| COMPSCI 330: Design and Analysis of Algorithms | $3 / 22 / 2016$ |
| :--- | ---: |
| Las Vegas and Monte Carlo Algorithms |  |
| Lecturer: Debmalya Panigrahi | Scribe: Tianqi Song |

## 1 Overview

In this lecture, first we cover the birthday paradox. Then, we will introduce Monte Carlo and Las Vegas algorithms. ${ }^{1}$

## 2 The Birthday Paradox

### 2.1 Problem Statement

There are $k$ people and each person has a random birthday from $n$ days. How large should $k$ be such that the chance that two people have the same birthday is at least $\alpha$ ?

### 2.2 Analysis

Let $X$ be the event that every person has a distinct birthday. The $i$ th person has ( $n-i$ ) options to have a distinct birthday and then the probability to have a distinct birthday is $\frac{n-i}{n}$, where $i \in\{0,1,2, \ldots, k-1\}$. We have:

$$
\begin{align*}
\operatorname{Pr}[X] & =\frac{n}{n}\left(\frac{n-1}{n}\right)\left(\frac{n-2}{n}\right) \ldots\left(\frac{n-(k-1)}{n}\right)  \tag{1}\\
& =1\left(1-\frac{1}{n}\right)\left(1-\frac{2}{n}\right) \ldots\left(1-\frac{k-1}{n}\right) \tag{2}
\end{align*}
$$

Because $e^{-x} \approx 1-x$ when $x$ is small:

$$
\begin{align*}
\operatorname{Pr}[X] & \approx 1 \cdot e^{-\frac{1}{n}} \cdot e^{-\frac{2}{n}} \ldots e^{-\frac{k-1}{n}}  \tag{3}\\
& \approx e^{-\frac{(1+(k-1))(k-1)}{n}}  \tag{4}\\
& =e^{-\frac{k(k-1)}{2 n}} \tag{5}
\end{align*}
$$

We just need to solve $\operatorname{Pr}[X] \approx e^{-\frac{k(k-1)}{2 n}} \leq 1-\alpha$. For example, when $\alpha=0.5$ and $n=365$, we have $k \geq 23$.

## 3 Monte Carlo Algorithms

Definition 1. A randomized algorithm is called a Monte Carlo algorithm if it may fail or return incorrect answers, but has runtime independent of the randomness.

[^0]
### 3.1 Global Minimum Cut

Definition 2. Let $G=(V, E)$ be an undirected, unit-capacity graph, where $|E|=m$ and $|V|=n$. The global minimum cut problem is to find the smallest non-trivial cut in $G$. In other words, a complete partitioning of $V$ into $(S, T)$ which minimizes:

$$
|\{(v, w) \mid(v, w) \in E \cap(S \times T)\}|
$$

where both $S$ and $T$ are nonempty.

### 3.1.1 Karger 1993

A Monte Carlo algorithm by Karger in 1993 gives the global min-cut, with high probability, in $O(m)$ time. This algorithm is based around edge contractions:

Definition 3. Let $G=(V, E)$ be a (multi)graph, and $(u, v) \in E$. Contracting edge ( $u, v$ ) produces a new multigraph where:

1. Both $u$ and $v$ are replaced by a single new supervertex $w=\{u, v\}$.
2. Each edge between $u$ and $v$ is replaced by a self-loop $(w, w)$. Edges which have $u$ or $v$ as an endpoint now have $w$ as that endpoint.

The resulting graph may include multi-edges and self-loops.
Edge contractions can simplify the graph while preserving a cut $(S, T)$.
Fact 1. Let $(S, T)$ be a cut in $G$. Let $(v, w)$ be an edge where both $v, w \in S$ or both $v, w \in T$. Then if $G^{\prime}$ is the graph after contracting $(v, w)$, then $|(S, T)|$ is the same in $G^{\prime}$ as in $G$.

Lemma 2. Let $G^{\prime}$ be the graph after contracting $e$ in $G$. Every cut in $G^{\prime}$ corresponds a cut of the same size in $G$, on the same supervertices.

Proof. This follows directly from Fact 1 . Note, however, that the opposite direction does not hold: not every cut in $G^{\prime}$ reappears as a cut in $G$. In particular, the cuts in $G$ which included $e$ disappear in $G^{\prime}$.

Contraction preserves the size for cuts which don't use the contracted edge. Intuitively, contraction doesn't produce smaller cuts in $G^{\prime}$, which means the minimum cut remains the minimum cut.

Corollary 3. Let $(S, T)$ be the global minimum cut of $G$, and let $G^{\prime}$ be the graph after contracting $e \notin(S, T)$. Then $(S, T)$ is the global minimum cut of $G^{\prime}$, and has the same size.

The algorithm is as follows:

```
Algorithm 1 (Karger 1993)
    \(G_{n} \leftarrow G\)
    for \(i=(n-1)\) to 2 do
        \(G_{i} \leftarrow G_{i+1}\)
        Contract \(e \in E_{i}\), picked uniformly at random, and remove all self-loops.
    return Cut represented by two remaining supervertices of \(G_{2}\).
```


### 3.1.2 Success Probability

First, we relate the size minimum cut to the number of edges:
Lemma 4. Let $\lambda$ be the size of the minimum cut. On a graph with i vertices, $|E| \geq(i \lambda) / 2$.
Proof. Let $\operatorname{deg}(v)$ be the degree of $v$. Summing over all vertices double counts every edge.

$$
\sum_{v \in V} \operatorname{deg}(v)=2|E|
$$

The degree cut of any vertex gives an upper bound on the size of the minimum cut.

$$
\sum_{v \in V} \operatorname{deg}(v) \geq i \lambda
$$

Rearranging:

$$
|E| \geq \frac{i \lambda}{2}
$$

Theorem 5. For any given global mincut ( $S, T=V \backslash S$ ), Karger's algorithm outputs ( $S, T$ ) with probability $\Omega\left(1 / n^{2}\right)$.
Proof. For any cut to be output by the algorithm, none of its edges must be contracted before the end $\left(G_{2}\right)$. Let $\lambda$ be the number of edges in global min-cut $(S, T)$. Let $G_{i}=\left(V_{i}, E_{i}\right)$ be the graph at $i$ vertices.
$(S, T)$ disappears in $G_{i}$ if one of its edges are contracted. Let $\mathscr{E}_{i}$ be the event that no $(S, T)$ edge is contracted in the contraction on $G_{i}$ to form $G_{i-1}$.

$$
\begin{aligned}
\operatorname{Pr}((S, T) \text { successfully output }) & =\operatorname{Pr}\left(\mathscr{E}_{n} \cap \mathscr{E}_{n-1} \cap \cdots \cap \mathscr{E}_{2}\right) \\
& =\operatorname{Pr}\left(\mathscr{E}_{n}\right) \operatorname{Pr}\left(\mathscr{E}_{n-1} \mid \mathscr{E}_{n}\right) \operatorname{Pr}\left(\mathscr{E}_{n-2} \mid \mathscr{E}_{n}, \mathscr{E}_{n-1}\right) \cdots
\end{aligned}
$$

Lemma 4 tells us that:

$$
\left|E_{i}\right| \geq \frac{i \lambda}{2}
$$

Since the edges are chosen uniformly from $E_{i}$ :

$$
\begin{aligned}
\operatorname{Pr}\left(\mathscr{E}_{i} \mid \cap_{k=n}^{i+1} \mathscr{E}_{k}\right) & =1-\frac{\lambda}{\left|E_{i}\right|} \\
& \geq 1-\frac{\lambda}{(i \lambda) / 2} \\
& =1-\frac{2}{i}
\end{aligned}
$$

Replacing:

$$
\begin{aligned}
\operatorname{Pr}((S, T) \text { successfully output }) & =\operatorname{Pr}\left(\mathscr{E}_{n}\right) \operatorname{Pr}\left(\mathscr{E}_{n-1} \mid \mathscr{E}_{n}\right) \operatorname{Pr}\left(\mathscr{E}_{n-1} \cap \mathscr{E}_{n} \mathscr{E}_{n-1}\right) \cdots \\
& \geq\left(1-\frac{2}{n}\right)\left(1-\frac{2}{n-1}\right) \cdots\left(1-\frac{2}{3}\right) \\
& =\frac{n-2}{n} \times \frac{n-3}{n-1} \times \frac{n-4}{n-2} \times \cdots \times \frac{3}{5} \times \frac{2}{4} \times \frac{1}{3} \\
& =\frac{2 \cdot 1}{n(n-1)} \\
& =\binom{n}{2}^{-1}=\Theta\left(1 / n^{2}\right)
\end{aligned}
$$

Corollary 6. There are at most $\binom{n}{2}$ unique minimum cuts in a graph.
How can we obtain a high probability bound? Repeat the algorithm $O\left(n^{2} \log n\right)$ times and output the smallest cut.

$$
\begin{aligned}
\operatorname{Pr}(\text { None of the outputs were global minimum cut }) & \leq \prod_{t=1}^{n^{2} c \log n}\left(1-\frac{1}{n^{2}}\right) \\
& \leq\left(e^{-1 / n^{2}}\right)^{n^{2} c \log n} \\
& =O\left(1 / n^{c}\right) \\
\operatorname{Pr}(\text { At least one output was global minimum cut }) & =\Omega(1-1 / \operatorname{poly}(n))
\end{aligned}
$$

### 3.1.3 Running Time

The contraction algorithm we described can be implemented in $O\left(m n^{2} \log n\right)$ time.

## 4 Las Vegas Algorithms

Definition 4. A randomized algorithm is called a Las Vegas algorithm if it always returns the correct answer, but its runtime bounds hold only in expectation.

In Las Vegas algorithms, runtime is at the mercy of randomness, but the algorithm always succeeds in giving a correct answer.

### 4.1 Randomized Quicksort Analysis

Recall that the randomized quicksort algorithm picks a pivot at random, and then partitions the elements into three sets: all the elements less than the pivot, all elements equal to the pivot, and all elements greater than the pivot.


We analyze the runtime using charging. In the computation tree shown above, we count the number times each element appears. The sum over all elements is equal to the sum of the sizes of all the sets in the computation tree, which is proportional to the runtime.

Claim 1. The number of times an element appears in sets in the computation tree is $O(\log n)$ in expectation.
Each set $S$ in the computation tree has three children. Let $S_{=p}$ be child set with elements equal to the pivot. Let $S_{>p}$ and $S_{<p}$ be defined similarly. We color the tree edges to these children as follows:

- Color the edge to $S_{=p}$ red.
- Color the edge to the larger of $S_{>p}$ and $S_{<p}$ black.
- Color the edge to the smaller of $S_{>p}$ and $S_{<p}$ blue.

Now, we trace the path of some element $e$ through the computation tree, and count the number of edges of each color along the path.

- The number of red edges is 1 . Only the last edge can be red.
- The number of blue edges is at most $\log _{2} n$, since the smaller of $S_{>p}$ and $S_{<p}$ has at most $\frac{1}{2} \cdot|S|$ elements.
- The number of black edges, using the definitions we have given, could be large.

What is the probability that an edge is black? That is, given that $e \in S$, what is the probability that the next edge in $e$ 's path is black? Suppose $e$ is the smallest element in the array.


For this element, any pivot in the right half of the array, illustrated with the shaded region in the image above, will cause $e$ 's next edge to black. Therefore, the probability that the next edge is black in this case is roughly $\frac{1}{2}$.

However, consider some element $e$ near the center of the array.


The shaded region shows which pivots will cause the next edge in $e$ 's path to be black. For this element, the probability that the next edge is black is nearly 1 . The probability will be high for all elements near the center of the array.

We need to modify our definition of black edges and blue edges to fix the problem. If $\left|S_{>p}\right|>\frac{3}{4}|S|$, color the edge to $S_{>p}$ black. Otherwise, color the edge to $S_{>p}$ blue. We color the edge to $S_{<p}$ using identical rules. Note that now it is possible for both edges to be blue (but not both to be black). We have the following:

- The number of red edges is still 1 .
- The number of blue edges is at most $\log _{4 / 3} n$.

To calculate the probability of an edge being black, we split the array into four equal parts:


If the pivot is selected from the shaded quadrants, one of the edges will be blue and the other will be black. If the pivot is selected from the non-shaded quadrants, both edges will be blue. Therefore, regardless of which element $e \in S$ we pick, the probability that the next edge on $e$ 's path is black is at most $\frac{1}{2}$.

Now, we can bound the number of black edges in a path. Let $X_{i}$ a random variable, equal to the number of black edges between the $i$ th and the $(i+1)$ st blue edge. As we have seen,

$$
\mathbb{E}\left[X_{i}\right] \leq \frac{1}{p_{\text {success }}} \leq \frac{1}{1 / 2}=2 .
$$

By linearity of expectation, we have

$$
\mathbb{E}[\text { total number of black edges }]=\mathbb{E}\left[\sum_{i} X_{i}\right]=\sum_{i} \mathbb{E}\left[X_{i}\right] \leq 2\left(\log _{4 / 3} n\right)=O(\log n) .
$$

We now need to bound the total running time, $T$. Let $T_{e}$ be the amount of running time charged to $e$ (which is proportional to the number of sets $e$ appears in).

$$
\mathbb{E}[T]=\mathbb{E}\left[\sum_{e} T_{e}\right]=\sum_{e} \mathbb{E}\left[T_{e}\right]=\sum_{e} O(\log n)=O(n \log n) .
$$


[^0]:    ${ }^{1}$ Some materials are from a note by Samuel Haney for this class in Fall 2014 and a note by Allen Xiao for COMPSCI 532 in Fall 2015.

