



PS Fotodesign Petra Schramböhrmer

Peter Neubäcker

The Celemony mastermind and CEO talks about the development of Local Sound Synthesis, integration and what a stone sounds like.

PETER NEUBÄCKER IS A qualified guitar-maker and studied classical homeopathy and astrology before founding a publishing house specialising in astrological subjects in 1984. However, his main interest remains the relationship between musical and mathematical structures and since 1986 he's been working in algorithmic composition and the development of computer programs primarily to assist in his research.

He began developing the idea behind Melodyne in 1997 and in 2000 founded Munich-based Celemony Software GmbH with the firm's software developer,

Carsten Gehle, now the technical director of the company, and Professor Dr Hildegard Sourgens, who is responsible for business development.

Celemony's only product to date, Melodyne, was introduced in 2001 and analyses audio data and allows you to edit individual parameters like pitch, phrasing and timing with complete freedom. Describing it an auto tuning program is short changing Melodyne's enormous capabilities and it remains one of the few exhibition floor demos that still has them standing open-mouthed in the aisles.

Describe the basic functions and features of Melodyne.

You can use Melodyne either as an audio editor for single sound files or as a multitrack recording and arranging program. The main difference between Melodyne and other recording software is that Melodyne detects the musical content of an audio file or signal and is therefore able to break it down into its musical components: individual notes. This enables the user to work at all times on a musical level. Although we have excellent algorithms for time stretching and pitch shifting, we'd like to think that users will just forget all about them and concentrate upon the music, manipulating their audio material as effortlessly as if each note were a MIDI note.

Describe the profile of your users?

They have a wide range of profiles, but broadly we can distinguish between those that see Melodyne as a production tool and those that see it more as a creative one. Many producers, for example, now routinely instruct Melodyne to detect every vocal or solo instrument line — and they do this before they've even decided whether or not they intend to process it — simply because once the material is in the Melodyne environment, if they do decide something needs correcting — the pitch or timing of a note, for example — they'll be able to do it with a few mouse clicks.

Those who see it as a creative tool are using Melodyne to change their audio material in ways that were unthinkable before, such as by transposing a saxophone's fundamentals down two octaves and its formants two octaves up, so that it sounds like no instrument known to man or beast...

What do you see as Melodyne's principal competitor?

The way we see it, Melodyne doesn't really have any competitors at the moment. It's true, people sometimes compare it to pitch correction tools like Autotune or to time stretching tools, but the resemblance never extends to more than one facet of Melodyne, while the methodology and indeed the whole philosophy of the two programs is invariably different.

Other tools are made for the manipulation in pitch or time of a material that is primarily considered to be 'frozen'. With Melodyne, you don't even think about pitch shifting or time stretching, because pitch and time are free parameters for you right from the start.

Besides, whatever modifications you do make to the pitch or timing of your material will almost invariably sound better within Melodyne, because Melodyne is able to detect the musical content of the signal, recognising where one note ends and the next note begins. Where you change the pitch of a note, for example, the authentic phrasing will be retained and the transitions to neighboring notes refashioned in a way that is musically convincing. When notes are stretched, the starting and ending transients that play a critical role (scarcely less important than that of timbre) in allowing the ear to distinguish one instrument or voice from another are left unchanged and only the steady-state, central portion of the note (the character of which is not disturbed by time-stretching) is extended. In this way, both the transients and the timbre of the note are retained, the only change being the one you actually intend: the note value (i.e. its length) is extended. Here, as always, Melodyne understands your musical purpose, eliminating unwanted side effects and artefacts.



Melodyne runs under Mac OS X and Windows XP and is currently offered in two versions: the studio edition offers an unlimited number of virtual tracks, 32-bit/192kHz support and stereo editing, while the medium-sized cre8 version offers 8 tracks up to 24-bit and 48kHz. The one-track Melodyne uno is designed for the uncomplicated editing of melodies.

What is the concept underlying Melodyne?

Melodyne was not originally conceived as a musical application at all. I regard myself primarily as an inventor and researcher; I wanted to know what sound is, how it functions and what exactly happens when someone speaks or sings. So I was sitting with a friend a few years ago and we asked ourselves what it would be like if you could free sound from time. We formulated the question in symbolic terms — what does a stone sound like? — because a stone has a form but no time. And it was out of this playful — half philosophical, half mad — speculation that the idea of treating a sound as though it were a stone and making time stand still for it emerged, and that was how Local Sound Synthesis, which is the basis of Melodyne, was born.

What is Local Sound Synthesis?

Local Sound Synthesis was developed based on the vision that 'sound' may exist independently of pitch and time. The usual approach of time stretching is that you have an original sound that you manipulate in some way to make it longer or shorter. From the view of Local Sound Synthesis, an original sound is a kind of landscape where you can find different sounds at different 'time places', and you can move through it back and forth at an arbitrary speed. Equally, the pitch of the sound at any location is arbitrary. Thus pitch, time and timbre become independent parameters, and an original recorded sound can be considered as a consistent sequence of sound shapes that can be used freely.

Unlike Granular Synthesis, Local Sound Synthesis treats sound as a continuum. If you move through the original sound very slowly using Granular Synthesis, you can hear the discontinuity as it jumps from grain to grain, so instead of a smooth progression, the sound becomes 'grainy'. Using Local Sound Synthesis, you can move through the original sound at whatever speed you please, and the sound will change, not in discrete steps, but continuously at the pace you determine. An impressive demonstration is provided by Melodyne's Scrubber, which is used to move through the source at varying speeds and in either direction. If you are using a MIDI controller, you could even determine the pace and direction of the

movement using one wheel and the pitch using a second.

Local Sound Synthesis is processed in the Time Domain only (in contrast to the Frequency Domain, which would be achieved by Fourier analysis and synthesis.) Thus this synthesis is very effective and fast, and a lot of voices can be rendered without overburdening the CPU. It always uses the original sound file as its sound source. The target sound is always generated on the fly using the original sound as the source of the current local sound. Thus, in Melodyne, Local Sound Synthesis is always applied in real-time and is non-destructive.

How easy is it to integrate Melodyne into productions?

Originally, Melodyne was a standalone application, which at many stages of the production process made it somewhat awkward to access. Now Melodyne integrates elegantly into all typical working environments. With the MelodyneBridge, we support VST, DXi and AudioUnits as well as offering ReWire, MIDI In/Out, SMPTE and much more. Now all users should be able to access Melodyne from within whatever their favorite production environment might happen to be.



Given sufficient processing power, could Melodyne also process polyphonic material?

The ability of a computer program to identify correctly to which part a particular note belongs in a passage of polyphony, sadly, has nothing to do with processing power; the problem is that there is no workable starting point. However much processing power they may have had at their disposal, no one has managed yet to come up with an algorithm for reducing a passage of polyphony accurately to its component melodic parts. But this, naturally, is a subject for further reflection.

How well Melodyne is able to cope in general with polyphonic material depends upon the nature of that material. Drum loops and similar percussive material could be considered polyphonic, and here it works well because the emphasis is more on beat separation; if each drum has a different pitch, that is another thing the program can pick up on. The same is true of a rhythm guitar, which has a fast attack and decay — there the prevailing pitch can be detected in most cases and you can modify both the pitch and the timing of such material quite well, but with real polyphonic source material, pitch analysis will not turn up anything useful and the output may be quite random.

What do Celemony and Melodyne names mean?

The name Celemony is a contraction of CELEstial harMONY, because in many cases the fundamental concepts informing our research are derived from the mathematical concept of the harmony of the spheres.

The word Melodyne is a contraction of MELOdy and DYNamic, and it expresses the fact that the program transforms audio material that was (in a sense) 'frozen' into something plastic and malleable. Something, in other words, that can be transformed dynamically.

What can we expect from Celemony in the future?

Naturally we will continue to improve Local Sound Synthesis and Melodyne and respond to the wishes of our customers. However, at the same time we are also working on other ideas; morphing, for example, fascinates me. I've developed a very good morphing technique that goes beyond the transition from one sound to another — capable, in other words, not only of interpolating between the sound of a guitar and that of a saxophone but also of extrapolating. This Spring we are also starting a joint venture with the sampling specialists Ueberschall that should result in the creation of VST instruments with Melodyne technology. I'm not prepared to say any more at present. I wouldn't want to spoil all the suspense. ■