

Largest minimal percolating sets in hypercubes under 2-bootstrap percolation

Eric Riedl

University of Notre Dame
Department of Mathematics

ebriedl@gmail.com

Submitted: Oct 10, 2009; Accepted: May 18, 2010; Published: May 25, 2010

Mathematics Subject Classification: 05D99

Abstract

Consider the following process, known as r -bootstrap percolation, on a graph G . Designate some initial infected set A and infect any vertex with at least r infected neighbors, continuing until no new vertices can be infected. We say A percolates if it eventually infects the entire graph. We say A is a minimal percolating set if A percolates, but no proper subset percolates. We compute the size of a largest minimal percolating set for $r = 2$ in the n -dimensional hypercube.

1 Introduction

In this paper, we consider the following process, known as r -bootstrap percolation. Designate an initial set A of infected vertices. Let $A_0 = A$. Then let A_t be the set of vertices in A_{t-1} union the set of vertices which have at least r neighbors in A_{t-1} . Set $\langle A \rangle = \cup_i A_i$, and call $\langle A \rangle$ the set of vertices infected by A . A set A percolates if it infects the entire graph. A percolating set A is said to be *minimal* if for all $v \in A$ the set $A \setminus v$ does not percolate. Let $E(G, r)$ be the largest size of a minimal percolating set and let $m(G, r)$ be the smallest size of a (necessarily minimal) percolating set. In this paper, we find $E(Q_n, 2)$, where Q_n is the n -dimensional hypercube and we use similar techniques to find bounds on $E([n]^d, 2)$ for all n and d . Since $r = 2$ for most of this paper, we write $E(G)$ for $E(G, 2)$ without ambiguity.

Bootstrap percolation was introduced in 1979 by Chalupa, Leath, and Reich [9] for its applications to dilute magnetic systems. For more information on the many physical applications of bootstrap percolation, see the survey article by Adler and Lev [1]. Arising naturally from the physical context is the following probabilistic problem. Let each vertex of G be initially infected independently with probability p . Then what is the probability that such a set percolates as a function of p ? In particular, if A is a randomly chosen

set, what is $p_c(G, r) = \inf\{p \mid \mathbb{P}(A \text{ percolates}) \geq 1/2\}$? Much work has been done on this question. Aizenmann and Lebowitz [2] and Cerf and Cirillo [7] did foundational work towards computing $p_c([n]^d, r)$ where $[n]^d$ is the $n \times \cdots \times n$ d -dimensional grid. Cerf and Manzo [8] proved that

$$p_c([n]^d, r) = \Theta\left(\frac{1}{\log^{(r-1)} n}\right)^{d-r+1},$$

where $\log^{(r)}(x)$ is $\log(\log(\cdots \log(x)))$ (r times). More precise asymptotics were found by Holroyd for $r = 2, d = 2$ [10] and Balogh, Bollobás, Duminil-Copin and Morris [4, 5] for general r and d . Balogh, Peres and Pete [6] determined p_c for infinite trees and relate it to the branching order.

Considerably less work has been done on finding $m(G, r)$ and $E(G, r)$. For $r \leq d$, it is known that

$$n^{r-1} \leq m([n]^d, r) \leq \frac{d^{r-1}}{r!} n^{r-1},$$

where the lower bound follows by Pete [12] and the upper bound by the method of Schonmann [14]. Balogh and Bollobás [3] prove that $m([n]^d, 2) = \lceil d(n-1)/2 \rceil$. Pete [12] finds an exact asymptotic for $m([n]^d, r)$ when $r = d$. Morris [11] shows that $\frac{4n^2}{33} \leq E([n]^2, 2) \leq \frac{n^2}{6}$ asymptotically, making progress on a question posed by Bollobás. In [13], an algorithm is presented for finding $m(T, r)$ and $E(T, r)$ for all finite trees T , and it is shown that if T is a finite tree with ℓ leaves, $m(T, r) \geq \frac{(r-1)|T|+1}{r}$, $E(T, r) \leq \frac{r|T|+(r-1)\ell}{r+1}$ and $\frac{E(T, r)-m(T, r)}{|T|} < \frac{r-1}{r^2}$. In this paper, we find $E(Q_n, 2)$ exactly, and show it to be on the order of $2^{n/4}$.

First we set some notation. We can represent the vertices of Q_n as strings of 0s and 1s of length n , with adjacent vertices being precisely those vertices which differ from each other in exactly one coordinate. We can also represent the vertices of the hypercube as the possible subsets of the set $\{1, \dots, n\}$. Recall that the automorphisms of the hypercube are all combinations of the $n!$ permutations of the dimensions and the 2^n reflections. We say that two subsets of the hypercube are isomorphic if there is an automorphism of the hypercube that takes one of them to the other.

In this paper, we prove the following main result.

Theorem 1. *Let $1 \leq r \leq 4$ be such that $n \equiv r \pmod{4}$. Then*

$$E(Q_n, 2) = \begin{cases} n+1 & 0 \leq n \leq 1 \\ n & 2 \leq n \leq 10 \\ (1+2^{r-4})2^{\lfloor \frac{n+3}{4} \rfloor} & n \geq 11 \end{cases}$$

Note that $\frac{E(Q_n, 2)}{E(Q_{n-1}, 2)}$ does not converge as $n \rightarrow \infty$, as it simply cycles around between four different values for large n .

For the case of grids, we modify our techniques to obtain the following result.

Corollary 2. *We have $E([n]^d, 2) \leq \lceil \frac{n}{2} \rceil^d (2 + \lfloor \frac{j}{3} \rfloor) 2^{\lfloor \frac{d-1}{3} \rfloor}$ where $d \equiv j \pmod{3}$, $1 \leq j \leq 3$.*

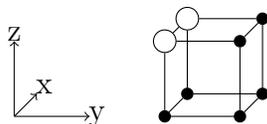


Figure 1: The subcube *01.

Note that combining Corollary 2 with a result coming from the proof of Theorem 14 of [11] we obtain

$$\left(\frac{1}{4}\right)^d \leq E([n]^d, 2) \leq \left(\frac{1}{2}\right)^{2d/3} n^d.$$

Note that this shows that $E([n]^d, 2) = o(n^d)$ if and only if $d = d(n) \rightarrow \infty$ as $n \rightarrow \infty$.

In Section 2 we review some basic facts about 2-neighbor bootstrap percolation on $[n]^d$. In Section 3 we describe a construction that is optimal in small dimensions and give a recursive upper bound for $E(Q_n, 2)$. In Section 4 we describe a construction that is optimal for higher dimensions, and prove optimality by classifying all of the isomorphism classes of largest minimal percolating sets. This gives our main result. In Section 5 we show $E([n]^d, 2) = O(2^{n/3})$ for all fixed n and $E(AQ_n, 2) = 2$ for the augmented hypercube AQ_n .

2 Basic facts about 2-percolation in hypercubes

Before proceeding, we summarize some basic definitions and facts about 2-percolation in Q_n and more generally, grids $P_{n_1} \times \cdots \times P_{n_d}$. The goal of this section is to obtain a description of the percolation process in terms of combining subcubes. The material in this section was proven by Balogh and Bollobás [3]. We say a set S is *closed* under percolation if $\langle S \rangle = S$. We call a subgraph G of a grid a *subgrid* if G is itself a grid. We call a subgraph G of a grid or hypercube a *subcube* if G is a hypercube.

Proposition 3. *The only subsets of the grid which are closed under percolation are those which are a union of disjoint subgrids that are distance at least three from each other.*

In the case of hypercubes, the only subgrids of Q_n are sub-hypercubes. We represent subcubes of Q_n as strings of 0's, 1's and *'s, where an * in position i means that the subcube contains vertices with both 0 and 1 in that position. In particular, the number of *'s is the dimension of the subcube. See Figure 1 for an example. We define the k th coordinate of the subcube to be the k th element of the string.

Proposition 4. *Let A and B be two subgrids of distance at most 2 from each other in a grid G . Then $\langle A \cup B \rangle$ is the smallest subgrid containing both A and B . Moreover, in the case $G = Q_n$ if A has coordinates a_1, \dots, a_n and B has coordinates b_1, \dots, b_n where $a_i, b_i \in \{0, 1, *\}$, then the coordinates of $\langle A \cup B \rangle$ are $a_i \vee b_i$, where $x \vee x = x$ and $x \vee y = *$ if $x \neq y$.*



Figure 2: Two different minimal percolating sets of size 3 in dimension 3.

Definition 5. Let A be given, and write $A = \cup_i C_i$, where each C_i is a set containing a single point, which is just a 0-dimensional subcube. Set $\mathcal{A}_0 = \cup_i \{C_i\}$. Then choose a sequence of sets of subgrids $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_k$ so that \mathcal{A}_t is identical to \mathcal{A}_{t-1} except that two subgrids $B, C \in \mathcal{A}_{t-1}$ within distance 2 of each other are replaced by the subgrid $\langle B \cup C \rangle$. Require \mathcal{A}_k to consist of a set of subgrids all of which are distance at least 3 from each other. Then $\mathcal{A}_0, \dots, \mathcal{A}_k$ is called an execution path of the percolation process.

For any execution path, we know that $\mathcal{A}_k = \{\langle A \rangle\}$, so \mathcal{A}_k is independent of execution path. We say a subset S of G is *internally spanned* by A if $\langle A \cap S \rangle = S$. Then each $B \in \mathcal{A}_i$ is internally spanned by the vertices that contributed to B in the execution path. Note that the two subgrids B and C which we combine at each step are not necessarily disjoint.

Proposition 6. Any percolating set of size at least 2 in Q_n will disjointly internally span two subcubes which together span the entire hypercube.

Proof. Choose an execution path and take the two hypercubes in \mathcal{A}_{k-1} . □

3 An initial construction and an upper bound

In this section, we give a simple lower bound that is sharp in low dimensions and a recursive upper bound, which is the key to our entire argument. First we give a simple construction to give an easy lower bound for $E(Q_n)$.

Proposition 7. Let $A = \{100\dots 0, 010\dots 0, 001\dots 0, \dots, 000\dots 01\}$. Then A is a minimal percolating set of size n for $n \geq 2$. Thus, $E(Q_n) \geq n$.

Proof. The set A clearly percolates, and if v_i is the vertex with a 1 in the i th coordinate, then $\langle A \setminus v_i \rangle$ will be the Q_{n-1} with a 0 in the i th coordinate. □

Note that for A given as in Proposition 7, $\langle A \setminus v \rangle$ will simply be a Q_{n-1} containing the empty set. Moreover, by changing v , we can make $\langle A \setminus v \rangle$ range over all such Q_{n-1} . Thus, for all v , $\langle A \setminus v \rangle$ is as large as it can be given that A is a minimal percolating set. It turns out that this property will complicate our goal of finding $E(Q_n)$ for higher dimensions and that this construction does not give the largest possible $E(Q_n)$ for these dimensions. However, it also will turn out that for $2 \leq n \leq 10$ this construction is optimal. See Section 4 for more details.

Before proving our recursive upper bound, we need a simple lemma. It is the analog of Lemma 7 of [11].

Lemma 8. *We have $E(Q_n) \geq E(Q_{n-1})$.*

Proof. For $n = 1$, the statement is obvious since $E(Q_0) = 1$ and $E(Q_1) = 2$.

Now suppose $n \geq 2$, and let A be a minimal percolating set in Q_{n-1} of largest possible size. We shall construct a minimal percolating set in Q_n of size at least $|A|$. Let P be a fixed sub- Q_{n-1} contained in Q_n , and view A as a subset of $P \subset Q_n$. Now, select a vertex $v \in A$. Let w be the unique neighbor of v which does not lie in P . Then $A \cup w$ percolates in Q_n .

We claim that $A \cup w \setminus u$ does not percolate for every $u \in A$ with $u \neq v$. This will complete the proof, since it will show that either $A \cup w$ is a minimal percolating set, or $A \cup w \setminus v$ is a minimal percolating set. By minimality of A , we know that $B = \langle A \setminus u \rangle$ will be a union of subcubes of distance at least 3 from each other and will have $P \setminus B$ nonempty. Since $v \in B$, $\langle B \cup w \rangle \cap P = B$, so $A \setminus u$ does not percolate. This concludes the proof. \square

We now turn our attention to proving a recursive upper bound. The general argument is very similar to an argument found in the proof of the upper bound of Theorem 11 in [3]. However, we include it here because we extract extra information from the proof. The proof relies heavily on the idea of viewing percolation as combining nearby subcubes, and it looks at the ways that the last two cubes in the process can be combined.

Proposition 9. *We have $E(Q_n) \leq \max\{E(Q_{n-1}) + 1, 2E(Q_{n-4})\}$.*

Proof. Let $A \subset Q_n$ be a minimal percolating set of size $E(Q_n)$. Since A percolates, we know that in any execution path, the final term \mathcal{A}_k will contain only Q_n itself. Hence, the penultimate term, \mathcal{A}_{k-1} will always consist of exactly two subcubes, say P and R , which together infect Q_n . Without loss of generality, let $\dim P \geq \dim R$. Now, $\dim P \leq n - 1$ by minimality of A . Among all execution paths, choose one which has $\dim P$ as large as possible. We divide into cases depending on $\dim P$.

Case 1 $\dim P = n - 1$. Then there must be a vertex of A outside of P , and that vertex plus $A \cap P$ will percolate, so R is simply a single point by minimality of A . Hence A is the union of one vertex and a set which minimally internally spans P , so $E(G) = |A| \leq E(Q_{n-1}) + 1$ in this case.

Case 2 $\dim P = n - 2$. Then we know that there cannot be a vertex of $A \cup R$ in $\{v \in Q_n \mid d(v, P) \leq 1\}$, since otherwise we could extend P to a cube of dimension $n - 1$. Thus, there must be a vertex v of A which has distance 2 from P (since every vertex has distance at most 2 from P) and $\langle P \cup v \rangle = Q_n$ (just write out the coordinates of P and v in the 0, 1, * notation). Hence, $|A| \leq E(Q_{n-2}) + 1$ in this case.

Case 3 $\dim P = n - 3$. Then we know that there cannot be a vertex of $A \cap R$ within distance 2 of P , as this would contradict maximality of $\dim P$. Hence, $A \cap R$ is contained in a subcube of Q_n of distance 3 from P , as the set of vertices which are distance 3 from P is a subcube of dimension $n - 3$. To see this, note that if, for

example, $P = 000* \dots *$, then the set of vertices of distance 3 from P is just $111* \dots *$. Thus, R is contained in this subcube, so $d(P, R) = 3$, which contradicts the fact that A percolates. Hence, this case cannot occur.

Case 4 $\dim P \leq n - 4$. Then by choice of R , we have $\dim R \leq n - 4$ as well. Now, P and R are both minimally internally spanned, so $|A \cap P|$ and $|A \cap R|$ are each at most $E(Q_{n-4})$. Hence, $|A| \leq 2E(Q_{n-4})$ in this case.

□

In fact, we can get more information from the proof, which we summarize in the following corollary.

Corollary 10. *If $E(Q_n) > E(Q_{n-1}) + 1$, then any minimal percolating set A of size $E(Q_n)$ has the form $A = A_1 \cup A_2$ where A_1 and A_2 are both minimal percolating sets in subcubes of dimension at most $n - 4$.*

This result gives a two other nice corollaries. The first is an order of growth upper bound on E .

Corollary 11. *We have $E(Q_n) = O(2^{n/4})$.*

The other is an exact calculation of E for small n .

Corollary 12. *We have $E(Q_0) = 1$, $E(Q_1) = 2$, and $E(Q_n) = n$ for $2 \leq n \leq 8$.*

Proof. When $n \leq 2$ the result is easy. For $n \geq 3$, recall that by Proposition 7 we have $E(Q_n) \geq n$, so it remains to show $E(Q_n) \leq n$. For $n = 3$ the result follows from Corollary 10, as it is not possible to have a subcube of dimension $n - 4$. For $4 \leq n \leq 8$, the result follows from Proposition 9, since $2E(Q_{n-4}) \leq n$ for these n . □

4 Jagged sets

Now, in light of Proposition 9, given $n > 8$, we wish to find minimal percolating sets of Q_{n-4} which we can use to create a minimal percolating set of twice the size in dimension n . The construction from Proposition 7 is unsuitable for this.

Proposition 13. *Suppose A is a minimal percolating set in Q_n with $A = B \cup C$ the disjoint union of two minimal percolating sets in subcubes of dimension $n - 4$. Then neither B nor C is isomorphic to the construction in Proposition 7.*

Proof. To see this, suppose we created a percolating set $A \subset Q_n$ which is a union of one copy of our initial construction B in dimension $n - 4$ and some minimal percolating set C in a Q_{n-4} , embedded into two different subcubes of Q_n of distance at most 2 from each other. Then in some execution path of the percolation process of A , the penultimate step will consist of precisely these two Q_{n-4} 's. To construct such an execution path, simply

combine subcubes in $\langle B \rangle$ with subcubes in $\langle B \rangle$ and subcubes in $\langle C \rangle$ with subcubes in $\langle C \rangle$ until $\langle B \rangle$ and $\langle C \rangle$ are the only two subcubes in \mathcal{A}_i for some i . Now, in our 0, 1, * notation, each Q_{n-4} will have exactly $n - 4$ *'s. Since $n > 8$, there must be at least one coordinate k in which both subcubes have *'s. Now, remove the (unique) vertex v from B so that the sub- Q_{n-5} $\langle B \setminus v \rangle$ does not have an * in the k th coordinate (we know such a vertex exists from the proof of Proposition 7). Then $\langle B \setminus v \rangle$ and $\langle C \rangle$ will still have distance at most 2 from each other, and will still span Q_n , so $B \cup C \setminus v$ will percolate. Thus, $B \cup C$ is not minimal. Hence, we cannot use our initial construction to create minimal percolating sets of size $n - 4 + E(Q_{n-4})$ in dimension n . \square

As the above proposition shows, our initial construction is not suitable for constructing large minimal percolating sets in high dimensions. Thus, we define a type of minimal percolating set which is suited to constructing large minimal percolating sets in high dimensions. It is analogous to Morris' [11] corner-avoiding minimal percolating sets.

Definition 14. We say that a minimal percolating set A in Q_n is jagged if for all $v \in A$, $\langle A \setminus v \rangle$ is disjoint from the $(n - 2)$ -dimensional subcube $* \dots * 00$.

Let $E'(Q_n)$ be the size of the largest jagged set in Q_n . Obviously $E'(Q_n) \leq E(Q_n)$. In the following lemma, we use jagged sets to construct large minimal percolating sets in higher dimensions.

Lemma 15. We have $E'(Q_n) \geq 2E'(Q_{n-4})$ for $n > 5$.

Proof. Suppose we have a minimal percolating set A in Q_{n-4} which is jagged. Then we construct a jagged minimal percolating set B of size $2|A|$ in Q_n . We build up our minimal percolating set B in two halves, B_1 and B_2 . For B_1 , we choose a jagged minimal percolating set isomorphic to A from the subcube

$$* \dots * * * 0001.$$

For B_2 , we choose a jagged minimal percolating set isomorphic to A from the subcube

$$* \dots * 00 * * 10.$$

Now, this set clearly percolates, as the two subcubes shown span all of Q_n and are distance two from each other. We claim that B is minimal.

To see this, suppose we remove a vertex v . By swapping coordinates $n - 3$ and $n - 4$ with coordinates $n - 5$ and $n - 6$ respectively, we can assume without loss of generality that v is from B_1 . Now, since A is jagged, $\langle B_1 \setminus v \rangle$ will be a union of subcubes of distance at least three from each other which have at least one 1 in the $(n - 5)$ -th and $(n - 4)$ -th coordinates. Then each subcube will have distance at least 3 from the others and from B_2 . Hence, $B \setminus v$ does not percolate, so B is minimal. Moreover, B is jagged because every vertex of $\langle B \setminus v \rangle$ will have either 01 or 10 in the last two coordinates. \square

Now, our construction from Proposition 7 is not jagged for $n > 2$, as $0 \dots 0$ is always infected by $A \setminus v$ for any v . However, we demonstrate a jagged minimal percolating set that uses almost as many vertices.

Lemma 16. *There exists a jagged percolating set of size $n - 1$ in dimension n for $n \geq 4$. Thus, $E'(Q_n) \geq n - 1$.*

Proof. Let the first $n - 2$ vertices of A be $\{v_i \mid 1 \leq i \leq n - 2\}$ where v_i is the vertex with a 1 in the i th position and the n th position and 0's everywhere else. Let the $(n - 1)$ st vertex of A be $1 \dots 10$. For example, when $n = 5$, we have $A = \{10001, 01001, 00101, 11110\}$.

The set clearly percolates. The last vertex is obviously necessary for percolation, as it is the only vertex without a 01 in the last two coordinates. Now, suppose we omit one of the other vertices. By permuting the first $n - 2$ coordinates, we can assume that we omit v_1 . Then all the vertices except the last combine to form $0 * \dots * 01$ which has distance 3 from $111\dots 110$, so the set does not percolate. Moreover, the set breaks up into two subcubes, each of which has either a 01 or a 10 in the last two coordinates, so it is jagged. \square

Corollary 17. *We have $E(Q_n) = \Theta(2^{n/4})$.*

Proof. By Proposition 9 we know $E(Q_n) = O(2^{n/4})$ and by Lemma 16 and Lemma 15 we know $E(Q_n) = \Omega(2^{n/4})$. The result follows. \square

In light of Lemma 15, we see that in order to find $E(Q_n)$ for large n , we need only find four sufficiently large consecutive integers with $E(Q_n) = E'(Q_n)$.

Corollary 18. *Suppose there exist four consecutive integers $j, \dots, j+3$ such that $E(Q_n) = E'(Q_n) > E(Q_{n-1})$ for $n \in \{j, \dots, j+3\}$. Then for all $n > j+3$, $E(Q_n) = E'(Q_n) = 2E(Q_{n-4})$.*

Because of the above Corollary, we need only deal with finitely many cases. We will show that $E(Q_n)$ and $E'(Q_n)$ have the values as given in the following chart. This will complete the proof of Theorem 1.

n	$E(Q_n)$	$E'(Q_n)$
2	2	2
3	3	3
4	4	4
5	5	4
6	6	5
7	7	6
8	8	8
9	9	9
10	10	10
11	12	12

Lemma 19. *The values for $E(Q_n)$ and $E'(Q_n)$ are as given in the chart. In particular, for $8 \leq n \leq 11$, we have $E(Q_n) = E'(Q_n) = 8 + \lfloor 2^{n-9} \rfloor$.*

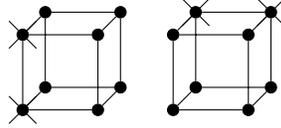


Figure 3: The only jagged minimal percolating set of size 4 in dimension 4.

Proof. We have assembled most of the ingredients of this proof already. The one major piece that we lack is a classification of minimal percolating sets of size n in dimensions $3 \leq n \leq 7$, so we start with this. In dimension 3, there are two isomorphism classes of minimal percolating set, one isomorphic to our initial construction and one jagged. By Corollary 10 we know that any minimal percolating set of size 4 in Q_4 must consist of one vertex plus a minimal percolating set in a sub- Q_3 . Thus, checking case-by-case, we find that there are two isomorphism classes of minimal percolating set in dimension 4, one (the one containing sets isomorphic to $\{0001, 0011, 1110, 1111\}$) that is jagged, and another (the one containing sets isomorphic to the one given in Proposition 7) that is not. Similar case-by-case checking shows that in dimension 5, the only minimal percolating sets of size 5 are isomorphic to our initial construction in Proposition 7. Using this, it is not hard to show that for $n \in \{6, 7\}$ it also holds that the only minimal percolating sets of size n in Q_n are isomorphic to our initial construction. This, together with Corollary 12 and Lemma 16 give the values of $E(Q_n)$ and $E'(Q_n)$ as shown in the chart for $n \leq 7$.

We now use the above classification to show that $E(Q_n)$ is at most the value in the table for $8 \leq n \leq 11$. Corollary 12 tells us that $E(Q_8) \leq 8$ (and is indeed equal to 8). For $9 \leq n \leq 11$, we know by Corollary 10 and Proposition 13 that $E(Q_n) < \max\{2E(Q_n) - 2, E(Q_{n-1}) + 1\}$, which gives the necessary upper bounds for $8 \leq n \leq 11$.

We now show that $E'(Q_n)$ is at least the value given in the table, which will complete the proof. In dimension 8, we can use Lemma 15 on the jagged set of size 4 in dimension 4 $\{0001, 0011, 1110, 1111\}$ to obtain the following jagged minimal percolating set in Q_8 : $\{00010001, 00110001, 11100001, 11110001, 00000110, 00001110, 11001010, 11001110\}$. In dimension 9, we can simply extend our jagged minimal percolating set in dimension 8 to a jagged minimal percolating set in dimension 9 by embedding Q_8 into Q_9 as $*****1$ and adding the vertex 00111110 to obtain $\{000100011, 001100011, 111000011, 111100011, 000001101, 000011101, 110010101, 110011101, 001111110\}$. In dimensions 10 and 11, we can directly apply Lemma 15 to the minimal percolating sets of sizes 5 and 6 in dimensions 6, and 7 respectively as given by Lemma 16 to obtain the desired result. \square

5 Variations

In this section we outline some partial results on generalizations of the original question, namely, grids and augmented hypercubes. We hope that this section leads to future study.

First, we find an upper bound on $E([n]^d, 2)$. Suppose we have an arbitrary grid

$P_{n_1} \times \cdots \times P_{n_d}$ with A a minimal percolating set in the grid. Then this grid will have many different subcubes in it, of many different dimensions. We find an upper bound on the number of elements of A contained in any subcube, in terms of the dimension d of the subcube.

Definition 20. Let $G(d)$ be the maximum over all possible grids, all possible d -dimensional subcubes of the grid, and all minimal percolating sets A in the grid of the number of elements of A contained in the subcube.

We know $G(d)$ is obviously finite because each cube has only finitely many vertices. We prove an upper bound on $G(d)$. The proof relies heavily on the idea of viewing percolation as combining nearby subcubes, and it looks at the ways that the last two cubes in the process can be combined. We set $G(d) = 0$ for $d < 0$. We have an obvious analogue of Lemma 8, whose proof is nearly identical, so we omit it.

Lemma 21. We have $G(d) \geq G(d - 1)$.

Proposition 22. For $d > 1$, $G(d) \leq \max\{G(d - 1) + 1, 2G(d - 3)\}$.

Proof. Let H be a grid, $Q \subset H$ a fixed d -dimensional hypercube and $A \subset H$ a minimal percolating set with $|A \cap Q| = G(d)$. Then since $d \geq 0$, $|A \cap Q| > 0$. Since A percolates, we know that the final term \mathcal{A}_s in any execution path will contain a set containing Q , whereas the first term will not. Thus, we can find a k such that \mathcal{A}_k contains a set containing Q , but \mathcal{A}_{k-1} does not. Hence, the term \mathcal{A}_{k-1} will contain some nonzero number of subgrids which intersect Q . We wish to consider the intersections of these subgrids with Q . Let C_1, \dots, C_j be the set of non-empty intersections of sets in \mathcal{A}_{k-1} with Q , and let C_1 be the largest cube. Note that the C_i are all subcubes of Q . Among all possible execution paths, select one with C_1 as large as possible. Among all execution paths with C_1 as large as possible, select one with k as small as possible. Now, if $j = 1$, then we know that $|A \cap Q| = |C_1| \leq G(d - 1)$, so we are done. Now suppose $j > 2$. Then at least one of the cubes C_i is not necessary to infect Q , since each step in the execution path involves combining only two cubes at a time. This contradicts minimality of k , since otherwise we could reduce k by not performing any of the steps that lead to forming the cubes not used in infecting Q . Thus, we may assume that $j = 2$ and the two subcubes C_1 and C_2 together infect Q_d . By choice of C_1 , $\dim C_1 \geq \dim C_2$. By minimality of A , $A \cap Q \subset C_1 \cup C_2$, and $\dim C_1 \leq d - 1$. We divide into cases depending on $\dim C_1$.

Case 1 $\dim C_1 = d - 1$. Then if $(C_2 \cap A) \setminus C_1$ is empty, we have $|A \cap Q| \leq G(d - 1)$ by induction, and we are done. Thus, suppose there is an $x \in (C_2 \cap A) \setminus C_1$. Then the cube spanned by x and C_1 is all of Q . By minimality of k , we know that $C_2 = \{x\}$, since otherwise we could find an execution path with smaller k by first combining cubes to infect C_1 and then combining C_1 with x . Thus, $|A \cap Q| \leq G(d - 1) + 1$.

Case 2 $\dim C_1 = d - 2$. As before, we are done if $(C_2 \cap A) \setminus C_1$ is empty, so as before let x be a vertex of $(C_2 \cap A) \setminus C_1$. Then x cannot have distance 1 from C_1 , as we could then find an execution path which would make C_1 larger, namely the path in

which we first combine cubes to create C_1 , then combine C_1 with the single vertex x . Thus, x has distance 2 from C_1 . As in case 1, we have by minimality of k that $C_2 = \{x\}$, and so $|A \cap Q| \leq G(d-2) + 1$.

Case 3 $\dim C_1 \leq d-3$. Then by definition, $|A \cap C_1|$ and $|A \cap C_2|$ are each at most $G(d-3)$. Hence, $|A \cap Q| \leq 2G(d-3)$ in this case. □

Since $G(0) = 1$, $G(1) = 2$, $G(2) = 2$, and $G(3) = 3$, we can use the above result to get an upper bound on $E([n]^d, 2)$, since $G(d)$ bounds how many vertices of a minimal percolating set of $[n]^d$ lie in any subcube of $[n]^d$.

Corollary 23. *For $d \geq 1$, we have $G(d) \leq (2 + \lfloor \frac{j}{3} \rfloor) 2^{\lfloor \frac{d-1}{3} \rfloor}$ where $d \equiv j \pmod{3}$, $1 \leq j \leq 3$.*

Corollary 24. *We have $E([n]^d, 2) \leq \lceil \frac{n}{2} \rceil^d G(d)$.*

Proof. Simply partition the grid into hypercubes and apply the previous corollary. □

Now we consider the augmented hypercube, a variation of the hypercube that can be used to model the topological structure of a large-scale parallel processing system. For $r = 2$, a simple inductive argument shows that any minimal percolating set has size 2, since any minimal percolating set must eventually infect two adjacent vertices in a sub- AQ_{n-1} , which will percolate. However, using the “wasted” edge-counting technique from [13], we see that $E(AQ_6, 7) \geq 14$, so percolation is nontrivial for larger r . As this example shows, percolation depends heavily on the structure of the graph, as the Augmented Hypercube has only twice the edges of the standard hypercube, yet percolation is quite different.

6 Conclusion

There are many related problems left to consider, and as we gain more understanding of the percolation process through studying this extremal problem, it is reasonable to hope that we will be able to make more progress on the original probabilistic question. We would like to know $m(G, r)$ and $E(G, r)$ for any finite graph. At the moment however, it seems that finding a general formula or algorithm is too ambitious. We suggest some graphs which we hope are more approachable than a general graph. First, as has been suggested beforehand (e.g. in [11]) it would be very interesting to generalize our hypercube results to $r > 2$.

Question 25. *Find $m(Q_n, r)$ and $E(Q_n, r)$ for $r \geq 3$.*

Exploring graphs similar to the hypercube is also likely to be interesting.

Question 26. *Find $m(G, r)$ and $E(G, r)$ for variations of the hypercube such as the augmented hypercube and the twisted hypercube.*

Finally, it would be interesting to improve our understanding of grids. Based on Corollary 23, we know that $\frac{E([n]^d, 2)}{n^d} \rightarrow 0$ as $d \rightarrow \infty$, independently of n . However, it would be interesting to study the question for fixed d , perhaps following Morris [11] or sharpening his result to find precise asymptotics for grids. Additionally, given the bounds

$$\left(\frac{1}{4}\right)^d \leq E([n]^d, 2) \leq \left(\frac{1}{2}\right)^{2d/3} n^d$$

from Corollary 2 and the proof of Theorem 14 of [11], it is natural to wish to investigate the behavior of $\alpha(n, d) = \left(\frac{E([n]^d, 2)}{n^d}\right)^{1/d}$.

Question 27. Find precise asymptotics for $E([n]^d, r)$.

Acknowledgements

I would like to thank Joe Gallian for bringing the problem to my attention and for all his help. I would like to thank Nathan Kaplan and Nathan Pflueger for all of their help with this paper. I would also like to thank Sasha Ovestsky Fradkin for her advice on this problem, and Phil Matchett Wood for his comments on a very early draft of this paper. This research was done at the University of Minnesota Duluth REU and was supported by the National Science Foundation (grant number DMS-0754106) and the National Security Agency (grant number H98230-06-1-0013).

References

- [1] J. Adler and U. Lev. Bootstrap percolation: visualizations and applications. *Braz. J. Phys.*, 33:641–644, 2003.
- [2] M. Aizenman and J. L. Lebowitz. Metastability effects in bootstrap percolation. *J. Phys. A*, 21(19):3801–3813, 1988.
- [3] J. Balogh and B. Bollobás. Bootstrap percolation on the hypercube. *Probability and Related Fields*, 134:624–648, 2006.
- [4] J. Balogh, B. Bollobás, H. Duminil-Copin, and R. Morris. The sharp metastability threshold for r -neighbor bootstrap percolation. *In preparation*.
- [5] J. Balogh, B. Bollobás, and R. Morris. Bootstrap percolation in three dimensions. *Ann. Probab.*, 37:1329–1380, 2009.
- [6] J. Balogh, Y. Peres, and G. Pete. Bootstrap percolation on infinite trees and non-amenable groups. *Combinatorics, Probability and Computing*, 15:715–730, 2006.
- [7] R. Cerf and E. Cirillo. Finite size scaling in three-dimensional bootstrap percolation. *Ann. Probab.*, 27(4):1837–1850, 1999.
- [8] R. Cerf and F. Manzo. The threshold regime of finite volume bootstrap percolation. *Stochastic Process. Appl.*, 101(1):69–82, 2002.

- [9] J. Chalupa, P. L. Leath, and G. R. Reich. Bootstrap percolation on a bethe lattice. *J. Phys. C.*, 12:L31–L35, 1979.
- [10] Alexander Holroyd. Sharp metastability threshold for two-dimensional bootstrap percolation. *Probab. Theory Related Fields*, 125(2):195–224, 2003.
- [11] Robert Morris. Minimal percolating sets in bootstrap percolation. *Electron. J. Combin.*, 16(1):Research Paper 2, 20, 2009.
- [12] Gabor Pete. Disease processes and bootstrap percolation. *Thesis for diploma at the Bolyai Institute, Jzsef Attila University, Szeged*, 1997.
- [13] Eric Riedl. Largest and smallest minimal percolating sets in trees. *Preprint*.
- [14] R. H. Schonmann. On the behaviour of some cellular automata related to bootstrap percolation. *Ann. Prob.*, 20:174–193, 1992.