Rachid Finge (00:01):

Welcome to the Made by Google Podcast. I'm your host, Rachid Finge, and this is the podcast where we bring you behind the scenes at Google, talking to the people who work on our devices and services. Whether that's your Pixel phone, pixel watch, Android, or one of our apps like the camera. If you've missed any of this, then why not check out one of our previous episodes, whether it's from this season or the first one, and subscribe to the Made by Google podcast to make sure you don't miss what's ahead. Now, today is about one of our apps, an app you might not even think about as an app. It's so integral to your phone that it's often referred to as the phone app. It's why we call a phone a phone, even if you do many other things on it than making calls. What's even funnier to me, the team that works on the phone by Google app is actually called the Dialer Team. It's a true blast from the past, and yet the dialer team works on the future of phone calls. So let's hear what they're up to from today's guest in the Made by Google podcast, he's the group product manager of the Dier team. Please welcome Jonathan Eccles.

Rachid Finge (01:11):

Jon, welcome to the Made by Google Podcast. Please tell us a little bit more about your role and how you ended up at Google.

Jonathan Eccles (01:17):

Awesome. Yeah, thanks a lot for having me. So I lead the product team for the phone by Google App. My role consists of working across a really, really talented team, building a lot of great things for, you know, in some cases, just making sure that basic calling feels amazing and feels perfect, but in other cases, really taking the calling experience beyond what anyone thinks is possible and whether that's ways to protect you or ways to make your calls more productive or expressive. That's something that we strive to build.

Rachid Finge (01:45):

Perfect. And how does someone become a, a product manager for maybe one of the most important apps on a smartphone?

Jonathan Eccles (01:52):

It's a funny journey here. I mean, I've had a weird obsession with audio, I guess my whole life, and this is sort of the highest reach pinnacle of audio is people talking on the phone. It's the thing that billions and billions of people do no matter where you are in the world and no matter, you know, what your socioeconomic status or anything, you know, for me it started actually with like music and listening to music and then making music and then recording music, and I got to be an audio geek from that. I got into monetization of audio and music and content. I was working at Pandora for a while and, you know, opened my mind to things like how do people listen when it's audio ads, when it's music content, when it's podcast content? I moved from there into Google Assistant, actually for conversational AI. But again, the case there was more around how do people use conversation and listening to arrive at media to play, whether it's music or video. And then that landed me here in the phone app. It's funny, of all the things I've done so far, this is the one that stands out as actually being the most exciting and interesting and just greenfield when it comes to that idea that it's the 150 year year old technology that seems totally boring. Yet under the surface, there's these incredible opportunities to just do things with how people listen and how people converse and how people consume audio that you know, people just don't think is possible yet.

Rachid Finge (03:12):

Today's guest works on making phone calls better for everyone. Jonathan Eccles calls himself an audio guy, which comes in handy when you're the group product manager of the dialer team. Jonathan plays guitar, bass drums, and even was a music instructor for a while in the past. He worked on the software and hardware that professionals use to record and produce music. He then ended up at Pandora before joining Google in 2019. As Jonathan will tell you, the dialer team does many more things than just making sure you can place and receive calls. In fact, it sometimes does the opposite, preventing you from getting nasty spam calls. And then there are many more advanced features that Jonathan is working on, some that you might not have heard about before. I hope you'll enjoy our conversation. So why is it exciting to work on a technology that's essentially 150 years old? How are you shaping the future?

Jonathan Eccles (04:11):

There's two main paths that we consider our vision working along. One of them is, you know, calling as ubiquitous calling is this incredible reach thing that everyone has built into their daily lives anywhere in the world and it has to work. And you know, the nice thing is actually if you look back decades and decades and the hundred years ago, like it's been really, really good at working. People could take any analog phone, dial a number and get to someone else. It was really reliable, it was really fast, and we need to hold ourselves to that same standard of perfection for the basics. You know, you have to be able to make and receive a call and you have to be able to be in a call and it has to work incredibly well. It has to be reliable, it has to be performant and fast, and it has to feel incredibly easy to use. So that's sort of path one is we've just gotta be perfect at that, but path two, which, you know, a lot of people don't think of as like a thing you can do underneath the calling experience is looking at all these places where calling as a behavior feels like something that's friction, feels like something that's annoying, feels like something that's interruptive, feels like something where you can't express yourself well enough. And we look across all these things and think to ourselves, are there ways we can elevate the phone call? And that's with great design that solves important user problems and in a lot of ways looking for AI to empower new ways to solve problems that people didn't think could be out there. And we're doing that in a lot of things, whether it's trying to solve spam, so that's just never a problem for people anymore on phones, trying to make your business calls more productive, trying to make your personal calls easier and more delightful. These are all things that we think AI can really supercharge for us in our future.

Rachid Finge (05:40):

I bet no one at Bell 150 years ago thought, you know what? AI is gonna be used in the in phone calls. So let's talk about that. And specifically about receiving calls, particularly from unknown callers on pixel phones in many countries. You can have the Google Assistant answer the phone, ask what a call is about. Why do we have that feature and how do you make it work?

Jonathan Eccles (06:00):

Yeah. You know, the history of call screen and spam on phone calls is very interesting. You know, it's one where you can take this base layer of trying to understand what people report a spam or what looks like spam and building machine learning models that figure out these phone numbers are spam and these ones aren't, but you know, we did call screen because you don't just stop at machine learning on identifying phone numbers. What we wanted to build was an extra layer between you and any of those suspicious or unknown numbers that could kind of facilitate that trust of the incoming call. And what I see looking forward is a place where this has solved forever. And I think it's possible, and I think it's possible sooner than a lot of people think, you know, if we look back to the start of call screen, it's this

thing that takes text to speech, it allows, you know, assistant to ask someone, Hey I'm an automated assistant and you know, I'd like to understand why you're calling and it takes transcription from the incoming caller. And then that turns into kind of two modes of protection. One is manual call screen whereas a user, I don't have to actually talk to incoming callers I may not want to talk to, I can just see what they're saying and make a decision about whether to pick up. And that's something that's available on Pixel and a lot of Android and around the world. And then there's a special flavor called automatic call screen. And that's something where we actually are allowing you to keep your phone in your pocket before you even knew it was ringing. And we can try to make a decision by picking up on your behalf to say, this sounds like a robocall, for example. And that's worked incredibly well. It's something that you know, it handles just so many different cases for spam and is just beloved by Pixel users.

Rachid Finge (07:37):

Now my pixel phone, whenever I get a spam call, my screen basically turns red. It will tell me like, you know, this could be a potential spam caller. How do you know that? Is that, is that also machine learning? And then what does it train on?

Jonathan Eccles (07:51):

Yeah, so when we're identifying spam calls, we're looking across the world's set of phone numbers and we have a lot of magic happening in our machine learning to make decisions about that. And that's looking across what users are reporting to us and and trends like that. And we're also using things like call screen to be able to figure out that kind of information too. But we see a lot of just runway ahead in learning more and more, especially with that idea of can we do more with conversational ai? We think that will uncrack a lot about as an interesting comparison, email spam is actually a much smaller problem than it was 10, 15 years ago. Sure. And that's because there's an idea that, you know, in a very, very good, trusted, private way, you can have your email assess the content of something that's coming in and make a decision. Right? That's a space that's really hard with phone calling today because it's happening in an analog format typically. And it's happening in real time. So we want to figure out are there ways we can break that barrier and actually achieve something that, you know, has been achieved with email where there's a level of understanding content to understand whether phone calls are safe.

Rachid Finge (09:02):

Since you mentioned email, I'm just curious, is there any, any synergy between the teams working against spam at Google? So maybe within the dialer team and within the Gmail team on how to fight spam? Do engineers ever talk to each other about, you know, new tricks and new ways to fight spam?

Jonathan Eccles (09:19):

Yeah, yeah. And folks, definitely the teams talk quite a lot. There's sort of shared methodologies, you know, I'm not nearly enough of an expert to say what kind of really smart things they're doing under the hood there, but I think what's really special is seeing across Google, there's this shared interest across so many groups to solve these problems of spam and user safety that, you know, people across every PA really try to band together and, and share what's working to think through how to solve the bigger problem.

Rachid Finge (09:49):

Now, another topic I wanted to talk about is a topic where I didn't know that it was the dialer team working on it because I thought you guys were, you know, mostly working on making it look good, but

you also work on making it sound good. So on Pixel seven and Pixel seven Pro, we have clear calling, reducing background noise on the side of the person you're speaking to so you can hear them better, maybe even if they're in some, some sort of storm or something like that. Could you tell us more about how clear calling came to be and how it works?

Jonathan Eccles (10:17):

Yeah, so hats off actually to, in the pixel org it's folks in that team who kind of have been the geniuses behind making clear calling and the AI behind that work. So you know, what, what I'm saying is, is really in credit to what they've come up with. It solves, you know, something that just so intuitively an important problem for people, which is when I'm on the phone I want to hear the conversation. Right? It's simple as that. What I mentioned before of the basics, just have to feel perfect. But there's some pretty interesting complexity that can happen under the hood of making these basics feel per perfect with clear calling. You know, you have a technology that's able to live entirely on the device. So it can happen with real time low latency and privacy that can happen by keeping everything on your device as the call is happening and processing what's in the audio so that we're actually running it through models that try to understand what are the noises that typically happen during calls. Maybe it's the sounds of dogs barking or breaking glass. And actually also models that are well tuned to understand this is the sound of human voices. And what's interesting here is noise canceling is not a new thing. There's been noise canceling for a while on headphones and phone calls and things like that, but that only works in kind of a way of trying to figure out what's ambient outside of a microphone and then what's coming in. This is actually another layer of that, of all the things that get through and things do get through very often, truly identifying the salient streams of the audio that are what should intentionally become part of the phone call. So it's a really exciting thing. I use it on my calls all the time. It works incredibly. And I think the team in pixels just done an insanely good job building that as a feature.

Rachid Finge (11:59):

Any clue why it's only on Pixel seven and seven Pro and not on the six series, for example?

Jonathan Eccles (12:04):

Yeah. Yeah. So in this case, pixel seven and Pixel seven Pro have, you know, the latest chip Tensor two. And what's happening here is actually requiring some pretty heavy duty processing with some really important needs around performance and latency. And that's something that you know, is possible today on seven and seven Pro. We wanna make sure anywhere that things come out, they're releasing with high quality and performance. And so that's the place where it works really, really well.

Rachid Finge (12:30):

So this is using machine learning I presume now to to learn, you know, what is noise and what isn't. So how do you go about that? For example, last week we were talking to the team that works on fall detection on pixel watch, and to my surprise, yes, there were stunted men involved who had to fall in order to train the system and the algorithm. So how do you get all the right kind of noises into the system and make sure that, you know, you have sort of this wide coverage of, of noise that that might happen during all the phone calls that happen every day.

Jonathan Eccles (13:04):

Yeah. at a high level, there's an interest in, you know, looking at the types of data we could test internally and creating these types of environments where we can train data. So you might imagine cases of, can

we look at types of calls that we might make internally for instance, and say, okay, let's find ones that are happening at a cafe where there's a bunch of people talking in a background. Let's look at cases where there's coffee grinders or wind or dogs barking. And the interesting thing about a lot of these is that you know, so often noise canceling works well on a static sound that's just there in one place in the background. But if you have someone that's moving in and out of different things and and shifting around and the positions, the sound and the profile of it's changing, that's where machine learning can actually do some interesting magic where you can start figuring out like, Hey, you have some, a stream nearby, there's pouring water on pebbles. And that's a very notable sound. And if you can actually train what that sounds like to a model, you can actually say, okay, well let's say you're moving around the room or moving around the forest or whatever and the stream's moving around you, there's still an interpretation that can figure out that's a stream. I know what a stream sounds like. I can get that out of the phone call.

Rachid Finge (14:12):

Amazing. There's some other things where the phone by Google app can help for example, calling customer service phone numbers where you get these lengthy menus you need to go through and Pixel can actually shorten that by quite a bit.

Jonathan Eccles (14:26):

Yeah, yeah. You know, there's a lot of things we hear from users that calling a business feels painful and a lot of folks just kind of assume that's what it's like to call a business, you know, customer service. Yeah. I'm gonna be stuck in a phone tree for a while. I'm gonna wait for five different options to get read out. Press four if you wanna talk to an agent about this and those are all things that people like resoundingly say like, that bugs me, but they aren't asking for a solution yet. We're providing one actually. And I think solving a lot of these problems for people and doing better and better at it. So some things that we do to help this experience end to end are, you know, at the outset of a call we can do something like wait times, where we look at as you're about to make a call is the time of day you're calling to this customer service line, the type of thing where there's a lot of traffic coming to it. Kind of in the same way that, you know, Google Maps has this type of data, right? The next phase is if you find yourself stuck in one of those phone trees and those phone trees can be really lengthy and really painful. Well, we do two interesting things. One is we transcribe it for you so that you don't have to think back and say, ''Oh, I think I heard the sixth option, but which one was the one where I needed to call about my account number Is that two or one you can look back''. And not only that, you can actually press those transcriptions as if you were pressing 1, 2, 3, 4, or 5

Rachid Finge (15:44):

On the screen?

Jonathan Eccles (15:45):

On the screen. Right. So that idea is how can we turn the concept of a phone tree almost into a graphical interface as if you went to their website. That becomes a super cool thing. And we take it a step further as, as we launched, you know, last year on Pixel where for a set of numbers that's constantly growing, we're not just transcribing real time. We actually know in advance what their phone tree looks like. So if you were about to spend five minutes going through all those different options and everything, maybe we can shrink that to 30 seconds. You say, okay, great, here's the selections. I press three here because I'm just reading it and I see the next ones already and I'll press five here and that's gonna take me to an

agent. Voila. Like, this is a nice kind of magical feeling, time saver. And then the last element that's super cool here is that moment of being put on hold. Maybe you broke through the phone tree or maybe you finally said, okay, talk to an agent. And then it says, okay, you're gonna be on hold for a while and you say, ''Ah, do I have to have my phone on speaker phone? Do I have to be waiting by it for the next 10 minutes or one hour''? And that's what hold for me solves with hold for me. You could have your own assistant hold on your behalf as ai, you can put your phone away in your pocket. And when an agent comes on the line that AI is actually able to figure out you've gone off hold. It's able to tell them, ''Hey, you're on your way'', and then your phone will ring to let you know you can come back. So these are all these incredible productivity builders that happen when you're on those types of business calls that just annoy the heck outta you.

Rachid Finge (17:12):

Amazing. I've never actually, because you know, I'm usually not at the agent's end. Do you have any clue of what the assistant says to an agent while, while they ring my phone?

Jonathan Eccles (17:21):

Oh, it's something I should know verbatim but I don't, I can't tell you the exact script pretty much. It's a very simple sort of message, kind of letting them know that they're an assistant. The thing talking is an assistant that's acting on behalf of the user and you know, just hold on one minute while we get them back. It's the type of message that is helpful and hopefully perceived as friendly. But you know, I think over time we see more and more on phone calls, people will need to get used more often to talking to ai. We think there's a lot more of that in the future and the more helpful it is, I think the more agents and users alike will feel comfortable in those conversations.

Rachid Finge (17:54):

So speaking of the future, what are things, because we described things that you can already use today, depending on where you live, what are things, and I know you cannot say everything, some things that you are working on right now to create a new decade of making phone calls even better and more exciting.

Jonathan Eccles (18:11):

Yeah, so I'm incredibly excited about a few things that we're experimenting with internally and, and thinking through ways to solve bigger and more important problems for, for our users. In particular, we think that a lot of the work we've done around spam, both for our machine learning of identifying spam numbers and also the work we've done on call screen, both manual call screen and automatic call screen to catch things like Robo callers. We think there's a world there that can go even further. It's early days, but that idea of multi-step, multi turn conversational AI could open a lot of interesting doors in terms of creating this protective and helpful layer at the front of every incoming call. We think that with AI used in the right responsible ways, in the right types of conversational multi-term bot methodologies, that there's a future where you should never, ever, ever be annoyed at the thought of your phone ringing. It should always be a moment where you assume it's something important or something delightful. Nothing should ever bother you. And at the same time, you should never, ever feel like you're ineffective or unproductive on any phone call. We think AI is what takes you there into the future and stay tuned this year. There'll be some things coming up in the world of solving for unwanted calls that I think are gonna be really exciting.

Rachid Finge (19:30):

This definitely sounds exciting indeed. And I'll let the listeners read between the lines, see if they can figure out what you're alluding to. That definitely sounds very, very appealing. Now Jon, we always close out every episode with a top tip for our listeners. What can they do to get the most out of the phone by Google App

Jonathan Eccles (19:53):

I'm gonna cheat and say two tips. If that's all right? Just because they're so damn unhelpful. First tip is everyone turn on clear calling. It's super easy in your settings. It's just one little flip of a switch and it will just make everything sound a bit better and it's super cool. Definitely. it's an easy one. Everyone on a Pixel seven and seven Pro should just go for it. It's great. And then tip number two is everyone should go in and turn on automatic call screen on your pixel phones. That's a feature, you know, if automatic call screen's available in the United States, but it will just change your experience where you don't feel like there's as many of those moments of your phone ringing and that awful robo call on the other end. These are super helpful things and it's something that, you know, if you've got a pixel, they're both incredibly easy to do.

Rachid Finge (20:40):

Well, those are two wonderful tips and definitely can vouch for clear calling. It's absolutely great. Jon, thank you so much for joining the Made by Google podcast. It's been great to learn more about the dialer team and that there's a long future ahead for a 150 year old technology.

Jonathan Eccles (20:57):

Thanks so much. It's been great to be here.

Rachid Finge (21:00):

Thanks again to Jonathan for taking the time today. I truly have a whole new appreciation for the phone by Google App and the people working on it. So even if phone calls have been around for a century and a half, it's clear that they're here to stay. We'll see a lot of innovation going forward. We're not done yet with this season of the May by Google podcast. So do subscribe on your favorite podcast platform. Maybe it is Google Podcasts and you'll hear more from us very soon. So thank you for listening. Take care and talk soon.