

Composable computation in discrete chemical reaction networks*

Eric E. Severson¹, David Haley¹, and David Doty¹

¹University of California, Davis. {eseverson, drhaley, doty}@ucdavis.edu.

Abstract

We study the composability of discrete chemical reaction networks (CRNs) that *stably compute* (i.e., with probability 0 of error) integer-valued functions $f : \mathbb{N}^d \rightarrow \mathbb{N}$. We consider *output-oblivious* CRNs in which the output species is never a reactant (input) to any reaction. The class of output-oblivious CRNs is fundamental, appearing in earlier studies of CRN computation, because it is precisely the class of CRNs that can be composed by simply renaming the output of the upstream CRN to match the input of the downstream CRN.

Our main theorem precisely characterizes the functions f stably computable by output-oblivious CRNs with an initial leader. The key necessary condition is that for sufficiently large inputs, f is the minimum of a finite number of nondecreasing *quilt-affine* functions. (An affine function is linear with a constant offset; a *quilt-affine* function is linear with a periodic offset).

1 Introduction

A foundational model of chemistry commonly used in natural sciences is that of chemical reaction networks (CRNs): finite sets of chemical reactions such as $A + B \rightarrow A + C$. The model is described as a continuous time, discrete state, Markov process [20]. A configuration of the system is a vector of non-negative integers specifying the molecular counts of the species (e.g., A , B , C), a reaction can occur only when all its reactants are present, and transitions between configurations correspond to reactions (e.g., when the above reaction occurs the count of B is decreased by 1 and the count of C increased by 1). CRNs are widely used to describe natural biochemical systems such as the intricate cellular regulatory networks responsible for the information processing within cells. Looking beyond the *scientific* goal of understanding natural CRNs, to the *engineering* goal of constructing programmable, autonomous smart molecules, artificial CRNs have been implemented using the physical primitive of nucleic-acid strand displacement cascades [8, 11, 23, 24].

Population protocols, a widely-studied model of distributed computing with very limited agents, are a restricted subset of CRNs (those with two reactants and two products in each reaction) that nevertheless capture many of the interesting features of CRNs. The key feature is the inability of agents (molecules) to control their schedule of communication (collisions). The decision problems solvable by population protocols have been studied extensively: they can simulate Turing machines with high probability in polylogarithmic time (with [5] or without [21] an initial leader), whereas requiring probability 0 of error limits the computable predicates to being semilinear [6].

* Authors supported by NSF grants 1619343 and 1844976.

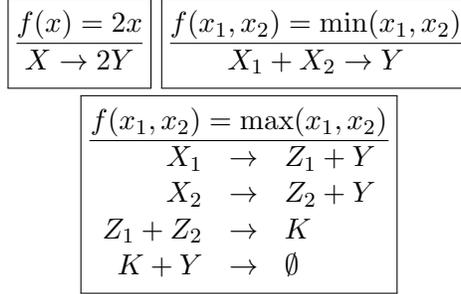


Figure 1: Functions stably computed by CRNs. Note max is computed as $x_1 + x_2 - \min(x_1, x_2)$.

1.1 Function computation

Computation of functions $f : \mathbb{N}^d \rightarrow \mathbb{N}$ was discussed briefly in the first population protocols paper [3, Section 3.4], which focused more on Boolean predicate computation, and it was defined formally first for CRNs [10, 17] and later for population protocols [7]. The class of functions stably computable in either model is the same: the semilinear functions [6, 10]. We use the CRN model because it is more natural for describing functions, but our results also apply to the population protocol model.

To represent an input $\mathbf{x} \in \mathbb{N}^d$, we start in a configuration with counts $\mathbf{x}(i)$ of species X_i for each $i \in \{1, \dots, d\}$, and count 1 of a “leader” species L .¹ A function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is *stably computable* by a CRN if a correct and *stable* configuration \mathbf{O} (i.e., on input \mathbf{x} the count of Y is $f(\mathbf{x})$ in all configurations reachable from \mathbf{O}) remains reachable no matter what reactions occur.²

See Fig. 1 for examples. It is known that a function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is stably computable by a CRN if and only if it is *semilinear*: intuitively, it is a piecewise affine function. (See Definition 2.6.)

1.2 Composability

Note a key difference between the CRNs for min and max in Fig. 1: the former only produces the output species Y , whereas the latter also contains reactions that consume Y . In one possible sequence of reactions for the max CRN, the inputs can be exhausted through the first two reactions before ever executing the last two reactions. In doing so, the count of Y overshoots its correct value of $\max(x_1, x_2)$ before the excess is consumed by the reaction $K + Y \rightarrow \emptyset$.

For this reason that the min CRN is more easily composed with a downstream CRN. For example, the function $2 \cdot \min(x_1, x_2)$ is stably computed by the reactions $X_1 + X_2 \rightarrow W$ (computing $w = x_1 + x_2$) and $W \rightarrow 2Y$ (computing $y = 2w$), renaming the output of the min CRN to match the input of the multiply-by-2 CRN. However, this approach does not work to compute $2 \cdot \max(x_1, x_2)$; changing Y to W in the four-reaction max CRN and adding the reaction $W \rightarrow 2Y$ can erroneously result in up to $2(x_1 + x_2)$ copies of Y being produced. Intuitively, the multiply-by-2 reaction $W \rightarrow 2Y$ competes with the upstream reaction $K + W \rightarrow \emptyset$ from the max CRN.

¹ The leader is discussed in Section 1.3. A CRN may ignore its leader, as in Fig. 1.

² We use this definition throughout the paper, but we mention here that it is equivalent to two other natural definitions. The first definition is that any fair sequence of reactions *will* take the CRN to such a correct stable configuration, where *fair* means that any configuration that is infinitely often reachable is eventually reached. The second definition is that a correct stable configuration is actually reached with probability 1.

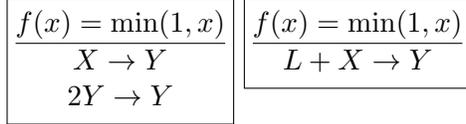


Figure 2: $\min(1, x)$ is stably computed by a *leaderless* non-output-oblivious CRN (left), and an output-oblivious CRN with a single leader L (right).

This motivates us to study the class of functions $f : \mathbb{N}^d \rightarrow \mathbb{N}$ stably computable by *output-oblivious* CRNs: those in which the output species Y appears only as a product, never as a reactant. We call such a function *obliviously-computable*. Any obliviously-computable function must be nondecreasing, otherwise reactions could incorrectly overproduce output (see Observation 2.1).

Obliviously-computable functions must also be semilinear, so it is reasonable to conjecture that a function is obliviously-computable if and only if it is semilinear and nondecreasing. In fact, this is true for 1D functions $f : \mathbb{N} \rightarrow \mathbb{N}$ (see Section 3). However, in higher dimensions, the function $\max : \mathbb{N}^2 \rightarrow \mathbb{N}$ is semilinear and nondecreasing, yet not obliviously-computable; its consumption of output turns out to be unavoidable. Assuming there is no leader, this is simple to prove: Since $\max(1, 0) = 1$, starting with one X_1 , a Y can be produced. Similarly, a Y can be produced starting with one X_2 . Then with one X_1 and one X_2 , these reactions can happen in parallel and produce two Y 's, too many since $\max(1, 1) = 1$. It is more involved to prove that even with a leader, it remains impossible to obliviously compute \max ; see Section 4.³

1.3 The role of the leader

Our model includes an initial leader, which is essential for our general constructions (see Sections 3 and 6). The class of stably computable functions is identical whether an initial leader is allowed or not [17], as is the class of stably computable predicates [6].

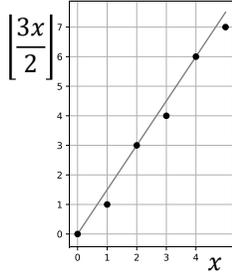
Interestingly, the class of obliviously-computable functions we study is provably larger when an initial leader is allowed. For example, consider the function $f(x) = \min(1, x)$ (see Fig. 2). f is stably computable with or without a leader, but only the construction with a leader is output-oblivious. Without using a leader, f is not obliviously-computable (see Observation 9.1).

Including the leader gives additional power to the model. This gives more power to our CRN constructions, but makes our impossibility results stronger. Fully classifying the obliviously-computable functions in a leaderless model remains an open question.

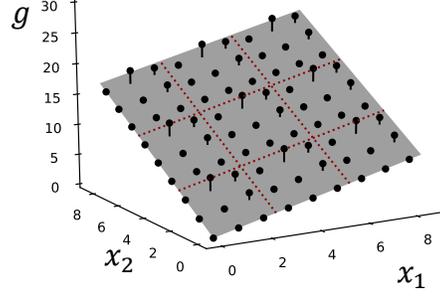
1.4 Contribution

Our main result, Theorem 5.2, provides a complete characterization of the class of obliviously-computable functions. It builds off a key definition: a *quilt-affine* function is a nondecreasing function that is the sum of a rational linear function and periodic function (formalized as Definition 5.1). For example, functions such as $\lfloor \frac{3x}{2} \rfloor$ are quilt-affine (see Fig. 3a). Such floored division functions are natural to the discrete CRN model ($\lfloor \frac{3x}{2} \rfloor$ is stably computed by reactions $X \rightarrow 3Z$, $2Z \rightarrow Y$). Fig. 3b shows a higher-dimensional quilt-affine function, with a “bumpy quilt” structure that motivates the name. Quilt-affine functions are also characterized by nonnegative periodic finite differences, a structure key to showing they are obliviously-computable (see Lemma 6.1).

³ This result was obtained independently by Chugg, Condon, and Hashemi [13].

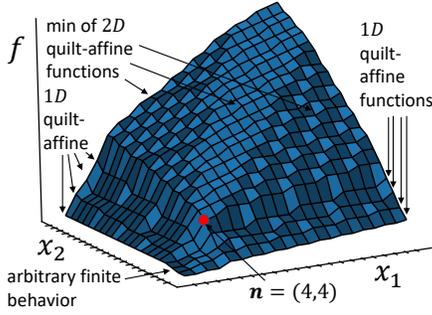


(a) A 1D (single-input) quilt-affine function $\lfloor \frac{3x}{2} \rfloor = \frac{3}{2}x + B(\bar{x} \bmod 2)$, where $B(\bar{0}) = 0$ and $B(\bar{1}) = -\frac{1}{2}$.

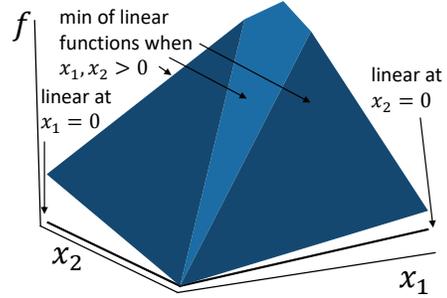


(b) A 2D quilt-affine function $g(\mathbf{x}) = (1, 2) \cdot \mathbf{x} + B(\bar{\mathbf{x}} \bmod 3)$, where $B(\bar{\mathbf{x}}) = 0$ except when $\bar{\mathbf{x}} \in \{(1, 2), (2, 2), (2, 1)\}$.

Figure 3: Examples of 1D and 2D quilt-affine functions.



(a) A 2D function satisfying Theorem 5.2.



(b) The scaling limit gives a 2D real-valued obviously-computable function from [9].

Figure 4: Example discrete and real-valued obviously-computable functions.

Theorem 5.2 states that a function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable if and only if

- i) [nondecreasing] f is nondecreasing,
- ii) [eventually-min] for sufficiently large inputs, f is the minimum of a finite number of quilt-affine functions, and
- iii) [recursive] every restriction $f_r : \mathbb{N}^{d-1} \rightarrow \mathbb{N}$ obtained by fixing some inputs to a constant value⁴ is obviously-computable (i.e., eventually the minimum of quilt-affine functions).

Condition (ii) characterizes f when all inputs are sufficiently large (greater than some $\mathbf{n} \in \mathbb{N}^d$), whereas condition (iii) characterizes f when some inputs are fixed to smaller values. See Fig. 4a for a representative example of an obviously-computable $f : \mathbb{N}^2 \rightarrow \mathbb{N}$. This pictured function has arbitrary nondecreasing values in the “finite region” where $\mathbf{x} < (4, 4)$, has eventual 1D quilt-affine behavior along the lines $x_1 = 0, 1, 2, 3$ and $x_2 = 0, 1, 2, 3$, and is the minimum of 3 different quilt-affine functions in the “eventual region” where $\mathbf{x} \geq (4, 4)$. This behavior generalizes naturally to higher dimensions.

⁴ Note that Theorem 5.2 defines fixed-input restrictions slightly differently; see Section 5 for an explanation.

The most technically sophisticated part of our result is the proof that the eventually-min condition (ii) of Theorem 5.2 is necessary; the main ideas of this proof are outlined in Section 7.1.

1.5 Related work

Chalk, Kornerup, Reeves, and Soloveichik [9] showed an analogous result in the *continuous* model of CRNs in which species amounts are given by nonnegative real concentrations. A consequence of their characterization is that any obviously-computable real-valued function is a minimum of linear functions when all inputs are positive. In Theorem 8.2, we demonstrate that the limit of “scalings” of a function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ satisfying our main Theorem 5.2 is in fact a function $\hat{f} : \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}_{\geq 0}$ satisfying the main theorem of [9] (see Fig. 4b). The discrete details lost in the scaling limit constitute precisely the unique challenges of proving Theorem 5.2 that are not handled by [9]. In particular, our function class can contain arbitrary finite behavior and repeated finite irregularities.

Returning to the *discrete* (a.k.a., *stochastic*) CRN model we study, Chugg, Condon, and Hashemi [13] independently investigated the special case of *two-input* functions $f : \mathbb{N}^2 \rightarrow \mathbb{N}$ computable by output-oblivious CRNs, obtaining a characterization equivalent to ours when restricted to 2D. Their characterization is phrased much differently, with specially constructed “fissure functions” to describe the function behavior across what we describe as *under-determined regions* (intuitively, thin “1D” regions bounded by parallel lines, where f cannot be described by a unique quilt-affine function, see Section 7). The ideas required to prove the 2D case are sophisticated and far from simple, yet unfortunately, these ideas do not extend straightforwardly to higher dimensions. The planarity of the 2D input space constrains the regions induced by separating hyperplanes (i.e., lines) in a strong way. Furthermore, the fact that there is only one nontrivial integer dimension smaller than 2 implies that the under-determined regions are simpler to reason about than in the case where they can have arbitrary dimension between 1 and d . Finally, even restricted to 2D, a notable aspect of our characterization is expressing f a minimum of quilt-affine functions, which are simple intrinsic building blocks that generalize immediately.

1.6 Other ways of composing computation

In Section 2.3 we show that a for a CRN \mathcal{C} be composable with downstream CRN \mathcal{D} by “concatenation” (renaming \mathcal{C} ’s output species to match \mathcal{D} ’s input species and ensuring all other species names are disjoint between \mathcal{C} and \mathcal{D}), it is (in a sense) necessary and sufficient for \mathcal{C} to be output-oblivious. There are other ways to compose computations, however.

A common technique (e.g. [19]) is for \mathcal{C} to detect when its output has changed and send a *restart* signal to \mathcal{D} . However, it is not obvious how to do this with function computation as defined in this paper, where \mathcal{D} changes \mathcal{C} ’s output by consuming it.

Another technique (e.g. [5]) is to set a *termination signal*, which is a sub-CRN that, with high probability, creates a copy of a signal species T , but not before \mathcal{C} has converged. T then “activates” the reactions of \mathcal{D} , so that \mathcal{D} will not consume the output of \mathcal{C} until it is safe to do so. However, this has some positive failure probability. In fact, if we require T to be guaranteed with probability 1 to be produced only after the CRN has converged, only constant functions can be stably computed. Worse yet, in the leaderless case, it is provably impossible to achieve this guarantee even with positive probability [16].

2 Preliminaries

2.1 Notation

\mathbb{N} denotes the set of nonnegative integers. For a set \mathcal{S} (of species), we write $\mathbb{N}^{\mathcal{S}}$ to denote the set of vectors indexed by the elements of \mathcal{S} (equivalently, functions $f : \mathcal{S} \rightarrow \mathbb{N}$). Vectors appear in boldface, and we reserve uppercase $\mathbf{A} \in \mathbb{N}^{\mathcal{S}}$ for such vectors indexed by species, and lowercase $\mathbf{a} \in \mathbb{N}^d, \mathbb{Z}^d, \mathbb{Q}^d, \mathbb{R}^d$ for vectors indexed by integers. $\mathbf{A}(S)$ or $\mathbf{a}(i)$ denotes the element indexed by $S \in \mathcal{S}$ or $i \in \{1, \dots, d\}$. We write $\mathbf{a} \leq \mathbf{b}$ to denote pointwise vector inequality $\mathbf{a}(i) \leq \mathbf{b}(i)$ for all i .

For $p \in \mathbb{N}_+$, $\mathbb{Z}/p\mathbb{Z}$ denotes the additive group of integers modulo p , whose elements are congruence classes. Generalizing to higher dimensions, $\mathbb{Z}^d/p\mathbb{Z}^d$ denotes the additive group of \mathbb{Z}^d modulo p , whose elements are congruence classes. For $\mathbf{x} \in \mathbb{N}^d$ where $d \geq 1$, we write $\bar{\mathbf{x}} \pmod p$ to denote the congruence class $\{\mathbf{x} + p\mathbf{z} : \mathbf{z} \in \mathbb{Z}^d\} \in \mathbb{Z}^d/p\mathbb{Z}^d$, also denoted $\bar{\mathbf{x}}$ when p is clear from context.

$\mathbb{R}_{\geq 0}^d$ denotes the nonnegative orthant in \mathbb{R}^d . We consider *regions* $R = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^d : T\mathbf{x} - \mathbf{h} \geq 0\}$ which are convex polyhedra given by a set of inequalities. $R \cap \mathbb{N}^d$ denotes all integer points in R , and for $\mathbf{x} \in \mathbb{N}^d$, $R \cap (\bar{\mathbf{x}} \pmod p)$ denotes the integer points in R in the same congruence class as \mathbf{x} .

2.2 Chemical reaction networks

We use the established definitions of stable function computation by (discrete) chemical reaction networks [4, 13]:

A *chemical reaction network* (CRN) $\mathcal{C} = (\mathcal{S}, \mathcal{R})$ is defined by a finite set \mathcal{S} of species and a finite set \mathcal{R} of reactions, where a reaction $(\mathbf{R}, \mathbf{P}) \in \mathbb{N}^{\mathcal{S}} \times \mathbb{N}^{\mathcal{S}}$ describes the counts of consumed *reactant* species and produced *product* species.⁵ For example, given $\mathcal{S} = \{A, B, C\}$, the reaction $((1, 0, 2), (0, 2, 1))$ would represent $A + 2C \rightarrow 2B + C$.

A configuration $\mathbf{C} \in \mathbb{N}^{\mathcal{S}}$ specifies the integer counts of all species. Reaction (\mathbf{R}, \mathbf{P}) is *applicable* to \mathbf{C} if $\mathbf{R} \leq \mathbf{C}$, and yields $\mathbf{C}' = \mathbf{C} - \mathbf{R} + \mathbf{P}$, so we write $\mathbf{C} \rightarrow \mathbf{C}'$. A configuration \mathbf{D} is *reachable* from \mathbf{C} if there exists a finite sequence of configurations such that $\mathbf{C} \rightarrow \mathbf{C}_1 \rightarrow \dots \rightarrow \mathbf{C}_n \rightarrow \mathbf{D}$; we write $\mathbf{C} \rightarrow^* \mathbf{D}$ to denote that \mathbf{D} is reachable from \mathbf{C} . Note this reachability relation is additive: if $\mathbf{A} \rightarrow^* \mathbf{B}$, then $\mathbf{A} + \mathbf{C} \rightarrow^* \mathbf{B} + \mathbf{C}$. This property is key in future proofs to show the reachability of configurations which overproduce output.

To compute a function⁶ $f : \mathbb{N}^d \rightarrow \mathbb{N}$, the CRN \mathcal{C} will include an ordered subset $\{X_1, \dots, X_d\} \subset \mathcal{S}$ of *input species*, an *output species* Y , and a *leader species* $L \in \mathcal{S}$. (Note that we consider removing the leader in Section 9).

The computation of $f(\mathbf{x})$ will start from an *initial configuration* $\mathbf{I}_{\mathbf{x}}$ encoding the input with $\mathbf{I}_{\mathbf{x}}(X_i) = \mathbf{x}(i)$ for all $i = 1, \dots, d$, along with a single leader $\mathbf{I}_{\mathbf{x}}(L) = 1$, and count 0 of all other species. A *stable* configuration \mathbf{C} has unchanged output $\mathbf{C}(Y) = \mathbf{D}(Y)$ for any configuration \mathbf{D} reachable from \mathbf{C} . The CRN \mathcal{C} *stably computes* $f : \mathbb{N}^d \rightarrow \mathbb{N}$ if for each initial configuration $\mathbf{I}_{\mathbf{x}}$ encoding any $\mathbf{x} \in \mathbb{N}^d$, and configuration \mathbf{C} reachable from $\mathbf{I}_{\mathbf{x}}$, there is a stable configuration \mathbf{O} reachable from \mathbf{C} with correct output $\mathbf{O}(Y) = f(\mathbf{x})$.

⁵ We do not limit ourselves to *bimolecular* (two input) reactions, but the higher-order reactions we use can easily be converted to have this form. For example, $3X \rightarrow Y$ is equivalent to two reactions $2X \leftrightarrow X_2$ and $X + X_2 \rightarrow Y$.

⁶We consider codomain \mathbb{N} without loss of generality, since $f : \mathbb{N}^d \rightarrow \mathbb{N}^l$ is stably computable if and only if each output component is stably computable by parallel CRNs.

2.3 Composition via output-oblivious CRNs

This section formally defines our notions of “composable computation with CRNs via concatenation of reactions” and “output-oblivious” CRNs that don’t consume their output, showing these notions to be essentially equivalent.

A CRN is *output-oblivious* if the output species Y is never a reactant⁷: for any reaction (\mathbf{R}, \mathbf{P}) , $\mathbf{R}(Y) = 0$. A function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is *obliviously-computable* if f is stably computed by an output-oblivious CRN.

We begin with an easy observation:

Observation 2.1. *An obliviously-computable function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ must be nondecreasing.*

Proof. Assume a CRN \mathcal{C} (with output species Y) stably computes f , but $f(\mathbf{a}) > f(\mathbf{b})$ for $\mathbf{a} \leq \mathbf{b}$. To stably compute $f(\mathbf{a})$, input configuration $\mathbf{I}_{\mathbf{a}} \rightarrow^* \mathbf{O}$ for some configuration with $\mathbf{O}(Y) = f(\mathbf{a})$. However, since $\mathbf{a} \leq \mathbf{b}$, that same sequence of reactions can be applied from the input configuration $\mathbf{I}_{\mathbf{b}} \geq \mathbf{I}_{\mathbf{a}}$. This overproduces Y since $f(\mathbf{a}) > f(\mathbf{b})$. Thus to stably compute $f(\mathbf{b})$, some reaction must consume Y as a reactant, so \mathcal{C} cannot be output-oblivious. \square

A CRN being output-oblivious was shown in [9] (for *continuous* CRNs) to be equivalent to being “composable via concatenation”, meaning renaming the output species of one CRN to match the input of another. This equivalence still holds in our discrete CRN model. This is formalized as Observation 2.2 and Lemma 2.3.

For CRNs \mathcal{C}_f stably computing $f : \mathbb{N}^d \rightarrow \mathbb{N}$ and \mathcal{C}_g stably computing $g : \mathbb{N} \rightarrow \mathbb{N}$, define the *concatenated* CRN $\mathcal{C}_{g \circ f}$ by combining species and reactions, with \mathcal{C}_f ’s output species as \mathcal{C}_g ’s input species and no other common species, plus a reaction $L \rightarrow L^f + L^g$ creating a copy of the leader from each of \mathcal{C}_f and \mathcal{C}_g .

We first observe that this composition works correctly if the *upstream* CRN \mathcal{C}_f is output-oblivious. Intuitively, the reactions from \mathcal{C}_g can only affect the reactions from \mathcal{C}_f via the common species W , but this output species of \mathcal{C}_f is never used as a reactant to stably compute $f(\mathbf{x})$.

Observation 2.2. *If \mathcal{C}_f stably computes $f : \mathbb{N}^d \rightarrow \mathbb{N}$, \mathcal{C}_g stably computes $g : \mathbb{N} \rightarrow \mathbb{N}$, and \mathcal{C}_f is output-oblivious, then the concatenated CRN $\mathcal{C}_{g \circ f}$ stably computes the composition $g \circ f : \mathbb{N}^d \rightarrow \mathbb{N}$.*

Note that the downstream CRN \mathcal{C}_g need not be output-oblivious, but if two output-oblivious CRNs are composed, then the composition $\mathcal{C}_{g \circ f}$ remains output-oblivious. More generally, g can take any number of inputs from output-oblivious CRNs, which act as modules for arbitrary feed-forward composition.

The converse shows that a composable CRN is essentially output-oblivious. If \mathcal{C}_f can be correctly composed with any downstream \mathcal{C}_g , then \mathcal{C}_f must function correctly even if downstream reactions from \mathcal{C}_g starve it of the common species W . Thus \mathcal{C}_f will still stably compute f if we remove all reactions with output W as a reactant, making it output-oblivious.

Lemma 2.3. *Let \mathcal{C}_f stably compute $f : \mathbb{N}^d \rightarrow \mathbb{N}$ such that for any \mathcal{C}_g stably computing $g : \mathbb{N} \rightarrow \mathbb{N}$, the concatenated CRN $\mathcal{C}_{g \circ f}$ stably computes the composition $g \circ f : \mathbb{N}^d \rightarrow \mathbb{N}$. Then \mathcal{C}_f still stably computes f if we remove all reactions using the output species as a reactant, making it output-oblivious.*

⁷ A more general definition in [13] of *output-monotonic* CRNs just requires no reaction to reduce the count of output species. This can be directly seen to classify the same set of functions, see Observation 2.4.

Proof. Let C_f (with output species W) stably compute $f : \mathbb{N}^d \rightarrow \mathbb{N}$. Let $g(x) = x$ be the identity function, stably computed by $W \rightarrow Y$. Assume the concatenated CRN $C_{g \circ f}$ stably computes $g \circ f = f$. Let C'_f be the output-oblivious CRN with all reactions using W as a reactant removed from C_f . We now show that C'_f also stably computes f .

For any $\mathbf{x} \in \mathbb{N}^d$, let $\mathbf{I}_{\mathbf{x}}$ be the initial configuration encoding \mathbf{x} in C'_f , and \mathbf{C} be a configuration reachable from $\mathbf{I}_{\mathbf{x}}$. Now we can naturally view $\mathbf{I}_{\mathbf{x}}$ and \mathbf{C} also as configurations in the concatenated CRN $C_{g \circ f}$. We first consider configuration \mathbf{D} reachable from \mathbf{C} by applying reaction $W \rightarrow Y$ $\mathbf{C}(W)$ until $\mathbf{D}(W) = 0$. Now because $C_{g \circ f}$ stably computes f , there exists a sequence of reactions α from \mathbf{D} to a stable configuration \mathbf{O} with $\mathbf{O}(Y) = f(\mathbf{x})$. α could contain reactions that use W as a reactant, but because $\mathbf{D}(W) = 0$, any such reactions must occur after additional W has been produced.

Before any such reactions using W as a reactant, we can insert more reactions $W \rightarrow Y$, reaching an intermediate configuration \mathbf{D}_2 again with $\mathbf{D}_2(W) = 0$. There must exist some new sequence of reactions α_2 from \mathbf{D}_2 to a stable configuration \mathbf{O}_2 with $\mathbf{O}_2(Y) = f(\mathbf{x})$. Repeating this process, we will eventually find a sequence of reactions β from \mathbf{C} to a correct stable configuration \mathbf{O}_n . Notice that we only have to “splice in” new reactions when W is produced, and this can only happen at most $f(\mathbf{x})$ times, so this process will terminate.

Thus we have demonstrated a sequence of reactions β in $C_{g \circ f}$ from \mathbf{C} reaching a stable correct output configuration \mathbf{O}_n without using any reactions using W as a reactant. $\mathbf{O}_n(Y) = f(\mathbf{x})$, so β contains precisely $f(\mathbf{x})$ copies of the reaction $W \rightarrow Y$. Ignoring these reactions then gives a sequence of reactions β' in C'_f from \mathbf{C} to a correct configuration \mathbf{O}' with $\mathbf{O}'(W) = f(\mathbf{x})$. Notice that \mathbf{O}_n being stable in $C_{g \circ f}$ implies \mathbf{O}' is stable in C'_f since no additional W can be produced. This shows C'_f stably computes f as desired. \square

We finally observe that the more general definition of *output-monotonic* CRNs (which cannot decrease the count of the output species) stably compute precisely the same set of functions as output-oblivious CRNs:

Observation 2.4. $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is stably computable by an output-oblivious CRN \iff f is stably computable by an output-monotonic CRN.

Proof. \implies : Any output-oblivious CRN must be output-monotonic.

\impliedby : If f was stably computed by an output-monotonic CRN \mathcal{C} which is not already output-oblivious, then there must be reactions of the form $Y + \dots \rightarrow Y + \dots$ with output species Y acting as a catalyst. \mathcal{C} can be made output-oblivious by replacing all such occurrences of Y as a catalyst by a new catalyst species Z that is always produced alongside Y . Since \mathcal{C} was output-monotonic, if a Y is ever produced, it cannot be consumed. Thus any reactions with Y as a catalyst are “turned on” the moment the first Y is produced and never turn off again. So it does not change the reachable configurations to irreversibly produce a Z alongside Y and use Z as the catalyst in place of Y . This output-oblivious CRN thus also stably computes f . \square

2.4 Semilinear functions

The functions stably computable by a CRN were shown in [10], building from work in [4], to be precisely the *semilinear* functions, which are defined based on semilinear sets⁸

⁸Semilinear sets have other common equivalent definitions [3]; the above definition is convenient for our proof.

Definition 2.5. A subset $S \subseteq \mathbb{N}^d$ is semilinear if S is a finite Boolean combination (union, intersection, complement) of threshold sets of the form $\{\mathbf{x} \in \mathbb{N}^d : \mathbf{a} \cdot \mathbf{x} \geq b\}$ for $\mathbf{a} \in \mathbb{Z}^d, b \in \mathbb{Z}$ and mod sets of the form $\{\mathbf{x} \in \mathbb{N}^d : \mathbf{a} \cdot \mathbf{x} \equiv b \pmod{c}\}$ for $\mathbf{a} \in \mathbb{Z}^d, b \in \mathbb{Z}, c \in \mathbb{N}_+$.

A semilinear function can be concisely defined as having a semilinear graph, but a more useful equivalent definition comes from Lemma 4.3 of [10]⁹:

Definition 2.6 ([10]). A function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is semilinear if f is the finite union of affine partial functions, whose domains are disjoint semilinear subsets of \mathbb{N}^d .

All functions discussed have been semilinear. For example, the function

$$\min(x_1, x_2) = \begin{cases} x_1, & \text{if } x_1 \leq x_2 \\ x_2, & \text{if } x_1 > x_2 \end{cases}$$

is semilinear with affine partial functions on disjoint domains which are defined by a single threshold and thus semilinear.

Similarly, the function

$$\left\lfloor \frac{3x}{2} \right\rfloor = \begin{cases} \frac{3}{2}x, & \text{if } x \text{ is even} \\ \frac{3}{2}x - \frac{1}{2}, & \text{if } x \text{ is odd} \end{cases}$$

is semilinear with affine partial functions on disjoint domains which are defined by parity (a single mod predicate) and thus semilinear. All quilt-affine $f : \mathbb{N}^d \rightarrow \mathbb{N}$ are semilinear by the same argument.

Lemma 2.7 ([10]). A function $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is stably computable $\iff f$ is semilinear.

3 Warm-up: One-dimensional case

For functions with one-dimensional input, the necessary conditions of being nondecreasing and semilinear are also sufficient.

Theorem 3.1. $f : \mathbb{N} \rightarrow \mathbb{N}$ is obviously-computable $\iff f$ is semilinear and nondecreasing.

Intuitively, the proof works as follows. We show semilinear, nondecreasing $f : \mathbb{N} \rightarrow \mathbb{N}$ must have the eventually quilt-affine structure in Fig. 5. From this structure, we define a CRN that uses auxiliary leader states to track the value of x (or $\bar{x} \pmod{p}$ once $x \geq n$), while outputting the correct finite differences from adding each input.

Proof. \implies : Lemma 2.7 and Observation 2.1.

\impliedby : If $f : \mathbb{N} \rightarrow \mathbb{N}$ is semilinear and nondecreasing, it will eventually be *quilt-affine* (generalized to higher-dimensional functions as Definition 5.1) and thus have periodic finite differences: for some $n \in \mathbb{N}$, period $p \in \mathbb{N}_+$, and finite differences $\delta_0, \dots, \delta_{p-1} \in \mathbb{N}$, then for all $x \geq n$, $f(x+1) - f(x) = \delta_{(\bar{x} \pmod{p})}$ (see Fig. 5).

Because f is semilinear, by Definitions 2.5 and 2.6, it can be represented as a disjoint union of affine partial functions, whose domains are semilinear sets, and thus represented as finite Boolean

⁹Lemma 4.3 in [10] has domains that are non-disjoint linear sets. We assume the domains are disjoint for convenience, making the domains semilinear sets.

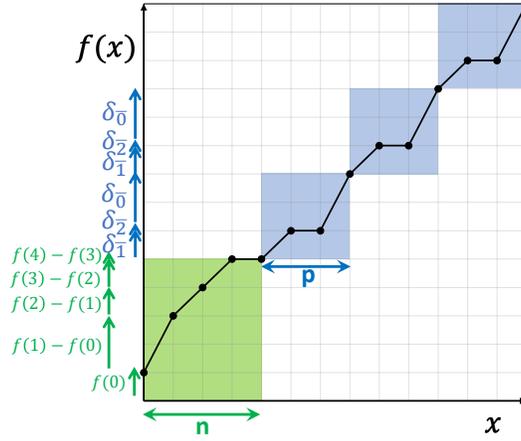
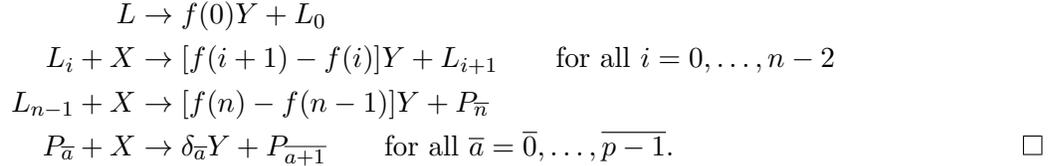


Figure 5: Every semilinear nondecreasing $f : \mathbb{N} \rightarrow \mathbb{N}$ is eventually quilt-affine, with periodic finite differences $\delta_{\bar{x}}$.

combinations of threshold $\{x \in \mathbb{N} : x \geq a\}$ and mod $\{x \in \mathbb{N} : x \equiv b \pmod{c}\}$ sets. Now take $n \in \mathbb{N}$ greater than all such a and $p = \text{lcm}(c)$ for all such c . Then for all $x \geq n$, f periodically cycles between affine partial functions. Because f is nondecreasing, these periodically-repeated affine partial functions must all have the same slope. This implies f is eventually quilt-affine, with periodic finite differences for all $x \geq n$ as claimed.

The CRN \mathcal{C} to stably compute f uses input species X , output species Y , leader L , and species $L_0, \dots, L_{n-1}, P_0, \dots, P_{p-1}$ corresponding to auxiliary “states” of the leader, i.e., exactly one of $L, L_0, \dots, L_{n-1}, P_0, \dots, P_{p-1}$ is present at any time. Intuitively the leader tracks how many input X it has seen, where the count past n wraps around mod p , and outputs the correct finite differences. The reactions of \mathcal{C} are as follows



In the 1D case, we can also characterize the functions obliviously-computable *without* a leader: they are semilinear and *superadditive*: meaning $f(x) + f(y) \leq f(x+y)$ for all x, y . (Theorem 9.2)

4 Impossibility result

The characterization of obliviously-computable functions as precisely semilinear and nondecreasing from Theorem 3.1 is insufficient in higher dimensions. As an example, consider the function $\max : \mathbb{N}^2 \rightarrow \mathbb{N}$, which is both semilinear and nondecreasing. We prove \max is not obliviously-computable via a more general lemma:

Lemma 4.1. *Let $f : \mathbb{N}^d \rightarrow \mathbb{N}$. If there exists an increasing sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots) \in \mathbb{N}^d$ such that for all $i < j$ there exists some $\Delta_{ij} \in \mathbb{N}^d$ with*

$$f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) > f(\mathbf{a}_j + \Delta_{ij}) - f(\mathbf{a}_j),$$

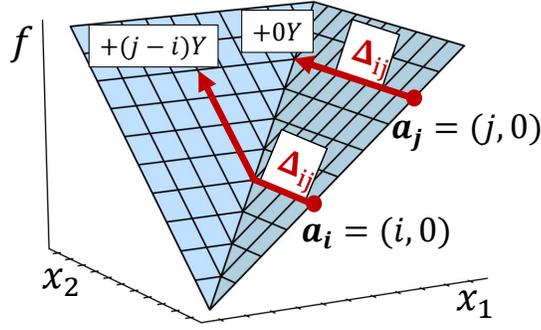


Figure 6: Lemma 4.1 applied to $f = \max(x_1, x_2)$.

then f is not obviously-computable.

Before proving Lemma 4.1, we use it to show \max is not obviously-computable.

For $f = \max(x_1, x_2)$, we let $\mathbf{a}_i = (i, 0)$ and $\Delta_{ij} = (0, j)$, so for $i < j$,

$$\max(i, j) - \max(i, 0) = j - i > \max(j, j) - \max(j, 0) = 0$$

as desired (see Fig. 6). Adding Δ_{ij} input after computing $f(\mathbf{a}_i)$ should produce $j - i$ additional output Y . However, adding Δ_{ij} input after computing $f(\mathbf{a}_j)$ should not. Lemma 4.1 uses this to show there exists a reaction sequence that overproduces Y , thus \max is not obviously-computable. We now prove Lemma 4.1.

Proof. Assume toward contradiction an output-oblivious CRN \mathcal{C} stably computes f . To stably compute each $f(\mathbf{a}_i)$, the initial configuration $\mathbf{I}_{\mathbf{a}_i} \rightarrow^* \mathbf{O}_i$ for some configuration with $\mathbf{O}_i(Y) = f(\mathbf{a}_i)$, giving a sequence of configurations $(\mathbf{O}_i)_{i=1}^\infty$. By Dickson's Lemma [15], any sequence of nonnegative integer vectors has a nondecreasing subsequence, so there must be $\mathbf{O}_i \leq \mathbf{O}_j$ for some $i < j$. By assumption there exists $\Delta_{ij} \in \mathbb{N}^d$ such that

$$f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) > f(\mathbf{a}_j + \Delta_{ij}) - f(\mathbf{a}_j)$$

Now consider the initial configuration $\mathbf{I}_{\mathbf{a}_i + \Delta_{ij}} \geq \mathbf{I}_{\mathbf{a}_i}$, so define the difference $\mathbf{D} = \mathbf{I}_{\mathbf{a}_i + \Delta_{ij}} - \mathbf{I}_{\mathbf{a}_i} \in \mathbb{N}^S$. Then the same sequence of reactions $\mathbf{I}_{\mathbf{a}_i} \rightarrow^* \mathbf{O}_i$ is applicable to $\mathbf{I}_{\mathbf{a}_i + \Delta_{ij}}$ reaching configuration $\mathbf{C}_i = \mathbf{O}_i + \mathbf{D}$, with $\mathbf{C}_i(Y) = \mathbf{O}_i(Y) = f(\mathbf{a}_i)$. Then to stably compute $f(\mathbf{a}_i + \Delta_{ij})$ there must exist a further sequence of reactions α from \mathbf{C}_i that produce an additional $f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i)$ copies of output Y .

By the same argument, from initial configuration $\mathbf{I}_{\mathbf{a}_j + \Delta_{ij}}$ the configuration $\mathbf{C}_j = \mathbf{O}_j + \mathbf{D}$ is reachable, with $\mathbf{C}_j(Y) = \mathbf{O}_j(Y) = f(\mathbf{a}_j)$. Then since $\mathbf{O}_i \leq \mathbf{O}_j$, we have $\mathbf{C}_i \leq \mathbf{C}_j$, so the same sequence of reactions α is applicable to \mathbf{C}_j , reaching some configuration \mathbf{C}'_j with an additional $f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i)$ copies of output Y , so

$$\mathbf{C}'_j(Y) = f(\mathbf{a}_j) + f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) > f(\mathbf{a}_j + \Delta_{ij})$$

Then $\mathbf{I}_{\mathbf{a}_j + \Delta_{ij}} \rightarrow^* \mathbf{C}'_j$ overproduces Y , so the output-oblivious CRN \mathcal{C} cannot stably compute $f(\mathbf{a}_j + \Delta_{ij})$. \square

Lemma 4.1 is our main technical tool used to show that a particular semilinear, nondecreasing function is not obviously-computable, the key challenge in the impossibility direction of Theorem 5.2.

5 Main result: Full-dimensional case

To formally state our main result, Theorem 5.2, we must first define *quilt-affine* functions as the sum of a linear and periodic function (see Fig. 3b):

Definition 5.1. *A nondecreasing function $g : \mathbb{N}^d \rightarrow \mathbb{Z}$ is quilt-affine (with period p) if there exists $\nabla_g \in \mathbb{Q}_{\geq 0}^d$ and $B : \mathbb{Z}^d / p\mathbb{Z}^d \rightarrow \mathbb{Q}$ such that*

$$g(\mathbf{x}) = \nabla_g \cdot \mathbf{x} + B(\bar{\mathbf{x}} \pmod{p}).$$

We call ∇_g the *gradient* of g , and the periodic function B the *periodic offset*. Without loss of generality we have the same period p along all inputs, since p could be the least common multiple of the periods along each input component. Note that $\nabla_g \cdot \mathbf{x}$ and B can each be rational, but the sum $g(x) \in \mathbb{Z}$ will be integer-valued. We allow g to have negative output for technical reasons¹⁰, but in the case that g is quilt-affine with nonnegative output (i.e. $g : \mathbb{N}^d \rightarrow \mathbb{N}$), there is a simple output-oblivious CRN construction to stably compute g . The intuitive idea is to use a single leader that reacts with every input species sequentially, tracks the periodic value $\bar{\mathbf{x}} \pmod{p}$, and outputs the correct changes in g (Lemma 6.1).

Our main result has a recursive condition where we fix the input of a function $f : \mathbb{N}^d \rightarrow \mathbb{N}$. For each $i = 1, \dots, d$ and $j \in \mathbb{N}$, define the *fixed-input restriction*¹¹ $f_{[\mathbf{x}(i) \rightarrow j]} : \mathbb{N}^d \rightarrow \mathbb{N}$ of f for all $\mathbf{x} \in \mathbb{N}^d$ by $f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) = f(\mathbf{x}(1), \dots, \mathbf{x}(i-1), j, \mathbf{x}(i+1), \dots, \mathbf{x}(d))$.

We can now formally state our main result:

Theorem 5.2. *$f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable \iff*

- i) [nondecreasing] f is nondecreasing,*
- ii) [eventually-min] there exist quilt-affine $g_1, \dots, g_m : \mathbb{N}^d \rightarrow \mathbb{Z}$ and $\mathbf{n} \in \mathbb{N}^d$ such that for all $\mathbf{x} \geq \mathbf{n}$, $f(\mathbf{x}) = \min_k(g_k(\mathbf{x}))$, and*
- iii) [recursive] all fixed-input restrictions $f_{[\mathbf{x}(i) \rightarrow j]}$ are obviously-computable.*

We first prove that these conditions imply f is obviously-computable via a general CRN construction in Section 6.

The nondecreasing condition (i) is necessary by Observation 2.1. It is immediate to see the recursive condition (iii) is also necessary:

Observation 5.3. *If $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable, then any fixed-input restriction $f_{[\mathbf{x}(i) \rightarrow j]} : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable.*

¹⁰ The quilt-affine functions that describe f for large inputs may be negative on inputs close to the origin.

¹¹ We define $f_{[\mathbf{x}(i) \rightarrow j]}$ to have domain \mathbb{N}^d because it is notationally convenient to have the same domain as f , but $f_{[\mathbf{x}(i) \rightarrow j]}$ only has relevant input in $d-1$ of its input components, making condition (iii) recursive.

Proof. Let the output-oblivious CRN \mathcal{C} stably compute f . We define the output-oblivious CRN \mathcal{C}' to “hardcode” the input $\mathbf{x}(i) = j$ by modifying the reactions of \mathcal{C} . Replace all instances of the leader L and input species X_i by L' and X'_i respectively, then add the initial reaction $L \rightarrow jX'_i + L'$. It is straightforward to verify that \mathcal{C}' stably computes $f_{[\mathbf{x}(i) \rightarrow j]}$. \square

Then the remaining work (and biggest effort of this paper) is to show the necessity of the eventually-min condition (ii): that every obviously-computable function can be represented as eventually a minimum of a finite number of quilt-affine functions, which is shown as Theorem 7.1. Its proof relies on f being semilinear, nondecreasing, and not having any “contradiction sequences” to apply Lemma 4.1. Thus the proof of Theorem 7.1 also yields the following alternative characterization to Theorem 5.2:

Theorem 5.4. $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable \iff f is semilinear, nondecreasing, and has no sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots)$ meeting the conditions of Lemma 4.1.

This gives a “negative characterization” identifying behavior obviously-computable functions must avoid, whereas Theorem 5.2, is a “positive characterization” describing the allowable behavior of such functions. We include Theorem 5.4, though it is less descriptive of the function, because it may be useful in other contexts.

6 Construction

First we show that any quilt-affine function with nonnegative range is stably computed by an output-oblivious CRN:

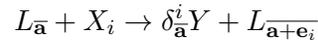
Lemma 6.1. Every quilt-affine function $g : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable.

Proof. Let $g : \mathbb{N}^d \rightarrow \mathbb{N}$ be quilt-affine with period p (recalling Definition 5.1). Notice that g has periodic finite differences. For each congruence class $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$ and input component $i = 1, \dots, d$, where \mathbf{e}_i is the i th standard basis vector, define

$$\delta_{\bar{\mathbf{a}}}^i = \nabla_g \cdot \mathbf{e}_i + B(\overline{\mathbf{a} + \mathbf{e}_i} \pmod p) - B(\bar{\mathbf{a}} \pmod p) \in \mathbb{N}.$$

Observe that for all $\mathbf{x} \in \bar{\mathbf{a}}$, $g(\mathbf{x} + \mathbf{e}_i) - g(\mathbf{x}) = \delta_{\bar{\mathbf{a}}}^i$. We now use these periodic finite differences to construct an output-oblivious CRN \mathcal{C} to stably compute g .

The CRN \mathcal{C} has input species X_1, \dots, X_d , output species Y , leader species L and p^d additional species $L_{\bar{\mathbf{a}}}$ for each $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$ corresponding to auxilliary “states” of the leader. The initial reaction $L \rightarrow g(\mathbf{0})Y + L_{\bar{\mathbf{0}}}$ is accompanied by dp^d reactions of the form



for each $i = 1, \dots, d$ and $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$. This CRN first creates $g(\mathbf{0})$ output, then sequentially outputs all finite differences, and is easily verified to stably compute g . \square

We now prove (in Lemma 6.2) one direction of Theorem 5.2: that conditions (i), (ii), and (iii) imply an output-oblivious CRN can stably compute f . Intuitively, by the eventually-min condition (ii) we compute $f(\mathbf{x})$ for $\mathbf{x} \geq \mathbf{n}$ by composing min and quilt-affine functions. If $\mathbf{x} \not\geq \mathbf{n}$, then $\mathbf{x}(i) = j$ for some input i and $j < \mathbf{n}(i)$. By the recursive condition (iii) we compute

$f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) = f(\mathbf{x})$ ¹². The key remaining insight is a trick (similar to a proof in [9]) to compose these pieces using minimum and indicator functions.

The proof of Lemma 6.2 then expresses f as such a minimum of finitely many pieces. We justify that f is obviously-computable by showing that each piece is obviously-computable, since by Observation 2.2 obviously-computable functions are closed under composition.

Lemma 6.2. *If $f : \mathbb{N}^d \rightarrow \mathbb{N}$ satisfies the conditions of Theorem 5.2, f is obviously-computable.*

Proof. Assume $f : \mathbb{N}^d \rightarrow \mathbb{N}$ satisfies the conditions of Theorem 5.2. Then by eventually-min condition (ii), there exist quilt-affine $g_1, \dots, g_m : \mathbb{N}^d \rightarrow \mathbb{Z}$ and $\mathbf{n} \in \mathbb{N}^d$ (without loss of generality assume $\mathbf{n} = (n, \dots, n)$) such that $f(\mathbf{x}) = \min_k(g_k(\mathbf{x}))$ for all $\mathbf{x} \geq \mathbf{n}$.

Let $\mathbf{x} \vee \mathbf{n} = (\max(\mathbf{x}(1), n), \dots, \max(\mathbf{x}(d), n))$ denote the componentwise max of \mathbf{x} and \mathbf{n} . Let $\mathbb{1}_{\mathbf{x}(i) > j} : \mathbb{N}^d \rightarrow \{0, 1\}$ denote the indicator function that is 1 \iff its input \mathbf{x} obeys $\mathbf{x}(i) > j$. Recall $f_{[\mathbf{x}(i) \rightarrow j]}$ is the fixed-input restriction setting input $\mathbf{x}(i) = j$. We claim that f can be expressed as

$$f(\mathbf{x}) = \min \left[f(\mathbf{x} \vee \mathbf{n}), \underbrace{f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) + \mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) \cdot f(\mathbf{x} \vee \mathbf{n})}_{\substack{i=1, \dots, d \\ j=0, \dots, n-1}} \right]. \quad (1)$$

We first show $f \geq \min[\dots]$ since for all $\mathbf{x} \in \mathbb{N}^d$, $f(\mathbf{x})$ is achieved by some term. If $\mathbf{x} \geq \mathbf{n}$, then $f(\mathbf{x}) = f(\mathbf{x} \vee \mathbf{n})$. If $\mathbf{x} \not\geq \mathbf{n}$, there must be $\mathbf{x}(i) = j$ for some $i = 1, \dots, d$ and $j = 0, \dots, n-1$, so $f(\mathbf{x}) = f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) = f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) + \mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) \cdot f(\mathbf{x} \vee \mathbf{n})$ since the indicator is 0.

We next show $f \leq \min[\dots]$ since $f(\mathbf{x}) \leq$ each term for all $\mathbf{x} \in \mathbb{N}^d$. $f(\mathbf{x}) \leq f(\mathbf{x} \vee \mathbf{n})$ since $\mathbf{x} \leq (\mathbf{x} \vee \mathbf{n})$ and f is nondecreasing. When $\mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) = 1$, we then have $f(\mathbf{x}) \leq f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) + \mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) \cdot f(\mathbf{x} \vee \mathbf{n})$. If $\mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) = 0$, then $\mathbf{x}(i) \leq j$ so $f(\mathbf{x}) \leq f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x})$ since f is nondecreasing. Thus equation 1 holds as claimed.

It remains to show that f is obviously-computable. From Observation 2.2, output-oblivious CRNs are closed under composition, and equation 1 gives a method to express f as a composition of functions. Thus it suffices to show that each piece is obviously-computable. Specifically, we show the functions $\min : \mathbb{N}^k \rightarrow \mathbb{N}$ (for any k), $f(\mathbf{x} \vee \mathbf{n}) : \mathbb{N}^d \rightarrow \mathbb{N}$, $f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) : \mathbb{N}^d \rightarrow \mathbb{N}$, and $c(a, b, \mathbf{x}) = a + \mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) \cdot b : \mathbb{N}^{d+2} \rightarrow \mathbb{N}$ are each obviously-computable. Implicit in the composed CRN to stably compute f as the composition from equation 1 is the ‘‘fan out’’ operation where reactions of the form $X_i \rightarrow X_i^1, \dots, X_i^m$ create multiple copies of species X_i to be used as independent inputs to multiple ‘‘modules’’ in this composition.

$\min : \mathbb{N}^k \rightarrow \mathbb{N}$ is obviously-computable:

Consider the CRN with single reaction $X_1, \dots, X_k \rightarrow Y$, the natural generalization of two-input min from Fig. 1.

$f(\mathbf{x} \vee \mathbf{n}) : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable:

By condition (ii), $f(\mathbf{x} \vee \mathbf{n}) = \min_k(g_k(\mathbf{x} \vee \mathbf{n}))$ since $\mathbf{x} \vee \mathbf{n} \geq \mathbf{n}$, so it suffices to show for each quilt-affine $g_k : \mathbb{N}^d \rightarrow \mathbb{Z}$ that $g_k(\mathbf{x} \vee \mathbf{n})$ is obviously-computable.

By condition (ii), $g_k(\mathbf{x} + \mathbf{n}) \geq f(\mathbf{x} + \mathbf{n}) \geq 0$ since $\mathbf{x} + \mathbf{n} \geq \mathbf{n}$. Then $g_k(\mathbf{x} + \mathbf{n}) : \mathbb{N}^d \rightarrow \mathbb{N}$ is still quilt-affine since that property is preserved by translation, but now has guaranteed nonnegative output. Thus by Lemma 6.1, $g_k(\mathbf{x} + \mathbf{n}) : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable.

¹²As a result, this construction is recursive, with an additional input being fixed at each level of the recursion, so the base case is simply a constant function.

Letting $(\mathbf{x} - \mathbf{n})_+ = (\max(\mathbf{x}(1) - n, 0), \dots, \max(\mathbf{x}(d) - n, 0))$, we then show the function $(\mathbf{x} - \mathbf{n})_+ : \mathbb{N}^d \rightarrow \mathbb{N}^d$ is obviously-computable via the CRN with reactions $(n+1)X_i \rightarrow nX_i + Y_i$ for each component $i = 1, \dots, d$.

Finally, because $\mathbf{x} \vee \mathbf{n} = (\mathbf{x} - \mathbf{n})_+ + \mathbf{n}$, we have shown $g_k(\mathbf{x} \vee \mathbf{n}) = g_k((\mathbf{x} - \mathbf{n})_+ + \mathbf{n})$ is obviously-computable as the composition of obviously-computable $g_k(\mathbf{x} + \mathbf{n})$ and $(\mathbf{x} - \mathbf{n})_+$.

$f_{[\mathbf{x}(i) \rightarrow j]}(\mathbf{x}) : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable:

This is precisely the assumed recursive condition (iii).

$c(a, b, \mathbf{x}) = a + \mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) \cdot b : \mathbb{N}^{d+2} \rightarrow \mathbb{N}$ is obviously-computable:

Consider the output-oblivious CRN (with input species A, B, X_1, \dots, X_d and output species Y) with two reactions $A \rightarrow Y$ and $(j+1)X_i + B \rightarrow (j+1)X_i + Y$. The A is all converted to Y , and $(j+1)$ copies of input species X_i catalyze the conversion of B to Y , which will only happen when $\mathbb{1}_{\{\mathbf{x}(i) > j\}}(\mathbf{x}) = 1$. Thus this stably computes $c(a, b, \mathbf{x})$ as desired. \square

7 Output-oblivious implies eventually min of quilt-affine functions

To complete the proof of Theorem 5.2, it remains to show the necessity of the eventually-min condition (ii):

Theorem 7.1. *If $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable, then there exist quilt-affine $g_1, \dots, g_m : \mathbb{N}^d \rightarrow \mathbb{Z}$ and $\mathbf{n} \in \mathbb{N}^d$ such that for all $\mathbf{x} \geq \mathbf{n}$, $f(\mathbf{x}) = \min_k(g_k(\mathbf{x}))$.*

For the remainder of Section 7, we fix an obviously-computable $f : \mathbb{N}^d \rightarrow \mathbb{N}$, and Section 7 is devoted to finding g_1, \dots, g_m and \mathbf{n} satisfying Theorem 7.1.

7.1 Proof outline

Section 7.2. Since $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable, f is semilinear (recall Definition 2.6), and we first consider all threshold sets used to define the semilinear domains of the affine partial functions that define f . Each threshold set defines a hyperplane, and we use these hyperplanes to define *regions* (see Fig. 8a and Fig. 8c). We consider regions as subsets of $\mathbb{R}_{\geq 0}^d$, so they are convex polyhedra with useful geometric properties.¹³

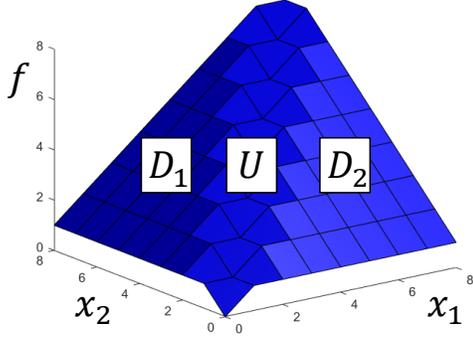
The regions partition¹⁴ the points in the domain \mathbb{N}^d . To prove Theorem 7.1, for each region R_k we will identify a quilt-affine function g_k (the *extension* of f from region R) such that $g(\mathbf{x}) = f(\mathbf{x})$ for all integer $\mathbf{x} \in R$. To ensure $f = \min_k(g_k)$, we further require that these quilt-affine extensions *eventually dominate* f (each $g_k(\mathbf{x}) \geq f(\mathbf{x})$ for sufficiently large \mathbf{x}). Also, because we only care about sufficiently large \mathbf{x} , we need only consider *eventual* regions which are unbounded in all inputs (for example regions 3, 4, and 5 in Fig. 8a).

As a simple motivating example, consider the semilinear, nondecreasing function

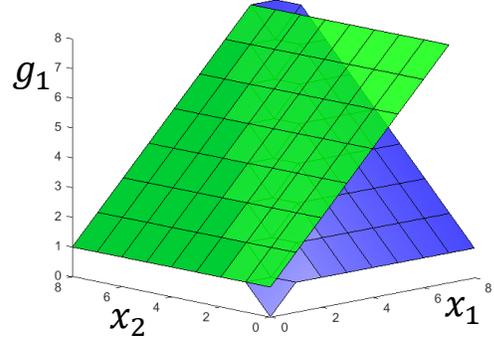
$$f(x_1, x_2) = \begin{cases} x_1 + 1, & \text{if } x_1 < x_2 \text{ (region } D_1) \\ x_2 + 1, & \text{if } x_1 > x_2 \text{ (region } D_2) \\ x_1 & \text{if } x_1 = x_2 \text{ (region } U) \end{cases}$$

¹³What we consider is a restricted case of a *hyperplane arrangement* [25], with well-studied combinatorial properties.

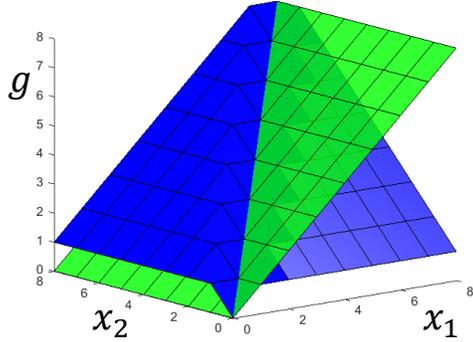
¹⁴Without loss of generality, we assume that the hyperplanes do not intersect \mathbb{N}^d , so that the partition is well-defined (see Fig. 8a).



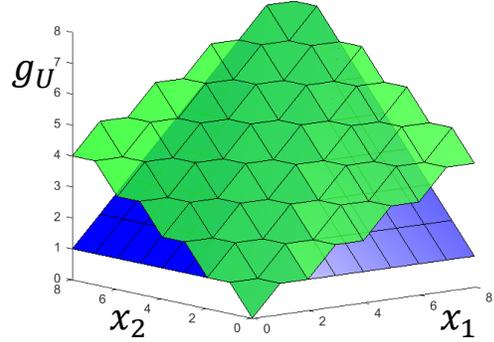
(a) Plot of f , whose domain has 3 regions: D_1 , D_2 , and U .



(b) g_1 (green) is the unique quilt-affine extension from region D_1 .



(c) g (green) is a quilt-affine extension from U , but $g < f$ on D_1 .



(d) g_U (green) is a quilt-affine extension from U , and $g_U \geq f$.

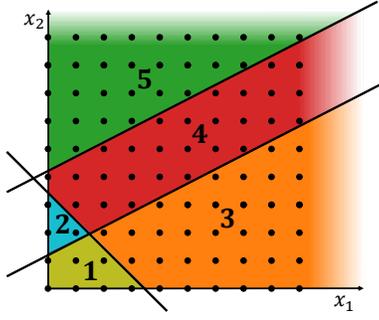
Figure 7: Obviously-computable f can be expressed as a min of quilt-affine functions.

As in Definition 2.6, f is piecewise-affine, with semilinear domains that happen to be only defined by threshold sets. These thresholds then partition the domain into three regions: D_1 , D_2 , and U (see Fig. 7a). For region D_1 , there is a unique quilt-affine extension $g_1(x_1, x_2) = x_1 + 1$ (note an affine function is the special case of a quilt-affine function with period 1). Also, g_1 eventually dominates f as desired, since $g_1(\mathbf{x}) \geq f(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{N}^2$ (see Fig. 7b). By symmetry, we have the same for region D_2 and its extension $g_2(x_1, x_2) = x_2 + 1$.

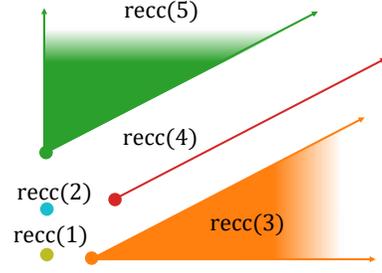
These desirable properties follow from D_1 and D_2 being “wide” regions that we define to be *determined* (formalized later). On the other hand, U is a “narrow” region that is *under-determined*. As a result, there is not a unique quilt-affine extension from U . For example, $g(x_1, x_2) = x_1$ is a quilt-affine extension, however, we do not have $g(\mathbf{x}) \geq f(\mathbf{x})$ for all sufficiently large \mathbf{x} (see Fig. 7c).

In order to identify a quilt-affine extension from U that does eventually dominate f , we will refer to the unique extensions g_1 and g_2 from regions D_1 and D_2 , which are *neighbors* of U (formalized later). We can construct a quilt-affine function with a gradient $(\frac{1}{2}, \frac{1}{2})$ that is the average of the gradients $(1, 0)$ of g_1 and $(0, 1)$ of g_2 . In particular, we can let $g_U(x_1, x_2) = \lceil \frac{x_1 + x_2}{2} \rceil$ (note this is a quilt-affine function with period 2, see Fig. 7d). We then have $f(\mathbf{x}) = \min [g_1(\mathbf{x}), g_2(\mathbf{x}), g_U(\mathbf{x})]$ for all $\mathbf{x} \geq \mathbf{n} = \mathbf{0}$ as guaranteed by Theorem 7.1.

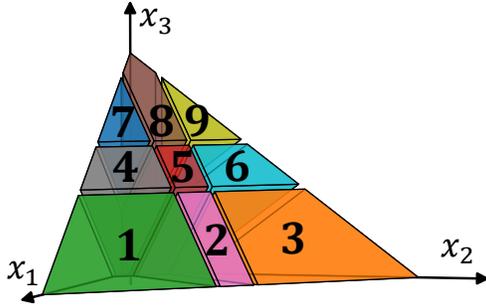
We now describe how we formalize the notion of a *determined region*, *under-determined region*,



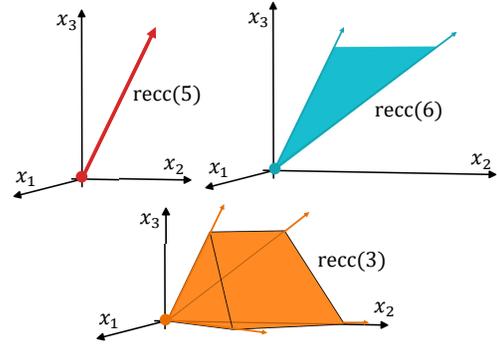
(a) Three threshold hyperplanes creating five regions. Regions 3 and 5 are determined, region 4 is under-determined but still eventual (unbounded in all input).



(b) The recession cones of all five regions. For finite regions, $\text{recc}(1) = \text{recc}(2) = \{\mathbf{0}\}$. Under-determined region 4 has a 1D recession cone, determined regions 3 and 5 have 2D recession cones.



(c) Two pairs of parallel threshold hyperplanes creating nine eventual regions. Regions 1,3,7,9 are determined. Region 5 is under-determined with 1D recession cone. Regions 2,4,6,8 are under-determined with 2D recession cones.



(d) $\text{recc}(5) \subseteq \text{recc}(6) \subseteq \text{recc}(3)$ so region 3 is a determined neighbor of under-determined region 5 and under-determined region 6. Also, region 6 is a neighbor of region 5.

Figure 8: Examples with domains \mathbb{N}^2 (top) and \mathbb{N}^3 (bottom), with threshold hyperplanes giving regions (left), which are classified by their recession cones (right).

and *neighbor*, for the general case of domain \mathbb{N}^d , where the regions are convex polyhedra in \mathbb{R}^d .

Section 7.3. To formally define determined regions, we identify the *recession cone* $\text{recc}(R) \subseteq \mathbb{R}^d$ of each region R : the set of vectors along infinite rays in R [22] (see Fig. 8b and Fig. 8d). A determined region D is defined as having a d -dimensional recession cone (see regions 3 and 5 in Fig. 8a and regions 1,3,7,9 in Fig. 8c). For determined regions, we can prove (see Lemmas 7.7 and 7.9) there is a unique quilt-affine extension, which eventually dominates f .

Section 7.4. Under-determined regions are then defined as having a recession cone with dimension $< d$ (see regions 1,2,4 in Fig. 8a and regions 2,4,5,6,8 in Fig. 8c). The above arguments do not work for under-determined regions. Instead, identify the *neighbors* of an under-determined region U as regions R with $\text{recc}(U) \subseteq \text{recc}(R)$ (see Fig. 8b and Fig. 8d). We consider the neighbors of U that are determined regions. The possible behavior of f on U is constrained by the unique extensions from these regions, and we can define an extension from U based on an averaging process. . (See Lemma 7.16).

7.2 Domain Decomposition

To identify the quilt-affine components g_k , the strategy will be to partition the domain \mathbb{N}^d into regions where the restriction of f to that region yields a quilt-affine function.

By Lemma 2.7, f is semilinear, so by Definition 2.6, f is the union of affine partial functions, whose disjoint domains are semilinear subsets of \mathbb{N}^d . This representation is not unique, so we now fix some arbitrary such representation of f . Recall by Definition 2.5, these semilinear domains are finite Boolean combinations of threshold and mod sets, so consider the collection \mathcal{T} of all threshold sets and the collection \mathcal{M} of all mod sets that defined any of these semilinear domains.

Let \mathcal{T} consist of l threshold sets $\{x \in \mathbb{N}^d : \mathbf{t}_i \cdot \mathbf{x} \geq h_i\}$ for each $i = 1, \dots, l$, where $\mathbf{t}_i \in \mathbb{Z}^d$ and $h_i \in \mathbb{Z}$. These thresholds are equivalently written $2\mathbf{t}_i \cdot \mathbf{x} > 2h_i - 1$ (since $\mathbf{t}_i \cdot \mathbf{x} \geq h_i \iff \mathbf{t}_i \cdot \mathbf{x} > h_i - \frac{1}{2}$), so we can assume without loss of generality that the boundary hyperplanes $H_i = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{t}_i \cdot \mathbf{x} = h_i\}$ contain no integer points (see Fig. 8a). These hyperplanes then partition the domain \mathbb{N}^d . For some $\mathbf{y} \in \mathbb{N}^d$, let $s_i = \text{sign}(\mathbf{t}_i \cdot \mathbf{y} - h_i) = \pm 1$ for each $i = 1, \dots, l$. Defining the threshold matrix $T \in \mathbb{Z}^{l \times d}$, offset vector $\mathbf{h} \in \mathbb{Z}^l$, and diagonal sign matrix $S \in \mathbb{Z}^{l \times l}$ as

$$T = \begin{pmatrix} \mathbf{t}_1^\top \\ \vdots \\ \mathbf{t}_l^\top \end{pmatrix} \quad \mathbf{h} = \begin{pmatrix} h_1 \\ \vdots \\ h_l \end{pmatrix} \quad S = \begin{pmatrix} s_1 & 0 & \dots & 0 \\ 0 & s_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & s_l \end{pmatrix}$$

then $S(T\mathbf{y} - \mathbf{h}) \geq \mathbf{0}$. This concise form will let us define the region of points that are in precisely the same threshold sets as \mathbf{y} (those that agree on the signs of the components of $T\mathbf{y} - \mathbf{h}$):

Definition 7.2. *Let S be a sign matrix: a diagonal matrix with diagonal entries $= \pm 1$. Then the region (induced by S) is defined as*

$$R = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^d : S(T\mathbf{x} - \mathbf{h}) \geq \mathbf{0}\}.$$

When referring to a region R , we use S_R to denote the sign matrix that induces R .

We consider nonnegative real vectors $\mathbf{x} \in \mathbb{R}^d$ rather than just $\mathbf{x} \in \mathbb{N}^d$, but we are only truly concerned with the integer points $R \cap \mathbb{N}^d$, and only consider regions where $R \cap \mathbb{N}^d \neq \emptyset$. Also, since each $\mathbf{y} \in \mathbb{N}^d$ induces a unique sign matrix as shown above, it follows that every $\mathbf{y} \in \mathbb{N}^d$ is contained in some unique region. The reason to consider $R \subset \mathbb{R}^d$ is that each region R is a convex polyhedron, with convenient properties from convex geometry (see Fig. 8a and Fig. 8c).

Now consider the collection \mathcal{M} , consisting of m mod sets $\{\mathbf{x} \in \mathbb{N}^d : \mathbf{a}_i \cdot \mathbf{x} \equiv b_i \pmod{c_i}\}$ for each $i = 1, \dots, m$, where $\mathbf{a}_i \in \mathbb{Z}^d, b_i \in \mathbb{Z}, c_i \in \mathbb{N}_+$. Then let the *global period* p be the least common multiple $\text{lcm}_i(c_i)$, so all elements of a congruence class $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$ are contained in precisely the same mod sets. Thus for a region R , the set $R \cap \bar{\mathbf{a}}$ is contained in precisely the same threshold and mod sets, so the restriction $f|_{R \cap \bar{\mathbf{a}}}$ is an affine partial function.

We now summarize this decomposition as a characterization of a semilinear function. Note that this applies to all semilinear functions, even those that are decreasing.

Lemma 7.3. *Let $f : \mathbb{N}^d \rightarrow \mathbb{N}$ be a semilinear function. Then there exist a finite set of regions $R_1, \dots, R_n \subseteq \mathbb{R}^d$ and a global period $p \in \mathbb{N}_+$ such that for each region R_i and congruence class $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$, there exist $\nabla_{R_i, \bar{\mathbf{a}}} \in \mathbb{Q}^d$ and $b_{R_i, \bar{\mathbf{a}}} \in \mathbb{Q}$ such that the restriction of f defined for all $\mathbf{x} \in R_i \cap \bar{\mathbf{a}}$ by*

$$f|_{R_i \cap \bar{\mathbf{a}}}(\mathbf{x}) = \nabla_{R_i, \bar{\mathbf{a}}} \cdot \mathbf{x} + b_{R_i, \bar{\mathbf{a}}}$$

is a (rational) affine partial function.

Notice the similarity in form between these affine partial functions in Lemma 7.3 and Definition 5.1 of quilt-affine functions. In fact, we will show the behavior of f on some regions has a unique quilt-affine structure. This will require the region to be “infinitely wide in all directions”, a notation we now make precise to define these determined regions.

7.3 Determined Regions

To formally define determined regions, we must first make the connection between regions and their *recession cones* (more information on recession cones can be found in [22]).

Definition 7.4. For a region R , the recession cone of R is

$$\text{recc}(R) = \{\mathbf{y} \in \mathbb{R}^d : \mathbf{x} + \lambda\mathbf{y} \in R \text{ for all } \mathbf{x} \in R, \lambda \in \mathbb{R}_{\geq 0}\}.$$

The recession cone corresponds to the directions (\mathbf{y}) one can proceed infinitely in a region R , and is a convex polyhedral cone (recall that a subset of \mathbb{R}^d is a *cone* if it is closed under positive scalar multiplication) (see Fig. 8b and Fig. 8d).

Recall by Definition 7.2 a region $R = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^d : S(T\mathbf{x} - \mathbf{h}) \geq \mathbf{0}\}$. It is possible to equivalently define

$$\text{recc}(R) = \{\mathbf{y} \in \mathbb{R}_{\geq 0}^d : S_R T\mathbf{y} \geq \mathbf{0}\}.$$

In other words, the recession cone is defined by the homogenized version of the same inequalities that defined R . One can easily verify this is equivalent to Definition 7.4.

We then say a region R is *determined* if $\dim \text{recc}(R) = d$. Otherwise, the region R is *under-determined*. We now make precise the idea that a determined region with full-dimensional recession cone is “infinitely wide in all directions”:

Lemma 7.5. Let D be a determined region. Then the recession cone $\text{recc}(D)$ contains open balls of arbitrarily large radius.

Proof. Since $\text{recc}(D)$ is a d -dimensional convex polyhedron, it has nonempty interior, so there exists some $\mathbf{x} \in \text{int}(\text{recc}(D))$ and an open ball $B_\epsilon(\mathbf{x}) \subset \text{recc}(D)$ of radius ϵ around \mathbf{x} contained in the cone. Since recession cones are closed under positive scalar multiplication, for any positive scalar c , the ball $B_{c\epsilon}(c\mathbf{x}) \subset \text{recc}(D)$. \square

We next make precise the idea that the function has a unique quilt-affine structure on a determined region.

Definition 7.6. An extension $g : \mathbb{N}^d \rightarrow \mathbb{Z}$ (of f) from a region R is a quilt-affine function that agrees with f on R : $f(\mathbf{x}) = g(\mathbf{x})$ for all integer $\mathbf{x} \in R \cap \mathbb{N}^d$.

We will now show there is a unique extension from each determined region. The construction of the regions yields a periodic piecewise-affine structure f restricted to a region (Lemma 7.3). In order for f to be nondecreasing, these affine gradients must all agree, which will let us uniquely describe a quilt-affine extension g using Definition 5.1.

Lemma 7.7. There is a unique extension $g : \mathbb{N}^d \rightarrow \mathbb{Z}$ from any determined region D .

Proof. By Definition 5.1, it suffices to show that there is a $\nabla_g \in \mathbb{Q}^d$ and $B : \mathbb{Z}^d/p\mathbb{Z}^d \rightarrow \mathbb{Q}$ such that, defining $g : \mathbb{N}^d \rightarrow \mathbb{Z}$ for all $\mathbf{x} \in \mathbb{N}^d$ as

$$g(\mathbf{x}) = \nabla_g \cdot \mathbf{x} + B(\bar{\mathbf{x}} \pmod p),$$

then for all $\mathbf{x} \in D \cap \mathbb{N}^d$, $g(\mathbf{x}) = f(\mathbf{x})$.¹⁵

By Lemma 7.3, for each $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$, there are $\nabla_{\bar{\mathbf{a}}} \in \mathbb{Q}^d$ and $b_{\bar{\mathbf{a}}} \in \mathbb{Q}$ such that

$$f|_{D \cap \bar{\mathbf{a}}}(\mathbf{x}) = \nabla_{\bar{\mathbf{a}}} \cdot \mathbf{x} + b_{\bar{\mathbf{a}}}.$$

Since D contains arbitrarily large open balls by Lemma 7.5, it contains points in all congruence classes, so all p^d constants $b_{\bar{\mathbf{a}}}$ are well-defined. Then we define the periodic offset function $B(\bar{\mathbf{x}} \pmod p) = b_{\bar{\mathbf{x}} \pmod p}$. This is almost what we require to prove the lemma, except that $\nabla_{\bar{\mathbf{a}}}$ depends on $\bar{\mathbf{a}}$. It remains to show that all vectors $\nabla_{\bar{\mathbf{a}}}$ are equal, so we can define the gradient $\nabla_g = \nabla_{\bar{\mathbf{a}}}$ for any $\bar{\mathbf{a}} \in \mathbb{Z}^d/p\mathbb{Z}^d$.

Assuming for the sake of contradiction that $\nabla_{\bar{\mathbf{a}}} \neq \nabla_{\bar{\mathbf{b}}}$ for some equivalence classes $\bar{\mathbf{a}}, \bar{\mathbf{b}}$, we will show that f cannot be nondecreasing. Since the recession cone $\text{recc}(D)$ is full-dimensional and $\nabla_{\bar{\mathbf{a}}} - \nabla_{\bar{\mathbf{b}}} \neq \mathbf{0}$, there must exist some $\mathbf{v} \in \text{recc}(D)$ such that $\nabla_{\bar{\mathbf{a}}} \cdot \mathbf{v} \neq \nabla_{\bar{\mathbf{b}}} \cdot \mathbf{v}$. Furthermore, by density of the rationals, we can assume $\mathbf{v} \in \mathbb{Q}^d$, and by scaling by the denominator we can assume $\mathbf{v} \in \mathbb{N}^d$. Now without loss of generality assume $\nabla_{\bar{\mathbf{a}}} \cdot \mathbf{v} > \nabla_{\bar{\mathbf{b}}} \cdot \mathbf{v}$.

Now pick some $\mathbf{y} \in D \cap \bar{\mathbf{a}}$, and some $\mathbf{z} \in D \cap \bar{\mathbf{b}}$ with $\mathbf{y} < \mathbf{z}$, which must exist because again D contains arbitrarily large open balls by Lemma 7.5. But since $\nabla_{\bar{\mathbf{a}}} \cdot \mathbf{v} > \nabla_{\bar{\mathbf{b}}} \cdot \mathbf{v}$, moving along \mathbf{v} , the function values in $\bar{\mathbf{a}}$ grow faster than in $\bar{\mathbf{b}}$ when moving along \mathbf{v} from \mathbf{y} and \mathbf{z} , respectively. Note that $\mathbf{y} + cp\mathbf{v}, \mathbf{z} + cp\mathbf{v} \in D$ by definition of $\mathbf{v} \in \text{recc}(D)$, and $cp\mathbf{v} \in p\mathbb{Z}^d$, so $\mathbf{y} + cp\mathbf{v} \in D \cap \bar{\mathbf{a}}$ and $\mathbf{z} + cp\mathbf{v} \in D \cap \bar{\mathbf{b}}$. Thus for some multiple cp of the period p , we must have

$$f(\mathbf{y} + cp\mathbf{v}) > f(\mathbf{z} + cp\mathbf{v})$$

but then f is not nondecreasing, since $\mathbf{y} + cp\mathbf{v} < \mathbf{z} + cp\mathbf{v}$.

Thus there is a uniquely determined gradient ∇_g . While there was not necessarily a unique choice for the period p , any valid choice will define the same function g . □

We next make precise why these extensions g_k can be used in the eventual-min that will define f to prove Theorem 7.1.

Definition 7.8. *An extension $g : \mathbb{N}^d \rightarrow \mathbb{Z}$ eventually dominates f if there exists $\mathbf{n} \in \mathbb{N}^d$ such that $f(\mathbf{x}) \leq g(\mathbf{x})$ for all $\mathbf{x} \geq \mathbf{n}$.*

We now show the uniquely defined extension g from a determined region D eventually dominates f . The idea is that if the extension g did not eventually dominate f , then we can apply Lemma 4.1 to show f is not obviously-computable.

For example, the function

$$\max(x_1, x_2) = \begin{cases} x_2, & \text{if } x_1 \leq x_2 \\ x_1, & \text{if } x_1 > x_2 \end{cases}$$

¹⁵Note that while p is not necessarily unique, any valid choice of p will result in the same function g .

is naturally identified by two determined regions, with unique extensions $g_1(\mathbf{x}) = (0, 1) \cdot \mathbf{x}$ and $g_2(\mathbf{x}) = (1, 0) \cdot \mathbf{x}$. These extensions do not eventually dominate f , and we already saw in Section 4 how Lemma 4.1 applies to \max . Thus the following lemma generalizes this example, finding a “contradiction sequence” $(\mathbf{a}_1, \mathbf{a}_2, \dots)$ to apply Lemma 4.1 whenever a determined extension g does not eventually dominate f :

Lemma 7.9. *The unique extension g from a determined region D eventually dominates f .*

Proof. Let g be the extension from some determined region D . Assume toward contradiction that g does not eventually dominate f , so for any point $\mathbf{n} \in \mathbb{N}^d$ there exists some “bad point” $\mathbf{b} \geq \mathbf{n}$ with $f(\mathbf{b}) > g(\mathbf{b})$. We will use Lemma 4.1 to show this implies f is not obviously-computable. To satisfy the Lemma conditions, we construct an increasing sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots) \in \mathbb{N}^d$ such that for all $i < j$ there exists some $\Delta_{ij} \in \mathbb{N}^d$ with

$$f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) > f(\mathbf{a}_j + \Delta_{ij}) - f(\mathbf{a}_j)$$

We will choose $(\mathbf{a}_1, \mathbf{a}_2, \dots) \in D \cap \bar{\mathbf{a}}$ to be points in D that are all in the same congruence class $\bar{\mathbf{a}} \pmod{p}$, and a sequence of vectors $(\mathbf{v}_1, \mathbf{v}_2, \dots) \in \mathbb{N}^d$ such that for all i , $\mathbf{a}_i + \mathbf{v}_i$ is a “bad point”: $f(\mathbf{a}_i + \mathbf{v}_i) > g(\mathbf{a}_i + \mathbf{v}_i)$, while for all $j > i$, $\mathbf{a}_j + \mathbf{v}_i \in D$, so $f(\mathbf{a}_j + \mathbf{v}_i) = g(\mathbf{a}_j + \mathbf{v}_i)$. Then for $i < j$, let $\Delta_{ij} = \mathbf{v}_i$, so

$$\begin{aligned} f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) &> g(\mathbf{a}_i + \mathbf{v}_i) - g(\mathbf{a}_i) \quad \text{since } \mathbf{a}_i + \mathbf{v}_i \text{ is a “bad point”} \\ &= g(\mathbf{a}_j + \mathbf{v}_i) - g(\mathbf{a}_j) \quad \text{since } g \text{ is quilt-affine and } \mathbf{a}_i \equiv \mathbf{a}_j \pmod{p} \\ &= f(\mathbf{a}_j + \Delta_{ij}) - f(\mathbf{a}_j) \end{aligned}$$

as desired, where $f = g$ for points in $\mathbf{a}_i, \mathbf{a}_j, \mathbf{a}_j + \mathbf{v}_i \in D$.

We now construct the sequences $(\mathbf{a}_1, \mathbf{a}_2, \dots)$ and $(\mathbf{v}_1, \mathbf{v}_2, \dots)$ recursively, to ensure for all $i < j$ that $\mathbf{a}_i \in \bar{\mathbf{a}}$, $\mathbf{a}_i + \mathbf{v}_i$ is a “bad point”, and $\mathbf{a}_j + \mathbf{v}_i \in D$. Let $\mathbf{a}_1 \in D$ arbitrarily, and the fixed congruence class $\bar{\mathbf{a}} = \bar{\mathbf{a}}_1$. For each \mathbf{a}_i , there must be a “bad point” above \mathbf{a}_i , so we recursively define each \mathbf{v}_i based on \mathbf{a}_i such that $\mathbf{a}_i + \mathbf{v}_i$ is this “bad point”: $f(\mathbf{a}_i + \mathbf{v}_i) > g(\mathbf{a}_i + \mathbf{v}_i)$.

Now we recursively define each \mathbf{a}_j based on \mathbf{a}_{j-1} and $\mathbf{v}_1, \dots, \mathbf{v}_{j-1}$ to ensure the sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots)$ is increasing, congruent, and $\mathbf{a}_j + \mathbf{v}_i \in D$ for all $i < j$. Since the recession cone $\text{recc}(D)$ contains open balls of arbitrary radius by Lemma 7.5, we can find a point \mathbf{a}_j in the recession cone above \mathbf{a}_{j-1} in the same congruence class $\bar{\mathbf{a}}$ such that the open ball $B_r(\mathbf{a}_j) \subset D$ for a large radius $r \geq \max_i |\mathbf{v}_i|$. This gives the desired condition $\mathbf{a}_j + \mathbf{v}_i \in D$ for all $i < j$.

By the above analysis, this increasing sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots)$ with $\Delta_{ij} = \mathbf{v}_i$ satisfies the conditions of Lemma 4.1, giving the contradiction that f is not obviously-computable. \square

The results of Lemmas 7.7 and 7.9 bring us close to proving Theorem 7.1. From each determined region D_1, \dots, D_q we have a quilt-affine extension $g_1, \dots, g_q : \mathbb{N}^d \rightarrow \mathbb{Z}$ that all eventually dominate f , so for some large enough $\mathbf{n} \in \mathbb{N}^d$ we have $f(\mathbf{x}) \leq \min_k g_k(\mathbf{x})$ for all $\mathbf{x} \geq \mathbf{n}$. Furthermore, $f(\mathbf{x}) = g_k(\mathbf{x})$ if $\mathbf{x} \in D_k$ for some determined region D_k . However, it is possible $f(\mathbf{x}) < \min_k g_k(\mathbf{x})$ for any \mathbf{x} that belong to an under-determined region. Since the bound \mathbf{n} can be arbitrarily large, we need only consider *eventual* under-determined regions that are unbounded in all inputs:

Definition 7.10. *A region R is eventual if for any $\mathbf{n} \in \mathbb{N}^d$, there exists some $\mathbf{x} \in \mathbb{N}^d \cap R$ such that $\mathbf{x} \geq \mathbf{n}$.*

To finish the proof of Theorem 7.1, it remains to show how to construct a quilt-affine extension g_U from each eventual under-determined region U , where g_U eventually dominates f .

7.4 Under-determined regions

Let U be an under-determined eventual region. The eventual condition implies that there is some $\mathbf{v} \in \text{recc}(U)$ strictly positive on all coordinates. Let $W = \text{span}(\text{recc}(U))$, which we call the *determined subspace* of U , with $1 \leq \dim W < d$. For example, for the “pizza slice” shaped region 6 in Fig. 8c, W is a 2D subspace, and region 6 is “infinitely wide” within the directions of W . This term reflects the fact that although the extensions from U are not unique, their values are uniquely defined *within* W . (For a determined region, its determined subspace is all of \mathbb{R}^d .)

Definition 7.11. *Region R is a neighbor of an under-determined region U if $\text{recc}(U) \subseteq \text{recc}(R)$.*

We will construct the extension from U by referencing the determined regions which are neighbors of U . Any under-determined eventual region will in fact have at least two determined neighbors (proved as Corollary 7.19 to the later Lemma 7.18). Geometrically, we can think of the under-determined recession cones as faces of each of the recession cones of the determined neighbors (see Fig. 8d).

Unlike in the proof of Lemma 7.7, the affine partial functions defining f within U do not need to have equal gradients. However, these gradients will be equal projected onto the subspace W . We now show a stronger statement, that this common gradient (projected onto W) agrees with the gradient of the extension from any determined neighbor D . Intuitively, if these gradients disagreed within W , then moving along directions in $\text{recc}(U)$, the differences between f on U and on D become arbitrarily large, contradicting that f must be nondecreasing.

Lemma 7.12. *Let U be an under-determined eventual region. Let D be a determined neighbor of U , with unique extension $g(\mathbf{x}) = \nabla_g \cdot \mathbf{x} + B(\bar{\mathbf{x}} \bmod p)$ given by Lemma 7.7. For any $\mathbf{u} \in U \cap \mathbb{N}^d$, consider the affine partial function $f|_{U \cap (\bar{\mathbf{u}} \bmod p)}(\mathbf{x}) = \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{x} + b_{\bar{\mathbf{u}}}$ given from Lemma 7.3. Then*

$$\text{proj}_W(\nabla_{\bar{\mathbf{u}}}) = \text{proj}_W(\nabla_g).$$

Proof. This proof uses similar techniques to the proof of Lemma 7.7: with two nonequal gradients, moving far enough along a recession cone direction to contradict the fact that f is nondecreasing.

Assume toward contradiction that for some $\mathbf{u} \in U \cap \mathbb{N}^d$, $\text{proj}_W(\nabla_{\bar{\mathbf{u}}}) \neq \text{proj}_W(\nabla_g)$. Then since $W = \text{span}(\text{recc}(U))$, we have $\nabla_{\bar{\mathbf{u}}} \cdot \mathbf{y} \neq \nabla_g \cdot \mathbf{y}$ for some $\mathbf{y} \in \text{recc}(U)$. Again, we can assume $\mathbf{y} \in \mathbb{N}^d$ by density of the rationals then scaling to clear denominators. Without loss of generality further assume $\nabla_{\bar{\mathbf{u}}} \cdot \mathbf{y} > \nabla_g \cdot \mathbf{y}$.

Now pick some $\mathbf{d} \in D \cap \mathbb{N}^d$ such that $\mathbf{d} \geq \mathbf{u}$, so $f(\mathbf{d}) \geq f(\mathbf{u})$ because f is nondecreasing. Then since $\mathbf{y} \in \text{recc}(U) \subset \text{recc}(D)$, for any $c \in \mathbb{N}$, $\mathbf{u} + c\mathbf{y} \in U \cap \bar{\mathbf{u}}$ and $\mathbf{d} + c\mathbf{y} \in D \cap \bar{\mathbf{d}}$. Since $f|_{U \cap \bar{\mathbf{u}}}(\mathbf{x}) = \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{x} + b_{\bar{\mathbf{u}}}$ and $f|_{D \cap \bar{\mathbf{d}}}(\mathbf{x}) = \nabla_g \cdot \mathbf{x} + B(\bar{\mathbf{d}})$, where $\nabla_{\bar{\mathbf{u}}} \cdot \mathbf{y} > \nabla_g \cdot \mathbf{y}$, for some large enough $c \in \mathbb{N}$, we have $f(\mathbf{u} + c\mathbf{y}) > f(\mathbf{d} + c\mathbf{y})$. But this contradicts the fact that f is nondecreasing, since $\mathbf{u} \leq \mathbf{d}$. \square

Lemma 7.12 constrains the behavior of f moving within the subspace W . The region U , however, could have a finite “width” in other directions. This motivates us to separate U into “strips”, partitioning its integer points to classes lying on translated versions of W :

Definition 7.13. *Let U be an under-determined region with $W = \text{span}(\text{recc}(U))$. The equivalence relation \equiv_W , where $\mathbf{x} \equiv_W \mathbf{y}$ if $\mathbf{x} - \mathbf{y} \in W$, partitions $U \cap \mathbb{N}^d$ into sets called strips. Thus a strip $I = \{\mathbf{x} \in U \cap \mathbb{N}^d : \mathbf{x} \equiv_W \mathbf{u}\}$ for some $\mathbf{u} \in U \cap \mathbb{N}^d$.*

For a strip I , we will consider the smallest affine set containing I : the *affine hull* $\text{aff}(I)$ [14]. Note that for every $\mathbf{u} \in I$, $\text{aff}(I) = \mathbf{u} + W = \{\mathbf{u} + \mathbf{w} : \mathbf{w} \in W\}$, and $\text{aff}(I)$ is a rational affine subspace. We next show a useful lemma about the distance from rational affine subspaces to surrounding integer points. This will be used to show there are only a finite number of strips, and will also be key to a trick in the later proof of Lemma 7.16.

Lemma 7.14. *Let $A \subset \mathbb{R}^d$ be a rational affine subspace, containing some point $\mathbf{x} \in \mathbb{N}^d$. Then there exists a constant $c > 0$ such that for any period $p^* \in \mathbb{N}_+$, for all $\mathbf{y} \in (\bar{\mathbf{x}} \bmod p^*)$ with $\mathbf{y} \notin A$, the distance $\text{dist}(\mathbf{y}, A) \geq cp^*$.*

Proof. Let A be a rational affine subspace containing $\mathbf{x} \in \mathbb{N}^d$. Let $\mathbf{y} \in \bar{\mathbf{x}} \bmod p^*$ with $\mathbf{y} \notin A$, so we can write $\mathbf{y} = \mathbf{x} + p^*\mathbf{v}$ for some $\mathbf{v} \in \mathbb{Z}^d$.

First we consider the case that A is a hyperplane, so we can write $A = \{\mathbf{z} \in \mathbb{R}^d : \mathbf{a} \cdot \mathbf{z} = b\}$ for some $\mathbf{a} \in \mathbb{Z}^d$ and $b \in \mathbb{Z}$. Then using a standard formula [12] for the distance from a point $\mathbf{y} \in \mathbb{R}^d$ to A :

$$\text{dist}(\mathbf{y}, A) = \frac{|\mathbf{a} \cdot \mathbf{y} - b|}{\|\mathbf{a}\|} = \frac{|\mathbf{a} \cdot (\mathbf{x} + p^*\mathbf{v}) - b|}{\|\mathbf{a}\|} = \frac{p^*|\mathbf{a} \cdot \mathbf{v}|}{\|\mathbf{a}\|} \geq \frac{p^*}{\|\mathbf{a}\|}$$

where the inequality follows from $|\mathbf{a} \cdot \mathbf{v}| \geq 1$: $\mathbf{a} \cdot \mathbf{v} \in \mathbb{Z}$ and $\mathbf{a} \cdot \mathbf{v} \neq 0$ since $\mathbf{y} \notin A$. The desired result then holds taking $c = 1/\|\mathbf{a}\|$, which depends only on A (and not on p^*).

Finally if the rational affine space A is not a hyperplane, it is the intersection of a finite number of hyperplanes, so A is contained in some rational hyperplane H and then by the above result $\text{dist}(\mathbf{y}, A) \geq \text{dist}(\mathbf{y}, H) \geq cp^*$. \square

Now we can show there are only a finite number of strips in each under-determined region.

Lemma 7.15. *The equivalence relation \equiv_W partitions $U \cap \mathbb{N}^d$ into a finite number of strips.*

Proof. Consider the set of unique strips $\{I_1, I_2, \dots\}$ each with some representative $\mathbf{u}_j \in I_j$. For each strip I_j , consider the affine hull $\text{aff}(I_j)$, which is a rational affine space, which are all parallel. For any $I_j \neq I_k$, since both $\text{aff}(I_j)$ and $\text{aff}(I_k)$ contain integer points, using Lemma 7.14 with $p = 1$ implies that $\text{dist}(\text{aff}(I_j), \text{aff}(I_k)) \geq c$ for some constant $c > 0$. This lower bound c is the same for all j, k because the $\text{aff}(I_j)$ are all parallel. The affine hulls of the strips being bounded away from each other will imply there can only be finitely many strips.

Since U is a convex polyhedron, we can write it as $U = H + \text{recc}(U) = \{\mathbf{h} + \mathbf{y} : \mathbf{h} \in H, \mathbf{y} \in \text{recc}(U)\}$, the sum of a bounded polytope H and the recession cone [14]. Thus each representative $\mathbf{u}_j = \mathbf{h}_j + \mathbf{y}_j$ for some $\mathbf{h}_j \in H$ and $\mathbf{y}_j \in \text{recc}(U)$, so $\mathbf{h}_j = \mathbf{u}_j - \mathbf{y}_j \in \text{aff}(I_j)$ (since $\mathbf{y}_j \in W$). Then $\text{dist}(\mathbf{h}_j, \mathbf{h}, k) \geq c$ for all $j \neq k$. Since all \mathbf{h}_j are contained in the bounded polytope H , there must be finitely many and thus finitely many strips. \square

Since there are a finite number of under-determined regions, by Lemma 7.15 there are a finite total number of strips, so it suffices to consider each strip separately and show there exists an extension from each strip that eventually dominates f .

We can build off Lemma 7.12 for a strip I to define an extension g_I from I that eventually dominates f . Intuitively, we will take the average gradient of the extensions from all determined neighbors. We crucially assume that the gradients of these extensions are not all the same. This will imply that their average grows faster than the minimum (and thus grows faster than f) moving away from $\text{aff}(I)$. It is then immediate that g_I will eventually dominate f sufficiently far from $\text{aff}(I)$, but requires a subtle trick to make this hold near $\text{aff}(I)$.

We choose g_I to have a larger period p^* that will guarantee (via Lemma 7.14) that points congruent to points in I are sufficiently far from $\text{aff}(I)$ (where $g_I \geq f$). The offsets in these congruence classes mod p^* are uniquely defined so $g_I = f$ on I . The remaining offsets are then maximized subject to the constraint that g_I is nondecreasing.

Finally, we must show g_I eventually dominates f on $\text{aff}(I)$. This is only nontrivial for $d \geq 3$. For example if I is a strip in the ‘‘pizza slice’’ shaped region 6 in Fig. 8c, we must argue that an extension from I will eventually dominate within the whole spanning plane. This is a generalization of Lemma 7.9, since I is essentially determined within $\text{aff}(I)$ (hence us calling W the determined subspace of U).

Lemma 7.16. *Let I be a strip of an under-determined eventual region U . Let D_1, \dots, D_m be the determined neighbors of U , with extensions g_1, \dots, g_m . Assume for all $\mathbf{z} \in W^\perp$, the gradients of the extensions along \mathbf{z} are not all equal: $\nabla_{g_i} \cdot \mathbf{z} \neq \nabla_{g_j} \cdot \mathbf{z}$ for some i, j . Then there exists an extension g_I from the strip I that eventually dominates f .*

Proof. Let D_1, \dots, D_m be determined neighbors of U , with a quilt-affine extension from each D_i $g_i(\mathbf{x}) = \nabla_{g_i} \cdot \mathbf{x} + B_{g_i}(\bar{\mathbf{x}} \bmod p)$. Let $\nabla_{\text{avg}} = \frac{1}{m} \sum_{i=1}^m \nabla_{g_i} \in \mathbb{Q}^d$. We will define $p^* \in \mathbb{N}_+$ and $B^* : \mathbb{Z}^d/p^*\mathbb{Z}^d \rightarrow \mathbb{Q}$ and let the extension be

$$g_I(\mathbf{x}) = \nabla_{\text{avg}} \cdot \mathbf{x} + B^*(\bar{\mathbf{x}} \bmod p^*),$$

which is quilt-affine with a potentially larger period $p^* = kp$ for some $k \in \mathbb{N}_+$, so $(\bar{\mathbf{x}} \bmod p^*) \subseteq (\bar{\mathbf{x}} \bmod p)$. To ensure that g_I still has integer outputs, we pick p^* such that $p^* \nabla_{\text{avg}} \in \mathbb{N}^d$. We will show later that p^* can be chosen large enough to make g_I eventually dominate f . Let $\mathbf{r} \in I$ be a fixed reference vector we will use to define B^* .

First we show that for each $\mathbf{u} \in I$, $B^*(\bar{\mathbf{u}} \bmod p^*)$ is uniquely defined so that $g_I(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in I \cap (\bar{\mathbf{u}} \bmod p^*)$ (in other words there are fixed values of B^* for inputs in I that will make g_I an extension of f from I). By Lemma 7.3 we have affine partial function $f|_{U \cap (\bar{\mathbf{u}} \bmod p)}(\mathbf{x}) = \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{x} + b_{\bar{\mathbf{u}}}$ and by Lemma 7.12 we have $\text{proj}_W(\nabla_{\bar{\mathbf{u}}}) = \text{proj}_W(\nabla_{g_i})$ for all gradients of determined neighbor extensions g_i . Thus we also have $\text{proj}_W(\nabla_{\bar{\mathbf{u}}}) = \text{proj}_W(\nabla_{\text{avg}})$, so $\nabla_{\bar{\mathbf{u}}} \cdot \mathbf{w} = \nabla_{\text{avg}} \cdot \mathbf{w}$ for all $\mathbf{w} \in W$. We then define $B^*(\bar{\mathbf{u}} \bmod p^*) = \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{r} - \nabla_{\text{avg}} \cdot \mathbf{r} + b_{\bar{\mathbf{u}}}$, which depends only on the congruence class $\bar{\mathbf{u}} \bmod p$ (but doesn’t depend on p^*). We can now verify that for any $\mathbf{x} \in I \cap (\bar{\mathbf{u}} \bmod p^*)$, where $\mathbf{x} - \mathbf{r} \in W$ by definition of the strip I , we have

$$\begin{aligned} g_I(\mathbf{x}) &= \nabla_{\text{avg}} \cdot \mathbf{x} + B^*(\bar{\mathbf{x}} \bmod p^*) \\ &= \nabla_{\text{avg}} \cdot (\mathbf{x} - \mathbf{r}) + \nabla_{\text{avg}} \cdot \mathbf{r} + \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{r} - \nabla_{\text{avg}} \cdot \mathbf{r} + b_{\bar{\mathbf{u}}} \\ &= \nabla_{\bar{\mathbf{u}}} \cdot (\mathbf{x} - \mathbf{r}) + \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{r} + b_{\bar{\mathbf{u}}} \quad \text{since } \mathbf{x} - \mathbf{r} \in W \\ &= \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{x} + b_{\bar{\mathbf{u}}} = f|_{U \cap (\bar{\mathbf{u}} \bmod p)}(\mathbf{x}) = f(\mathbf{x}) \end{aligned}$$

g_I is currently a partial function, only defined on the set $I^* = (I + p^*\mathbb{Z}^d) \cap \mathbb{N}^d$ of points congruent mod p^* to some $\mathbf{u} \in I$. For all other congruence classes $\bar{\mathbf{a}} \bmod p^* \in \mathbb{Z}^d/p^*\mathbb{Z}^d$ such that $\bar{\mathbf{a}} \cap I = \emptyset$, we will define $B^*(\bar{\mathbf{a}})$ to be as large as possible while still having g_I be nondecreasing. For g_I to be nondecreasing, $g_I(\mathbf{x}) \leq \min_{\mathbf{y} \geq \mathbf{x}} g_I(\mathbf{y})$ for all $\mathbf{x} \in \mathbb{N}^d$. We maximize $B^*(\bar{\mathbf{a}})$ such that for all $\mathbf{x} \in \bar{\mathbf{a}}$,

$$g_I(\mathbf{x}) = \min_{\mathbf{y} \in I^*, \mathbf{y} \geq \mathbf{x}} g_I(\mathbf{y})$$

Observe that since the finite differences above each \mathbf{x} are periodic (as observed formally to prove Lemma 6.1), this required offset $B^*(\bar{\mathbf{a}})$ depends only on the congruence class $\bar{\mathbf{a}}$ of \mathbf{x} .

Now in order to show that g_I eventually dominates f , we claim it suffices to show that g_I eventually dominates f on I^* : for some $\mathbf{n} \in \mathbb{N}^d$, $g_I(\mathbf{x}) \geq f(\mathbf{x})$ for all $\mathbf{x} \in I^*$ with $\mathbf{x} \geq \mathbf{n}$. If this holds, then for any $\mathbf{x} \notin I^*$ with $\mathbf{x} \geq \mathbf{n}$, we have $g_I(\mathbf{x}) = g_I(\mathbf{y})$ for some $\mathbf{y} \in I^*$, $\mathbf{y} \geq \mathbf{x}$, so

$$g_I(\mathbf{x}) = g_I(\mathbf{y}) \geq f(\mathbf{y}) \geq f(\mathbf{x})$$

showing that g_I eventually dominates f as long as g_I eventually dominates f on I^* .

We will next show g_I eventually dominates f for \mathbf{x} sufficiently far from $\text{aff}(I)$ by comparing g_I to the extension $g_j(\mathbf{x}) = \nabla_{g_j} \cdot \mathbf{x} + B_{g_j}(\bar{\mathbf{x}} \bmod p)$ from some determined neighbor D_j . Let $\mathbf{x} \in I^*$ and let g_j be the extension of any determined neighbor D_j . Writing $\mathbf{x} = \mathbf{r} + \mathbf{w} + \mathbf{z}$ for the fixed reference $\mathbf{r} \in I$, $\mathbf{w} = \text{proj}_W(\mathbf{x} - \mathbf{r}) \in W$ and $\mathbf{z} = \text{proj}_{W^\perp}(\mathbf{x} - \mathbf{r}) \in W^\perp$, we have

$$\begin{aligned} g_I(\mathbf{x}) - g_j(\mathbf{x}) &= \nabla_{\text{avg}} \cdot (\mathbf{r} + \mathbf{w} + \mathbf{z}) + B^*(\bar{\mathbf{x}} \bmod p^*) - \nabla_{g_j} \cdot (\mathbf{r} + \mathbf{w} + \mathbf{z}) - B_{g_j}(\bar{\mathbf{x}} \bmod p) \\ &= \nabla_{\text{avg}} \cdot \mathbf{z} - \nabla_{g_j} \cdot \mathbf{z} + [\nabla_{\text{avg}} \cdot \mathbf{r} - \nabla_{g_j} \cdot \mathbf{r} + B^*(\bar{\mathbf{x}} \bmod p^*) - B_{g_j}(\bar{\mathbf{x}} \bmod p)] \end{aligned}$$

Notice that the term $[\dots]$ depends only on j and $\bar{\mathbf{x}} \bmod p$ (since B^* was uniquely defined on I^* based only on $\bar{\mathbf{x}} \bmod p$). Thus minimizing over all finitely many j and $\bar{\mathbf{x}} \bmod p$ gives some (possibly negative) lower bound $-q \in \mathbb{Q}$ (which crucially does not depend on the choice p^*) such that

$$g_I(\mathbf{x}) - g_j(\mathbf{x}) \geq \nabla_{\text{avg}} \cdot \mathbf{z} - \nabla_{g_j} \cdot \mathbf{z} - q$$

for all $\mathbf{x} \in I^*$ and extensions g_j .

Now we use the crucial assumption that for any $\mathbf{z} \in W^\perp$, $\nabla_{g_i} \cdot \mathbf{z} \neq \nabla_{g_j} \cdot \mathbf{z}$ for some i, j . Considering first unit vectors $\mathbf{v} \in W^\perp$ with $\|\mathbf{v}\| = 1$, then there is some j minimizing $\nabla_{g_j} \cdot \mathbf{v}$ with $\nabla_{\text{avg}} \cdot \mathbf{v} - \nabla_{g_j} \cdot \mathbf{v} > 0$. We claim that there exists $\epsilon > 0$ such that $\nabla_{\text{avg}} \cdot \mathbf{v} - \nabla_{g_j} \cdot \mathbf{v} \geq \epsilon$ for all such \mathbf{v} and their corresponding j . If not, then there is a sequence (\mathbf{v}_i) of unit vectors with $\nabla_{\text{avg}} \cdot \mathbf{v}_i - \nabla_{g_j} \cdot \mathbf{v}_i \rightarrow 0$. Since the unit ball is compact, there must be a subsequence of (\mathbf{v}_i) converging to some \mathbf{v} , which implies $\nabla_{\text{avg}} \cdot \mathbf{v} - \nabla_{g_j} \cdot \mathbf{v} = 0$. This completes the claim that such $\epsilon > 0$ exists. Then as long as $\|\mathbf{z}\| \geq q/\epsilon$, we have for some j (which minimizes $\nabla_{g_j} \cdot \mathbf{z}$)

$$g_I(\mathbf{x}) - g_j(\mathbf{x}) \geq \|\mathbf{z}\| \left(\nabla_{\text{avg}} \cdot \frac{\mathbf{z}}{\|\mathbf{z}\|} - \nabla_{g_j} \cdot \frac{\mathbf{z}}{\|\mathbf{z}\|} \right) - q \geq \|\mathbf{z}\| \epsilon - q \geq 0$$

Since $\mathbf{x} = \mathbf{r} + \mathbf{w} + \mathbf{z}$ for some $\mathbf{r} + \mathbf{w} \in \text{aff}(I)$ and $\mathbf{z} \in W^\perp$, we have $\|\mathbf{z}\| = \text{dist}(\mathbf{x}, \text{aff}(I))$. Thus we have shown for \mathbf{x} sufficiently far from $\text{aff}(I)$, $g_I(\mathbf{x}) \geq g_j(\mathbf{x})$ for some quilt-affine g_j which itself eventually dominates f (by Lemma 7.9). Crucially this bound $\|\mathbf{z}\| \geq q/\epsilon$ did not depend on p^* , so we will now use Lemma 7.14 to choose p^* large enough that $\text{dist}(\mathbf{x}, \text{aff}(I)) \geq q/\epsilon$ for all $\mathbf{x} \in I^*$ with $\mathbf{x} \notin \text{aff}(I)$. Since such $\mathbf{x} \in (\bar{\mathbf{u}} \bmod p^*)$ for some $\mathbf{u} \in I \subset \text{aff}(I)$ and $\text{aff}(I)$ is a rational affine subspace, by Lemma 7.14 there is some bound $c > 0$ (depending only on $\text{aff}(I)$) such that $\text{dist}(\mathbf{x}, \text{aff}(I)) \geq cp^*$. Thus we choose a large enough multiple $p^* = kp$ such that $cp^* \geq q/\epsilon$.

We have shown that $g_I(\mathbf{x})$ eventually dominates $f(\mathbf{x})$ for all $\mathbf{x} \in I^*$ with $\mathbf{x} \notin \text{aff}(I)$. It finally remains to show that g_I eventually dominates f on $\text{aff}(I)$. This is true for the same reasons as Lemma 7.9 (because g_I is a quilt-affine extension of f from I) following the same proof strategy. In the proof of Lemma 7.9 we assumed toward contradiction a sequence of “bad points”, and used the fact that a determined region D contained arbitrarily large open balls to construct a contradiction sequence for Lemma 4.1. Now we are only showing g_I eventually dominates f on $\text{aff}(I)$, so all “bad points” would be in $\text{aff}(I)$. We can construct a contradiction sequence $(\mathbf{a}_i) \in I$ by the same

argument. Since $W = \text{span}(\text{recc}(U))$, by the same argument as Lemma 7.5, $\text{recc}(U)$ contains arbitrarily large open balls within the subspace W . This is sufficient to ensure $(\mathbf{a}_i) \in I$, since the sequence of vectors (\mathbf{v}_i) with $\mathbf{a}_i + \mathbf{v}_i$ being “bad points” are in W because all “bad points” are in $\text{aff}(I)$.

Thus g_I eventually dominates f everywhere, as desired. \square

Lemma 7.16 crucially assumed that along any $\mathbf{z} \in W^\perp$, the gradients of all extensions from determined neighbors are not equal. If this does not hold, f could fail to be obviously-computable. For example, consider the function

$$f(x_1, x_2) = \begin{cases} x_1 + x_2 + 1, & \text{if } x_1 \neq x_2 \\ x_1 + x_2, & \text{if } x_1 = x_2 \end{cases} \quad (2)$$

which is a single affine function, depressed by 1 along the diagonal $x_1 = x_2$. f is semilinear and nondecreasing. The two determined regions where $x_1 > x_2$ and $x_1 < x_2$ have the same quilt-affine extension $((1, 1) \cdot \mathbf{x} + 1)$, which eventually dominates f . On the underdetermined region, which here consists of just the single strip where $x_1 = x_2$, f is strictly smaller. There does not actually exist given a quilt-affine extension from this strip that eventually dominates f . (One can show directly f is not obviously-computable by Lemma 4.1, with $\mathbf{a}_i = (i, 0)$ and $\Delta_{ij} = (0, j)$).

The remaining case thus serves to disallow general versions of this counterexample. Lemma 7.16 assumed for all $\mathbf{z} \in W^\perp$, $\nabla_{g_i} \cdot \mathbf{z} \neq \nabla_{g_j} \cdot \mathbf{z}$ for some i, j . We will now consider the negation: that for some $\mathbf{z} \in W^\perp$ we have $\nabla_{g_i} = \nabla_{g_j}$ for all i, j . To proceed in this case, we will need to be able to identify the *neighbor of U in the direction of \mathbf{z}* .

For example, consider the under-determined eventual region 5 in Fig. 8c. The determined subspace W is 1D, while the orthogonal complement W^\perp is a 2D. For each $\mathbf{z} \in W^\perp$, the neighbor in the direction of \mathbf{z} will correspond to one of the 8 other pictured regions.

We now identify which threshold hyperplanes can distinguish a region from its neighbors. Recall the threshold hyperplanes $H_i = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{t}_i \cdot \mathbf{x} = h_i\}$ for $i = 1, \dots, l$. We now show that some of these hyperplanes must be parallel to all vectors in $\text{recc}(U)$:

Lemma 7.17. *Let U be an under-determined eventual region. Then there exists some threshold hyperplane $H_i = \{\mathbf{x} \in \mathbb{R}^d : \mathbf{t}_i \cdot \mathbf{x} = h_i\}$ such that $\mathbf{t}_i \cdot \mathbf{y} = 0$ for all $\mathbf{y} \in \text{recc}(U)$.*

Proof. Assume toward contradiction that for all \mathbf{t}_i , there exists some $\mathbf{y}_i \in \text{recc}(U)$ such that $\mathbf{t}_i \cdot \mathbf{y}_i \neq 0$, so $s_i(\mathbf{t}_i \cdot \mathbf{y}_i) > 0$, where s_i is the i th sign that defined the region U , and $s_j(\mathbf{t}_j \cdot \mathbf{y}_i) \geq 0$ for all $j = 1, \dots, l$ since $\mathbf{y}_i \in \text{recc}(U)$. Then let $\mathbf{y} = \sum_{i=1}^l \mathbf{y}_i \in \text{recc}(U)$ since $\text{recc}(U)$ is closed under addition, so $s_i(\mathbf{t}_i \cdot \mathbf{y}) > 0$ for all $i = 1, \dots, l$.

Recall that $\text{recc}(U) = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^d : (\forall i) s_i(\mathbf{t}_i \cdot \mathbf{x}) \geq 0\}$ is a closed convex polyhedron that is the intersection of closed half-spaces. Then $\text{int}(\text{recc}(U)) = \{\mathbf{x} \in \mathbb{R}_{> 0}^d : (\forall i) s_i(\mathbf{t}_i \cdot \mathbf{x}) > 0\}$ is the intersection of the respective open half-spaces, and we have $\mathbf{y} \in \text{int}(\text{recc}(U))$. Thus because $\text{int}(\text{recc}(U))$ is nonempty, $\text{recc}(U)$ must be full-dimensional, contradicting that U is an under-determined region. \square

We call such hyperplanes *neighbor-separating hyperplanes* for reasons that will be made clear shortly. If $\mathbf{t}_i \cdot \mathbf{y} = 0$ for all $\mathbf{y} \in \text{recc}(U)$, we also have $\mathbf{t}_i \cdot \mathbf{y} = 0$ for all $\mathbf{y} \in W$, so by definition $\mathbf{t}_i \in W^\perp$. Then let $L_U = \{i \in \{1, \dots, l\} \mid \mathbf{t}_i \in W^\perp\}$ be the subset of labels of all such neighbor-separating hyperplanes for U .

For example, in Fig. 8c, for under-determined region 5, all four hyperplanes are neighbor-separating hyperplanes. For under-determined region 6, only the pair of horizontally oriented hyperplanes are neighbor-separating.

Recalling Definition 7.2, let $S_U = \text{diag}(s_1, \dots, s_l)$ be the sign matrix that defined U . For $\mathbf{z} \in W^\perp$, define $R_{\mathbf{z}}$, the *neighbor of U in the direction of \mathbf{z}* , by a related sign matrix $S_{\mathbf{z}} = \text{diag}(s'_1, \dots, s'_l)$, where if $i \in L_U$ and $\text{sign}(\mathbf{t}_i \cdot \mathbf{z}) = -s_i$, then let $s'_i = -s_i$, but otherwise $s'_i = s_i$ for all other $i = 1, \dots, l$. Intuitively, for all neighbor-separating hyperplanes, $R_{\mathbf{z}}$ is on the same side as the direction \mathbf{z} , but is otherwise identical to U .

The following lemma justifies the use of the word “neighbor” in the previous definition, and further shows that such a neighbor is “more determined” (recession cone has higher dimension) than U .

Lemma 7.18. *Let U be an under-determined eventual region with $W = \text{span}(\text{recc}(U))$, let $\mathbf{z} \in W^\perp$, and let the region $R_{\mathbf{z}}$ be the neighbor of U in the direction of \mathbf{z} . Then $R_{\mathbf{z}} \cap \mathbb{N}^d$ is nonempty, $R_{\mathbf{z}}$ is a neighbor of U and furthermore $\dim \text{recc}(U) < \dim \text{recc}(R_{\mathbf{z}})$.*

Proof. We first show that $R_{\mathbf{z}}$ is a neighbor of U , i.e., that

$$\text{recc}(U) = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^d : S_U T \mathbf{x} \geq 0\} \subset \text{recc}(R_{\mathbf{z}}) = \{\mathbf{x} \in \mathbb{R}_{\geq 0}^d : S_{\mathbf{z}} T \mathbf{x} \geq 0\}.$$

Let $\mathbf{x} \in \text{recc}(U)$. Then for all $i \in L_U$, $s_i \cdot \mathbf{x} = 0$, so $s'_i \cdot \mathbf{x} \geq 0$, and otherwise $s'_i = s_i$, so $S_{\mathbf{z}} T \mathbf{x} \geq 0$. This implies $\mathbf{x} \in \text{recc}(R_{\mathbf{z}})$, i.e., $\text{recc}(U) \subset \text{recc}(R_{\mathbf{z}})$, proving the claim that $R_{\mathbf{z}}$ is a neighbor of U .

Next we argue that $\dim \text{recc}(U) < \dim \text{recc}(R_{\mathbf{z}})$. Now similar to the proof of Lemma 7.17, for each $i \notin L_U$ there exists some $\mathbf{y}_i \in \text{recc}(U)$ such that $\mathbf{t}_i \cdot \mathbf{y}_i \neq 0$, so $s_i(\mathbf{t}_i \cdot \mathbf{y}_i) > 0$. Then taking $\mathbf{y} = \sum_{i \notin L_U} \mathbf{y}_i$ we have $\mathbf{y} \in \text{recc}(U)$, and also $s_i(\mathbf{t}_i \cdot \mathbf{y}) > 0$ for all $i \notin L_U$.¹⁶

We will show for some $\epsilon > 0$, $\mathbf{y} + \epsilon \mathbf{z} \in \text{recc}(R_{\mathbf{z}})$, which will imply $\mathbf{z} \in \text{span}(\text{recc}(R_{\mathbf{z}}))$ so $\dim \text{recc}(R_{\mathbf{z}}) \geq \dim \text{recc}(U) + 1$ since $\mathbf{z} \in W^\perp$. Intuitively, to show this, we perturb the vector \mathbf{y} by a slight amount in the direction of \mathbf{z} to be on the correct side of all neighbor-separating hyperplanes, while remaining on the same side of all other hyperplanes. Formally, for all $i \in L_U$, we have $s'_i(\mathbf{t}_i \cdot \mathbf{z}) \geq 0$ by construction of s'_i . This might not hold for $i \notin L_U$, but in that case $s'_i(\mathbf{t}_i \cdot \mathbf{y}) > 0$. Thus we can pick some small enough $\epsilon > 0$ such that $s'_i(\mathbf{t}_i \cdot (\mathbf{y} + \epsilon \mathbf{z})) \geq 0$ for all $i \notin L_U$. For $i \in L_U$, we have $\mathbf{t}_i \cdot \mathbf{y} = 0$ (by definition of L_U since $\mathbf{y} \in \text{recc}(U)$), so we also have $s'_i(\mathbf{t}_i \cdot (\mathbf{y} + \epsilon \mathbf{z})) \geq 0$ and thus $\mathbf{y} + \epsilon \mathbf{z} \in \text{recc}(R_{\mathbf{z}})$. This concludes the claim that $\dim \text{recc}(U) < \dim \text{recc}(R_{\mathbf{z}})$.

Finally, we argue that $R_{\mathbf{z}} \cap \mathbb{N}^d$ is nonempty, so the region $R_{\mathbf{z}}$ is meaningfully defined. We can further assume that the vector $\mathbf{y}^+ = \mathbf{y} + \epsilon \mathbf{z} \in \mathbb{N}^d$ (again by density assuming that the pieces are rational and then scaling up to clear denominators). Now consider a point $\mathbf{u} \in U \cap \mathbb{N}^d$, so $S_U(T\mathbf{u} - \mathbf{h}) \geq 0$, and consider moving along \mathbf{y}^+ . For all i such that $s_i \neq s'_i$, we have $s'_i(\mathbf{t}_i \cdot \mathbf{y}^+) > 0$. Thus for all sufficiently large constants c , $S_{\mathbf{z}}(T(\mathbf{u} + c\mathbf{y}^+) - \mathbf{h}) \geq 0$ so $\mathbf{u} + c\mathbf{y}^+ \in R_{\mathbf{z}}$. Intuitively, the path from U along the vector \mathbf{y}^+ will eventually remain in the region $R_{\mathbf{z}}$. This shows that the neighbor of U in the direction of \mathbf{z} is well-defined. \square

For under-determined eventual region U , by repeatedly applying Lemma 7.18 using directions $\pm \mathbf{z}$ for any $\mathbf{z} \in W^\perp$, we can show that determined neighbors must actually exist:

Corollary 7.19. *An under-determined eventual region U has at least 2 determined neighbors.*

¹⁶ In fact \mathbf{y} can be shown to be in the *relative interior* of $\text{recc}(U)$, where the relative interior is the interior within the affine hull. [14]

We are now ready to consider the remaining case left after Lemma 7.16, when all determined neighbor gradients agree along some $\mathbf{z} \in W^\perp$.

Lemma 7.20. *Let I be a strip of an under-determined eventual region U . Let D_1, \dots, D_m be the determined neighbors of U , with extensions g_1, \dots, g_m . Assume there exists $\mathbf{z} \in W^\perp$ such that the gradients of the extensions along \mathbf{z} are all equal: $\nabla_{g_i} \cdot \mathbf{z} = \nabla_{g_j} \cdot \mathbf{z}$ for all i, j . Let $R_{\mathbf{z}}$ be the neighbor of U in the direction \mathbf{z} , with extension $g_{\mathbf{z}}$. Then $g_{\mathbf{z}}$ is also an extension from I : $g_{\mathbf{z}}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in I$, so taking $g_{\mathbf{z}}$ gives an extension from I which eventually dominates f .*

Proof. Let I be a strip of under-determined eventual region U , with extensions g_1, \dots, g_m from determined neighbors D_1, \dots, D_m respectively. Let $\mathbf{z} \in W^\perp$ such that $\nabla_{g_i} \cdot \mathbf{z} = \nabla_{g_j} \cdot \mathbf{z}$ for all i, j . We will consider the regions $R_{\mathbf{z}}$ and $R_{-\mathbf{z}}$ which are the neighbors of U in the directions \mathbf{z} and $-\mathbf{z}$. Recall by Lemma 7.18 that $\dim \text{recc}(U) < \dim \text{recc}(R_{\pm\mathbf{z}})$. The proof will proceed by induction on the codimension: $d - \dim \text{recc}(U)$. For U with codimension 1, $R_{\pm\mathbf{z}}$ must be determined regions with unique extensions. In general, $R_{\pm\mathbf{z}}$ could be under-determined, but will have lower codimension, so by the inductive hypothesis we assume $R_{\pm\mathbf{z}}$ have extensions which eventually dominate f (considering this Lemma alongside Lemma 7.16). Thus there exist quilt-affine extensions $g_{\mathbf{z}}$ and $g_{-\mathbf{z}}$ (from $R_{\mathbf{z}}$ and $R_{-\mathbf{z}}$) which eventually dominate f . Note these may have a larger period p^* as used in the proof of Lemma 7.16. We assume $g_{\mathbf{z}}$ and $g_{-\mathbf{z}}$ have common period p^* by taking the least common multiple if necessary. Thus we can write $g_{\mathbf{z}}(\mathbf{x}) = \nabla_{g_{\mathbf{z}}} \cdot \mathbf{x} + B_{g_{\mathbf{z}}}(\bar{\mathbf{x}} \bmod p^*)$ and $g_{-\mathbf{z}}(\mathbf{x}) = \nabla_{g_{-\mathbf{z}}} \cdot \mathbf{x} + B_{g_{-\mathbf{z}}}(\bar{\mathbf{x}} \bmod p^*)$.

Now by assumption $\nabla_{g_i} \cdot \mathbf{z} = \nabla_{g_j} \cdot \mathbf{z}$ for any determined neighbors D_i, D_j . Also by Lemma 7.12, $\text{proj}_W(\nabla_{g_i}) = \text{proj}_W(\nabla_{g_j})$. Thus all determined gradients agree along $\text{span}(W, \mathbf{z})$. The regions $R_{\pm\mathbf{z}}$ are either determined, or their determined neighbors are among D_1, \dots, D_m (by transitivity of the neighbor relation). Regardless, we can say $\text{proj}_{\text{span}(W, \mathbf{z})}(\nabla_{g_{\mathbf{z}}}) = \text{proj}_{\text{span}(W, \mathbf{z})}(\nabla_{g_{-\mathbf{z}}})$. For this proof, we will consider the affine space $A = \text{aff}(I) + \text{span}(\mathbf{z}) = \{\mathbf{i} + \mathbf{w} + c\mathbf{z} : \mathbf{i} \in I, \mathbf{w} \in W, c \in \mathbb{R}\}$ containing all points reachable from I by vectors in $\text{span}(W, \mathbf{z})$. We now claim that $g_{\mathbf{z}}(\mathbf{x}) = g_{-\mathbf{z}}(\mathbf{x})$ for all $\mathbf{x} \in A \cap \mathbb{N}^d$. The gradients along directions in A ($\text{span}(W, \mathbf{z})$) were already shown to be equal. Thus if $g_{\mathbf{z}}(\mathbf{x}) \neq g_{-\mathbf{z}}(\mathbf{x})$ (without loss of generality $g_{\mathbf{z}}(\mathbf{x}) < g_{-\mathbf{z}}(\mathbf{x})$), then $g_{\mathbf{z}}(\mathbf{y}) < g_{-\mathbf{z}}(\mathbf{y})$ for all congruent $\mathbf{y} \in \bar{\mathbf{x}} \bmod p^*$. However, $g_{-\mathbf{z}}$ is an extension of f from $R_{-\mathbf{z}}$, so we have $g_{\mathbf{z}}(\mathbf{y}) < f(\mathbf{y})$ for all $\mathbf{y} \in R_{-\mathbf{z}} \cap (\bar{\mathbf{x}} \bmod p^*)$. This contradicts the fact that $g_{\mathbf{z}}$ eventually dominates f , and completes the claim that $g_{\mathbf{z}}(\mathbf{x}) = g_{-\mathbf{z}}(\mathbf{x})$ on $A \cap \mathbb{N}^d$.

Thus we must have $g_{\mathbf{z}}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in A \cap R_{\mathbf{z}} \cap \mathbb{N}^d$ (since $g_{\mathbf{z}}$ is an extension from $R_{\mathbf{z}}$) and $\mathbf{x} \in A \cap R_{-\mathbf{z}} \cap \mathbb{N}^d$ (since $g_{\mathbf{z}} = g_{-\mathbf{z}}$, the extension from $R_{-\mathbf{z}}$). We now show also that $g_{\mathbf{z}}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in I$. Assume toward contradiction that $g_{\mathbf{z}}(\mathbf{u}) \neq f(\mathbf{u})$ for some $\mathbf{u} \in I$. By Lemma 7.3 we have affine partial function $f|_{U \cap (\bar{\mathbf{u}} \bmod p)}(\mathbf{x}) = \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{x} + b_{\bar{\mathbf{u}}}$ and by Lemma 7.12 we have $\text{proj}_W(\nabla_{\bar{\mathbf{u}}}) = \text{proj}_W(\nabla_{g_{\mathbf{z}}})$. Then for any $\mathbf{x} \in I \cap (\bar{\mathbf{u}} \bmod p^*)$, $\mathbf{x} = \mathbf{u} + \mathbf{w}$ for some $\mathbf{w} \in W$ by definition of I , so

$$\begin{aligned} g_{\mathbf{z}}(\mathbf{x}) - f(\mathbf{x}) &= \nabla_{g_{\mathbf{z}}} \cdot (\mathbf{u} + \mathbf{w}) + B_{g_{\mathbf{z}}}(\bar{\mathbf{u}} \bmod p^*) - \nabla_{\bar{\mathbf{u}}} \cdot (\mathbf{u} + \mathbf{w}) - b_{\bar{\mathbf{u}}} \\ &= \underbrace{(\nabla_{g_{\mathbf{z}}} \cdot \mathbf{w} - \nabla_{\bar{\mathbf{u}}} \cdot \mathbf{w})}_{=0 \text{ since } \mathbf{w} \in W} + (\nabla_{g_{\mathbf{z}}} \cdot \mathbf{u} + B_{g_{\mathbf{z}}}(\bar{\mathbf{u}} \bmod p^*)) - (\nabla_{\bar{\mathbf{u}}} \cdot \mathbf{u} + b_{\bar{\mathbf{u}}}) \\ &= g_{\mathbf{z}}(\mathbf{u}) - f(\mathbf{u}) \end{aligned} \tag{3}$$

In other words, if $g_{\mathbf{z}}(\mathbf{u}) \neq f(\mathbf{u})$, they are also unequal for all \mathbf{x} within the strip I on the entire congruence class $\bar{\mathbf{u}}$.

If we had $g_{\mathbf{z}}(\mathbf{u}) < f(\mathbf{u})$, then $g_{\mathbf{z}}(\mathbf{x}) < f(\mathbf{x})$ for all $\mathbf{x} \in I \cap (\bar{\mathbf{u}} \bmod p^*)$ by (3), which contradicts that $g_{\mathbf{z}}$ eventually dominates f .

The other case is that $g_{\mathbf{z}}(\mathbf{u}) > f(\mathbf{u})$, so again by (3), we have $g_{\mathbf{z}}(\mathbf{x}) > f(\mathbf{x})$ for all $\mathbf{x} \in I \cap (\bar{\mathbf{u}} \bmod p^*)$. (This is the behavior of our example (2), which will be shown to not be obviously-computable by the following general argument). Here, similar to the proof of Lemma 7.9, we will apply Lemma 4.1 by creating a contradiction sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots) \in \mathbb{N}^d$ such that for all $i < j$ there exists some $\Delta_{ij} \in \mathbb{N}^d$ with

$$f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) > f(\mathbf{a}_j + \Delta_{ij}) - f(\mathbf{a}_j).$$

To do this, we will find a sequence $(\mathbf{a}_1, \mathbf{a}_2, \dots) \in A \cap R_{\mathbf{z}} \cap (\bar{\mathbf{u}} \bmod p^*)$, so $g_{\mathbf{z}}(\mathbf{a}_i) = f(\mathbf{a}_i)$ for all i . We will then find another sequence $(\mathbf{v}_1, \mathbf{v}_2, \dots) \in \mathbb{N}^d$ such that $\mathbf{a}_i + \mathbf{v}_i \in I \cap (\bar{\mathbf{u}} \bmod p^*)$ for all i , implying $g_{\mathbf{z}}(\mathbf{a}_i + \mathbf{v}_i) > f(\mathbf{a}_i + \mathbf{v}_i)$. Also, we need that for all $i < j$, $\mathbf{a}_i + \mathbf{v}_j \in A \cap R_{-\mathbf{z}} \cap (\bar{\mathbf{u}} \bmod p^*)$, so $g_{\mathbf{z}}(\mathbf{a}_i + \mathbf{v}_j) = f(\mathbf{a}_i + \mathbf{v}_j)$. If these are both true, then choosing $\Delta_{ij} = \mathbf{v}_j$ for $i < j$ gives

$$\begin{aligned} & f(\mathbf{a}_i + \Delta_{ij}) - f(\mathbf{a}_i) \\ &= g_{\mathbf{z}}(\mathbf{a}_i + \mathbf{v}_j) - g_{\mathbf{z}}(\mathbf{a}_i) \\ &= g_{\mathbf{z}}(\mathbf{a}_j + \mathbf{v}_j) - g_{\mathbf{z}}(\mathbf{a}_j) \quad \text{since } g_{\mathbf{z}} \text{ is quilt-affine and } \mathbf{a}_i \equiv \mathbf{a}_j \pmod{p^*} \\ &> f(\mathbf{a}_j + \mathbf{v}_j) - f(\mathbf{a}_j) = f(\mathbf{a}_j + \Delta_{ij}) - f(\mathbf{a}_j). \end{aligned}$$

Lemma 4.1 then tells us that f is not obviously-computable, a contradiction. It remains to show that such sequences $\mathbf{a}_i \in R_{\mathbf{z}}$ and $\mathbf{v}_i \in \mathbb{N}^d$ can be found satisfying $\mathbf{a}_i + \mathbf{v}_i \in I \cap (\bar{\mathbf{u}} \bmod p^*)$ for all i , and for all $i < j$, $\mathbf{a}_i + \mathbf{v}_j \in A \cap R_{-\mathbf{z}} \cap (\bar{\mathbf{u}} \bmod p^*)$.

Now from the proof of Lemma 7.18, we take the same $\mathbf{y} \in \text{recc}(U)$ and perturbed $\mathbf{y}^+ = \mathbf{y} + \epsilon \mathbf{z} \in \text{recc}(R_{\mathbf{z}})$ that were defined in that proof. Recall we showed for $\mathbf{x} \in U$, for all large enough c , $\mathbf{x} + c\mathbf{y}^+ \in R_{\mathbf{z}}$. Likewise, we also have $\mathbf{y}^- = \mathbf{y} - \epsilon \mathbf{z} \in \text{recc}(R_{-\mathbf{z}})$ (taking ϵ small enough to work for both $R_{\mathbf{z}}$ and $R_{-\mathbf{z}}$) and we can assume (by density of rationals and scaling up denominators) that $\mathbf{y}^+, \mathbf{y}^- \in \mathbb{N}^d$.

Pick $c \in \mathbb{N}$ large enough that $\mathbf{u} + cp^*\mathbf{y}^+ \in R_{\mathbf{z}}$ and $\mathbf{u} + cp^*\mathbf{y}^- \in R_{-\mathbf{z}}$. Then for all $i \in \mathbb{N}_+$, let $\mathbf{a}_i = \mathbf{u} + icp^*\mathbf{y}^+$ and $\mathbf{v}_i = icp^*\mathbf{y}^-$. Since $\mathbf{y}^+, \mathbf{y}^- \in \text{span}(W, \mathbf{z})$, we have $\mathbf{a}_i \in A$ for all i , and the multiple of p^* ensures all points are in $(\bar{\mathbf{u}} \bmod p^*)$ as desired. Finally, we can check that

$$\mathbf{a}_i + \mathbf{v}_i = \mathbf{u} + icp^*\mathbf{y}^+ + icp^*\mathbf{y}^- = \mathbf{u} + 2icp^*\mathbf{y} \in I$$

since $\mathbf{y}^+ + \mathbf{y}^- = 2\mathbf{y} \in \text{recc}(U)$. Also, for $i < j$ we have

$$\mathbf{a}_i + \mathbf{v}_j = \mathbf{u} + icp^*\mathbf{y}^+ + jcp^*\mathbf{y}^- = \mathbf{u} + (j-i)cp^*\mathbf{y}^- + icp^*\mathbf{y}$$

Note that $j-i \geq 1$, $\mathbf{u} + (j-i)cp^*\mathbf{y}^- \in R_{-\mathbf{z}}$, and $icp^*\mathbf{y} \in \text{recc}(R_{-\mathbf{z}})$. Thus

$$\mathbf{a}_i + \mathbf{v}_j = \mathbf{u} + (j-i)cp^*\mathbf{y}^- + icp^*\mathbf{y} \in R_{-\mathbf{z}}$$

as required.

Thus Lemma 4.1 gives a contradiction that f cannot be obviously-computable. We reached this contradiction by assuming that $g_{\mathbf{z}}(\mathbf{u}) \neq f(\mathbf{u})$ for some $\mathbf{u} \in I$. Thus we conclude that $g_{\mathbf{z}}(\mathbf{x}) = f(\mathbf{x})$ for all $\mathbf{x} \in I$, so $g_{\mathbf{z}}$ is an extension from I that eventually dominates f . \square

For any strip I of an under-determined eventual region U , one of the cases from Lemmas 7.16 or 7.20 applies to show there exists an extension from I that eventually dominates f . There are only finitely many such strips (Lemma 7.15), so alongside the unique extensions from the determined regions (Lemmas 7.7 and 7.9), we have identified a finite collection g_1, \dots, g_m of quilt-affine functions to complete the proof of Theorem 7.1.

8 Comparison to continuous case

In [9], the authors classified the power of output-oblivious *continuous* CRNs to stably compute real-valued functions $f : \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}_{\geq 0}$. We can generalize to also consider such functions by introducing the following natural scaling:

Definition 8.1. For a function $f : \mathbb{N}^d \rightarrow \mathbb{N}$, the ∞ -scaling $\hat{f} : \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}_{\geq 0}$ is given by

$$\hat{f}(\mathbf{z}) = \lim_{c \rightarrow \infty} \frac{f(\lfloor c\mathbf{z} \rfloor)}{c}.$$

Note this limit may not exist for arbitrary $f : \mathbb{N}^d \rightarrow \mathbb{N}$, but it will exist for all obviously-computable f .

The next theorem shows that in this scaling limit, our output-oblivious function class exactly corresponds to the real-valued function class from [9] (see Fig. 4b).

Theorem 8.2. If $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is obviously-computable, then the ∞ -scaling $\hat{f} : \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}_{\geq 0}$ is obviously-computable by a continuous CRN. Furthermore, every function obviously-computable by a continuous CRN is the ∞ -scaling of some function obviously-computable by a discrete CRN.

Proof. To prove the first statement, let $f : \mathbb{N}^d \rightarrow \mathbb{N}$ be obviously-computable. We will show the ∞ -scaling $\hat{f} : \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}_{\geq 0}$ satisfies the main classification of [9]: that \hat{f} is superadditive, positive-continuous, and piecewise rational-linear.

First we prove that for any quilt-affine $g : \mathbb{N}^d \rightarrow \mathbb{Z}$, the ∞ -scaling \hat{g} is nonnegative and rational-linear. From Definition 5.1, we can express $g(\mathbf{x}) = \nabla_g \cdot \mathbf{x} + B(\bar{\mathbf{x}} \bmod p)$ for $\nabla_g \in \mathbb{Q}_{\geq 0}^d$ and $B : \mathbb{Z}^d/p\mathbb{Z}^d \rightarrow \mathbb{Q}$. Then for any $\mathbf{z} \in \mathbb{R}_{\geq 0}^d$,

$$\hat{g}(\mathbf{z}) = \lim_{c \rightarrow \infty} \frac{\nabla_g \cdot \lfloor c\mathbf{z} \rfloor + B(\overline{\lfloor c\mathbf{z} \rfloor} \bmod p)}{c} = \nabla_g \cdot \mathbf{z},$$

since B is bounded. Because $\nabla_g \in \mathbb{Q}_{\geq 0}^d$, \hat{g} is nonnegative and rational-linear.

Now by the eventually-min condition (ii) of Theorem 5.2, there exists quilt-affine $g_1, \dots, g_m : \mathbb{N}^d \rightarrow \mathbb{Z}$ and $\mathbf{n} \in \mathbb{N}^d$ such that $f(\mathbf{x}) = \min_k(g_k(\mathbf{x}))$ for all $\mathbf{x} \geq \mathbf{n}$. Then for any $\mathbf{z} \in \mathbb{R}_{> 0}^d$, $\lfloor c\mathbf{z} \rfloor \geq \mathbf{n}$ for large enough c , so

$$\hat{f}(\mathbf{z}) = \lim_{c \rightarrow \infty} \frac{f(\lfloor c\mathbf{z} \rfloor)}{c} = \lim_{c \rightarrow \infty} \frac{\min_k(g_k(\lfloor c\mathbf{z} \rfloor))}{c} = \min_k \left(\lim_{c \rightarrow \infty} \frac{g_k(\lfloor c\mathbf{z} \rfloor)}{c} \right) = \min_k(\hat{g}_k(\mathbf{z})), \quad (4)$$

where we pass the limit through the min function because min is continuous. Since $\hat{g}_k(\mathbf{z}) = \nabla_{g_k} \cdot \mathbf{z}$ are all rational-linear, on the domain $\mathbb{R}_{> 0}^d$, $\hat{f}(\mathbf{z}) = \min_k(\hat{g}_k(\mathbf{z}))$ is continuous and piecewise rational-linear.

We will now generalize this argument to show on the full domain $\mathbb{R}_{\geq 0}^d$, \hat{f} is piecewise rational-linear and *positive-continuous*: for each subset $S \subseteq \{1, \dots, d\}$, \hat{f} is continuous on domain $D_S = \{\mathbf{z} \in \mathbb{R}_{\geq 0}^d : \mathbf{z}(i) = 0 \iff i \in S\}$. Fix any such S . By repeatedly applying the recursive condition (iii) of Theorem 5.2, the fixed-input restriction $f_{[(\forall i \in S) \mathbf{x}(i) \rightarrow 0]}$ fixing input coordinates in S to 0, is obviously-computable. Then by the eventually-min condition (ii), there exists quilt-affine g_1^S, \dots, g_m^S and $\mathbf{n} \in \mathbb{N}^d$ such that $f_{[(\forall i \in S) \mathbf{x}(i) \rightarrow 0]}(\mathbf{x}) = \min_k(g_k^S(\mathbf{x}))$ for all $\mathbf{x} \geq \mathbf{n}$, but it sufficient for $\mathbf{x}(i) \geq \mathbf{n}(i)$ for all $i \notin S$. Now let $\mathbf{z} \in D_S$, so $\mathbf{z}(i) = 0 \iff i \in S$. Then for large enough c , $\lfloor c\mathbf{z} \rfloor(i) \geq \mathbf{n}(i)$ for all $i \notin S$, so

$$f(\lfloor c\mathbf{z} \rfloor) = f_{[(\forall i \in S) \mathbf{x}(i) \rightarrow 0]}(\lfloor c\mathbf{z} \rfloor) = \min_k(g_k^S(\lfloor c\mathbf{z} \rfloor)).$$

Now repeating equation (4), we have $\hat{f}(\mathbf{z}) = \min_k(\hat{g}_k^S(\mathbf{z}))$, so \hat{f} is continuous and piecewise rational-linear on D_S . This holds for all $S \subseteq \{1, \dots, d\}$, so \hat{f} is positive-continuous and piecewise rational linear.

It remains to show that \hat{f} must be superadditive: $\hat{f}(\mathbf{a}) + \hat{f}(\mathbf{b}) \leq \hat{f}(\mathbf{a} + \mathbf{b})$ for all $\mathbf{a}, \mathbf{b} \in \mathbb{R}_{\geq 0}^d$. Let $\mathbf{a}, \mathbf{b} \in \mathbb{R}_{\geq 0}^d$ with $\mathbf{a} + \mathbf{b} \in D_S$ for a domain D_S defined as above. Then

$$\hat{f}(\mathbf{a} + \mathbf{b}) = \min_k(\hat{g}_k^S(\mathbf{a} + \mathbf{b})) = \hat{g}_i^S(\mathbf{a} + \mathbf{b}) = \hat{g}_i^S(\mathbf{a}) + \hat{g}_i^S(\mathbf{b})$$

for some minimizing rational-linear \hat{g}_i^S . It remains to show $\hat{g}_i^S(\mathbf{a}) \geq \hat{f}(\mathbf{a})$ (and by symmetry $\hat{g}_i^S(\mathbf{b}) \geq \hat{f}(\mathbf{b})$). This is immediate if $\mathbf{a} \in D_S$ since $\hat{f}(\mathbf{a}) = \min_k(\hat{g}_k^S)$. Otherwise if $\mathbf{a} \notin D_S$, assume toward contradiction that $\hat{f}(\mathbf{a}) > \hat{g}_k^S(\mathbf{a})$. Then for some small enough $\epsilon > 0$, we also have $\hat{f}(\mathbf{a}) > \hat{g}_k^S(\mathbf{a} + \epsilon\mathbf{b})$. Observing that $\mathbf{a} + \epsilon\mathbf{b} \in D_S$, then $\hat{g}_k^S(\mathbf{a} + \epsilon\mathbf{b}) \geq \hat{f}(\mathbf{a} + \epsilon\mathbf{b})$. But then $\hat{f}(\mathbf{a}) > \hat{f}(\mathbf{a} + \epsilon\mathbf{b})$, a contradiction since \hat{f} must be nondecreasing as the ∞ -scaling of the nondecreasing function f .

Thus \hat{f} is semilinear, positive-continuous, and piecewise rational-linear as desired.

Next, to prove the second statement, let $\hat{f} : \mathbb{R}_{\geq 0}^d \rightarrow \mathbb{R}_{\geq 0}$ be any semilinear, positive-continuous, and piecewise rational-linear function. We will show that there exists some obviously-computable $f : \mathbb{N}^d \rightarrow \mathbb{N}$ such that its ∞ -scaling is \hat{f} .

On each domain $D_S = \{\mathbf{z} \in \mathbb{R}_{\geq 0}^d : \mathbf{z}(i) = 0 \iff i \in S\}$ for $S \subseteq \{1, \dots, d\}$, $\hat{f}|_{D_S}$ is superadditive, continuous, and piecewise rational-linear. By Lemma 8 in [9], $\hat{f}|_{D_S}$ can be written as the minimum of a finite number of rational linear functions $\hat{g}_k^S(\mathbf{z}) = \nabla_{g_k} \cdot \mathbf{z}$. For each \hat{g}_k^S , we will identify a quilt-affine g_k^S with gradient ∇_{g_k} . In particular, we can define $g_k^S : \mathbb{N}^d \rightarrow \mathbb{N}$ for all $\mathbf{x} \in \mathbb{N}^d$ by $g_k^S(\mathbf{x}) = \lfloor \nabla_{g_k} \cdot \mathbf{x} \rfloor$, which will be quilt-affine.

Now for all S and integer $\mathbf{x} \in D_S \cap \mathbb{N}^d$, define $f(\mathbf{x}) = \min_k(g_k^S(\mathbf{x}))$. From the above proof it follows that \hat{f} is the ∞ -scaling of f . It is also straightforward to verify that f is obviously-computable by satisfying Theorem 5.2. f is nondecreasing, satisfying condition (i), because \hat{f} was semilinear and thus nondecreasing. f satisfies eventually-min condition (ii) since for all $\mathbf{x} \geq (1, \dots, 1)$, $\mathbf{x} \in D_\emptyset = \mathbb{R}_{> 0}^d$, so $f(\mathbf{x}) = \min_k(g_k^\emptyset(\mathbf{x}))$. For all other $S \neq \emptyset$, the fixed-input restriction $f_{[(\forall i \in S) \mathbf{x}(i) \rightarrow 0]}(\mathbf{x}) = \min_k(g_k^S(\mathbf{x}))$. It follows that f satisfies recursive condition (iii), because any fixed-input restriction will be eventually-min of quilt-affine functions.

Thus any function \hat{f} obviously-computable by a continuous CRN is the ∞ -scaling limit of some f obviously-computable by a discrete CRN. \square

9 Leaderless one-dimensional case

In this section we show a characterization of 1D functions $f : \mathbb{N} \rightarrow \mathbb{N}$ that are obliviously-computable *without* a leader. The general case for leaderless oblivious computation in higher dimensions remains open.

Note that the following observation applies to any number of dimensions. We say $f : \mathbb{N}^d \rightarrow \mathbb{N}$ is *superadditive* if $f(\mathbf{x}) + f(\mathbf{y}) \leq f(\mathbf{x} + \mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{N}^d$.

Observation 9.1. *Every f obliviously-computable by a leaderless CRN is superadditive.*

Proof. Let \mathcal{C} be a leaderless CRN stably computing f . We prove the observation by contrapositive. Suppose f is not superadditive. Then there are $\mathbf{x}, \mathbf{z} \in \mathbb{N}^d$ such that $f(\mathbf{x}) + f(\mathbf{z}) > f(\mathbf{x} + \mathbf{z})$. Recall $\mathbf{I}_{\mathbf{w}}$ is the initial configuration of \mathcal{C} representing input \mathbf{w} . Let $\alpha_{\mathbf{x}}$ be a sequence of reactions applied to $\mathbf{I}_{\mathbf{x}}$ to produce $f(\mathbf{x})$ copies of Y , and let $\alpha_{\mathbf{z}}$ be a sequence of reactions applied to $\mathbf{I}_{\mathbf{z}}$ to produce $f(\mathbf{z})$ copies of Y .

Since \mathcal{C} is leaderless, $\mathbf{I}_{\mathbf{x}+\mathbf{z}} = \mathbf{I}_{\mathbf{x}} + \mathbf{I}_{\mathbf{z}}$. Thus we can apply $\alpha_{\mathbf{x}}$ to $\mathbf{I}_{\mathbf{x}+\mathbf{z}}$, followed by $\alpha_{\mathbf{z}}$, producing $f(\mathbf{x}) + f(\mathbf{z})$ copies of Y . Since this is greater than $f(\mathbf{x} + \mathbf{z})$, to stably compute f , \mathcal{C} must have a reaction consuming Y , so it is not output-oblivious. Since \mathcal{C} was arbitrary, f cannot be obliviously-computable. \square

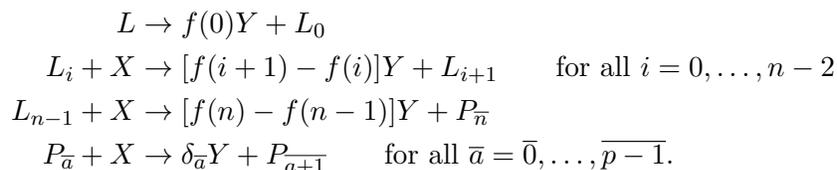
This added condition of superadditivity gives us the 1D leaderless characterization.

Theorem 9.2. *For any $f : \mathbb{N} \rightarrow \mathbb{N}$, f is obliviously-computable by a leaderless CRN $\iff f$ is semilinear and superadditive.*

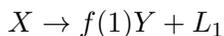
Proof. \implies : By Lemma 2.7 and Observation 9.1.

\impliedby : If f is superadditive, then f is also nondecreasing (since $f(x+1) \geq f(x) + f(1) \geq f(x)$). Then as in the Proof of Theorem 3.1, f is eventually quilt-affine, so there exist $n \in \mathbb{N}$, period $p \in \mathbb{N}_+$, and finite differences $\delta_{\bar{0}}, \dots, \delta_{\overline{p-1}} \in \mathbb{N}$, such that for all $x \geq n$, $f(x+1) - f(x) = \delta_{(\bar{x} \bmod p)}$. Also, without loss of generality assume p divides n , so $\bar{n} \bmod p = \bar{0}$.

The new CRN construction is motivated by trying to simply remove the leader species L from the construction used in Theorem 3.1. Recall that set of reactions was

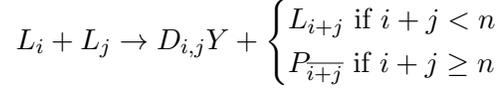


Since f is superadditive, we must have $f(0) = 0$. We then remove the species L and L_0 , and the two reaction that contain them, and add the first reaction



If this reaction only occurred once, this would still correctly compute f . Otherwise, however, there will be multiple ‘‘auxiliary leader species’’ from $\{L_1, \dots, L_{n-1}, P_{\bar{0}}, \dots, P_{\overline{p-1}}\}$ in the system. To correctly compute f , we must introduce pairwise reactions between these species that reduce the count of auxiliary leaders and add a corrective difference.

For all $i, j \in \{1, \dots, n-1\}$, the reaction between L_i and L_j is



where $D_{i,j} := f(i+j) - f(i) - f(j) \geq 0$ by superadditivity, and is the difference between how much output Y was released in the reactions that produced L_i and L_j and how much should have been produced from the input that led to L_i and L_j .

We have similar reactions between L_i and $P_{\bar{a}}$ for all $i \in \{1, \dots, n-1\}$ and $\bar{a} \in \{\bar{0}, \dots, \overline{p-1}\}$



where $D_{i,\bar{a}} := f(i+n+a) - f(i) - f(n+a) \geq 0$ by superadditivity. The reaction sequences that produced L_i consumed i copies of input X , and those that produced $P_{\bar{a}}$ consumed $n+a+kp$ for some $k \in \mathbb{N}$, so we have undercounted by $f(i+n+a+kp) - f(i) - f(n+a+kp) = D_{i,\bar{a}}$ since the periodic differences cancel.

Finally, the reactions between $P_{\bar{a}}$ and $P_{\bar{b}}$ for all $\bar{a}, \bar{b} \in \{\bar{0}, \dots, \overline{p-1}\}$ are



where $D_{\bar{a},\bar{b}} := f(n+a+n+b) - f(n+a) - f(n+b) \geq 0$ by superadditivity, and this gives the corrective difference in output by a similar argument.

Note that the rest of the reactions used in Theorem 3.1 are not strictly necessary, since if all input X undergoes the first reaction $X \rightarrow f(1)Y + L_1$, the corrective difference reactions will then reduce the count of auxiliary leader species down to 1, while outputting the correct differences to produce precisely $f(x)$ output. \square

10 Conclusion

An obvious question is the computational power of output-oblivious CRNs without an initial leader. A leaderlessly-obliviously-computable function must be superadditive, which is a strictly stronger condition than being nondecreasing. The continuous result [9] had the same restriction of superadditivity, so our “scaling limit” reduction to their function class (Theorem 8.2) shows our main function class is already “almost superadditive.” We also showed in the 1D case, $f : \mathbb{N} \rightarrow \mathbb{N}$ is leaderlessly-obliviously-computable if and only if f is semilinear and superadditive (Theorem 9.2).

Does adding the additional constraint of superadditivity to our full result (Theorem 5.2) classify leaderlessly-obliviously-computable $f : \mathbb{N}^d \rightarrow \mathbb{N}$? If this were true, a proof would require modifying our construction (Section 6) to eliminate the leader L . We successfully modified the 1D construction (Theorem 3.1) to remove the leader in proving Theorem 9.2, but it has been difficult to extend the same ideas to our much more complicated general construction.

An initial leader can also help make computation faster [5, 7, 21]. Many recent results in population protocols have shown time upper and lower bounds for computational tasks such as leader election and function/predicate computation [1, 2, 7, 18, 19, 21]. These techniques, however, are not at all designed to handle the constraint of output-obliviousness. It would be interesting to study how this constraint affects the time required for computation.

Acknowledgements. We thank Anne Condon, Cameron Chalk, Niels Kornerup, Wyatt Reeves, and David Soloveichik for discussing their related work with us and contributing early ideas.

References

- [1] Dan Alistarh, James Aspnes, David Eisenstat, Rati Gelashvili, and Ronald L. Rivest. Time-space trade-offs in molecular computation. In *SODA 2017: Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2017.
- [2] Dan Alistarh, James Aspnes, and Rati Gelashvili. Space-optimal majority in population protocols. In *SODA 2018: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2221–2239. SIAM, 2018.
- [3] Dana Angluin, James Aspnes, Zoë Diamadi, Michael Fischer, and René Peralta. Computation in networks of passively mobile finite-state sensors. *Distributed Computing*, 18:235–253, 2006. Preliminary version appeared in PODC 2004.
- [4] Dana Angluin, James Aspnes, and David Eisenstat. Stably computable predicates are semi-linear. In *PODC 2006: Proceedings of the twenty-fifth annual ACM symposium on Principles of distributed computing*, pages 292–299, New York, NY, USA, 2006. ACM Press.
- [5] Dana Angluin, James Aspnes, and David Eisenstat. Fast computation by population protocols with a leader. *Distributed Computing*, 21(3):183–199, September 2008. Preliminary version appeared in DISC 2006.
- [6] Dana Angluin, James Aspnes, David Eisenstat, and Eric Ruppert. The computational power of population protocols. *Distributed Computing*, 20(4):279–304, 2007.
- [7] Amanda Belleville, David Doty, and David Soloveichik. Hardness of computing and approximating predicates and functions with leaderless population protocols. In *ICALP*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 141:1–141:14, 2017.
- [8] Luca Cardelli. Strand algebras for DNA computing. *Natural Computing*, 10(1):407–428, 2011.
- [9] Cameron Chalk, Niels Kornerup, Wyatt Reeves, and David Soloveichik. Composable rate-independent computation in continuous chemical reaction networks. In *Computational Methods in Systems Biology*, 2018.
- [10] Ho-Lin Chen, David Doty, and David Soloveichik. Deterministic function computation with chemical reaction networks. *Natural Computing*, 13(4):517–534, 2014. Preliminary version appeared in DNA 2012.
- [11] Yuan-Jyue Chen, Neil Dalchau, Niranjana Srinivas, Andrew Phillips, Luca Cardelli, David Soloveichik, and Georg Seelig. Programmable chemical controllers made from DNA. *Nature Nanotechnology*, 8(10):755–762, 2013.
- [12] Ward Cheney and David Kincaid. *Linear Algebra: Theory and Applications*. Jones and Bartlett Publishers, 2010.

- [13] Ben Chugg, Anne Condon, and Hooman Hashemi. Output-oblivious stochastic chemical reaction networks. In *OPODIS 2018: Proceedings of the 22nd International Conference on Principles of Distributed Systems*, 2018.
- [14] Jesús A De Loera, Raymond Hemmecke, and Matthias Köppe. *Algebraic and geometric ideas in the theory of discrete optimization*, volume 14. SIAM, 2013.
- [15] Leonard E. Dickson. Finiteness of the odd perfect and primitive abundant numbers with n distinct prime factors. *American Journal of Mathematics*, 35(4):413–422, 1913.
- [16] David Doty and Mahsa Eftekhari. Efficient size estimation and impossibility of termination in uniform dense population protocols. In *PODC 2019: Proceedings of the 38th ACM Symposium on Principles of Distributed Computing*, 2019. to appear.
- [17] David Doty and Monir Hajiaghayi. Leaderless deterministic chemical reaction networks. *Natural Computing*, 14(2):213–223, 2015. Preliminary version appeared in DNA 2013.
- [18] David Doty and David Soloveichik. Stable leader election in population protocols requires linear time. *Distributed Computing*, 31(4):257–271, 2018. Special issue of invited papers from DISC 2015.
- [19] Leszek Gasieniec and Grzegorz Stachowiak. Fast space optimal leader election in population protocols. In *SODA 2018: Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2653–2667, 2018.
- [20] Daniel T. Gillespie. Exact stochastic simulation of coupled chemical reactions. *Journal of Physical Chemistry*, 81(25):2340–2361, 1977.
- [21] Adrian Kosowski and Przemyslaw Uznanski. Brief announcement: Population protocols are fast. In *PODC 2018: Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, pages 475–477. ACM, 2018.
- [22] Tyson Rockafeller. *Convex Analysis*, chapter 8. Princeton University Press, 1970.
- [23] David Soloveichik, Georg Seelig, and Erik Winfree. DNA as a universal substrate for chemical kinetics. *Proceedings of the National Academy of Sciences*, 107(12):5393, 2010. Preliminary version appeared in DNA 2008.
- [24] Niranjana Srinivas, James Parkin, Georg Seelig, Erik Winfree, and David Soloveichik. Enzyme-free nucleic acid dynamical systems. *Science*, 358(6369):eaal2052, 2017.
- [25] Richard Stanley. An introduction to hyperplane arrangements, 2006. URL: <http://www.cis.upenn.edu/~cis610/sp06stanley.pdf>.