

An Account of my Oral Qualifying Examination

by Jacob D. Baron, ABD

March 3, 2013

On Thursday, January 31, 2013, at 3:08 pm, I passed my oral qualifying examination to advance to candidacy for the PhD.¹ This note details the 88 minutes that led to that happy moment.

Four eminent mathematicians comprised my committee: **K**, **B**, **Z** and **T**, for our purposes. **B** and **T** arrived first, several minutes early. Both politely declined the two boxes of cookies I had brought, so I nibbled at them myself. **K** arrived right on time, at 1:40, and remarked slyly on the absence of **Z**, who places a high value on punctuality. Then, seeing the cookies, he exclaimed “a bribe!”—to which I replied that I am not above such things. When **Z** arrived, a few minutes late,² we began.

B asked the first line of questions. He started with what should have been easy: State Roth’s 1/4-theorem on the discrepancy of the arithmetic progressions in $[n]$, and compare to Szemerédi’s theorem. Embarrassingly, I messed up the statements of both theorems: In Roth’s I replaced $\Omega(n^{1/4})$ by $O(n)$ for some reason, and for Szemerédi I stated van der Waerden instead. An auspicious beginning, clearly. **K** corrected me on both counts, but I think he was convinced that I knew what I was talking about—it was just my nerves that didn’t. I mentioned the difference between Szemerédi and van der Waerden, that the first is a “density version” of the second. I think something like this is basically what **B** was looking for in the comparison of Roth vs. Szemerédi, since we didn’t return to that as far as I remember.

B’s next question was much harder: Explain intuitively why the exponent in Roth’s theorem should be 1/4. I had no clue, so I filibustered by changing the subject to counting the arithmetic progressions in $[n]$. My immediate bound was $O(n^3)$, which is correct but far from the truth. **B** and **K** asked for a tighter bound, so after maybe five minutes of stumbling with summations I came to $O(n^2 \log n)$.³

Returning to the question of the reason for 1/4, I blew aimless hot air for a minute before **B** realized I didn’t know. He told me what he was looking for, which I no longer remember precisely (it was something like, “most APs have length about \sqrt{n} , and any coloring will have square root error”), and we moved on.⁴

¹ Actually, candidacy came several days later, after a minor tussle with the university bureaucracy.

² In fact, **Z** had told me beforehand that he would be a few minutes late as he had to get lunch, but this did not satisfy **K**. Perhaps **Z** gets a B+?

³ **K** remarked to me after the exam that he thought I was slow here, but I thought I was actually all right. The question was on the easier side, though.

⁴ Sometime during **B**’s questioning, I don’t remember exactly when, **B** and **Z** got into a lively debate between themselves about a conjectured improvement to $O(\sqrt{n})$ in **B**’s famous theorem on discrepancy. I don’t remember

Z took the next turn, which he opened by asking for the statement of Ramsey's theorem. I gave it in greatest generality, ("the finite version, for you," I said), with any number of colors, any size of the sets being colored, and cliques of any size. Naturally this is not at all what he wanted, which he demonstrated with his next question: What is $R(2, 2)$? (This question, I think and hope, was mostly a joke.) I answered, 2, and then added, incorrectly and totally unnecessarily, that $R(2, n) = 2$ for all n . **K** raised an eyebrow, and I quickly corrected myself, embarrassed again. **Z** next asked for $R(3, 3)$, which I gave, $R(4, 4)$, for which I was off by 1, and and $R(5, 5)$, for I gave the right ballpark ("40's or 50's"). **K** had been growing more and more impatient through this whole process, and finally blurted out, "But let me ask you: Do you care?" Of course I don't, and I said so. "Then you fail automatically," exclaimed **Z**.

My next exchange with **Z** was perhaps the highlight of the entire exam. "How would you go about proving," he asked, "that $R(3, 3) \leq 6$?" Immediately I began to outline the standard proof—pick a vertex... . Once he saw that I knew it, he pounced: "NO! That's how **B** would prove that $R(3, 3) \leq 6$!" The right way to prove this, apparently, is by brute force, examining every graph on six vertices and checking that it has either a triangle or an antitriangle. Naturally this process should be computer-assisted. So **Z** had me spend the next 20 minutes writing Maple code to generate the graphs and do this check. Mercifully he did not insist on perfect Maplese. Towards the end of this ordeal, **K**, who through the whole thing had been taking less and less care to hide his intense lack of interest, blurted out that I clearly knew what I was doing so we should get a move on, and that this wasn't on my syllabus anyway. But **Z** insisted, "I have my time, I have my time," to a resigned eyeroll from **K**.

When I finally finished the Maple code, **Z** asked one last question: If I were to try to prove that $R(4, 4) \leq 18$ via the code I had just written, why would I fail? The answer of course is that there are $2^{\binom{18}{2}}$ graphs to check—way too many. This segued nicely into **K**'s first question, which I believe he came up with on the spot. We had remarked during the code writing that I was generating labeled graphs on $[n]$, not unlabeled ones—the former being easy to count, the latter utterly intractable. But if we abandon hope for an exact count of the unlabeled graphs, the asymptotics are easy—and that was his question. "You can actually guess the asymptotics," he said. The key insight, which **B** prodded me to after a few minutes, is that a.a.s. $\mathbb{G}_{n, \frac{1}{2}}$ has trivial automorphism group.⁵ Thus, for basically every labeled graph on $[n]$, each of the $n!$ relabelings produces a different (though isomorphic) graph. In other words, for basically every unlabeled graph on n vertices, exactly $n!$ labeled graphs on $[n]$ give rise to it after dropping the labels. This suggests that there are asymptotically $\frac{2^{\binom{n}{2}}}{n!}$ unlabeled graphs on n . It's not a proof, of course (we would have to worry about the "almost"), but it's what **K** was looking for.

K's next question, which took the bulk of his time and which I would consider the climax of the exam, was the following. Given any connected graph G with m edges, prove that the expected number of edges in a uniformly chosen connected spanning subgraph of G is at least $m/2$. Within

how exactly this came up, but it led after a moment to **K** announcing, loudly, "Gentlemen! You both pass! But we're testing Jake..."

⁵ They had me wondering about the size of the automorphism group of a generic graph on $[n]$, and I was floundering. At some point **B** stuck his index finger up in front of him, shaking it slightly. I asked if he was trying to communicate something; he said yes. "You're not being very successful," I said, to laughter from **K**. "One? One what?" I soon figured out that his one finger was for the one (trivial) automorphism of a generic graph.

seconds I gave the correct idea, which is that since connectivity is an increasing graph property, a random connected subgraph of G should be on average at least as large as a random arbitrary subgraph of G . From this intuition I knew I should use a correlation inequality, but it took me a few minutes and some stumbling to nail down a proof, with Harris's inequality.

K's next question was whether the inequality I had just proved is tight, i.e. whether there exists some $\epsilon > 0$ such that for any connected G , $\mathbb{E}|H| \geq m/2 + \epsilon$ for H uniformly chosen from the connected spanning subgraphs of G . “It must be,” I said—“look at the complete graph.” I didn’t exactly know how this would work in the moment, but **K** was satisfied (“good!”) and we moved on. After the exam I worked it out: Consider choosing from G a subgraph H uniformly, not necessarily connected. Then

$$\mathbb{E}|H| = \frac{m}{2} \geq \Pr(H \text{ connected}) \left[\frac{\mathbb{E}(|H| \mathbf{1}_{\{H \text{ connected}\}})}{\Pr(H \text{ connected})} \right].$$

The quantity in brackets is the quantity we’re worried about, the expected size of a uniformly chosen connected spanning subgraph of G . Letting $G = K_n$, so that $H = \mathbb{G}_{n,1/2}$, we have $\Pr(H \text{ connected}) \rightarrow 1$ as $n \rightarrow \infty$ (since $1/2 \gg \frac{\log n}{n}$, $\mathbb{G}_{n,1/2}$ is basically always connected), and so the quantity in brackets decreases to $m/2$. This shows the tightness of the inequality up to \geq vs. $>$, which is all **K** cared about.

Finally it was **T**’s turn. Interestingly, despite that **T** was representing my secondary topic of set theory, I felt after the exam that I performed best on **T**’s questions. His time had more of a straightforward question–answer character than the rest of the exam, in which I spent a greater proportion of the time hemming and hawing. Anyway, **T** began by asking for the statement and proof of the Δ -system lemma. This was easy, as I had studied it just a few days prior. I initially forgot the finiteness assumption on the sets, but remembered it when it came up in the proof. **T** next asked whether the lemma remains true if we replace $|A| < \omega \forall A \in \mathcal{H}$ by $|A| \leq \omega$ in its hypotheses. Here the answer is clearly no, which I quickly showed by taking \mathcal{H} to be the branches of a complete binary branching tree (which seems to be a counterexample to almost everything). **K** then interjected, seemingly curious himself, “What’s the largest Δ -system you can find in that gadget?” “Two?,” I said hesitantly, which **T** confirmed. **T** then continued with the next logical question: How could you modify the hypotheses in another way to get a version of the Δ -system lemma that allows for countably infinite sets in \mathcal{H} ? My first thought, which I gave, was to kick ω_1 up to ω_2 for the size of \mathcal{H} . **T** was nearly satisfied with this, but wanted one more thing... “the continuum hypothesis!,” I replied, to his satisfaction. Mercifully, **T** didn’t ask for a proof of the generalized Δ -system lemma, deeming it too technical.

T switched gears, to ultrafilters, for his next line of questioning. He first asked for a proof that there exists a nonprincipal ultrafilter on ω . “Zorn’s lemma,” I said immediately, shortly followed by the standard proof. When he saw I knew where I was going, he moved on: “You used the axiom of choice,” he said. “Can you do it without AC?” “No,” I replied, “because a nonprincipal ultrafilter is a monster, and you need the axiom of choice to conjure monsters.” He was understandably not satisfied with this answer, and asked for something more precise. My next try was to claim that a nonprincipal ultrafilter would be a nonmeasurable subset of $[0, 1]$, which everyone knows would not exist without the axiom of choice. Okay, he responded, but can you prove that the ultrafilter would in fact be nonmeasurable? At this point I finally remembered the trick he was fishing for. “I can wave a wand from probability theory?” I asked. “Of course! You’re a probabilist—or, a probabilistic combinatorialist,” he replied. The wand I waved was Kolmogorov’s 0-1 law, which

implies that if measurable the ultrafilter would have measure 0 or 1, but if the “flip all bits” map is measure-preserving, it must have measure $\frac{1}{2}$, contradiction.

T next switched gears once again, this time to recursion theory. He first asked for the definitions of the class of primitive recursive functions and the class of recursive functions, and for an explanation of why the containment is strict. I gave the definitions with no trouble, and for a witness to the strict containment I mentioned the Ackermann function. He asked me to expand on this (“the Ackerman function is just a name,” he said), so I gave the standard argument about bounding in advance the maximum possible growth rate of a primitive recursive function based on the number of times recursion is used in its definition. I gave an extremely hand-wavey rendition of this argument, but apparently a satisfactory one.⁶

T’s next question was whether the set of recursive permutations of ω forms a group. It does, of course, and I said so—for multiplication, you have composition, and for inverses, the unbounded search operator. His follow-up question was much harder: Is the set of primitive recursive permutations a group? I had some trouble with this. Multiplication is still fine, but inverses may not be. I suspected the answer is no (not a group), but couldn’t immediately see how to prove this. Sensing I was stuck, **T** suggested that I think once again about growth rates.⁷ This gave me the answer: There are permutations that send enormous things to small things, which could be p.r. with non-p.r. inverses. This is not a proof, of course, but it satisfied **T**.

We were running up on time at this point, and I sensed I was in the clear. I had managed to muddle through most questions and hadn’t hugely bungled anything. But **T** tossed me one last question as the professors were about to ask me to leave: “Give an example of a theorem in finite combinatorics whose proof requires reference to infinite sets.” This question was obviously a shot at **Z**,⁸ whose hostility to the infinite is well-known. I sensed this was not supposed to be hard, but I was drained and anxious to hear what I was confident would be good news. So after thinking for a minute and not coming up with anything, I asked—“there exists such a theorem?” To this **T** declared, as **Z** had earlier, that I would fail automatically. **T** gave as an example the strengthened finite Ramsey theorem, whose unprovability in PA is the Parris-Harrington theorem.⁹

I left the room to allow the committee to deliberate. After a surprisingly short time—what felt to me like no longer than a minute—I was greeted by **B**’s smiling face. Hands were shaken, forms were signed, congratulations were given. I slept soundly that night, for the first time in weeks.

⁶ It occurred to me after the exam that the Ackermann argument is probably not what **T** was looking for as a proof of the strict containment. He had covered this material at the end of the previous semester in his course on Descriptive Set Theory, and there he mentioned that the Ackermann argument for this point is unnecessarily technical, because there’s an easier one—you can diagonalize out of the primitive recursive functions, but not the recursive functions.

⁷ Actually, before making this suggestion, **T** asked the following doozy of a question: “Does there exist a group of primitive recursive permutations on ω ?” Suspicious, I answered deliberately: “Yes, the trivial group.” “Oh sorry sorry sorry! That’s not what I meant,” **T** replied, to general laughter. I still have no idea what he did mean.

⁸ If there had been any doubt, **T** confirmed this to me later.

⁹ **T** also told me later that **Z** remarked to him during the deliberations that this is a beautiful theorem, even if, in his view, false.