Liquid Channel FAQ
The following information is taken directly from the Liquid Channel web site found at:
http://www.ffliquid.com

Q: What kind of technology is used by Focusrite in The Liquid Channel™?

A: The technology involves Dynamic Convolution techniques and a new liquid pre-amp which together replicate vintage pre-amps and compressors.

Q: What is Dynamic Convolution?

A: The convolution process has been defined as: 'The term given to the mathematical technique for determining a system output, given an input signal and a system impulse response.' What that means is if you know what is coming in to your system, and you can control your system's impulse response, you can exactly define the system's output. In other words, you can replicate a compressor's sound, or even a mic-pre's sound if you add a suitable analogue mic-pre hardware circuit. Put another way, Focusrite have found a way to precisely replicate the way in which any classic compressor or mic-pre ever made affects sound.

Q: How is this different from modelling and other 'simulations' we've seen and heard before?

A: Modelling looks at the way a device works and then relies on the generation of code to try to emulate the typical way in which a device would respond, usually in a certain limited set of situations. Convolution, on the other hand, records data about the way a device behaves and then replicates that. To use a simple analogy, it's like the difference between sampling and synthesis; if you want a REAL violin sound triggered from your keyboard you sample a violin, if you want a modelled sound which recalls the real sound of a violin, synthesis will generate a similar violin-like waveform.

Q: Why have modelling devices never succeeded in exactly nailing the way a compressor or mic-pre responds?

A: The problem with a compressor is that it is a dynamic processor. That is, it is reacting to changes in input signal, and varies its response according to those changes. The problem with a mic-pre is similar – it is constantly interacting with whichever microphone is feeding signal to the pre, and it is the combination of pre and mic that characterises the sound. Dynamic convolution plus liquid hardware enable these phenomenally complex interactive relationships to be replicated.

Q: What does the impulse response/convolution process involve?

A: A while back, Focusrite set about driving a huge set of impulse responses into the best collection of vintage and modern compressors and EQs ever assembled. The impulse response device they used for this process is called, with good reason, 'The Replicator'. This mysterious black box outputs an impulse – a very narrow (time-wise) voltage spike of amplitude which contains an infinite number of frequencies. The impulse spike is sent to the device you wish to replicate. By measuring the output of the device itself, Focusrite's R and D team were then able to calculate exactly what the device has done to the spike, hence calculate and reproduce every aspect of the device that relates to frequency- and time-related parameters; frequency response, headroom, distortion, everything!

One impulse, of course, tells you how the processor will react to a particular input level, so it's necessary to sample the data of the device you wish to replicate firstly with a spike that drives the box into distortion, then with a fractionally lower amplitude spike, then lower again, all the way down to tiny spikes down in the noise floor. Once you have all this data recorded you can repeat any change in input level, i.e. replicate the response to all input source types. Then all you have to do is apply convolution for every parameter setting combination, and you have the genetic blueprint of the device ready to go.

Q: So The Liquid Channel™ can reproduce the sound of any compressor and any mic-pre ever made?

A: Yes, and more. Once you have The Liquid Channel™ you can mix and match the sound of your dream pre's and compressors into user memories to reproduce any combination you want.

[1]

Q: How on earth can one machine deal with the sheer weight of calculations required to produce ALL those responses for every group of parameter settings of ALL those mic-pre's AND compressors?

A: Good question. It took the world's fastest audio-implemented SHARC chip technology to be able to crunch the mind-boggling numbers. That, and a huge number of patient hours replicating the sound of the classic units from audio history.

Q: So everything is pre-programmed?

A: No, all the convolution programming has already been done for you, but of course the impulse responses have to process the audio in real time inside The Liquid Channel™.
Q: Does The Liquid Channel™ allow me to replicate my own choices of pre and compressor?

A: No, that's pretty specialised stuff and best left to the Focusrite R and D team, but the unit ships with 40 mic-pre's and 40 compressors ready to go in the box. You can make up your own combinations and store them in user memories, with or without EQ.

Q: Are there 'user memories' in addition to these?

A: Yes, the mic-pre and compressor 'building blocks' can be combined into 99 program memories. All parameters including mic-pre gain, EQ and compressor settings etc. are also stored within the program memory.

Q: What if I want the sound of a specific mic-pre or compressor, which isn't one of the chosen devices?

A: Focusrite plan to make the sound of further pre's and compressors available as downloads from a new dedicated Liquid Channel™ website – www.ffliquid.com. The free LiquidControl™ application software then allows you to load the sound of the pre's and compressors that you crave into The Liquid Channel™ via the USB port on the rear panel. You can also save program memories externally to your PC or Mac, and even edit the unit remotely via USB! All parameters will be editable on-screen remotely, even mic-pre gain settings for example, and can be transferred from session to session in e.g. a Pro Tools folder.

Q: What if I should overwrite my classic compressor by mistake?

A: No problem; you can always reload a back-up set from your PC or Mac, or download the original factory settings from www.ffliquid.com.

Q: How can convolution replicate the interaction of mic-pre and microphone?

A: On its own it can't. You need a separate analogue circuit to be present and to work with the convolution engine.

Q: So why do mic-pre's present such a problem?

A: Mic-pre's have always had to connect to the source microphone, but it's an interactive system that isn't 100% efficient. Mic amps have been designed since the 1920s to suit a wide variety of different types of mics – passive carbon dynamics, then coil-based designs, then valve amplifiers, large diaphragms, phantom powered condensers etc. Hence, different vintages and types of mic amp, will vary dramatically in terms of the way that their input has been designed. For example, the range of electronic/transformer front ends that have been used over the years exhibit a wide-ranging set of impedances, and this is why an analogue front end needs to be included. If a specific mic is not being loaded by the analogue circuit just as it was by an original vintage device, then the sound from that microphone will be different.

Q: So there's no real mic-pre standard?

A: Exactly. Take a transformer for instance. It has two coils of wire, the first coil generates a magnetic field, and this then passes into the second coil – which in itself is not a fixed transfer mechanism, there's a lot of variation in transformers. What impedance appears at the input of the pre is also a key factor – when you connect a mic it has an output impedance of its own. The two sides (mic and pre) react, and frequency-related level can vary wildly as a result. Capacitances also interact as both mic and transformer have capacitances that vary, so HF roll off may occur for example, or you may get an HF peak (the famous Focusrite 'airiness' typified by the ISA range for example). Older mics designed for broadcast applications often roll off at e.g. 12 kHz, since before 1970 few people cared about HF matters. (Designers used to just roll off at 12 kHz to filter out problems above this threshold.)

Q: So how do you design one mic-pre circuit that can reproduce all the variables within this wide range?

A: The only way to accommodate the full range of different designs is to allow huge flexibility in the resistance and capacitance parameters in a custom transformer designed specifically for that flexibility. Hence, The Liquid Channel™ physically changes analogue circuitry as well as using dynamic convolution technology to create mic-pre replicas.

Q: What about electronic or tube mic pres that do not include a transformer?

A: The Liquid Channel™'s transformer is auto-switched out when an electronic transformerless mic-pre is chosen by the user (this is indicated on the front panel). Focusrite has built in the variations required to exactly reproduce the vagaries of any electronic mic-pre. The capacitance and resistance are then varied in the circuit, and dynamic convolution technology is used to replicate the full range of electronic pre's. Tube replication is also covered 100% – this is taken care of by the dynamic convolution process. Whatever artefacts were present in a classic vintage tube piece are also present in The Liquid Channel™.

Q: So this is really a hybrid technology that allows total freedom over exact reproduction of the sound of analogue pres and compressors?

A: Yes, the sound of every opto, and every VCA compressor, every transformer-balanced, electronically balanced (including tube pre's) can be replicated exactly, because each device's every response has been precisely mapped.

Q: Why is it necessary to have additional circuitry for the mic-pre and not the compressor? Surely, if the convolution DSP is as thorough as you say, there should be no need for further processing.

A: As mentioned above, the interaction between the individual microphone and pre-amp is a key factor in the sound of the pre as a whole. (The ISA 428 and ISA 430 MK II have switchable impedance values that the user can implement to specifically tailor the character and response of the device for this very reason.) By including a 'Liquid' pre-amp circuit containing a flexible signal path (transformer or electronic) and variable impedance value, The Liquid Channel™ can match that of the classic mic-pre to ensure that the interaction with the microphone is identical. This issue isn't something that affects a compressor but the DSP processing required is nonetheless immense. The user's ability to affect the threshold and ratio of the compressor means that there are additional responses needed for the side-chain to account for the numerous variations in character (types of 'knee', presence of 'over-compression' etc.)

Q: Is the transformer the traditional Focusrite Lundahl™ transformer? Or another famous brand like a Jensen™?

A: No, it's a brand new custom precision-wound Focusrite 'Liquid' transformer, designed and built in the UK by Focusrite's R and D department to be extremely flexible; transparent or coloured as required.

Q: If I connect different mics to the plethora of mic-pres that The Liquid Channel™ offers, will each of the mics sound different?

A: Of course. The results will be the same as if you were connecting your collection of favourite mics to the various pres. Of course, if you're just modelling in software this is simply impossible- how can a particular mic. interact in its distinctive way in real time with a particular mic-pre if the mic-pre is in fact absent?!

Q: What about interaction between the mic-pre's and compressors?

A: The Liquid Channel™'s dynamic convolution DSP is separate for both the pre-amp and compressor. That is to say, the device acts exactly as the separate units would, with the same signal leaving the mic-pre and entering the compressor. So, interaction is identical to the hardware equivalent without the extra cost/size and weight/plugging in required, but with all the reliability of a first class digital audio device.

Q: What about EQ?

A: Focusrite's R and D department in England decided that a truly Liquid Channel™ strip should also include a flexible EQ. So they created a brand new digital EQ, modelled on the curves of the fabled ISA 110. This EQ is a single Focusrite British EQ design that is the perfect complement to the range of mic-pre's and compressors available.

Q: Will the EQ sound exactly like an ISA 110's EQ?

A: It's based around the 110 design, but since it doesn't use convolution technology it won't be exactly the same. If you want the original sound of a Focusrite EQ unit with the same classic analogue footprint as the Focusrites of the 1980s, buy an ISA 430 MKII or an ISA 220.

Q: Can I put the EQ in front of the compressor?

A: Yes you can. You can also drive the compressor from the parametric mid-range section to use The Liquid Channel™ as a de-esser – more liquid than liquid! The EQ is editable, programmable directly from the front panel, and includes high and low shelving and parametric mid-range across a huge range of frequencies.

Q: I understand the unit is a recording channel, i.e. mono. But what if I want to record in stereo – can I chain two units together?

A: Yes, all you need is a standard RCA (phono) cable to transmit the data between units. The pre's, compressors and EQ will then all function as a perfect stereo pair, even if you are operating from the LiquidControl™ software application. With two units linked together, The Liquid Channel™ also becomes ideal for stereo mix-down and mastering applications.

Q: Is there any way to run a super-short signal path from the mic-pre to the output?

A: Yes. By not selecting Comp or EQ in, the signal will pass through the analogue front-end, A/D and mic-pre section of the DSP, then straight through to the AES digital output (or through the D/A to the analogue output).

Q: What if I record a great vocal performance only to find out later that I drove the pre too hard and caused clipping?

A: The Liquid Channel™ includes a feature called Session Saver which will automatically prevent this if you enable it. The Liquid Channel™ notices that digital overs are in danger of occurring, and if it sees significant danger it turns down the analogue pre-amp gain – a 1 dB reduction for any level above 0 dBfs.

Q: How flexible is my record path?

A: Very. You can record in the following ways:

* Analogue to digital: (mic. connects to mic-pre via the balanced XLR connector, through the A/D converter, through the pre and compressor convolution processors, and then exits via the AES D/A which is included as standard.) The D/A can be used to monitor post-DSP as a super-low latency feed to bypass DAW delays if you wish.
* Digital to digital: The digital input can be re-routed into the front end, feeding into the mic-pre and/or compressor convolution engines as desired.
* Digital to analogue: as Digital to digital above, but the balanced XLR analogue outputs are used.
* Analogue to analogue: as Analogue to digital above, but the balanced XLR analogue outputs are used.

Q: What are the A/D and D/A specifications?

A: The format is the world professional standard AES-EBU, and all sample rates from 44.1 kHz – 192 kHz are supported, in and out, as standard. Hence you can digitally input a pre-recorded instrument or line level signal into the pre./ or direct to the compressor.

Q: What A/D encoders are featured?

A: The latest 192 kHz/120 dB spec AKM™ 5394s.

Q: I notice there's an extra parameter called 'Harmonics'. What does it do?

A: Warm is good, everybody tells us so. Of course The Liquid Channel™ will perfectly re-create all the classic analogue warmth that is present in vintage pieces, along with all the other sonic artefacts. But imagine if your favourite mic-pre was for some reason warmer than most other pre's of that type? (Since many vintage pieces were hand-built, there is often some component tolerance variation for example.) This parameter means it's actually possible to add in extra warmth to 'tune in' The Liquid Channel™ to your 'special' unit...

Q: What if I want to compress pre-A/D to get that "driven hard" sound?

A: There's no need; you can now optimize the level using the gain encoder, then add in the "high gain" warmth via the Harmonics pot. as described above. Certain classic pre's which benefit from being driven hard will also be replicated 'driven hard' at source.

Q: What is the latency of the unit?

A: Worst case scenario (analogue in, analogue out, all sections in circuit, better than 2 milliseconds. In other words, insignificant.

Q: Does the fact that responses are measured in increments mean that the ear hears differences between the replica and the original?

A: Not at all. Firstly, the response measurement accounts for non-linear behaviour of the vintage units by sending impulse trains – literally single impulses at decreasing levels – from peak excitation to approximate noise floor level, separated by time divisions to allow for system reset. This means that the dynamic nature of the devices being replicated is taken on board unlike with most other devices that just assume linearity for ease of design.
Secondly, these 'response filters' are then applied proportionally to every sample of audio. Hence, at 96 kHz, the signal is being processed with 96,000 different dynamic responses each second! It's true that this method isn't 100% continuous, but the density of data in such a thorough system is such that any tiny errors present are not in any way audible.

Q: What about any extra knobs on the vintage models other than those present on The Liquid Channel™? How can The Liquid Channel™ accurately replicate without them?

A: The only controllable parameters available on classic mic-pres/compressors are those featured on The Liquid Channel™ – threshold, ratio, attack, release, gain. (If anything, some models have less controls, with some vintage compressors just offering gain controls for example. Focusrite do also plan to have more than one replica for certain devices which may require them.) In addition, The Liquid Channel™ extends the user's control further by allowing an extra dial to change the percentage of the mic-pre's second, third and fifth order harmonic distortion, and hence the overall warmth, of the mic-pre. This extra control allows the user to account for any variation between individual vintage pre's of a particular type. The Liquid Channel™ also features a Setup menu to configure the unit's parameters as per the original unit, or to allow The Liquid Channel™'s parameters to be fully edited, even if this was not possible on the original unit.

Q: What components change on the analogue PCB when switching pre's?

A: The transformer and the matrix of relays, which switch resistors/capacitors.

Q: How is the transformer manipulated? Are there actual variations between the primaries and secondaries?

A: Transformer variation is actually partly taken care of by the resistor/caps variation. The transformer is huge and is big enough not to load down the circuit – hence the transformer (1 to 1 type) is utterly transparent but can be configured not to be if the original device to be replicated requires it to be coloured in the analogue domain. Primary/secondary variation specifically is largely taken care of by convolution.

Q: Do the replicas sound worse at 44.1 kHz than at 192 kHz?

A: No. Sample rate does not affect replication quality; we are always 32-bit floating, and e.g. ADC quality is far more affected by lower/higher sample rates. N.B. We do not sample rate convert because we don't need to as we have all different sample rates already stored on the SHARC chip. We have low speed (44.1 kHz) and high speed replica data High is actually only at 96 kHz – nothing higher is required, since 96 kHz is high enough when applied to replication, and any further improvements are way out of even the potentially significant psycho-acoustic realm. This however is not true in the world of ADCs for example where 192 kHz vs. 96 kHz is an audible difference.

Q: Why is there no tube? If the transformer is required in the analogue circuit how come a tube is not also required?

A: Convolution handles and effectively replicates all tube characteristics. However, there are additional benefits to using transformers over and above their warmth; better CMRR and the transformer's direct impact on the connected mic. directly for example. This latter point is why we need a transformer in circuit to replicate the mic-pre- the interactivity with the mic. is key for a transformer in a way in which it is not key for a tube.

Q: Are there any audible (distortion) differences between applying an "HG" replica versus using a regular level replica and then adding 2nd/3rd/5th order distortion via the dedicated Harmonics encoder afterwards?

A: Yes. At full gain on an old mic-pre. you may have 60 dB of gain at 1 kHz but only 40 dB at 10kHz. Distortion is affected in a similar way. Third order harmonic distortion is reduced at low gain, so this may be present when using an HG replica, but not present if you just add second order distortion after the fact. This is one reason why we include HG settings. There are also differences re: different loadings on the transformer occurring when driving high gain in at the front end, hence mic interaction changes occur. For example, HF roll-off etc may change.

Q: Why do we only use one set of impulses when creating the replicas? Surely we need to replicate all combinations of threshold, ratio parameters etc?

A: Convolution, using a single set of impulses, exactly replicates the sound of the signal path at all frequencies/levels. However, the FF Liquid process is actually more complicated than this. After replicating using convolution, we then measure the compressor curve at different ratio/threshold positions. Then we measure the attack and release characteristics, as well as the RMS vs. peak detection of the side-chain signal in order to see if it discriminating more towards peak or RMS. (If you put e.g. a drum kit thru a peak-detecting compressor such as a Focusrite) the compressor side-chain will follow the curve of the signal that's coming in and compress in a manner that follows that curve. An RMS compressor will compress the signal against the average level and ignore the peaks. Hence a drum kit is smoothly compressed by a Focusrite piece, but an RMS unit will result in attacky/toppy compression with many transients which have more energy/are less smoothly compressed.

We then glue all this info together to get a compressor that acts so that when the signal enters it gets rectified, goes into peak/RMS depending on what the original vintage unit requires, then enters the curve circuit to recreate the original side-chain, then applies that to the convolution data. This cannot be done at the impulse stage since the impulse maths will then look at the amplifier rather than the compressor.

Q: Can you route the digital input into the pre?

A: Yes, you can route the digital input to the pre. In fact it is always routed there but you can select the 'no pre replica' to bypass the pre and route directly to the compressor.