



Paul Balister · Béla Bollobás · Robert Morris · Paul Smith

Universality for monotone cellular automata

Received November 6, 2022; revised September 10, 2024

Abstract. In this paper we study monotone cellular automata in d dimensions. We develop a general method for bounding the growth of the infected set when the initial configuration is chosen randomly, and then use this method to prove a lower bound on the critical probability for percolation that is sharp up to a constant factor in the exponent for every ‘critical’ model. This is one of three papers that together confirm the Universality Conjecture of Bollobás, Duminil-Copin, Morris and Smith.

Keywords: bootstrap percolation, cellular automata, universality, threshold.

Contents

1. Introduction	4725
2. The resistance of an update family	4732
3. The main theorem, and an outline of the proof	4736
4. Bounding trees	4743
5. Icebergs	4753
6. The size of an iceberg	4760
7. One-step hierarchies	4774
8. The proof of Theorem 3.1	4781
A. Simple calculations	4790
References	4793

1. Introduction

The Universality Conjecture of Bollobás, Duminil-Copin, Morris and Smith [17] states that every d -dimensional monotone cellular automaton (see Definition 1.1) is a member of

Paul Balister: Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK;
paul.balister@maths.ox.ac.uk

Béla Bollobás: Department of Pure Mathematics and Mathematical Statistics,
 University of Cambridge, Cambridge CB3 0WA, UK; Department of Mathematical Sciences,
 University of Memphis, Memphis, TN 38152, USA; bb12@cam.ac.uk

Robert Morris: IMPA, 22460-320 Rio de Janeiro, Brazil; rob@impa.br

Paul Smith: London, UK; paulsmith@cantab.net

Mathematics Subject Classification 2020: 60C05 (primary); 60K35 (secondary).

one of $d + 1$ ‘universality classes’, which are characterized by their behaviour on sparse random sets. More precisely, it states that the critical probability for percolation in $\mathbb{Z}_n^d$ of each such automaton is either bounded away from zero, or is of order $(\log_{(r-1)} n)^{-\Theta(1)}$ for some $r \in \{1, \dots, d\}$, where $\log_{(r)}$ denotes an r -times iterated logarithm. The automata with non-trivial phase transitions were characterized in [6, 7], and for those with trivial phase transitions an upper bound on the critical probability was proved in terms of an explicit parameter $r = r(\mathcal{U})$. In this paper we prove a matching lower bound, and hence complete the proof of the Universality Conjecture.

The mathematical study of general cellular automata was initiated over 70 years ago,¹ by von Neumann (see [70, 74, 77, 78]) and Ulam [73], who were inspired by the then-incipient theory of computing, and by the model of neural networks that had been introduced a few years earlier by McCulloch and Pitts [62]. They can be used to simulate the behaviour of a wide variety of biological, chemical and physical processes, and have become important objects of study in many areas of applied mathematics. For example, some of the most extensively studied models in statistical physics are (probabilistic) cellular automata, including the (kinetic) Ising model of ferromagnetism [32, 40, 59] and kinetically constrained models of the liquid-glass transition [21, 41, 72] (see Sections 1.1 and 1.2). Perhaps surprisingly, an important first step in understanding the metastable behaviour of these models is often to study the corresponding monotone model, in which sites can only change state in one direction (see, for example, [24, 38, 46, 47, 60, 63]).²

In this paper we consider general monotone cellular automata in d dimensions. The study of such processes originated in the work of Chalupa, Leath and Reich [25] in 1979, who introduced the family of models now known as ‘ r -neighbour bootstrap percolation’ (see below), and many other special cases have since been studied (see, e.g., [35, 36, 42, 55, 65, 76]). The following definition, however, is due to Bollobás, Smith and Uzzell [18], who were the first to study completely general models in two dimensions.

Definition 1.1. Let $\mathcal{U} = \{X_1, \dots, X_m\}$ be a finite collection of finite, non-empty subsets of $\mathbb{Z}^d \setminus \{\mathbf{0}\}$. Let $A \subset \mathbb{Z}_n^d$ be a set of initially *infected* sites, set $A_0 = A$, and define

$$A_{t+1} = A_t \cup \{x \in \mathbb{Z}_n^d : x + X \subset A_t \text{ for some } X \in \mathcal{U}\}$$

for each $t \geq 0$. The $\mathcal{U}$ -closure of A is the set $[A]_{\mathcal{U}} := \bigcup_{t \geq 0} A_t$ of eventually infected sites.

We call this model $\mathcal{U}$ -bootstrap percolation, and $\mathcal{U}$ the *update family* of the process. Informally, a vertex x (that is not in A) becomes infected as soon as some set $X \in \mathcal{U}$ translated by x is entirely infected, and infected sites remain infected forever.

¹However, according to von Neumann [77], it was already the case in 1951 that ‘Automata have been playing a continuously increasing, and have by now attained a very considerable, role in the natural sciences. This is a process that has been going on for several decades.’

²The reason that this might be surprising is that both the Ising model and kinetically constrained models exhibit vastly more complex behaviour than their monotone analogues. We refer the reader to the books [19, 66] for a general introduction to metastability in statistical physics.

We are interested in understanding the behaviour of the $\mathcal{U}$ -bootstrap process when the initial set A is chosen randomly. To be precise, let us say that A is a p -random subset of $\mathbb{Z}_n^d$ if each vertex of $\mathbb{Z}_n^d$ is included in A independently with probability p , and write $\mathbb{P}_p$ for the associated product probability measure. Our main aim is to determine the behaviour (as $n \rightarrow \infty$) of the *critical probability*

$$p_c(\mathbb{Z}_n^d, \mathcal{U}) := \inf \{p \in (0, 1] : \mathbb{P}_p([A]u = \mathbb{Z}_n^d) \geq 1/2\} \quad (1)$$

for each d -dimensional update family $\mathcal{U}$.

The most extensively-studied of these models is the classical r -neighbour bootstrap process, in which vertices are infected at time t if they have at least r already-infected neighbours. The update family $\mathcal{N}_r^d$ of this model consists of the $\binom{2d}{r}$ subsets of size r of the $2d$ nearest neighbours of the origin. Over the past 40 years a great deal of work has been done on this model, and its behaviour is now well-understood, both in two dimensions [4, 50, 52, 53] and in higher dimensions [10, 11, 22, 23, 68]. In particular, it was proved by Holroyd [53] in the case $d = r = 2$, and by Balogh, Bollobás, Duminil-Copin and Morris [10, 11] in general, that

$$p_c(\mathbb{Z}_n^d, \mathcal{N}_r^d) = \left(\frac{\lambda(d, r) + o(1)}{\log_{(r-1)} n} \right)^{d-r+1} \quad (2)$$

as $n \rightarrow \infty$ for each $2 \leq r \leq d$, where $\lambda(d, r) > 0$ is an explicit constant. If $r > d$, on the other hand, then it is not hard to see that $p_c(\mathbb{Z}_n^d, \mathcal{N}_r^d) \rightarrow 1$ as $n \rightarrow \infty$, since an uninfected copy of $\{0, 1\}^d$ cannot be infected in the process. Similar results have been proved for some other families of automata in two dimensions [16, 35–37], and weaker bounds are known for a small number of specific three-dimensional update families [13, 76].

For general $\mathcal{U}$ -bootstrap processes the challenges are rather different, and the first bounds for arbitrary update families³ were obtained by Bollobás, Smith and Uzzell [18], who proposed a partition of the family of all two-dimensional models into three classes, which they called ‘subcritical’, ‘critical’ and ‘supercritical’. They proved that $p_c(\mathbb{Z}_n^2, \mathcal{U})$ is polynomial in n for supercritical families, and poly-logarithmic in n for critical families. Balister, Bollobás, Przykucki and Smith [9] then proved that the critical probability for subcritical two-dimensional families is bounded away from zero, confirming a conjecture from [18]. For critical two-dimensional families, much more precise bounds were obtained by Bollobás, Duminil-Copin, Morris and Smith [17], who determined the critical probability up to a constant factor.

One of the key insights of Bollobás, Smith and Uzzell [18] was that, at least in two dimensions, the typical behaviour of the $\mathcal{U}$ -bootstrap process should, roughly speaking, be determined by its action on discrete half-spaces; specifically, whether or not discrete half-spaces are ‘stable’ under the action of $\mathcal{U}$. To be precise, for each $u \in \mathbb{S}^{d-1}$, let us write $\mathbb{H}_u^d := \{x \in \mathbb{Z}^d : \langle x, u \rangle < 0\}$ for the discrete half-space with normal u , where

³Slightly earlier, Duminil-Copin and Holroyd [35] had obtained very precise results for a large family of ‘balanced’ and ‘symmetric’ two-dimensional monotone cellular automata.

$\mathbf{S}^{d-1} = \{x \in \mathbb{R}^d : \|x\| = 1\}$ is the unit sphere in $\mathbb{R}^d$. Now, given a d -dimensional update family $\mathcal{U}$, we define the set of *stable directions* of $\mathcal{U}$ to be

$$\mathcal{S}(\mathcal{U}) := \{u \in \mathbf{S}^{d-1} : [\mathbb{H}_u^d]_{\mathcal{U}} = \mathbb{H}_u^d\}. \quad (3)$$

Note that $u \in \mathcal{S}(\mathcal{U})$ if and only if $X \not\subset \mathbb{H}_u^d$ for every $X \in \mathcal{U}$.

The following classification of d -dimensional update families was proposed by Bollobás, Smith and Uzzell [18] in the case $d = 2$, by Balister, Bollobás, Przykucki and Smith [9] for subcritical families, and by Bollobás, Duminil-Copin, Morris and Smith [17] in general. Given a sphere $\mathbf{S} \subset \mathbb{R}^d$ of arbitrary dimension and a set $\mathcal{T} \subset \mathbb{R}^d$, we write $\text{int}_{\mathbf{S}}(\mathcal{T})$ for the interior of $\mathcal{T}$ in $\mathbf{S}$ with respect to the topology induced by geodesic distance.

Definition 1.2. A d -dimensional update family $\mathcal{U}$ with stable set $\mathcal{S} = \mathcal{S}(\mathcal{U})$ is

- *supercritical* if $H \cap \mathcal{S} = \emptyset$ for some open hemisphere $H \subset \mathbf{S}^{d-1}$;
- *critical* if there exists a hemisphere $H \subset \mathbf{S}^{d-1}$ such that $\text{int}_{\mathbf{S}^{d-1}}(H \cap \mathcal{S}) = \emptyset$ and if $H \cap \mathcal{S} \neq \emptyset$ for every open hemisphere $H \subset \mathbf{S}^{d-1}$;
- *subcritical* if $\text{int}_{\mathbf{S}^{d-1}}(H \cap \mathcal{S}) \neq \emptyset$ for every hemisphere $H \subset \mathbf{S}^{d-1}$.

To motivate this definition, let us mention that if $\mathcal{U}$ is supercritical, then there is a direction (the centre of the empty open hemisphere H) in which a large ‘droplet’ (see Definition 3.2) can grow without needing the ‘help’ of any additional infected sites, and if $\mathcal{U}$ is subcritical then there is a (bounded, convex) polytope whose sides are perpendicular to interior points of $\mathcal{S}$. Note that every $\mathcal{U}$ is either supercritical, critical or subcritical.

It was conjectured by Balister, Bollobás, Przykucki and Smith [9], and proved by the present authors in [6, 7], that $p_c(\mathbb{Z}^d, \mathcal{U}) > 0$ if and only if $\mathcal{U}$ is subcritical.⁴ Moreover, in [7] we proved upper bounds on $p_c(\mathbb{Z}_n^d, \mathcal{U})$ for all critical and supercritical models. In this paper we prove matching lower bounds for all critical models in d dimensions, and hence complete the proof of the following ‘universality’ theorem, which confirms a conjecture of Bollobás, Duminil-Copin, Morris and Smith [17], and implies that *every* $\mathcal{U}$ -bootstrap process ‘resembles’ (in its large-scale behaviour) one of the r -neighbour models.

Theorem 1.3. *Let $\mathcal{U}$ be a d -dimensional update family.*

- If $\mathcal{U}$ is supercritical then $p_c(\mathbb{Z}_n^d, \mathcal{U}) = n^{-\Theta(1)}$.*
- If $\mathcal{U}$ is critical then there exists $r \in \{2, \dots, d\}$ such that*

$$p_c(\mathbb{Z}_n^d, \mathcal{U}) = \left(\frac{1}{\log_{(r-1)} n} \right)^{\Theta(1)}.$$

- If $\mathcal{U}$ is subcritical then $p_c(\mathbb{Z}^d, \mathcal{U}) > 0$.*

⁴ $p_c(\mathbb{Z}^d, \mathcal{U})$ is defined as in (1), except the $\mathcal{U}$ -bootstrap process is run on the infinite graph $\mathbb{Z}^d$ instead of the finite torus $\mathbb{Z}_n^d$. An alternative proof of the fact that $p_c(\mathbb{Z}^d, \mathcal{U}) > 0$ for all subcritical update families $\mathcal{U}$ was found recently by Hartarsky and Szabó [51].

We will define the quantity $r = r(\mathcal{U})$ explicitly in Section 2, and use it to state a more precise version of the theorem above (see Theorem 2.6). We will then be able to state the main result of this paper (Theorem 3.1), which implies the lower bound in Theorem 1.3 for critical update families. It moreover does so with a universal exponent, depending only on d , which is sharp for the 2-neighbour bootstrap process. The upper bounds in parts (a) and (b) of Theorem 1.3 were proved in [7, 18], and part (c) was proved in [6, 9].

As mentioned above, in the case $d = 2$ a more precise version of Theorem 1.3 was proved in [17, 18], and it is natural to ask whether one could obtain a similar result in higher dimensions. Unfortunately, however, it is proved in [8] that for critical models with $2 \leq r < d$ and for supercritical models with $d \geq 2$, the exponent in the critical probability (if it exists) is in general uncomputable.⁵ The exponent in the case $d = r = 2$ is computable (it was determined in [17], though computing it is NP-hard, see [49]), and we expect the same to be true for all $d = r \geq 2$. Determining the correct exponent for all d -dimensional update families with $r(\mathcal{U}) = d$ seems likely to be very difficult, however, and we consider this to be the central open problem in the area.

Problem 1.4. Determine $\log p_c(\mathbb{Z}_n^d, \mathcal{U})$ up to a factor of $1 + o(1)$ as $n \rightarrow \infty$ for every d -dimensional critical update family $\mathcal{U}$ with $r(\mathcal{U}) = d$.

All previous lower bounds for critical models with $r \geq 3$ (see [10, 11, 15, 22, 23, 54, 76]) have used the method introduced by Cerf and Cirillo [22] in their groundbreaking work on the 3-neighbour bootstrap process. In their proof, they coupled this model with a sequence of adjacent two-dimensional models, using the fact that the most ‘difficult’ directions for growth come in pairs of the form $\{u, -u\}$. One of the main problems that we face is that our models do not (in general) have any form of symmetry, and we have therefore had to develop a new and completely different method for controlling the growth. This new method involves a delicate induction on the ‘size’ of an ‘iceberg’ (see Sections 3 and 7–8), and was inspired by Holroyd’s method of hierarchies [53].

In Sections 2 and 3 we shall state our main theorem precisely, and give a detailed sketch of its proof. The proof itself is given in Sections 4–8. Before embarking on this journey, however, let us provide some further motivation for Theorem 1.3 by briefly discussing some potential applications of our method to two (non-monotone) spin models that have been extensively studied in the statistical physics literature: the Ising model of ferromagnetism, and kinetically constrained models of the liquid-glass transition.

1.1. Kinetically constrained models

A glass is a disordered material that nevertheless behaves mechanically like a solid, and is formed by rapidly cooling a viscous liquid. Understanding this liquid-glass transition is an

⁵More precisely, it is shown in [8] that for each Turing machine program P and each $1 \leq r < d$, there exists a d -dimensional update family $\mathcal{U}$ with $r(\mathcal{U}) = r$ such that if P halts then the exponent of $\mathcal{U}$ is at most $2/3$ and if P does not halt then the exponent is at least 1. If the exponents were always computable, then this would violate the undecidability of the halting problem.

important open problem in condensed matter physics (see for example [5, 29]). Kinetically constrained models were introduced in the 1980s (see [39], or the reviews [41, 67]) in order to model the liquid-glass transition, and exhibit several key properties of cooled liquids near the glass transition point.

Perhaps the simplest way to understand a kinetically constrained model is as a biased random walk on the family of percolating sets⁶ in $\mathcal{U}$ -bootstrap percolation on $\mathbb{Z}^d$. More specifically, the state (either ‘empty’ or ‘occupied’) of each site $x \in \mathbb{Z}^d$ updates at rate 1 as long as the set $x + X$ is entirely empty for some $X \in \mathcal{U}$; when it updates, it becomes empty with probability p , and occupied with probability $1 - p$, independently of its current state (and all other events). Observe that if the initial set A_0 of empty sites is chosen to be p -random, then this process $(A_t)_{t \geq 0}$ is stationary. Well-studied examples of kinetically constrained models include the East model (see, e.g., [28]), whose update family consists of the single set $\{-1\}$, and the r -facilitated Friedrickson–Andersen model, introduced in [39], which corresponds to r -neighbour bootstrap percolation.

Over the past few years, there have been some dramatic advances in our understanding of general kinetically constrained models in two dimensions, mirroring those in the study of $\mathcal{U}$ -bootstrap percolation. One of the main quantities of interest for a kinetically constrained model (starting from equilibrium) is the infection time of the origin⁷

$$\tau(\mathbb{Z}^d, \mathcal{U}) := \inf \{t \geq 0 : \mathbf{0} \in A_t\},$$

which, by the fundamental work of Cancrini, Martinelli, Roberto and Toninelli [20], is almost surely finite whenever $p > p_c(\mathbb{Z}^d, \mathcal{U})$. Moreover, a lower bound on the mean infection time is given (up to a constant factor) by the median infection time of the origin in the corresponding $\mathcal{U}$ -bootstrap process (see [61, Lemma 4.3]). This bound is sometimes sharp (for example, for the 2-neighbour model, see [48]), but for many other update families (for example, for the Duarte model, see [58]) the more complex behaviour of kinetically constrained models (in particular, the existence of ‘energy barriers’) causes $\tau(\mathbb{Z}^d, \mathcal{U})$ to be significantly larger than the bound given by the $\mathcal{U}$ -bootstrap process.

Motivated by the connection with $\mathcal{U}$ -bootstrap percolation, a number of conjectures were made in [64] regarding the rate of growth of the mean infection time as $p \rightarrow 0$ for critical update families in two dimensions, and for supercritical update families in d dimensions. These have now all been either proved or disproved (see [46, 47, 58, 60]), and in two dimensions the situation is extremely well-understood, with the full universality picture determined (see [44, 45, 48] for the most recent developments). More precisely, both the critical and supercritical families need to be partitioned into two different universality classes [46, 47], and then further refined for logarithmic corrections [44, 45].

⁶More precisely, this is true if the initial set of empty sites percolates for the $\mathcal{U}$ -bootstrap process. If the empty sites at time zero are chosen to be p -random, then this holds almost surely for every $p > 0$ if $\mathcal{U}$ is not subcritical, by Theorem 1.3 (specifically, by the upper bounds proved in [7]).

⁷The mean of this quantity also provides a lower bound on the relaxation time; see [60, Section 2.2].

In higher dimensions, corresponding results are known only for the r -neighbour model, the main obstruction to studying more general models being the lack of tools for controlling the corresponding $\mathcal{U}$ -bootstrap models. Having developed such tools in this paper and in [7], it now seems feasible that substantial progress could be made on this problem. For example, the following lower bound is a straightforward consequence of the results of this paper. The statement depends on the parameter $r(\mathcal{U})$, which will be defined explicitly in Section 2 (see Definition 2.4). We write $\exp_{(r)}$ for an r -times iterated exponential, so $\exp_{(0)} n = n$ and $\exp_{(r)} n = \exp(\exp_{(r-1)} n)$ for each $r \geq 1$.

Theorem 1.5. *Let $\mathcal{U}$ be a critical d -dimensional update family. Then*

$$\mathbb{E}[\tau(\mathbb{Z}^d, \mathcal{U})] \geq \exp_{(r-1)}(p^{-c})$$

for some constant $c = c(d) > 0$, where $r = r(\mathcal{U}) \in \{2, \dots, d\}$.

The deduction of Theorem 1.5 from our results is standard (see, e.g., [61]), and we therefore leave the details to the interested reader. It seems reasonable to conjecture that the bound in Theorem 1.5 is sharp up to the value of the constant c for all critical families (even though this is not true for supercritical families, see [58]), and it may perhaps even be possible to prove a matching upper bound using the techniques of [20, 60, 61], together with those introduced in [7]. However, there are likely to be very significant technical challenges involved in making such an approach rigorous.

1.2. The Ising model of ferromagnetism

Introduced over 100 years ago, in an attempt to model Curie's temperature,⁸ the Ising model has become one of the most extensively-studied models in statistical physics. The equilibrium (Gibbs) measures of the model are now relatively well-understood, and a wide array of powerful tools have been developed for their study; we refer the reader to the excellent surveys by Duminil-Copin [31, 32], and the recent papers [1–3, 26, 27, 33, 71], for more details.

The kinetic (or stochastic) Ising model, on the other hand, which was one of the original motivations for the introduction of the r -neighbour bootstrap process in [25], and also the early rigorous work in [4], is rather less well-understood, despite the recent breakthroughs [56, 57] on the mixing time in the high temperature regime. For example, consider the case of the zero-temperature kinetic Ising model on $\mathbb{Z}^d$ with zero external field. In this setting, the state (+ or −) of each site of $\mathbb{Z}^d$ updates at random times to agree with the majority of its (nearest) neighbours, with ties broken randomly. It is a longstanding (possibly folklore) conjecture that, for all $d \geq 2$, if the initial states are chosen independently with probability $p > 1/2$ of being +, then the system almost surely fixates at +. However, this has only been proved when p is sufficiently close to 1, by Fontes, Schonmann and Sidoravicius [38], and when $p \rightarrow 1/2$ (slowly) as $d \rightarrow \infty$,

⁸This is the temperature above which a magnetic substance loses its ferromagnetic properties.

by Morris [63]. We remark that results on the r -neighbour bootstrap process (from [4] and [12], respectively) play a crucial role in the proofs in [38] and [63].

The r -neighbour bootstrap process has also found applications in the study of the metastable behaviour of the stochastic Ising model at small positive temperatures and with non-zero external field. For example, the relaxation of the system, starting with all states equal, was studied by Schonmann [69] (who confirmed a conjecture of Aizenman and Lebowitz [4]) when the temperature is small but fixed and the external field tends to zero, and by Dehghanpour and Schonmann [30] and Cerf and Manzo [24] when the external field is fixed and the temperature tends to zero.

We expect that the techniques introduced in this paper will similarly find applications in the study of the stochastic Ising model with more general (finite-range) interactions. For example, it was conjectured in [64] that for any critical d -dimensional update family $\mathcal{U}$, fixation occurs almost surely in the $\mathcal{U}$ -Ising dynamics⁹ if p is sufficiently close to 1. There are two main obstructions to using the method of [38] to prove the conjecture: the first, covering the sites in state ‘−’ with well-separated $\mathcal{U}$ -closed droplets, is dealt with by the method of this paper; the second, showing that a (−)-droplet disappears in polynomial time, seems to require additional new ideas (see [14] for more details). Similarly, it seems plausible that the results of [24] could be extended to models with finite-range interactions, and we would expect the techniques introduced in this paper to play an important role in the proof of any such generalization.

2. The resistance of an update family

In this section we shall define the parameter $r = r(\mathcal{U})$ of a d -dimensional update family (see Definition 2.4), which we call the *resistance* of $\mathcal{U}$. This will then allow us to state a refined version of Theorem 1.3 (see Theorem 2.6). To finish the section, we will describe some of the key new ideas and techniques that will be introduced in this paper.

The idea behind the definition of the resistance $r(\mathcal{U})$ is quite simple: we will define the ‘induced’ resistance of a point $u \in \mathbf{S}^{d-1}$ (with respect to the stable set $\mathcal{S}(\mathcal{U})$) inductively, by considering the resistance of a small copy of $\mathbf{S}^{d-2}$ centred at u . The resistance of $\mathcal{S}(\mathcal{U})$ will then be defined using the hemisphere in which the maximum resistance is smallest. The details are somewhat technical, however, so the reader may find it helpful to have in mind pictures of the stable sets of the 2- and 3-neighbour models in three dimensions. For $\mathcal{N}_2^3$, the stable set is the six points $\{\pm e_1, \pm e_2, \pm e_3\}$, and for $\mathcal{N}_3^3$ the stable set is the union of the three great circles orthogonal to the standard basis vectors.

First we introduce a family of objects called ‘**S**-stable sets’, which are subsets of the sphere $\mathbf{S}$, and generalize the notion of a stable set of an update family.

⁹In these dynamics, which were introduced in [64], the state of a site x updates to $+$ at rate 1 if there exists $X \in \mathcal{U}$ such that $x + X$ is entirely $+$, and similarly for $-$. When $\mathcal{U} = \mathcal{N}_d^d$, this is equivalent to the zero-temperature nearest-neighbour Ising model, and therefore in this case the conjecture follows from the theorem of Fontes, Schonmann and Sidoravicius [38].

Definition 2.1. Let $\mathcal{T} \subset \mathbf{S}^{d-1}$ and let $\mathbf{S}$ be a sphere (of arbitrary size and dimension) embedded in $\mathbf{S}^{d-1}$. We say that $\mathcal{T}$ is $\mathbf{S}$ -stable if there exists a finite collection $\mathcal{H}_1, \dots, \mathcal{H}_m$ of finite families of closed hemispheres of $\mathbf{S}$ such that

$$\mathcal{T} \cap \mathbf{S} = \bigcap_{i=1}^m \bigcup_{H \in \mathcal{H}_i} H.$$

In particular, the stable set of a d -dimensional update family is $\mathbf{S}^{d-1}$ -stable.

Lemma 2.2. If $\mathcal{U}$ is a d -dimensional update family, then

$$\mathcal{S}(\mathcal{U}) = \bigcap_{X \in \mathcal{U}} \bigcup_{x \in X} \{u \in \mathbf{S}^{d-1} : \langle x, u \rangle \geq 0\}.$$

In particular, $\mathcal{S}(\mathcal{U})$ is $\mathbf{S}^{d-1}$ -stable.

Proof. This is nothing more than the observation that a direction $u \in \mathbf{S}^{d-1}$ is unstable for $\mathcal{U}$ if and only if there exists $X \in \mathcal{U}$ such that $\langle x, u \rangle < 0$ for every $x \in X$. ■

Now, given a sphere $\mathbf{S} \subset \mathbf{S}^{d-1}$, $u \in \mathbf{S}$ and $\eta > 0$, define the subsphere¹⁰

$$S_\eta(\mathbf{S}, u) := \{v \in \mathbf{S} : \|u - v\| = \eta\}.$$

We say that two sets $A, B \subset \mathbb{R}^d$ are *equivalent*, and write $A \equiv B$, if one can be obtained from the other by a composition of translations, dilations and rotations. The following simple lemma, which was proved in [7], says that the intersections of a stable set with any two sufficiently small spheres centred at a given point are equivalent.

Lemma 2.3 ([7, Lemma 2.3]). Let $\mathbf{S} \subset \mathbf{S}^{d-1}$ be a sphere, let $\mathcal{T} \subset \mathbf{S}^{d-1}$ be $\mathbf{S}$ -stable, and let $u \in \mathbf{S}$. Then there exists $\eta_0 = \eta_0(\mathbf{S}, \mathcal{T}, u) > 0$ such that, for all $0 < \eta < \eta_0$,

$$\mathcal{T} \cap S_\eta(\mathbf{S}, u) \equiv \mathcal{T} \cap S_{\eta_0}(\mathbf{S}, u)$$

and $\mathcal{T}$ is $S_\eta(\mathbf{S}, u)$ -stable.

We shall often use Lemma 2.3 implicitly by writing $S_\eta(\mathbf{S}, u)$ without specifying η . This should always be taken to mean that $\eta < \eta_0$, where $\eta_0 = \eta_0(\mathbf{S}, \mathcal{T}, u)$ (for the $\mathcal{T}$ currently under consideration) is given by Lemma 2.3.

We are now ready to define the resistance of a d -dimensional update family. For each $k \in \{0, \dots, d-1\}$, let us write $\mathcal{C}^k$ for the set of all k -dimensional spheres embedded in $\mathbf{S}^{d-1}$, and for each $\mathbf{S} \in \mathcal{C}^k$, let us write $\mathcal{B}(\mathbf{S})$ for the set of all $\mathbf{S}$ -stable sets.

Definition 2.4. For each $1 \leq k \leq d$ and each $\mathbf{S} \in \mathcal{C}^{k-1}$, we define two functions

$$\begin{aligned} \rho^{k-1}(\mathbf{S}; \cdot, \cdot) : \mathcal{B}(\mathbf{S}) \times \mathbf{S} &\rightarrow \{0, 1, \dots, k\}, \\ r^k(\mathbf{S}; \cdot) : \mathcal{B}(\mathbf{S}) &\rightarrow \{1, \dots, k+1\} \end{aligned}$$

¹⁰Here, and throughout the paper, we write $\|\cdot\|$ for the Euclidean norm on $\mathbb{R}^d$.

inductively as follows. Set $r^0 \equiv 1$, and let $1 \leq k \leq d$, $\mathbf{S} \in \mathcal{C}^{k-1}$, $\mathcal{T} \in \mathcal{B}(\mathbf{S})$ and $u \in \mathbf{S}$. We define the *induced resistance of u with respect to $\mathcal{T}$ in $\mathbf{S}$* to be

$$\rho^{k-1}(\mathbf{S}; \mathcal{T}, u) := \begin{cases} r^{k-1}(S_\eta(\mathbf{S}, u); \mathcal{T}) & \text{if } u \in \mathcal{T}, \\ 0 & \text{otherwise,} \end{cases} \quad (4)$$

and the *resistance of $\mathcal{T}$ in $\mathbf{S}$* to be

$$r^k(\mathbf{S}; \mathcal{T}) := \min_H \max_{u \in H} \rho^{k-1}(\mathbf{S}; \mathcal{T}, u) + 1, \quad (5)$$

where the minimum is taken over all open hemispheres $H \subset \mathbf{S}$.

Now, given a d -dimensional update family $\mathcal{U}$, define the *resistance of $\mathcal{U}$* to be

$$r(\mathcal{U}) := r^d(\mathbf{S}^{d-1}; \mathcal{S}(\mathcal{U})).$$

In words, the resistance of $\mathcal{U}$ is determined by the largest induced resistance of a direction in the ‘easiest’ hemisphere of $\mathbf{S}^{d-1}$, and the induced resistance of a stable direction u is given by the resistance of a small sphere centred at u . It was shown in [7, Observation 2.6] that, for each $1 \leq r \leq d + 1$, the update family $\mathcal{N}_r^d$ has resistance r .

The following lemma, which was proved in [7], shows that $r(\mathcal{U})$ refines Definition 1.2.

Lemma 2.5 ([7, Lemma 2.7]). *If $\mathcal{U}$ is a d -dimensional update family, then:*

- (a) $\mathcal{U}$ is supercritical if and only if $r(\mathcal{U}) = 1$.
- (b) $\mathcal{U}$ is critical if and only if $r(\mathcal{U}) \in \{2, \dots, d\}$.
- (c) $\mathcal{U}$ is subcritical if and only if $r(\mathcal{U}) = d + 1$.

Lemma 2.5 follows in a straightforward way from Definition 2.4 (see [7, Section 2] for the details). Having defined $r(\mathcal{U})$, we can now state a more precise universality theorem. Note in particular that, by Lemma 2.5, it implies Theorem 1.3.

Theorem 2.6. *Let $\mathcal{U}$ be a d -dimensional update family.*

- (a) *If $\mathcal{U}$ is supercritical then $p_c(\mathbb{Z}_n^d, \mathcal{U}) = n^{-\Theta(1)}$.*
- (b) *If $\mathcal{U}$ is critical then*

$$p_c(\mathbb{Z}_n^d, \mathcal{U}) = \left(\frac{1}{\log_{(r-1)} n} \right)^{\Theta(1)},$$

where $r = r(\mathcal{U})$.

- (c) *If $\mathcal{U}$ is subcritical then $p_c(\mathbb{Z}^d, \mathcal{U}) > 0$.*

The upper bounds in (a) and (b) were proved in [7], and part (c) was proved in [6, 51]. In the next section we shall state the precise lower bound that will be proved in this paper, and provide a detailed outline of our approach. First, however, let us briefly describe a few of the key new ideas and techniques that will be introduced during the proof.

To begin, observe that, by (5), there exists a finite set

$$\mathcal{T} \subset \{u \in \mathbf{S}^{d-1} : \rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u) \geq r - 1\} \quad (6)$$

intersecting every open hemisphere of $\mathbf{S}^{d-1}$. Our main task (see Proposition 3.5) will be to bound the probability that a ‘ $\mathcal{T}$ -droplet’ D (see Definition 3.2) is ‘internally spanned’ (see Definition 3.4). Roughly speaking, this is the event that the $\mathcal{U}$ -closure of $D \cap A$ (where A is the set of initially-infected vertices) ‘connects’ the faces of D . Intuitively, the most likely way in which this can occur is via the growth of a droplet inside D . It will therefore be necessary to control the growth of the infected set on the faces of this smaller droplet (and on the faces of droplets growing on these faces, and so on). In order to do so, we first construct (in Section 4) a tree of stable directions: for each $u \in \mathcal{T}$, there exists a set of directions (cf. (6)), each with induced resistance at least $r - 2$, that intersects every open hemisphere of $\mathcal{S}_\eta(\mathbf{S}^{d-1}, u)$, and so on. There is, however, a significant subtlety in doing this, which we discuss in Section 3.4.

Our main tools for controlling growth on the faces of droplets (that are themselves growing on the faces of droplets) are ‘icebergs’, which we introduce in this paper, and which are analogues of droplets for growth that is assisted by a collection of half-spaces.¹¹ Icebergs play a central role in our proof (see, for example, our main induction hypothesis, Definition 8.1), and in order to work with them we will need to develop an entirely new set of tools (see Sections 5 and 6). Developing these tools required a number of significant innovations; we mention in particular the ‘iceberg containers’ introduced in Section 5, and the subadditivity lemma for iceberg containers proved in Section 6. These will be described in more detail in Section 3.3.

We would like to emphasize that, although our icebergs are named after their (much simpler) two-dimensional counterparts from [17], the difficulties that we encounter in dealing with the d -dimensional versions are completely different, and have no analogue in two dimensions. In particular, it turns out that the usual subadditivity lemma for droplets in two dimensions *is not true* in higher dimensions, and recovering this situation (which, at first sight, appears to be catastrophic) requires a series of substantial conceptual innovations. To be more precise, we actually prove a subadditivity lemma for iceberg containers using a subtle and somewhat surprising notion of size (see Definition 6.13). The definition of iceberg containers (see Definition 5.5) is moreover delicate and (necessarily) somewhat complicated, and was designed very carefully in order to allow us to prove this subadditivity lemma.

In order to bound the probability that a droplet is internally spanned, we need to bound the probability of all possible ways in which growth can occur, not only the simplest. As mentioned above, in all previous work on models with $r \geq 3$ this was done using the Cerf–Cirillo method, which is not available to us due to the lack of symmetry. We have therefore

¹¹These objects are the higher-dimensional analogues of the two-dimensional icebergs introduced in [17]. To see why they were given this name, the reader should imagine the half-spaces as being an ocean.

developed, in Sections 7 and 8, a new method, which was inspired by Holroyd's method of hierarchies [53] (see Lemmas 7.2 and 7.4), but is used here in a very different way, and for a very different purpose. To be more precise, Holroyd used his method to group the possible paths of a growing droplet in the 2-neighbour model into a bounded collection of 'types', which then allowed him to use a simple union bound. Our hierarchies, on the other hand, are used to handle droplets at *very* different scales, and to deal with the lack of symmetry. This is also the first time that hierarchies have been used in a setting where $r(\mathcal{U}) \geq 3$.¹² We remark that we also do not use a union bound over hierarchies, but instead a more subtle inductive argument (see Lemma 8.3).

3. The main theorem, and an outline of the proof

Our task in the remainder of the paper is to prove the following theorem, which implies the lower bound in Theorem 2.6 for critical update families.

Theorem 3.1. *Let $\mathcal{U}$ be a critical d -dimensional update family. Then*

$$p_c(\mathbb{Z}_n^d, \mathcal{U}) \geq \left(\frac{1}{\log_{(r-1)} n} \right)^{d-1+o(1)},$$

where $r = r(\mathcal{U})$ is the resistance of $\mathcal{U}$.

The constant $d - 1$ in the exponent is sharp for the 2-neighbour model, for which the threshold (see (2)) was first determined by Aizenman and Lebowitz [4]. Recall also that, by the results of [8], the best possible exponent is uncomputable in general if $r < d$. Determining the correct exponent for all d -dimensional update families with $r(\mathcal{U}) = d$ also seems likely to be extremely challenging, and would require significant new ideas.

The rest of this section will be taken up by a detailed outline of the proof of Theorem 3.1. In particular, we shall introduce (informally) the main new tools in the proof: bounding trees, icebergs, and one-step hierarchies. First, however, we shall introduce the fundamental notion of an 'internally spanned $\mathcal{T}$ -droplet', which will then allow us to state the key bound that we shall prove on the probability that a $\mathcal{T}$ -droplet is internally spanned (see Proposition 3.5). To avoid repetition, let us fix for the rest of the paper a critical d -dimensional update family $\mathcal{U} = \{X_1, \dots, X_m\}$ with $r(\mathcal{U}) = r$.

3.1. Internally spanned droplets

We begin by defining droplets, which are just discrete polytopes. Note that we do not insist that droplets are finite.

¹²While hierarchies were used in [10, 11] to study the r -neighbour model, the authors of those papers first used the symmetry of the model to reduce to a setting whose behaviour resembled that of the 2-neighbour model, and then used hierarchies in a similar way to Holroyd's original paper.

Definition 3.2. Given a finite set $\mathcal{T} \subset \mathbb{S}^{d-1}$, a $\mathcal{T}$ -droplet is a non-empty set of the form

$$D = \bigcap_{u \in \mathcal{T}} \{x \in \mathbb{Z}^d : \langle x, u \rangle \leq a_u\}$$

for some collection $\{a_u \in \mathbb{R} : u \in \mathcal{T}\}$.

Note that if $\mathcal{T} \subset \mathcal{S}(\mathcal{U})$, then a $\mathcal{T}$ -droplet is closed under the $\mathcal{U}$ -bootstrap process; in particular, this holds if $\mathcal{T}$ is chosen as in (6). We can therefore use such droplets to bound the growth of the process. In order to relate percolation to droplets, we shall use an idea from [4, 22], which we generalize to our setting as follows.

Suppose that $[A]_{\mathcal{U}} = \mathbb{Z}_n^d$, and run the $\mathcal{U}$ -bootstrap process with initial set A , except infecting only one new site in each step. In each step, the diameter¹³ of the largest ‘component’ K of infected sites at most triples (once it is larger than some constant), and so at some point in the process K will have diameter $\Theta(\log n)$, and moreover (if we choose the definition of component correctly) will be infected using only sites of $K \cap A$. In particular, if D is the minimal $\mathcal{T}$ -droplet containing K , then $K \subset [D \cap A]_{\mathcal{U}}$.

We shall, roughly speaking, follow this strategy (see Lemma 8.12), using the following notion of connectivity. Let us define $R = R(\mathcal{U})$, the *range* of $\mathcal{U}$, to be

$$R := 2 \cdot \max_{x \in X \in \mathcal{U}} \|x\|. \quad (7)$$

Definition 3.3. We say that a set $K \subset \mathbb{R}^d$ is *strongly connected* if it is connected in the graph G with vertex set $\mathbb{R}^d$ and edge set

$$E(G) = \{xy : \|x - y\| \leq R\}.$$

We also say that two sets $K, K' \subset \mathbb{R}^d$ are *strongly connected to each other* if there exist $x \in K$ and $y \in K'$ with $\|x - y\| \leq R$.

Now, observe that if we run the $\mathcal{U}$ -bootstrap process on a set A , except infecting only one vertex in each step, and if K is a strongly connected component of infected sites at time t , then $K \subset [K \cap A]_{\mathcal{U}}$. To see this, simply note that if $X \in \mathcal{U}$ and $x \in \mathbb{Z}^d$, then x is strongly connected to each vertex of $x + X$. If $\mathcal{T}$ is a set of stable directions, then it follows that the minimal $\mathcal{T}$ -droplet D containing K has the following key property.

Definition 3.4. A $\mathcal{T}$ -droplet D is *internally $\mathcal{T}$ -spanned* if there exists a set $K \subset D \cap A$ such that $[K]_{\mathcal{U}}$ is strongly connected and D is the minimal $\mathcal{T}$ -droplet containing $[K]_{\mathcal{U}}$.

Note that the event that D is internally $\mathcal{T}$ -spanned depends on the set A , and also on the update family $\mathcal{U}$ (which we fixed above). We can now state our key result about droplets, which (by the argument above, together with a simple union bound over droplets) implies a slightly weaker version of Theorem 3.1, with a worse exponent.

¹³The diameter $\text{diam}(\mathcal{X})$ of a set $\mathcal{X} \subset \mathbb{R}^d$ is the supremum of $\|x - y\|$ over all $x, y \in \mathcal{X}$.

Proposition 3.5. *There exist $\varepsilon > 0$ and $\delta > 0$, and a finite set $\mathcal{T} \subset \mathcal{S}(\mathcal{U})$ intersecting every open hemisphere of $\mathbf{S}^{d-1}$, such that the following holds. Let D be a $\mathcal{T}$ -droplet with*

$$\text{diam}(D) \leq \exp_{(r-2)}(p^{-\varepsilon}).$$

Then D is internally $\mathcal{T}$ -spanned with probability at most $\exp(-\delta \cdot \text{diam}(D))$.

We remark that in the proof of Theorem 3.1 we shall actually use Proposition 8.4, which is a slightly stronger and more technical variant of this result.

Before describing how we shall prove Proposition 3.5, let us begin by explaining why the methods used in previous works fail for general models. As mentioned above, all previous lower bounds for models with $r \geq 3$ (see [10, 11, 22, 23, 54, 76]) used a technique that was introduced by Cerf and Cirillo [22] during their work on the r -neighbour process. Note that in that setting one can take $\mathcal{T} = \{\pm e_1, \dots, \pm e_d\}$, and thus $\mathcal{T}$ -droplets are cuboids. To bound the probability that a cuboid D is internally $\mathcal{T}$ -spanned, they partitioned D into $(d-1)$ -dimensional slabs of thickness 2, coupled the r -neighbour process in D with independent $(r-1)$ -neighbour processes in these slabs, noting that each vertex has only one neighbour in D outside its own slab, and used induction on r to control the process inside the slabs. Observe that symmetry plays a crucial role in the coupling.

In the definition of $\mathcal{U}$ we made no assumptions of symmetry, and one of our main challenges in proving Proposition 3.5 will be to deal with the lack of such an assumption. Indeed, the strategy described above of partitioning into lower-dimensional slabs would not work for us, since we cannot expect to have two opposite stable directions, and therefore growth inside slabs can be assisted by $\mathcal{T}$ -droplets growing outside those slabs and ‘breaking in’ – such $\mathcal{T}$ -droplets are sometimes called ‘savers’.

3.2. One-step hierarchies

In order to deal with the lack of symmetry, we introduce a novel ‘one-step’ variant of the ‘hierarchies’ technique that was pioneered by Holroyd [53] in his work on the 2-neighbour model, and which has since become a fundamental tool in the study of bootstrap percolation (see for example [10, 17, 36]). Hierarchies are normally used in order to carefully analyse the possible growth paths of droplets, and thus to prove sharp threshold results for symmetric models with $r(\mathcal{U}) = 2$. An important innovation of [17] was to apply them for a different purpose: to deal with the lack of symmetry in general two-dimensional models.¹⁴ Here we again use them for this purpose, but the difficulties involved in doing so are very different, and we shall need to modify the method much more dramatically. In particular, in all previous applications of the method of hierarchies the droplets being studied had size $p^{-\Theta(1)}$. Here, on the other hand, we will need to ‘track’ the growth of droplets (and, more generally, icebergs) from size $p^{-O(1)}$ all the way up to the ‘critical’ scale, $\exp_{(r-2)}(p^{-\Theta(1)})$.

¹⁴To be precise, the authors of [17] only used hierarchies for unbalanced models, whose growth is asymptotically one-dimensional.

In the usual hierarchies method (including in [17]) one uses a union bound over hierarchies, but when $r(\mathcal{U}) \geq 3$ there would be far too many hierarchies for this to work. In order to overcome this problem we introduce a simpler family of objects, which we call ‘one-step hierarchies’ (see Section 7). Roughly speaking, these consist either of two icebergs that are sufficiently large and disjointly ‘iceberg spanned’,¹⁵ or of two icebergs, one inside the other, where the smaller iceberg is iceberg spanned and also a certain ‘crossing event’ holds (see Lemma 7.2). The construction of these objects relies heavily on the theory of icebergs that we develop in Sections 5 and 6 (see also Section 3.3), but otherwise follows the familiar strategy that was introduced in [53].

The second main innovation of our new technique is the way in which we use these one-step hierarchies to bound the probability that an iceberg is spanned. We do so by replacing the simple union bound used in previous applications by a carefully designed induction (see Definition 8.1 and Lemma 8.3). Intuitively, this allows us to ‘pay’ for the number of choices for the hierarchies by allowing our bound on the probability that an iceberg is spanned to increase gradually with the diameter. More precisely, our bound will increase (roughly speaking) from p^x when the diameter x is small, to e^{-x} at the critical scale $x = \exp_{(r-2)}(p^{-\Theta(1)})$, see (84). We will apply the induction hypothesis to the icebergs in our one-step hierarchy, and also to a smaller iceberg whose existence we deduce from the crossing event (see Lemma 7.4).

The reader may be wondering why we worked with icebergs in the discussion above, rather than droplets. The reason is that when $r \geq 4$, we shall need to apply the strategy outlined above for droplets growing ‘on the sides’ of larger droplets, and for droplets growing on the sides of those droplets, and so on. In this more complex setting the role of droplets is played by icebergs, and in order to carry out our approach using one-step hierarchies we will need to develop a new general method for using icebergs to bound growth in the $\mathcal{U}$ -bootstrap process ‘assisted’ by certain unions of half-spaces.

3.3. Icebergs

Consider a sequence of droplets, each growing on the sides of the droplets that came before it in the sequence. In order to study the growth in such situations, it will be convenient to replace each droplet in the sequence by a half-space $\mathbb{H}_{u_i}^d$, where u_i is the normal to the side on which we are growing, and consider the process ‘assisted’ by the set $\mathbb{H}_{u_1}^d \cup \dots \cup \mathbb{H}_{u_k}^d$. We will choose the directions $u_1, \dots, u_k$ carefully so that this set is closed under the $\mathcal{U}$ -bootstrap process (see Proposition 4.10). In Sections 5 and 6 we shall develop a theory of ‘icebergs’, which play the role of droplets in this setting.

As mentioned above, icebergs are named after some similar (but much simpler) two-dimensional objects that were introduced in [17] in order to control the growth on the sides of droplets in models with ‘drift’ (such as the Duarte model). In higher dimensions

¹⁵The definition of an iceberg being ‘iceberg spanned’ is introduced in Section 5 (see Definitions 5.5 and 5.9). It is somewhat more complicated than the corresponding notion for droplets.

icebergs are significantly more complicated objects, and they will play a *very* different (and *much* more fundamental) role in our proof than they did in [17]. Indeed, they are the central object of study in this paper: all of our results in Sections 5–8, including our main induction hypothesis (see Definition 8.1), are about properties of icebergs.

Very roughly speaking, our strategy for controlling icebergs follows the approach that was first introduced in the groundbreaking work of Aizenman and Lebowitz [4], and that has been used in essentially all work on bootstrap percolation since then. In particular, we will need to design an ‘iceberg spanning algorithm’ (see Definition 5.10), and prove analogues for icebergs of the classical ‘extremal’ and ‘Aizenman–Lebowitz’ lemmas (see Lemmas 6.18 and 6.19). However, as always, the difficulty of carrying out this approach lies in the details, and here the obstructions to doing so are much greater than usual.

In this subsection we will concentrate on the central obstruction that we shall need to overcome in the proof of Theorem 3.1: designing and proving a suitable ‘subadditivity’ lemma for icebergs. For the 2-neighbour model such a result is immediate: it simply says that if two (axis-parallel) rectangles R_1 and R_2 intersect, then the semi-perimeter of the smallest rectangle that contains both is at most the sum of the semi-perimeters of R_1 and R_2 . A similar result also holds for arbitrary polygons in two dimensions; see [17, Lemma 6.4] for a simple proof. For polytopes in three or more dimensions, however, the corresponding subadditivity lemma turns out to be false!

Overcoming this obstacle is the main topic of Sections 5 and 6, and required a number of significant conceptual innovations, as well as the development of an entirely new and substantially more involved set of tools. The result we eventually manage to prove is somewhat complicated (see Lemma 6.14): it is a subadditivity lemma for a certain measure (the ‘iceberg diameter’) of the size of the ‘iceberg container’ of the ‘ u -closure’ of a set of vertices K . To unpack this a little, let us take these definitions in reverse order. First, the u -closure of K is the set of vertices infected in the $\mathcal{U}$ -bootstrap process with the help of the set¹⁶ $\mathbb{H}_{u_1}^d \cup \dots \cup \mathbb{H}_{u_k}^d$, starting from set K (see Definition 5.1). Next, the iceberg container of a set (see Definition 5.5) is an iceberg that is chosen very carefully so that it is not too large, and has several useful properties (in particular, various kinds of monotonicity). Finally, the ‘iceberg diameter’ (see Definition 6.13) of an iceberg (J, u) is δ^k times the ‘weighted diameter’ of a certain (again, carefully chosen) polytope that contains J , where $\delta > 0$ is a small constant and k is a function of the ‘type’ u of the iceberg (roughly speaking, it is the number of half-spaces used when constructing J). The weighted diameter (see Definition 6.2), which plays a key role in the proof, measures the size of a polytope by taking a suitable linear sum of the distances of the faces to the origin. Moreover, and crucially, we can control its interaction with the definition of iceberg containers, and also with two particular families of polytopes (see Definition 6.6 and several of the lemmas in Section 6, in particular Lemmas 6.7 and 6.11).

We remark that our subadditivity lemma also plays a key role in the construction of our one-step hierarchies (see Lemma 7.2), and that our induction hypothesis (Definition 8.1)

¹⁶Here $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k = u$ is a directed path in the bounding tree of $\mathcal{U}$; see Sections 3.4 and 4.

will be defined using the iceberg diameter. We consider the techniques developed in the proof of Lemma 6.14 to be of significant independent interest, and expect them to have many further applications in the study of high-dimensional cellular automata.

3.4. Bounding trees

An important point, which we have so far avoided discussing, is that of exactly *which* stable directions and half-spaces we use to define our droplets and icebergs. This question will be the subject of Section 4, but let us briefly describe here the basic idea. First, choosing the set $\mathcal{T}$ in Proposition 3.5 is straightforward, since the definition of $r(\mathcal{U})$ implies (see Definition 2.4) that for every open hemisphere $H \subset \mathbf{S}^{d-1}$, there exists a direction $u \in H$ such that

$$\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}(\mathcal{U}), u) \geq r - 1.$$

We can take $\mathcal{T}$ to be any finite collection of these directions that intersects every open hemisphere. Now, for each $u \in \mathcal{T}$, we need to choose a collection of directions that we can use to control growth on the u -side of a $\mathcal{T}$ -droplet. We again use Definition 2.4, which implies that for every open hemisphere $H \subset \mathbf{S}_u := S_\eta(\mathbf{S}^{d-1}, u)$, we have

$$\rho^{d-2}(\mathbf{S}_u; \mathcal{S}(\mathcal{U}), v) \geq r - 2 \quad (8)$$

for some direction $v \in H$. We choose a finite collection of these directions that intersects every open hemisphere of $S_\eta(\mathbf{S}^{d-1}, u)$, and use these directions (together with $-u$) to define the icebergs that we will use to control the $\mathcal{U}$ -bootstrap process assisted by $\mathbb{H}_u^d$.

Now, however, comes a subtle and important point. If $r \geq 4$, then for each v chosen above, we need to choose a collection of directions that we can use to control growth on the v -side of an iceberg that is itself growing on the u -side of a $\mathcal{T}$ -droplet. Using Definition 2.4 in a naive way, we would obtain, for each open hemisphere $H \subset S_\eta(\mathbf{S}_u, v)$, a direction $w \in H$ such that

$$\rho^{d-3}(S_\eta(\mathbf{S}_u, v); \mathcal{S}(\mathcal{U}), w) \geq r - 3.$$

However, choosing a collection of such directions that intersects every open hemisphere of $S_\eta(\mathbf{S}_u, v)$ is not sufficient. The issue is that $S_\eta(\mathbf{S}_u, v)$ has too small a dimension, which means that droplets formed by these directions together with $-v$ may be infinite, and we need them to be finite (see Lemma 4.12). What we would like to do instead is to choose a direction in each hemisphere of $S_\eta(\mathbf{S}^{d-1}, v)$, but a priori we do not know that there exist directions with suitable resistances here. In order to resolve this issue, we use a property of stable sets (Lemma 4.3) that was proved in [7], and which plays a key role in the proof of the upper bounds in Theorem 2.6. This lemma implies that if $v \in \mathbf{S}_u$ satisfies (8), then

$$\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}(\mathcal{U}), v) \geq r - 2,$$

and hence, for every open hemisphere $H \subset \mathbf{S}_v := S_\eta(\mathbf{S}^{d-1}, v)$, there exists $w \in H$ with

$$\rho^{d-2}(\mathbf{S}_v; \mathcal{S}(\mathcal{U}), w) \geq r - 3.$$

We choose a finite collection of these directions that intersects every open hemisphere of $\mathbf{S}_v$, and use these directions (together with $-u$ and $-v$) to define the icebergs that control growth on the v -side of an iceberg that is growing on the u -side of a $\mathcal{T}$ -droplet.

Iterating this process gives a tree of stable directions that we use to define our icebergs; however, there are two other subtle points that we must consider when choosing this tree: the directions along each path from the root to a leaf must be linearly independent, and the union of any collection of half-spaces in these directions must be $\mathcal{U}$ -closed. These conditions will allow us define icebergs (for this we will need the boundaries of any such collection of half-spaces to intersect), and to show that they are $\mathcal{U}$ -closed, respectively. To deal with the requirement that the directions along a path are linearly independent (which only causes problems when $r \geq 5$), we shall prove a slightly technical lemma, which allows us to avoid the subspace spanned by the ancestors of a vertex in the tree when choosing its children (see Lemma 4.4). To prove that the union of half-spaces is closed is simpler: we just need to take η sufficiently small (see Proposition 4.10).

3.5. The lattice Λ

To finish this section, let us note a slightly subtle technical point: it will be convenient in the proof to treat the origin in $\mathbb{R}^d$ as different from the origin in $\mathbb{Z}^d$, since we shall need to consider shifted (continuous) half-spaces whose boundaries do not necessarily contain a lattice point (see Section 7). Equivalently, we shall think of the lattice $\mathbb{Z}^d$ as being translated by some unknown element of $\mathbb{R}^d$. The following lemma records the fact that shifting $\mathbb{Z}^d$ in this way does not affect the stable set $\mathcal{S}(\mathcal{U})$.

Lemma 3.6. *Let $u \in \mathbf{S}^{d-1}$ and $a \in \mathbb{R}$, and define $\mathbb{H}_u^d(a) := \{x \in \mathbb{Z}^d : \langle x, u \rangle < a\}$. Then*

$$[\mathbb{H}_u^d(a)]_{\mathcal{U}} = \mathbb{H}_u^d(a) \iff [\mathbb{H}_u^d]_{\mathcal{U}} = \mathbb{H}_u^d \iff u \in \mathcal{S}(\mathcal{U}).$$

Proof. The second equivalence is the definition (3) of $\mathcal{S}(\mathcal{U})$, so it suffices to prove, for each $X \in \mathcal{U}$, that $X \subset \mathbb{H}_u^d$ if and only if $x + X \subset \mathbb{H}_u^d(a)$ for some $x \in \mathbb{Z}^d \setminus \mathbb{H}_u^d(a)$. One direction is immediate, since if $\langle x + y, u \rangle < a$ and $\langle x, u \rangle \geq a$, then $\langle y, u \rangle < 0$. In the other direction, we are similarly done if there exists $x \in \mathbb{Z}^d$ with $\langle x, u \rangle = a$, so assume not, and suppose that $X \subset \mathbb{H}_u^d$. Recalling that X is finite, set

$$\gamma := \max \{\langle y, u \rangle : y \in X\},$$

and note that $\gamma < 0$. Now, by increasing a if necessary (without changing $\mathbb{H}_u^d(a)$) we may assume that there exists $x \in \mathbb{Z}^d \setminus \mathbb{H}_u^d(a)$ with $\langle x, u \rangle < a - \gamma$. But now $\langle x + y, u \rangle < a$ for every $y \in X$, and hence $x + X \subset \mathbb{H}_u^d(a)$, as required. ■

Since the properties stated in Section 2 only depended on the stable set $\mathcal{S}(\mathcal{U})$, it follows from Lemma 3.6 that they all hold equally well in the ‘shifted’ setting. We shall therefore assume that the $\mathcal{U}$ -bootstrap process takes place on a lattice

$$\Lambda := y + \mathbb{Z}^d \subset \mathbb{R}^d$$

for some arbitrary $y \in \mathbb{R}^d$, and redefine $\mathbb{H}_u^d$ to be the discrete half-space in this lattice, i.e., $\mathbb{H}_u^d := \{x \in \Lambda : \langle x, u \rangle < 0\}$. It will only be important in a few places that we allow $y \notin \mathbb{Z}^d$, and these will be pointed out to the reader when they occur.

4. Bounding trees

Recall that in Section 3 we fixed integers $d \geq r \geq 2$, and also a d -dimensional update family $\mathcal{U} = \{X_1, \dots, X_m\}$ with $r(\mathcal{U}) = r$. The purpose of this section is to select the various finite subsets of $\mathcal{S} = \mathcal{S}(\mathcal{U})$ whose elements will define the faces of the droplets and icebergs that we shall later use to bound growth under the action of the $\mathcal{U}$ -bootstrap process. These sets of stable directions will be constructed in Proposition 4.10 via a directed tree T of depth $r - 1$, rooted at $\mathbf{0}$ (the origin) and with all other vertices elements of $\mathbf{S}^{d-1}$. This tree will have the following properties:

- (a) $\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u) \geq r - i$ for every vertex $u \in V(T)$ at distance i from the root;
- (b) if $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_{r-1}$ is a directed path in T , then the set $\{u_1, \dots, u_{r-1}\}$ is linearly independent;
- (c) the children of the vertex $u \in V(T)$ (if they exist) all lie in a sufficiently small $(d - 1)$ -sphere $S_\eta(\mathbf{S}^{d-1}, u) \subset \mathbf{S}^{d-1}$ centred at u .

The vertices of T at depth 1 from the root will form the set $\mathcal{T}$ in Proposition 3.5, and for each path $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k$ in T with u_k not a leaf, the set of children of u_k , together with $\{-u_1, \dots, -u_k\}$, will form the set of directions used to define ‘ u -iceberg droplets’ (see Definition 5.2). Icebergs will be the central subject of Sections 5 and 6.

To underline the importance of the tree T defined in this section, let us remark that, once T has been chosen, the only other property of $\mathcal{U}$ that will be used in the proof of Theorem 3.1 is its range R , which was defined in (7).

We divide the (somewhat technical) definition of T into two parts. In the first, we define the abstract concept of a ‘bounding tree’ without any reference to stable sets. In the second, we introduce the connection between stable sets and bounding trees. The reader will see that we in fact define a pair (T, η) , where T is a tree and η plays the same role as the function whose existence (for a given stable set) is guaranteed by Lemma 2.3. Since we wish to define bounding trees without reference to stable sets, we allow an arbitrary function η in Definition 4.1. We remark that the function η will be discarded almost as soon as our final tree T has been constructed (in Proposition 4.10), since all of the information that we will need in later sections will be encoded in the tree T .

Given a sphere $\mathbf{S} \subset \mathbf{S}^{d-1}$ and a set $B \subset \mathbf{S}$, we say that B is an $\mathbf{S}$ -*bounding set* if $B \cap H \neq \emptyset$ for every open hemisphere $H \subset \mathbf{S}$. We also write $N_{\vec{T}}(v)$ to denote the out-neighbourhood of a vertex v in a directed tree T .

Definition 4.1. A *bounding tree* is a pair $\mathcal{B} = (T, \eta)$, where T is a finite, directed, rooted tree, with root vertex $\mathbf{0} \in \mathbb{R}^d$, all other vertices elements of $\mathbf{S}^{d-1}$, and all edges directed away from the root, and the *magnification function* $\eta: \{1, \dots, s - 1\} \rightarrow (0, \infty)$, for some $1 \leq s \leq d - 1$, is a decreasing function such that the following hold:

- (a) Every leaf of T is at (graph) distance s from $\mathbf{0}$.
 (b) If $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_s$ is a directed path in T , then

$$u_{i+1} \in \mathbf{S}_{u_i} := S_{\eta(i)}(\mathbf{S}^{d-1}, u_i) \quad (9)$$

for each $1 \leq i \leq s-1$, and the set $\{u_1, \dots, u_s\}$ is linearly independent.

- (c) If $u \in V(T)$ is not a leaf, then $N_T^{\rightarrow}(u)$ is an $\mathbf{S}_u$ -bounding set, where $\mathbf{S}_0 := \mathbf{S}^{d-1}$.

The parameter s is the *depth* of $\mathcal{B}$.

Thus, if $\mathcal{B} = (T, \eta)$ is a bounding tree, then the non-root vertices are elements of $\mathbf{S}^{d-1}$, and the children of a vertex u of depth i in T all lie in (and intersect every open hemisphere of) a sphere of radius $\eta(i)$ around u in $\mathbf{S}^{d-1}$. The reader should think of the function η as decreasing sufficiently rapidly so that the children of vertices of greater depth lie much closer to their parents than vertices closer to the root.

Definition 4.2. A bounding tree $\mathcal{B} = (T, \eta)$ of depth $r-1$ is $\mathcal{S}$ -good if the following holds. For every directed path $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_{r-1}$ in T ,

$$\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u_i) \geq r-i \quad (10)$$

for each $1 \leq i \leq r-1$.

Observe in particular that if $\mathcal{B}$ is $\mathcal{S}$ -good then $\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u) \geq 1$ for every vertex $u \in V(T)$, and hence every vertex of the tree T is a stable direction of $\mathcal{U}$.

4.1. Locally inherited resistance

As explained at the end of Section 3, we will use two technical lemmas involving stable sets in order to complete our construction of an $\mathcal{S}$ -good bounding tree. The first of those lemmas was proved in [7].

Lemma 4.3 ([7, Lemma 4.2]). *If $\mathcal{T}$ is an $\mathbf{S}^{d-1}$ -stable set, then*

$$\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{T}, v) \geq \rho^{d-2}(S_{\eta}(\mathbf{S}^{d-1}, u); \mathcal{T}, v) \quad (11)$$

for every $u \in \mathbf{S}^{d-1}$ and $v \in S_{\eta}(\mathbf{S}^{d-1}, u)$.

The second of the two technical lemmas (Lemma 4.4 below) will provide us with a lower bound on the right-hand side of (11). For the statement, it will be convenient to work in a copy of the sphere $\mathbf{S}^{d-2}$, rather than the sphere $\mathbf{S}_u = S_{\eta}(\mathbf{S}^{d-1}, u)$. We shall later deduce a bound for $\mathbf{S}_u$ using the homothety φ_u that maps the centre of $\mathbf{S}_u$ to the origin, and $\mathbf{S}_u$ to the sphere $\mathbf{S}^{d-1} \cap \{u\}^{\perp}$.

Let us write $\mathbf{S}^{d-2}$ to denote an arbitrary $(d-2)$ -dimensional sphere in $\mathbf{S}^{d-1}$, centred at the origin. Recall from Definition 2.1 that if $\mathcal{T}$ is an $\mathbf{S}^{d-2}$ -stable set, then

$$\mathcal{T} \cap \mathbf{S}^{d-2} = \bigcap_{i=1}^t \bigcup_{u \in Y_i} H_u \quad (12)$$

for some finite collection $Y_1, \dots, Y_t$ of finite families of vectors in $\mathbf{S}^{d-2}$, where H_u is the closed hemisphere of $\mathbf{S}^{d-2}$ centred at u ; that is,

$$H_u := \{v \in \mathbf{S}^{d-2} : \langle u, v \rangle \geq 0\}.$$

For each $\mathbf{S}^{d-2}$ -stable set $\mathcal{T}$, let us choose a minimal representation as in (12),¹⁷ and define the *set of centres of $\mathcal{T}$ with respect to $\mathbf{S}^{d-2}$* to be

$$C(\mathbf{S}^{d-2}; \mathcal{T}) := Y_1 \cup \dots \cup Y_t.$$

In particular, if $\mathcal{T} = \mathcal{S}$ then the set of centres can be taken to be (a subset of) the elements of the update sets (cf. Lemma 2.2). Finally, recall that B is an $\mathbf{S}^{d-2}$ -bounding set if $B \cap H \neq \emptyset$ for every open hemisphere $H \subset \mathbf{S}^{d-2}$.¹⁸

The following second technical lemma will allow us to construct a bounding tree whose paths are linearly independent.

Lemma 4.4. *Let $\mathcal{T}$ be an $\mathbf{S}^{d-2}$ -stable set, and let $W \subset \mathbf{S}^{d-1}$ be such that $W^\perp \neq \{\mathbf{0}\}$ and $W^\perp \cap \mathbf{S}^{d-1} \subset \mathbf{S}^{d-2}$. If $C(\mathbf{S}^{d-2}; \mathcal{T}) \subset W^\perp$, then there exists an $\mathbf{S}^{d-2}$ -bounding set B with*

$$\rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, v) \geq r^{d-1}(\mathbf{S}^{d-2}; \mathcal{T}) - 1 \quad (13)$$

for every $v \in B$, and such that $B \cap \langle W \rangle = \emptyset$.

Lemma 4.4 will be used to construct the neighbourhood of a vertex u in our $\mathcal{S}$ -good bounding tree. In our application we shall set $\mathcal{T} := \varphi_u(\mathcal{S} \cap \mathbf{S}_u)$ and $W := \{u_1, \dots, u_s\}$, where φ_u is the homothety that maps $\mathbf{S}_u$ to $\mathbf{S}^{d-2}$, and $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_s = u$ is the directed path in T from the root to u . We remark that the existence of an $\mathbf{S}^{d-2}$ -bounding set whose elements satisfy (13) follows easily from Definition 2.4, so the content of the lemma is that we can choose B avoiding the span of W .

In order to prove Lemma 4.4, we need another lemma from [7], whose statement requires a little extra notation. Given a set $W \subset \mathbf{S}^{d-1}$, let us write

$$\mathbf{S}(W) := \mathbf{S}^{d-1} \cap W^\perp,$$

and define the *projection* $\pi(u, W^\perp)$ of a point $u \in \mathbf{S}^{d-1} \setminus \langle W \rangle$ onto $\mathbf{S}(W)$ to be the unique element of $\mathbf{S}(W)$ such that

$$u = w + \lambda \cdot \pi(u, W^\perp)$$

for some $w \in \langle W \rangle$ and $\lambda > 0$. Note that $\varphi_u(x) = \pi(x, \{u\}^\perp)$, and if $x \in W^\perp$, then

$$\operatorname{sgn}(\langle x, u \rangle) = \operatorname{sgn}(\langle x, \pi(u, W^\perp) \rangle), \quad (14)$$

since $\langle x, u \rangle = \lambda \cdot \langle x, \pi(u, W^\perp) \rangle$. The following is the additional lemma that we need.

¹⁷By ‘minimal’ we mean that no proper subset of the collection of H_u gives the same $\mathcal{T} \cap \mathbf{S}^{d-2}$.

¹⁸In particular, an $\mathbf{S}^{d-2}$ -bounding set can be infinite.

Lemma 4.5 ([7, Lemma 4.5]). *Let $\mathcal{T}$ be an $\mathbf{S}^{d-2}$ -stable set, and let $W \subset \mathbf{S}^{d-1}$. If $C(\mathbf{S}^{d-2}; \mathcal{T}) \subset W^\perp$, then*

$$\rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, u) = \rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, v),$$

for all $u, v \in \mathbf{S}^{d-2} \setminus \langle W \rangle$ such that $\pi(u, W^\perp) = \pi(v, W^\perp)$.

The proof of Lemma 4.5 in [7] is not too difficult: the key observation is that in a representation of $\mathcal{T} \cap \mathbf{S}^{d-2}$ as in (12), the boundary of each hemisphere H_u must either pass through both u and v , or miss both (and hence miss an open ball around each).

With Lemma 4.5 in hand, we are now ready to prove Lemma 4.4.

Proof of Lemma 4.4. We begin by constructing an $\mathbf{S}(W)$ -bounding set B_W such that (13) holds for every $v \in B_W$. We shall later show (in (16)) how to augment B_W to obtain an $\mathbf{S}^{d-2}$ -bounding set B , recalling that $\mathbf{S}(W) \subset \mathbf{S}^{d-2}$ by assumption.

To construct B_W , observe first that, by (5), for each $u \in \mathbf{S}(W)$, there exists $w_u \in \mathbf{S}^{d-2}$ such that $\langle u, w_u \rangle > 0$ and

$$\rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, w_u) \geq r^{d-1}(\mathbf{S}^{d-2}; \mathcal{T}) - 1. \quad (15)$$

Note that $w_u \notin \langle W \rangle$, since $u \in \mathbf{S}(W) \subset W^\perp$. We may therefore set $v_u := \pi(w_u, W^\perp)$, and define

$$B_W := \{v_u : u \in \mathbf{S}(W)\}.$$

To see that B_W is an $\mathbf{S}(W)$ -bounding set, observe that $\langle u, v_u \rangle > 0$ for every $u \in \mathbf{S}(W)$, since $\langle u, w_u \rangle > 0$ by construction, and by (14). Moreover, by Lemma 4.5 and (15), we have

$$\rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, v_u) = \rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, w_u) \geq r^{d-1}(\mathbf{S}^{d-2}; \mathcal{T}) - 1,$$

since $v_u, w_u \notin \langle W \rangle$. It follows that (13) holds for every $v \in B_W$, as claimed.

Now let $Y := \{x \in \langle W \rangle : \|x\| = 1/2\}$, and for each $x \in Y$, define

$$B_x := \{v \in \mathbf{S}^{d-2} : v - x \in W^\perp \text{ and } \pi(v, W^\perp) \in B_W\}.$$

We claim that the set

$$B := B_W \cup \bigcup_{x \in Y} B_x \quad (16)$$

is an $\mathbf{S}^{d-2}$ -bounding set, that $B \cap \langle W \rangle = \emptyset$, and that (13) holds for every $v \in B$.

To show that $B \cap \langle W \rangle = \emptyset$, recall that $B_W \subset W^\perp$ and that $x \in \langle W \rangle$ and $\|x\| = 1/2$ for each $x \in Y$. Since $B_x \subset \mathbf{S}^{d-2} \cap (x + W^\perp)$, it follows that $B_x \cap \langle W \rangle = \emptyset$.

Next, to show that (13) holds for $v \in B$, observe that, by Lemma 4.5,

$$\rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, v) = \rho^{d-2}(\mathbf{S}^{d-2}; \mathcal{T}, \pi(v, W^\perp)),$$

since $v \in B$ and $B \cap \langle W \rangle = \emptyset$. Now, since $\pi(v, W^\perp) \in B_W$ for every $v \in B$, and since (13) holds for every element of B_W , it follows that (13) holds for v .

Our only remaining task is therefore to verify that B is an $\mathbf{S}^{d-2}$ -bounding set. To do so, we need to show that for each $u \in \mathbf{S}^{d-2}$, there exists $v \in B$ such that $\langle u, v \rangle > 0$. If $u \in \langle W \rangle$ then set $x := u/2$, and note that $x \in Y$. Now, observe that B_x is non-empty, because B_W is an $\mathbf{S}(W)$ -bounding set and $W^\perp \neq \{0\}$, and so B_W is non-empty. We may therefore choose an arbitrary element $v \in B_x$, and observe that $v = x + y$ for some $y \in W^\perp$, so $\langle u, v \rangle = \langle u, x \rangle = 1/2 > 0$, as required. On the other hand, if $u \notin \langle W \rangle$, then set $x := \pi(u, W^\perp) \in \mathbf{S}(W)$. Since B_W is an $\mathbf{S}(W)$ -bounding set, there exists $v \in B_W$ with $\langle x, v \rangle > 0$, and since $v \in \mathbf{S}(W)$, it follows by (14) that $\langle u, v \rangle > 0$. Thus B is an $\mathbf{S}^{d-2}$ -bounding set, as claimed, and this completes the proof of the lemma. ■

4.2. Constructing the bounding tree

We are now ready to show that there exists an $\mathcal{S}$ -good bounding tree. However, the existence of such trees will not be sufficient for our purposes; the tree also needs to be chosen so that the set $\mathbb{H}_{u_1}^d \cup \dots \cup \mathbb{H}_{u_{r-1}}^d$ is $\mathcal{U}$ -closed for every path $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_{r-1}$ in T . We shall show in Proposition 4.10 that a tree with this property exists. The slightly technical statement of Lemma 4.8 is designed to facilitate the induction step in the proof of Proposition 4.10. In particular, the following variant of Definition 4.2 will be required.

Definition 4.6. A bounding tree $\mathcal{B} = (T, \eta)$ of depth s is *partially $\mathcal{S}$ -good* if, for every directed path $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_s$ in T , we have

$$\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u_i) \geq r - i \quad (17)$$

for each $1 \leq i \leq s$, and moreover, for every $x \in X \in \mathcal{U}$ and $1 \leq i < s$, we have either

$$u_i \in \{x\}^\perp \quad \text{or} \quad \|u_i - \{x\}^\perp\| > r \cdot \eta(i). \quad (18)$$

Observe that a partially $\mathcal{S}$ -good bounding tree of depth $r - 1$ is $\mathcal{S}$ -good. Note also that condition (18) is satisfied automatically as long as we choose $\eta(i)$ sufficiently small depending on u_i (see Definition 4.1).

We first show that there exists a partially $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$ of depth 1, which is the base case of our induction.

Lemma 4.7. *There exists a partially $\mathcal{S}$ -good bounding tree of depth 1.*

Proof. To construct this tree, we simply take as our collection of neighbours of $\mathbf{0}$ a finite $\mathbf{S}^{d-1}$ -bounding set $B \subset \mathcal{S}$ such that $\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u) \geq r - 1$ for all $u \in B$. The existence of such a set follows from our assumption that $r(\mathcal{U}) = r$; indeed, by (5) the set

$$\{\{v \in \mathbf{S}^{d-1} : \langle u, v \rangle > 0\} : u \in \mathbf{S}^{d-1} \text{ and } \rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u) \geq r - 1\}$$

is an open cover of $\mathbf{S}^{d-1}$, and we can let B be the set of centres of the open hemispheres in a finite subcover. Note that the set $\{u\}$ is automatically linearly independent for each $u \in B$, the function η has empty domain, so does not need to be specified, and condition (18) is empty when $s = 1$. ■

We now use the technical lemmas from Section 4.1 to prove the induction step.

Lemma 4.8. *Every partially $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$ of depth $1 \leq s \leq r - 2$ can be extended to a partially $\mathcal{S}$ -good bounding tree $\mathcal{B}' = (T', \eta')$ of depth $s + 1$ with $\eta'(s)$ arbitrarily small. In particular, there exists an $\mathcal{S}$ -good bounding tree.*

Proof. Fix $1 \leq s \leq r - 2$ and a partially $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$ of depth s . Set $\eta'(i) = \eta(i)$ for every $1 \leq i \leq s - 1$, and choose $\eta'(s)$ sufficiently small. To construct the bounding tree $\mathcal{B}' = (T', \eta')$, we need to choose a neighbourhood

$$N_{T'}^{\rightarrow}(u) \subset \mathbf{S}_u := S_{\eta'(s)}(\mathbf{S}^{d-1}, u)$$

for each leaf u of T . This neighbourhood must be an $\mathbf{S}_u$ -bounding set, and each of its elements must satisfy (17) and must not lie in the span of the elements on the path in T from $\mathbf{0}$ to u . As we noted above, condition (18) in Definition 4.6 is automatically satisfied as long as we choose $\eta'(s)$ sufficiently small (depending on T).

To construct these neighbourhoods, let $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_s = u$ be a directed path in T , and observe that, by (4) and since $\mathcal{B}$ is partially $\mathcal{S}$ -good, we have

$$r^{d-1}(\mathbf{S}_u; \mathcal{S}) = \rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u) \geq r - s. \quad (19)$$

We shall use Lemma 4.4 to prove the following claim.

Claim 4.9. *There exists a finite $\mathbf{S}_u$ -bounding set B_u with*

$$\rho^{d-2}(\mathbf{S}_u; \mathcal{S}, v) \geq r^{d-1}(\mathbf{S}_u; \mathcal{S}) - 1 \quad (20)$$

for every $v \in B_u$, and such that $B_u \cap \langle W \rangle = \emptyset$, where $W := \{u_1, \dots, u_s\}$.

Proof of Claim 4.9. Set $\mathcal{T} := \varphi_u(\mathcal{S} \cap \mathbf{S}_u)$, where $\varphi_u: \mathbf{S}_u \rightarrow \mathbf{S}(\{u\})$ is the homothety defined by $\varphi_u(x) = \pi(x, \{u\}^\perp)$. We will apply Lemma 4.4 to the set $\mathcal{T}$, and then map the $\mathbf{S}(\{u\})$ -bounding set given by the lemma to a subset of $\mathbf{S}_u$ using φ_u^{-1} .

Observe first that $\varphi_u(\mathbf{S}_u) = \mathbf{S}(\{u\})$ is a copy of $\mathbf{S}^{d-2}$, that $\mathbf{S}(W) \subset \mathbf{S}(\{u\})$, since $u \in W$, and that $W^\perp \neq \{\mathbf{0}\}$, since $|W| = s < r \leq d$. In order to apply Lemma 4.4, we therefore only need to show that $\mathcal{T}$ is an $\mathbf{S}(\{u\})$ -stable set, and that $C(\mathbf{S}(\{u\}); \mathcal{T}) \subset W^\perp$ for some valid choice of the set of centres of $\mathcal{T}$ with respect to $\mathbf{S}(\{u\})$.

To choose the set $C(\mathbf{S}(\{u\}); \mathcal{T})$, observe first that, by Lemma 2.2, we have

$$\mathcal{S} \cap \mathbf{S}_u = \bigcap_{X \in \mathcal{U}} \bigcup_{x \in X} (H_x \cap \mathbf{S}_u), \quad (21)$$

where $H_x = \{v \in \mathbf{S}^{d-1} : \langle x, v \rangle \geq 0\}$. Moreover, since $\mathcal{U}$ is a finite collection of finite sets and $\eta'(s)$ is sufficiently small, for each $x \in X \in \mathcal{U}$ we have either $u \in \{x\}^\perp$ or $\mathbf{S}_u \cap \{x\}^\perp = \emptyset$. We claim that if $u \in \{x\}^\perp$ then $x \in W^\perp$. Indeed, if $u_i \notin \{x\}^\perp$ for some $1 \leq i < s$ then, since $\mathcal{B}$ is partially $\mathcal{S}$ -good, we have $\|u_i - \{x\}^\perp\| > r \cdot \eta'(i)$, by (18). Recalling that $u = u_s$, it follows that

$$\|u - \{x\}^\perp\| \geq \|u_i - \{x\}^\perp\| - \|u_i - u_s\| > r \cdot \eta'(i) - \sum_{j=i}^{s-1} \eta'(j) > 0,$$

since $\|u_j - u_{j+1}\| = \eta'(j)$ for each $1 \leq j \leq s-1$, and since η' is decreasing and $s \leq r-1$. This contradicts our assumption that $u \in \{x\}^\perp$, and hence $x \in W^\perp$, as claimed.

We have shown that if $x \notin W^\perp$, then $\mathbf{S}_u \cap \{x\}^\perp = \emptyset$, which in turn implies that $H_x \cap \mathbf{S}_u \in \{\emptyset, \mathbf{S}_u\}$. Now, if $H_x \cap \mathbf{S}_u = \emptyset$ then we do not need to include x in (21), and if $H_x \cap \mathbf{S}_u = \mathbf{S}_u$ then we do not need to include X in (21). Therefore

$$\mathcal{S} \cap \mathbf{S}_u = \bigcap_{X' \in \mathcal{U}'} \bigcup_{x \in X'} (H_x \cap \mathbf{S}_u), \quad (22)$$

where $\mathcal{U}'$ is the family of sets $X' = \{x \in X : H_x \cap \mathbf{S}_u \neq \emptyset\}$ such that $X \in \mathcal{U}$ and $H_x \cap \mathbf{S}_u \neq \mathbf{S}_u$ for every $x \in X$. In particular, $x \in W^\perp$ for every $x \in X' \in \mathcal{U}'$.

Applying φ_u to both sides of (22), we obtain

$$\mathcal{T} = \varphi_u(\mathcal{S} \cap \mathbf{S}_u) = \bigcap_{X' \in \mathcal{U}'} \bigcup_{x \in X'} \varphi_u(H_x \cap \mathbf{S}_u).$$

Now, by (14), for each $x \in X' \in \mathcal{U}'$ we have

$$\varphi_u(H_x \cap \mathbf{S}_u) = \{\varphi_u(v) : v \in \mathbf{S}_u, \langle x, v \rangle \geq 0\} = \{z \in \mathbf{S}(\{u\}) : \langle x, z \rangle \geq 0\}, \quad (23)$$

since $\varphi_u(v) = \pi(v, \{u\}^\perp)$ and $x \in \{u\}^\perp$. It follows that $\mathcal{T}$ is an $\mathbf{S}(\{u\})$ -stable set, and that there exists a valid choice for $C(\mathbf{S}(\{u\}); \mathcal{T})$ with

$$C(\mathbf{S}(\{u\}); \mathcal{T}) \subset \bigcup_{X' \in \mathcal{U}'} X' \subset W^\perp,$$

as claimed. By Lemma 4.4, it follows that there exists an $\mathbf{S}(\{u\})$ -bounding set B' with

$$\rho^{d-2}(\mathbf{S}(\{u\}); \mathcal{T}, v) \geq r^{d-1}(\mathbf{S}(\{u\}); \mathcal{T}) - 1 \quad (24)$$

for every $v \in B'$, and such that $B' \cap \langle W \rangle = \emptyset$.

Observe that $B := \varphi_u^{-1}(B')$ is an $\mathbf{S}_u$ -bounding set, since B' is an $\mathbf{S}(\{u\})$ -bounding set and φ_u maps hemispheres of $\mathbf{S}_u$ to hemispheres of $\varphi_u(\mathbf{S}_u) = \mathbf{S}(\{u\})$ (as in (23)). Moreover, (20) holds for every $v \in B$, since (24) holds for every $v \in B'$, and φ_u is a homothety mapping $\mathcal{S} \cap \mathbf{S}_u$ to $\mathcal{T}$, so

$$\rho^{d-2}(\mathbf{S}_u; \mathcal{S}, v) = \rho^{d-2}(\varphi_u(\mathbf{S}_u); \mathcal{T}, \varphi_u(v)) \geq r^{d-1}(\varphi_u(\mathbf{S}_u); \mathcal{T}) - 1 = r^{d-1}(\mathbf{S}_u; \mathcal{S}) - 1$$

for every $v \in B$. Observe also that if $x \in B \cap \langle W \rangle$, then

$$\varphi_u(x) = \pi(x, \{u\}^\perp) = w + \lambda x \in B' \cap \langle W \rangle$$

for some $w \in \langle u \rangle$ and $\lambda > 0$, since $u \in W$. But $B' \cap \langle W \rangle = \emptyset$, and thus also $B \cap \langle W \rangle = \emptyset$.

Finally, since B is an $\mathbf{S}_u$ -bounding set, the set

$$\{\{w \in \mathbf{S}_u : \langle w, v \rangle > 0\} : v \in B\}$$

is an open cover of $\mathbf{S}_u$. Choose a finite subcover, and let B_u be the set of centres of the open hemispheres of $\mathbf{S}_u$ in this subcover. Since $B_u \subset B$, it follows that B_u is a finite $\mathbf{S}_u$ -bounding set, $B_u \cap \langle W \rangle = \emptyset$, and (20) holds for every $v \in B_u$, as required. ■

Now, define $\mathcal{B}' = (T', \eta')$ by adding, for each leaf u of T , the set B_u as the out-neighbourhood of u in T' . We claim that $\mathcal{B}'$ is a partially $\mathcal{S}$ -good bounding tree. To see that $\mathcal{B}'$ is a bounding tree, observe that, by Definition 4.1, we only need to check that if $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_s = u$ is a directed path in T , then the set $\{u_1, \dots, u_s, v\}$ is linearly independent for every $v \in B_u$, which follows because $B_u \cap \langle u_1, \dots, u_s \rangle = \emptyset$, and that B_u is a finite $\mathbf{S}_u$ -bounding set, which is the case by construction. To see that $\mathcal{B}'$ is partially $\mathcal{S}$ -good, observe first that, by (19) and (20), we have

$$\rho^{d-2}(\mathbf{S}_u; \mathcal{S}, v) \geq r - s - 1$$

for every $v \in B_u \subset \mathbf{S}_u$. Now, by Lemma 4.3, we have

$$\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, v) \geq \rho^{d-2}(\mathbf{S}_u; \mathcal{S}, v),$$

and hence (17) holds for v . Since $\mathcal{B}$ was assumed to be partially $\mathcal{S}$ -good and $\eta'(s)$ was chosen sufficiently small, it follows that $\mathcal{B}'$ is partially $\mathcal{S}$ -good, as required.

Finally, to deduce that there exists an $\mathcal{S}$ -good bounding tree, we simply apply the first part of the lemma $r - 2$ times to the partially $\mathcal{S}$ -good bounding tree of depth 1 given by Lemma 4.7. We obtain a partially $\mathcal{S}$ -good bounding tree of depth $r - 1$, which is therefore $\mathcal{S}$ -good, as required. ■

We can now complete the proof of the main result of this section, which states that there exists an $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$ with the property that unions of half-spaces along paths in T are $\mathcal{U}$ -closed. This property will play an important role in allowing us to control the growth of the $\mathcal{U}$ -bootstrap process on the side of a droplet using icebergs (see Section 5). In order to construct a tree with this property, we need to choose a magnification function η that decreases sufficiently quickly, depending not only on the stable set $\mathcal{S}$, but also on the range R of the update family $\mathcal{U}$ (see (7)). The flexibility in Lemma 4.8 of choosing $\eta'(s)$ arbitrarily small will be useful for this purpose.

In order to prove the next lemma when $\Lambda \neq \mathbb{Z}^d$, it will be useful to write

$$\mathbf{H}_u^d := \{x \in \mathbb{R}^d : \langle x, u \rangle < 0\}$$

for the continuous half-space with normal u . We remark that $\mathbf{H}_u^d$ does not depend on Λ and that $\mathbb{H}_u^d = \mathbf{H}_u^d \cap \Lambda$.

Proposition 4.10. *There exists an $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$, not depending on Λ , such that, for every directed path $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_{r-1}$ in T , and every $X \in \mathcal{U}$,*

$$X \not\subset \mathbf{H}_{u_1}^d \cup \dots \cup \mathbf{H}_{u_{r-1}}^d.$$

In particular, the set $\mathbb{H}_{u_1}^d \cup \dots \cup \mathbb{H}_{u_{r-1}}^d$ is $\mathcal{U}$ -closed.

Proof. We shall build the bounding tree $\mathcal{B}$ inductively using Lemma 4.8. Our induction statement is that, for each $1 \leq s \leq r - 1$, there exists a partially $\mathcal{S}$ -good bounding tree $\mathcal{B}_s = (T_s, \eta_s)$ of depth s such that the following holds:

(+) For any path $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_s$ in T_s , if

$$x \in \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_s}^d$$

for some $x \in X \in \mathcal{U}$, then $x \in \mathbf{H}_{u_s}^d$.

(It is perhaps worth emphasizing here that there is no symmetry between the u_i , since they are being chosen successively depending on those directions already chosen.)

Let us first deduce the lemma from the induction hypothesis with $s = r - 1$. Note that $\mathcal{B}_{r-1}$ is an $\mathcal{S}$ -good bounding tree, since it is partially $\mathcal{S}$ -good by construction and has depth $r - 1$. In particular, this implies that $\rho^{d-1}(\mathbf{S}^{d-1}; \mathcal{S}, u_{r-1}) \geq 1$, by (10), and thus $u_{r-1} \in \mathcal{S}$. Now, let $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_{r-1}$ be a directed path in T_{r-1} , and suppose that $X \subset \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_{r-1}}^d$ for some $X \in \mathcal{U}$. By (+), it follows that $X \subset \mathbf{H}_{u_{r-1}}^d$, which contradicts¹⁹ the fact that u_{r-1} is a stable direction. This contradiction implies that there does not exist an update family $X \in \mathcal{U}$ with $X \subset \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_{r-1}}^d$, as required.

To deduce that $\mathbb{H}_{u_1}^d \cup \cdots \cup \mathbb{H}_{u_{r-1}}^d$ is $\mathcal{U}$ -closed, observe that if there exists $y \in \Lambda$ with

$$y + X \subset \mathbb{H}_{u_1}^d \cup \cdots \cup \mathbb{H}_{u_{r-1}}^d \quad \text{and} \quad y \notin \mathbb{H}_{u_1}^d \cup \cdots \cup \mathbb{H}_{u_{r-1}}^d, \quad (25)$$

then $X \subset \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_{r-1}}^d$. Indeed, for each $x \in X$ there exists $i \in [r - 1]$ such that $\langle x + y, u_i \rangle < 0$ and $\langle y, u_i \rangle \geq 0$, and therefore $\langle x, u_i \rangle < 0$.

To prove the induction statement, recall first that, by Lemma 4.7, there exists a partially $\mathcal{S}$ -good bounding tree $\mathcal{B}_1 = (T_1, \eta_1)$ of depth 1. Since the extra condition of property (+) holds automatically when $s = 1$, this suffices to prove the base case.

Now, let $1 \leq s \leq r - 2$, and suppose that we have constructed a partially $\mathcal{S}$ -good bounding tree $\mathcal{B}_s = (T_s, \eta_s)$ of depth s that satisfies (+). By Lemma 4.8, there exists a partially $\mathcal{S}$ -good bounding tree $\mathcal{B}_{s+1} = (T_{s+1}, \eta_{s+1})$ of depth $s + 1$ that extends $\mathcal{B}_s$, and with $\eta_{s+1}(s)$ arbitrarily small. Let $L(T_s)$ denote the set of leaves of T_s , and observe that

$$\max \{ \langle x, u \rangle : x \in \mathbf{H}_u^d \cap B_R(\mathbf{0}), u \in L(T_s) \} < 0,$$

where $B_R(\mathbf{0}) := \{x \in \mathbb{Z}^d : \|x\| \leq R\}$ is the discrete Euclidean ball of radius R centred at $\mathbf{0}$, since $\langle x, u \rangle < 0$ for each $x \in \mathbf{H}_u^d$, and since the sets $B_R(\mathbf{0})$ and $L(T_s)$ are finite.²⁰ Hence, by the continuity of $\langle x, \cdot \rangle$, we may choose $\delta > 0$ such that

$$\max \{ \langle x, v \rangle : x \in \mathbf{H}_u^d \cap B_R(\mathbf{0}), u \in L(T_s), v \in \mathbf{S}^{d-1}, \|u - v\| \leq \delta \} < 0. \quad (26)$$

Choose $\mathcal{B}_{s+1}$ with $\eta_{s+1}(s) \leq \delta$, and note that this choice of $\mathcal{B}_{s+1}$ has not depended on Λ .

It remains to prove that $\mathcal{B}_{s+1}$ satisfies (+), so let $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_{s+1}$ be a path in T_{s+1} , and suppose that

$$x \in \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_{s+1}}^d$$

¹⁹Recall that, by Lemma 3.6, the set $\mathcal{S}$ does not depend on our choice of the lattice Λ .

²⁰We stress that this definition of $B_R(\mathbf{0})$ as a subset of $\mathbb{Z}^d$ rather than as a subset of Λ is intentional, and is because X itself is a subset of $\mathbb{Z}^d$, and our aim is to prove (+). The final assertion of the proposition (which transfers from $\mathbf{H}_u^d$ to $\mathbb{H}_u^d$) has already been proved above, around (25).

for some $x \in X \in \mathcal{U}$. Then either $x \in \mathbf{H}_{u_{s+1}}^d$ (in which case we are done), or by the induction hypothesis we have $x \in \mathbf{H}_{u_s}^d$. Now, note that $x \in B_R(\mathbf{0})$, by (7), that $u_s \in L(T_s)$, and that $\|u_s - u_{s+1}\| = \eta_{s+1}(s) \leq \delta$. Therefore, by (26), if $x \in \mathbf{H}_{u_s}^d$, then $\langle x, u_{s+1} \rangle < 0$, and so $x \in \mathbf{H}_{u_{s+1}}^d$, as required. This completes the induction step, and hence the proof of the lemma. ■

4.3. The tree T

Let us fix, for the rest of the paper (and, in particular, for the rest of the proof of Theorem 3.1), an $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$ satisfying the conclusion of Proposition 4.10, and with $\eta(1)$ sufficiently small. (In fact, with the exception of the proof of Lemma 4.12 below, we shall only need the tree, T , and can ignore from now on the magnification function η , which was only needed in the inductive construction of T .)

We shall use the following notation and conventions in connection with T :

- Frequently we stipulate that $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_k$ is a path in T ; when we do so, we always mean a directed path, we allow $k = 0$, and we set $u_0 := \mathbf{0}$.
- We write $d_T(\mathbf{0}, u)$ for the graph distance (or, equivalently, the depth) of u from the root $\mathbf{0}$ in T . Thus, $d_T(\mathbf{0}, u) \leq r - 1$ for all $u \in V(T)$.
- If $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_k$ is a path in T , then we write

$$[\mathbf{0}, u_k]_T := \{\mathbf{0}, u_1, \dots, u_k\}.$$

The following sets will be used in the next section to define our ‘icebergs’.

Definition 4.11. Set $M_T(\mathbf{0}) := M_T^+(\mathbf{0}) := N_T^{\rightarrow}(\mathbf{0})$. If u is a non-leaf vertex of T with $u \neq \mathbf{0}$, then define

$$M_T(u) := N_T^{\rightarrow}(u) \cup \{-u\} \quad (27)$$

and

$$M_T^+(u) := N_T^{\rightarrow}(u) \cup \{-u_1, \dots, -u_k\},$$

where $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_k = u$ is the path in T from $\mathbf{0}$ to u .

The following property of $M_T(u)$ will be useful in Sections 5 and 6.

Lemma 4.12. *Let u be a non-leaf vertex of T . Then $M_T(u)$ is an $\mathbf{S}^{d-1}$ -bounding set. Moreover, if u is not the root of T , then $N_T^{\rightarrow}(u)$ is not an $\mathbf{S}^{d-1}$ -bounding set.*

Proof. It is straightforward to show that $N_T^{\rightarrow}(u)$ is not an $\mathbf{S}^{d-1}$ -bounding set, since it does not intersect the open hemisphere centred at $-u$. Indeed, if u is not the root, then by (9) every element of $N_T^{\rightarrow}(u)$ lies on a sphere of radius at most $\eta(1)$ centred at u . Since we assumed that $\eta(1)$ is sufficiently small, it follows that $\langle u, v \rangle > 0$ for each $v \in N_T^{\rightarrow}(u)$.

To show that $M_T(u)$ is an $\mathbf{S}^{d-1}$ -bounding set, we let $H := \{x \in \mathbf{S}^{d-1} : \langle x, v \rangle > 0\}$ for some $v \in \mathbf{S}^{d-1}$ be an arbitrary open hemisphere, and we claim that $M_T(u) \cap H$ is non-empty. Note that we are done immediately if $u = \mathbf{0}$, by Definition 4.1, or if $\langle u, v \rangle < 0$,

since $-u \in M_T(u)$. We are also done if $u = v$, because in that case (as observed above) we have $N_T^{\rightarrow}(u) \subset H$, since we assumed that $\eta(1)$ is sufficiently small.

We may therefore assume that $v = \lambda u + w$ with $\lambda \geq 0$ and $\mathbf{0} \neq w \in \{u\}^\perp$. Recall from (9) that $\mathbf{S}_u = S_\eta(\mathbf{S}^{d-1}, u)$ for some $\eta > 0$. We claim that

$$\{x \in \mathbf{S}_u : \langle x, w \rangle > 0\} \subset H. \quad (28)$$

Note that the left-hand side of (28) is an open hemisphere in $\mathbf{S}_u$, since $\mathbf{0} \neq w \in \{u\}^\perp$, and hence has non-empty intersection with $N_T^{\rightarrow}(u)$, since $N_T^{\rightarrow}(u)$ is an $\mathbf{S}_u$ -bounding set, by Definition 4.1. It therefore suffices to prove (28). To do so, let $x = \mu u + y \in \mathbf{S}_u$ with $y \in \{u\}^\perp$ be such that $\langle x, w \rangle > 0$. Then

$$\langle x, v \rangle = \mu\lambda + \langle x, w \rangle > 0,$$

since $\lambda, \mu \geq 0$, and hence $x \in H$, as required. This proves (28), and the lemma follows. ■

5. Icebergs

In the next two sections we shall define a family of $\mathcal{U}$ -closed structures (‘icebergs’) that generalize droplets, and develop a set of tools that we will use to control the growth of the $\mathcal{U}$ -bootstrap process acting on the side of a droplet. In particular, in Section 6 we will prove a subadditivity property for icebergs (see Lemma 6.14), and use it to deduce an ‘Aizenman–Lebowitz-type’ lemma and an extremal lemma (Lemmas 6.18 and 6.19, respectively). These three results are the basic tools that we will need in order to control the growth of a droplet using icebergs. This general approach can be traced back to the groundbreaking work of Aizenman and Lebowitz [4] on the 2-neighbour model. However, we emphasize once again that the results in these sections are not simply technical generalizations of the results in two dimensions (or for the r -neighbour model), and that the development of these tools will require a number of very substantial innovations.

In this section we lay the groundwork for the proofs of these key lemmas by making a number of important definitions, and proving various basic properties. In particular, we define icebergs (see Definition 5.2), ‘iceberg containers’ (see Definition 5.5), and the notion of a ‘ u -iceberg spanned iceberg’ (see Definition 5.9). We shall also define the ‘ u -iceberg spanning algorithm’ (see Definition 5.10), which will play a role analogous to that of the rectangles process of Aizenman and Lebowitz [4].

The basic purpose of icebergs is to enable us to control growth on the side of a droplet, which may itself be growing on the side of a (higher-dimensional) droplet, and so on. In order to do so, we work in a more ‘generous’ setting in which entire half-spaces are assumed to be already infected; in other words, we consider growth that is ‘assisted’ by certain unions of half-spaces. In this section (and also in Section 6), these unions of half-spaces will always be of the following form: if $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_k$ is a path in T ,²¹ then

²¹Recall that we fixed in Section 4.3 an $\mathcal{S}$ -good bounding tree $\mathcal{B} = (T, \eta)$ satisfying the conclusion of Proposition 4.10.

we define

$$H_T(u_k) := \bigcup_{i=1}^k \mathbb{H}_{u_i}^d = \bigcup_{i=1}^k \{x \in \Lambda : \langle x, u_i \rangle < 0\},$$

so, in particular, $H_T(\mathbf{0}) = \emptyset$. We remark that in Section 7 we will need to consider a slight generalization of the objects defined in this section, in which the boundaries of the half-spaces above are all translated by some element of $\mathbb{R}^d$ (which might not be in $\mathbb{Z}^d$). It is for this reason that we work in the shifted lattice Λ , instead of in $\mathbb{Z}^d$.

Running the $\mathcal{U}$ -bootstrap process with the assistance of the half-spaces in $H_T(u)$ leads us to the following definition.

Definition 5.1. Given a vertex u of T and a finite set $K \subset \Lambda \setminus H_T(u)$, let

$$[K]_u := [K \cup H_T(u)]_{\mathcal{U}} \setminus H_T(u)$$

be the u -closure of K . We say that K is u -closed if $[K]_u = K$.

Recall from Section 3 that we control the growth of the infected set using $\mathcal{T}$ -droplets, which are $\mathcal{U}$ -closed (and hence $\mathbf{0}$ -closed) if $\mathcal{T} = N_T^{\rightarrow}(\mathbf{0})$. For other vertices of T we shall control the growth using icebergs, which are defined as follows.

Definition 5.2. Let u be a non-leaf vertex of T . A u -iceberg droplet is a non-empty set of the form $D \setminus H_T(u)$, where D is an $N_T^{\rightarrow}(u)$ -droplet. An iceberg is a pair $\mathcal{J} = (J, u)$, where J is a u -iceberg droplet.

If we wish to emphasize u , then we shall say that $\mathcal{J}$ is a u -iceberg, or that it is an iceberg of type u . Crucially, u -iceberg droplets are u -closed, as will be shown shortly in Lemma 5.4. Before that, let us observe that they are finite.

Lemma 5.3. *If u is a non-leaf vertex of T and (J, u) is an iceberg, then J is finite.*

Proof. Recalling Definition 4.11, note that every u -iceberg droplet is also an $M_T^+(u)$ -droplet, and that $M_T(u) \subset M_T^+(u)$. Since, by Lemma 4.12, $M_T(u)$ is an $\mathbf{S}^{d-1}$ -bounding set, it follows that u -iceberg droplets are finite, as claimed. ■

Next we prove the key fact that u -iceberg droplets are u -closed, which follows from the property of the tree T guaranteed by Proposition 4.10. The proof is slightly complicated by the fact that we are working in Λ rather than $\mathbb{Z}^d$.

Lemma 5.4. *If u is a non-leaf vertex of T and (J, u) is an iceberg, then J is u -closed.*

Proof. Suppose, for a contradiction, that we have

$$x \notin J \cup H_T(u) \quad \text{and} \quad x + X \subset J \cup H_T(u) \tag{29}$$

for some $X \in \mathcal{U}$ and some $x \in \Lambda$, and recall from Section 4.3 that we chose T so that the conclusion of Proposition 4.10 holds, and in particular so that for every path $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_{r-1}$ in T , the set $\mathbb{H}_{u_1}^d \cup \cdots \cup \mathbb{H}_{u_{r-1}}^d$ is $\mathcal{U}$ -closed. Note that, by Definitions 3.2

and 5.2, we have

$$\begin{aligned} J \cup H_T(u) &= \left(\bigcap_{v \in N_T^{\rightarrow}(u)} (\mathbb{H}_v^d + a_v) \right) \cup H_T(u) \\ &= \bigcap_{v \in N_T^{\rightarrow}(u)} ((\mathbb{H}_v^d + a_v) \cup H_T(u)) \end{aligned} \quad (30)$$

for some collection $\{a_v \in \mathbb{R}^d : v \in N_T^{\rightarrow}(u)\}$. It follows from (29) and (30) that there exists $v \in N_T^{\rightarrow}(u)$ such that $(\mathbb{H}_v^d + a_v) \cup H_T(u)$ contains $x + X$ but does not contain x .

We claim that if $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_{r-1}$ is a path in T containing u and v , then

$$X \subset \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_{r-1}}^d, \quad (31)$$

which contradicts our choice of T , by Proposition 4.10. To prove (31), let $y \in X$, and recall that either $x + y \in \mathbb{H}_v^d + a_v$, or $x + y \in \mathbb{H}_w^d$ for some $\mathbf{0} \neq w \in [\mathbf{0}, u]_T$. We claim that in the former case we have $y \in \mathbf{H}_v^d$, and in the latter case we have $y \in \mathbf{H}_w^d$.

Indeed, if $x + y \in \mathbb{H}_v^d + a_v$, then $\langle x + y - a_v, v \rangle < 0$, and also $\langle x - a_v, v \rangle \geq 0$, since $x \notin \mathbb{H}_v^d + a_v$. It follows that $\langle y, v \rangle < 0$, and so $y \in \mathbf{H}_v^d$. Similarly, if $x + y \in \mathbb{H}_w^d$, then $\langle x + y, w \rangle < 0$, and $\langle x, w \rangle \geq 0$, since $x \notin H_T(u)$. Thus $\langle y, w \rangle < 0$, and so $y \in \mathbf{H}_w^d$.

In either case, it follows that $y \in \mathbf{H}_{u_1}^d \cup \cdots \cup \mathbf{H}_{u_{r-1}}^d$, and so this proves that (31) holds. This contradicts our choice of T , and hence completes the proof of the lemma. ■

We are now ready for an important (and slightly more technical) definition: that of the u -iceberg container of a set $K \subset \Lambda \setminus H_T(u)$. The definition is somewhat delicate, and has been chosen so that the corresponding droplets are u -closed (see Lemma 5.6), and so that it exhibits various types of ‘monotonicity’ (see Lemmas 5.7, 5.8 and 5.14). We remark that iceberg containers will also play a central role in Sections 6 and 7.

Definition 5.5. Let $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_k$ be a path in T , with u_k not a leaf, and let $K \subset \Lambda \setminus H_T(u_k)$ be a finite, non-empty set. For each $0 \leq i \leq k$, let J_i be the minimal u_i -iceberg droplet containing K . We say that a sequence

$$0 = i_0 < i_1 < \cdots < i_\ell \leq k$$

is (K, u_k) -permissible if, for each $0 \leq j < \ell$, the set J_{i_j} is strongly connected to $\mathbf{H}_{u_{i_{j+1}}}^d$.²² Now define $\mathcal{J}_{u_k}(K)$, the u_k -iceberg container of K , to be the iceberg (J_i, u_i) , where i is the largest integer that appears in any (K, u_k) -permissible sequence.

One should think of the permissible sequences in Definition 5.5 as allowing us to take the type of the iceberg container $\mathcal{J}_{u_k}(K)$ to lie as far from the root $\mathbf{0}$ as possible,

²²Recall from Definition 3.3 that two sets are strongly connected to each other if they are at distance at most R . Recall also from Section 4.3 that $u_0 = \mathbf{0}$, and note that the sequence $i_0 = 0$ of length 1 is always permissible, and that the type of a u_k -iceberg container is an element of the set $[\mathbf{0}, u_k]_T = \{\mathbf{0}, u_1, \dots, u_k\}$.




Fig. 1. The u -iceberg container of a set K . Three iceberg-droplets are shown: the minimal u_i -iceberg droplet J_i containing K for each $i \in \{0, 1, 2\}$. Because J_0 is strongly connected to $\mathbf{H}_{u_1}^d$ and J_1 is strongly connected to $\mathbf{H}_{u_2}^d$, the u_2 -iceberg container of K is (J_2, u_2) .

while still allowing us to maintain control over the size of J_i . We are able to maintain this control because J_{i_j} being strongly connected to $\mathbf{H}_{u_{i_j+1}}^d$ ensures that J_{i_j+1} cannot be too much larger than J_{i_j} (see Lemma 6.11). This fact will be crucial in the proof of our subadditivity lemma, Lemma 6.14.

An important property of u -iceberg containers is that their iceberg droplets are u -closed. This is a simple consequence of Lemma 5.4.

Lemma 5.6. *Let u be a non-leaf vertex of T , and let $\mathcal{J}_u(K) = (J, v)$ be the u -iceberg container of a finite set $K \subset \Lambda \setminus H_T(u)$. Then $J \subset \Lambda \setminus H_T(u)$, and J is u -closed.*

Proof. Let $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k = u$ be the path in the tree T from the root to u , and suppose that $v = u_i$. By Definition 5.5, it follows that $J = J_i$, the minimal u_i -iceberg droplet containing K , and in particular $J \subset \Lambda \setminus H_T(u_i)$, by Definition 5.2. Moreover, again by Definition 5.5, there exists a (K, u) -permissible path ending in i , but there does not exist one ending in j for any $i < j \leq k$, and therefore J is not strongly connected to the set $\mathbf{H}_{u_{i+1}} \cup \dots \cup \mathbf{H}_{u_k}$. It follows that $J \subset \Lambda \setminus H_T(u)$.

To prove that J is u -closed, suppose, for a contradiction, that we have

$$x \notin J \cup H_T(u) \quad \text{and} \quad x + X \subset J \cup H_T(u)$$

for some $X \in \mathcal{U}$ and some $x \in \Lambda$. Recall that (by our choice of T) the set $H_T(u)$ is $\mathcal{U}$ -closed, and therefore does not contain the set $x + X$. It follows that $x + X$ intersects J , and hence, since J is not strongly connected to $\mathbf{H}_{u_{i+1}} \cup \dots \cup \mathbf{H}_{u_k}$,

$$(x + X) \cap (\mathbf{H}_{u_{i+1}} \cup \dots \cup \mathbf{H}_{u_k}) = \emptyset.$$

But it now follows that $x + X \subset J \cup H_T(u_i)$, which means that J is not u_i -closed. This contradicts Lemma 5.4, and therefore proves the lemma. ■

The following simple but very useful monotonicity property provides further motivation for Definition 5.5.

Lemma 5.7. *Let u be a non-leaf vertex of T , and let $K \subset K' \subset \Lambda \setminus H_T(u)$ be finite sets. If $\mathcal{J}_u(K)$ is a v -iceberg and $\mathcal{J}_u(K')$ is a w -iceberg, then $d_T(\mathbf{0}, v) \leq d_T(\mathbf{0}, w)$.*

Proof. The lemma follows immediately from the observation that any (K, u) -permissible sequence is also (K', u) -permissible, which in turn follows from the fact that for any $u' \in [0, u]_T$, the minimal u' -iceberg droplet containing K is a subset of the minimal u' -iceberg droplet containing K' . ■

We shall also need the following monotonicity-type properties.

Lemma 5.8. *Let $0 \rightarrow u_1 \rightarrow \dots \rightarrow u_k$ be a path in T , with u_k not a leaf, and let $K \subset \Lambda \setminus H_T(u_k)$ be a finite set.*

- (a) *If $\mathcal{J}_{u_k}(K)$ is a u_i -iceberg, then $\mathcal{J}_{u_j}(K) = \mathcal{J}_{u_k}(K)$ for every $i \leq j \leq k$.*
- (b) *If $\mathcal{J}_{u_k}([K]_{u_k})$ is a u_i -iceberg, then*

$$[K]_{u_j} = [K]_{u_k} \quad \text{and} \quad \mathcal{J}_{u_j}([K]_{u_j}) = \mathcal{J}_{u_k}([K]_{u_k})$$

for every $i \leq j \leq k$.

Proof. In order to prove (a), we shall show that the (K, u_j) -permissible sequences are the same for every $i \leq j \leq k$. Indeed, if the sequence $0 = i_0 < \dots < i_\ell \leq k$ is (K, u_k) -permissible, then $i_\ell \leq i \leq j$, since $\mathcal{J}_{u_k}(K)$ is a u_i -iceberg. Since the iceberg droplets $J_0, \dots, J_j$ are the same in both cases, it follows that a sequence is (K, u_k) -permissible if and only if it is (K, u_j) -permissible. It follows immediately that $\mathcal{J}_{u_j}(K) = \mathcal{J}_{u_k}(K)$.

To deduce (b) from (a), it suffices to prove that $[K]_{u_j} = [K]_{u_k}$ for every $i \leq j \leq k$. To do so, we need to show that the set $[K]_{u_j} \cup H_T(u_k)$ is $\mathcal{U}$ -closed, and that the set $[K]_{u_j} \cap H_T(u_k)$ is empty. Both of these properties will follow from the following claim:

$$[K]_{u_j} \text{ is not strongly connected to } \mathbf{H}_{u_{j+1}}^d \cup \dots \cup \mathbf{H}_{u_k}^d. \quad (32)$$

To prove (32), let $\mathcal{J}_{u_k}([K]_{u_k}) = (J, u_i)$, and observe first that $\mathcal{J}_{u_j}([K]_{u_k}) = (J, u_i)$, by part (a), and therefore J is u_j -closed, by Lemma 5.6. Since $K \subset J$, it follows that $[K]_{u_j} \subset [J]_{u_j} = J$. Now, to deduce (32) simply note that J is not strongly connected to $\mathbf{H}_{u_{j+1}}^d \cup \dots \cup \mathbf{H}_{u_k}^d$, by Definition 5.5, since otherwise there would be a $([K]_{u_k}, u_k)$ -permissible sequence containing an element strictly larger than i .

It follows immediately from (32) and Definition 5.1 that $[K]_{u_j} \cap H_T(u_k)$ is empty. Finally, to show that $[K]_{u_j} \cup H_T(u_k)$ is $\mathcal{U}$ -closed, suppose for a contradiction that

$$x \notin [K]_{u_j} \cup H_T(u_k) \quad \text{and} \quad x + X \subset [K]_{u_j} \cup H_T(u_k)$$

for some $X \in \mathcal{U}$ and $x \in \Lambda$. Recall that $[K]_{u_j} \cup H_T(u_j)$ is $\mathcal{U}$ -closed (by definition) and that $H_T(u_k)$ is $\mathcal{U}$ -closed (by our choice of T), and observe that therefore $x + X$ must intersect both $[K]_{u_j}$ and $\mathbf{H}_{u_{j+1}}^d \cup \dots \cup \mathbf{H}_{u_k}^d$. But this contradicts (32), and hence proves that $[K]_{u_j} \cup H_T(u_k)$ is $\mathcal{U}$ -closed, as claimed.

As noted above, this implies that $[K]_{u_j} = [K]_{u_k}$, and applying part (a) it follows that we also have $\mathcal{J}_{u_j}([K]_{u_j}) = \mathcal{J}_{u_k}([K]_{u_k})$, as required. ■

Having introduced iceberg containers, we are now ready for the following definition, which generalizes to the iceberg setting the property of being internally $\mathcal{T}$ -spanned (see Definition 3.4), and which will play a central role in Sections 7 and 8.

Definition 5.9. Let u be a non-leaf vertex of T . An iceberg $\mathcal{J} = (J, v)$ is said to be u -iceberg spanned if there exists $K \subset J \cap A$ such that $[K]_u$ is strongly connected and $\mathcal{J} = \mathcal{J}_u([K]_u)$. We denote the event that $\mathcal{J}$ is u -iceberg spanned by $\text{Ice}_u^\times(\mathcal{J})$.

Note that $\text{Ice}_u^\times(\mathcal{J})$ is an increasing event, and that if $\mathcal{J} = (J, v)$ is u -iceberg spanned then $v \in [0, u]_T$, since v is the type of a u -iceberg container. Note also that we may assume that $K \subset \Lambda \setminus H_T(u)$, since this does not affect the u -closure $[K]_u$.

Our next task is to define a *spanning algorithm*, which will be one our tools for understanding the event that an iceberg is u -iceberg spanned. The following algorithm generalizes the spanning algorithm of [17] to the iceberg setting.²³

Definition 5.10 (The u -iceberg spanning algorithm). Let u be a non-leaf vertex of T , let K be a finite subset of $\Lambda \setminus H_T(u)$, and set $\mathcal{K}^0 := \{\{x\} : x \in K\}$. Set $t := 0$, and repeat the following steps until STOP:

- If there are two sets $K_1, K_2 \in \mathcal{K}^t$ such that the set

$$[K_1]_u \cup [K_2]_u$$

is strongly connected, then set

$$\mathcal{K}^{t+1} := (\mathcal{K}^t \setminus \{K_1, K_2\}) \cup \{K_1 \cup K_2\},$$

and set $t := t + 1$.

- Otherwise set $T := t$ and STOP.

The output of the algorithm is the u -iceberg span of K ,

$$\langle K \rangle_u := \{\mathcal{J}_u([K']_u) : K' \in \mathcal{K}^T\}.$$

The alert reader may perhaps already have guessed that the u -iceberg span of K is simply the collection of u -iceberg containers of the strongly connected components of $[K]_u$. The following lemma records this simple but important fact.

Lemma 5.11. Let u be a non-leaf vertex of T , and let $K \subset \Lambda \setminus H_T(u)$ be a finite set. Then

$$\langle K \rangle_u = \{\mathcal{J}_u(L_1), \dots, \mathcal{J}_u(L_\ell)\},$$

where $L_1, \dots, L_\ell$ are the strongly connected components of $[K]_u$.

Before proving Lemma 5.11, let us first observe that if $K \subset \Lambda \setminus H_T(u)$ is strongly connected, then $[K]_u$ is also strongly connected. Indeed, this follows from the fact that $H_T(u)$ is $\mathcal{U}$ -closed, by Proposition 4.10, and therefore every vertex of $\Lambda \setminus H_T(u)$ that is infected by the $\mathcal{U}$ -process with initial set $K \cup H_T(u)$, must be strongly connected to some already-infected vertex of $\Lambda \setminus H_T(u)$. We shall use this fact to show that every set $[K']_u$ that occurs in the u -iceberg spanning algorithm is strongly connected.

²³More precisely, setting $u = \mathbf{0}$ in the algorithm below gives [17, Definition 6.14].

Proof of Lemma 5.11. We need to show that the sets $\{[K']_u : K' \in \mathcal{K}^T\}$ at the end of the algorithm are precisely the strongly connected components of $[K]_u$. To see this, note first that $[K']_u$ is strongly connected for every $0 \leq t \leq T$ and $K' \in \mathcal{K}^t$, since this is true at time $t = 0$, and if the set $[K_1]_u \cup [K_2]_u$ is strongly connected for some $K_1, K_2 \in \mathcal{K}^t$, then the set $[K_1 \cup K_2]_u$ is also strongly connected, by the observation above. Moreover, for any two (distinct) sets $K', K'' \in \mathcal{K}^T$, the sets $[K']_u$ and $[K'']_u$ are not strongly connected to each other, since the algorithm stopped at step T .

It therefore remains to show that $[K]_u = \bigcup_{i=1}^k [K_i^T]_u$. Clearly $[K]_u \supset \bigcup_{i=1}^k [K_i^T]_u$, so suppose for a contradiction that $x \in [K]_u$ for some $x \notin \bigcup_{i=1}^k [K_i^T]_u$. It follows, by Definition 5.1, that $x \notin H_T(u)$, and that $x + X \subset \bigcup_{i=1}^k [K_i^T]_u \cup H_T(u)$ for some $X \in \mathcal{U}$. Since $[K_i^T]_u \cup H_T(u)$ is $\mathcal{U}$ -closed for each i (by definition), the set $x + X$ must intersect at least two of the sets $[K_i^T]_u$. But this is impossible, by the definitions of R and of being strongly connected, so we have the required contradiction. ■

Let us also record, for future convenience, the following connection between Definitions 5.9 and 5.10, which is an immediate consequence of Lemma 5.11.

Lemma 5.12. *Let u be a non-leaf vertex of T . Then an iceberg $\mathcal{J} = (J, v)$ is u -iceberg spanned if and only if $\langle K \rangle_u = \{\mathcal{J}\}$ for some $K \subset J \cap A$.*

Proof. By Lemma 5.11, it follows that $\langle K \rangle_u = \{\mathcal{J}\}$ if and only if $[K]_u$ is strongly connected and $\mathcal{J} = \mathcal{J}_u([K]_u)$. The fact that this holds for some $K \subset J \cap A$ is precisely the definition of the event that $\mathcal{J}$ is u -iceberg spanned. ■

The following lemma is another easy (but important) consequence of the u -iceberg spanning algorithm. The proof is essentially identical to that of [17, Lemma 6.20].

Lemma 5.13. *Let u be a non-leaf vertex of T and let $K \subset \Lambda \setminus H_T(u)$ be a finite set with $|K| \geq 2$. If $[K]_u$ is strongly connected, then there exists a partition $K = K_1 \cup K_2$ into non-empty sets such that $[K_1]_u$, $[K_2]_u$ and $[K_1]_u \cup [K_2]_u$ are all strongly connected.*

Proof. Run the u -iceberg spanning algorithm on K and consider the penultimate step. Observe that, for each $0 \leq t \leq T$, the sets in $\mathcal{K}^t$ form a partition of K , and the set $[K']_u$ is strongly connected for every $K' \in \mathcal{K}^t$ (see the proof of Lemma 5.11).

Now, since $[K]_u$ is strongly connected, by Lemma 5.11 we have $\langle K \rangle_u = \{\mathcal{J}_u([K]_u)\}$, and therefore

$$\mathcal{K}^{T-1} = \{K_1, K_2\}$$

for some disjoint, non-empty sets $K_1, K_2 \subset K$ with $K_1 \cup K_2 = K$. By the observations above, both $[K_1]_u$ and $[K_2]_u$ are strongly connected. Finally, since K_1 and K_2 combine at the final step, the set $[K_1]_u \cup [K_2]_u$ must also be strongly connected, as required. ■

To finish the section, let us note another monotonicity property, which will be used in Section 8, and which is a straightforward corollary of Lemma 5.8.

Lemma 5.14. *Let $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k$ be a path in T , with u_k not a leaf. If the iceberg $\mathcal{J} = (J, u_i)$ is u_k -iceberg spanned, then $\mathcal{J}$ is also u_j -iceberg spanned for every $i \leq j \leq k$.*

Proof. By Definition 5.9, if $\mathcal{J}$ is u_k -iceberg spanned, then there exists a set $K \subset J \cap A$ such that $[K]_{u_k}$ is strongly connected and $\mathcal{J} = \mathcal{J}_{u_k}([K]_{u_k}) = (J, u_i)$. But then, by Lemma 5.8, it follows that $[K]_{u_j} = [K]_{u_k}$ is strongly connected and $\mathcal{J}_{u_j}([K]_{u_j}) = \mathcal{J}$ for every $i \leq j \leq k$, and hence $\mathcal{J}$ is u_j -iceberg spanned, as claimed. ■

6. The size of an iceberg

In this section we shall prove our main subadditivity lemma for icebergs, Lemma 6.14. This lemma is an analogue for icebergs of the following inequality from [17, 18]: if $\mathcal{T} \subset \mathbf{S}^1$ is a finite set, and if $D_1, D_2 \subset \mathbb{Z}^2$ are two $\mathcal{T}$ -droplets with $D_1 \cap D_2 \neq \emptyset$, then

$$\text{diam}(D) \leq \text{diam}(D_1) + \text{diam}(D_2), \quad (33)$$

where D is the minimal $\mathcal{T}$ -droplet containing $D_1 \cup D_2$. This inequality has several crucial consequences for the study of two-dimensional (critical) update families, including the so-called ‘Aizenman–Lebowitz lemma’, which implies the existence of an internally spanned ‘critical’ droplet, and the ‘extremal lemma’, which gives a lower bound on $|D \cap A|$ for a droplet D that is internally spanned by the set A .

Variants of these lemmas that hold in the more general setting of icebergs (see Lemmas 6.18 and 6.19) will be required in Section 8. Unfortunately, proving an analogue of (33) for icebergs turns out to be significantly more difficult when $d \geq 3$; we therefore begin by briefly discussing some of the challenges that will be faced in this section, and how we shall go about overcoming them.

The first obvious problem we encounter is that icebergs may have different types, and icebergs of different types have different shapes, and in such cases it is not obvious what the generalization of (33) should be. However, there is a more fundamental problem, which is that, even restricting to the setting of droplets, the direct generalization of (33) to dimensions 3 and higher is false: that is, there exists a finite set $\mathcal{T} \subset \mathbf{S}^2$ and a pair $D_1, D_2 \subset \mathbb{Z}^3$ of intersecting $\mathcal{T}$ -droplets such that

$$\text{diam}(D) > \text{diam}(D_1) + \text{diam}(D_2), \quad (34)$$

where D is the minimal $\mathcal{T}$ -droplet containing $D_1 \cup D_2$.²⁴

In order to address this second problem, we shall introduce in Definition 6.2 a new notion of the ‘size’ of a $\mathcal{T}$ -droplet D , which we denote by $\text{wd}_{\mathcal{T}}(D)$ and call the $\mathcal{T}$ -weighted diameter of D . This function will be chosen so that $\text{wd}_{\mathcal{T}}(D) = \Theta_{\mathcal{T}}(\text{diam}(D))$ for every droplet D (see Lemma 6.4), and so that a subadditivity inequality similar to (33) holds (see Lemma 6.7). We shall then use the function $\text{wd}_{\mathcal{T}}$ to define a measure, $\text{diam}^*(\mathcal{J})$, of the size of an iceberg $\mathcal{J}$ (see Definition 6.13).

²⁴To give a concrete example in three dimensions, consider tetrahedra D_1 and D_2 with corners $\{(1, \pm 2c, 0), (-1, 0, \pm 2c)\}$, and $\{(1, 0, \pm 2c), (-1, \pm 2c, 0)\}$. Each has diameter close to 2 when c is small, but $(3, 0, 0)$ and $(-3, 0, 0)$ are both contained in the smallest $\mathcal{T}$ -droplet containing $D_1 \cup D_2$.

The main result of the section, Lemma 6.14, is a subadditivity inequality for the measure $\text{diam}^*(\mathcal{J})$, but the statement is (necessarily) rather subtle, involving the u -iceberg containers introduced in the previous section (see Definition 5.5), as well as another notion of container (see Definition 6.6) which is carefully chosen to work well with u -closures and (K, u) -permissible sequences (see Lemmas 6.11 and 6.12). Once we have proved Lemma 6.14, it will be relatively straightforward²⁵ to deduce Lemmas 6.18 and 6.19, our ‘Aizenman–Lebowitz’ and ‘extremal’ lemmas for icebergs, respectively.

In order to define $\text{wd}_{\mathcal{T}}$, we need the following simple geometric lemma. Recall that an $\mathbf{S}$ -bounding set is a set $\mathcal{T} \subset \mathbf{S}$ that intersects every open hemisphere of $\mathbf{S}$.

Lemma 6.1. *Let $\mathcal{T} = \{u_1, \dots, u_k\} \subset \mathbf{S}^{d-1}$ be an $\mathbf{S}^{d-1}$ -bounding set. Then there exist real numbers $\lambda_1, \dots, \lambda_k > 0$ such that $\sum_{i=1}^k \lambda_i u_i = \mathbf{0}$ and $\sum_{i=1}^k \lambda_i = 1$.*

Proof. Suppose, for a contradiction, that $\mathbf{0}$ is not in the interior of $\text{conv}(\mathcal{T})$, the convex hull of $\mathcal{T}$. Since the interior of $\text{conv}(\mathcal{T})$ is also convex, it follows by the hyperplane separation theorem that there exists $v \in \mathbf{S}^{d-1}$ such that $\langle u, v \rangle \leq 0$ for all $u \in \mathcal{T}$, contradicting our assumption that $\mathcal{T}$ meets every open hemisphere of $\mathbf{S}^{d-1}$. It follows that $\mathbf{0}$ is in the interior of $\text{conv}(\mathcal{T})$, and the conclusion of the lemma follows easily. ■

For each finite $\mathbf{S}^{d-1}$ -bounding set $\mathcal{T} \subset \mathbf{S}^{d-1}$, we may use Lemma 6.1 to choose a set $\Xi_{\mathcal{T}} = \{\lambda_v : v \in \mathcal{T}\}$, with $\lambda_v > 0$ for all $v \in \mathcal{T}$, such that

$$\sum_{v \in \mathcal{T}} \lambda_v v = \mathbf{0} \quad \text{and} \quad \sum_{v \in \mathcal{T}} \lambda_v = 1. \quad (35)$$

We consider the sets $\Xi_{\mathcal{T}}$ to be fixed for the rest of the paper. Moreover, for each non-leaf vertex u of T , define

$$\Xi_u := \Xi_{M_T(u)},$$

where the set $M_T(u)$, defined in (27), is an $\mathbf{S}^{d-1}$ -bounding set by Lemma 4.12.

Definition 6.2. Let $\mathcal{T} \subset \mathbf{S}^{d-1}$ be a finite $\mathbf{S}^{d-1}$ -bounding set. For each bounded, non-empty set $\mathcal{X} \subset \mathbb{R}^d$, define the $\mathcal{T}$ -weighted diameter of $\mathcal{X}$ to be

$$\text{wd}_{\mathcal{T}}(\mathcal{X}) := \sum_{v \in \mathcal{T}} \lambda_v \cdot d_v(\mathcal{X}),$$

where $d_v(\mathcal{X}) := \sup_{y \in \mathcal{X}} \langle y, v \rangle$, and $\Xi_{\mathcal{T}} = \{\lambda_v : v \in \mathcal{T}\}$ is the set chosen in (35).

Moreover, for each non-leaf vertex u of T , define $\text{wd}_u := \text{wd}_{M_T(u)}$.

In order to motivate the definition above, let us note first that if $\mathcal{X} \subset \Lambda$ and D is the minimal $\mathcal{T}$ -droplet containing $\mathcal{X}$, then $\text{wd}_{\mathcal{T}}(D) = \text{wd}_{\mathcal{T}}(\mathcal{X})$. Indeed, D consists of exactly the points $y \in \Lambda$ such that $\langle y, v \rangle \leq d_v(\mathcal{X})$ for every $v \in \mathcal{T}$. Observe also that $\text{wd}_{\mathcal{T}}$ is increasing, in the sense that if $\mathcal{X} \subset \mathcal{Y}$ then $\text{wd}_{\mathcal{T}}(\mathcal{X}) \leq \text{wd}_{\mathcal{T}}(\mathcal{Y})$.

²⁵Unlike in the case $d = 2$, however, to do so we will require one additional technical lemma, whose proof is surprisingly challenging; see Lemma 6.17.

As further motivation, the following lemma shows that the $\mathcal{T}$ -weighted diameter is (as one would hope) invariant under translations. This property follows easily from (35).

Lemma 6.3. *Let $\mathcal{T} \subset \mathbf{S}^{d-1}$ be a finite $\mathbf{S}^{d-1}$ -bounding set, let $\mathcal{X} \subset \mathbb{R}^d$ be bounded and non-empty, and let $x \in \mathbb{R}^d$. Then $\text{wd}_{\mathcal{T}}(\mathcal{X}) = \text{wd}_{\mathcal{T}}(\mathcal{X} + x)$.*

Proof. Define $f(x) := \text{wd}_{\mathcal{T}}(\mathcal{X} + x)$, and observe that

$$f(x) = \sum_{v \in \mathcal{T}} \lambda_v (d_v(\mathcal{X}) + \langle x, v \rangle),$$

where $d_v(\mathcal{X}) = \sup_{y \in \mathcal{X}} \langle y, v \rangle$ and $\Xi_{\mathcal{T}} = \{\lambda_v : v \in \mathcal{T}\}$, as in Definition 6.2. Now, recalling that $\sum_{v \in \mathcal{T}} \lambda_v v = \mathbf{0}$, by (35), it follows that $f(x) = f(\mathbf{0})$, as claimed. ■

The next useful property of the $\mathcal{T}$ -weighted diameter is that $\text{wd}_u(\mathcal{X}) = \Theta(\text{diam}(\mathcal{X}))$ for every non-leaf vertex u of T , where the implicit constants depend only on T .

Lemma 6.4. *There exists $C > 0$ depending only on T such that the following holds. If u is a non-leaf vertex of T , then*

$$\frac{1}{C} \cdot \text{diam}(\mathcal{X}) \leq \text{wd}_u(\mathcal{X}) \leq \text{diam}(\mathcal{X})$$

for every bounded and non-empty set $\mathcal{X} \subset \mathbb{R}^d$.

Proof. By Lemma 6.3, we may assume without loss of generality that $\mathbf{0} \in \mathcal{X}$. It follows that $d_v(\mathcal{X}) = \sup_{y \in \mathcal{X}} \langle y, v \rangle \leq \text{diam}(\mathcal{X})$, and hence

$$\text{wd}_u(\mathcal{X}) = \sum_{v \in M_T(u)} \lambda_v \cdot d_v(\mathcal{X}) \leq \left(\sum_{v \in M_T(u)} \lambda_v \right) \cdot \text{diam}(\mathcal{X}) = \text{diam}(\mathcal{X}),$$

since $\lambda_v > 0$ for all $v \in M_T(u)$ and $\sum_{v \in M_T(u)} \lambda_v = 1$, by (35).

To prove the lower bound on $\text{wd}_u(\mathcal{X})$, observe first that

$$\text{wd}_u(\mathcal{X}) = \sum_{v \in M_T(u)} \lambda_v \cdot d_v(\mathcal{X}) \geq \Delta_u(\mathcal{X}) \cdot \min \Xi_u, \quad (36)$$

where $\Delta_u(\mathcal{X}) = \max \{d_v(\mathcal{X}) : v \in M_T(u)\}$, and that

$$\mathcal{X} \subset \bigcap_{v \in M_T(u)} \{x \in \mathbb{R}^d : \langle x, v \rangle \leq \Delta_u(\mathcal{X})\}. \quad (37)$$

Note that $\Delta_u(\mathcal{X}) \geq 0$, since $\mathbf{0} \in \mathcal{X}$, and set

$$F_u := \bigcap_{v \in M_T(u)} \{x \in \mathbb{R}^d : \langle x, v \rangle \leq 1\}.$$

It follows from (36) and (37) that

$$\text{diam}(\mathcal{X}) \leq \Delta_u(\mathcal{X}) \cdot \text{diam}(F_u) \leq \frac{\text{diam}(F_u)}{\min \Xi_u} \cdot \text{wd}_u(\mathcal{X}) \leq C \cdot \text{wd}_u(\mathcal{X}),$$

where

$$C := \max_{u \in T} \frac{\text{diam}(F_u)}{\min \Xi_u}$$

is finite, by Lemma 4.12, and depends only on T , as required. $\blacksquare$

Let us note here, for future reference, the following immediate corollary of Lemma 6.4.

Lemma 6.5. *There exists $C > 0$ depending only on T such that the following holds. If u and v are non-leaf vertices of T , then*

$$\text{wd}_u(\mathcal{X}) \leq C \cdot \text{wd}_v(\mathcal{X})$$

for every bounded and non-empty set $\mathcal{X} \subset \mathbb{R}^d$.

Proof. By two applications of Lemma 6.4, we have

$$\text{wd}_u(\mathcal{X}) \leq \text{diam}(\mathcal{X}) \leq C \cdot \text{wd}_v(\mathcal{X}). \quad \blacksquare$$

In the proofs of the subadditivity lemmas it will be convenient to work with the following continuous analogue of droplets, which we shall refer to as ‘regions’.

Definition 6.6. Let $\mathcal{T} \subset \mathbb{S}^{d-1}$ be a finite and non-empty set. A $\mathcal{T}$ -region is a non-empty set of the form

$$\bigcap_{u \in \mathcal{T}} \{x \in \mathbb{R}^d : \langle x, u \rangle \leq a_u\}$$

for some collection $\{a_u \in \mathbb{R} : u \in \mathcal{T}\}$.

Now, given a non-leaf vertex u of T and a bounded, non-empty set $\mathcal{X} \subset \mathbb{R}^d$, define

- (a) $F_u(\mathcal{X})$ to be the minimal $M_T(u)$ -region containing $\mathcal{X}$;
- (b) $G_0(\mathcal{X}) := F_0(\mathcal{X})$, and if u is not the root then

$$G_u(\mathcal{X}) := \{x \in F_u^*(\mathcal{X}) : \langle x, u \rangle \geq 0\},$$

where $F_u^*(\mathcal{X})$ is the minimal $N_T^{\rightarrow}(u)$ -region containing $\mathcal{X}$.

Note that if $\{x \in \mathcal{X} : \langle x, u \rangle \geq 0\}$ is non-empty then both $F_u(\mathcal{X})$ and $G_u(\mathcal{X})$ are bounded $M_T(u)$ -regions. Observe also that if $\mathcal{X} \cap H_T(u) = \emptyset$ and J is the minimal u -iceberg droplet containing $\mathcal{X}$, then by Definition 5.2 we have $J \subset G_u(\mathcal{X})$.

The following is our first subadditivity lemma for d -dimensional droplets. Recalling from (34) that the corresponding statement for $u = \mathbf{0}$ with wd_0 replaced by diam is false in general, the lemma provides further motivation for Definition 6.2.

Lemma 6.7. *Let u be a non-leaf vertex of T , and let $\mathcal{X}_1, \mathcal{X}_2 \subset \mathbb{R}^d$ be bounded sets such that $\mathcal{X}_1 \cap \mathcal{X}_2 \neq \emptyset$. Then*

$$\text{wd}_u(F_u(\mathcal{X}_1 \cup \mathcal{X}_2)) \leq \text{wd}_u(\mathcal{X}_1) + \text{wd}_u(\mathcal{X}_2).$$




Fig. 2. A u_1 -iceberg droplet J_1 and a $\mathbf{0}$ -iceberg droplet J_2 . Here, $\mathbf{0} \rightarrow u_1 \rightarrow u_2$ is a path in T . Also shown are the regions $G_{u_1}(J_1 \cup J_2)$ and $G_{u_2}(J_1 \cup J_2)$ (the latter being clipped at its extremities for space).

Proof. By Lemma 6.3, we may assume that $\mathbf{0} \in \mathcal{X}_1 \cap \mathcal{X}_2$, and hence $d_v(\mathcal{X}_1) \geq 0$ and $d_v(\mathcal{X}_2) \geq 0$ for each $v \in M_T(u)$. Let $\mathcal{X} := F_u(\mathcal{X}_1 \cup \mathcal{X}_2)$, and observe that

$$d_v(\mathcal{X}) = \sup_{x \in \mathcal{X}} \langle x, v \rangle = \max \{d_v(\mathcal{X}_1), d_v(\mathcal{X}_2)\} \leq d_v(\mathcal{X}_1) + d_v(\mathcal{X}_2)$$

for each $v \in M_T(u)$. By Definition 6.2, it follows that

$$\text{wd}_u(\mathcal{X}) \leq \sum_{v \in M_T(u)} \lambda_v \cdot (d_v(\mathcal{X}_1) + d_v(\mathcal{X}_2)) = \text{wd}_u(\mathcal{X}_1) + \text{wd}_u(\mathcal{X}_2),$$

as required. ■

We shall use Lemma 6.7 in the (significantly more difficult) proof of our main subadditivity lemma for icebergs, Lemma 6.14. We shall also use the following bound on the $M_T(u)$ -weighted diameter of $F_v(\mathcal{X})$, which is an easy consequence of Lemma 6.5.

Lemma 6.8. *There exists $C > 0$ depending only on T such that the following holds. If u and v are non-leaf vertices of T , then*

$$\text{wd}_u(F_v(\mathcal{X})) \leq C \cdot \text{wd}_u(\mathcal{X})$$

for every bounded and non-empty set $\mathcal{X} \subset \mathbb{R}^d$.

Proof. Note that $F_v(\mathcal{X})$ is bounded, since $M_T(v)$ is $\mathbf{S}^{d-1}$ -bounding, by Lemma 4.12. Therefore, by Lemma 6.5, we have

$$\text{wd}_u(F_v(\mathcal{X})) \leq C \cdot \text{wd}_v(F_v(\mathcal{X})).$$

Now, by Definitions 6.2 and 6.6, we have $\text{wd}_v(F_v(\mathcal{X})) = \text{wd}_v(\mathcal{X})$, since $F_v(\mathcal{X})$ consists of exactly the points $y \in \mathbb{R}^d$ such that $\langle y, v \rangle \leq d_v(\mathcal{X})$ for every $v \in M_T(v)$. Hence, applying Lemma 6.5 again, we obtain

$$\text{wd}_v(F_v(\mathcal{X})) = \text{wd}_v(\mathcal{X}) \leq C \cdot \text{wd}_u(\mathcal{X}).$$

■

The following property of the function G_u , which follows from Lemma 5.4, will also be useful in the proof of Lemma 6.14.

Lemma 6.9. *Let u be a non-leaf vertex of T , and let $K \subset \Lambda \setminus H_T(u)$ be a finite set. Then $G_u([K]_u) = G_u(K)$, and in particular $[K]_u \subset G_u(K)$.*

Proof. By Definition 6.6, we are required to show that

$$\max \{ \langle x, v \rangle : x \in [K]_u \} = \max \{ \langle x, v \rangle : x \in K \} \quad (38)$$

for each $v \in N_T^{\rightarrow}(u)$. To do so, let J be the minimal u -iceberg droplet containing K , and (recalling Definition 5.2) observe that

$$\max \{ \langle x, v \rangle : x \in J \} = \max \{ \langle x, v \rangle : x \in K \}$$

for every $v \in N_T^{\rightarrow}(u)$. Now, since J is u -closed, by Lemma 5.4, and $K \subset J$, we have

$$\max \{ \langle x, v \rangle : x \in [K]_u \} \leq \max \{ \langle x, v \rangle : x \in [J]_u \} = \max \{ \langle x, v \rangle : x \in J \},$$

and (38) follows, since $K \subset [K]_u$. This proves that $G_u([K]_u) = G_u(K)$, and it follows that $[K]_u \subset G_u(K)$, since $\langle x, u \rangle \geq 0$ for every $x \in [K]_u$. ■

We need two more technical lemmas for the proof of Lemma 6.14. The following preliminary inequality will be used in the proof of each of them. We shall write $O(1)$ to denote a function that is bounded by a constant depending on both T and R .

Lemma 6.10. *Let u be a non-root, non-leaf vertex of T , let $\mathcal{X} \subset \mathbb{R}^d$ be a bounded and non-empty set, and suppose that $\mathcal{X}$ is strongly connected to $\mathbf{H}_u^d$. Then*

$$\text{wd}_u(G_u(\mathcal{X})) \leq \text{wd}_u(\mathcal{X}) + O(1).$$

Proof. Recall from (27) that $-u \in M_T(u)$. By Definition 6.6, and since $\mathcal{X}$ is strongly connected to $\mathbf{H}_u^d$, it follows that

$$d_{-u}(F_u(\mathcal{X})) = d_{-u}(\mathcal{X}) \geq -R,$$

where $d_v(\mathcal{X}) = \sup_{y \in \mathcal{X}} \langle y, v \rangle$, as in Definition 6.2. Since $d_v(G_u(\mathcal{X})) \leq d_v(F_u(\mathcal{X}))$ for all $v \in N_T^{\rightarrow}(u)$, also by Definition 6.6, it follows from Definition 6.2 that

$$\text{wd}_u(G_u(\mathcal{X})) \leq \text{wd}_u(F_u(\mathcal{X})) + O(1).$$

But $\text{wd}_u(F_u(\mathcal{X})) = \text{wd}_u(\mathcal{X})$, so this completes the proof. ■

The first of our two technical lemmas allows us to control the change in $\text{wd}_v(G_v(K))$ as v varies along a (K, u) -permissible sequence.

Lemma 6.11. *There exists $C > 0$ depending only on T such that the following holds. Let $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k$ be a path in T , with u_k not a leaf, let $K \subset \Lambda \setminus H_T(u_k)$ be a finite, non-empty set, and let $0 = i_0 < \dots < i_\ell \leq k$ be a (K, u_k) -permissible sequence. Then*

$$\text{wd}_{u_{i_j}}(G_{i_j}) \leq C \cdot \text{wd}_{u_{i_{j+1}}}(G_{i_{j+1}}) + O(1) \quad (39)$$

and

$$\text{wd}_{u_{i_j+1}}(G_{i_j+1}) \leq C \cdot \text{wd}_{u_{i_j}}(G_{i_j}) + O(1) \quad (40)$$

for each $0 \leq j < \ell$, where $G_i := G_{u_i}(K)$ for each $0 \leq i \leq k$.

Proof. Starting with (40), the key fact we need is that J_{i_j} , the minimal u_{i_j} -iceberg droplet containing K , is strongly connected to $\mathbf{H}_{u_{i_j+1}}^d$, which follows from Definition 5.5 since $0 = i_0 < \dots < i_\ell \leq k$ is a (K, u_k) -permissible sequence. Since $J_{i_j} \subset G_{i_j}$, by Definitions 5.2 and 6.6, this means that G_{i_j} is also strongly connected to $\mathbf{H}_{u_{i_j+1}}^d$.

By Lemma 6.10 applied with $\mathcal{X} = G_{i_j}$, it follows that

$$\text{wd}_{u_{i_j+1}}(G_{u_{i_j+1}}(G_{i_j})) \leq \text{wd}_{u_{i_j+1}}(G_{i_j}) + O(1). \quad (41)$$

Now, note that $K \subset J_{i_j} \subset G_{i_j}$, and so

$$\begin{aligned} \text{wd}_{u_{i_j+1}}(G_{i_j+1}) &= \text{wd}_{u_{i_j+1}}(G_{u_{i_j+1}}(K)) \\ &\leq \text{wd}_{u_{i_j+1}}(G_{u_{i_j+1}}(G_{i_j})). \end{aligned} \quad (42)$$

Moreover, by Lemma 6.5 we have

$$\text{wd}_{u_{i_j+1}}(G_{i_j}) \leq C \cdot \text{wd}_{u_{i_j}}(G_{i_j}) \quad (43)$$

for some $C > 0$ depending only on T . Combining (41)–(43) gives (40).

To prove (39), we shall compare $\text{wd}_{u_{i_j+1}}(G_{i_j+1})$ to $\text{wd}_{u_0}(G_0)$, and then use the fact that we have already proved (40) for all $0 \leq j < \ell$. To be precise, observe first that

$$\text{wd}_{u_{i_j+1}}(G_{i_j+1}) \geq \text{wd}_{u_{i_j+1}}(K) \geq \frac{1}{C'} \cdot \text{wd}_0(K) = \frac{1}{C'} \cdot \text{wd}_0(G_0) \quad (44)$$

for some $C' > 0$ depending only on T . Indeed, the first inequality holds because (as noted above) $K \subset G_{i_j+1}$, and the second inequality holds by Lemma 6.5. For the final step, recall from Definition 6.6 that $G_0 = G_0(K)$ is the minimal $N_T^{\rightarrow}(\mathbf{0})$ -region containing K , and that $M_T(\mathbf{0}) = N_T^{\rightarrow}(\mathbf{0})$, so the equality holds by Definition 6.2. Now, by (40) (applied j times), we have

$$\text{wd}_{u_{i_j}}(G_{i_j}) \leq C^j \cdot \text{wd}_0(G_0) + O(1), \quad (45)$$

and combining (44) and (45) gives (39), as required. $\blacksquare$

The final lemma that we need for the proof of Lemma 6.14 allows us to bound the ‘size’ of the u_k -closure of a set in terms of the ‘size’ of its u_i -closure.

Lemma 6.12. *There exists $C > 0$ depending only on T such that the following holds. Let $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k$ be a path in T , with u_k not a leaf, and let $K \subset \Lambda \setminus H_T(u_k)$ be a finite, non-empty set. Then*

$$\text{wd}_{u_i}(G_{u_i}([K]_{u_k})) \leq C \cdot \text{wd}_{u_i}(G_{u_i}([K]_{u_i})) + O(1)$$

for all $0 \leq i \leq k$.

Proof. We shall in fact prove that

$$\text{wd}_{u_i}(G_{u_i}([K]_{u_{j+1}})) \leq C \cdot \text{wd}_{u_i}(G_{u_i}([K]_{u_j})) + O(1) \quad (46)$$

for each $i \leq j < k$, which is clearly sufficient to prove the lemma.

If $[K]_{u_{j+1}} = [K]_{u_j}$ then there is nothing to prove, so assume that $[K]_{u_{j+1}} \neq [K]_{u_j}$. We claim that in this case $[K]_{u_j}$ is strongly connected to $\mathbf{H}_{u_{j+1}}^d$. To see this, note first that, by Definition 5.1, if $[K]_{u_j}$ does not intersect $\mathbf{H}_{u_{j+1}}^d$, then $[K]_{u_j} \subset [K]_{u_{j+1}}$. It follows that if $[K]_{u_{j+1}} \neq [K]_{u_j}$, then $x + X \subset [K]_{u_j} \cup H_T(u_{j+1})$ for some $x \notin [K]_{u_j} \cup H_T(u_{j+1})$ and $X \in \mathcal{U}$. Now, the sets $H_T(u_{j+1})$ and $[K]_{u_j} \cup H_T(u_j)$ are $\mathcal{U}$ -closed, by Proposition 4.10 and Definition 5.1, so $x + X$ must intersect both $[K]_{u_j}$ and $\mathbf{H}_{u_{j+1}}^d$. By Definition 3.3, this implies that $[K]_{u_j}$ is strongly connected to $\mathbf{H}_{u_{j+1}}^d$, as claimed.

Now, observe that $K \subset G_{u_i}([K]_{u_j})$, and therefore $[K]_{u_{j+1}} \subset G_{u_{j+1}}(G_{u_i}([K]_{u_j}))$, by Lemma 6.9, so

$$G_{u_i}([K]_{u_{j+1}}) \subset G_{u_i}(G_{u_{j+1}}(G_{u_i}([K]_{u_j}))) =: \mathcal{X}. \quad (47)$$

In order to bound $\text{wd}_{u_i}(\mathcal{X})$, we claim that, for each $v \in N_T(u_i)$, we have

$$d_v(\mathcal{X}) \leq C \cdot \text{wd}_{u_i}(G_{u_i}([K]_{u_j})) + O(1). \quad (48)$$

To prove (48), let us write $\mathcal{Y} := G_{u_{j+1}}(G_{u_i}([K]_{u_j}))$, so that $\mathcal{X} = G_{u_i}(\mathcal{Y})$. Now, observe that $d_v(\mathcal{X}) \leq d_v(\mathcal{Y})$, since $\mathcal{X}$ is contained in the minimal $N_T^{\rightarrow}(u_i)$ -region containing $\mathcal{Y}$, by Definition 6.6. It follows that

$$d_v(\mathcal{X}) \leq d_v(\mathcal{Y}) \leq C \cdot \text{wd}_{u_i}(\mathcal{Y}) \leq C^2 \cdot \text{wd}_{u_{j+1}}(\mathcal{Y}), \quad (49)$$

by Definition 6.2 and Lemma 6.5, for some C depending only on T .

Now, since $[K]_{u_j}$ is strongly connected to $\mathbf{H}_{u_{j+1}}^d$ and $i \leq j$, so $\langle x, u_i \rangle \geq 0$ for every $x \in [K]_{u_j}$, it follows that $G_{u_i}([K]_{u_j})$ is also strongly connected to $\mathbf{H}_{u_{j+1}}^d$, and thus

$$\text{wd}_{u_{j+1}}(G_{u_{j+1}}(G_{u_i}([K]_{u_j}))) \leq \text{wd}_{u_{j+1}}(G_{u_i}([K]_{u_j})) + O(1),$$

by Lemma 6.10. Hence, by Lemma 6.5, we obtain

$$\text{wd}_{u_{j+1}}(\mathcal{Y}) = \text{wd}_{u_{j+1}}(G_{u_{j+1}}(G_{u_i}([K]_{u_j}))) \leq C \cdot \text{wd}_{u_i}(G_{u_i}([K]_{u_j})) + O(1),$$

which, together with (49), proves (48).

If $u_i = \mathbf{0}$ then (46) now follows from (47) and (48). On the other hand, if $u_i \neq \mathbf{0}$ then observe that

$$d_{-u_i}(\mathcal{X}) = d_{-u_i}(G_{u_i}(\mathcal{Y})) \leq 0,$$

since $\langle x, u_i \rangle \geq 0$ for every $x \in G_{u_i}(\mathcal{Y})$, by Definition 6.6. Together with (47) and (48), this implies (46), and hence completes the proof of the lemma. $\blacksquare$

We are almost ready to state our subadditivity lemma for icebergs. In order to do so, we need one more important definition. Let $\delta > 0$ be a sufficiently small constant, depending on T , and fixed for the rest of the paper.

Definition 6.13. Let $\mathcal{J} = (J, u)$ be an iceberg and let $k := d_T(\mathbf{0}, u)$. Then we define $\text{diam}^*(\mathcal{J})$, the *iceberg diameter* of $\mathcal{J}$, as follows:

$$\text{diam}^*(\mathcal{J}) := \delta^k \text{wd}_u(G_u(J)).$$

The next lemma is our subadditivity lemma for icebergs. The lemma will be used in the proofs of both the Aizenman–Lebowitz lemma for icebergs (Lemma 6.18) and the extremal lemma for icebergs (Lemma 6.19), and also in the proof of Lemma 7.2.

Lemma 6.14. *Let u be a non-leaf vertex of T , and let K_1 and K_2 be finite non-empty subsets of $\Lambda \setminus H_T(u)$ such that $[K_1]_u$, $[K_2]_u$ and $[K_1]_u \cup [K_2]_u$ are all strongly connected. Then*

$$\text{diam}^*(\mathcal{J}_u([K_1 \cup K_2]_u)) \leq \text{diam}^*(\mathcal{J}_u([K_1]_u)) + \text{diam}^*(\mathcal{J}_u([K_2]_u)) + O(1). \quad (50)$$

Proof. Let $\mathbf{0} \rightarrow u_1 \rightarrow \cdots \rightarrow u_k = u$ be the path from $\mathbf{0}$ to u in T . Let

$$\mathcal{J}_1 := \mathcal{J}_u([K_1]_u) = (J_1, u_{i_1}) \quad \text{and} \quad \mathcal{J}_2 := \mathcal{J}_u([K_2]_u) = (J_2, u_{i_2}) \quad (51)$$

for some $0 \leq i_1, i_2 \leq k$, and assume without loss of generality that $i_1 \geq i_2$. Now let

$$\mathcal{J} = \mathcal{J}_u([K_1 \cup K_2]_u) = (J, u_\ell),$$

and observe, by Lemma 5.7, that $\ell \geq \max\{i_1, i_2\}$, since $[K_1]_u \cup [K_2]_u \subset [K_1 \cup K_2]_u$. We claim that we may in fact assume that $\ell = k$. Indeed, this follows from Lemma 5.8, which implies that

$$\mathcal{J}_{u_\ell}([K_1 \cup K_2]_{u_\ell}) = \mathcal{J}_{u_k}([K_1 \cup K_2]_{u_k}),$$

since $\mathcal{J}$ is a u_ℓ -iceberg, and similarly

$$\mathcal{J}_{u_\ell}([K_j]_{u_\ell}) = \mathcal{J}_{u_k}([K_j]_{u_k}) \quad \text{and} \quad [K_j]_{u_\ell} = [K_j]_{u_k}$$

for each $j \in \{1, 2\}$, so $[K_1]_{u_\ell}$, $[K_2]_{u_\ell}$ and $[K_1]_{u_\ell} \cup [K_2]_{u_\ell}$ are all strongly connected.

Recalling Definition 6.13, it follows that (50) is equivalent to

$$\delta^k \text{wd}_{u_k}(G_{u_k}(J)) \leq \delta^{i_1} \text{wd}_{u_{i_1}}(G_{u_{i_1}}(J_1)) + \delta^{i_2} \text{wd}_{u_{i_2}}(G_{u_{i_2}}(J_2)) + O(1). \quad (52)$$

We divide the proof of (52) into two claims, using the set $[K_1 \cup K_2]_{u_{i_1}}$ as a stepping stone between the sets $[K_1]_{u_{i_1}}$ and $[K_2]_{u_{i_2}}$, and the set $[K_1 \cup K_2]_{u_k}$.

Claim 6.15.

$$\delta^{i_1} \text{wd}_{u_{i_1}}(G_{u_{i_1}}([K_1 \cup K_2]_{u_{i_1}})) \leq \delta^{i_1} \text{wd}_{u_{i_1}}(G_{u_{i_1}}(J_1)) + \delta^{i_2} \text{wd}_{u_{i_2}}(G_{u_{i_2}}(J_2)) + O(1).$$

Claim 6.16.

$$\delta^k \text{wd}_{u_k}(G_{u_k}(J)) \leq \delta^{i_1} \text{wd}_{u_{i_1}}(G_{u_{i_1}}([K_1 \cup K_2]_{u_{i_1}})) + O(1).$$

Note that (52) follows immediately from Claims 6.15 and 6.16, and therefore, by the observations above, proving these two claims is sufficient to prove Lemma 6.14.

Proof of Claim 6.15. Recalling Definition 6.6, set

$$G_1 := G_{u_{i_1}}(J_1), \quad G_2 := G_{u_{i_2}}(J_2), \quad \text{and} \quad G'_2 := F_{u_{i_1}}(G_2) = F_{u_{i_1}}(G_{u_{i_2}}(J_2)).$$

Our first subclaim is that

$$\delta^{i_2} \text{wd}_{u_{i_2}}(G_2) \geq \delta^{i_1} \text{wd}_{u_{i_1}}(G'_2). \quad (53)$$

If $i_1 = i_2$ then $G'_2 = G_2$ because in this case G_2 is already an $M_T(u_{i_1})$ -region, and (53) is immediate. On the other hand, if $i_1 > i_2$ then we have

$$\delta^{i_2} \text{wd}_{u_{i_2}}(G_2) \geq \delta^{i_2+1/2} \text{wd}_{u_{i_2}}(G'_2) \geq \delta^{i_2+1} \text{wd}_{u_{i_1}}(G'_2) \geq \delta^{i_1} \text{wd}_{u_{i_1}}(G'_2),$$

where the first inequality follows from Lemma 6.8 and the second from Lemma 6.5, and since δ is sufficiently small (depending on T). Thus in either case (53) holds, and hence

$$\delta^{i_1} \text{wd}_{u_{i_1}}(G_1) + \delta^{i_2} \text{wd}_{u_{i_2}}(G_2) \geq \delta^{i_1} (\text{wd}_{u_{i_1}}(G_1) + \text{wd}_{u_{i_1}}(G'_2)). \quad (54)$$

Our second subclaim is that

$$\text{wd}_{u_{i_1}}(G_1) + \text{wd}_{u_{i_1}}(G'_2) \geq \text{wd}_{u_{i_1}}(F_{u_{i_1}}(G_1 \cup G'_2)) - O(1). \quad (55)$$

This will follow from two applications of Lemma 6.7, and our assumption that the sets $[K_1]_u$ and $[K_2]_u$ are strongly connected to each other. Observe first that $[K_1]_u \subset J_1 \subset G_1$, since J_1 is, by (51) and Definition 5.5, the minimal u_{i_1} -iceberg droplet containing $[K_1]_u$, and therefore does not intersect $\mathbb{H}_{u_{i_1}}^d$. Similarly, we have $[K_2]_u \subset J_2 \subset G_2 \subset G'_2$, and so G_1 and G'_2 are strongly connected to each other, since $[K_1]_u$ and $[K_2]_u$ are strongly connected to each other by assumption. Let B be a set of diameter at most R such that $G_1 \cap B \neq \emptyset$ and $G'_2 \cap B \neq \emptyset$. Applying Lemma 6.7 to the sets G_1 and B gives

$$\text{wd}_{u_{i_1}}(F_{u_{i_1}}(G_1 \cup B)) \leq \text{wd}_{u_{i_1}}(G_1) + \text{wd}_{u_{i_1}}(B),$$

and applying Lemma 6.7 to the sets $F_{u_{i_1}}(G_1 \cup B)$ and G'_2 gives

$$\text{wd}_{u_{i_1}}(F_{u_{i_1}}(G_1 \cup G'_2)) \leq \text{wd}_{u_{i_1}}(F_{u_{i_1}}(G_1 \cup B)) + \text{wd}_{u_{i_1}}(G'_2),$$

since $G_1 \cup G'_2 \subset F_{u_{i_1}}(G_1 \cup B) \cup G'_2$. Combining these two inequalities, and noting that $\text{wd}_{u_{i_1}}(B) \leq R$, by Lemma 6.4, we obtain the inequality (55).

Our third subclaim, which will complete the proof of Claim 6.15, is that

$$G_{u_{i_1}}([K_1 \cup K_2]_{u_{i_1}}) \subset F_{u_{i_1}}(G_1 \cup G'_2), \quad (56)$$

and hence

$$\text{wd}_{u_{i_1}}(G_{u_{i_1}}([K_1 \cup K_2]_{u_{i_1}})) \leq \text{wd}_{u_{i_1}}(F_{u_{i_1}}(G_1 \cup G'_2)). \quad (57)$$

To prove (56), first we show that

$$G_{u_{i_1}}(G_1 \cup G'_2) \subset F_{u_{i_1}}(G_1 \cup G'_2). \quad (58)$$

If $u_{i_1} = \mathbf{0}$, then $F_{u_{i_1}} = G_{u_{i_1}}$ and (58) holds trivially, so suppose that u_{i_1} is not the root. Now, recalling Definition 6.6, observe that $G_1 = G_{u_{i_1}}(J_1)$ is bounded and $F_{u_{i_1}}^*(J_1)$ is unbounded, by Lemma 4.12, and therefore

$$G_1 \cap \{x \in \mathbb{R}^d : \langle x, u_{i_1} \rangle = 0\} \neq \emptyset.$$

By Definition 6.6, this is enough to prove (58).

To deduce (56) from (58), it suffices to observe that

$$G_{u_{i_1}}([K_1 \cup K_2]_{u_{i_1}}) = G_{u_{i_1}}(K_1 \cup K_2) \subset G_{u_{i_1}}(G_1 \cup G_2'),$$

where the first step follows from Lemma 6.9, and the second step follows since (as we observed in the proof of the second subclaim) we have $K_1 \subset G_1$ and $K_2 \subset G_2'$.

Claim 6.15 now follows by combining (54), (55) and (57). ■

Claim 6.16 is a relatively straightforward application of Lemmas 6.11 and 6.12.

Proof of Claim 6.16. Recall that J is, by Definition 5.5, the minimal u_k -iceberg droplet containing $[K_1 \cup K_2]_{u_k}$, and by Definition 5.2 is therefore contained in the minimal $N_T^{\rightarrow}(u_k)$ -droplet containing $[K_1 \cup K_2]_{u_k}$. By Definition 6.6, it follows that $G_{u_k}(J) \subset G_{u_k}([K_1 \cup K_2]_{u_k})$, and hence

$$\text{wd}_{u_k}(G_{u_k}(J)) \leq \text{wd}_{u_k}(G_{u_k}([K_1 \cup K_2]_{u_k})). \quad (59)$$

This proves the claim if $i_1 = k$, so let us assume that $i_1 < k$. Observe that

$$\text{wd}_{u_{i_1}}(G_{u_{i_1}}([K_1 \cup K_2]_{u_k})) \leq C \cdot \text{wd}_{u_{i_1}}(G_{u_{i_1}}([K_1 \cup K_2]_{u_{i_1}})) + O(1), \quad (60)$$

for some C depending only on T , by Lemma 6.12. Now define

$$G^{(j)} := G_{u_j}([K_1 \cup K_2]_{u_k})$$

for each $0 \leq j \leq k$. Since we chose δ sufficiently small depending on T , and since $i_1 < k$, by (59) and (60) it will suffice to show that

$$\delta^{1/2} \text{wd}_{u_k}(G^{(k)}) \leq \text{wd}_{u_{i_1}}(G^{(i_1)}) + O(1). \quad (61)$$

To prove (61) we shall apply Lemma 6.11 along two $([K_1 \cup K_2]_{u_k}, u_k)$ -permissible sequences leading to i_1 and k respectively.

To be precise, by Definition 5.5 and since $\mathcal{J}_{u_k}([K_1 \cup K_2]_{u_k}) = (J, u_k)$, there exists a $([K_1 \cup K_2]_{u_k}, u_k)$ -permissible sequence containing k . Moreover, since $\mathcal{J}_{u_k}([K_1]_{u_k}) = (J_1, u_{i_1})$, there exists a $([K_1]_{u_k}, u_k)$ -permissible sequence containing i_1 . Now, recall that if $K \subset K'$ then any (K, u_k) -permissible sequence is also (K', u_k) -permissible (see the proof of Lemma 5.7), so there also exists a $([K_1 \cup K_2]_{u_k}, u_k)$ -permissible sequence containing i_1 . By Lemma 6.11, it follows that

$$\text{wd}_{u_k}(G^{(k)}) \leq C^k \text{wd}_0(G^{(0)}) + O(1) \leq C^{k+i_1} \text{wd}_{u_{i_1}}(G^{(i_1)}) + O(1).$$

Since we chose δ sufficiently small depending on T , this proves (61), and hence, by (59) and (60), this suffices to prove Claim 6.16. ■

As observed earlier, Claims 6.15 and 6.16 together imply (52), and thus together they complete the proof of Lemma 6.14. ■

We shall next deduce our Aizenman–Lebowitz-type and extremal lemmas for icebergs from the u -iceberg spanning algorithm and Lemma 6.14. In order to do so, we need the following lemma, which provides a uniform upper bound on the size of the u -iceberg container of the u -closure of a *single* vertex. The proof of the lemma is (perhaps surprisingly) non-trivial, and relies on Lemma 6.11.

Lemma 6.17. *There exists a constant $C > 0$ such that the following holds. If u is a non-leaf vertex of T and $x \in \Lambda \setminus H_T(u)$, then $\text{diam}^*(\mathcal{J}_u([\{x\}]_u)) \leq C$.*

Proof. Set $K := [\{x\}]_u$, and note that $K \subset \Lambda \setminus H_T(u)$, by Definition 5.1, and that K is non-empty and finite, by Lemmas 5.3 and 5.4, since it is contained in the minimal u -iceberg droplet containing x . Let $\mathcal{J}_u(K) = (J, v)$ so, by Definition 6.13, we are required to bound $\text{wd}_v(G_v(J))$.

Observe that $G_v(J) = G_v(K)$, since J is the minimal v -iceberg droplet containing K . Now, since by Definition 5.5 there exists a (K, u) -permissible sequence ending in v , we have

$$\text{wd}_v(G_v(J)) = \text{wd}_v(G_v(K)) \leq C^k \text{wd}_0(G_0(K)) + O(1),$$

by Lemma 6.11. Since $G_0(K) = F_0(K)$ and $\text{wd}_0(F_0(K)) = \text{wd}_0(K)$, the lemma will hold if we can prove a uniform bound on the $\mathbf{0}$ -weighted diameter of $K = [\{x\}]_u$.

To bound $\text{wd}_0(K)$, let $\mathcal{J}_u(\{x\}) = (J', w)$ be the u -iceberg container of $\{x\}$, and observe that J' is u -closed, by Lemma 5.6, and therefore that $K \subset J'$. Note also that $J' \subset G_w(J') = G_w(\{x\})$, since J' is the minimal w -iceberg droplet containing x . Now, since there exists an $(\{x\}, u)$ -permissible sequence ending in w , we have

$$\text{wd}_w(K) \leq \text{wd}_w(G_w(\{x\})) \leq C^k \text{wd}_0(G_0(\{x\})) + O(1),$$

by Lemma 6.11. Noting that $G_0(\{x\}) = F_0(\{x\})$ and $\text{wd}_0(F_0(\{x\})) = 0$, since $\{x\}$ is an $N_{\vec{T}}(\mathbf{0})$ -region, it follows, using Lemma 6.5, that

$$\text{wd}_0(K) \leq C \cdot \text{wd}_w(K) = O(1),$$

as required. ■

We can now prove our Aizenman–Lebowitz-type lemma for icebergs. We remark that, while the lemma is a straightforward consequence of the u -iceberg spanning algorithm and Lemmas 5.11, 6.14 and 6.17, the statement is slightly more complicated than that of the corresponding lemma in [17] (which was about $\mathcal{T}$ -internally spanned $\mathcal{T}$ -droplets). This is because, while an iceberg being u -iceberg spanned implies the existence of similar such icebergs at smaller scales, it does *not* imply that the corresponding iceberg droplets are nested if the icebergs are of different types. However, the lemma gives us control over the location of the elements of A that are used to span our icebergs, and this will suffice for our purposes. Let $\lambda > 0$ be a sufficiently large constant (in particular, larger than the bound in Lemma 6.17, and the constant implicit in the $O(1)$ -term in Lemma 6.14).

Lemma 6.18 (Aizenman–Lebowitz lemma for icebergs). *Let u be a non-leaf vertex of T , and let $K \subset \Lambda \setminus H_T(u)$ be a finite set. Suppose that $\langle K \rangle_u = \{\mathcal{J}\}$ for some iceberg $\mathcal{J}$, and let $\lambda \leq m \leq \text{diam}^*(\mathcal{J})$. Then there exists a set $K' \subset K$ and an iceberg $\mathcal{J}'$ with*

$$m \leq \text{diam}^*(\mathcal{J}') \leq 3m$$

such that $\langle K' \rangle_u = \{\mathcal{J}'\}$.

Proof. Apply the u -iceberg spanning algorithm (Definition 5.10) to K . We claim that

$$f(t) := \max \{ \text{diam}^*(\mathcal{J}_u([L]_u)) : L \in \mathcal{K}^t \}$$

satisfies $f(t) \leq 2f(t-1) + O(1)$ for every $t \in [T]$, and hence that $f(t) \leq 3f(1)$ if $f(1) \geq \lambda$, by our choice of λ . To see this, suppose that the sets K_1 and K_2 are united in step t of the algorithm, and observe that $K_1 \cup K_2 \subset K \subset \Lambda \setminus H_T(u)$, and that the sets $[K_1]_u$, $[K_2]_u$ and $[K_1]_u \cup [K_2]_u$ are all strongly connected (see the proof of Lemma 5.11). It follows, by Lemma 6.14, that

$$\text{diam}^*(\mathcal{J}_u([K_1 \cup K_2]_u)) \leq \text{diam}^*(\mathcal{J}_u([K_1]_u)) + \text{diam}^*(\mathcal{J}_u([K_2]_u)) + O(1),$$

so $f(t) \leq 2f(t-1) + O(1)$, as claimed.

Now, recall that $\mathcal{K}^0 = \{\{x\} : x \in K\}$ and observe that, by Lemma 6.17 and our choice of λ , we have $\text{diam}^*(\mathcal{J}_u([\{x\}]_u)) < \lambda$ for every $x \in K$. Since $\langle K \rangle_u = \{\mathcal{J}\}$ and we have $\lambda \leq m \leq \text{diam}^*(\mathcal{J})$, it follows that there exists a minimum $t \in [T]$ such that $f(t) \geq m$, that is, such that there exists a set $K' \in \mathcal{K}^t$ with

$$\text{diam}^*(\mathcal{J}_u([K']_u)) \geq m.$$

Now, since $m \geq \lambda$, and hence $f(t) \leq 3f(t-1)$, the minimality of t implies that

$$m \leq \text{diam}^*(\mathcal{J}_u([K']_u)) \leq 3m,$$

as required. Finally, note that $\langle K' \rangle_u = \{\mathcal{J}_u([K']_u)\}$ by Lemma 5.11, since the set $[K']_u$ is strongly connected. ■

Next we prove our extremal lemma, which will be used in Section 8 to prove the base case of our induction hypothesis (see Definition 8.1 and Lemma 8.2). This lemma, which generalizes [17, Lemma 6.17] to the iceberg setting, follows easily from the iceberg spanning algorithm, together with Lemmas 6.14 and 6.17.

Lemma 6.19 (Extremal lemma for icebergs). *Let u be a non-leaf vertex of T and let $\mathcal{J} = (J, v)$ be a u -iceberg spanned iceberg. Then $|J \cap A| = \Omega(\text{diam}^*(\mathcal{J}))$.*

Proof. Since $\mathcal{J}$ is u -iceberg spanned, by Definition 5.9 there exists a set $K \subset J \cap A$ such that $[K]_u$ is strongly connected and $\mathcal{J} = \mathcal{J}_u([K]_u)$. If we apply the u -iceberg spanning algorithm to the set K , then the algorithm begins with $|K|$ sets containing the individual elements of K , and finishes with the single set K . At each step of the algorithm the number

of sets in the collection decreases by 1, and we claim that the quantity

$$\sum_{K' \in \mathcal{K}^t} \text{diam}^*(\mathcal{J}_u([K']_u))$$

increases (additively) by at most $O(1)$, by Lemma 6.14. To see this, recall (cf. the proof of Lemma 6.18) that if K_1 and K_2 are united at time t then the sets $[K_1]_u$, $[K_2]_u$ and $[K_1]_u \cup [K_2]_u$ are all strongly connected, and $K_1 \cup K_2 \subset K$. Moreover, we know that $K \subset J \subset \Lambda \setminus H_T(u)$, by Lemma 5.6, because $\mathcal{J}$ is the u -iceberg container of $[K]_u$. The inequality (50) therefore follows from Lemma 6.14, and this proves the claim.

Thus, comparing these quantities at times $t = 0$ and $t = |K| - 1$, it follows that

$$\text{diam}^*(\mathcal{J}) \leq \sum_{x \in K} \text{diam}^*(\mathcal{J}_u([\{x\}]_u)) + O(|K|) = O(|K|),$$

since $\text{diam}^*(\mathcal{J}_u([\{x\}]_u)) = O(1)$ for every $x \in \Lambda \setminus H_T(u)$ by Lemma 6.17. This implies that $|J \cap A| \geq |K| = \Omega(\text{diam}^*(\mathcal{J}))$, as required. ■

We finish this section with two lemmas about the size and number of icebergs. The first, which provides an upper bound on $|J|$ in terms of $\text{diam}^*(\mathcal{J})$, where $\mathcal{J} = (J, u)$, will allow us to bound the number of choices for the set $J \cap A$.

Lemma 6.20. *There exists $C > 0$ depending only on T such that the following holds. Let u be a non-leaf vertex of T and let $\mathcal{J} = (J, u)$ be an iceberg. Then*

$$|J| \leq (C \cdot \text{diam}^*(\mathcal{J}) + 1)^d.$$

Proof. Observe that there exists $x \in \mathbb{R}^d$ such that J is contained in $x + [0, \text{diam}(J)]^d$. Hence, since $\text{diam}(J) \leq C \cdot \text{wd}_u(J)$ for some C depending only on T , by Lemma 6.4, and $\text{wd}_u(J) \leq \delta^{-d} \text{diam}^*(\mathcal{J})$, by Definition 6.13, the lemma follows. ■

Our final lemma of the section allows us to count icebergs themselves, given limited information about their position and size.

Lemma 6.21. *There exists $C > 0$ depending only on T such that the following holds. Let u be a non-leaf vertex of T , and let $\mathcal{J} = (J, u)$ be an iceberg with $\text{diam}^*(\mathcal{J}) \geq 2$. Then the number of icebergs $\mathcal{J}' = (J', v)$ with v a non-leaf vertex of T such that*

$$J \cap J' \neq \emptyset \quad \text{and} \quad \text{diam}^*(\mathcal{J}') \leq \text{diam}^*(\mathcal{J})$$

is at most $\text{diam}^(\mathcal{J})^C$.*

Proof. Recall from Definition 5.2 that the iceberg $\mathcal{J}'$ is uniquely defined by the choices of $v \in T$ and the set $\{d_w(J') : w \in N_T^{\rightarrow}(v)\}$, where $d_w(J') := \max_{x \in J'} \langle x, w \rangle$. Let us fix $y \in J \cap J'$ and observe that, by Lemma 6.20, there are at most $(C \cdot \text{diam}^*(\mathcal{J}) + 1)^d$ choices for y . We claim that there exists C' , depending only on T , such that

$$J' \subset y + [-C' \text{diam}^*(\mathcal{J}), C' \text{diam}^*(\mathcal{J})]^d. \quad (62)$$

Indeed, this follows as in the proof of Lemma 6.20, since we have

$$\text{diam}(J') \leq C \cdot \text{wd}_v(J') \leq C \delta^{-d} \cdot \text{diam}^*(\mathcal{J}'),$$

where the first inequality holds by Lemma 6.4, and the second by Definition 6.13, and since $J' \subset G_v(J')$ for any v -iceberg droplet J' .

It follows from (62) that, given $y \in J \cap J'$, for each $w \in N_T^{\rightarrow}(v)$ we have at most $(2C' \text{diam}^*(\mathcal{J}) + 1)^d$ choices for the $x \in J'$ that maximizes $\langle x, w \rangle$. Putting the pieces together, and recalling that $\text{diam}^*(\mathcal{J}) \geq 2$, the number of choices for $\mathcal{J}'$ is at most

$$\sum_{v \in V(T)} (C \cdot \text{diam}^*(\mathcal{J}) + 1)^d (2C' \text{diam}^*(\mathcal{J}) + 1)^{d \cdot |N_T^{\rightarrow}(v)|} \leq (\text{diam}^*(\mathcal{J}))^{C''}$$

for some C'' depending only on T , as required. $\blacksquare$

7. One-step hierarchies

Hierarchies, originally introduced by Holroyd [53] in the context of the 2-neighbour model in two dimensions, and subsequently developed further in [10, 11, 17, 34–36, 43, 50], are a tool for encoding the formation of an internally spanned droplet D by a relatively small number of internally spanned subdroplets and ‘crossing’ events (see Definition 7.1). Roughly speaking, this is achieved by running the spanning algorithm (for droplets) with initial set $D \cap A$, and recording only some of the sets that appear during the evolution of the algorithm. One significant problem with using hierarchies in our setting is that we must deal with droplets (and icebergs of various types) on very different scales: from polynomial in p , up to towers of exponentials of height $r - 2$.

We overcome this issue by using a different (and, in fact, much simpler) notion of hierarchy, which we call ‘one-step hierarchies’. The reader who is familiar with the various notions of hierarchies used previously in the literature should think of these as consisting of a root vertex together with its neighbour(s) (see Lemma 7.2). There will be enough leeway in our induction hypothesis in Section 8 that we can afford to apply Lemma 7.2 at each step rather than building up larger hierarchies as in previous works [10, 17, 53].

We remark that for update families with resistance $r = 2$, we do not need the lemmas of this section. In those cases, Lemmas 6.18 and 6.19 suffice to prove Theorem 3.1 (see Lemma 8.2). Moreover, for such models the only non-leaf vertex of T is the root, in which case both lemmas are statements about $N_T^{\rightarrow}(\mathbf{0})$ -droplets. Thus, to prove the case $r = 2$ of Theorem 3.1 we do not require any of the icebergs framework that we developed in the previous two sections, and in fact the result follows in much the same way as the lower bound for critical two-dimensional models in [18].

We begin by defining precisely the event that one iceberg ‘grows’ into a larger iceberg. We shall only be interested in this event when the two icebergs have the same type.

Definition 7.1. Let u be a non-leaf vertex of T , and let $\mathcal{J} = (J, u)$ and $\mathcal{J}' = (J', u)$ be icebergs such that $J' \subset J$. Then we define

$$\Delta_u(\mathcal{J}', \mathcal{J}) := \{\langle K \rangle_u = \langle \mathcal{J} \rangle \text{ for some } K \subset J' \cup (J \cap A)\}.$$

We are now ready to state and prove our only hierarchies-type lemma. It says, roughly speaking, that if $\mathcal{J} = (J, u)$ is a u -iceberg spanned iceberg (see Definition 5.9), then one of two possibilities occurs. The first is that there exist two disjointly²⁶ u -iceberg spanned icebergs, $\mathcal{J}'_1$ and $\mathcal{J}'_2$, whose iceberg droplets intersect J , and such that $\text{diam}^*(\mathcal{J}'_1)$ and $\text{diam}^*(\mathcal{J}'_2)$ are both reasonably large, and $\text{diam}^*(\mathcal{J}'_1) + \text{diam}^*(\mathcal{J}'_2)$ is not much smaller than $\text{diam}^*(\mathcal{J})$. The second is that there exists a u -iceberg spanned iceberg $\mathcal{J}' = (J', v)$ whose iceberg droplet J' intersects J , with $\text{diam}^*(\mathcal{J}')$ not much smaller than $\text{diam}^*(\mathcal{J})$, such that either $v \neq u$, or the event $\Delta_u(\mathcal{J}', \mathcal{J})$ occurs.

The precise (and slightly technical) statement of this ‘one-step hierarchies’ lemma is as follows. Recall from Lemmas 6.18 and 6.19 that λ is a sufficiently large constant.

Lemma 7.2. *Let u be a non-leaf vertex of T , let $\mathcal{J} = (J, u)$ be a u -iceberg spanned u -iceberg, and let $\lambda \leq y \leq \text{diam}^*(\mathcal{J})/4$. Then one of the following two events occurs:*

- (a) *There exist icebergs $\mathcal{J}'_1 = (J'_1, u_1)$ and $\mathcal{J}'_2 = (J'_2, u_2)$ such that $J'_1 \cap J \neq \emptyset$ and $J'_2 \cap J \neq \emptyset$, satisfying*

$$y \leq \text{diam}^*(\mathcal{J}'_2) \leq \text{diam}^*(\mathcal{J}'_1) \leq \text{diam}^*(\mathcal{J}) - 3y \quad (63)$$

and

$$\text{diam}^*(\mathcal{J}) \leq \text{diam}^*(\mathcal{J}'_1) + \text{diam}^*(\mathcal{J}'_2) + 2y, \quad (64)$$

and such that $\mathcal{J}'_1$ and $\mathcal{J}'_2$ are disjointly u -iceberg spanned.

- (b) *There exists an iceberg $\mathcal{J}' = (J', v)$ with $J' \cap J \neq \emptyset$ satisfying*

$$y \leq \text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \leq 3y$$

such that $\text{Ice}_u^x(\mathcal{J}')$ occurs. If $v = u$, then also $J' \subset J$ and $\Delta_u(\mathcal{J}', \mathcal{J})$ occurs.

Observe that in (b) the two events $\text{Ice}_u^x(\mathcal{J}')$ and $\Delta_u(\mathcal{J}', \mathcal{J})$ automatically occur disjointly, since for each fixed $\mathcal{J}'$ the two events depend on disjoint subsets of A .

Proof of Lemma 7.2. We claim that there exists a sequence

$$J \cap A \supset K_0 \supset K_1 \supset \cdots \supset K_m$$

such that setting

$$\mathcal{J}_i := \mathcal{J}_u([K_i]_u)$$

for $0 \leq i \leq m$, we have the following: $[K_0]_u$ is strongly connected and $\mathcal{J}_0 = \mathcal{J}$, $|K_m| = 1$, and for every $1 \leq i \leq m$, $[K_i]_u$, $[K_{i-1} \setminus K_i]_u$ and $[K_i]_u \cup [K_{i-1} \setminus K_i]_u$ are all strongly connected, and

$$\text{diam}^*(\mathcal{J}_u([K_i]_u)) \geq \text{diam}^*(\mathcal{J}_u([K_{i-1} \setminus K_i]_u)). \quad (65)$$

²⁶Two increasing events E and F (depending on the set A) are said to *occur disjointly* if there exist disjoint ‘witness’ sets $X, Y \subset A$ such that $\{X \subset A\} \Rightarrow E$ and $\{Y \subset A\} \Rightarrow F$.

To construct the sequence, first we choose K_0 such that $[K_0]_u$ is strongly connected and $\mathcal{J} = \mathcal{J}_u([K_0]_u)$, which we can do, by Definition 5.9, because $\mathcal{J}$ is u -iceberg spanned. Now, supposing $K_{i-1} \subset J \subset \Lambda \setminus H_T(u)$ has already been chosen for some $i \geq 1$ and is such that $|K_{i-1}| \geq 2$ and $[K_{i-1}]_u$ is strongly connected, we apply Lemma 5.13 to K_{i-1} . This gives a subset K_i of K_{i-1} such that $[K_i]_u$, $[K_{i-1} \setminus K_i]_u$ and $[K_i]_u \cup [K_{i-1} \setminus K_i]_u$ are all strongly connected. Moreover, without loss of generality this subset satisfies (65).

Now, let $k \geq 1$ be minimal such that

$$\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}_k) \geq y.$$

Note that such a k exists because $\text{diam}^*(\mathcal{J}) \geq 4y$, and since we have $|K_m| = 1$ and $K_m \subset J \subset \Lambda \setminus H_T(u)$, so $\text{diam}^*(\mathcal{J}_m) \leq \lambda \leq y$, by Lemma 6.17. Suppose first that

$$\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}_k) \geq 3y. \quad (66)$$

We claim that in this case the conditions of (a) hold with

$$\mathcal{J}'_1 = (J'_1, u_1) := \mathcal{J}_k = \mathcal{J}_u([K_k]_u) \quad (67)$$

and

$$\mathcal{J}'_2 = (J'_2, u_2) := \mathcal{J}_u([K_{k-1} \setminus K_k]_u). \quad (68)$$

Indeed, $\mathcal{J}'_1$ and $\mathcal{J}'_2$ are disjointly u -iceberg spanned by their definitions in (67) and (68), since $[K_k]_u$ and $[K_{k-1} \setminus K_k]_u$ are both strongly connected by construction, and since $K_k \subset K_{k-1} \subset A$. To prove (64), observe that

$$\text{diam}^*(\mathcal{J}_{k-1}) \leq \text{diam}^*(\mathcal{J}'_1) + \text{diam}^*(\mathcal{J}'_2) + O(1) \quad (69)$$

by Lemma 6.14, since the sets $[K_k]_u$, $[K_{k-1} \setminus K_k]_u$ and $[K_k]_u \cup [K_{k-1} \setminus K_k]_u$ are strongly connected, and $K_{k-1} \subset J \subset \Lambda \setminus H_T(u)$. It follows that

$$\text{diam}^*(\mathcal{J}) \leq \text{diam}^*(\mathcal{J}_{k-1}) + y \leq \text{diam}^*(\mathcal{J}'_1) + \text{diam}^*(\mathcal{J}'_2) + 2y,$$

where the first step follows from by the minimality of k , and the second from (69) and the bound $y \geq \lambda$, and therefore (64) holds as claimed. To prove (63), firstly observe that

$$\text{diam}^*(\mathcal{J}'_1) \geq \text{diam}^*(\mathcal{J}'_2) \geq \text{diam}^*(\mathcal{J}_{k-1}) - \text{diam}^*(\mathcal{J}_k) - O(1) \geq y,$$

where the first inequality holds by (65), the second by (69) and since $\mathcal{J}'_1 = \mathcal{J}_k$, and the third since $\text{diam}^*(\mathcal{J}_{k-1}) - \text{diam}^*(\mathcal{J}_k) \geq 2y$, which follows from (66) and the minimality of k . Together with (66), this gives (63). Finally, both J'_1 and J'_2 intersect J because both K_k and $K_{k-1} \setminus K_k$ are contained in J . Thus, $\mathcal{J}'_1$ and $\mathcal{J}'_2$ satisfy the conditions of (a).

Now suppose that (66) does not hold, so we instead have

$$y \leq \text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}_k) \leq 3y. \quad (70)$$

We claim that the conditions of (b) are satisfied with

$$\mathcal{J}' = (J', v) := \mathcal{J}_k = \mathcal{J}_u([K_k]_u). \quad (71)$$

Indeed, $\mathcal{J}'$ is u -iceberg spanned by its definition in (71), since $[K_k]_u$ is strongly connected by construction, and $K_k \subset A$. The bounds on $\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}')$ are the same as those in (70), so they hold automatically, and we have $J' \cap J \neq \emptyset$ for the same reason as in part (a): the set K_k is a subset of J . These are all the conditions that apply when $v \neq u$, so to deal with the remaining two conditions we assume in addition that $v = u$. These conditions are in fact straightforward: the inclusion $J' \subset J$ follows from $K_k \subset K_0$ and from J' and J being the smallest u -iceberg droplets containing $[K_k]_u$ and $[K_0]_u$ respectively (see Definition 5.5), and the event $\Delta_u(\mathcal{J}', \mathcal{J})$ is an immediate consequence of the event $\text{Ice}_u^\times(\mathcal{J})$ and Lemma 5.12, since we may take the set K in Definition 7.1 to be equal to any $K \subset J \cap A$ such that $\langle K \rangle_u = \{\mathcal{J}\}$. ■

The ‘crossing’ event $\Delta_u(\mathcal{J}', \mathcal{J})$ appeared naturally in the proof of Lemma 7.2, but it is not immediately clear how we can bound its probability. We will do so using our final deterministic property of icebergs, Lemma 7.4, which will allow us to deduce from the event $\Delta_u(\mathcal{J}', \mathcal{J})$ the occurrence of another event that we know better how to control: namely, the existence of a certain iceberg spanned iceberg. In the proof of Lemma 7.4 we shall use the following simple property of a pair of nested u -icebergs.

Lemma 7.3. *Let u be a non-leaf vertex of T , let $k := d_T(\mathbf{0}, u)$, and let $\mathcal{J} = (J, u)$ and $\mathcal{J}' = (J', u)$ be u -icebergs with $J' \subset J$. Then*

$$\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \leq \delta^k \cdot \max \{d_v(J) - d_v(J') : v \in N_T^{\rightarrow}(u)\}.$$

Proof. By Definitions 6.2 and 6.13, we have

$$\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') = \delta^k \sum_{v \in M_T(u)} \lambda_v (d_v(G_u(J)) - d_v(G_u(J'))),$$

where $\Xi_u = \{\lambda_v : v \in M_T(u)\}$. Now, by Definition 6.6, we have

$$d_v(G_u(J)) = d_v(J) \quad \text{and} \quad d_v(G_u(J')) = d_v(J')$$

for each $v \in N_T^{\rightarrow}(u)$. Moreover, if $u \neq \mathbf{0}$, then $d_{-u}(G_u(J)) = d_{-u}(G_u(J')) = 0$, since $N_T^{\rightarrow}(u)$ is not an $\mathbf{S}^{d-1}$ -bounding set, by Lemma 4.12. It follows that

$$\begin{aligned} \text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') &= \delta^k \sum_{v \in N_T^{\rightarrow}(u)} \lambda_v (d_v(J) - d_v(J')) \\ &\leq \delta^k \cdot \max \{d_v(J) - d_v(J') : v \in N_T^{\rightarrow}(u)\}, \end{aligned}$$

as required, since $\lambda_v \geq 0$ for every $v \in M_T(u)$, and $\sum_{v \in M_T(u)} \lambda_v = 1$. ■

Up to this point, we have defined icebergs (and their related structures and events) with respect to unions of half-spaces passing through the origin. However, in our final deterministic lemma we need to allow the role of the origin to be played by other elements of $\mathbb{R}^d$. As noted at the beginning of Section 5, this is why we have worked in the shifted lattice Λ , instead of in $\mathbb{Z}^d$.

In order to state and prove Lemma 7.4, let us generalize some of the definitions of Sections 5 and 6 by replacing the origin by some arbitrary element $x \in \mathbb{R}^d$. First, if $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_k = u$ is a path in T , then we define

$$H_T(u, x) := \bigcup_{i=1}^k (\mathbf{H}_{u_i}^d + x) \cap \Lambda. \quad (72)$$

Let u be a non-leaf vertex of T , and let $K \subset \Lambda \setminus H_T(u, x)$ be a finite, non-empty set.

- The (u, x) -iceberg closure of K is

$$[K]_{u,x} := [K \cup H_T(u, x)]_u \setminus H_T(u, x). \quad (73)$$

- An x -shifted u -iceberg is a triple $\mathcal{J} = (J, u, x)$, where J is a (u, x) -iceberg droplet, i.e., a non-empty set of the form $D \setminus H_T(u, x)$, where D is an $N_T^{\rightarrow}(u)$ -droplet.
- The (u, x) -iceberg spanning algorithm is defined as in Definition 5.10, except replacing the u -iceberg closure by the (u, x) -iceberg closure. The output of the algorithm is the (u, x) -iceberg span of K ,

$$\langle K \rangle_{u,x} := \{\mathcal{J}_{u,x}([K']_u) : K' \in \mathcal{K}^T\},$$

where $\mathcal{J}_{u,x}(K)$, the (u, x) -iceberg container of K , is defined as in Definition 5.5, except with J_i defined to be the minimal (u_i, x) -iceberg droplet containing K , and with each half-space $\mathbf{H}_{u_i}^d$ replaced by $\mathbf{H}_{u_i}^d + x$.

- The event $\text{Ice}_{u,x}^{\times}(\mathcal{J})$ holds (and we say that $\mathcal{J}$ is (u, x) -iceberg spanned) if there exists $K \subset J \cap A$ such that $[K]_{u,x}$ is strongly connected and $\mathcal{J} = \mathcal{J}_{u,x}([K]_{u,x})$.²⁷
- If $\mathcal{J} = (J, u, x)$ is an x -shifted u -iceberg, and $d_T(\mathbf{0}, u) = k$, then

$$\text{diam}_x^*(\mathcal{J}) := \delta^k \text{wd}_u(G_{u,x}(J)), \quad (74)$$

where $G_{u,x}(J) = \{z \in F_u^*(J) : \langle z, u \rangle \geq \langle x, u \rangle\}$.

Observe that, because we were working in the lattice $\Lambda = y + \mathbb{Z}^d$, where $y \in \mathbb{R}^d$ was arbitrary, all of the results proved in Sections 5 and 6 also hold in this setting.

Lemma 7.4. *Let u be a vertex of T with $d_T(\mathbf{0}, u) \leq r - 3$, and let $\mathcal{J} = (J, u)$ and $\mathcal{J}' = (J', u)$ be icebergs with $J' \subset J$ and $\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \geq \lambda$.*

If the event $\Delta_u(\mathcal{J}', \mathcal{J})$ occurs, then there exist $x \in \mathbb{R}^d$ and an x -shifted iceberg $\mathcal{J}''$ with

$$\text{diam}_x^*(\mathcal{J}'') \geq \delta^3 \cdot (\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}')) \quad (75)$$

such that $\langle K \rangle_{v,x} = \{\mathcal{J}''\}$ for some $K \subset (J \cap A) \setminus J'$ and $v \in \{u\} \cup N_T^{\rightarrow}(u)$.

Proof. By Definition 7.1, the event $\Delta_u(\mathcal{J}', \mathcal{J})$ implies that there exists a set $K' \subset J' \cup (J \cap A)$ such that $\langle K' \rangle_u = \{\mathcal{J}\}$. Suppose first $K' \subset (J \cap A) \setminus J'$, and observe that the

²⁷By Lemma 5.12, this is equivalent to the event that $\langle K \rangle_{u,x} = \{\mathcal{J}\}$ for some $K \subset J \cap A$.

iceberg $\mathcal{J}'' = \mathcal{J}$ has the required properties. Indeed, (75) holds trivially with $x = \mathbf{0}$, and $\langle K' \rangle_{u, \mathbf{0}} = \langle K' \rangle_u = \{\mathcal{J}''\}$, as required. We may therefore assume that $K' \cap J' \neq \emptyset$.

We begin by choosing $v \in N_T^{\rightarrow}(u)$ and $x \in \mathbb{R}^d$. To do so, note first that, by Lemma 7.3, there exists $v \in N_T^{\rightarrow}(u)$ such that

$$\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \leq \delta^k \cdot (d_v(J) - d_v(J')), \quad (76)$$

where $k := d_T(\mathbf{0}, u)$. Moreover, v is not a leaf of T , since $k \leq r - 3$ and T has depth $r - 1$. Note that the set²⁸

$$S := \{z \in J : \langle z, v \rangle > d_v(J')\}$$

is finite, by Lemma 5.3, and choose $a \in S$ with $\langle a, v \rangle$ minimal.

Now, let $\mathbf{0} \rightarrow u_1 \rightarrow \dots \rightarrow u_{k-1} \rightarrow u \rightarrow v$ be the path in T from the root to v , and recall from Definition 4.1 that the directions $\{u_1, \dots, u_{k-1}, u, v\}$ are linearly independent. Since there are at most $r - 2 \leq d - 2$ linear constraints, it follows that there exists an element $x \in \mathbb{R}^d$ such that $\langle x - a, v \rangle = 0$ and $\langle x, w \rangle = 0$ for every $w \in \{u_1, \dots, u_{k-1}, u\}$.

Having chosen $v \in N_T^{\rightarrow}(u)$ and $x \in \mathbb{R}^d$, we now define an x -shifted iceberg

$$\mathcal{J}'' := \mathcal{J}_{v,x}(L) = (J'', w, x), \quad (77)$$

where L is any strongly connected component of $[K' \cap S]_{v,x}$ such that $d_v(L) \geq d_v(J)$. We claim that such a set L exists, and that $\mathcal{J}''$ has the required properties. To prove this, we need the following simple claim, which follows from our choice of x .

Claim 7.5. $[K']_u \setminus H_T(v, x) \subset [K' \cap S]_{v,x}$.

Proof of Claim 7.5. Note first that $H_T(v, x) = H_T(u) \cup ((\mathbf{H}_v^d + a) \cap \Lambda)$, by (72) and our choice of x . Noting that $K' \subset (K' \cap S) \cup (\mathbf{H}_v^d + a)$, since $K' \subset J$, we obtain

$$[K']_u \subset [K' \cup H_T(u)]_u \subset [(K' \cap S) \cup H_T(v, x)]_u.$$

Removing the set $H_T(v, x)$ from both sides, and recalling (73), the claim follows. ■

Now, recall that $\langle K' \rangle_u = \{\mathcal{J}\}$ and $\mathcal{J} = (J, u)$, and note that therefore $\mathcal{J} = \mathcal{J}_u([K']_u)$, by Lemma 5.11. By Definition 5.5, it follows that J is the minimal u -iceberg droplet containing $[K']_u$. Since $v \in N_T^{\rightarrow}(u)$, this implies that

$$\langle y, v \rangle = d_v([K']_u) = d_v(J) \geq \langle x, v \rangle \quad (78)$$

for some $y \in [K']_u$, the inequality holding because $\langle x, v \rangle = \langle a, v \rangle$ and $a \in J$.

To deduce that L exists, observe that, by Claim 7.5, we have

$$[K']_u \subset [K' \cap S]_{v,x} \cup (\mathbf{H}_v^d + x),$$

²⁸We think of S as the ‘strip’ between J' and J in direction v . Note that S is non-empty, by (76) and since $\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \geq \lambda > 0$.

since $[K']_u \cap H_T(u) = \emptyset$. It follows that $y \in [K' \cap S]_{v,x}$ since, by (78), we have $\langle y, v \rangle \geq \langle x, v \rangle > \langle z, v \rangle$ for all $z \in \mathbf{H}_v^d + x$. In particular, if L is the strongly connected component of $[K' \cap S]_{v,x}$ containing y , then $d_v(L) \geq \langle y, v \rangle \geq d_v(J)$.

Next, observe that $L = [L \cap K' \cap S]_{v,x}$, since L is a strongly connected component, so must contain every site of $K' \cap S$ that is used to infect it. Set

$$K := L \cap K' \cap S \subset (J \cap A) \setminus J',$$

so $L = [K]_{v,x}$ is strongly connected. Recalling from (77) that $\mathcal{J}'' = \mathcal{J}_{v,x}(L)$, it follows by Lemma 5.11 that $\langle K \rangle_{v,x} = \{\mathcal{J}''\}$, as required.

It remains to verify the lower bound (75) on $\text{diam}_x^*(\mathcal{J}'')$. To do so, recall (74), and observe first that

$$\text{diam}_x^*(\mathcal{J}'') \geq \delta^{k+1} \cdot \text{wd}_w(J'') \geq \delta^{k+2} \cdot \text{diam}(J''), \quad (79)$$

where the first inequality holds because $J'' \subset G_{w,x}(J'')$ and $d_T(\mathbf{0}, w) \leq k+1$, and the second by Lemma 6.4. To deduce (75), we shall use (76) and the following claim.

Claim 7.6. *L is strongly connected to $\mathbf{H}_v^d + x$.*

Proof of Claim 7.6. Observe first that, by Claim 7.5, we have

$$[K' \cap S]_{v,x} = [[K']_u \cap S]_{v,x}. \quad (80)$$

We shall show that $([K']_u \cap S) \cup (\mathbf{H}_v^d + x)$ is strongly connected, which will in turn imply that every strongly connected component of $[K']_u \cap S$ is strongly connected to $\mathbf{H}_v^d + x$. By (80), it will then follow that every strongly connected component of $[K' \cap S]_{v,x}$ is strongly connected to $\mathbf{H}_v^d + x$, as required.

To show that $([K']_u \cap S) \cup (\mathbf{H}_v^d + x)$ is strongly connected, recall that $\langle K' \rangle_u = \{\mathcal{J}\}$, which implies that $[K']_u$ is strongly connected, by Lemma 5.11. It also follows from $\langle K' \rangle_u = \{\mathcal{J}\}$ that $[K']_u \subset J$, and therefore $[K']_u \setminus (\mathbf{H}_v^d + x) \subset S$, by the definitions of S and x . Since $K' \cap J'$ is non-empty and $J' \subset \mathbf{H}_v^d + x$, it follows that the set

$$[K']_u \cup (\mathbf{H}_v^d + x) = ([K']_u \cap S) \cup (\mathbf{H}_v^d + x)$$

is strongly connected, as required. ■

By Claim 7.6, and since $d_v(L) \geq d_v(J)$, it follows that

$$\text{diam}(J'') \geq d_v(J) - \langle x, v \rangle - R,$$

and hence, by (79), and noting that $\langle x, v \rangle \leq d_v(J') + R$, we obtain

$$\text{diam}_x^*(\mathcal{J}'') \geq \delta^{k+2} (d_v(J) - d_v(J') - 2R).$$

Combining this with (76), it follows that

$$\text{diam}_x^*(\mathcal{J}'') \geq \delta^3 \cdot (\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}')),$$

as required, since $\text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \geq \lambda$. ■

Remark 7.7. Observe that, in the proof above, we took x equal either to $\mathbf{0}$ (when $v = u$), or to a deterministic function of the pair (J, J') and the vertex $v \in N_T^{\rightarrow}(u)$. This fact will be used in the next section in order to bound the number of choices for $\mathcal{J}''$.

8. The proof of Theorem 3.1

Our toolbox of deterministic lemmas is now complete, and we are ready to begin assembling the parts of the proof of Theorem 3.1. The proof will be by induction on r , and also by induction on the diameter of the iceberg.

To begin, we need to fix some constants. In view of Lemma 6.19, choose $c > 0$ such that, whenever $u \in V(T)$ is not a leaf²⁹ and $\mathcal{J} = (J, v)$ is u -iceberg spanned, we have

$$|J \cap A| \geq c \cdot \text{diam}^*(\mathcal{J}). \quad (81)$$

Next, select an arbitrary sequence $\varepsilon(2), \dots, \varepsilon(r)$ such that

$$\frac{1}{d-1} > \varepsilon(2) > \dots > \varepsilon(r) > 0. \quad (82)$$

The function

$$x(s) := \exp_{(s-2)}(p^{-\varepsilon(s)}), \quad (83)$$

defined for each $2 \leq s \leq r$, will be (up to a constant factor) the maximum diameter of the icebergs of type u that we consider when $d_T(\mathbf{0}, u) = r - s$.

Recall from Section 6 that λ is a sufficiently large constant, which is allowed to depend on all other constants (in particular, it will depend on $\varepsilon(2)$ and c).³⁰ We emphasize that p is not considered to be a constant, and will always be assumed to be sufficiently small. Now, let $q: \mathbb{R} \rightarrow \mathbb{R}$ be the function defined by $q(x) = 1$ for all $x < 1$, and

$$q(x) := \exp\left(-x \cdot \frac{\log(1/p)}{\log_{(r-1)}(\lambda \cdot x)}\right) \quad (84)$$

for every $x \geq 1$. Thus, up to constant factors in the exponents, $q(x)$ interpolates between p^x when $x = O(1)$, and e^{-x} when x is close to $x(r)$. This will be our bound on the probability that an iceberg of diameter x is u -iceberg spanned (for appropriate u).

When we introduced the iceberg diameter $\text{diam}^*(\mathcal{J})$ in Definition 6.13, we ‘gave away’ a factor of δ^k , which allowed us to prove Lemma 6.14 even when the icebergs involved were of different types. We now give away another factor of δ^k (though in this case a factor of 2^k would be more than sufficient), which will allow us to use the trivial bound on the probability of the event $\Delta_u(\mathcal{J}', \mathcal{J})$ whenever $\mathcal{J}$ and $\mathcal{J}'$ are icebergs of different types. To that end, if $\mathcal{J} = (J, u)$ is an iceberg and $k := d_T(\mathbf{0}, u)$, then we define

$$\text{diam}^{**}(\mathcal{J}) := \delta^k \text{diam}^*(\mathcal{J}). \quad (85)$$

We can now state our induction hypothesis.

²⁹Recall that the tree T was fixed in Section 4.3.

³⁰We shall need to assume that $\lambda > \exp_{(r-1)}(c^{-1}(1 - \varepsilon(2)(d-1))^{-1})$.

Definition 8.1 (The induction hypothesis). For each $2 \leq s \leq r$ and each $x \leq x(s)$, let $\text{IH}(s, x)$ denote the following statement:

(*) For each $u \in V(T)$ with $d_T(\mathbf{0}, u) = r - s$, and every iceberg $\mathcal{J}$ with $\text{diam}^*(\mathcal{J}) \leq x$,

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J})) \leq q(\text{diam}^{**}(\mathcal{J})).$$

We shall prove the induction hypothesis in two steps, as follows.

Lemma 8.2. *The statement $\text{IH}(s, x(2))$ holds for each $2 \leq s \leq r$.*

Lemma 8.3. *For each $3 \leq s \leq r$ and each $x(2) < x \leq x(s)$, we have*

$$\text{IH}(s-1, x(t-1)) \cap \text{IH}(s, x-1) \cap (\text{IH}(s+1, x) \cap \cdots \cap \text{IH}(r, x)) \implies \text{IH}(s, x),$$

where $3 \leq t \leq s$ is such that $x(t-1) < x \leq x(t)$.

Before proving Lemmas 8.2 and 8.3, let us quickly observe that together they imply that $\text{IH}(s, x)$ holds for all $2 \leq s \leq r$ and all $x \leq x(s)$, and show that this is sufficient to prove Proposition 3.5. In fact, we shall prove the following slightly stronger proposition, which moreover specifies the set $\mathcal{T}$ explicitly (this will turn out to be useful later).

Proposition 8.4. *Let $\mathcal{J} = (J, \mathbf{0})$ be a $\mathbf{0}$ -iceberg such that $\text{diam}^*(\mathcal{J}) \leq x(r)$. Then*

$$\mathbb{P}_p(\text{Ice}_\mathbf{0}^\times(\mathcal{J})) \leq \exp(-\text{diam}^*(\mathcal{J})).$$

Proof. Suppose Lemmas 8.2 and 8.3 hold. We shall prove, by induction on t , that $\text{IH}(s, x)$ holds for every $t \leq s \leq r$ and all $x \leq x(t)$. The case $t = 2$ holds by Lemma 8.2, so let $3 \leq t \leq r$ and suppose that we have already proved that $\text{IH}(s, x)$ holds for every $t-1 \leq s \leq r$ and all $x \leq x(t-1)$. By Lemma 8.3 and induction on x , it follows that $\text{IH}(r, x)$ holds for all $x \leq x(t)$. Now, again by Lemma 8.3 and induction on x , it follows that $\text{IH}(r-1, x)$ holds for all $x \leq x(t)$. Repeating this argument $r-t+1$ times, we deduce that $\text{IH}(s, x)$ holds for every $t \leq s \leq r$ and all $x \leq x(t)$, as claimed.

It follows that $\text{IH}(s, x)$ holds for every $2 \leq s \leq r$ and $x \leq x(s)$, and so, in particular, $\text{IH}(r, x(r))$ holds. Therefore, if $\mathcal{J}$ is a $\mathbf{0}$ -iceberg with $\text{diam}^*(\mathcal{J}) \leq x(r)$, then

$$\mathbb{P}_p(\text{Ice}_\mathbf{0}^\times(\mathcal{J})) \leq q(\text{diam}^{**}(\mathcal{J})) = q(\text{diam}^*(\mathcal{J})) \leq \exp(-\text{diam}^*(\mathcal{J})),$$

where the first step holds by $\text{IH}(r, x(r))$, applied with $u = \mathbf{0}$, the second because $\mathcal{J}$ is a $\mathbf{0}$ -iceberg, so $\text{diam}^{**}(\mathcal{J}) = \text{diam}^*(\mathcal{J})$, and the third by (83) and (84), since $\varepsilon(r) < 1$. ■

Proposition 3.5 follows easily from Proposition 8.4, but let us spell out the details for completeness.

Proof of Proposition 3.5. Set $\mathcal{T} := N_T^{\rightarrow}(\mathbf{0})$ and $\varepsilon := \varepsilon(r)$, and let D be an $N_T^{\rightarrow}(\mathbf{0})$ -droplet with $\text{diam}(D) \leq x(r) = \exp_{(r-2)}(p^{-\varepsilon})$. Note that D is a $\mathbf{0}$ -iceberg droplet, by Definition 5.2, set $\mathcal{J} := (D, \mathbf{0})$, and observe that

$$\text{diam}^*(\mathcal{J}) = \text{wd}_\mathbf{0}(G_\mathbf{0}(D)) = \text{wd}_\mathbf{0}(D) \leq \text{diam}(D) \leq x(r), \quad (86)$$

where the first step is Definition 6.13, the second follows from Definitions 6.2 and 6.6, the third holds by Lemma 6.4, and the fourth holds by assumption. Hence, by Proposition 8.4, we have

$$\mathbb{P}_p(\text{Ice}_0^\times(\mathcal{J})) \leq \exp(-\text{diam}^*(\mathcal{J})) \leq \exp(-\delta' \cdot \text{diam}(D))$$

for some sufficiently small constant $\delta' > 0$, where the final inequality holds because $\text{diam}^*(\mathcal{J}) = \text{wd}_0(D)$, by (86), and by Lemma 6.4.

To complete the proof, we shall show that the event $\text{Ice}_0^\times(\mathcal{J})$ is the same as the event that the droplet D is internally $N_T^{\rightarrow}(\mathbf{0})$ -spanned. To see this, recall from Definition 5.9 that, since $\mathcal{J} = (D, \mathbf{0})$, the event $\text{Ice}_0^\times(\mathcal{J})$ holds if and only if there exists $K \subset D \cap A$ such that $[K]_0$ is strongly connected and $\mathcal{J}_0([K]_0) = (D, \mathbf{0})$. Note that $[K]_0 = [K]_u$, by Definition 5.1, and that $\mathcal{J}_0([K]_0) = (J, \mathbf{0})$, where J is the minimal $N_T^{\rightarrow}(\mathbf{0})$ -droplet containing $[K]_0$, by Definitions 5.2 and 5.5. It follows that $\text{Ice}_0^\times(\mathcal{J})$ holds if and only if there exists $K \subset D \cap A$ such that $[K]_u$ is strongly connected and D is the minimal $N_T^{\rightarrow}(\mathbf{0})$ -droplet containing $[K]_u$. By Definition 3.4, this is the same as the event that D is internally $N_T^{\rightarrow}(\mathbf{0})$ -spanned, as required. ■

In order to prove Proposition 8.4, it remains to prove Lemmas 8.2 and 8.3. We begin with Lemma 8.2, which follows easily from our extremal lemma, Lemma 6.19, and the simple bound on the volume of an iceberg given by Lemma 6.20.

Proof of Lemma 8.2. Let $2 \leq s \leq r$, and let $u \in V(T)$ with $d_T(\mathbf{0}, u) = r - s$, so in particular u is not a leaf. Let $\mathcal{J} = (J, v)$ be an iceberg with $\text{diam}^*(\mathcal{J}) \leq x(2) = p^{-\varepsilon(2)}$. We are required to show that

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J})) \leq q(\text{diam}^{**}(\mathcal{J})). \quad (87)$$

The bound holds trivially if $\text{diam}^{**}(\mathcal{J}) < 1$, since $q(x) = 1$ for every $0 \leq x < 1$. We may therefore assume that $\text{diam}^{**}(\mathcal{J}) \geq 1$, and hence

$$|J| \leq (C \cdot \text{diam}^*(\mathcal{J}) + 1)^d \leq \lambda \cdot (\text{diam}^*(\mathcal{J}))^d,$$

by Lemma 6.20, and since λ is sufficiently large. Now, recall from (81) (which followed from Lemma 6.19) that if $\mathcal{J}$ is u -iceberg spanned, then $|J \cap A| \geq c \cdot \text{diam}^*(\mathcal{J})$. Hence, by Markov's inequality, $\mathcal{J}$ is u -iceberg spanned with probability at most

$$\binom{\lambda x^d}{cx} p^{cx} = O(x^{d-1} p)^{cx} \leq O(p^{1-\varepsilon(2)(d-1)})^{cx} = p^{\Omega(x)} \leq q(x), \quad (88)$$

where $x = \text{diam}^*(\mathcal{J})$, the second step holds since $x \leq p^{-\varepsilon(2)}$, the third holds because $\varepsilon(2) < 1/(d-1)$, and the last follows from (84), since λ is sufficiently large.

Finally, note that $x = \text{diam}^*(\mathcal{J}) \geq \text{diam}^{**}(\mathcal{J})$, and that q is decreasing. It follows that $q(x) \leq q(\text{diam}^{**}(\mathcal{J}))$, and hence (87) follows from (88). ■

The proof of Lemma 8.3 (the induction step) is significantly more challenging. The idea is to apply Lemma 7.2 to the iceberg $\mathcal{J}$, and then bound the probabilities of the

various events using Lemma 7.4 and the induction hypothesis. We shall also make use of the following fundamental lemma of van den Berg and Kesten [75]. If E and F are increasing events, then we write $E \circ F$ for the event that E and F occur disjointly.

Lemma 8.5 (The van den Berg–Kesten inequality). *Let E and F be any two increasing events and let $p \in (0, 1)$. Then*

$$\mathbb{P}_p(E \circ F) \leq \mathbb{P}_p(E) \mathbb{P}_p(F).$$

Three straightforward inequalities will be needed in order to show that the bounds we obtain are sufficiently strong. The proofs of these inequalities, Lemmas 8.6–8.8 below, can be found in Appendix A.

Let $C > 0$ be a sufficiently large constant (depending on T), and fix a sufficiently small constant $\eta > 0$ with $(1 + \eta)\varepsilon(3) < \varepsilon(2)$. The first inequality will be used when event (a) of Lemma 7.2 occurs. We write $\beta \gg 1$ to denote that $\beta \rightarrow \infty$ as $p \rightarrow 0$.

Lemma 8.6. *Let $r \geq 3$, and let $\alpha \geq \beta \gg 1$ and $\gamma \in \mathbb{R}$ satisfy*

$$\gamma \leq C(\log \alpha)^{1+\eta} \leq C^2 \beta.$$

Then

$$q(\alpha) \cdot q(\beta) \cdot (\alpha + \beta + \gamma)^{4C} \leq q(\alpha + \beta + \gamma).$$

The second inequality will be used when event (b) of Lemma 7.2 occurs, but only when the icebergs $\mathcal{J}$ and $\mathcal{J}'$ of that lemma are of different types.

Lemma 8.7. *Let $\alpha \geq \gamma \gg 1$. Then*

$$q(\alpha) \cdot (\alpha + \gamma)^{3C} \leq q(\delta(\alpha + \gamma)).$$

The third inequality will also be used when event (b) of Lemma 7.2 occurs and the icebergs of that lemma are of the same type.

Lemma 8.8. *Let $r \geq 3$ and $\alpha, \gamma \gg 1$ be such that*

$$C^{-1}(\log(\alpha + \gamma))^{1+\eta} \leq \gamma \leq C(\log(\alpha + \gamma))^{1+\eta}.$$

Then

$$q(\alpha) \cdot q(\delta^4 \gamma) \cdot (\alpha + \gamma)^{5C} \leq q(\alpha + \gamma).$$

We are now ready to prove the induction step, Lemma 8.3.

Proof of Lemma 8.3. Let $3 \leq s \leq r$, let $x(2) < x \leq x(s)$, and let $3 \leq t \leq s$ be such that $x(t-1) < x \leq x(t)$. Our aim is to show that $\text{IH}(s, x)$ holds, so let $u \in V(T)$ be a vertex with $d_T(\mathbf{0}, u) = k := r - s$, and let $\mathcal{J} = (J, w)$ be an iceberg with $\text{diam}^*(\mathcal{J}) \leq x$. Note in particular that u is not a leaf of T . Under the assumption that

$$\text{IH}(s-1, x(t-1)) \cap \text{IH}(s, x-1) \cap (\text{IH}(s+1, x) \cap \cdots \cap \text{IH}(r, x)) \quad (89)$$

holds, we must prove

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J})) \leq q(\text{diam}^{**}(\mathcal{J})). \quad (90)$$

To begin, let us eliminate cases that are either trivial or dealt with by (89). First, by Lemma 8.2, we may assume that $\text{diam}^*(\mathcal{J}) > x(2)$. Next, we may assume that $w \in [\mathbf{0}, u]_T$, because if not then $\mathcal{J}$ cannot be u -iceberg spanned, so $\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J})) = 0$. Third, note that the event $\text{Ice}_u^\times(\mathcal{J})$ implies $\text{Ice}_w^\times(\mathcal{J})$, by Lemma 5.14, and therefore if $d_T(\mathbf{0}, w) = r - s'$ for some $s + 1 \leq s' \leq r$, then

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J})) \leq \mathbb{P}_p(\text{Ice}_w^\times(\mathcal{J})) \leq q(\text{diam}^{**}(\mathcal{J}))$$

by IH(s', x), which holds by (89). We may therefore assume that $w = u$.

Suppose then that $\mathcal{J} = (J, u)$ and that $\text{Ice}_u^\times(\mathcal{J})$ occurs. Apply Lemma 7.2 to $\mathcal{J}$, with

$$y := (\log \text{diam}^*(\mathcal{J}))^{1+\eta}, \quad (91)$$

observing that $\lambda \leq y \leq \text{diam}^*(\mathcal{J})/4$, since $\text{diam}^*(\mathcal{J}) > x(2) \gg 1$ as $p \rightarrow 0$. The remainder of the proof is divided into three claims, according to whether event (a) or event (b) in Lemma 7.2 occurs, and in the case of event (b), whether or not the new iceberg $\mathcal{J}'$ is also of type u . We refer to these three events as E_1 , E_2 and E_3 respectively. That is:

- E_1 is the event that there exist icebergs $\mathcal{J}_1 = (J_1, u_1)$ and $\mathcal{J}_2 = (J_2, u_2)$ such that $J_1 \cap J \neq \emptyset$ and $J_2 \cap J \neq \emptyset$ satisfying

$$y \leq \text{diam}^*(\mathcal{J}_2) \leq \text{diam}^*(\mathcal{J}_1) \leq \text{diam}^*(\mathcal{J}) - 3y \quad (92)$$

and

$$\text{diam}^*(\mathcal{J}) \leq \text{diam}^*(\mathcal{J}_1) + \text{diam}^*(\mathcal{J}_2) + 2y, \quad (93)$$

and such that $\mathcal{J}_1$ and $\mathcal{J}_2$ are disjointly u -iceberg spanned.

- E_2 is the event that there exists an iceberg $\mathcal{J}' = (J', v)$ with $v \neq u$, $J' \cap J \neq \emptyset$,

$$y \leq \text{diam}^*(\mathcal{J}) - \text{diam}^*(\mathcal{J}') \leq 3y, \quad (94)$$

and such that $\text{Ice}_u^\times(\mathcal{J}')$ occurs.

- E_3 is the event that there exists an iceberg $\mathcal{J}' = (J', u)$ such that $J' \subset J$, the bounds (94) hold, and the events $\text{Ice}_u^\times(\mathcal{J}')$ and $\Delta_u(\mathcal{J}', \mathcal{J})$ both occur.

We shall show in the next three claims that

$$\mathbb{P}_p(E_i) \leq q(\text{diam}^{**}(\mathcal{J}))/3$$

for each $i \in \{1, 2, 3\}$. Since, by Lemma 7.2, the event $\text{Ice}_u^\times(\mathcal{J})$ implies that the event $E_1 \cup E_2 \cup E_3$ holds, these three bounds together imply (90) via the union bound. The three claims below will therefore complete the proof of the lemma.

Claim 8.9. $\mathbb{P}_p(E_1) \leq q(\text{diam}^{**}(\mathcal{J}))/3$.

To prove the claim we shall use the union bound over pairs $(\mathcal{J}_1, \mathcal{J}_2)$ of icebergs, bounding the number of choices for $\mathcal{J}_1$ and $\mathcal{J}_2$ using Lemma 6.21, and the probability that the icebergs $\mathcal{J}_1$ and $\mathcal{J}_2$ are disjointly u -iceberg spanned using the van den Berg–Kesten inequality (Lemma 8.5) and $\text{IH}(s, x - 1)$ (which is one of our assumptions in (89)).

Proof of Claim 8.9. First, set

$$\alpha := \text{diam}^*(\mathcal{J}_1), \quad \beta := \text{diam}^*(\mathcal{J}_2) \quad \text{and} \quad \gamma := \text{diam}^*(\mathcal{J}) - \alpha - \beta.$$

By Lemma 6.21, it follows that we have at most

$$(\text{diam}^*(\mathcal{J}))^{2C} \leq (\delta^k(\alpha + \beta + \gamma))^{3C}$$

choices for the icebergs $\mathcal{J}_1$ and $\mathcal{J}_2$. Indeed, the inequality follows since $\text{diam}^*(\mathcal{J}) = \alpha + \beta + \gamma$ and since $\text{diam}^*(\mathcal{J}) > x(2) \gg 1$. Next, observe that

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J}_i)) \leq q(\text{diam}^{**}(\mathcal{J}_i))$$

for each $i \in \{1, 2\}$, by $\text{IH}(s, x - 1)$, since $\text{diam}^*(\mathcal{J}_i) \leq \text{diam}^*(\mathcal{J}) - 3y$ and $\text{diam}^*(\mathcal{J}) \leq x$. Moreover, recalling (85) and that q is a decreasing function, we have

$$q(\text{diam}^{**}(\mathcal{J}_1)) \leq q(\delta^k \alpha) \quad \text{and} \quad q(\text{diam}^{**}(\mathcal{J}_2)) \leq q(\delta^k \beta),$$

since if $\mathcal{J}_i$ is u -iceberg spanned then $u_i \in [0, u]_T$.

Thus, by the van den Berg–Kesten inequality (Lemma 8.5), and applying the union bound over all choices of $\mathcal{J}_1$ and $\mathcal{J}_2$, we obtain

$$\mathbb{P}_p(E_1) \leq \sup_{\alpha, \beta, \gamma} (\delta^k(\alpha + \beta + \gamma))^{3C} \cdot q(\delta^k \alpha) \cdot q(\delta^k \beta), \quad (95)$$

where the supremum is taken over the valid choices of α , β and γ .

We shall bound (95) using Lemma 8.6, applied to $\delta^k \alpha$, $\delta^k \beta$ and $\delta^k \gamma$. We may do so, since

$$\alpha \geq \beta \geq y = (\log \text{diam}^*(\mathcal{J}))^{1+\eta} \gg 1,$$

by (92) and since $\text{diam}^*(\mathcal{J}) > x(2)$, and

$$\gamma \leq 2y = 2(\log(\alpha + \beta + \gamma))^{1+\eta} \leq 3(\log \alpha)^{1+\eta} \leq 3y \leq 3\beta,$$

by (92) and (93), and noting that $\beta + \gamma \leq 3\alpha$. It follows, by Lemma 8.6 and (95), that

$$\mathbb{P}_p(E_1) \leq \sup_{\alpha, \beta, \gamma} \frac{q(\delta^k(\alpha + \beta + \gamma))}{(\delta^k(\alpha + \beta + \gamma))^C} \leq \frac{q(\text{diam}^{**}(\mathcal{J}))}{3},$$

as claimed, since $\text{diam}^{**}(\mathcal{J}) = \delta^k(\alpha + \beta + \gamma) \gg 1$. ■

We next turn to the events E_2 and E_3 . In both events we have an iceberg $\mathcal{J}' = (J', v)$ with $J' \cap J \neq \emptyset$ such that (94) holds and the event $\text{Ice}_u^\times(\mathcal{J}')$ occurs. We shall take a union bound over choices for $\mathcal{J}'$, bounding the number of choices using Lemma 6.21, and the

probability that $\mathcal{J}'$ is u -iceberg spanned using the $\text{IH}(s, x - 1)$. For the event E_3 , an extra argument using Lemma 7.4 will be needed to bound the probability of $\Delta_u(\mathcal{J}', \mathcal{J})$.

To save repetition, let us perform the common part of the argument here, before separating into cases. Set

$$\alpha := \text{diam}^*(\mathcal{J}') \quad \text{and} \quad \gamma := \text{diam}^*(\mathcal{J}) - \alpha, \quad (96)$$

and observe that, by Lemma 6.21, we have at most

$$(\text{diam}^*(\mathcal{J}))^C \leq (\delta^k(\alpha + \gamma))^{2C} \quad (97)$$

choices for the iceberg $\mathcal{J}'$, where the inequality follows since $\text{diam}^*(\mathcal{J}) = \alpha + \gamma > x(2)$. Next, observe that

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J}')) \leq q(\text{diam}^{**}(\mathcal{J}')) \quad (98)$$

by $\text{IH}(s, x - 1)$, since $\text{diam}^*(\mathcal{J}') \leq \text{diam}^*(\mathcal{J}) - \gamma$, by (94), and $\text{diam}^*(\mathcal{J}) \leq x$.

We are now ready to bound the probability of the event E_2 . The reader may be concerned that we ought to lose here, because $\text{diam}^*(\mathcal{J}') < \text{diam}^*(\mathcal{J})$. We are saved, however, by the extra factor of δ^k in $\text{diam}^{**}(\mathcal{J})$ (compared with $\text{diam}^*(\mathcal{J})$), and the fact that v is strictly closer to $\mathbf{0}$ than u , which imply that $\text{diam}^{**}(\mathcal{J}') > \text{diam}^{**}(\mathcal{J})$.

Claim 8.10. $\mathbb{P}_p(E_2) \leq q(\text{diam}^{**}(\mathcal{J}))/3$.

Proof of Claim 8.10. Observe first that $d_T(\mathbf{0}, v) \leq k - 1$, since if $\mathcal{J}' = (J', v)$ is u -iceberg spanned then $v \in [0, u]_T$, and $v \neq u$. We therefore have $\text{diam}^{**}(\mathcal{J}') \geq \delta^{k-1}\alpha$, by (85). Since q is a decreasing function, it follows from (98) that

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J}')) \leq q(\text{diam}^{**}(\mathcal{J}')) \leq q(\delta^{k-1}\alpha).$$

Hence, taking a union bound over the choices of $\mathcal{J}'$, and recalling (97), we obtain

$$\mathbb{P}_p(E_2) \leq \sup_{\alpha, \gamma} (\delta^k(\alpha + \gamma))^{2C} \cdot q(\delta^{k-1}\alpha), \quad (99)$$

where the supremum is taken over valid values of α and γ .

We shall bound (99) using Lemma 8.7, applied to $\delta^{k-1}\alpha$ and $\delta^{k-1}\gamma$. The conditions of the lemma are satisfied because $\gamma \leq \gamma \leq 3\gamma$ and $\alpha \geq \text{diam}^*(\mathcal{J}) - 3\gamma$, by (94), and therefore $\alpha \geq \gamma \gg 1$, since $\gamma = (\log \text{diam}^*(\mathcal{J}))^{1+\eta}$ and $\text{diam}^*(\mathcal{J}) > x(2) \gg 1$. Thus, by (99) and Lemma 8.7, we have

$$\mathbb{P}_p(E_2) \leq \sup_{\alpha, \gamma} \frac{q(\delta^k(\alpha + \gamma))}{(\delta^k(\alpha + \gamma))^C} \leq \frac{q(\text{diam}^{**}(\mathcal{J}))}{3},$$

since $\text{diam}^{**}(\mathcal{J}) = \delta^k(\alpha + \gamma) \gg 1$. ■

The proof of the final claim is similar, but the details are somewhat more complicated because we also need to bound the probability of the event $\Delta_u(\mathcal{J}', \mathcal{J})$. This will be done by first applying Lemma 7.4 to deduce that there exists a (v, x) -spanned x -shifted iceberg $\mathcal{J}''$

for some $x \in \mathbb{R}^d$ and $v \in \{u\} \cup N_{\vec{T}}(u)$, and then by bounding the probability that such an iceberg exists by taking a union bound over choices of $\mathcal{J}''$. We shall use Lemma 6.21 and Remark 7.7 to bound the number of choices, and $\text{IH}(s', x(t-1))$ for some $s' \in \{s-1, s\}$ (which holds by (89)) to bound the probability that $\mathcal{J}''$ is (v, x) -spanned.

Claim 8.11. $\mathbb{P}_p(E_3) \leq q(\text{diam}^{**}(\mathcal{J}))/3$.

Proof of Claim 8.11. We shall again use (97) and (98) to count the possible icebergs $\mathcal{J}'$ and to bound the probability of the event $\text{Ice}_u^x(\mathcal{J}')$ for each of them. Our task will therefore be to bound the probability of the event $\Delta_u(\mathcal{J}', \mathcal{J})$ for fixed $\mathcal{J}$ and $\mathcal{J}' = (J', u)$ such that $J' \subset J$ and the bounds (94) hold.

Suppose that the event $\Delta_u(\mathcal{J}', \mathcal{J})$ holds, and apply Lemma 7.4. Observe that the conditions of the lemma are satisfied, since $d_T(\mathbf{0}, u) = r - s$ and $s \geq 3$, by assumption, and since $J' \subset J$ and $\gamma \geq y \geq \lambda$, by (94) and (96), and since $\text{diam}^*(\mathcal{J}) > x(2) \gg 1$. By Lemma 7.4, it follows that there exists $x \in \mathbb{R}^d$, an x -shifted iceberg $\mathcal{J}''$, and a set $K \subset (J \cap A) \setminus J'$, such that

$$\text{diam}_x^*(\mathcal{J}'') \geq \delta^3 \gamma \quad \text{and} \quad \langle K \rangle_{v,x} = \{\mathcal{J}''\}$$

for some $v \in \{u\} \cup N_{\vec{T}}(u)$. Recall that this implies that the event $\text{Ice}_{v,x}^x(\mathcal{J}'')$ holds.

To bound the probability that there exists such an iceberg, observe first that, by Lemma 6.18, we may assume that

$$\delta^3 \gamma \leq \text{diam}_x^*(\mathcal{J}'') \leq \gamma. \quad (100)$$

Indeed, if $\text{diam}_x^*(\mathcal{J}'') > \gamma$ then, by Lemma 6.18, there exists a set $K' \subset K$ and an x -shifted iceberg $\mathcal{J}'''$ with $\gamma/3 \leq \text{diam}_x^*(\mathcal{J}''') \leq \gamma$ such that $\langle K' \rangle_{v,x} = \{\mathcal{J}'''\}$.

Now, to count the number of choices for the iceberg $\mathcal{J}'' = (J'', w)$ with $J'' \cap J \neq \emptyset$ and $\text{diam}^*(\mathcal{J}'') \leq \gamma \leq \text{diam}^*(\mathcal{J})$, we would like to apply Lemma 6.21. The lemma gives a bound on the number of icebergs, not x -shifted icebergs; however, by Remark 7.7, x is a deterministic function of the icebergs $\mathcal{J}$ and $\mathcal{J}'$ (which were fixed above) and the vertex $v \in \{u\} \cup N_{\vec{T}}(u)$. There are therefore only $|N_{\vec{T}}(u)| + 1$ possible choices for x (given $\mathcal{J}$ and $\mathcal{J}'$), and so we have at most

$$(|N_{\vec{T}}(u)| + 1) \cdot (\text{diam}^*(\mathcal{J}))^C \leq (\delta^k(\alpha + \gamma))^{2C} \quad (101)$$

choices for the iceberg $\mathcal{J}''$, since $\text{diam}^*(\mathcal{J}) = \alpha + \gamma \gg 1$.

Next we use the induction hypothesis to bound the probability of the event $\text{Ice}_{v,x}^x(\mathcal{J}'')$. Set $s' := r - d_T(\mathbf{0}, v)$, and observe that $s' \in \{s-1, s\}$, since $v \in \{u\} \cup N_{\vec{T}}(u)$. Now, recalling from (91), (94), (96) and (100) that

$$\text{diam}_x^*(\mathcal{J}'') \leq \gamma \leq 3y = 3(\log \text{diam}^*(\mathcal{J}))^{1+\eta},$$

and that $\text{diam}^*(\mathcal{J}) \leq x(t) = \exp_{(t-2)}(p^{-\varepsilon(t)})$, it follows that

$$\text{diam}_x^*(\mathcal{J}'') \leq 3(\exp_{(t-3)}(p^{-\varepsilon(t)}))^{1+\eta} \leq \exp_{(t-3)}(p^{-\varepsilon(t-1)}) - 1 = x(t-1) - 1, \quad (102)$$

since $\varepsilon(t) < \varepsilon(t-1)$, and by our choice of η .

Now, recall from (89) that $\text{IH}(s-1, x(t-1))$ and $\text{IH}(s, x-1)$ both hold, and also that $x > x(t-1)$, by our choice of t . Hence $\text{IH}(s', x(t-1)-1)$ also holds, which together with (100) and (102) implies that

$$\mathbb{P}_p(\text{Ice}_{v,x}^\times(\mathcal{J}'')) \leq q(\text{diam}_x^{**}(\mathcal{J}'')) \leq q(\delta^{k+4}\gamma),$$

since q is decreasing and $d_T(\mathbf{0}, v) \leq k+1$, so $\text{diam}_x^{**}(\mathcal{J}'') \geq \delta^{k+1} \text{diam}_x^*(\mathcal{J}'') \geq \delta^{k+4}\gamma$.

By the union bound, and recalling (101), it follows that

$$\mathbb{P}_p(\Delta_u(\mathcal{J}', \mathcal{J})) \leq (\delta^k(\alpha + \gamma))^{2C} \cdot q(\delta^{k+4}\gamma). \quad (103)$$

Moreover, since $\text{diam}^{**}(\mathcal{J}') = \delta^k\alpha$ by (96) and because $\mathcal{J}'$ is an iceberg of type u , it follows from (98) that

$$\mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J}')) \leq q(\text{diam}^{**}(\mathcal{J}')) = q(\delta^k\alpha). \quad (104)$$

Now, for each fixed $\mathcal{J}'$, the events $\text{Ice}_u^\times(\mathcal{J}')$ and $\Delta_u(\mathcal{J}', \mathcal{J})$ are independent, because they depend on disjoint subsets of A . Hence, by the union bound, and using (97), (103) and (104), we obtain

$$\begin{aligned} \mathbb{P}_p(E_3) &\leq \sum_{\mathcal{J}'} \mathbb{P}_p(\text{Ice}_u^\times(\mathcal{J}')) \cdot \mathbb{P}_p(\Delta_u(\mathcal{J}', \mathcal{J})) \\ &\leq \sup_{\alpha, \gamma} (\delta^k(\alpha + \gamma))^{4C} \cdot q(\delta^k\alpha) \cdot q(\delta^{k+4}\gamma), \end{aligned} \quad (105)$$

where the sum is over all valid choices of $\mathcal{J}'$.

Finally, we bound the right-hand side of (105) using Lemma 8.8, applied to $\delta^k\alpha$ and $\delta^k\gamma$. To check that the conditions of the lemma are satisfied, recall from (91), (94) and (96) that $y = (\log(\alpha + \gamma))^{1+\eta}$ and $y \leq \gamma \leq 3y$. Moreover, $\alpha \gg \gamma \geq y \gg 1$, so by (105) and Lemma 8.8, we have

$$\mathbb{P}_p(E_3) \leq \sup_{\alpha, \gamma} \frac{q(\delta^k(\alpha + \gamma))}{(\delta^k(\alpha + \gamma))^C} \leq \frac{q(\text{diam}^{**}(\mathcal{J}))}{3}$$

as required, since $\text{diam}^{**}(\mathcal{J}) = \delta^k(\alpha + \gamma) \gg 1$. ■

As observed earlier, the event $\text{Ice}_u^\times(\mathcal{J})$ implies that the event $E_1 \cup E_2 \cup E_3$ holds, by Lemma 7.2, and therefore (90) follows from Claims 8.9–8.11. This completes the proof of the lemma. ■

Now that Proposition 8.4 has been proved, all that remains before we can complete the proof of Theorem 3.1 is to verify that if $[A]_{\mathcal{U}} = \mathbb{Z}_n^d$ then $\mathbb{Z}_n^d$ contains a $\mathbf{0}$ -iceberg spanned $\mathbf{0}$ -iceberg of diameter roughly $x(r)$. We therefore now switch our setting to the torus, and for the rest of the proof we assume that $A \subset \mathbb{Z}_n^d$.

The following Aizenman–Lebowitz-type lemma is not, strictly speaking, a consequence of Lemma 6.18, because of the wrap-around effect of the torus. However, since we only need the lemma for droplets, the proof is particularly simple.

Lemma 8.12. *Suppose that $[A]\mathcal{U} = \mathbb{Z}_n^d$ for some $n \gg x(r)$. Then there exists a $\mathbf{0}$ -iceberg spanned $\mathbf{0}$ -iceberg $\mathcal{J}$ with*

$$x(r)/3 \leq \text{diam}^*(\mathcal{J}) \leq x(r).$$

Proof. We run the $\mathbf{0}$ -iceberg spanning algorithm on $\mathbb{Z}_n^d$ with initial set A . Observe that, since $[A]_0 = [A]\mathcal{U} = \mathbb{Z}_n^d$, the algorithm cannot terminate before

$$\text{diam}^*(\mathcal{J}_0([K']\mathcal{U})) \geq x(r)$$

for some $t \in \mathbb{N}$ and $K' \in \mathcal{K}^t$. Now, by Lemma 6.14, the quantity

$$\max \{ \text{diam}^*(\mathcal{J}_0([K']\mathcal{U})) : K' \in \mathcal{K}^t \}$$

at most triples at each step, provided that this maximum is at least an absolute constant (cf. the proof of Lemma 6.18). It follows that for some $t \leq T$, there exists $K' \in \mathcal{K}^t$ such that $x(r)/3 \leq \text{diam}^*(\mathcal{J}) \leq x(r)$, where $\mathcal{J}_0([K']\mathcal{U}) = \{\mathcal{J}\}$, as required. ■

We are finally ready to complete the proof of Theorem 3.1.

Proof of Theorem 3.1. We are required to prove that for any $0 < \varepsilon < \varepsilon(r)$, if

$$p \leq (\log_{(r-1)} n)^{-1/\varepsilon}, \quad (106)$$

then $\mathbb{P}_p([A]\mathcal{U} = \mathbb{Z}_n^d) \rightarrow 0$ as $n \rightarrow \infty$. Since, by (82), the constant $\varepsilon(r)$ may be chosen arbitrarily close to $1/(d-1)$, this suffices to prove the theorem.

Let A be a p -random subset of $\mathbb{Z}_n^d$, and suppose that $[A]\mathcal{U} = \mathbb{Z}_n^d$. By Lemma 8.12, there exists a $\mathbf{0}$ -iceberg spanned $\mathbf{0}$ -iceberg $\mathcal{J}$ such that

$$\frac{1}{3} \cdot \exp_{(r-2)}(p^{-\varepsilon(r)}) \leq \text{diam}^*(\mathcal{J}) \leq \exp_{(r-2)}(p^{-\varepsilon(r)}).$$

Now, the number of $\mathbf{0}$ -icebergs in $\mathbb{Z}_n^d$ is $n^{O(1)}$, since we have at most n^d choices for each face, and each is $\mathbf{0}$ -iceberg spanned with probability at most

$$\exp(-\text{diam}^*(\mathcal{J})) \leq \exp\left(-\frac{1}{3} \cdot \exp_{(r-2)}(p^{-\varepsilon(r)})\right),$$

by Proposition 8.4. Thus, by (106) and since $\varepsilon < \varepsilon(r)$, we finally see that

$$\mathbb{P}_p([A] = \mathbb{Z}_n^d) \leq n^{O(1)} \cdot \exp\left(-\frac{1}{3} \cdot \exp_{(r-2)}(p^{-\varepsilon(r)})\right) \rightarrow 0$$

as $n \rightarrow \infty$, and this completes the proof of Theorem 3.1. ■

Appendix A. Simple calculations

This appendix contains the proofs of three simple bounds that were used in Section 8, namely Lemmas 8.6–8.8. For the reader's convenience, we restate each of these here, as Lemmas A.1–A.3, respectively. All constants used in this appendix should be interpreted

as they were used in Section 8. Recall that p is assumed to be sufficiently small, so in particular $\log(1/p) \gg C$.

Lemma A.1. *Let $r \geq 3$, and let $\alpha \geq \beta \gg 1$ and $\gamma \in \mathbb{R}$ satisfy*

$$\gamma \leq C(\log \alpha)^{1+\eta} \leq C^2 \beta.$$

Then

$$q(\alpha) \cdot q(\beta) \cdot (\alpha + \beta + \gamma)^{4C} \leq q(\alpha + \beta + \gamma). \quad (107)$$

Proof. The left-hand side of (107) is increasing in γ and the right-hand side is decreasing in γ (since q is decreasing), so we may assume that $\gamma = C(\log \alpha)^{1+\eta} > 0$. Recalling the definition (84) of $q(x)$, we are required to show that

$$\frac{\alpha}{\log_{(r-1)}(\lambda \alpha)} + \frac{\beta}{\log_{(r-1)}(\lambda \beta)} \geq \frac{\alpha + \beta + \gamma}{\log_{(r-1)}(\lambda(\alpha + \beta + \gamma))} + \frac{4C \log(\alpha + \beta + \gamma)}{\log(1/p)}. \quad (108)$$

To prove (108), observe first that if

$$\log_{(r-1)}(\lambda(\alpha + \beta + \gamma)) \geq C^3 \cdot \log_{(r-1)}(\lambda \beta), \quad (109)$$

then, since $\gamma \leq C^2 \beta$, we have

$$\frac{\beta}{\log_{(r-1)}(\lambda \beta)} - \frac{\beta + \gamma}{\log_{(r-1)}(\lambda(\alpha + \beta + \gamma))} \geq \frac{\beta}{2 \log_{(r-1)}(\lambda \beta)}.$$

Now, since $\beta \geq C^{-1}(\log \alpha)^{1+\eta} \gg 1$ and $r \geq 2$, we have

$$\frac{\beta}{\log_{(r-1)}(\lambda \beta)} \geq \log \alpha \gg \frac{\log(\alpha + \beta + \gamma)}{\log(1/p)},$$

so in this case (108) holds. On the other hand, if (109) fails to hold, then it follows that $\beta \geq (\log \alpha)^{2C}$, and therefore

$$\frac{\beta}{(\log \alpha)^C} \geq (\log \alpha)^C \gg \frac{\gamma}{\log_{(r-1)}(\lambda(\alpha + \beta + \gamma))} + \frac{4C \log(\alpha + \beta + \gamma)}{\log(1/p)},$$

since $\gamma \leq C(\log \alpha)^{1+\eta}$. It will therefore suffice to show that

$$\frac{\alpha}{\log_{(r-1)}(\lambda \alpha)} - \frac{\alpha}{\log_{(r-1)}(\lambda(\alpha + \beta + \gamma))} \geq \frac{\beta}{(\log \alpha)^C}. \quad (110)$$

To see this, note first that, by the mean value theorem, we have

$$\log_{(r-1)}(\lambda(\alpha + \beta + \gamma)) - \log_{(r-1)}(\lambda \alpha) = \frac{\beta + \gamma}{(\alpha + \xi) \log(\lambda(\alpha + \xi)) \cdots \log_{(r-2)}(\lambda(\alpha + \xi))}$$

for some $\xi \in (0, \beta + \gamma)$. Now, observe that the right-hand side is at least

$$\frac{\beta}{C^4 \alpha \cdot (\log \alpha)^{r-2}} \geq \frac{\beta}{\alpha \cdot (\log \alpha)^{C-3}},$$

since $\alpha + \xi < \alpha + \beta + \gamma \leq C^3 \alpha$ and $\gamma \geq 0$, and recalling that $\alpha \gg 1$ and C is a large constant. Rearranging, we obtain (110), and as noted above this proves (108). ■

Lemma A.2. *Let $\alpha \geq \gamma \gg 1$. Then*

$$q(\alpha) \cdot (\alpha + \gamma)^{3C} \leq q(\delta(\alpha + \gamma)).$$

Proof. Using (84), the conclusion of the lemma may be rewritten as

$$\frac{\alpha}{\log_{(r-1)}(\lambda\alpha)} \geq \frac{\delta(\alpha + \gamma)}{\log_{(r-1)}(\lambda\delta(\alpha + \gamma))} + \frac{3C \log(\alpha + \gamma)}{\log(1/p)}. \quad (111)$$

Observe that $\log_{(r-1)}(\delta\lambda\alpha) \geq 4\delta \cdot \log_{(r-1)}(\lambda\alpha)$, since $\alpha \gg 1$ and $\delta < 1/4$, and therefore

$$\frac{\alpha}{\log_{(r-1)}(\lambda\alpha)} - \frac{\delta(\alpha + \gamma)}{\log_{(r-1)}(\lambda\delta(\alpha + \gamma))} \geq \frac{\alpha}{2 \log_{(r-1)}(\lambda\alpha)},$$

since $\alpha \geq \gamma$. The bound (111) now follows easily. $\blacksquare$

Lemma A.3. *Let $r \geq 3$ and $\alpha, \gamma \gg 1$ be such that*

$$C^{-1}(\log(\alpha + \gamma))^{1+\eta} \leq \gamma \leq C(\log(\alpha + \gamma))^{1+\eta}. \quad (112)$$

Then

$$q(\alpha) \cdot q(\delta^4\gamma) \cdot (\alpha + \gamma)^{5C} \leq q(\alpha + \gamma).$$

Proof. Using (84), the conclusion of the lemma may be rewritten as

$$\frac{\alpha}{\log_{(r-1)}(\lambda\alpha)} + \frac{\delta^4\gamma}{\log_{(r-1)}(\lambda\delta^4\gamma)} \geq \frac{\alpha + \gamma}{\log_{(r-1)}(\lambda(\alpha + \gamma))} + \frac{5C \log(\alpha + \gamma)}{\log(1/p)}.$$

Since $\gamma > 0$ and $5C < \log(1/p)$, this will follow if we can prove that

$$\frac{\delta^4\gamma}{\log_{(r-1)}(\lambda\delta^4\gamma)} - \frac{\gamma}{\log_{(r-1)}(\lambda(\alpha + \gamma))} \geq \log(\alpha + \gamma). \quad (113)$$

To prove (113), note that, since $\alpha, \gamma \gg 1$, it follows from (112) that

$$\delta^4 \cdot \log_{(r-1)}(\lambda(\alpha + \gamma)) \geq 2 \cdot \log_{(r-1)}(\lambda\delta^4\gamma),$$

and therefore

$$\frac{\delta^4\gamma}{\log_{(r-1)}(\lambda\delta^4\gamma)} - \frac{\gamma}{\log_{(r-1)}(\lambda(\alpha + \gamma))} \geq \frac{\delta^4\gamma}{2 \cdot \log_{(r-1)}(\lambda\delta^4\gamma)}.$$

Since $\gamma \geq C^{-1}(\log(\alpha + \gamma))^{1+\eta}$, by (112), we obtain (113), as required. $\blacksquare$

Acknowledgements. The authors would like to thank Hugo Duminil-Copin for a very helpful discussion at the outset of this project, and for many other interesting conversations over the years.

Funding. P.B. and B.B. were partially supported by NSF grant DMS 1855745, R.M. by FAPERJ (Proc. E-26/202.993/2017) and CNPq (Proc. 303681/2020-9), and by ERC Starting Grant 680275 MALIG, and P.S. by ERC Starting Grant 676970, by ISF grants 1147/14 and 1207/15, and by a CNPq bolsa PDJ.

References

- [1] Aizenman, M., Duminil-Copin, H.: [Marginal triviality of the scaling limits of critical 4D Ising and \$\phi_4^4\$ models](#). Ann. of Math. (2) **194**, 163–235 (2021) Zbl [1489.60150](#) MR [4276286](#)
- [2] Aizenman, M., Duminil-Copin, H., Sidoravicius, V.: [Random currents and continuity of Ising model's spontaneous magnetization](#). Comm. Math. Phys. **334**, 719–742 (2015) Zbl [1315.82004](#) MR [3306602](#)
- [3] Aizenman, M., Duminil-Copin, H., Tassion, V., Warzel, S.: [Emergent planarity in two-dimensional Ising models with finite-range interactions](#). Invent. Math. **216**, 661–743 (2019) Zbl [1417.82022](#) MR [3955708](#)
- [4] Aizenman, M., Lebowitz, J. L.: Metastability effects in bootstrap percolation. J. Phys. A **21**, 3801–3813 (1988) Zbl [0656.60106](#) MR [0968311](#)
- [5] Arceri, F., Landes, P. F., Berthier, L., Biroli, G.: [A statistical mechanics perspective on glasses and aging](#). In: Encyclopedia of complexity and systems science, Springer, 1–68 (2022)
- [6] Balister, P., Bollobás, B., Morris, R., Smith, P.: [Subcritical monotone cellular automata](#). Random Structures Algorithms **64**, 38–61 (2024) Zbl [1536.81030](#) MR [4672996](#)
- [7] Balister, P., Bollobás, B., Morris, R., Smith, P.: The critical length for growing a droplet. Mem. Amer. Math. Soc. (to appear)
- [8] Balister, P., Bollobás, B., Morris, R., Smith, P.: Uncomputability of critical probabilities for monotone cellular automata. In preparation
- [9] Balister, P., Bollobás, B., Przykucki, M., Smith, P.: [Subcritical \$\mathcal{U}\$ -bootstrap percolation models have non-trivial phase transitions](#). Trans. Amer. Math. Soc. **368**, 7385–7411 (2016) Zbl [1342.60159](#) MR [3471095](#)
- [10] Balogh, J., Bollobás, B., Duminil-Copin, H., Morris, R.: [The sharp threshold for bootstrap percolation in all dimensions](#). Trans. Amer. Math. Soc. **364**, 2667–2701 (2012) Zbl [1238.60108](#) MR [2888224](#)
- [11] Balogh, J., Bollobás, B., Morris, R.: [Bootstrap percolation in three dimensions](#). Ann. Probab. **37**, 1329–1380 (2009) Zbl [1187.60082](#) MR [2546747](#)
- [12] Balogh, J., Bollobás, B., Morris, R.: [Majority bootstrap percolation on the hypercube](#). Combin. Probab. Comput. **18**, 17–51 (2009) Zbl [1198.60041](#) MR [2497373](#)
- [13] Blanquicett, D.: [Anisotropic bootstrap percolation in three dimensions](#). Ann. Probab. **48**, 2591–2614 (2020) Zbl [1456.60250](#) MR [4152652](#)
- [14] Blanquicett, D.: [Fixation for two-dimensional \$\mathcal{U}\$ -Ising and \$\mathcal{U}\$ -voter dynamics](#). J. Statist. Phys. **182**, article no. 21, 24 pp. (2021) Zbl [1469.60307](#) MR [4205282](#)
- [15] Blanquicett, D.: [The \$d\$ -dimensional bootstrap percolation models with axial neighbourhoods](#). Stochastic Process. Appl. **174**, article no. 104383, 13 pp. (2024) Zbl [1554.60117](#) MR [4746579](#)
- [16] Bollobás, B., Duminil-Copin, H., Morris, R., Smith, P.: [The sharp threshold for the Duarte model](#). Ann. Probab. **45**, 4222–4272 (2017) Zbl [1454.60141](#) MR [3737910](#)
- [17] Bollobás, B., Duminil-Copin, H., Morris, R., Smith, P.: [Universality for two-dimensional critical cellular automata](#). Proc. London Math. Soc. (3) **126**, 620–703 (2023) Zbl [1538.60156](#) MR [4550150](#)
- [18] Bollobás, B., Smith, P., Uzzell, A.: [Monotone cellular automata in a random environment](#). Combin. Probab. Comput. **24**, 687–722 (2015) Zbl [1371.60170](#) MR [3350030](#)
- [19] Bovier, A., den Hollander, F.: [Metastability](#). Grundlehren Math. Wiss. 351, Springer, Cham (2015) Zbl [1339.60002](#) Zbl [1339.60002](#) MR [3445787](#)
- [20] Cancrini, N., Martinelli, F., Roberto, C., Toninelli, C.: [Kinetically constrained spin models](#). Probab. Theory Related Fields **140**, 459–504 (2008) Zbl [1139.60343](#) MR [2365481](#)
- [21] Cancrini, N., Martinelli, F., Roberto, C., Toninelli, C.: [Kinetically constrained models](#). In: New trends in mathematical physics, Springer, 741–752 (2009) Zbl [1176.82019](#)

- [22] Cerf, R., Cirillo, E. N. M.: [Finite size scaling in three-dimensional bootstrap percolation](#). *Ann. Probab.* **27**, 1837–1850 (1999) Zbl 0960.60088 MR 1742890
- [23] Cerf, R., Manzo, F.: [The threshold regime of finite volume bootstrap percolation](#). *Stochastic Process. Appl.* **101**, 69–82 (2002) Zbl 1075.82010 MR 1921442
- [24] Cerf, R., Manzo, F.: [Nucleation and growth for the Ising model in \$d\$ dimensions at very low temperatures](#). *Ann. Probab.* **41**, 3697–3785 (2013) Zbl 1286.60091 MR 3161463
- [25] Chalupa, J., Leath, P. L., Reich, G. R.: [Bootstrap percolation on a Bethe lattice](#). *J. Phys. C* **12**, L31–L35 (1979)
- [26] Chelkak, D., Duminil-Copin, H., Hongler, C., Kemppainen, A., Smirnov, S.: [Convergence of Ising interfaces to Schramm’s SLE curves](#). *C. R. Math. Acad. Sci. Paris* **352**, 157–161 (2014) Zbl 1294.82007 MR 3151886
- [27] Chelkak, D., Smirnov, S.: [Universality in the 2D Ising model and conformal invariance of fermionic observables](#). *Invent. Math.* **189**, 515–580 (2012) Zbl 1257.82020 MR 2957303
- [28] Chleboun, P., Faggionato, A., Martinelli, F.: [Time scale separation and dynamic heterogeneity in the low temperature East model](#). *Comm. Math. Phys.* **328**, 955–993 (2014) Zbl 1294.82026 MR 3201217
- [29] Debenedetti, P. G., Stillinger, F. H.: [Supercooled liquids and the glass transition](#). *Nature* **410**, 259–267 (2001)
- [30] Dehghanpour, P., Schonmann, R. H.: [Metropolis dynamics relaxation via nucleation and growth](#). *Comm. Math. Phys.* **188**, 89–119 (1997) Zbl 0882.60094 MR 1471333
- [31] Duminil-Copin, H.: Random currents expansion of the Ising model. In: *European Congress of Mathematics (Berlin, 2016)*, European Mathematical Society, Zürich, 869–889 (2018) Zbl 1403.82005 MR 3890455
- [32] Duminil-Copin, H.: [Lectures on the Ising and Potts models on the hypercubic lattice](#). In: *Random graphs, phase transitions, and the Gaussian free field*, Springer Proc. Math. Statist. 304, Springer, 35–161 (2020) Zbl 1447.82007 MR 4043224
- [33] Duminil-Copin, H., Goswami, S., Raoufi, A.: [Exponential decay of truncated correlations for the Ising model in any dimension for all but the critical temperature](#). *Comm. Math. Phys.* **374**, 891–921 (2020) Zbl 1475.82003 MR 4072233
- [34] Duminil-Copin, H., Hartarsky, I.: [Sharp metastability transition for two-dimensional bootstrap percolation with symmetric isotropic threshold rules](#). *Probab. Theory Related Fields* **190**, 445–483 (2024) Zbl 07924625 MR 4797373
- [35] Duminil-Copin, H., Holroyd, A.: Finite volume bootstrap percolation with threshold rules on $\mathbb{Z}^2$: balanced case. Unpublished manuscript (2012)
- [36] Duminil-Copin, H., van Enter, A. C. D.: [Sharp metastability threshold for an anisotropic bootstrap percolation model](#). *Ann. Probab.* **41**, 1218–1242 (2013) Zbl 1356.60166 MR 3098677
- [37] Duminil-Copin, H., van Enter, A. C. D., Hulshof, T.: [Higher order corrections for anisotropic bootstrap percolation](#). *Probab. Theory Related Fields* **172**, 191–243 (2018) Zbl 1404.60149 MR 3851832
- [38] Fontes, L. R., Schonmann, R. H., Sidoravicius, V.: [Stretched exponential fixation in stochastic Ising models at zero temperature](#). *Comm. Math. Phys.* **228**, 495–518 (2002) Zbl 1004.82013 MR 1918786
- [39] Fredrickson, G. H., Andersen, H. C.: [Kinetic Ising model of the glass transition](#). *Phys. Rev. Lett.* **53**, 1244–1247 (1984)
- [40] Friedli, S., Velenik, Y.: *Statistical mechanics of lattice systems*. Cambridge University Press, Cambridge (2018) Zbl 1407.82001 MR 3752129
- [41] Garrahan, J. P., Sollich, P., Toninelli, C.: [Kinetically constrained models](#). In: *Dynamical heterogeneities in glasses, colloids, and granular media*, Oxford University Press, 341–369 (2011)
- [42] Gravner, J., Griffeath, D.: [Threshold growth dynamics](#). *Trans. Amer. Math. Soc.* **340**, 837–870 (1993) Zbl 0791.58053 MR 1147400

- [43] Gravner, J., Holroyd, A. E., Morris, R.: [A sharper threshold for bootstrap percolation in two dimensions](#). Probab. Theory Related Fields **153**, 1–23 (2012) Zbl 1254.60092 MR 2925568
- [44] Hartarsky, I.: [Refined universality for critical KCM: upper bounds](#). Comm. Math. Phys. **405**, article no. 13, 68 pp. (2024) Zbl 07798465 MR 4694420
- [45] Hartarsky, I., Marêché, L.: [Refined universality for critical KCM: lower bounds](#). Combin. Probab. Comput. **31**, 879–906 (2022) Zbl 1507.60130 MR 4472293
- [46] Hartarsky, I., Marêché, L., Toninelli, C.: [Universality for critical KCM: infinite number of stable directions](#). Probab. Theory Related Fields **178**, 289–326 (2020) Zbl 1464.60081 MR 4146539
- [47] Hartarsky, I., Martinelli, F., Toninelli, C.: [Universality for critical KCM: finite number of stable directions](#). Ann. Probab. **49**, 2141–2174 (2021) Zbl 1490.60272 MR 4317702
- [48] Hartarsky, I., Martinelli, F., Toninelli, C.: [Sharp threshold for the FA-2f kinetically constrained model](#). Probab. Theory Related Fields **185**, 993–1037 (2023) Zbl 1523.60166 MR 4556287
- [49] Hartarsky, I., Mezei, T. R.: [Complexity of two-dimensional bootstrap percolation difficulty: algorithm and NP-hardness](#). SIAM J. Discrete Math. **34**, 1444–1459 (2020) Zbl 1462.68070 MR 4117299
- [50] Hartarsky, I., Morris, R.: [The second term for two-neighbour bootstrap percolation in two dimensions](#). Trans. Amer. Math. Soc. **372**, 6465–6505 (2019) Zbl 1455.60130 MR 4024528
- [51] Hartarsky, I., Szabó, R.: [Subcritical bootstrap percolation via Toom contours](#). Electron. Commun. Probab. **27**, article no. 55, 13 pp. (2022) Zbl 1515.60313 MR 4510850
- [52] Hartarsky, I., Teixeira, A.: [Bootstrap percolation is local](#). arXiv:2404.07903v1 (2024)
- [53] Holroyd, A. E.: [Sharp metastability threshold for two-dimensional bootstrap percolation](#). Probab. Theory Related Fields **125**, 195–224 (2003) Zbl 1042.60065 MR 1961342
- [54] Holroyd, A. E.: [The metastability threshold for modified bootstrap percolation in \$d\$ dimensions](#). Electron. J. Probab. **11**, no. 17, 418–433 (2006) Zbl 1112.60080 MR 2223042
- [55] Holroyd, A. E., Liggett, T. M., Romik, D.: [Integrals, partitions, and cellular automata](#). Trans. Amer. Math. Soc. **356**, 3349–3368 (2004) Zbl 1095.60003 MR 2052953
- [56] Lubetzky, E., Sly, A.: [Cutoff for the Ising model on the lattice](#). Invent. Math. **191**, 719–755 (2013) Zbl 1273.82014 MR 3020173
- [57] Lubetzky, E., Sly, A.: [Information percolation and cutoff for the stochastic Ising model](#). J. Amer. Math. Soc. **29**, 729–774 (2016) Zbl 1342.60173 MR 3486171
- [58] Marêché, L., Martinelli, F., Toninelli, C.: [Exact asymptotics for Duarte and supercritical rooted kinetically constrained models](#). Ann. Probab. **48**, 317–342 (2020) Zbl 1451.60106 MR 4079438
- [59] Martinelli, F.: [Lectures on Glauber dynamics for discrete spin models](#). In: Lectures on probability theory and statistics (Saint-Flour, 1997), Lecture Notes in Math. 1717, Springer, Berlin, 93–191 (1999) Zbl 1051.82514 MR 1746301
- [60] Martinelli, F., Morris, R., Toninelli, C.: [Universality results for kinetically constrained spin models in two dimensions](#). Comm. Math. Phys. **369**, 761–809 (2019) Zbl 1419.82037 MR 3962008
- [61] Martinelli, F., Toninelli, C.: [Towards a universality picture for the relaxation to equilibrium of kinetically constrained models](#). Ann. Probab. **47**, 324–361 (2019) Zbl 1466.60210 MR 3909971
- [62] McCulloch, W. S., Pitts, W.: [A logical calculus of the ideas immanent in nervous activity](#). Bull. Math. Biophys. **5**, 115–133 (1943) Zbl 0063.03860 MR 0010388
- [63] Morris, R.: [Zero-temperature Glauber dynamics on \$\mathbb{Z}^d\$](#) . Probab. Theory Related Fields **149**, 417–434 (2011) Zbl 1236.60095 MR 2776621
- [64] Morris, R.: [Bootstrap percolation, and other automata](#). Eur. J. Combin. **66**, 250–263 (2017) Zbl 1376.82090 MR 3692148
- [65] Mountford, T. S.: [Critical length for semi-oriented bootstrap percolation](#). Stochastic Process. Appl. **56**, 185–205 (1995) Zbl 0821.60092 MR 1325219

- [66] Olivieri, E., Vares, M. E.: [Large deviations and metastability](#). Encyclopedia Math. Appl. 100, Cambridge University Press (2005) Zbl [1075.60002](#) MR [2123364](#)
- [67] Ritort, F., Sollich, P.: [Glassy dynamics of kinetically constrained models](#), Adv. Phys. **52**, 219–342 (2003)
- [68] Schonmann, R. H.: On the behavior of some cellular automata related to bootstrap percolation. Ann. Probab. **20**, 174–193 (1992) Zbl [0742.60109](#) MR [1143417](#)
- [69] Schonmann, R. H.: Slow droplet-driven relaxation of stochastic Ising models in the vicinity of the phase coexistence region. Comm. Math. Phys. **161**, 1–49 (1994) Zbl [0796.60103](#) MR [1266068](#)
- [70] Shannon, C. E.: [Von Neumann’s contributions to automata theory](#). Bull. Amer. Math. Soc. **64**, 123–129 (1958) Zbl [0080.00418](#) MR [0092240](#)
- [71] Smirnov, S.: [Conformal invariance in random cluster models. I. Holomorphic fermions in the Ising model](#). Ann. of Math. (2) **172**, 1435–1467 (2010) Zbl [1200.82011](#) MR [2680496](#)
- [72] Toninelli, C.: Interacting particle systems with kinetic constraints. SpringerBriefs in Mathematical Physics (to appear)
- [73] Ulam, S.: Random processes and transformations. In: Proceedings of the International Congress of Mathematicians (Cambridge, MA, 1950), Vol. 2, American Mathematical Society, Providence, RI, 264–275 (1952) Zbl [0049.09511](#) MR [0045334](#)
- [74] Ulam, S.: [John von Neumann, 1903–1957](#). Bull. Amer. Math. Soc. **64**, 1–49 (1 plate) (1958) Zbl [0080.00411](#) MR [0091904](#)
- [75] van den Berg, J., Kesten, H.: [Inequalities with applications to percolation and reliability](#). J. Appl. Probab. **22**, 556–569 (1985) Zbl [0571.60019](#) MR [0799280](#)
- [76] van Enter, A. C. D., Fey, A.: [Metastability thresholds for anisotropic bootstrap percolation in three dimensions](#). J. Statist. Phys. **147**, 97–112 (2012) Zbl [1243.82038](#) MR [2922761](#)
- [77] von Neumann, J.: The general and logical theory of automata. In: Cerebral mechanisms in behavior: the Hixon symposium, Wiley, New York, 1–31; discussion, 32–41 (1951) MR [0045446](#)
- [78] von Neumann, J.: [Theory of self-reproducing automata](#). Univ. of Illinois Press, Champaign, 1966