Hello again, Ars Technica readers. This is the second installment of a three-part interview with George Church – who is one of the most influential people in the worlds of synthetic biology and genomics. If you haven't yet heard part one, there's a link to it on the page where you found this player, and I strongly suggest that you go back and listen to that installment before this one.

And with that - back to my conversation with George Church.

Rob Reid: Let's talk about what makes CRISPR exceptional versus previous editing methods, both quantitatively in terms of price performance, and speed' and also qualitatively, in terms of things that it's good at that its predecessors weren't good at.

George Church: The two main things would be price and efficiency. Almost nobody makes their own CRISPR. They purchase the enzyme in the form of a gene, and they purchase the RNA that goes with the enzyme custom. So, you pay some custom price, which varies depending on how smart you are at getting the right vendor, but it's in the order of hundreds of dollars, and maybe the TALENs are $1,000 or something like that.

Rob Reid: To clarify for those who don't know, TALEN is a fairly new approach to gene editing whose popularity peaked right before CRISPR. And it sounds like CRISPR is cheaper but not radically cheaper than TALENs. Clearly, then, CRISPR's advantages go far beyond cost.

George Church: Then efficiency is the other thing. There was a nice study early on. The same sites were targeted by TALENs and CRISPR. On average, it was about a factor of four. Our paper was the first where we compared TALENs and CRISPR in the same site, and we had a factor of 20, but that was one anecdote. This paper did nine different sites, and it was an average of four.

Rob Reid: So, the efficiency went up by somewhere between four to 20X, which is substantial. And I sense that efficiency has a very precise definition here. Does it refer to yield?

George Church: Once you get it into the cell, what's the probability that cell gets altered?

Rob Reid: In the way that you want?

George Church: In the way you want.

Rob Reid: Okay, so CRISPR improves efficiency by 4X or more. What odds of a successful alteration does that map to in raw percentage terms?

George Church: It depends on the tissue. It depends on your organism, the particular way you do it. It can get up into the 90% range, but some tissues are still in the 1% range. So, that factor four was going from, say, 10% on average to 40% on average.

Rob Reid: So, 4X is clearly a very respectable improvement. But you've seen much bigger improvements on other fronts.

George Church: Yeah. We get excited about that factor four. But then there's that factor three million, which is the reduction in cost of sequencing.

Rob Reid: So, there is substantial room for improvement – not only in efficiency, but in the elegance of the process Because you've been quoted saying that, in some ways, CRISPR is almost more like vandalism than true editing. And a moment ago, you compared it to bashing. How does CRISPR fall short of an ideal system that would do true editing?

George Church: Well, you can get it to put in what you want. You put the donor DNA, and then you cut the recipient, and then you hope the donor DNA defuses to the site where the recipient's been cut to act as a bandage to repair it. But the host can't wait. The host cell is saying, "Hey, there's a double-strand break. I'm freaking out. Let's fix it," and they fix it by a process called non-homologous end joining, which typically removes a base, or adds a base or two or three or 100, and you get some random thing. The previous set of gene therapies were additive. You're genetically missing a gene; your gene therapy adds it back. So you need it to be able to subtract.

Then the third thing you want it to do is whatever you want. You want to be able to do something very precise. Let's say sickle cell has a particular base pair change – you want to change it back. And that's neither addition nor subtraction. It's precise editing. And if we had precise editing that was cheap and effective, we'd do everything that way. If you edit a document, you don't want to say, "Oh, I can only remove some random number of letters from a word," right? You could say, "Well, that word's wrong. Let's just change it to some other word," or, "Let's just delete a few letters." You want to change it to exactly what you want.

Rob Reid: In a sense, for all its great improvements, CRISPR is still kind of like being able to swap out a paragraph? Like, I can take out a paragraph and glue in a more perfect paragraph, but I can't dive into a paragraph that only has two wrong letters and precisely change those two letters?

George Church: It’s just inefficient. Many of the projects that my lab wants to do –and many other labs – is not change one gene, you want to change many of them. And it's so inefficient at doing one that it's hard to even imagine doing multiple ones.

Rob Reid: The reason is – Just to make sure I'm understanding this correctly – if you've got a 40% change of making any one particular change successfully, your odds of making two changes is 40% of 40%, or just 16%. And your odds of making three changes is what? About 6-1/2%. So if you want to make 50 changes, your odds of pulling that off are 40% to the 50th power, which is essentially zero.

George Church: Yeah.

Rob Reid: So, to do hundreds of changes, you really need to push your efficiency up to like 99% or something. Plus, you want to become surgical as opposed to additive and subtractive, which seems like a tall order, but given that we've already had 10 different types of editing since the '80s, we'll probably have 10 more in the next couple of decades. And at some point, we'll get there, right?

George Church: It'll probably be fast, too. I mean, TALENs lasted for about three years before they were replaced by something else. It's already been five years since our first publication of CRISPR as a technology, in 2013. We're overdue.

Rob Reid: And I'll add that CRISPR almost seemed to come out of nowhere. And not just to outsiders, because I also know lots of people who are very deep in the field who just didn't see it coming. Now, for you, as a super-deep insider, how clear or opaque is the pipeline of future editing technologies today?

George Church: Some of these things look faster to people that are far away from it than people who are close to it. For example, fluorescent next-gen sequencing looks like it came out of nowhere. It came very quickly, but I watched it slowly develop from our first paper in 1999. But for most people, it suddenly popped up in 2005, Okay? But that's still pretty fast. With CRISPR, to me, it was on a continuum with all these other methods like Smithies and Capecchi and the meganucleases and so forth. And so it looked slow to me. It was another factor of four. We already had several factors of four from Smithies and Capecchi's stuff, which was one in 10,000.

Rob Reid: It's interesting. I had an interview with Rodney Brooks a couple weeks ago, and when we were talking about AI, he made a similar point – which is that while neural networks and back propagation seemed like sudden and radical developments on the outside, the field had actually

been kind of pregnant with them for years and years and years. So when they did finally happen, it did seem like a completely discontinuous change to us non-AI people, but it was actually nothing of the sort.

Anyway, moving on to the practical applications of editing, I'd like to talk about xenotransplantation. Partly because it's such a cool word, and I feel smart whenever I say it, but mainly because it has extraordinary potential for humanity. And it's being enabled by these advances in editing, so let's talk about what xenotransplantation is, why progress in it came to a grinding halt about 15 years ago, and how these breakthroughs in editing that we're getting now might get it back on the fast track.

George Church: It was a great idea, about 20 years ago, that if you look at pig organs, they're all about the right size and shape to replace all the organs we get currently by going between people. For a while, it seemed like the problem was finding a match. But it's mainly there just aren't enough people that are dying under circumstances where they can donate organs. Usually, they've got multiple organ failure, or they're really old, and they're not good donors.

Rob Reid: And millions of people die awaiting a transplant.

George Church: Exactly. I mean, there are millions of people that could benefit from transplants. Some of them aren't even considered seriously right now because there's such a shortage. They're going to say, "Well, you're in such a long queue, that we're not even going to call you part of the queue. You have a kind of disease which will just recur if we give you an organ."

And 20 years ago, the idea of the pig was there, and I think it was quickly appreciated that when you transplant, you get a very acute reaction, which goes beyond the transplant reaction you get between humans.

Rob Reid: By acute reaction, you mean the body rejects it?

George Church: Rejects it very fast, in hours not in weeks. And that's because pigs have different sugar on the surface of their cells. Then they figured out a way to fix that. Well, that's one gene, and that did get rid of the really fast reaction, but you still have all kinds of slower reactions, and there were clearly multiple things wrong. And the other thing that killed it 20 years ago was that it also slowly dawned on people all the pig cells, in every pig in the world, were producing viruses.

Rob Reid: The cells themselves were producing viruses?

George Church: And it wasn't because the pigs were infected. It was built into the pig genome. So a natural property of being a pig is that you make viruses. And it's not just pigs, many mammals do this. So this is a worst case scenario. Put an organ from a pig into a human, you save the human's life temporarily, but the recipient is immune-suppressed because that's the way you do transplantation.

Rob Reid: You put them on drugs to knock their immune system down so they won't reject the organ.

George Church: But now they also won't reject the viruses.

Rob Reid: So, now there's a swarm of pig viruses in our patient.

George Church: That are going like gangbusters because they're immune-compromised, and that means they can start evolving.

Rob Reid: The viruses can start evolving?

George Church: The viruses start evolving in that patient, and then they could infect healthcare workers and family members. You would not want to be the provider of that technology.

Rob Reid: Nor would you want to be the person changing their sheets or serving that poor, ailing patient breakfast.

George Church: So that caused a slowdown.

Rob Reid: And it was a billion dollar effort in the '90s, wasn't it?

George Church: Yeah. Novartis and many other companies were investing more than a billion, and it was worth it. It was considered a good effort. Then, when we invented CRISPR, we suddenly got – within a few weeks " three different pioneers, people that really had lived through this era and had dreamt the dream for a long time. They each came to us independently and said, "Hey. You know, we'd love to have you use your CRISPR. Because we know there's not one gene. There's lots of genes we've got to change, maybe dozens. Do you think CRISPR's up to it?"

Rob Reid: They knew roughly where on the pig genome these viruses were lurking?

George Church: The viruses and the immune functions and the blood clotting and so forth. Either they knew or they knew somebody knew. And so they recruited us, and we started looking at the problem. And we said, "Yeah. This looks like a good choice for CRISPR." Although, it was still pretty challenging. At the time, we and others were changing one or two genes

at time with CRISPR, and we didn't even know how many viruses there were.

So the first thing we did was sequence a bunch of pig genomes to figure out how many viruses there were. And in the pig cells that we were going to test this on, there were 62 viruses. So we had to make 62 at once. And we just went through the calculation of how inefficient CRISPR is, but we said, "Well, we've got to try it anyway. It seems like it's going to be the product of a bunch of inefficient steps .4 to the 62 power. Doesn't seem very hopeful, but we're definitely going to fail if we don't try."

And at turned out it was much easier than we expected, embarrassingly easy. At 14 days in the incubator and almost no human effort, and then a little bit of PCR screening, all 62 of the viruses were now damaged in the way we wanted them to be damaged in the polymerase gene, but that was just the first part. Those were done in abnormal cells. People have a tendency to do CRISPR on cells that replicate well that are abnormal. And so we had to redo it on normal cells, so we could put it into a pig and develop baby pigs. And that worked.

Rob Reid: So, there are now baby pigs.

George Church: The first pigs in the world that have no functional endogenous retro viruses.

Rob Reid: When were they born?

George Church: They were born in early 2017.

Rob Reid: Where are they?

George Church: We're developing them separately in America and China because it's actually hard to move animals from one country to another.

Rob Reid: And you're in both countries because you started a company with a Chinese co-founder, right? What's the startup's name again?

George Church: eGenesis is a company that we founded after we started this project, along with Luhan Yang, who was a graduate student and a post-doc in my lab and one of the co-inventors of CRISPR as a technology. She is now chief scientific officer there. And they're not only piglets anymore, they're grown, sexually mature animals having their own babies. And so we think that particular step is not damaging to any step in the life of the pig. In parallel, we made a lot of changes that improved the immune quality and compatibility with humans and clotting, a number of things that make them more compatible. And now we're putting the other

final pig that will have all of these things in one pig, and then those will be used for testing in non-human primates.

Rob Reid: If things keep moving along as they currently are, how long do you think it'll be before the first pig organ transplants will be made into humans who desperately need them?

George Church: We need to get the non-human primate trials done, but that takes about six months. Then it could be year and a half, two years before we're doing human clinical trials. People can get fixed. Just because they're clinical trials doesn't mean you can't be fixing a lot of people.

Rob Reid: So human trials could be just a few years off, and if this works ... I mean, tell me if I'm being overly optimistic here. But it sounds like if this works, the organ shortage is over.

George Church: Yeah. It sounds crazy, but yeah, that's what it means. And it's even better than that – It sounds funny to even say that, but it's better in the sense that because these organs are engineered to be compatible, for the first time in history, you don't have to find a match. You're essentially engineering a universal organ. So you could have stockpiles. You could have them ready. And they're fresher too because one of the problems with transplants is you'll get an organ in as fast as you possibly can from some motorcycle victim or something. And it was in pretty good shape when it was taken out of the victim, although not perfect because you can't take it out until the victim is declared dead, and then there's the delay of shipment. Then the surgeon has to decide, "Well, it's a little gray. It's not so red. Should I use it?" Then you take a chance of hurting the patient you're transplanting it into. This would eliminate that.

But it's even better than that, in that because they're universal and they're engineered, you can now make organs that are resistant to pathogens, resistant to senescence, resistant to cancer and compatible with cryopreservation. So you could even have these things frozen, ready to use immediately, so you don't have to have a living pig in the surgical ward.

Rob Reid: So if you're in the unfortunate situation of desperately needing a kidney transplant, you not only get off dialysis and get a completely new lease on life, but now you're never going to die of kidney cancer. I mean, you could die of another cancer, but this organ, at least, is pretty safe.

George Church: It opens up that door. This is not a promise. We've made animals that are resistant to certain pathogens. We've made animals that are resistant to senescence and cancer. It would be very hard to engineer a

human to have these enhancements, but if you're going to transplant the organ anyway, you might as well give it the best organ you can.

Rob Reid: This whole development is very well-timed because the organ shortage is probably about to get a lot worse. Albeit for a very positive reason, which is that self-driving cars are expected to radically reduce the number of traffic deaths, which are our main source of organs, over the coming decade.

George Church: This already happened ... We reduced the number of accidents already with seat belts and airbags.

Rob Reid: Now, speaking of pig transplants, you're vegan, right?

George Church: I am.

Rob Reid: Does the notion of raising pigs for the purpose of accessing their organs sit awkwardly with your reasons for becoming vegan?

George Church: Well, I'm vegan for a variety of reasons. The main one is health. I think many people could benefit from a low cholesterol diet. I benefited more than I did from statins. That's number one. Now, there's humanitarian components too, but I considered that one pig saving 10 people makes a whole lot more sense than one pig producing bacon and pork chops that are going to hurt 10 people. But, all that said, the ultra-vegan in me would say, "This is a stop-gap procedure that is going to save millions of lives," but let's also start getting organs from human tissue culture. Meaning, lab-produced organs, which we are actively working on, and we've been working on since before the pig. If those arrive in time, fine. If they don't, then for years we'll be depending on pigs, and maybe forever. But the point is that as a humanitarian, I think that we need to try everything we can to prevent human suffering. And minimize pig suffering. We want to get as many organs as we can per pig and arrange it so the pigs have no pain, so they have a happy life.

Rob Reid: Now, there's another even more ambitious editing project, which is the bacteria which is resistant to viruses.

George Church: Ah, yes.

Rob Reid: Let's talk about that.

George Church: The trick here is a very simple idea, which is you change the genetic code – that is to say how you go from DNA to RNA to proteins. Which is a triplet code. There's four bases, A, C, G, T – and they can be combined in 64 ways to make triplets.

Rob Reid: A quick review for those whose biology is rusty. The total number of three-character strings that DNA can make using its four-letter alphabet is four to the third power, or 64. And each of those 64 combinations represents one of the 20 amino acids that life runs on, or the stop command. So a total of 21. Which leaves plenty of room for redundancy. Because, for instance, four of the 64 combinations code for the amino acid glycine. You don't necessarily need four. Two code for glutamine. Three combinations indicate a stop command, et cetera.

George Church: Right, exactly. You can move the code around in a way that doesn't hurt the cell. You can take TAG and turn it into a TAA, and it still does the stop codon function. You've freed up-

Rob Reid: You've freed up TAG entirely because you have two other codons that can say, "Stop." As for TAG ...

George Church: You can eliminate its use from the whole genome. That way, when the viruses come in, they don't know what to make of that, and four out of six viruses are completely baffled. They can't replicate.

Rob Reid: They're in an environment in which this triplet is just not used at all, and that throws their metabolism off so much that-

George Church: They cannot replicate. Even if you've put in viruses at very high titers, like 10 to the 12th trillion viruses, you get nothing out. So they are really, really resistant to these viruses – but only four out of six. We made that genome. We tested it, and it's amazing.

Rob Reid: This is E. coli?

George Church: This is E. coli, which is an industrial microorganism, which has a virus problem. That required 321 changes out of four million base pairs. We are redoing that, so instead of four out of six, it'll be all viruses. Same trick can be used in every organism that has a virus problem.

Rob Reid: How far along are you with this E. coli project?

George Church: We're mostly done with it. That was done almost entirely by CRISPR, and we're now scaling up to do that in other organisms, so industrial microorganisms, dairy microbes, agricultural plants and animals, and human cells for transplantation and pharmaceutical manufacturing.

Rob Reid: You've made this strain of E. coli essentially bulletproof. Where do you take it from here?

George Church: We're starting to do it in human cells. Mammalian cells, like human cells, are used for manufacturing a lot of pharmaceuticals like human proteins

that have to be properly coded with carbohydrates. That's also used for manufacturing vaccines and for transplantations. Each of those has either a risk of getting viral contamination, which would be really bad for a therapeutic. Or if it's being put into a human, it's getting exposed to viruses all the time. So there are reasons to make human cells that are multi-virus resistant.

Rob Reid: Let's move on to part three of our agenda, which is synthesis. The synthesis of oligos, how has that accelerated, improved, and cheapened since the Human Genome Project? I understand, by some metrics, synthesis has gone even faster than sequencing.

George Church: To give some points on it, in my lab, when I was a student, we ordered a piece of DNA that was 10 long – 10 As, Cs, Gs, and Ts in a row – for almost $4,000. Today, for $4,000, we can essentially get a human genome's worth of oligonucleotides. We can get billions of base pairs for $4,000. It's gone from 10 to three billion, more like three hundred million fold rather than three million fold.

Rob Reid: Wow. And when oligos are made, it's in sequences that are extremely short compared to any genome, certainly compared to the three billion letters in a human genome, right?

George Church: Two hundred is kind of a limit, but then you can assemble those into bigger pieces.

Rob Reid: Why do you have to start with these relatively teeny building blocks? Why do they top out around 200 letters?

George Church: As you go from 10 to 100, you're accumulating errors. The efficiency for a cycle might be 99.5, and then to the hundredth power, you start getting just a lot of errors. But then, as you get to 200, it starts getting even worse and worse as if it's kind of folding up on itself, and things are becoming less efficient.

Rob Reid: Now, this might be a bit of a false separation, but I think of those short oligo building blocks as being synthesis, and then the assembly of lots and lots of them into a single, long strand as being a second thing, as being assembly. Now, obviously, it's the combination of these two steps that result in long strands of DNA, but it seems like different entities do the mass production of short oligos versus the assembly of long strands. So, it's kind of two different steps, two different expertises. Is that right?

George Church: There are some companies that specialize in the short oligos. And most of the really big constructs, the things that are in the multimillion base range, are being done by academic labs.

Rob Reid: Do you expect the price curve for oligos to continue to drop at this blistering rate for years to come?

George Church: I do, but I can't guarantee it. I think both the reading and writing curves have quite a bit left, many factors of 10.

Rob Reid: Now, you said that the very skilled artisanal academic labs can now create DNA strands stretching into the millions of letters. Given that virus genomes generally top out in the very low hundreds of thousands of letters and bacterial genomes start getting into the millions, it sounds like we're starting to become capable of synthesizing bacterial genomes. Is that correct?

George Church: Bacteria and microorganisms, let's say. E. coli and yeast are probably the two most ambitious projects right now. And most of the assembly is done in vivo. The little building blocks, as you call them, the short 200-MERs are done by organic chemistry, so it's very far from biology. Then the little assembly is done in vitro, so it's biochemistry. Then once you get them in the two to four kilo based range, then they're assembled in living cells. Essentially by the kind of recombination we were talking about earlier – which we might call editing – assembly is a form of editing. And it's assembled, and then put into the right place in the E. coli or the yeast genome.

Rob Reid: Is the maximum length of assembled strands growing in rapid and compounding ways that parallel the price drops in gene sequencing and synthesis?

George Church: Maximum strand length is hard to say when you get to assembly in vivo because you're back editing. So the metric can't be the length of the thing you're changing, it has to be the fraction of base pairs you're changing. What's interesting is that's largely limited by design. We can't currently design a new chromosome that is more than, say, 4% different from an old chromosome that will still work and do something new.

Rob Reid: Because our understanding of the metabolic pathways and the underlying dynamics simply aren't rich enough?

George Church: That's right. Evolution typically only changes a couple of base pairs at a time per cell generation, and here we are changing 4%. That's pretty amazing, better than evolution, but still what engineering and evolution have in common is a lot of trial and error.

END SHARED ELEMENT OF SEGMENT TWO

And so, Ars listeners, we come to the end of the second excerpt of three from my interview with George Church. As mentioned before, if you can't wait to hear the

rest of it, you can just head on over to my site, at after-on.com. Or, type the words After On into your favorite podcast player, and scroll through the episodes to find this one, which originally ran on April 3rd. Or, you can join me tomorrow here on Ars, when we'll continue with third and final part of this interview.