Walker-Breaker games

Lisa Espig [*] Alan Frieze [†] Michael Krivelevich[‡] Wesley Pegden[§]

June 23, 2015

Abstract

We introduce and analyze the *Walker-Breaker* game, a variant of Maker-Breaker games where Maker is constrained to choose edges of a walk or path in a given graph G, with the goal of visiting as many vertices of the underlying graph as possible.

1 Introduction

Maker-Breaker games were introduced by Erdős and Selfridge [5] as a generalisation of Tic-Tac-Toe. Since then there have been many results on variations on this theme. In a standard version, played on the complete graph K_n, Maker and Breaker take turns acquiring edges, with Maker trying to build a particular structure (e.g., a clique) in her own edges, and with Breaker trying to prevent this. See the recent book by Beck [1] for a comprehensive analysis of Maker-Breaker games.

We consider the following variant on the standard Maker-Breaker game. In this variant, the *Walker-Breaker game*, the "Walker" acquires the edges of a walk consecutively; i.e., at any given moment of the game we have her positioned at some vertex v of a graph G and on her turn, she moves along an edge e of G that is (i) incident with v and (ii) has not been acquired by Breaker. If she has not already acquired e, then she is now considered to have acquired it. On Breaker's move, he can acquire any edge not already owned by Walker. In some cases we will allow him to acquire β edges in one move; in this case the *bias* of the game is $1 : \beta$.

In this paper, we consider Walker-Breaker games where Walker's goal is to visit as many vertices as she can. Breaker's goal is to reduce the number of vertices that she visits. The game ends when there is no path from Walker's current position to an unvisited vertex along edges not acquired by Breaker.

We also consider a variant of this game, the *PathWalker-Breaker game*, in which Walker cannot revisit any previously visited vertices; this game ends when there is no path from Walker's current position to an unvisited vertex along edges not acquired by Breaker, and vertices not previously visited by Walker. (Obviously, Walker can visit at least as many vertices in the Walker-Breaker game on a graph as in the PathWalker-Breaker game on the same graph). In this situation, we sometimes refer to Walker as PathWalker to avoid ambiguity.

In a fictional scenario, Walker represents a missionary who is traversing a network, trying to convert as many people ($\equiv$ vertices) to his beliefs. Breaker represents the devil, whose only way to block Walker is to burn untraversed edges of the network.

Our first Theorem can be seen as a strengthening of the result of Hefetz, Krivelevich, Stojaković and Szabó [8], that in a Maker-Breaker game on K_n, Maker can construct a Hamilton path in $n - 1$ moves.

Theorem 1. *Under optimum play in the* $1 : 1$ *PathWalker-Breaker game on* K_n $(n > 5)$*, PathWalker visits all but two vertices.*

[*] Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA15213, USA. Research supported in part by NSF grant ccf1013110, Email lespig@andrew.cmu.edu.

[†] Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA15213, USA. Research supported in part by NSF grant ccf1013110, Department of Mathematical Sciences, Email alan@random.math.cmu.edu.

[‡] School of Mathematical Sciences, Raymond and Beverly Sackler Faculty of Exact Sciences, Tel Aviv University, Tel Aviv, 69978, Israel. Research supported in part by a USA-Israel BSF grant and by a grant from the Israel Science Foundation. Email krivelev@post.tau.ac.il.

[§] Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA15213, USA. Research supported in part by NSF grant, Email wes@math.cmu.edu.

The bias has a substantial effect on the PathWalker-Breaker game:

Theorem 2. *Under optimum play in the* $1:\beta$ *PathWalker-Breaker game on* K_n *for* $1<\beta=O(1)$*, PathWalker visits all but* s *vertices for* $c_1\log n\le s\le c_2\log n$*, for constants* c_1,c_2 *depending on* β*.*

In the Walker-Breaker game, the effect of bias is not so drastic:

Theorem 3. *Under optimum play in the* $1:\beta$ *Walker-Breaker game on* K_n $(n>2\beta^2)$*, Walker visits* $n-2\beta+1$ *vertices. Here* $1\le\beta=O(1)$*.*

For the sake of a graph which is not complete, consider the cube Q_n, which is the graph on the vertex set $\{0,1\}^n$, where two strings are adjacent iff they have Hamming distance 1.

Theorem 4. *In optimum play in the Walker-Breaker game on* Q_n*, Walker visits at least* 2^{n-2} *vertices, and at most* 2^{n-1} *vertices.*

Finally, we consider a one-player game the *Random-Walker game*, in which the moves of Walker are made according to a random walk on the edges not acquired by Breaker. Breaker acquires one edge per move and he has the goal of minimizing the typical number of vertices visited by RandomWalker. The game ends when there is no path between the position of RandomWalker and an unvisited vertex along edges not acquired by Breaker. For a graph G, we let the *co-degree* of vertices u,v be their number of common neighbors.

Theorem 5. *If* G *has minimum co-degree at least* αn *for some absolute constant* $\alpha>0$ *then under optimum play (by Breaker), RandomWalker visits at least* $n-c\log n$ *vertices of* G *w.h.p., for a constant* c *depending on* α*.*

Theorem 6. *If* G *has minimum degree at least* αn *for some absolute constant* $\alpha>0$ *then under optimum play (by Breaker), RandomWalker visits at most* $n-c\log n$ *vertices w.h.p., for any constant* $c<\alpha$*.*

1.1 Some notation

We let V_t (resp. U_t) denote the set of vertices that have been visited (resp. not visited) by Walker after Walker has made t moves. Walker is at vertex v_t after t moves. The graph induced by Breaker's edges is denoted by Γ_B and the graph induced by Walker's edges is denoted by Γ_W.

2 Proof of Theorem 1

Here at time t the graph Γ_W is a path P_t. To show that who goes first does not matter, we assume that Breaker goes first for any lower bound on the number of visited vertices, and that Walker goes first for any upper bound on the number of visited vertices.

2.1 Lower bound

Walker's strategy is as follows: If $|U_t|>2$ and Breaker chooses f_t and $f_t\cap V_{t-1}=\emptyset$ then Walker moves to $v_t\in f_t$. Otherwise, Walker moves to an arbitrary vertex. So long as Walker is able to follow this strategy, we will have after each Walker move that:

$$\text{Every Breaker edge contains a member of } V_t. \tag{1}$$

We now check that this strategy is feasible for $|U_t|>2$. We begin with the first type of Breaker's edge, which is disjoint from V_t. Fix t and let $v_{t-1}=x$ and $v_t=y$. Suppose that Breaker chooses an edge (b_1,b_2) where $b_1,b_2\notin V_t$ and such that for $i=1,2$, (y,b_i) is a Breaker edge. Assume that this is the first time this situation happens. Suppose next that (y,b_i) is the s_i-th edge chosen by Breaker. Assume that $s_1<s_2$. We now have a contradiction to (1) after the choice of x. For after x is chosen, (y,b_1) is a Breaker edge that does not contain a member of V_{t-1}.

We now consider the case where Breaker's edge is incident with V_t. (1) implies that Breaker's choice is at most the second edge between v_t and U_t. In particular, $|U_t|>2$ implies that Walker can move to an unvisited vertex, and Walker will succeed at visiting all but 2 vertices of the graph.

2.1.1 Upper bound

Breaker plays arbitrarily until his move at time $n-4$, when $|U_{n-4}|=4$. In his next two moves he chooses the two edges of a matching in U_{n-4}. After these two moves (with one Walker move in between), it is again Walker's turn, and 3 unvisited vertices remain. Regardless of which vertex of U_{n-3} Walker might move to next, that vertex will already be adjacent along one of Breaker's 2 matching edges to a vertex in U_{n-2}; thus, with one additional move, Breaker will ensure that both edges from v_{n-2} to U_{n-2} are occupied by Breaker.

3 Proof of Theorem 2

3.1 Lower bound

We assume that Breaker goes first and describe Walker's strategy. We suppose that Walker is at some vertex x and describe the next sequence of moves B,W,B,W (Breaker,Walker,Breaker,Walker.) We call such a sequence a *round*. We keep track of two sets L, R that partition the set of unvisited vertices U_t. Let $\beta_R(v)$ be the number of Breaker edges $(v, z), z \in R$. At the outset of the game, $L = \varnothing$. Before Breaker's first move in each round, we move vertices satisfying $\beta_R(v) \geq \alpha|R|$ from from R to L one by one, (updating R each time), until no such vertices remain, for $\alpha = \frac{1}{3(\beta+1)}$.

We describe Walker's strategy for a round as follows. Suppose that Breaker has made his first move of the round and let $R = \{w_1, w_2, \ldots, w_r\}$ at the end of this move. We assume that $\beta_R(w_i) \geq \beta_R(w_{i+1}), 1 \leq i < r$.

For his first move of the round, Walker moves to a vertex $z \in R$ such that none of $(z, w_j), 1 \leq j \leq \beta+1$, is a Breaker edge. Breaker's response consists of just β edges; Walker's final move is to move from z to one of the vertices w_j $1 \leq j \leq \beta+1$.

We will prove that it is possible to follow this strategy until most vertices have been visited. This proof is based on two ingredients:

Claim 1. *There is a constant c_β such that so long as Walker follows this strategy and so long as $|R| > c_\beta \log n$, at most β vertices are moved from R to L in any given round.*

Claim 2. *There is a constant C_β such that so long as Walker follows this strategy, we will have $|L| \leq C_\beta \log n$.*

Let us first see how these two claims imply that Walker can follow this strategy until she has visited all but $c_2 \log n$ vertices for some c_2 depending on β.

We first check that Walker can always move to a suitable intermediate vertex z. After Breaker's opening move of the round, at most $(\beta+1)\alpha|R| + \beta = |R|/3 + \beta$ vertices of R can be Breaker neighbors of $w_1, w_2, \ldots, w_{\beta+1} \in R$. Moreover, the fact that x was in R at the beginning of the previous round, together with Claim 1, means that if $|R| > c_\beta \log n$, then x has at most $\beta + \alpha(|R| + \beta) \leq \beta + |R|/3$ neighbors in R, leaving at least $\frac{1}{3}|R| - 2\beta$ choices for z. So, Walker will be able to move to such a z as long as $|R| > \max(6\beta, c_\beta \log n)$. For $\beta = O(1)$, Claim 2 now implies that Walker can follow this strategy until all but $(C_\beta + c_\beta) \log n$ vertices have been visited.

It remains to prove Claims 1 and 2. We do this via a simpler box game.

3.1.1 Box game

We analyze Walker's strategy via a *box game*, similar to the box game of Chvátal and Erdős [3]. There will now be only one player, whom we call BREAKER. Any move by Breaker in the Walker-Breaker game will be modelled by a BREAKER move in this box Game.

Consider a sequence $b_1 \geq b_2 \geq \cdots \geq b_n$ of non-negative integers. The box game is played as follows: At the beginning of each turn of the game, BREAKER has a *loss* phase, in which he may, at his option, designate terms with value at least $\alpha = \frac{1}{3(\beta+1)}$ times the remaining number of terms in the sequence as *lost*.

Following the loss phase, BREAKER increases each of the first β terms of the sequence by an amount up to 4β. In addition he also increases terms b_i for $i > \beta$ by a total amount up to 4β. After this he deletes one of the currently largest $\beta+1$ terms $b_1, b_2, \ldots, b_{\beta+1}$ of the sequence, and up to one other term from anywhere in the sequence. (At any point, the sets of remaining, lost, and deleted terms of the original sequence form a partition of the terms of the original sequence.)

The relevance of this game stems from the following:

Claim 3. *If Breaker can play the PathWalker-Breaker game on a graph with n vertices against Walker which is following the strategy described earlier such that $|L|$ increases by ℓ_t in each round t, then in this box game, beginning with the all zeroes sequence of length n, BREAKER can maintain that ℓ_t terms become lost in turn t.*

Proof. At any point, the sequence (b_i) represents Breaker degrees of vertices not yet visited by Walker and not in L, in decreasing order. In any round, Breaker places 2β edges in the graph, which increases the total degrees by 4β; thus, BREAKER can produce the exact same resulting (degree) sequence with an allowed alteration of the terms of the sequence. In any round, Walker will visit two vertices, the second of which is a vertex of among the highest possible $\beta+1$ degrees; this corresponds to the deletion of the two terms on each turn of the box game. If in each loss phase of the box game, BREAKER loses as many terms as possible, then BREAKER will lose exactly as many terms as vertices that enter L in the PathWalker-Breaker game. $\square$

Note that our definition of the box game allows much more freedom to BREAKER than is necessary for Claim 3. This extra freedom does however simplify the analysis, by enabling us to decouple consideration of the first β "boxes" from the rest. In particular, call a term in the box game sequence a *tail term* if it is not among the β largest. We will prove Claims 1 and 2 by proving the following lemma regarding the box game:

Lemma 1. *There is a constant A_β such that after any number of steps t in the box game, the tail terms are all at most $A_\beta \log t$.*

First let us observe that Lemma 1 implies both Claims 1 and 2, via Claim 3.

Proof of Claim 1. Let $c_\beta = A_\beta/\alpha$, and suppose Breaker can play the PathWalker-Breaker game on a graph with n vertices against Walker which is following the strategy described earlier, and achieve that on some turn, more than β vertices become lost, i.e. move from R to L. Claim 3 implies that he can play the box game and achieve that on some turn, more than β terms b_j become lost, which would be a contradiction since $|R| > \frac{A_\beta}{\alpha}\log n$ at the beginnning of the turn and $b_j \le A_\beta \log t \le A_\beta \log n$ for $j > \beta$ implies that $b_j < \alpha\,|R|$ for $j > \beta$; in particular, only the β largest terms can become lost on any given turn. $\square$

Proof of Claim 2. It suffices to prove that the claim holds so long as $|R| > 2c_\beta \log n$ (with $c_\beta = A_\beta/\alpha$, as before), since Claim 2 will then remain true even if all remaining vertices in R are moved to L. In particular, we may assume that in any given round, only vertices from among the β maximum Breaker-degree vertices of R become lost.

Using Claim 3, we carry out our analysis in the box game. At the beginning of a turn, some terms b_i for $i \le \beta$ may become lost. If the term b_i became one of the β largest (for the last time) at turn t_0 when r_0 terms remained, and became lost at turn $t_1 = t_0 + k$, when a total of r_1 terms remain, then Lemma 1 implies that the term had at most $A_\beta \log t_0$ balls at turn t_0. To become lost at time t_1 requires the term to have at least αr_1 balls; thus, we have that

$$\alpha r_1 - A_\beta \log t_0 \le 4\beta k.$$

Since $r_1 \ge r_0 - k(\beta+2)$, we have

$$k \ge \frac{\alpha r_1 - A_\beta \log t_0}{4\beta} \ge \frac{\alpha r_1}{8\beta} \ge \frac{\alpha(r_0 - (\beta+2)k)}{8\beta},$$

since $r_1 \ge 2c_\beta \log n \ge 2c_\beta \log t_0$. In particular,

$$k \ge \frac{\alpha r_0}{8\beta + \alpha(\beta+2)} \ge \frac{\alpha r_0}{10\beta}. \tag{2}$$

In particular, when a term enters the largest β for the last time, it takes at least a number of steps k which is a constant fraction of the number r_0 of remaining terms when it entered, before it can become lost. Once it becomes lost, say, when there are r_1 remaining terms, some other term enters the largest β terms, and to become lost this term requires at least a number of terms to become lost which is a constant fraction of r_1. Continuing in this manner, we see the terms $r_1, r_2, \ldots,$ of this sequence must satisfy (with k as in (2)) $r_{i+1} \le r_i - k \le r_i(1 - \frac{\alpha}{10\beta})$ by (2), since at least one term is deleted on each turn of the box game. In particular, with $r_0 = n$, we have that this sequence can have at most $\frac{\log n}{\log\left(\frac{10\beta}{10\beta-\alpha}\right)}$ terms. We can have β different such sequences producing lost terms (one for each of the β initially largest terms), giving a maximum of

$$\frac{\beta}{\log\left(\frac{10\beta}{10\beta-\alpha}\right)}\log n$$

lost terms produced. $\square$

We will use the following lemma to prove Lemma 1:

Lemma 2. *Suppose that $e_1 \geq e_2 \geq \cdots \geq e_s$ and $b_1 \geq b_2 \geq \cdots \geq b_r$ are two states of the box game where $s \leq r$ and $e_i \leq b_i$ for all $1 \leq i \leq s$, and $f : \mathbb{N} \to \mathbb{N}$ is an arbitrary function. If BREAKER has a strategy in the first state to ensure that for some t, some tail term has value $\geq f(t)$, then he has a strategy in the second state to ensure that for some t, some tail term is $\geq f(t)$ by the turn t.*

Proof. BREAKER simply mimics his strategy for the sequence $\{e_i\}$ with the sequence $\{b_i\}$. Terms deleted or lost for the game on the first sequence are deleted or lost, respectively, for the game on the second sequence. (Note that the freedom BREAKER has to choose not to lose terms is important here.) $\square$

We are now ready to prove Lemma 1.

Proof of Lemma 1. Lemma 2 implies that it suffices to prove Lemma 1 for the case where, on each turn, after the loss phase, only one term is deleted by BREAKER, and this term is the $(\beta+1)$'st largest term. Under this assumption, the terms $b_{\beta+1}, b_{\beta+2}, \ldots$ are reproducing the classical box game of Chvátal and Erdős [3], see also Hamidoune and Las Vergnas [6], since the largest of these terms is deleted on each turn. In particular, a simple potential function argument shows that $b_{\beta+1} \leq A_\beta \log t$ throughout, for $A_\beta = \frac{1}{\log\left(\frac{4\beta}{4\beta-1}\right)}$. $\square$

3.2 Upper bound

Breaker's strategy is as follows: Breaker chooses a vertex $w_1 \notin \{v_1, v_2\}$. He will spend the next $(n-1)/\beta$ moves making sure that Walker cannot visit w_1. In a move, Breaker claims the edge from w_1 to v_t, if necessary, plus $\beta - 1$ other edges incident with w_1. This takes approximately $n_1 = (n-1)/\beta$ moves. Breaker then chooses another unvisited vertex w_2 and spends approximately $n_2 = (n-1-n_1)/\beta$ moves protecting w_2. It takes only n_2 rather than n_1 moves because Walker cannot use n_1 of the edges to w_2, because she has visited the other endpoint. Continuing in this manner Breaker protects w_k in n_k moves where

$$n_k = \frac{n-1-(n_1+n_2+\cdots+n_{k-1})}{\beta} = \frac{n-1}{\beta}\left(1-\frac{1}{\beta}\right)^{k-1}.$$

It follows from this that

$$n_1 + n_2 + \cdots n_k = (n-1)\left(1-\left(1-\frac{1}{\beta}\right)^k\right).$$

Thus we can take $k = c_1 \log n$ where $c_1 = 1/\log(\beta/(\beta-1))$. This will be our value of c_1 in Theorem 2.

This completes the proof of Theorem 2.

4 Proof of Theorem 3

In this section, Walker is not constrained to a path (her walk may use an edge more than once).

4.1 Lower bound

Walker builds a tree T in a depth first manner. She starts at the root v_1 at depth 0. All depth/parent/child statements are with respect to this root. A vertex $v \in T$ will have a parent $w = \pi(v)$ where the depth of v is one more than the depth of w. If Walker is at vertex x and there is a vertex $y \in U_t$ such that Breaker has not claimed the edge (x, y) then Walker moves to y. We let $x = \pi(y)$. Otherwise, if no such move is possible, Walker moves to $\pi(x)$ and repeats the search for $y \in U_t$ on her next move. The game is over when Walker finds herself at v_1 and all edges v_1 to U_t have been taken by Breaker.

Suppose that the game ends with $|U_t| = k$. Then Walker has made $2(n-k-1)$ moves. Each edge of T has been traversed twice, once in a forward direction and once in a backwards direction. Breaker has captured at least $k(n-k)$ edges between T and U_t. We therefore have

$$k(n-k) \leq 2\beta(n-k-1).$$

It follows from this that $k < 2\beta$. This shows that Walker visits at least $n - 2\beta + 1$ vertices.

4.2 Upper bound

The argument here has some similarities to that in Section 3.2. Breaker's strategy is as follows: Assuming Walker goes first and claims an edge $\{v_1, v_2\}$, Breaker chooses a vertex $w_1 \notin \{v_1, v_2\}$. He will spend the next $(n-1)/\beta$ moves making sure that Walker cannot visit w_1. In a move, he claims the edge from w_1 to v_t, if necessary, plus $\beta - 1$ other edges incident with w_1. This takes approximately $(n-1)/\beta$ moves. Then he chooses w_2, not visited and protects it from being visited in the same way. He does this for $w_1, w_2, \ldots, w_{\beta-1}$. Altogether, this takes up at most $(\beta-1)\lceil (n-1)/\beta \rceil$ moves, leaving at least $\lfloor (n-1)/\beta \rfloor + 1$ vertices unvisited. Breaker then chooses β unvisited vertices $y_1, y_2, \ldots, y_\beta$ (possible since $n > 2\beta^2$) and a move consists of capturing the edges (v_t, y_i), $i = 1, 2, \ldots, \beta$. This protects $y_1, y_2, \ldots, y_\beta$ and so Walker visits at most $n - 2\beta + 1$ vertices. This completes the proof of Theorem 3.

5 Proof of Theorem 4

5.1 Lower bound

We use a similar argument to that in Section 4.1. Walker builds a Depth First Search tree T. Again, the edges between T and U_t will all be Breaker's edges. Suppose now that T has k vertices. Then

$$2(k-1) \geq e(T, U_t) \geq k(n - \log_2 k).$$

The lower bound follows from Harper's theorem [7]. It follows that

$$\log_2 k \geq n - 2 + \frac{2}{k}$$

and so at least 2^{n-2} vertices are visited by Walker.

5.2 Upper bound

Suppose that Walker goes first and assume w.l.o.g. that she starts at $(0, 0, \ldots, 0)$ and then moves to $(0, 1, \ldots, 0)$. Breaker will not allow her to visit any vertex whose first component is 1. When Walker moves to $(0, x_2, x_3, \ldots, x_n)$, Breaker acquires the edge $((0, x_2, x_3, \ldots, x_n), (1, x_2, x_3, \ldots, x_n))$. Breaker can acquire the edge $((0, 0, \ldots, 0), (1, 0, \ldots, 0)$ on his last move, if not before. It follows that at most 2^{n-1} vertices are visited by Walker. This completes the proof of Theorem 4.

6 Proof of Theorem 5

Here we will assume that Walker does a random walk on a graph G. When at a vertex v she chooses a random neighbor w for which the edge (v, w) is not a Breaker edge.

Let $\beta = \alpha^3/36$ and consider the first $t_0 = 4\beta^{-1} n \log n$ moves. We will show that Walker will w.h.p. visit the required number of vertices within this time. Let G_t be the subgraph of G induced by the edges not acquired by Breaker after t moves. Let L_t be the set of vertices incident with more than $\alpha n/3$ Breaker edges after the completion of t moves by Breaker. Clearly $|L_t| \leq C_0 \log n$, where $C_0 = 12(\alpha\beta)^{-1}$.

Recall that v_t denotes the current vertex being visited by Walker. We re-define, for the purposes of this proof, U_t to be the set of vertices that are not in L_t and are currently unvisited. Then

(a) If $v_t \notin L_t$ then the probability that Walker visits U_t within two steps is at least $\beta|U_t|/n$. To see this let $Z = |N(v_{t+1}) \cap U_t|$. Then $\mathbf{E}(Z) \geq \alpha|U_t|/3$. This is because v_t and any $w \in U_t$ have at least $\alpha n - 2\alpha n/3 = \alpha n/3$ common neighbors in G_t. Thus, if $\bar{Z} = |U_t| - Z$ then $\mathbf{E}(\bar{Z}) \leq (1 - \alpha/3)|U_t|$. It follows from the Markov inequality that $\mathbf{Pr}(\bar{Z} \geq (1 - \alpha^2/9)|U_t|) \leq \frac{1}{1+\alpha/3}$ and so $\mathbf{Pr}(Z \geq \alpha^2|U_t|/9) \geq \frac{\alpha}{3+\alpha}$. Finally observe that $\mathbf{Pr}(v_{t+2} \in U_t \mid Z) \geq (Z-1)/n$, where we have subtracted 1 to account for Breaker's next move.

(b) We divide our moves up into periods $A_1, B_1, A_2, B_2, \ldots$, where A_j is a sequence of moves taking place entirely outside L_t and B_j is a sequence of moves entirely within L_t. During a time period A_j, the probability this period ends is at most $\frac{C_0 \log n}{\alpha n}$. So the number of time periods is dominated by the binomial $Bin(4\beta^{-1} n \log n, C_0 \log n/(\alpha n))$ and so with probability $1 - o(n^{-3})$ the number of periods is less than $5C_0(\alpha\beta)^{-1} \log^2 n$.

(c) We argue next that

$$\text{with probability } 1 - o(n^{-3}) \text{ each } B_j \text{ takes up at most } O(\log^6 n) \text{ moves.} \tag{3}$$

Suppose that B_j begins with a move from $v \notin L_t$ to $w \in L_t$. Let $L^* = L_t \cup \{v\}$ and let H^* denote the subgraph induced by the edges contained in L^* that have not been acquired by Breaker. Walker's moves in period B_j constitute a random walk on (part) of the graph H^*. This is not quite a simple random walk, since H^* changes due to the fact that Breaker can delete some of the edges available to Walker. Nevertheless, Walker will always be in a component of H^* containing v_t. This is because Walker has arrived at the current vertex via a walk from v_t. Now consider running this walk for $C_1 \log^5 n$ steps, where C_1 is some sufficiently large constant. Observe that Breaker can claim at most $C_0^2 \log^2 n$ edges inside this component of H^*. Hence there will be an interval of length $C_2 \log^3 n$, $C_2 = C_1/C_0^2$ where Breaker does not claim any edge inside H^*. This means that in this interval we perform a simple random walk on a connected graph with at most $1 + C_0 \log n$ vertices. If we start this interval at a certain vertex x, then we are done if the random walk visits v. It follows from Brightwell and Winkler [2] that the expected time for the walk to visit v can be bounded by $C_0^3 \log^3 n$. So, if $C_2 > 2C_0^3$ then v will be visited with probability at least 1/2.

Suppose that time has increased from the time t when B_j began to t' when v is first re-visited. If $v \notin L_{t'}$ then B_j is complete. If however $v \in L_{t'}$ then we know that v is incident with at most $\alpha n/3 + C_1 \log^5 n$ Breaker edges. So the probability that Walker leaves $L_{t'}$ in her next step is at least

$$\frac{d_G(v) - (\alpha n/3 + C_1 \log^5 n)}{d_G(v)} \geq \frac{1}{2}. \tag{4}$$

So the probability that B_j ends after $C_1 \log^5 n$ steps is at least 1/4. Suppose on the other hand that B_j does not end and that we return to v for kth time where $k \leq 20 \log n$. The effect of this is to replace $C_1 \log^5 n$ in (4) by $kC_1 \log^5 n$. This does not however affect the final inequality. So if C_1 is sufficiently large, the probability that B_j does not end after $20C_1 \log^6 n$ steps is at most $(3/4)^{20 \log n} = o(n^{-4})$. Estimate (3) follows immediately.

(d) Combining the discussion in (b), (c) we see that w.h.p. $\left|\bigcup_j B_j\right| = O(\log^8 n)$, which is negligible compared with t_0; i.e., Walker spends almost all of her time outside L_{t_0}. Let $X_i, 1 \leq i \leq k = n - C_0 \log n$, be the time needed to add the ith vertex to the list of vertices visited by Walker. (Here we exclude any time spent in $\bigcup_j B_j$.) It follows from (a) that $X_i/2$ is dominated by a geometric random variable with probability of success $\frac{(n-i-C_0 \log n)\beta}{n}$. This is true regardless of $X_1, X_2, \ldots, X_{i-1}$. So $(X_1 + \cdots + X_k)/2$ is dominated by the sum of independent geometric random variables. Furthermore, $\mathbf{E}(X_1 + \cdots + X_k) \leq \frac{2}{\beta} n \log n$ and it is not difficult to show that $X_1 + \cdots + X_k \leq t_0$ w.h.p. Indeed, the variance of a geometric random variable with probability of success p is given by $1/p^2 - 1/p \leq 1/p^2$. So, by the Chebyshev inequality,

$$\mathbf{Pr}(X_1 + \cdots + X_k > t_0) \leq \frac{\beta^2}{4n^2 \log^2 n} \sum_{i=1}^{k} \frac{4n^2}{\beta^2(n - i - C_0 \log n)^2} = O\left(\frac{1}{\log^2 n}\right).$$

This completes the proof of Theorem 5.

7 Proof of Theorem 6

We assume that Walker chooses a vertex a_0 to start at and then Breaker chooses an edge to acquire.

Breaker's strategy will be to choose an arbitrary unvisited vertex v_1 and protect it by always on his turn taking the edge (v_1, w) where w is the current vertex being visited by Walker, if $(v_1, w) \in E(G)$. If Breaker has already acquired (v_1, w) or $(v_1, w) \notin E(G)$ then he will choose an unacquired edge incident with v_1. This continues until Breaker has acquired all of the edges incident with v_1. He then chooses v_2 and protects it. This continues until there are no unvisited vertices to protect.

After Breaker has protected $v_1, v_2, \ldots, v_{k-1}$ and while he is protecting v_k, Walker finds herself doing a random walk on a dense graph with $n-k$ vertices. Let the moves spent protecting v_k be denoted by *round* k. (More precisely, round k consists of the moves, after round $k-1$, until Breaker has acquired all edges incident with v_k.)

Fix $k = O(\log n)$ and let ζ_k be the number of unvisited, unprotected vertices when Breaker begins protecting v_k. Because Breaker has taken the edges incident with $v_1, v_2, \ldots, v_{k-1}$, it will take at most $n-k$ more moves to protect v_k. If w is an unvisited, unprotected vertex at the start of the round, then it remains unvisited with probability at least $\left(1 - \frac{1}{\alpha n - k}\right)^{n-k-1} = e^{-1/\alpha} + O(1/(n-k))$. It follows that $\mathbf{E}(\zeta_{k+1}) \approx \zeta_k / e^{1/\alpha}$. To show that it is close to this w.h.p. we proceed as follows: Suppose that we throw $n-k$ balls randomly into $\alpha n - k$ boxes, of which ζ_k are special. Then ζ_{k+1} dominates the number of empty special boxes. Chernoff bounds are applicable in this case due to negative association (see Dubhashi and Ranjan [4]) and so if $\zeta_k \gg \log n$, ζ_{k+1} will be concentrated around its mean.

It follows that w.h.p. $\zeta_k \approx n e^{-k/\alpha}$ for $k \leq (1-\epsilon)\alpha \log n$ where $0 < \epsilon < 1$ is a positive constant. Thus Breaker will w.h.p. be able to protect $(1-\epsilon)\alpha \log n$ vertices and we can choose any $c < \alpha$ in Theorem 6.

This completes the proof of Theorem 6.

8 Further Questions

Some natural questions spring to mind:

- How large a cycle can Walker make under the various conditions?
- Suppose the goal is to visit as many edges as possible: what can be achieved under various game conditions?
- Which subgraphs can Walker make? How large a clique can she make? Observe that PathWalker cannot even make a triangle.
- What if we allow Walker to have b moves to Breaker's one move?
- What happens if Breaker is also a walker?

Acknowledgement: We thank Dennis Clemens for pointing out an error in an earlier version of Theorem 1.

References

[1] J. Beck, Combinatorial Games: Tic-Tac-Toe Theory, Cambridge University Press, 2008.

[2] G. Brightwell and P. Winkler, Maximum hitting time for random walks on graphs, *Random Structures and Algorithms* 1 (1990) 263–276.

[3] V. Chvátal and P. Erdős, Biased positional games, *Annals of Discrete Mathematics* 2 (1978) 221–229.

[4] D. Dubhashi and D. Ranjan, Balls and bins: a study in negative dependence, *Random Structures and Algorithms* 13 (1998) 99-124

[5] P. Erdős and J. Selfridge, On a combinatorial game, *Journal of Combinatorial Theory B* (1973) 298–301.

[6] Y. Hamidoune and M. Las Vergnas, A solution to the box game, *Discrete Mathematics* 65 (1987) 157–171.

[7] L. Harper, Optimal numberings and isoperimetric problems on graphs, *Journal of Combinatorial Theory* 1 (1966) 385–394.

[8] D. Hefetz, M. Krivelevich, M. Stojaković and T. Szabó, Fast winning strategies in Maker-Breaker games. *Journal of Combinatorial Theory B* 99 (2009), 39–47.

[9] D. Levin, Y. Peres, and E. Wilmer, Markov Chains and Mixing Times, American Mathematical Society (2008).