k-cuts on a Path

Xing Shi Cai[1(✉)], Luc Devroye[2], Cecilia Holmgren[1], and Fiona Skerman[1]

[1] Mathematics Department, Uppsala University, 75237 Uppsala, Sweden
{xingshi.cai,cecilia.holmgren,fiona.skerman}@math.uu.se

[2] School of Computer Science, McGill University, Montréal, QC H3A 2A7, Canada
lucdevroye@gmail.com

Abstract. We define the (random) k-cut number of a rooted graph to model the difficulty of the destruction of a resilient network. The process is as the cut model of Meir and Moon [14] except now a node must be cut k times before it is destroyed. The first order terms of the expectation and variance of $\mathcal{X}_n$, the k-cut number of a path of length n, are proved. We also show that $\mathcal{X}_n$, after rescaling, converges in distribution to a limit $\mathcal{B}_k$, which has a complicated representation. The paper then briefly discusses the k-cut number of general graphs. We conclude by some analytic results which may be of interest.

Keywords: Cutting · k-cut · Network · Record · Permutation

1 Introduction and Main Results

1.1 The k-cut Number of a Graph

Consider $\mathbb{G}_n$, a connected graph consisting of n nodes with exactly one node labeled as the *root*, which we call a *rooted* graph. Let k be a positive integer. We remove nodes from the graph as follows:

1. Choose a node uniformly at random from the component that contains the root. Cut the selected node once.
2. If this node has been cut k times, remove the node together with edges attached to it from the graph.
3. If the root has been removed, then stop. Otherwise, go to step 1.

We call the (random) total number of cuts needed to end this procedure the k-cut number and denote it by $\mathcal{K}(\mathbb{G}_n)$. (Note that in traditional cutting models, nodes are removed as soon as they are cut once, i.e., $k = 1$. But in our model, a node is only removed after being cut k times.)

One can also define an edge version of this process. Instead of cutting nodes, each time we choose an edge uniformly at random from the component that contains the root and cut it once. If the edge has been cut k-times then we

This work is supported by the Knut and Alice Wallenberg Foundation, the Swedish Research Council, and the Ragnar Söderbergs foundation.

P. Heggernes (Ed.): CIAC 2019, LNCS 11485, pp. 112–123, 2019.
https://doi.org/10.1007/978-3-030-17402-6_10

remove it. The process stops when the root is isolated. We let $\mathcal{K}_e(\mathbb{G}_n)$ denote the number of cuts needed for the process to end.

Our model can also be applied to botnets, i.e., malicious computer networks consisting of compromised machines which are often used in spamming or attacks. The nodes in $\mathbb{G}_n$ represent the computers in a botnet, and the root represents the bot-master. The effectiveness of a botnet can be measured using the size of the component containing the root, which indicates the resources available to the bot-master [6]. To take down a botnet means to reduce the size of this root component as much as possible. If we assume that we target infected computers uniformly at random and it takes at least k attempts to fix a computer, then the k-cut number measures how difficult it is to completely isolate the bot-master.

The case $k = 1$ and $\mathbb{G}_n$ being a rooted tree has aroused great interests among mathematicians in the past few decades. The edge version of one-cut was first introduced by Meir and Moon [14] for the uniform random Cayley tree. Janson [12,13] noticed the equivalence between one-cuts and records in trees and studied them in binary trees and conditional Galton-Watson trees. Later Addario-Berry, Broutin, and Holmgren [1] gave a simpler proof for the limit distribution of one-cuts in conditional Galton-Watson trees. For one-cuts in random recursive trees, see [7,11,15]. For binary search trees and split trees, see [9,10].

1.2 The k-cut Number of a Tree

One of the most interesting cases is when $\mathbb{G}_n = \mathbb{T}_n$, where $\mathbb{T}_n$ is a rooted tree with n nodes.

There is an equivalent way to define $\mathcal{K}(\mathbb{T}_n)$. Imagine that each node is given an alarm clock. At time zero, the alarm clock of node v is set to ring at time $T_{1,v}$, where $(T_{i,v})_{i\geq 1, v\in\mathbb{T}_n}$ are i.i.d. (independent and identically distributed) Exp(1) random variables. After the alarm clock of node v rings the i-th time, we set it to ring again at time $T_{i+1,v}$. Due to the memoryless property of exponential random variables (see [8, pp. 134]), at any moment, which alarm clock rings next is always uniformly distributed. Thus, if we cut a node that is still in the tree when its alarm clock rings, and remove the node with its descendants if it has already been cut k-times, then we get exactly the k-cut model. (The random variables $(T_{i,v})_{i\geq 1}$ can be seen as the holding times in a Poisson process $N(t)_v$ of parameter 1, where $N(t)_v$ is the number of cuts in v during the time $[0, t]$ and has a Poisson distribution with parameter t.)

How can we tell if a node is still in the tree? When node v's alarm clock rings for the r-th time for some $r \leq k$, and no node above v has already rung k times, we say v has become an *r-record*. And when a node becomes an r-record, it must still be in the tree. Thus, summing the number of r-records over $r \in \{1, \dots, k\}$, we again get the k-cut number $\mathcal{K}(\mathbb{T}_n)$. One node can be a 1-record, a 2-record, etc., at the same time, so it can be counted multiple times. Note that if a node is an r-record, then it must also be a i-record for $i \in \{1, \dots, r-1\}$.

To be more precise, we define $\mathcal{K}(\mathbb{T}_n)$ as a function of $(T_{i,v})_{i\geq 1, v\geq 1}$. Let

$$G_{r,v} \stackrel{\text{def}}{=} \sum_{i=1}^{r} T_{i,v},$$

i.e., $G_{r,v}$ is the moment when the alarm clock of node v rings for the r-th time. Then $G_{r,v}$ has a gamma distribution with parameters $(r, 1)$ (see [8, Theorem 2.1.12]), which we denote by Gamma(r). Let

$$I_{r,v} \stackrel{\text{def}}{=} [\![G_{r,v} < \min\{G_{k,u} : u \in \mathbb{T}_n, u \text{ is an ancestor of } v\}]\!], \tag{1.1}$$

where $[\![\cdot]\!]$ denotes the Iverson bracket, i.e., $[\![S]\!] = 1$ if the statement S is true and $[\![S]\!] = 0$ otherwise. In other words, $I_{r,v}$ is the indicator random variable for node v being an r-record. Let

$$\mathcal{K}_r(\mathbb{T}_n) \stackrel{\text{def}}{=} \sum_{v\in\mathbb{T}_n} I_{r,v}, \qquad \mathcal{K}(\mathbb{T}_n) \stackrel{\text{def}}{=} \sum_{r=1}^{k} \mathcal{K}_r(\mathbb{T}_n).$$

Then $\mathcal{K}_r(\mathbb{T}_n)$ is the number of r-records and $\mathcal{K}(\mathbb{T}_n)$ is the total number of records.

1.3 The k-cut Number of a Path

Let $\mathbb{P}_n$ be a one-ary tree (a path) consisting of n nodes labeled $1, \ldots, n$ from the root to the leaf. To simplify notations, from now on we use $I_{r,i}, G_{r,i}$, and $T_{r,i}$ to represent $I_{r,v}, G_{r,v}$ and $T_{r,v}$ respectively for a node v at depth i.

Let $\mathcal{X}_n \stackrel{\text{def}}{=} \mathcal{K}(\mathbb{P}_n)$ and $\mathcal{X}_{n,r} = \mathcal{K}_r(\mathbb{P}_n)$. In this paper, we mainly consider $\mathcal{X}_n$ and we let $k \geq 2$ be a fixed integer.

The first motivation of this choice is that, as shown in Sect. 4, $\mathbb{P}_n$ is the fastest to cut among all graphs. (We make this statement precise in Lemma 4.) Thus $\mathcal{X}_n$ provides a universal stochastic lower bound for $\mathcal{K}(\mathbb{G}_n)$. Moreover, our results on $\mathcal{X}_n$ can immediately be extended to some trees of simple structures: see Sect. 4. Finally, as shown below, $\mathcal{X}_n$ generalizes the well-known record number in permutations and has very different behavior when $k = 1$, the usual cut-model, and $k \geq 2$, our extended model.

The name record comes from the classic definition of *records* in random permutations. Let $\sigma_1, \ldots, \sigma_n$ be a uniform random permutation of $\{1, \ldots, n\}$. If $\sigma_i < \min_{1\leq j<i} \sigma_j$, then i is called a *(strictly lower) record*. Let $\mathcal{R}_n$ denote the number of records in $\sigma_1, \ldots, \sigma_n$. Let $W_1, \ldots, W_n$ be i.i.d. random variables with a common continuous distribution. Since the relative order of $W_1, \ldots, W_n$ also gives a uniform random permutation, we can equivalently define σ_i as the rank of W_i. As gamma distributions are continuous, we can in fact let $W_i = G_{k,i}$. Thus, being a record in a uniform permutation is equivalent to being a k-record and $\mathcal{R}_n \stackrel{\mathcal{L}}{=} \mathcal{X}_{n,k}$. Moreover, when $k = 1$, $\mathcal{R}_n \stackrel{\mathcal{L}}{=} \mathcal{X}_n$.

Starting from Chandler's article [5] in 1952, the theory of records has been widely studied due to its applications in statistics, computer science, and physics. For more recent surveys on this topic, see [2].

A well-known result of $\mathcal{R}_n$ (and thus also $\mathcal{X}_{n,k}$) [16] is that $(I_{k,j})_{1\le j\le n}$ are independent. It follows from the Lindeberg–Lévy–Feller Theorem that

$$\frac{\boldsymbol{E}\left[\mathcal{R}_n\right]}{\log n} \to 1, \qquad \frac{\mathcal{R}_n}{\log n} \overset{a.s.}{\to} 1, \qquad \mathcal{L}\left(\frac{\mathcal{R}_n - \log n}{\sqrt{\log n}}\right) \overset{d}{\to} \mathcal{N}(0,1),$$

where $\mathcal{N}(0,1)$ denotes the standard normal distribution.

In the following, Theorem 1 gives the expectation of $\mathcal{X}_{n,r}$ which implies that the number of one-records dominates the number of other records. Subsequently Theorems 2 and 3 estimate the variance and higher moments of $\mathcal{X}_{n,1}$.

Theorem 1. *For all fixed $k \in \mathbb{N}$,*

$$\boldsymbol{E}\left[\mathcal{X}_{n,r}\right] \sim \begin{cases} \eta_{k,r} n^{1-\frac{r}{k}} & (1 \le r < k), \\ \log n & (r = k), \end{cases}$$

where the constants $\eta_{k,r}$ are defined by

$$\eta_{k,r} \overset{\text{def}}{=} \frac{(k!)^{\frac{r}{k}}}{k-r} \frac{\Gamma\left(\frac{r}{k}\right)}{\Gamma(r)},$$

where $\Gamma(z)$ denotes the gamma function. Therefore $\boldsymbol{E}\left[\mathcal{X}_n\right] \sim \boldsymbol{E}\left[\mathcal{X}_{n,1}\right]$. Also, for $k = 2$,

$$\boldsymbol{E}\left[\mathcal{X}_n\right] \sim \boldsymbol{E}\left[\mathcal{X}_{n,1}\right] \sim \sqrt{2\pi n}.$$

Theorem 2. *For all fixed $k \in \{2, 3, \dots\}$,*

$$\boldsymbol{E}\left[\mathcal{X}_{n,1}(\mathcal{X}_{n,1} - 1)\right] \sim \boldsymbol{E}\left[\left(\mathcal{X}_{n,1}\right)^2\right] \sim \gamma_k n^{2-\frac{2}{k}},$$

where

$$\gamma_k = \frac{\Gamma\left(\frac{2}{k}\right)(k!)^{\frac{2}{k}}}{k-1} + 2\lambda_k,$$

and

$$\lambda_k = \begin{cases} \dfrac{\pi \cot\left(\frac{\pi}{k}\right) \Gamma\left(\frac{2}{k}\right)(k!)^{\frac{2}{k}}}{2(k-2)(k-1)} & k > 2, \\ \dfrac{\pi^2}{4} & k = 2. \end{cases}$$

Therefore

$$\operatorname{Var}\left(\mathcal{X}_{n,1}\right) \sim \left(\gamma_k - \eta_{k,1}^2\right) n^{2-\frac{2}{k}}.$$

In particular, when $k = 2$

$$\operatorname{Var}\left(\mathcal{X}_{n,1}\right) \sim \left(\frac{\pi^2}{2} + 2 - 2\pi\right) n.$$

Theorem 3. *For all fixed $k \in \{2, 3, \dots\}$ and $\ell \in \mathbb{N}$*

$$\limsup_{n\to\infty} \boldsymbol{E}\left[\left(\frac{\mathcal{X}_{n,1}}{n^{1-\frac{1}{k}}}\right)^{\ell}\right] \le \rho_{k,\ell} \stackrel{\text{def}}{=} \ell! \Gamma\left(\ell + 1 - \frac{\ell}{k}\right)^{-1} \left(\frac{\pi}{k}(k!)^{1/k} \sin\left(\frac{\pi}{k}\right)^{-1}\right)^{\ell}.$$

The upper bound is tight for $\ell = 1$ since $\rho_{k,1} = \eta_{k,1}$.

The above theorems imply that the correct rescaling parameter should be $n^{1-\frac{1}{k}}$. However, unlike the case $k = 1$, when $k \ge 2$ the limit distribution of $\mathcal{X}_n/n^{1-\frac{1}{k}}$ has a rather complicated representation $\mathcal{B}_k$ defined as follows: Let $U_1, E_1, U_2, E_2, \dots$ be mutually independent random variables with $E_j \stackrel{\mathcal{L}}{=} \mathrm{Exp}(1)$ and $U_j \stackrel{\mathcal{L}}{=} \mathrm{Unif}[0,1]$. Let

$$\begin{aligned} S_p &\stackrel{\text{def}}{=} \left(k! \sum_{1 \le s \le p} \left(\prod_{s \le j < p} U_j\right) E_s\right)^{\frac{1}{k}}, \\ B_p &\stackrel{\text{def}}{=} (1 - U_p) \left(\prod_{1 \le j < p} U_j\right)^{1-\frac{1}{k}} S_p, \\ \mathcal{B}_k &\stackrel{\text{def}}{=} \sum_{1 \le p} B_p, \end{aligned}$$

where we use the convention that an empty product equals one.

Remark 1. An equivalent recursive definition of S_p is

$$S_p = \begin{cases} k! E_1 & (p = 1), \\ \left(U_{p-1} S_{p-1}^k + k! E_p\right)^{\frac{1}{k}} & (p \ge 2). \end{cases}$$

Theorem 4. *Let $k \in \{2, 3, \dots\}$. Let $\mathcal{L}(\mathcal{B}_k)$ denote the distribution of $\mathcal{B}_k$. Then*

$$\mathcal{L}\left(\frac{\mathcal{X}_n}{n^{1-\frac{1}{k}}}\right) \xrightarrow{d} \mathcal{L}(\mathcal{B}_k).$$

Thus, by Theorems 1, 2 and 3, the convergence also holds in L^p for all $p > 0$ and

$$\boldsymbol{E}[\mathcal{B}_k] = \eta_{k,1}, \qquad \boldsymbol{E}\left[\mathcal{B}_k^2\right] = \gamma_k, \qquad \boldsymbol{E}\left[\mathcal{B}_k^p\right] \in [\eta_{k,1}^p, \rho_{k,p}] \qquad (p \in \mathbb{N}).$$

Remark 2. It is easy to see that $\mathcal{X}^{\mathrm{e}}_{n+1} \stackrel{\text{def}}{=} \mathcal{K}_e(P_{n+1}) \stackrel{\mathcal{L}}{=} \mathcal{X}_n$ by treating each edge on a length $n + 1$ path as a node on a length n path.

The rest of the paper is organized as follows: Sect. 2 sketches the proofs for the moment results Theorems 1, 2, and 3. Section 3 deals with the distributional result Theorem 4. Section 4 discusses some easy results for general graphs. Finally, Sect. 5 collects analytic results used in the proofs, which may themselves be of interest. For detailed proofs, see the full version of this paper [3]. For k-cuts in complete binary trees, see our follow-up paper [4].

2 The Moments

2.1 The Expectation

Lemma 1. *Uniformly for all* $i \geq 1$ *and* $r \in \{1, \dots, k\}$,

$$\boldsymbol{E}\left[I_{r,i+1}\right] = \left(1 + O\left(i^{-\frac{1}{2k}}\right)\right) \frac{(k!)^{\frac{r}{k}}}{k} \frac{\Gamma\left(\frac{r}{k}\right)}{\Gamma(r)} i^{-\frac{r}{k}}.$$

Proof. By (1.1), $\boldsymbol{E}\left[I_{r,i+1}\right] = \boldsymbol{P}\left\{G_{k,1} > G_{r,i+1}, \dots, G_{k,i} > G_{r,i+1}\right\}$. Conditioning on $G_{r,i+1} = x$ yields $\boldsymbol{E}\left[I_{r,i+1}\right] = \int_0^\infty x^{r-1} \mathrm{e}^{-x} / \Gamma(r) \boldsymbol{P}\left\{G_{k,1} > x\right\}^i \,\mathrm{d}x$. Lemma 1 thus follows from Lemma 7.

Proof (Proof of Theorem 1). A simply computation shows that for $a \in (0,1)$

$$\sum_{1 \leq i \leq n} \frac{1}{i^a} = \frac{1}{1-a} n^{1-a} + O(1).$$

It then follows from Lemma 1 that for $r \in \{1, \dots, k-1\}$.

$$\boldsymbol{E}\left[\mathcal{X}_{n,r}\right] = \sum_{0 \leq i < n} \boldsymbol{E}\left[I_{r,i+1}\right] = \frac{(k!)^{\frac{r}{k}}}{k} \frac{\Gamma\left(\frac{r}{k}\right)}{\Gamma(r)} \frac{1}{1-\frac{r}{k}} n^{1-\frac{r}{k}} + O\left(n^{1-\frac{r}{k}-\frac{1}{2k}}\right) + O(1).$$

When $r = k$, $\boldsymbol{E}\left[\mathcal{X}_{n,k}\right] = \boldsymbol{E}\left[\mathcal{R}_n\right] \sim \log(n)$ is already well-known.

2.2 The Variance

In this section we prove Theorem 2.

Let $E_{i,j}$ denote the event that $[I_{1,i+1} I_{1,j+1} = 1]$. Let $A_{x,y}$ denote the event that $[G_{1,i+1} = x \cap G_{1,j+1} = y]$. Then conditioning on $A_{x,y}$

$$E_{i,j} = \left[\bigcap_{1 \leq s \leq i} G_{k,s} > x \vee y\right] \cap \left[G_{k,i+1} > y\right] \cap \left[\bigcap_{i+2 \leq s \leq j} G_{k,s} > y\right],$$

where $x \vee y \stackrel{\text{def}}{=} \max\{x, y\}$. Since conditioning on $A_{x,y}$, $G_{k,i+1} \stackrel{\mathcal{L}}{=} \text{Gamma}(k-1) + x$, $G_{k,s} \stackrel{\mathcal{L}}{=} \text{Gamma}(k)$ for $s \notin \{i+1, j+1\}$, and all these random variables are independent, we have

$$\boldsymbol{P}\left\{E_{i,j} | A_{x,y}\right\} = \boldsymbol{P}\left\{G_{k-1,1} + x > y\right\} \boldsymbol{P}\left\{G_{k,1} > x \vee y\right\}^i \boldsymbol{P}\left\{G_{k,1} > y\right\}^{j-i-1}.$$

It follows from $G_{1,i+1} \stackrel{\mathcal{L}}{=} G_{1,j+1} \stackrel{\mathcal{L}}{=} \text{Exp}(1)$ that

$$\begin{aligned} \boldsymbol{P}\left\{E_{i,j}\right\} &= \int_0^\infty \int_y^\infty \mathrm{e}^{-x-y} \boldsymbol{P}\left\{E_{i,j} | A_{x,y}\right\} \,\mathrm{d}x \,\mathrm{d}y \\ &\quad + \int_0^\infty \int_0^y \mathrm{e}^{-x-y} \boldsymbol{P}\left\{E_{i,j} | A_{x,y}\right\} \,\mathrm{d}x \,\mathrm{d}y \\ &\stackrel{\text{def}}{=} A_{1,i,j} + A_{2,i,j}. \end{aligned}$$

Thus Theorem 2 follows from $\mathcal{X}_{n,1}(\mathcal{X}_{n,1} - 1) = 2 \sum_{1 \leq i < j \leq n} I_{1,i} I_{1,j}$ and the following two lemmas whose proofs rely on Lemmas 8, 9, 10.

Lemma 2. *Let $k \in \{2, 3, \dots\}$. We have*

$$A_{2,i,j} = \left(1 + O\left(j^{-\frac{1}{2k}}\right)\right) \frac{(k!)^{\frac{2}{k}}}{k} \Gamma\left(\frac{2}{k}\right) j^{-\frac{2}{k}}.$$

Lemma 3. *Let $k \in \{2, 3, \dots\}$. Let $a = i$ and $b = j - i - 1$. Then for all $a \geq 1$ and $b \geq 1$,*

$$A_{1,i,j} = \xi_k(a, b) + O\left(\left(a^{-\frac{1}{2k}} + b^{-\frac{1}{2k}}\right)\left(a^{-\frac{2}{k}} + b^{-\frac{2}{k}}\right)\right),$$

where

$$\xi_k(a, b) \stackrel{\text{def}}{=} \int_0^\infty \int_y^\infty \exp\left(-a\frac{x^k}{k!} - b\frac{y^k}{k!}\right) \mathrm{d}x\, \mathrm{d}y.$$

2.3 Higher Moments

The computations of higher moments of $\mathcal{X}_{n,1}$ are rather complicated. However, an upper bound is readily available. Let $1 \leq i_1 < i_2 < \cdots < i_\ell \leq n$. Then

$$\boldsymbol{E}\left[I_{1,i_1} I_{1,i_2} \cdots I_{1,i_\ell}\right] \leq \boldsymbol{E}\left[I_{1,i_1}\right] \boldsymbol{E}\left[I_{1,i_2-i_1}\right] \cdots \boldsymbol{E}\left[I_{1,i_\ell - i_{\ell-1}}\right].$$

The above inequality holds since if i_j is a one-record in the whole path, then it must also be a one-record in the segment $(i_{j-1}+1, \dots, i_j)$ ignoring everything else, and what happens in each of such segments are independent. Theorem 3 follows easily from this observation.

3 Convergence to the *k*-cut Distribution

By Theorem 1 and Markov's inequality, $\mathcal{X}_{n,r}/n^{1-\frac{1}{k}} \xrightarrow{p} 0$ for $r \in \{2, \dots, k\}$. So it suffices to prove Theorem 4 for $\mathcal{X}_{n,1}$ instead of $\mathcal{X}_n$. Throughout Sect. 3, unless otherwise emphasized, we assume that $k \geq 2$.

The idea of the proof is to condition on the positions and values of the k-records, and study the distribution of the number of one-records between two consecutive k-records.

We use $(R_{n,p})_{p \geq 1}$ to denote the k-record values and $(P_{n,p})_{p \geq 1}$ the positions of these k-records. To be precise, let $R_{n,0} \stackrel{\text{def}}{=} 0$, and $P_{n,0} \stackrel{\text{def}}{=} n+1$; for $p \geq 1$, if $P_{n,p-1} > 1$, then let

$$R_{n,p} \stackrel{\text{def}}{=} \min\{G_{k,j} : 1 \leq j < P_{n,p-1}\},$$
$$P_{n,p} \stackrel{\text{def}}{=} \operatorname{argmin}\{G_{k,j} : 1 \leq j < P_{n,p-1}\},$$

i.e., $P_{n,p}$ is the unique positive integer which satisfies that $G_{k,P_{n,p}} \leq G_{k,i}$ for all $1 \leq i < P_{n,p-1}$; otherwise let $P_{n,p} = 1$ and $R_{n,p} = \infty$. Note that $R_{n,1}$ is simply the minimum of n i.i.d. Gamma(k) random variables.

According to $(P_{n,p})_{p\geq 1}$, we split $\mathcal{X}_{n,1}$ into the following sum

$$\mathcal{X}_{n,1} = \sum_{1\leq j\leq n} I_{1,j} = \mathcal{X}_{n,k} + \sum_{1\leq p}\sum_{1\leq j} [\![P_{n,p-1} > j > P_{n,p}]\!]\, I_{1,j} \stackrel{\text{def}}{=} \mathcal{X}_{n,k} + \sum_{1\leq p} B_{n,p}. \tag{3.1}$$

Figure 1 gives an example of $(B_{n,p})_{p\geq 1}$ for $n = 12$. It depicts the positions of the k-records and the one-records. It also shows the values and the summation ranges for $(B_{n,p})_{p\geq 1}$.

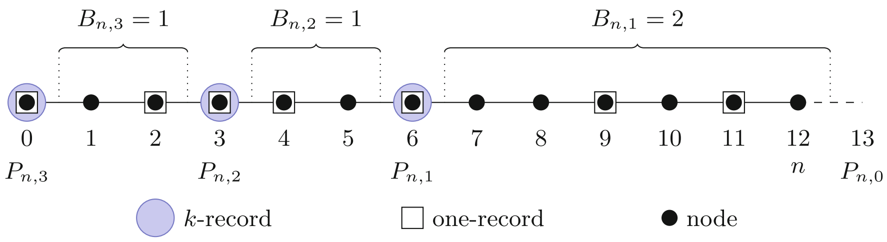

Fig. 1. An example of $(B_{n,p})_{p\geq 1}$ for $n = 12$.

Recall that $T_{r,j} \stackrel{\mathcal{L}}{=} \text{Exp}(1)$, is the lapse of time between the alarm clock of j rings for the $(r-1)$-st time and the r-th time. Conditioning on $(R_{n,p}, P_{n,p})_{n\geq 1, p\geq 1}$, for $j \in (P_{n,p}, P_{n,p-1})$, we have

$$\boldsymbol{E}\left[I_{1,j}\right] = \boldsymbol{P}\left\{T_{1,j} < R_{n,p} \,|G_{k,j} > R_{n,p-1}\right\}.$$

Then the distribution of $B_{n,p}$ conditioning on $(R_{n,p}, P_{n,p})_{n\geq 1, p\geq 1}$ is simply that of

$$\text{Bin}\left(P_{n,p-1} - P_{n,p} - 1, \boldsymbol{P}\left\{T_{1,j} < R_{n,p} \,|G_{k,j} > R_{n,p-1}\right\}\right),$$

where $\text{Bin}(m, p)$ denotes a binomial (m, p) random variable. When $R_{n,p-1}$ is small and $P_{n,p-1} - P_{n,p}$ is large, this is roughly

$$\text{Bin}\left(P_{n,p-1} - P_{n,p}, \boldsymbol{P}\left\{T_{1,j} < R_{n,p}\right\}\right) \stackrel{\mathcal{L}}{=} \text{Bin}\left(P_{n,p-1} - P_{n,p}, 1 - e^{-R_{n,p}}\right). \tag{3.2}$$

Therefore, we first study a slightly simplified model. Let $(T^*_{r,j})_{r\geq 1, j\geq 1}$ be i.i.d. Exp(1) which are also independent from $(T_{r,j})_{r\geq 1, j\geq 1}$. Let

$$I^*_j \stackrel{\text{def}}{=} [\![T^*_{1,j} < \min\{G_{k,i} : 1 \leq i \leq j\}]\!], \qquad \mathcal{X}^*_n \stackrel{\text{def}}{=} \sum_{1\leq j\leq n} I^*_j.$$

We say a node j is an *alt-one-record* if $I^*_j = 1$. As in (3.1), we can write

$$\mathcal{X}^*_n = \sum_{1\leq j\leq n} I^*_j = \sum_{1\leq p}\sum_{1\leq j} [\![P_{n,p-1} > j \geq P_{n,p}]\!]\, I^*_j \stackrel{\text{def}}{=} \sum_{1\leq p} B^*_{n,p}.$$

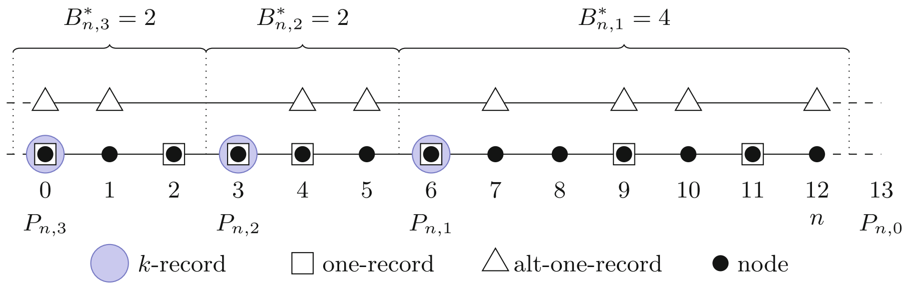

Fig. 2. An example of $(B^*_{n,p})_{p\ge 1}$ for $n = 12$.

Then conditioning on $(R_{n,p}, P_{n,p})_{n\ge 1,p\ge 1}$, $B^*_{n,p}$ has exactly the distribution as (3.2). Figure 2 gives an example of $(B^*_{n,p})_{p\ge 1}$ for $n = 12$. It shows the positions of alt-one-records, as well as the values and the summation ranges of$(B^*_{n,p})_{p\ge 1}$.

The main part of the proof for Theorem 4 consist of showing the following

Proposition 1. *For all fixed $p \in \mathbb{N}$ and $k \ge 2$,*

$$\mathcal{L}\left(\left(\frac{B^*_{n,1}}{n^{1-\frac{1}{k}}}, \dots, \frac{B^*_{n,p}}{n^{1-\frac{1}{k}}}\right)\right) \xrightarrow{d} \mathcal{L}\left((B_1, \dots B_p)\right),$$

which implies by the Cramér–Wold device that

$$\mathcal{L}\left(\sum_{1\le j\le p} \frac{B^*_{n,j}}{n^{1-\frac{1}{k}}}\right) \xrightarrow{d} \mathcal{L}\left(\sum_{1\le j\le p} B_j\right),$$

Then we can prove that p can be chosen large enough so that $\sum_{p<j} B^*_{n,j}/n^{1-\frac{1}{k}}$ is negligible. Thus,

$$\mathcal{L}\left(\frac{\mathcal{X}^*_n}{n^{1-\frac{1}{k}}}\right) \overset{\text{def}}{=} \mathcal{L}\left(\frac{\sum_{1\le j} B^*_{n,j}}{n^{1-\frac{1}{k}}}\right) \xrightarrow{d} \mathcal{L}\left(\sum_{1\le j} B_j\right) \overset{\text{def}}{=} \mathcal{L}\left(\mathcal{B}_k\right).$$

Following this, we can use a coupling argument to show that $\mathcal{X}_{n,1}/n^{1-\frac{1}{k}}$ and $\mathcal{X}^*_n/n^{1-\frac{1}{k}}$ converge to the same limit, which finishes the proof of Theorem 4.

4 Some Extensions

4.1 A Lower Bound and an Upper Bound for General Graphs

Let $\mathcal{G}_n$ be the set of rooted graphs with n nodes. It is obvious that $\mathbb{P}_n$ is the easiest to cut among all graphs in $\mathcal{G}_n$. We formalize this by the following lemma:

Lemma 4. *Let $k \in \mathbb{N}$. For all $\mathbb{G}_n \in \mathcal{G}_n$, $\mathcal{X}_n \stackrel{\text{def}}{=} \mathcal{K}(\mathbb{P}_n) \preceq \mathcal{K}(\mathbb{G}_n)$. Therefore,*

$$\min_{\mathbb{G}_n \in \mathcal{G}_n} \boldsymbol{E}\mathcal{K}(\mathbb{G}_n) \geq \boldsymbol{E}\mathcal{X}_n \sim \begin{cases} \frac{(k!)^{\frac{1}{k}}}{k-1} \Gamma\left(\frac{1}{k}\right) n^{1-\frac{1}{k}} & (k \geq 2), \\ \log n & (k = 1), \end{cases}$$

by Theorem 1.

The most resilient graph is obviously $\mathbb{K}_n$, the complete graph with n vertices. Thus, we have the following upper bound:

Lemma 5. *Let $k \in \mathbb{N}$.*

(i) Let $Y \stackrel{\mathcal{L}}{=} \text{Gamma}(k)$, $Z \stackrel{\mathcal{L}}{=} \text{Poi}(Y)$, and $W \stackrel{\mathcal{L}}{=} Z \wedge k$, i.e., $W \stackrel{\mathcal{L}}{=} \min\{Z, k\}$. Then

$$\mathcal{L}\left(\frac{\mathcal{K}(\mathbb{K}_n)}{n}\right) \xrightarrow{d} \mathcal{L}\left(\boldsymbol{E}\left[W|Y\right]\right) = \mathcal{L}\left(\frac{\Gamma(k+1, Y) - e^{-Y} Y^{k+1}}{k!} + k\right),$$

where $\Gamma(\ell, z)$ denotes the upper incomplete gamma function. Note that when $k = 1$, the right-hand-side is simply Unif$[0, 1]$.

(ii) For all $\mathbb{G}_n \in \mathcal{G}_n$, $\mathcal{K}(\mathbb{G}_n) \preceq \mathcal{K}(\mathbb{K}_n)$. Therefore,

$$\max_{\mathbb{G}_n \in \mathcal{G}_n} \boldsymbol{E}\mathcal{K}(\mathbb{G}_n) \leq \boldsymbol{E}\mathcal{K}(\mathbb{K}_n) \sim k\left(1 - \frac{1}{2^{2k}}\binom{2k}{k}\right) n.$$

4.2 Path-Like Graphs

If a graph $\mathbb{G}_n$ consists of only long paths, then the limit distribution $\mathcal{K}(\mathbb{G}_n)$ should be related to $\mathcal{B}_k$, the limit distribution of $\mathcal{K}(\mathbb{P}_n)/n^{1-\frac{1}{k}}$ (see Theorem 4). We give two simple examples with $k \in \{2, 3, \dots\}$.

Example 1 (Long path). Let $(\mathbb{G}_n)_{n \geq 1}$ be a sequence of rooted graphs such that $\mathbb{G}_n$ contains a path of length $m(n)$ starting from the root with $n - m(n) = o(n^{1-\frac{1}{k}})$. Since it takes at most $k(n - m(n))$ cuts to remove all the nodes outside the long path,

$$\mathcal{K}(P_{m(n)}) \preceq \mathcal{K}(\mathbb{G}_n) \preceq \mathcal{K}(P_{m(n)}) + ko\left(n^{1-1/k}\right).$$

Thus, by Lemma 4, this implies that $\mathcal{K}(\mathbb{G}_n)/n^{1-\frac{1}{k}}$ converges in distribution to $\mathcal{B}_k$.

5 Some Auxiliary Results

Lemma 6. *Let $G_k \stackrel{\mathcal{L}}{=} \text{Gamma}(k)$. Let $\alpha \stackrel{\text{def}}{=} \frac{1}{2}\left(\frac{1}{k} + \frac{1}{k+1}\right)$ and $x_0 \stackrel{\text{def}}{=} m^{-\alpha}$. Then uniformly for all $x \in [0, x_0]$,*

$$\boldsymbol{P}\left\{G_k > x\right\}^m = \left(\frac{\Gamma\left(k, x\right)}{\Gamma(k)}\right)^m = \left(1 + O\left(m^{-\frac{1}{2k}}\right)\right) \exp\left(-\frac{m x^k}{k!}\right),$$

where $\Gamma(\ell, z)$ denotes the upper incomplete gamma function.

Lemma 7. *Let* $G_k \stackrel{\mathcal{L}}{=} \mathrm{Gamma}(k)$. *Let* $a \geq 0$ *and* $b \geq 1$ *be fixed. Then uniformly for* $m \geq 1$,

$$\int_0^\infty x^{b-1} \mathrm{e}^{-ax} \boldsymbol{P}\left\{G_k > x\right\}^m \,\mathrm{d}x = \left(1 + O\left(m^{-\frac{1}{2k}}\right)\right) \frac{(k!)^{\frac{b}{k}}}{k} \Gamma\left(\frac{b}{k}\right) m^{-\frac{b}{k}}.$$

Lemma 8. *For* $a > 0$, $b > 0$ *and* $k \geq 2$,

$$\begin{aligned} \xi_k(a,b) &\stackrel{\text{def}}{=} \int_0^\infty \int_y^\infty \mathrm{e}^{-ax^k/k! - by^k/k!} \,\mathrm{d}x \,\mathrm{d}y \\ &= \frac{\Gamma\left(\frac{2}{k}\right)}{k} \left(\frac{k!}{a}\right)^{\frac{2}{k}} F\left(\frac{2}{k}, \frac{1}{k}; 1 + \frac{1}{k}; -\frac{b}{a}\right), \end{aligned}$$

where F *denotes the hypergeometric function. In particular,*

$$\xi_2(a,b) = \arctan\left(\sqrt{\frac{b}{a}}\right)(ab)^{-\frac{1}{2}}.$$

Lemma 9. *For* $a > 0$, $b > 0$ *and* $k \geq 2$,

$$(a+b)^{-\frac{2}{k}} \leq \frac{k}{\Gamma\left(\frac{2}{k}\right)(k!)^{\frac{2}{k}}} \xi_k(a,b) \leq a^{-\frac{2}{k}} + b^{-\frac{2}{k}}.$$

Moreover, $\xi_k(a,b)$ *is monotonically decreasing in both* a *and* b.

Lemma 10. *For* $k \geq 2$, *let*

$$\lambda_k \stackrel{\text{def}}{=} \int_0^1 \int_0^{1-s} \xi_k(s,t) \,\mathrm{d}t \,\mathrm{d}s.$$

Then

$$\lambda_k = \begin{cases} \dfrac{\pi \cot\left(\frac{\pi}{k}\right) \Gamma\left(\frac{2}{k}\right)(k!)^{\frac{2}{k}}}{2\,(k-2)\,(k-1)} & k > 2, \\ \dfrac{\pi^2}{4} & k = 2. \end{cases}$$

References

1. Addario-Berry, L., Broutin, N., Holmgren, C.: Cutting down trees with a Markov chainsaw. Ann. Appl. Probab. **24**(6), 2297–2339 (2014)
2. Ahsanullah, M.: Record Values-Theory and Applications. University Press of America Inc., Lanham (2004)
3. Cai, X.S., Devroye, L., Holmgren, C., Skerman, F.: k-cut on paths and some trees. ArXiv e-prints, January 2019

4. Cai, X.S., Holmgren, C.: Cutting resilient networks - complete binary trees. arXiv e-prints, November 2018
5. Chandler, K.N.: The distribution and frequency of record values. J. R. Stat. Soc. Ser. B. **14**, 220–228 (1952)
6. Dagon, D., Gu, G., Lee, C.P., Lee, W.: A taxonomy of botnet structures. In: Twenty-Third Annual Computer Security Applications Conference (ACSAC 2007), pp. 325–339 (2007)
7. Drmota, M., Iksanov, A., Moehle, M., Roesler, U.: A limiting distribution for the number of cuts needed to isolate the root of a random recursive tree. Random Struct. Algorithms **34**(3), 319–336 (2009)
8. Durrett, R.: Probability: Theory and Examples, Cambridge Series in Statistical and Probabilistic Mathematics, vol. 31, 4th edn. Cambridge University Press, Cambridge (2010)
9. Holmgren, C.: Random records and cuttings in binary search trees. Combin. Probab. Comput. **19**(3), 391–424 (2010)
10. Holmgren, C.: A weakly 1-stable distribution for the number of random records and cuttings in split trees. Adv. Appl. Probab. **43**(1), 151–177 (2011)
11. Iksanov, A., Möhle, M.: A probabilistic proof of a weak limit law for the number of cuts needed to isolate the root of a random recursive tree. Electron. Comm. Probab. **12**, 28–35 (2007)
12. Janson, S.: Random records and cuttings in complete binary trees. In: Mathematics and Computer Science III, Trends Math, pp. 241–253. Birkhäuser, Basel (2004)
13. Janson, S.: Random cutting and records in deterministic and random trees. Random Struct. Algorithms **29**(2), 139–179 (2006)
14. Meir, A., Moon, J.W.: Cutting down random trees. J. Austral. Math. Soc. **11**, 313–324 (1970)
15. Meir, A., Moon, J.: Cutting down recursive trees. Math. Biosci. **21**(3), 173–181 (1974)
16. Rényi, A.: Théorie des éléments saillants d'une suite d'observations. Ann. Fac. Sci. Univ. Clermont-Ferrand No. **8**, 7–13 (1962)