A geometric Ramsey theorem of Suk

The IMOTC has apart from interesting problem solving sessions for the students in the camp, also included sessions that include what I would call modules featuring interesting theorems that usually aren't, but (in my opinion, certainly) can be introduced to gifted high school kids. These notes feature one such result - or at least, a simplified version, of a theorem of Andrew Suk which made a remarkable improvement on an old standing conjecture of Erdős and Szekeres. I was supposed to deliver 4 lectures to the participants of the IMOTC 2020 and I had intended to go over this result in those lectures, but the current COVID-19 situation has made physical lecturing impossible, hence these notes.

These notes are self-contained, and do not assume much from the students, other than a brief acquaintance with calculus (which they all have in plenty!). I have resorted to a conversational mode of transcribing so as to suggest a lecturing style. I have also cut the material into chapters with each chapter roughly containing what i would have done in the corresponding lecture.

Mathematics as is practiced by its practitioners (the mathematicians) is both similar to one's thinking process in competition exams, such as the IMO, and also opposite to it in the way one works with what one can, as an artist would. These lectures are MY attempt to deconstruct Suk's proof; if *you* were to discover this proof, could you have done it? Here is one way you could have - thats the perspective of these lectures (and notes).

I hope the kids at the IMOTC find the notes interesting, and get a sense of what a fantastic result this is.

Niranjan Balachandran,
IIT Bombay.

For starters: Stirling's formula

This lecture will appear as a bit of a digression, not completely related to the main objective of these lectures; I will get to them later. But as the title of this chapter suggests, this is a good round of mathematical houvre d'oeuvres. The only thing I will assume here is a very brief acquaintance with basic calculus.

The main purpose of this lecture is to prove

Theorem 1. ***(Stirling's formula)****: There is an absolute constant $C > 0$ such that*

$$\lim_{n\to\infty} \frac{n!e^n}{n^{n+\frac{1}{2}}} = C.$$

To put it somwewhat differently, a good 'approximation' to $n!$ would be $C\left(\frac{n}{e}\right)^n\sqrt{n}$. Actually, one can compute this C to be $\sqrt{2\pi}$. But that needs more work, and would not be necessary for our needs later (so this is not entirely a digression - it is more like preparatory work), so we will prove it in the form stated above.

I am assuming here that you are all familiar with the notion of limits. I will only assume one result regarding the existence of a limit, viz., the following

Theorem 2. *If $\{a_n\}_{n\geq 1}$ is a monotone increasing sequence, i.e., $a_n \leq a_{n=1}$ for all n, and if the sequence is bounded above, i.e., $a_n \leq M$ for all n, then $\lim_{n\to\infty} a_n$ exists.*

An analogous statement with monotone decreasing sequences bounded below follows by considering the negation of such a sequence.

So, how does one go about proving Stirling's formula?

Anyone moderately familiar with how one might approach this, is naturally drawn to considering $\log n!$ so that $\log n! = \displaystyle\sum_{k=1}^{n} \log k$.

One of the first important ideas we derive from calculus is that (Riemann) integrals are finite sums in the limit, so finite sums can be approximated by integrals. Now it is easy to see that since

$$\int_1^n \log x \; dx \leq \sum_{k=1}^{n} \log k \leq \int_0^{n+1} \log x \; dx \tag{1}$$

and since $\int_1^n \log x \; dx = n\log n - n$ (this is a simple calculus exercise - integration by parts!), a good starting point for an approximation to $n!$ is

$$n! \approx \left(\frac{n}{e}\right)^n.$$

Important remark: Stirling's approximation is a multiplicative approximation, not an additive one.

From (1) it is easy to see that $\lim_{n\to\infty} \frac{\log n!}{\log(n/e)^n} = 1$ but how about

$$\lim_{n\to\infty} \frac{n!}{(n/e)^n}?$$

Let's go back to (1). Calculating the integrals gives

$$\begin{aligned} n(\log n - 1) \quad &\leq \quad \log n! \quad \leq \quad (n+1)(\log(n+1) - 1) && (2)\\ &= \quad n(\log n - 1) + \log n + (n+1)\log\left(1 + \frac{1}{n}\right) && (3) \end{aligned}$$

and this suggests that we look at $x_n := \frac{n!}{(n/e)^n}$.

Let $b_n := \frac{x_{n-1}}{x_n}$ with $x_0 = 1$. Write

$$\begin{aligned} x_n^{-1} &= b_1 \cdots b_n \\ &= x_1^{-1} \prod_{j=2}^{n} \frac{1}{e}\left(1 + \frac{1}{j-1}\right)^{j-1} \\ &= \frac{1}{e^2} \prod_{j=1}^{n-1}\left(1 + \frac{1}{j}\right)^{j} \end{aligned}$$

Since

$$\lim_{n\to\infty}\left(1 + \frac{1}{n}\right)^n = e$$

we write

$$x_{n+1}^{-1} = \frac{1}{e}\prod_{j=1}^{n}(1+\varepsilon_j)$$

where

$$\varepsilon_j := \frac{1}{e}\left(1 + \frac{1}{j}\right)^j - 1.$$

This changes our focus to a question of the following general type: When does $\prod_{k=1}^{n}(1+\varepsilon_k)$ have a finite limit as $n \to \infty$? Here, in this more general question (with an abuse of notation where I have used the same expression ε_k) we shall think of ε_j as small positive quantities. Basic calculus again tells us that this limit exists if and only if the sum $\sum_{k=1}^{n}\log(1+\varepsilon_j)$ has a finite limit as $n \to \infty$.

Okay. So how do we estimate/approx the last sum? Again, calculus comes to help us. The buzzword now is 'Taylor series expansion'. If you don't know this well, don't worry too much; you have all seen some 'series expansions' of functions - $\sin x, \cos x, e^x$ etc. - those are infinite series expansions, but the general Taylor series is a finite summation, with an 'error' term. A general caveat here: It is very important in the branch of mathematics called Analysis to examine the rules for manipulating infinite sums/products, but the general principle of thinking is based on the same intuitive principles you use while working with finite sums. If this is a bit fuzzy, then it is meant to be, since I am not going to get into the details in Analysis - you may learn that later sometime. But for now, what we do with these is of greater importance.

I shall also need another very useful notation - called the asymptotic Big-Oh notation. Given functions f, g we say that $f(x) = O(g(x))$ for 'small' x (which means in a small vicinity of 0) if $|f(x)| \leq C|g(x)|$ for some fixed constant $C > 0$ and all $x \in (-a, a)$ for some small fixed $a > 0$. We could also have this only for $x \in [0, a)$ but the overall utility remains teh same. One of the main advantages of using this notation is that it allows us to think of inequalities as if they were equalities. As you will see, there are many utilities of this notation but for now, this is one for starters.

Right! So we all know

$$\log(1+x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \cdots$$

Indeed, start with

$$\frac{1}{1+x} = 1 - x + x^2 - \cdots$$

and integrate term-by-term. Again, there are many things here that need justification. For instance, the geometric series expansion only works when $|x| < 1$ and integrating an infinite series term-by-term needs a lot of justification, along with knowing when one can do so! But, as I said before, let us gloss over that for now.

A Taylor series expansion for $\log(1+x)$ is

$$\log(1+x) = x - \frac{x^2}{2} + R_1(x) \tag{4}$$

where the error function $R(x)$ satisfies $|R_1(x)| = O(x^3)$ whenever $|x| < 1/2$, say. We sometimes write this more tersely as

$$\log(1+x) = x - \frac{x^2}{2} + O(x^3).$$

Similarly, we have

$$e^x = 1 + x + O(x^2). \tag{5}$$

You may treat the above statements facts, if you are troubled by my lack of explanation.

Now, let us use this to get back to the previous expression. Recall $\varepsilon_j = \frac{1}{e}\left(1 + \frac{1}{j}\right)^j - 1$, so by (4)

$$\begin{aligned}
\log\left(1 + \frac{1}{j}\right)^j &= j \log\left(1 + \frac{1}{j}\right) \\
&= j\left(\frac{1}{j} - \frac{1}{2j^2} + O\left(\frac{1}{j^3}\right)\right) \\
&= 1 - \frac{1}{2j} + O\left(\frac{1}{j^2}\right).
\end{aligned}$$

Using (4) and (5) we have

$$\left(1+\frac{1}{j}\right)^{j}=\exp\left(1-\frac{1}{2j}+O\left(\frac{1}{j^2}\right)\right)=e\left(1-\frac{1}{2j}+O\left(\frac{1}{j^2}\right)\right)$$

so that

$$\varepsilon_j:=\frac{1}{e}\left(1+\frac{1}{j}\right)^{j}-1=-\frac{1}{2j}+O(\frac{1}{j^2}).$$

Unfortunately, this is of little help to our earlier observation on when the product $\prod_{k=1}^{n}(1+\varepsilon_k)$ has a limit, since $\sum_{j\geq 1}\frac{1}{2j}$ has an infinite sum (Why?! This is a good exercise to think about in case you haven't seen this before). So, all our work now hasn't cut anything of value...

Or maybe it has. Recall that what we were considering was to try to show if x_n^{-1} has a limit and our calculations have only shown that $x_n\to 0$ as $n\to\infty$ (Why? Follow the trail of arguments here). But that is not really Stirling's formula; in fact, all this shows is, $n!$ is much larger than $(\frac{n}{e})^n$ for n large. How much larger is it?

To see if we can make a better estimate, let $y_n=\frac{n!}{e^{-n}n^{n+s}}$ for some positive $s>$. This is basically an attempt to see if we can improve on our first approximation by a polynomial bit. At the moment, we haven't committed to the precise value of s, but as we will see, our analysis will tell us, that if we are to make some headway in this analysis, then the correct value of s will pop out as a byproduct.

Let us again proceed as before. Let

$$c_j=\frac{y_{j-1}}{y_j}=e\left(\frac{j}{j-1}\right)^{j+s-1}=e^2b_j\left(\frac{j}{j-1}\right)^{s}$$

with $y_0=1$. Let us go the whole nine yards again:

$$y_n^{-1}=c_1\cdots c_n=\frac{1}{y_1}\prod_{j=2}^{n-1}b_j\left(\frac{j+1}{j}\right)^{s}=\frac{1}{y_1}\prod_{j=2}^{n-1}b_j\left(1+\frac{1}{j}\right)^{s}$$

Since

$$\left(1+\frac{1}{j}\right)^{s}=1+s\frac{1}{j}+O(\frac{1}{j^2})$$

from the Newton binomial formula if you wish. This, again is its corresponding Taylor series. Plugging this back, we have

$$\begin{aligned} b_j\left(1+\frac{1}{j}\right)^{s} &= \left(1-\frac{1}{2j}+O(\frac{1}{j^2})\right)\left(1+\frac{s}{j}+O(\frac{1}{j^2})\right) \\ &= \left(1+(s-\frac{1}{2})\frac{1}{j}+O(\frac{1}{j^2})\right) \end{aligned}$$

so, setting $s=\frac{1}{2}$ gives

$$\lim y_n^{-1}=\lim_{n\to\infty}\prod_{j=1}^{n}c_j=e^{-1}\lim_{n\to\infty}\prod_{j=1}^{n}\left(1+O(\frac{1}{j^2})\right)$$

and now, this last term has limit as $n \to \infty$ by the exact same reason as elucidated earlier. This establishes the theorem. It would be a good exercise to fill in the details that I have not mentioned (and by that I mean, going over the calculations, not the proofs of the slightly harder statements of analysis).

As we mentioned earlier this does not establish the value of the limit, which is also part of the statement of Stirling's theorem. But that will not be necessary for us.

LECTURE 2: The Erdős-Szekeres appetizer and the Pór-Valtr salad

Let us get into 'where it all started' - what Erdős called the 'Happy Ending Theorem'. Esther Klein had made the following two observations:

1. For any set of 5 points in **general position** (no three points are collinear) in the plane, some four of these points are the vertices of a convex quadrilateral. (Almost trivial)

2. For any set of 9 points in general position in the plane, some five of these points are the vertices of a convex pentagon. (needs some work)

Both these are also best possible, in the sense, that there are configurations with fewer points than the ones of the respective statements where the conclusion does not hold.

The Erdős-Szekeres theorem is the following:

Theorem 3. *For any $n \geq 4$, there exists $N(n)$ such that the following holds. For any set of $N(n)$ points in the plane in general position, one can locate n among these points that form the vertices of a convex n-gon. Moreover,*

$$2^{n-2} < N(n) \leq \binom{2n-4}{n-2}.$$

By Stirling's formula, the upper bound above is of the order $\frac{4^n}{\sqrt{n}}$, so the upper bound is almost as large as the square of the lower bound.

There are many proofs of this theorem, but we shall see the one that is based on a latter paper of the same authors. To get there, we make the following definition (as they do).

Definition 4. *Suppose $P_1, \dots, P_k$ are in general position in the plane. By $x_i := x(P_i)$ (resp. $y_i := y(P_i)$) we mean the x-coordinate (resp. y-coordinate) of P_i. Without loss of generality, suppose $x_1 \leq \cdots \leq x_k$. We say that the ordered set $(P_1, \dots, P_k)$ is a k-cup if for all $2 \leq i \leq k-1$,*

$$\frac{y_i - y_{i-1}}{x_i - x_{i-1}} \leq \frac{y_{i+1} - y_i}{x_{i+1} - x_i}.$$

We say that the points form a k-cap, if for all $2 \leq i \leq k-1$

$$\frac{y_i - y_{i-1}}{x_i - x_{i-1}} \geq \frac{y_{i+1} - y_i}{x_{i+1} - x_i}.$$

For simplicity, we shall refer to the ratio $\frac{y_i - y_{i-1}}{x_i - x_{i-1}}$ by $s(P_{i-1}, P_i)$.

The motivation for these definitions is quite self-explanatory. An n-cup or an n cap form the vertices of a convex n-gon, and it is natural to think of coordinatising the points. The theorem of Erdős and Szekeres that establishes the upper bound in the aforementioned theorem is

Theorem 5. *Suppose $k, \ell \geq 2$. Let $f(k, \ell)$ denote the least integer such that any set of $f(k, \ell$ points in the plane either contains a k-cup or an ℓ-cap. Then*

$$f(k, \ell) = \binom{k+\ell-4}{k-2} + 1.$$

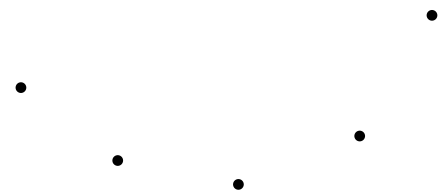

5-cup

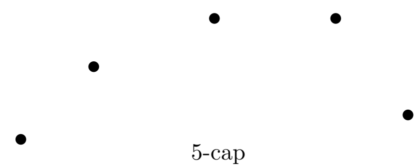

5-cap

Figure 1: Examples of a 5-cup and a 5-cap.

The theorem consists of two parts: the upper bound and the lower bound (to establish equality).

The proof of the upper bound is now a classic, so I will do it rather quickly. We'll denote the binomial coefficient by $\phi(k, \ell)$ for simplicity. We shall induct on $k + \ell$. If k or ℓ equals 2 then the statement is trivial, so suppose $k, \ell \geq 3$. Suppose we have a set $\mathcal{P}$ of $\phi(k, \ell) + 1$ points in general position with no k-cups or ℓ-caps. Since $\phi(k, \ell) > \phi(k-1, \ell), \phi(k, \ell-1)$, there are $k-1$-cups, as well as $\ell - 1$ caps, by induction. Let A denote the set of all the last points from all possible $(k-1)$-cups in $\mathcal{P}$. Then $\mathcal{P} \setminus A$ has no $k-1$-cups or ℓ caps, so again, by induction, $|\mathcal{P} \setminus A| \leq \phi(k-1, \ell)$ which means that $|A| \geq \phi(k, \ell) = 1 - \phi(k-1, \ell) = \phi(k, \ell - 1) + 1$. So it follows that there is an $\ell - 1$ cap $\mathcal{C} = (P_1, \ldots, P_{\ell-1})$ in this order, contained in A. But since each point of A is the end point of a $(k-1)$-cup, there is a $(k-1)$-cup $\mathcal{C}' = (Q_1, \ldots, Q_{k-1} = P_1)$ in this order.

Now we are through; if $s(Q_{k-2}, Q_{k-1}) \leq s(P_1, P_2)$ then $(Q_1, \ldots, Q_{k-1}, P_2)$ is a k-cup, otherwise $(Q_{k-2}, P_1, \ldots, P_{\ell-1})$ is an ℓ-cap.

The lower bound requires the construction of a set of $\phi(k, \ell)$ points with neither a k-cup, nor an ℓ-cap, and this is again constructed inductively. The crux of the proof lies in the observation that since $\phi(k, \ell) = \phi(k-1, \ell) + \phi(k, \ell - 1)$, one can imagine sets $\mathcal{P}_1, \mathcal{P}_2$ of sizes $\phi(k-1, \ell), \phi(k, \ell - 1)$ respectively without either $k-1$-cups or ℓ-caps (resp. k-caps or $\ell-1$-cups). Without loss of generality, assume that the points of $\mathcal{P}_1$ and $\mathcal{P}_2$ are $\{(i, y_i^{(1)})\}$ and $\{(i, y_i^{(2)})\}$ respectively. Then for a large enough m and $\varepsilon > 0$ small enough, the set$\mathcal{P}$ consisting of the points $(i, \varepsilon y_i^{(1)})$ and $(m + i, m + \varepsilon y_i^{(2)})$ will do the job. I leave it as a simple exercise for you guys to see why it is the case.

Clearly, this theorem establishes the upper bound in theorem **??**. What about the lower bound? This again follows the idea of the inductive construction of the lower bound in theorem **??**, but it is a little more involved, so I will not get into that for now.

This leads us to the statement of the Erdős-Szekeres conjecture:

Conjecture 6. *$N(n) = 2^{n-2} + 1$ where $N(n)$ is the function described in the statement of theorem* **??**.

In other words, they believed that the lower bound that they obtained was just short of the correct value for $N(n)$.

This conjecture was made in 1935, and it still remains unproven. But what was more drastic is that the upper bound remained unshaken until in 2017 when Andrew Suk proved a remarkable result which is the main focus of the next course of lectures. I will not show you a proof of his exact result since it is a little more laborious, but I will show you a version that will still make you appreciate what a huge improvement it is, on the known results till that point. Indeed, suppose $\varepsilon > 0$ be small. Then Suk's theorem shows

Theorem 7. *For n sufficiently large,*

$$N(n) \leq 2^{(1+\varepsilon)n}.$$

The exact statement has a slightly more messy expression in the exponent rather than $(1+\varepsilon)n$ which makes it somewhat stronger than the version stated here, but the proof of the exact result only needs sharper estimates and not newer ideas than the ones we shall see. Also note, this does not get us actually closer to the original conjecture of Erdős and Szekeres, but this is already a huge leap in that direction.

As we seek to improve upon the upper bound, we reframe the Erdős-Szekeres theorem in the following manner. Suppose $n \geq 4$ is a fixed integer. Given a configuration $\mathcal{P}$ of N points in the plane in general position, determine $\mathfrak{c}_n(\mathcal{P})$, the number of convex n-gons that occur with vertices amongst the points of $\mathcal{P}$. The Erdős-Szekeres theorem states that if $N \geq \binom{2n-4}{n-2}$, then $\mathfrak{c}_n(\mathcal{P}) \geq 1$. Can we get a better lower bound in general?

As I remarked earlier, the upper bound has remained unshakeable for a long time, so one of the approaches of mathematicians is to look at what *they can do instead with techniques they know*[1]. Let us now try and address the aforementioned question, and towards that end, we'll need a couple of definitions.

A triple of points (P_1, P_2, P_3) is said to have positive orientation if the clockwise ordering of the vertices of the triangle formed by these points is (P_1, P_2, P_3), otherwise we say that the orientation of the triple is negative.

Definition 8. *The tuples $\mathcal{P}_1 = (P_1, \ldots, P_t)$ and $\mathcal{P}_2 = (P'_1, \ldots, P'_t)$ of points in general position are said to be of the* **same order type** *if for any $1 \leq i < j < k \leq t$, the triples of points (P_i, P_j, P_k) and (P'_i, P'_j, P'_k) have the same orientation.*

The objective of this definition is to get to the following more general notion.

[1]This is best captured by the words of John Pierce, the father of modern satellite communication technology, 'We do what we *can* do, not what we think we *should* or what we *want* to do.' (emphases, mine).

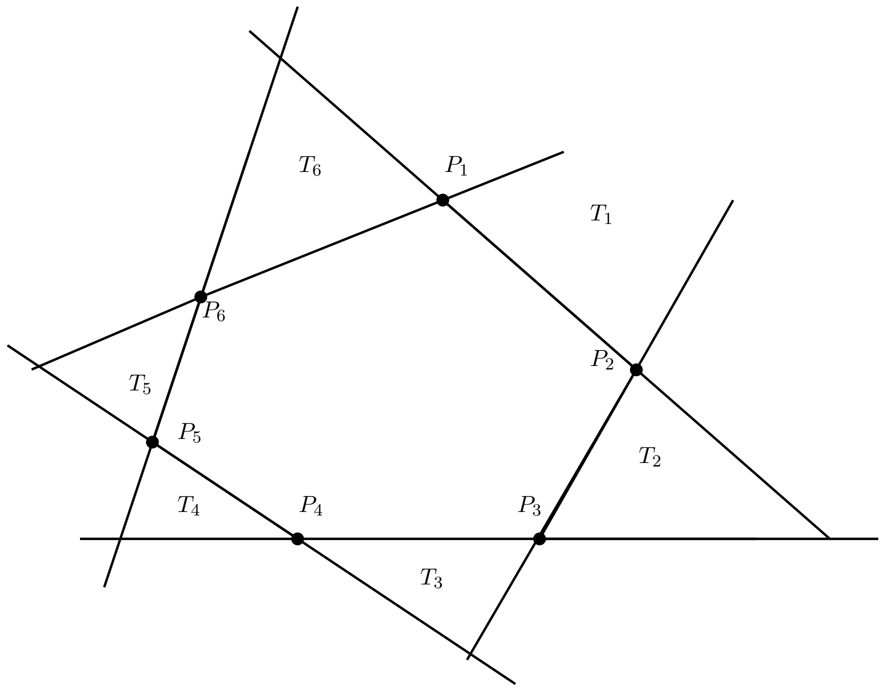

Figure 2: The regions T_i, for $i = 1, \ldots, 6$. It is important to remark that the regions T_i are not necessarily triangles.

Definition 9. *Suppose $\mathcal{P}$ is a finite set of points in general position in the plane. A partition $\mathcal{P} = X_1 \cup \cdots \cup X_n$ is called a n-clustering if*

- $|X_1| = \cdots = |X_n|$.
- *All n-tuples $(x_1, \ldots, x_n)$ with $x_i \in X_i$ have the same order type.*

If all the n-tuples form convex n-gons, then we call this a convex n-clustering.

In other words, if we have a convex n-clustering inside of a point set $\mathcal{P}$, then we have at least $\prod_{i=1}^{k} |X_i|$ different convex n-gons in $\mathcal{P}$.

So how do we manage to locate a convex n-clustering in $\mathcal{P}$? Towards this end, Pór & Valtr (2002) start with a simple but very useful observation. Suppose k is even, and let X be a set of k points in convex position. Let Y be a subset of X of size $k/2$ consisting of every alternate point of X, so that there are two choices for the set Y. Let us give this kind of set a name; following Pór and Valtr, we shall say that Y *supports* X in the aforementioned scenario. The following is the key observation.

Observation 0.1. *Suppose $Y = \{P_1, \ldots, P_k\}$ are points (forming a convex k-gon) in clockwise order. Denote by T_i, the region outside the convex hull of Y bound by the lines determined by $P_{i-1}P_i, P_iP_{i+1}$, and $P_{i+1}P_{i+2}$ (see figure 2 for an illustration with $k = 6$). If Y supports a $2k$-gon whose remaining points (apart from the points of Y) are $P'_1, \ldots, P'_k$, then each of the regions T_i contains exactly one of the P'_i.*

Suppose $\mathcal{P}$ is a set of N points in general position in the plane. The upper bound of the Erdős-Szekeres theorem says that among any 4^k points of $\mathcal{P}$ there is at least one convex k-gon. This gives us a basic idea; each set of 4^k points counts at least one towards $\mathfrak{c}_n(\mathcal{P})$. This is a basic double counting method, but it is much better perceived when put in a probabilistic language.

Suppose we pick a random (uniformly random) set S of 4^k points from the points of $\mathcal{P}$, and let $c(S)$ denote the number of convex k-gons in S. One of the most useful things to do with any random variable (integral valued, for instance) is to compute its expected value. Erdős-Szekeres tells us that $C := c(S) \geq 1$ for all S, so in particular $\mathbb{E}(C) \geq 1$. Let $\mathcal{C}$ denote the set of convex k-gons in $\mathcal{P}$. Then

$$1 \leq \mathbb{E}(C) = \sum_{X \in \mathcal{C}} \mathbb{P}(X \subset S) = |\mathcal{C}| \frac{\binom{N-k}{4^k-k}}{\binom{N}{4^k}} \quad \Rightarrow \quad |\mathcal{C}| \geq \frac{\binom{N}{4^k}}{\binom{N-k}{4^k-k}} \geq \frac{\binom{N}{4^k}}{\binom{4^k}{k}}.$$

There are $\binom{N}{k/2}$ different subsets of $\mathcal{P}$ of size $k/2$, and at least $\frac{\binom{N}{4^k}}{\binom{4^k}{k}}$ k-subsets of $\mathcal{P}$ in convex position, so by a simple double count argument, there is some $(k/2)$-subset Y such that Y supports at least

$$\begin{aligned} \frac{\binom{N}{4^k}}{\binom{4^k}{k}\binom{N}{k/2}} &> \frac{(N-k)^{k/2} \cdot (k/2)!}{\binom{4^k}{k} \cdot k!} \\ &> \frac{(N-k)^{k/2}}{4^{k^2}} \end{aligned}$$

convex k-subsets of $\mathcal{P}$.

To keep things from getting uglier notationally, let us write $k = 2\ell$. Let $Y = (P_1, \dots, P_\ell)$ in clockwise order. By observation **??**, if the regions T_i determined by the points P_i contains t_i points of $\mathcal{P}$, then each 2ℓ-gon supported by Y must have its remaining points coming one each from each of the T_i, so the number of such 2ℓ-gons supported by Y is at most $\prod_{i=1}^{\ell} t_i$.

Hence,

$$\prod_{i=1}^{\ell} t_i > \frac{(N-2\ell)^\ell}{4^{4\ell^2}}.$$

Now, to get hold of a convex clustering, if $t_i \geq t_0$ then we have a lower bound on the number of convex ℓ-gons. Now, it is not possible to prove that all the t_i are all large. But, here is another computational idea: If t_0 denotes the $(\ell/2)^{th}$ largest among the t_i then $\prod_{i=1}^{\ell} t_i \leq t_0^{\ell/2} N^{\ell/2}$. What this suggests is: Set $n = \ell/2$, and this analysis gets us a bound on the number of convex n-gons.

Piecing all these together, (writing $n = 2\ell$), we have

$$\begin{aligned} t_0 &> \left(\frac{(N-4n)^{2n}}{2^{32n^2} \cdot N^n}\right)^{\frac{1}{n}} \\ &> \frac{N-8k}{2^{32n}} \\ &> \frac{N}{2^{32n}} - 1. \end{aligned}$$

So what have we done so far? If we collect those indices i such that $t_i \geq t_0$, and pick subsets $S_i \subset T_i \cap \mathcal{P}$ for those i, then these sets, along with the corresponding P_i (from the set Y) together describe a convex n-clustering with each of the clusters containing $\lceil \frac{N}{2^{32n}} \rceil$ points of $\mathcal{P}$. This is the Pór-Valtr theorem:

Theorem 10. *(Pór-Valtr, 2002) 'Positive Fraction Erdős-Szekeres theorem') For an $n \geq 3$ if $\mathcal{P}$ is a set of N points in general position in the plane, then it contains a convex n-clustering of size at least $\frac{k \cdot N}{2^{32n}}$.*

LECTURE 3: Main course - Preparations towards Suk's proof

Before we get to the main course, let me make one small comment regarding the result from the previous lecture. The definition of cups and caps is actually a 'local' definition, i.e., it is determined by two consecutive pairs of points (consecutive is determined by the ordering of the points according to their x-coordinate). This suggests a more combinatorial rephrasing of the aforementioned theorem.

Imagine coloring triangles formed by these points by the following rule. If P_1, P_2, P_3 are three points with $x_1 < x_2 < x_3$ ($x_1 = x(P_1)$, etc.) then let us color the triangle $P_1P_2P_3$ red if P_1, P_2, P_3 form a 3-cup and blue otherwise. What characterizes the cups-and-caps theorem is the following *transitivity property*: If $P_1P_2P_3, P_2P_3P_4$ are both red (resp. blue) then $P_1P_2P_4$ and $P_2P_3P_4$ are also red (resp. blue). The proof of the cups-and-caps theorem can then be recast into the following purely combinatorial theorem (which absorbs the geometric nature of the setting). Here's a bit of terminology first; a coloring of triples of a set $[n]$ is said to be it transitive if whenever xyz, yzw have the same color then xyw, yzw also have the same color as xyz. By a *monochromatic k-clique* we mean a set $x_1, \ldots, x_k$ such that all the triples $x_ix_jx_k$ have the same color.

Theorem 11. *Suppose $k, \ell \geq 3$ are integers. Denote by $N = g(k, \ell)$ the smallest integer such that for any transitive 2-coloring of the triples of $[N]$, Red and Blue, say, there is either a red k-clique or a blue ℓ-clique. Then*

$$g(k, \ell) = \binom{k + \ell - 4}{k - 2} + 1.$$

I leave it to you guys to make this translation of language. The same proof works - just verify it yourselves.

Before we launch into our unpacking of Suk's proof, I shall set some terminology, which should not be difficult to formalize for the non-litigiously minded. I shall henceforth refer to points lying above and below a line. With that notion, let us make a more geometric definition of cups and caps. A k-cap is a set of k points $P_1, \ldots, P_k$ such that for each P_i there is some line ℓ_i such that $P_i \in \ell_i$ and all the other P_j lie on or below (resp. above) ℓ_i. This description is more a local one in the sense that to check if a set if a cap, one needs to associate a reference line to each of these points such that this property is satisfied.

Let us now attempt to improve upon the Erdős-Szekeres bound. Suppose for a suitable k (which we shall determine later), we use the Pór-Valtr theorem to get a convex k-clustering with each of the parts having several points. We will attempt to pick subsets $\mathcal{P}_1, \ldots, \mathcal{P}_r$ such that hopefully, $\cup_i \mathcal{P}_i$ will describe the vertices of a convex n-gon.

A k-gon admits a subset of at least $k/2$ which is either a cup or cap, so without loss of generality, we assume that (for the set $\mathcal{P}$ of size N) there is a convex k-clustering Y forming a k-cap, with each set of the cluster having size at least $\frac{N}{2^{64k}}$. If our purported $\mathcal{P}$ is a union of $\mathcal{P}_i$, then the most natural thing would be that to get each $\mathcal{P}_i$ as a cap, of a suitable size t, say. But the more

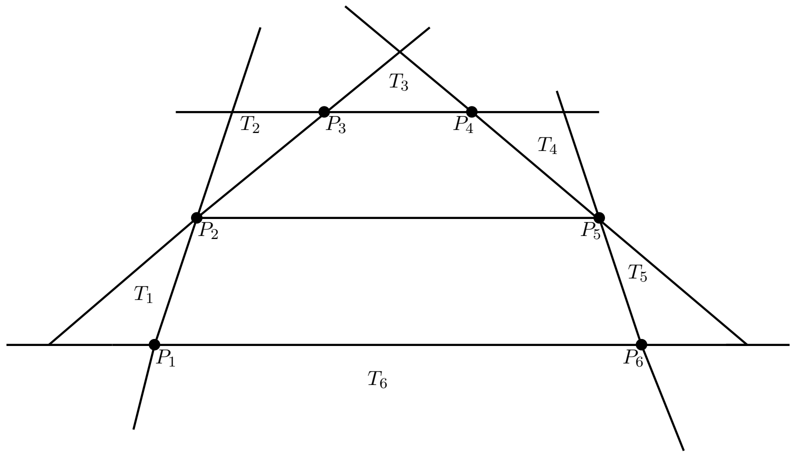

Figure 3: The regions T_i and the associated line segment $\overline{\ell_i}$. In this example, we illustrate it with $k = 6$ and $\overline{\ell_3} = \overline{P_2P_5}$.

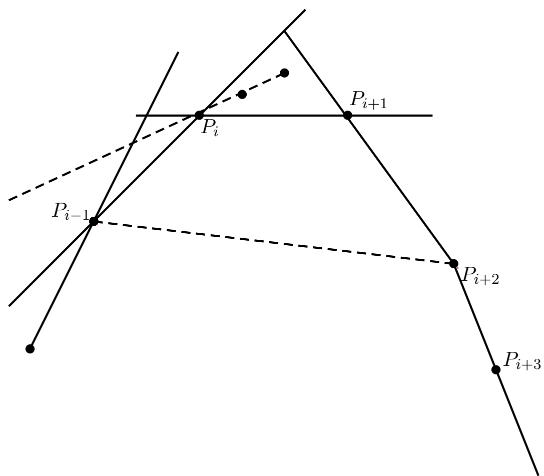

Figure 4: The line joining two of the points in $\mathcal{T}_i$ does not meet the line segment $\overline{\ell_i}$.

immediate question is: How does one guarantee that the union of these $\mathcal{P}_i$ is also a cap? At this juncture come Suk's (in my opinion) first real breakthrough observations:

- Restrict attention to the *line segments* $\overline{P_{i-1}P_{i+2}}$. If we can pick a subset $\mathcal{T}_i \subset T_i$ such that for any $Q, Q' \in \mathcal{T}_i$, $QQ'P_{i+2}P_{i-1}$ is a convex quadrilateral, then the line QQ' does not intersect the line segment $\overline{\ell_i} := \overline{P_{i-1}P_{i+2}}$.

- Suppose $\mathcal{P}_i \subset \mathcal{T}_i, \mathcal{P}_j \subset \mathcal{T}_j$ are caps in the regions T_i, T_j respectively. Let Q, Q' be distinct points in $\mathcal{P}_j$. Then if the line QQ' intersects the region T_i, we must have $|i - j| = 1$. This is easily seen once by observing that if say $i < j - 1$, then the region T_i lies below the line $P_{i-1}P_i$. A similar remark is applicable for $i > j + 1$.

How large a set $\mathcal{T}_i$ can we get? First, one can imagine doing this greedily, i.e., picking points sequentially so that the desired property is not destroyed by the addition of the next point. But a more theoretic way of looking at this would be: when do points Q, Q' along with P_{i-1}, P_{i+2} NOT form a convex quadrilateral? This is precisely when one of the points (say Q') lines *inside the triangle* $QP_{i-1}P_{i+2}$.

This suggests a relation among the points of T_i. For points $Q \neq Q'$ write $Q' \prec Q$ if Q' lies inside the triangle $QP_{i-1}P_{i+2}$. This relation makes T_i into a poset, and what we want for $\mathcal{T}_i$ is an *antichain* in this poset.

So, with these preparatory ideas, we make the first stab. We will call the regions T_i, T_{i+1} as successive sets. Suppose there are non-consecutive regions $T_{i_1}, \ldots, , T_{i_r}$ (for some r; we will tie them all together later) with each of these regions admitting antichains $\mathcal{T}_i$ of size L (again, TBDL - to be determined). Then within these antichains we can find caps $\mathcal{P}_i$ of size t each then our preceding observations tell us that the union $\cup_{i=1}^{r} \mathcal{P}_i$ is also a cap. Indeed, for each point $P \in \mathcal{P}_i$, since $\mathcal{P}_i$ is a cap, there is some line passing through P such that all the other points of $\mathcal{P}_i$ lie below it. We may choose this line to be the one passing through P and one of its adjacent points in the cap (within T_i). But since the line joining any two of the points here must lie *above* the regions T_j for all the other j, it follows that all the other points of $\cup \mathcal{P}_i$ also lie below this line. And that establishes that we have a cap on our hands now.

How about size? And also, how do we guarantee that each of these sets $\mathcal{P}_i$ is a cap? That is easily gotten again through the Erdős-Szekeres result. Indeed, if $|\mathcal{T}_i| > \phi(n, t)$ - where $\phi(n, t)$ is the binomial coefficient in the bound (see lecture 2) - then either we have an n-cup, which means we are done, or we have a t-cap $\mathcal{P}_i \subset \mathcal{T}_i$.

And how does one get this large a $\mathcal{T}_i$? We know that each T_i has size at least $\frac{N}{2^{64k}}$. But how do we guarantee a large antichain in T_i? In fact, you cannot guarantee it. What if all the points of T_i form a *chain* in T_i?

But at this point, we are reminded of Dilworth's theorem:

Theorem 12. *(Dilworth) If* $(\mathfrak{P}, \prec)$ *is a finite poset then either there is an antichain of size* t *or a chain of size* $\frac{|\mathfrak{P}|}{t}$[2].

This is a highly nontrivial theorem in extremal combinatorics, and while you must have seen the statement before, I am not sure if you have seen its proof. I'll come back to it later.

Let us tentatively set the size of $\mathcal{T}_i$ to be $|T_i|^\alpha$ for some $\alpha > 0$, and $k = n^\beta + 2$ with $0 < \beta < 1$. If

$$\left(\frac{N}{2^{64k}}\right)^\alpha > \binom{t+n-4}{n-2} \tag{6}$$

$$n^\beta t \geq n, \tag{7}$$

then we have the following statement: If there are non-consecutive regions T_i such that the corresponding posets admits 'large' antichains (here the notion of large is, an antichain of size $|T_i|^\alpha$)

[2] Technically, one should keep floor and ceilings everywhere for this to read exactly but I shall routinely ignore them for convenience.

then in that case, if the numerics above work out well, then we have a possible improvement of the Erdős-Szekeres bound!

That sounds like a huge wish list, so let us see if at least the numerics work out. The second inequality suggests that we take $t = n^{1-\beta}$. Now Stirling's formula gives us that the RHS of the first inequality is approximately

$$\begin{aligned}\binom{n+n^{1-\beta}-4}{n-2} &= \frac{(n+n^{1-\beta}-4)!}{(n-2)!(n^{1-\beta}-2)!} \\ &\approx \frac{(n+n^{1-\beta}-4)^{n+n^{1-\beta}-4}\cdot\sqrt{n+n^{1-\beta}-4}}{(n-2)^{n-2}\cdot\sqrt{n-2}\cdot(n^{1-\beta}-2)^{n^{1-\beta}-2}\cdot\sqrt{n^{1-\beta}-2}} \\ &\approx \frac{n^{n+n^{1-\beta}-4}\cdot\sqrt{n}}{n^{(n-2)+(1-\beta)(n^{1-\beta}-2)}\cdot n^{1-\beta/2}} \\ &= 2^K\end{aligned}$$

where

$$\begin{aligned}K &= \Big(n+n^{1-\beta}-4+1/2-(n-2)-(1-\beta)n^{1-\beta}+2(1-\beta)-1+\beta/2\Big)\log_2 n \\ &\leq \beta n^{1-\beta}\log_2 n.\end{aligned}$$

But on the LHS of the first inequality of our wish list, we have $\left(\frac{N}{2^{64k}}\right)^\alpha$, so if we set $N = 2^{(1+\varepsilon)n}$, this gives us

$$\left(\frac{N}{2^{64k}}\right)^\alpha \geq 2^{(1+\varepsilon)\alpha n-64\alpha n^\beta}$$

so that if

$$(1+\varepsilon)\alpha n - 64\alpha n^\beta > \beta n^{1-\beta}\log_2 n$$

we are through. But this holds for n sufficiently large, so indeed what we were trying will work!

Of course, before we all start patting our backs, let me remind you that all we have done is, under a certain situation, our scheme will get us through. But what do we do if the dream situation does not happen? Let us write that out more explicitly: Let us color the region T_i Red if the corresponding poset admits an antichain of size $|T_i|^\alpha$ and Blue if the poset admits a chain of order $|T_i|^{1-\alpha}$. Our dream scenario was, the presence of at least $n^{1-\beta}$ non adjacent Red T_i's.

At this juncture, it is necessary that t is smaller than k, and since our computation earlier did not need any specific value for α or β, this forces us to make one commitment; let us set $\beta = \frac{2}{3}$, so that $t = n^{1/3}$.

So, what if this dream scenario does not pan out? Then we can conclude that there are at least $n^{1/3}$ *consecutive* $T_j, T_{j+1}, \ldots$ such that in each of these sets, there is a chain Q_i of size at least $2^{(1+\varepsilon)n-65n^{2/3}}$.

This is a good place to stop.

LECTURE 4: Dessert - The cherry on the cake à la Suk

So, a quick recap on where we are perched currently. We are now in the situation where there are $n^{1/3}$ consecutive regions with each of these having chains of size at least $|T_i|^{1-\alpha} \geq \left(\frac{2^{(1+\varepsilon)n}}{2^{65n^{2/3}}}\right)^{1-\alpha}$. Let us call these chains $\mathcal{Q}_1, \mathcal{Q}_2, \ldots$ for convenience and let us label the points inside $\mathcal{Q}_i$ according to $\prec$ as $Q_1^{(i)}, Q_2^{(i)}, \ldots$.

Following Suk, let us call a subset $Y \subset \mathcal{Q}_i$ a *right cap* if $P_i \cup Y$ is in convex position, and $Y \subset \mathcal{Q}_{i-1}$ a *left cap* if $P_i \cup Y$ is in convex position. Figure 2 illustrates this for convenience. If the line $P_{i-1}P_i$ is a vertical line, then a left cap in $\mathcal{Q}_{i-1}$ is precisely a cap, and a right cap in $\mathcal{Q}_i$ is a cup in the usual sense. There are two important consequences to note following this definition:

- Since each $\mathcal{Q}_i$ is a chain, by the definition of the poset, every triple of points in $\mathcal{Q}_i$ is either a left-cap or a right-cap, but cannot be both.
- If $Q_y^{(i)} \prec Q_u^{(i)} \prec Q_v^{(i)} \prec Q_w^{(i)}$ are distinct points in $\mathcal{Q}_i$, and if $Q_y^{(i)}, Q_u^{(i)}, Q_v^{(i)}$ form a left (right) cap, and $Q_u^{(i)}, Q_v^{(i)}, Q_w^{(i)}$ also form a left (resp. right) cap, then $Q_y^{(i)}, Q_u^{(i)}, Q_w^{(i)}$ and $Q_y^{(i)}, Q_v^{(i)}, Q_w^{(i)}$ also are left (resp. right) caps. Consequently, by the combinatorial version of the Erdős-Szekeres bound, it follows that if $|\mathcal{Q}_i| > \phi(k, \ell)$ then $\mathcal{Q}_i$ has either a k-left cap or an ℓ-right cap.

Consider the chains $\mathcal{Q}_{i-1}$ and $\mathcal{Q}_i$. Here, is Suk's second critical observation:

Observation 0.2. *If $\mathcal{P} \subset \mathcal{Q}_{i-1}$ is a k-left cap, and $\mathcal{P}' \subset \mathcal{Q}_i$ is an ℓ-right cap, then $\mathcal{P} \cup \mathcal{P}'$ is a set of $k + \ell$ points in convex position.*

Figure 2 gives an illustration of this with a 3-left cap and a 3-right cap in consecutive regions.

To prove this observation, we shall show that every four points in this union forms a convex quadrilateral. It is an easy and well-known fact that if a set of points has the property that every four among them are the vertices of a convex quadrilateral, then the points are in convex position.

Consider a set Q_1, Q_2, Q_3, Q_4. If all of these are in $\mathcal{P}$ or all in $\mathcal{P}'$ then by assumption they are in convex position. So, suppose that, without loss of generality, that $Q_1, Q_2 \in \mathcal{P}$ and $Q_3, Q_4 \in \mathcal{P}'$.

Since the sets $\mathcal{Q}_i$ are chains with respect to $\prec$, a moment's reflection[3] will tell you that the line Q_1Q_2 does not meet T_i and similarly, the line Q_3Q_4 does not meet the region T_{i-1}. Note that while in general the regions T_i are not necessarily triangles, the ones in the 'middle' like the ones we are dealing with currently, are triangular. Consequently, these lines do the requisite job, i.e., for each of these points, these lines have the property that the other points are all on the same side of this line.

[3] Many authors from the sixties and seventies have used this phrase rather nonchalantly, but the ambiguousness that comes from this phrase has led to its disuse in recent times. This is sometimes referred to, maybe with a hint of derision, as the 'Reflection Principle'.

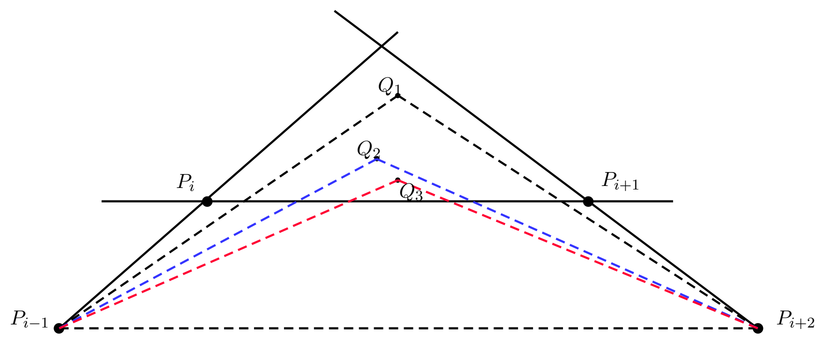

Figure 5: Points $Q_1 \prec Q_2 \prec Q_3$ in $\mathcal{Q}_i$.

Now, the other case. Suppose $Q_1, Q_2, Q_3 \in \mathcal{Q}_{i-1}$ and $Q_4 \in \mathcal{Q}_i$. Assume that $Q_1 \prec Q_2 \prec Q_3$. Note that the line parallel to $P_{i+1}P_{i+2}$ passing through Q_4 has all the other Q_i on one side of this line. Since Q_i for $1 \leq i \leq 3$ are in the chain $\mathcal{Q}_{i-1}$, the lines formed by a pair of these three all touch the line segment $\overline{\ell_{i-1}}$. In particular, both the points P_i and Q_4 lie on the same side of the region bounded by the lines formed by these three pairs. This establishes the observation.

Let us see where this takes us. Start with $\mathcal{Q}_1$. By assumption, $|\mathcal{Q}_1| \geq \left(\frac{N}{2^{64k}}\right)^{1-\alpha}$. If this is greater than $\phi(K, n)$ for some K, then without loss of generality, $\mathcal{Q}_1$ contains a K-left cap. Now, Suk's ingenious idea is this: Consider $\mathcal{Q}_2$; by the previous observation, if there is an $(n-K)$-right cap here, then the union of this along with the K-left cap in $\mathcal{Q}_1$ gives us n points in convex position. So, if we wish to build upon this, then we should force a larger left cap in $\mathcal{Q}_2$. So, the natural formulation is: If $|\mathcal{Q}_2| > \phi(2K, n-K)$ then either we are done or there is a $2K$-left cap in $\mathcal{Q}_2$.

Let us now continue this a little longer. For the s^{th} chain $\mathcal{Q}_s$, we wish to force $|\mathcal{Q}_s| > \phi(sK, n-(s-1)K)$ then either we are done, by having n points in convex position coming from the union of the $(s-1)K$-left cap in $\mathcal{Q}_s$ and the $n-(s-1)K$-right cap in $\mathcal{Q}_s$, or there is an sK-left cap in $\mathcal{Q}_s$. We can continue this till $n^{1/3}$ times, so if $K \gtrsim n^{2/3}$, and all the inequalities we have imposed hold, then we are through!

But again, let us get to work. We need

$$\left(\frac{2^{(1+\varepsilon)n}}{2^{65n^{2/3}}}\right)^{1-\alpha} > \binom{sK + (n-(s-1)K) - 4}{sK - 2} = \binom{n+K-4}{sK-2} \text{ for all } 1 \leq s \leq n^{1/3}. \tag{8}$$

Let us simplify things again for computations' sake. Set $K = n^\gamma$ for some $\frac{2}{3} \leq \gamma < 1$. Since

$$\binom{n+K-4}{sK-2} \quad < \quad 2^{n+K-4}$$

we may as well require that

$$\left(\frac{2^{(1+\varepsilon)n}}{2^{65n^{2/3}}}\right)^{1-\alpha} \geq 2^{n+n^{2/3}-4}.$$

We are just setting $K = n^{2/3}$. Note that so far, our choice for α has been uncommitted.

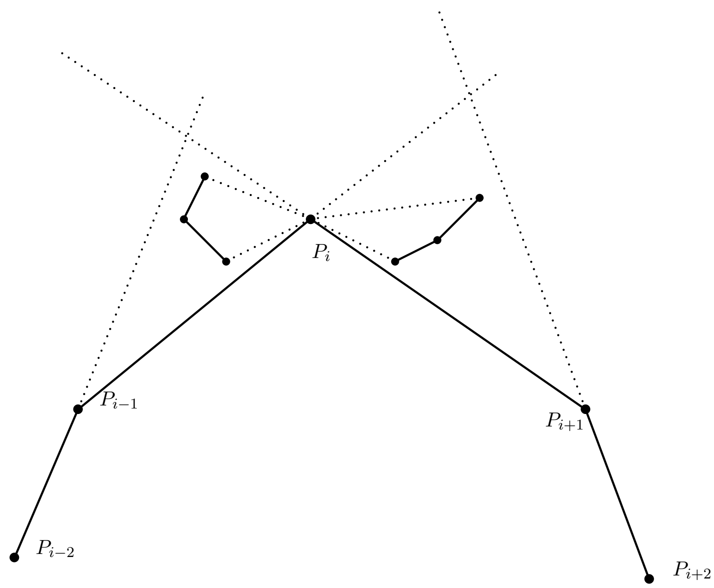

Figure 6: A 3-left cap and a 3-right cap in consecutive regions. The union of these two forms a convex 6-gon.

Now, the final assault. Let us take $\log_2$ on both sides to simplify our requirement. We need

$$(1-\alpha)\left((1+\varepsilon)n - 65n^{2/3}\right) \geq n + n^{2/3} - 4.$$

Here, we think of $\varepsilon > 0$ as fixed, but arbitrarily small (less than $1/2$, say). But now, we get to pick α! Choose $\alpha > 0$ small (for instance $\alpha = \varepsilon/10$, say); then it suffices to show

$$n\left(1 + \frac{9\varepsilon}{10} - \frac{\varepsilon^2}{10}\right) - 64n^{2/3} \geq n + n^{2/3}$$

or more simplistically, if we have

$$\frac{4\varepsilon n}{5} > 65n^{2/3}$$

then we are through. But this certainly holds for all large enough n, so we are through with room to spare!

Actually, our choices here were quite arbitrary. In a sense, asymptotics works best when you choose values that are easy to work with. But, even with that caveat, there was an abundance of room here, and indeed, Suk proves (with some more care in making the estimates) that $N(n) = 2^{n+6n^{2/3}\log n}$ works for the same proof. This is stronger than what we have in our calculations.

That finishes the proof of Suk's theorem.

How good a result is this? It is interesting to know that before this result, the best known bound was

$$N(n) \leq \frac{7}{16}\binom{2n-4}{n-2}$$

which barely makes a dent on the original bound. This one brings the exponent very close to n.

But then, how close is this to proving the Erdős-Szekeres conjecture? There, we are still 'infinitely' far off. But maybe some day...