

The Mix in *The Last of Us*

GANG IESD co-chair Kenny Young speaks to Naughty Dog's Phillip Kovats and Jonathan Lanier about the GANG award winning mix in *The Last of Us*.



KY: Hi guys, thanks for agreeing to do this. Could you introduce yourselves for the record?

PK: I'm Phil Kovats the audio lead here at Naughty Dog for *The Last of Us*.

JL: I'm Jonathan Lanier, and I'm the audio programmer here at Naughty Dog. The only one! I work on all our projects, flipping between whatever is in the pipeline.

KY: Do you like being the only audio coder - do you think you could do more if you had more people?

JL: You can always do more with more people, sure, but it's never that simple. I mean, obviously, when you are the only one working on something you don't have to worry about butting heads with people or dividing up the work. It definitely simplifies things in that regard. But, on the other hand, there are limits to my bandwidth because there's only one of me, and there are therefore times where we have ideas about what we could do, but we have to fall back to Plan B where the design isn't perhaps as dynamic as we'd like it to be or we have to pass on something and put it off until the next game. I don't think we're at a point where we really suffer as a result of that, but it'll be interesting to see if we can keep pace with the bar, or

if it's ever going to get complicated enough that we need to put more people on it. That's an interesting question!

But, in general, and this was one of the main points I wanted to put across in my talk at GDC this year, my feeling is that the entire industry under-represents technology for audio because they are so used to being told that they can't have any support - they can't have this, they can't have that. And I want to change that - I really would like to see more acceptance of the idea that we need technology in order to keep raising the bar for audio, and we shouldn't look at it as our enemy, we should look at it as something that can help us. We need to look beyond what we can get just using out of the box audio engines, because games are a dynamic medium and we need to be able to find better ways of tying our sound models in to the dynamic behavior of games, and not just environmental audio models but also the sound design, music and everything else.

KY: You mentioned that you can save ideas for the next game. I find that interesting in that that's not necessarily something I can do in my work, because when we make a new IP at Media Molecule it's quite often so radically different that we're creating a bunch of new technology for it from scratch. What were the benefits of working on, say, three *Uncharted* games, iterating on audio features as you go, and then moving on to *The Last of Us*? Ok, so it's a different genre, and it introduces the survival horror gameplay on top of the action shooter gameplay of *Uncharted*, but I can definitely see a relationship between that iteration and the polished presentation that Naughty Dog's games are known for.

JL: It's really interesting that you bring that up, because I think that's a great dovetail into talking about the mix in *The Last of Us*. There's definitely a lot of similarities at some level between *Uncharted* and *The Last of Us* - they're both 3rd person games, they're both cinematic in style, and in both games we want to have exposition and critical dialogue happening during gameplay. There are certain traits of Naughty Dog's games that are common to both, but as you pointed out it's also a completely different genre of game, and that radically affected the mix. You know, I actually looked at the different sections of my code and the different models, and there wasn't really one section that went untouched. I did a code diff at one point and it turned out that I'd pretty much modified and changed more subsystems for *The Last of Us* than I had since the beginning of the *Uncharted* series. The nature of the mix in *The Last of Us* really forced us to rip things up that we had taken for granted because they just weren't going to work - we had to make a lot of technology changes to get the mix to work right.

KY: One of the things that really stood out to us about the game when the IESD were reviewing and judging all the nominees for the Best Mix award was the fact that, quite often, the only thing going on in a scene is the audio! Like, even when the player-character is cowering in a corner and you can see very little, you can nonetheless still hear all of the infected everywhere, identify which type of infected they are, and even be quite specific about their location - whether they are two rooms over or are just outside the door. How did it come to be that you ended up working on such an audio-centric game?

PK: There were a couple of things that started that. One of them was that, even before I came onboard with Naughty Dog in August 2010, I believe you guys had a talk with Bruce Straley (Game Director) and Neil Druckmann (Creative Director) about propagation?

JL: Yeah, that was one of the very first things they wanted to know about. It was something we'd wanted to do a couple of games back on *Uncharted*, but we kept putting it off...

PK: It's one of those things that Bruce had evangelized. And, yeah, we started talking about it during *Uncharted 2* when I was a Senior Sound Designer on that game.

JL: They didn't know it was reverb-portal propagation, or how we would actually do it. But they came to us and they were like "we want to be able to tell if someone is in the next room - how do we do that, can sound help us? We really want to have this spooky feeling that there's something coming from the space next door. We want to be able to lead the player in to spaces..."



PK: And there was a lot of work required to pull that off. They created this temp level that was a convenience store where you could walk around inside and outside, and there were different rooms to explore, and Jonathan had started preliminary work on the propagation model. We already had occlusion and obstruction and reverb environments, and so we could set the different environments to sound differently. We mapped out scenarios where the player could be in a room with two entry or exit points and still be able to tell which door an enemy was closer to - basically, tactical information via your ear. And that was one thing that we felt was missing from this style of game - a lot of games punt on this for one reason or another, and it

can take away from the scariness. We felt that if a creature sounds like they are right on your ass when they are way over the other side of the room then it really takes away from the tension and feeling of dread, and I guess what we were going for was more of an emotive kind of experience.

JL: This was a critical factor in our decision not to use High Dynamic Range (HDR) audio. I've got a lot of respect for the DICE team, they're awesome, and when they came up with the concept of HDR audio a few years back they definitely raised the bar. We knew that dynamic range was going to be an issue, so our first thought was like "Hey, what about HDR? DICE did it, why can't we?" So we started thinking about it and doing some brainstorming. What we quickly realized was that the whole point of HDR is to make it easy to hear the quiet things - you want to be able to scale up the dynamic range of the quiet sounds in times when you're not in the thick of a firefight. But in our genre of game, making it easier to hear the quiet things makes it less scary. It turns out that things are scarier when you're straining to hear them! We knew that dynamic range would still be problematic, so that's when Phil and I came up with the idea of Parametric Dynamic Range (PDR), where we drive the dynamic range, not based on perceived loudness but instead based on the tension and the mood of the game state. That allowed us to have areas with exploration and exposition where audio that needed to be audible, such as the dialogue, would be brought up. Then when you were in an area where you were meant to be scared, we could set everything exactly where it needed to be so that you would strain to hear the things that were scary. I think that was really critical in getting it to work right.

KY: It's quite a confusing concept because, traditionally, certainly in games, you don't associate the things you want or need the player to hear with them being quiet or hard to hear! So it's quite an inspired way of thinking. You're obviously benefiting from the fact that the soundscape in your game is allowing for that - you have room in the mix to go for that level of subtlety.

PK: That's something that was present from very early on, even when I first got the pitch from Bruce and Neil. They were referencing projects like *No Country for Old Men*, *I Am Legend*, *Book of Eli*, *The Road*, which were very linear story-driven movies, but they drew the viewer in by the sense of space, or lack thereof, and tension especially.

KY: There's that scene in the motel in *No Country for Old Men*, where Bardem's character is walking down the corridor...

PK: ...yeah! He's walking down the hallway, turns off the light and he takes his shoes off - that was our benchmark! That was literarily our benchmark scene. The golden rung that we were shooting for was to have this ability to draw players in. At E3 2012 we had a behind the scenes demo in a theatre for the press and the moderators, and I was there to answer questions from the press, but what I did a lot during those three days whilst Bruce and others were showing off the game, was I sat there and watched the audience. And in those moments in the hotel where they were sneaking around and they were looking to see where people were, I could physically look at the viewers and watch them quietly lean forward in their seats. And I knew at that point we were on to something, because that's the point where people jump, right? When something

loud happens, they jump back into their seat, it gives them a shock, and they're immediately engrossed. So, the point of the dynamic range was to have the player inhabit that world, and everything that was happening to Joel and Ellie within it.

KY: Well, it worked!

PK: *[laughs]* ...you know, there are a lot of games out there that do a lot with volume. And there are some good mixes, but they're all very different games - you can't necessarily say "this game should be like that game." But what we truly wanted to do in this case, a survival action game, was we wanted to take the player and let them be in that emotive state of Ellie and Joel, and let them experience this world. So when it was supposed to be beautiful, it was beautiful, and when it was supposed to be horrific, it was horrific. And we really sculpted that emotional arc with the sound throughout the game.



KY: When you have an experience like seeing people play the game and their reactions to it before the game has shipped - how much does that influence what you then go on to do? Did you come back from E3 wanting to change work you'd already done or was it more the case that it cemented what was already in your heads and you were able to just go to town on that for the rest of the experience?

PK: I would say the latter, because we had already made a lot of decisions and assumptions about how we wanted to do this. We'd had internal play tests, and we knew what was working and what wasn't. And I had worked with Jonathan, Steve Johnson, and Derrick Espino on how

we wanted this game to sound, the specific audio quality of this game. And what we found out was that we were truly on the right track - we didn't have to go back and redo the aesthetic, we had to say "OK. Well, this is the hole we dug for ourselves, so let's keep digging!"

JL: I think it's worth pointing out that when we started out we didn't know how to make this game. We'd never done a game in this genre before. So it was a learning experience for us, and it was a problem solving exercise. We'd think "how would you do that in a movie" and "how would you replicate that in a game." Some cinematic techniques simply can't be replicated in a game because games are dynamic. So, being the audio programmer, my job was to find technology solutions that allow us to dynamically make creative mix adjustments to match the dynamic gameplay. I'm abusing the word "dynamic" slightly, but it's an important word because it describes the difference between a linear medium like film and a medium like games that's constantly changing. In the case of HDR, DICE really broke the mold, because they rejected the prior convention that a mix has to be a static thing or a series of static snapshots - the mix can be a dynamic thing that lives and breathes with the game. My light bulb moment for HDR mixing wasn't the HDR technique itself as much as it was "Hey! Dynamic mixing!", and so we came up with the PDR technique because that's what we felt this game needed. This goes back to your earlier comment about each game being different. I think different genres of games have different demands for the soundscape, and relying on any out of the box solution, bolting it on and saying "this is the one we have to use" is not the right thing to do. I think the right thing to do is to creatively figure out what the soundscape needs and then try to find ways of addressing those needs. Hopefully in the future we'll have a whole palette of dynamic mix techniques to use, because I think this is just the beginning of a whole new era of dynamic mixing - there's no limit to the ideas that people can come up with for that.

KY: The other interesting thing is that HDR is now available to anyone who wants to use it via middleware. Which is great! But just because you're using middleware doesn't mean you don't need an audio coder, because you're still going to hit those occasions where the middleware, as close to the bleeding edge as it is, is still only giving you preset solutions which don't necessarily allow you to do what is best for the game.

PK: I think there's something to be said about that. There's nothing wrong with using middleware - different tools can get the job done in many different ways. And more full-featured tools allow you to do a bit more. But, the idea of having a dedicated audio coder on staff allows you to look at what you are doing and mold the technology to those creative challenges.

KY: I think there are mixed messages about middleware. I got into this industry half-way through the PlayStation 2's lifecycle, and it felt like that was a point where the emphasis was very much on developing tools that meant the audio designer was empowered such that they didn't need to bother a coder - they could do it all themselves. That remains the promise of contemporary middleware solutions.

PK: I think that's largely because in the past it was felt like there was no love for audio, and the audio teams were not getting the support they needed, and so along comes middleware to the rescue, right? Well, that only gets you so far, because, yeah, you can create these sounds, but they still have to be integrated into the game in such a way that the behaviors of those sounds are implemented to perform the way your specific game requires and make it sound appropriate and interesting.

JL: I think it's kind of ironic, the whole "get the programmers out the loop" trend that seems to come from evangelization of middleware over in-house audio programming support. I actually subscribe to that philosophy, but I don't think that it means in-house audio programming is unnecessary! When I came to Naughty Dog more than ten years ago, when the sound lead wanted to get something running in the game he had to go find the game programmer associated with that feature and get them to wire up the sound. I was shouting "No, no, no, this is wrong, wrong, wrong, bad, bad, bad! We need to get programmers out of the loop!" Programmers won't have the time or the desire to edit a hard-coded volume number for a sound designer, that's absurd! Sound designers should be able to wire up most of the sounds in the game themselves - sounds that are bound to objects, sounds that are bound to the environment, these should be placed in a level editor and no programmer should have to wire this stuff up. Programming should be for special cases, not for 99% of the sounds in a game. I don't mind being cut out of the implementation loop as a programmer because there's plenty of audio programming tasks for me to do that are going to be much more helpful to the sound designers. My utility to Phil and Bruce Swanson (*Uncharted's* audio lead) is much greater when I'm helping them solve big picture problems with technology. So, whenever we can find a systemic solution to something where I can add a knob, I give them a knob! I'm like "Hey guys, here's a new knob for that thing you wanted, now you can tweak it." So, I give them knobs and then they do the work, and then my job is to just add more knobs! But there are still plenty of knobs to add, it's not like I'm out of work to do.

PK: Yeah, so it's not that we're against middleware. I think the idea is that middleware is just a starting point. The sound designer needs tools that allow them to be creative. But, in addition, what we've found here is that the model we need to make these intricate and integrated games is a very collaborative one. So we have to be able to work with the programmers, designers and artists, to be able to sculpt exactly the frame that we're looking for. Without that kind of integration and collaboration there's really not going to be that cohesive experience that you're looking for. For example, I have to work closely with the environmental art team to make sure that the collision properties they are setting works for the audio - so, that's things like correct surface types mapped onto the walkable surfaces, and that the walls are set to a certain type of collision thickness. If there's a chain link fence we need to be able to hear through that; if there's a pole in the middle of the room, I don't want that to obstruct the sound, especially if there's a gun behind it. In that scenario, we mark certain features as 'high SPL hear-through', so if someone's behind it their voice gets muffled but if they're shooting a gun with that kind of high SPL then it shouldn't be muffled. We work with the programmers besides Jonathan to get our physics system working right so all the bumpables (physics items) work OK. Here's a real special case - one of the lead programmers, Jason Gregory, took over

conversational dialogue for the game so, in essence, we had two audio programmers on *The Last of Us*. It was a passion project for Jason to create a fact and rule based conversational dialogue system that allowed us to get a lot of really nice contextual dialogue throughout the game.

JL: And I was happy to let him do it - I had plenty of other tasks to do, so that was great! Jason really did a great job with the new dialogue management system.

PK: And we were working with the designers to make sure that we understood the kind of tactical situations that they wanted players to be in so we could think about how we wanted to do these audio setups. Where were the story beats, the emotional beats? A good point to make here is that the sound in the game actually follows Ellie's emotional arc, and not Joel's. So, when I would work with Neil, the director, or the designers, I would try to figure out "what is Ellie supposed to be feeling right now?" We would sculpt the ambience, and the tones based upon those situations.



KY: Related to that, what were the conscious "pillars" of the mix in *The Last of Us*? You've mentioned the use of propagation, and dynamic range, and you mentioned Ellie's perspective just there. But were there any other high level, guiding principles?

PK: There were a few. One pillar of Naughty Dog's approach to mixing is that we start mixing with the placement of the very first sound in the game. We mix as we go - we don't subscribe to the "this is just gonna be temp" mentality. Even a "sketch" needs to be appropriate and work within its given context, otherwise people won't get it. And if people don't get it then you're

not doing your job. So we mix as we go - there's not this moment where we say "right, that's all the sounds in the game now - let's start mixing!"

KY: Right. That notion of "the final mix" is a very film-like approach to mixing. But when you're dealing with systems which are much more global and occur throughout the entire gameplay experience then the process of iterating on the mix takes a lot more time. So, as you play through the game and experience many of those moments, and different combinations of those moments, you tweak those models so that they perform robustly across the vast majority of the gameplay experience that players are likely to have. So, for me, the "final mix" is more a conscious attempt to sanity-check and sign off on a process that has been ongoing for months.

PK: I'll let you in to a little secret. We talked about this before during *Uncharted 2* - if we did not mix as we go, our games would not sound as good. Because the final mix for *The Last of Us*, this 16 hour game, was done in about 20 hours. And not over a couple of days! *[laughs]*

KY: Right! I'm sure that was a pretty stressful experience at the end of development. But you also probably wouldn't have benefitted from much more time because you wouldn't be able to make any radical changes, because a radical change would probably require a significant tweak to the technology, and that is a no-no at that point in development. It's not safe - you can't afford to make changes that you're unable to test and verify.

JL: It's also worth pointing out that the more dynamic your mix techniques are - HDR, PDR or whatever - the less sense it makes to mix at the end, because you've got knobs tied to other knobs tied to other knobs! It's not like you have a flat mixing board - you tweak one knob at a high level and you've just adjusted a dozen things underneath it that you might not even be hearing at the point you made the change.

KY: A nice way to think about this approach is that as you finish the game you should be tweaking the mix less, not more! You don't want to disrupt this model you have been honing and perfecting over time. In that regard it is the opposite of a film mix.

Whilst we're talking about linear mixing, it's a fairly common problem for the cinematics in a game to be at a noticeably different volume level to the in-game/gameplay level, even in big budget AAA titles that otherwise have high production values. I can understand how this might happen - cinematics are a prime candidate for being outsourced, or even 'insourced', and this division of labor can have undesirable knock-ons if not carefully managed. You guys outsourced the cinematics in *The Last of Us*, right?

PK: Yeah, it was Soundelux at the time, but they're over at Formosa Group now. Shannon Potter was our sound supervisor there.

KY: So, how do you ensure that the work you're getting back is going to fit in with your work? What's the communication process like to ensure that this kind of problem doesn't happen?

PK: As with everything else, it starts with being very collaborative - understanding each other and what our goals are. One of our goals was most definitely to have a seamless experience between the cinematics and the game audio. Even though we knew that they were going to be doing the cinematics, and they had a lot of material to work on, we had them hold off on doing final mixes till later because I hadn't quite worked out what the volume and dynamic range for the game was going to be. But once we had that, I got together some in-game examples and played them to Shannon and her mixer Chad Bedell - the first couple of mixes they did I was actually there with them to go over the process of matching the audio. But once they had finished a pass on a mix they'd send it back to us and I'd give them a couple of notes, so there was a little iteration - it didn't take a long time or anything, but that time we did take was crucial because that's what allowed us to have that seamlessness between the cinematics and the game.

JL: Actually, we have no option but to make them match. Naughty Dog likes their blends to be seamless - we generally play the in-game ambience through a cinematic so that they don't bump. So, the mix of the cinematic has to match the game because it's essentially layered on top.

PK: With only a few exceptions where a cinematic isn't set in a playable game location.

JL: And we also have a camera file that we edit to handle camera cuts and perspective cuts, so that the ambience matches the camera moves in the cinematic. It's a pretty complicated setup, I wish it were simpler!

PK: Yeah, me too...

JL: *[laughs]*

KY: So, you're editing metadata that is moving a virtual microphone about in the environment?

PK: Yup. Oh yeah.

KY: *[laughs]*

PK: Those are fun days when we get to do that.

KY: It's great attention to detail though. I just love that having decided to let the in-game ambiances run through the cinematics you didn't just stop there, you came up with a solution to the problem this imposes. The ambience is virtual and dynamic even if the rest of the cinematic audio is totally linear - it's crazy! I love it!

PK: But sometimes it's the opposite, because if you were to cut the mic every time the camera cuts or changes perspective it would drive you insane.

KY: Ah, right. Yeah, there are a lot of games with in-game cutscenes where the mic is slavishly tied to the camera, and the cuts are criminal - it really jars. But I suppose you have the flexibility to place the mic wherever you want.

PK: Yes. I've got to give a shoutout to Neil on this, because he - and this kicked me in the shins at the time, because I just wanted to be done - but he actually would not approve the final mixes until the camera metadata was done and the correct backgrounds were playing through the cinematics properly. Because he wanted to make sure that his scene played out the way he wanted his scene to play. And, you know what? That's frickin' awesome. Because when you're in the thick of it and you have 5000 other things to think about, you don't want something else piled on to you but, at the same time, it's so important to have all the pieces of the puzzle in place before you sign off on it.

KY: Man, I'm so glad I asked you about that because that answer was a lot more interesting than I was expecting! OK, so, going back to the pillars of the mix in *The Last of Us*, because we kinda went off on one there...

PK: [laughs] OK, so... just as with any Naughty Dog game, dialogue is king. So we wanted to make sure that we could hear the dialogue and understand what was going on, but it had to do so within the constraints of the environmental audio technology that we were trying to create.

KY: Could you talk a bit more about that - what were the challenges you faced here?

PK: There were a couple of things we did a bit differently for dialogue in *The Last of Us*. Early on we decided that we were going to create a more natural fall-off model. We didn't have a curve editor tool, so it was all numbers in a text file, but we had the ability to change the fall-off curve for each individual sound at will. We worked for a long time on the fall-off curves for dialogue, taking in to account the size of the maps, the different setups in each area, whether it was interior or exterior, all these different considerations, so that you could really hear the human enemies in the mix. But then once we started working with Neil Druckmann, the creative director, on the sound of the infected and how scary they were going to be we found that the kind of falloff curves we'd been applying to the human enemy dialogue didn't work on the infected's vocalizations - it just wasn't very scary to be able to hear the infected from far away. It communicated that there were infected present but that there was nothing to worry about, and that really diminished the power and meaning of those sounds. The infected are at their scariest when they are on your ass - with a character like the clicker, it's just one bite and you're done, it's a one hit kill - so what we wanted was for the player to associate the sound of the infected with an immediate threat. So, they said to us "yeah, everything we told you about the dialogue for the game - that doesn't work for the infected. Figure it out." So, Jonathan and I had to go back to the drawing board and talk about how we were going to make the dialogue not behave like dialogue!

JL: Yeah, we couldn't just split them out on different fall-off curves. If we'd done that then you'd get weird situations like hearing your buddy reacting to getting attacked but not hearing

the infected doing the attacking! So we knew we had to make the falloff curves match up when you are interacting with the infected. But we couldn't just drop the buddy dialogue to match the infected vocalizations because the buddy dialogue had to be loud and audible to make sure any exposition was coming across. And we didn't want any mix inconsistencies as a result of abrupt changes in state, like your buddy sounds loud and then all of a sudden they sound quiet as soon as you've wiped out all the infected. So that was where the whole parametric dynamic range technique came in to play, because what we decided to do instead was split the curves out but we would blend them over time to match each other when you were interacting and then unblend them back when the infected were gone. That way we could tune the two curves separately, and the blend would happen slow enough for it not to be noticeable. Even when you know it's there it's very hard to hear because it's very subtle.

PK: We got lucky actually, because there was a consistent rule in the design of the game that there could only be human enemies or infected in a given area - there were never both at the same time. Left Behind [*The Last of Us* DLC] kinda changed that a little bit.

JL: Yeah, that kinda broke that...

PK: [*laughs*] ...we had to adapt that a little bit! But then we got lucky like that...

JL: It wasn't luck! I mean, we took advantage of a...

PK: Ok, ok! We took advantage of a "feature"! [*laughs*]

JL: Well, you know, they asked us to do something impossible - "make this louder and quieter at the same time." We were like "You can't do that!" And they said "well, that's not good enough. It's an audio problem, you're audio guys - go figure it out." We were really mad about that! It was like "How dare they ask us to do something impossible - don't they understand it's impossible?" Thing is, you can't fault them for that because they're creatives - they're focused on how they want the player to feel in a given situation.

KY: Well, to their credit, being given a problem to figure out is a lot better than them suggesting a really bad solution!

PK: Yeah, and there's another side to that coin. They can go to an environment artist with a specific solution to a problem like "hey, we need some rocks placed over there because we don't want to see that area any more" then the artist puts in some rocks and it's done. But creating solutions to these kind of complex audio problems is an iterative process that takes time. Audio's one of those things that people aren't necessarily able to give specific instructions on, but they know when something's not working right and they're quick to point that out. Whilst we were trying to get this stuff to work someone sent out an e-mail to the whole team, all in capitals, that said "ALL DIALOGUE MUST BE HEARD AT ALL TIMES FOR ALL REASONS."

JL: And we were like "Err... no."

PK: Yeah, we explained “You don’t understand the process we’re going through right now.” They were getting frustrated. They were trying to create these AI systems and listen to the dialogue being controlled by those, but that was conflicting with the work we were doing developing how the dialogue would get mixed. And this is what led to the idea of the ‘dialogue sweetener’, which is another piece of custom tech. Basically, the problem was that there were areas of certain maps, such as Hunter City in the hotel where you meet Henry and Sam, and they wanted these characters to wander off and converse together but, at the same time, the player was not required to follow them. But there was important exposition that it was important for the player to understand, and so if this couldn’t be heard that was a big problem.

JL: The issue was down to the obstruction tech. If one of the characters walked around a corner, or went downstairs, or into another room, their voice would be obstructed. Even with long and loud fall off curves for these scenarios, you couldn’t hear them because of the obstruction. So they asked us “can’t you just turn it off?” But we explained that we couldn’t because then the dialogue would sound wrong, and this would break the sense of immersion. And they were like “We don’t care - this is really important to the story, it’s more important we hear the dialogue”, and that’s when they sent that e-mail. I was in panic, and I went to Phil saying “we have to fix this - they’re going to make us turn off all the obstruction on all the dialogue and that’s going to ruin the immersion, everything we’ve been trying to do will be lost.” We had a matter of hours to prove to them that we could fix it. So I put my thinking cap on... If you play a sound dry, and you don’t filter it, then it doesn’t sound obstructed or distant. But when you roll off the high frequencies it sounds muffled and quiet and you can’t hear it. So, I came up with the idea of leaving the direct path audio alone and fudging the reflected path - a really wide-band filter on the reflected path so it doesn’t roll off, and keeping the early reflections really tight to give the overall impression that the sound has bounced around a corner. That way it would sound distant but it would still be intelligible. So we tried it, and it worked, and they let us run with it. It’s not perfect, so we didn’t enable it everywhere. But, at the end of the day, it was a way of cheating to make sure that principal characters could be heard even when they were obstructed. When they’re not obstructed, we fade it out, cross-faded relative to the obstruction level. So if they walk around a corner into line of sight then the sweetener magically fades away, and if they walk around a corner the sweetener magically fades in. And in that way the exposition can happen, and it’s still intelligible, but it still sounds like they got obstructed and went around the corner.

KY: I think that’s a great example of mixing - you’re doing something that isn’t necessarily naturalistic, but it’s necessary to make the player experience as good as it can be. It’s the kind of mix decision that would be relatively effortless to achieve in a film, but in a game you need to develop and implement a system to do it that can work dynamically. Like you were saying earlier, we’re at the beginning of working out how to do these dynamic mix things, and it’s great to get a bit of insight into this.

PK: I know some people don’t like that comparison, but if you make these adventure style games with human characters you can’t help but compare them to films. I know that we’re not

making feature films, but they make some really great decisions and solutions to problems that we can learn from.

JL: There's a language there that works.

PK: There is. And if we look towards that - not trying to necessarily emulate it - but just look towards it, and try to figure out what we like and what we're wanting to sound like, that gives us a goal. We can actually shoot for something that works better.

KY: Here's an interesting example... last month [Develop magazine had an article](#) on the audio in *Forza 5* where Nick Wiswell discussed giving WIP video footage of the game to Skywalker Sound during pre-production, asking them to do a pass on it as if it were a movie, and the results directly inspiring their sound design and approach to mixing. I thought that was really clever - bypass the problems inherent in hiring a film sound designer to work directly on a game by getting them to concentrate on what they do best, in total ignorance of whether it was achievable or not in a game, in order to inspire the game's sound designers. For the right kind of project, that's definitely an approach that people could emulate and learn from.

PK: Here's a fundamental philosophy with Naughty Dog that both Bruce Swanson, *Uncharted's* audio lead, and I subscribe to. We both have backgrounds in post, but I think we both consider ourselves to be game sound people at this point. There tends to be an overarching philosophy from video game sound designers that every sound they create is this amazing moment. They are called a "sound designer", so they try to design these awesome "designed sounds" no matter what it is. But what makes a really great sound designer is understanding that it's not just about the sounds you are creating, in fact it's rarely ever about that, it's about the context. Like, what are you trying to achieve with the sound you are creating? And I think that's the perspective you are talking about right there - usually in film it's more about the director's vision, what they are trying to achieve, what they are trying to get the viewer to feel. So, context is what it is all about - if I was going to teach a video game sound designer anything, I'd sit them down and ask them "what are you trying to do, what is the scene, what are you trying to achieve, what is the focus?"

KY: Right. That is the first question. You shouldn't be doing anything until you have answered that. Related to this, there's an interesting dichotomy that I encountered on my most recent project. I've come to think of myself now as a "context junky" - I'm aware that I pretty much refuse to work if I don't have context because I know that what I'm making has no real value if I don't understand its role intimately from all angles. Like, there is no point in just ticking the box of creating a sound because the game needs a sound. However, on *Tearaway* I was involved from pretty much day one, and what caught me out was that at this early stage of development there is very little context because the game is way too nebulous and fragile. And that was so hard for me - I mean, I'm used to finishing stuff! I was depressed! For months!

PK: Right! And I think that's something we need to go through. I was on *The Last of Us* from pretty much the beginning - I think I missed out on the first couple of months - Bruce Swanson

had done a couple of animatics and a concept level piece for it, which were awesome, so he had helped sculpt out the original soundscape for the game. But once I got on board there was a lot of iteration, and I will admit there was a lot of “throw away” sound required at the beginning of the project. Even though we kind of knew what the game was they were trying to make, once they start making it, it inevitably changed, and stuff got tweaked and remade, so you’ve got to be able to adjust and roll with that. To be on a project early allows you to go through that process of trying different ideas and honing in on the ones that work best.

KY: The analogy that I like to make for this is that it’s like moss...

[silence]

KY: Moss is this weird, magic stuff - if you have a pitched roof, and you look out and see a piece of moss, it didn’t used to be there! It grows seemingly out of nothing. Your ideas are like little specs of dirt that find themselves on the roof of your mind - most of them are going to get washed away in the rains of the project. But some of that dirt is going to get stuck in a crevice and give way to an idea that will sprout into a fragile little colony of moss. Even if that moss withers and dies, it leaves behind some fertile soil for more moss to grow on. Eventually you have healthy little clump of moss! [laughs]

PK: That’s exactly right! I think the biggest example of that from *The Last of Us* was probably our melee sounds. There was a lot of discussion about those. There was a point where it just wasn’t working. We were trying to do something different with the gameplay - we were slowing down the player - you couldn’t jump, there weren’t auto-firing weapons, there weren’t waves of enemies coming at you, it was a more exploratory kind of gameplay experience. It was actually more about what the player couldn’t do than what they could do. So, there ended up being a lot of iteration on the hand-to-hand, melee combat for the game. They put some stuff in, then took it out, then changed it, then changed it again, then took it out and started over. For the sound, we tried everything from *Indiana Jones* style punches, all the way down to just the subtle kinda slap sounds you can get when you hit your knuckles against your other hand! We scaled so many times on what these encounters were going to sound like because the experience kept changing - cameras would change, the feeling would change, and the speed would change, stuff like this. But what it came down to was that we knew we wanted the hand-to-hand combat to be very visceral, intimate and very brutal. But it was almost like the less you heard, or the more realistic we made it, the better it was. We didn’t sweeten a lot of it - there are hardly any whooshes or whatever, most of it is just cloth and the sound of knuckles. So, we had to go down a lot of roads only to find that they didn’t go anywhere, until we found that the more reductive and simplistic we were then the worse it sounded. In the best possible way!

KY: In general, “reductive” is quite a good way to describe the mix in *The Last of Us*. Fundamentally, it manages to pull you in by being very subtle and giving the player space to appreciate that subtlety. Which really grounds the player in the world that you created - you could only be playing *The Last of Us* when you hear the game. It’s a game which often works without its visuals, and you can’t often say about a game, especially on a big AAA console title. I

think that's testament to how great both the audio and the audio technology is. It's really inspiring.

PK: Thank you! I'd just like to make one last point, which is that it's a testament to the vision that Neil and Bruce had as well. They were willing to let us try different things which we really felt that we could do but that there was no test for.

KY: Right - so if they hadn't come to you in the beginning and asked what you could do, or if you had been unable to deliver on that, the game would have ended up very different. It's because you were able to run with an idea, prototype it to a point where it felt it had legs that you ended up actually inspiring them and giving them confidence. It's only through that kind of collaboration and mutual inspiration that an experience as special as *The Last of Us* can come in to being.



JL: Whilst Phil rightfully gets a lot of credit for being the lead and an all round awesome dude, a significant contributing factor is that he never gives up. Sometimes we'd get in to arguments as a result of this, but in the end it's usually for the best. One of us will win, or we'll compromise and come up with something even better. Phil refused to be constrained by what was available in the Naughty Dog audio toolset at the beginning of the project. One thing I've noticed in some sound people is that they have been so cowed and put down that they default to Plan B as a starting point. They don't ask for things the game needs because they don't think they'll get it. If Phil had taken that approach there is no way this game would have ended up sounding the way it does. Phil pushed me to give him what he needed, and he had no fear in doing so. At Naughty Dog we have a culture where you are allowed to do that - we can walk in and talk to our co-presidents Evan or Christophe, we can talk to any of the other leads, we can talk to any

of the other people, we have an open door policy. Even beyond the collaboration, the whole “no fear” approach of attempting to do things different from the norm because the game demands it is absolutely critical. We did so many things we’d never done before, and we learnt so much just doing this one game. I would encourage anyone else in the industry that if they want to make games that do something different that they should not be afraid to push the bar, and never give up.

PK: We have an amazing team. I was really lucky to work with a great group of guys and gals who signed on for this really crazy and occasionally painful *[laughs]* mission to make something different. I’m really proud of what they were able to accomplish. Winning these awards and being noticed for it is humbling to say the least. But we’re not done! *[laughs]* We’re not done - we already have ideas! The bar has to be raised again! We get inspired by other people and what they do and try to achieve. I want to be inspired, and I want to inspire.

JL: It was so humbling going to GDC this year and meeting people from DICE, Rockstar, Guerrilla and Media Molecule - those guys are so awesome. There’s not enough collaboration in this industry across developers and their technology - we all work in our own bubbles and our own games, and we don’t always make the time to get out there and network and learn the techniques that other people are using, but there’s a lot of cool ideas out there and they need to be shared!

KY: Phil, Jonathan - thank you for sharing. And thank you for your time.

