2 : (A.state = IDLE) \land (\exists B \in N(A) : B.state \in \{FOLLOWER, ROOT\}) \rightarrow$

Find a port p for which CONNECTED(p) = TRUE and $READ(p, state) \in \{FOLLOWER, ROOT\}$. WRITE(\bot , parent, p). WRITE(\bot , state, FOLLOWER). We use an **adversary** to model timing and progress. Four levels of concurrency:

1) Sequential. At most one amoebot can be active at a time.



3) Synchronous. Arbitrary sets of amoebots can be active in each discrete round.





4) Asynchronous. Arbitrary sets of amoebots can be simultaneously active.



 $time \rightarrow$

The adversary can only activate amoebots with enabled actions. Three levels of fairness:

- 1) Strongly Fair. Every amoebot that is enabled infinitely often is activated infinitely often.
- 2) Weakly Fair. Every continuously enabled amoebot is eventually activated.
- 3) Unfair. Some enabled amoebot is eventually activated.

The rest of this talk will primarily focus on unfair asynchronous adversaries, the most general of all adversarial activation models.

Informally: the adversary can activate any amoebot with something to do whenever it wants to.



Asynchronous Hexagon Formation

<u>Problem</u>: Reconfigure any connected amoebot system as a regular hexagon, assuming there is a unique seed amoebot initially in the system.



Asynchronous Hexagon Formation

We formulate a Hexagon-Formation algorithm in terms of actions based on [DGRSS 2015].

```
Algorithm 1 Hexagon-Formation for Amoebot A
 1: \alpha_1 : (A.state \in {IDLE, FOLLOWER}) \land (\exists B \in N(A) : B.state \in {SEED, RETIRED}) \rightarrow
         WRITE(\perp, parent, NULL).
 2:
         WRITE(\perp, state, ROOT).
 3:
         WRITE(\perp, dir, GETNEXTDIR(counter-clockwise)).
 4:
 5: \alpha_2 : (A.\mathtt{state} = \mathtt{IDLE}) \land (\exists B \in N(A) : B.\mathtt{state} \in \{\mathtt{FOLLOWER}, \mathtt{ROOT}\}) \rightarrow
         Find a port p for which CONNECTED(p) = TRUE and READ(p, state) \in \{FOLLOWER, ROOT\}.
 6:
         WRITE(\perp, parent, p).
 7:
         WRITE(\perp, state, FOLLOWER).
 8:
 9: \alpha_3: (A.shape = CONTRACTED) \land (A.state = ROOT) \land (\forall B \in N(A): B.state \neq IDLE)
           \wedge (\exists B \in N(A) : (B.state \in {SEED, RETIRED}) \wedge (B.dir is connected to A)) \rightarrow
10:
         WRITE(\perp, dir, GETNEXTDIR(clockwise)).
11:
         WRITE(\perp, state, RETIRED).
12:
13: \alpha_4: (A.shape = CONTRACTED) \land (A.state = ROOT) \land (the node adjacent to A.dir is empty) \rightarrow
         EXPAND(A.dir).
14:
15: \alpha_5 : (A.\mathtt{shape} = \mathtt{EXPANDED}) \land (A.\mathtt{state} \in \{\mathtt{FOLLOWER}, \mathtt{ROOT}\}) \land (\forall B \in N(A) : B.\mathtt{state} \neq \mathtt{IDLE})
           \wedge (A has a tail-child B : B.shape = CONTRACTED) \rightarrow
16:
         if READ(\bot, state) = ROOT then WRITE(\bot, dir, GETNEXTDIR(counter-clockwise)).
17:
         Find a port p \in \text{TAILCHILDREN}() s.t. \text{READ}(p, \text{shape}) = \text{CONTRACTED}.
18:
         Let p' be the label of the tail-child's port that will be connected to p after the pull handover.
19:
         WRITE(p, parent, p').
20:
         \operatorname{PULL}(p).
21:
22: \alpha_6 : (A.\mathtt{shape} = \mathtt{EXPANDED}) \land (A.\mathtt{state} \in \{\mathtt{FOLLOWER}, \mathtt{ROOT}\}) \land (\forall B \in N(A) : B.\mathtt{state} \neq \mathtt{IDLE})
           \wedge (A has no tail-children) \rightarrow
23:
         if READ(\bot, state) = ROOT then WRITE(\bot, dir, GETNEXTDIR(counter-clockwise)).
24:
         CONTRACT(TAIL).
25:
```

Asynchronous Hexagon Formation

Theorem. Assuming geometric space, assorted orientations, and constant-size memory, Hexagon-Formation (HF) is correct under any (i.e., an unfair asynchronous) adversary.

Outline of analysis:

- HF is correct under an unfair sequential adversary.
- Enabled actions of HF remain enabled despite concurrent executions.
- Enabled actions of HF are executed identically in sequential and asynchronous executions.
- The action executions of any asynchronous execution of HF can be serialized.
- Any asynchronous execution of HF terminates.





Sufficient Conditions for Concurrent Correctness

Our analysis of Hexagon-Formation shows that any algorithm satisfying

- 1. correctness under an unfair sequential adversary,
- 2. enabled actions remaining enabled despite concurrent action executions, and
- 3. enabled action executions remaining successful/unaffected by concurrent action executions must also be correct under an unfair asynchronous adversary, the most general of all.

Notably, this is true without using the Lock/Unlock operations, demonstrating that while useful, locks are not always necessary for designing correct concurrent amoebot algorithms.

We've shown how individual algorithms can be designed with actions that are ambivalent to the effects of concurrency, but this can be difficult to achieve in general.

Here, we use **locks** to lift correct sequential algorithms to the asynchronous setting.



This is the best of both worlds: the ease of designing algorithms in the sequential setting and the relevance of correct execution in the more realistic concurrent setting.

This is too optimistic and may be impossible to guarantee in general, so instead we only consider algorithms \mathcal{A} that obey certain conventions.

Key Ideas:

- On activation, an amoebot *A* first attempts to lock its neighborhood.
- If successful, its locked neighbors cannot move or change their memory contents.
- So A can evaluate its guards and perform its actions as if things were sequential (sort of).
- Failed locking attempts and expansions have no effect on the rest of the system.

Key Issue: Locks can't stop amoebots from expanding into an acting amoebot's neighborhood!

Example. "If I have no neighbors, then expand in the 'forward' direction."





We introduce a set of conventions that must be satisfied for the framework to apply.

Convention 1: Any execution of an enabled action must succeed in the sequential setting.

Convention 2: All compute operations must precede at most one movement operation.

Convention 3: **Monotonicity**. Action executions cannot be affected by (unlocked) amoebots that concurrently enter the acting amoebot's neighborhood.



<u>Static</u> algorithms (i.e., those that don't involve movement) are trivially monotonic. This includes most of the leader election algorithms.

<u>Open Question</u>: What amoebot algorithms satisfy monotonicity?

Algorithm 2 Concurrency Control Framework for Amoebot A **Input**: Algorithm $\mathcal{A} = \{ [\alpha_i : q_i \to ops_i] : i \in \{1, \ldots, m\} \}$ satisfying Conventions 1–3. 1: Set $q' \leftarrow (A.act = TRUE)$ and $ops' \leftarrow$ "Do: try: Set $\mathcal{L} \leftarrow \text{LOCK}()$ to attempt to lock A and its neighbors. catch lock-failure do abort. if A.awaken = TRUE then for all amoebots $B \in \mathcal{L}$ do WRITE B.act \leftarrow TRUE. WRITE A.awaken \leftarrow FALSE, UNLOCK(\mathcal{L}), and abort. for all actions $[\alpha_i : g_i \to ops_i] \in \mathcal{A}$ do Perform CONNECTED and READ operations to evaluate guard q_i w.r.t. \mathcal{L} . Evaluate q_i in private memory to determine if α_i is enabled. if no action is enabled then WRITE A.act \leftarrow FALSE, UNLOCK(\mathcal{L}), and abort. Choose an enabled action $\alpha_i \in \mathcal{A}$ and perform its compute phase in private memory. Let W_i be the set of WRITE operations and M_i be the movement operation in ops_i based on its compute phase; set $M_i \leftarrow$ NULL if there is none. if M_i is EXPAND (say, from node u into node v) then try: Perform the EXPAND operation and WRITE A.awaken \leftarrow TRUE. **catch expand-failure do** UNLOCK(\mathcal{L}) and abort. for all amoebots $B \in \mathcal{L}$ do WRITE $B.act \leftarrow TRUE$. for all $(B.x \leftarrow x_{val}) \in W_i$ do WRITE $B.x \leftarrow x_{val}$ and WRITE B.awaken \leftarrow TRUE if M_i is NULL or EXPAND then UNLOCK each amoebot in \mathcal{L} . else if M_i is CONTRACT (say, from nodes u, v into node u) then UNLOCK each amoebot in \mathcal{L} that is adjacent to node v but not to node u. Perform the CONTRACT operation. UNLOCK each remaining amoebot in \mathcal{L} . else if M_i is PUSH (say, A is pushing B) then WRITE A.awaken \leftarrow TRUE and B.awaken \leftarrow TRUE. Perform the PUSH operation. UNLOCK(\mathcal{L}). else if M_i is PULL (say, A in nodes u, v is pulling B into node v) then WRITE B.awaken \leftarrow TRUE. UNLOCK each amoebot in \mathcal{L} (except B) that is adjacent to node v but not to node u. Perform the PULL operation. UNLOCK each remaining amoebot in \mathcal{L} ."

[Failure Point 1] Attempt to Lock.

- Evaluate all guards of actions in \mathcal{A} .
- If there is an enabled action, emulate its compute phase (Connected, Read, and Write operations) in private memory.
 - [Failure Point 2] If there's an Expand operation, attempt to do so.
- Perform any Write operations decided on 5. in the compute phase.
- Handle any other movement types. ← 6.

Other bookkeeping throughout: Unlocks, setting act/awaken bits, etc.

32: return $\mathcal{A}' = \{ [\alpha' : q' \to ops'] \}.$

2:

3:

4:

5:

6:

7:

8:

9:

10:

11:

12:

13:

14:

15:

16:17:

18:

19:

20:

21:

22:

23:

24:

25:26:

27:

28:

29:

30:

31:

- 1. <u>Validity</u>. Any execution of an enabled action succeeds in the sequential setting.
- 2. <u>Computing Before Moving</u>. Compute operations precede movement operations.
- 3. <u>Monotonicity</u>. Action executions are not affected by (unlocked) amoebots that concurrently enter the acting amoebot's neighborhood.

Theorem. Consider any algorith in \mathcal{A} satisfying Conventions 1-3 and let \mathcal{A}' be the algorithm obtained by the concurrency control framework. If \mathcal{A} terminates under any sequential execution, then every asynchron pus execution of \mathcal{A}' terminates in an outcome that some sequential execution of \mathcal{A} also terminates in.



Open Questions

- When existing amoebot algorithms are standardized according to the canonical model's hierarchy of assumptions and described in action semantics, how do they compare?
- How should the canonical model be extended to address fault tolerance? To date, crash faults have been considered for specific problems [DFPSV 2018, DRW 2021].
- We know static algorithms satisfy monotonicity what about algorithms with movement?
- Do there exist algorithms that are not correct under an asynchronous adversary but are compatible with our concurrency control framework?
- Are there other, less restrictive sufficient conditions for asynchronous correctness?

Thank you!

sops.engineering.asu.edu

jdaymude.github.io

J <u>@joshdaymude</u>





