



CDP FUNCTION GROUPS

The CDP suite consists of a number of Function Groups, all of which are listed here. Most of these are the names of programs, each containing a number of functions; however, newer single-function programs have also been allocated to a group, for reference purposes. A few Function Groups consist solely of single-function programs.

FUNCTION GROUP DESCRIPTIONS

BLUR – blur spectral data to diffuse clarity

Some of the BLUR functions operate across a number of time-windows, for example time-averaging the spectrum (**BLUR**), shuffling the windows (**SHUFFLE**), or moving along the windows (**DRUNK** [c.f. **EXTEND DRUNK**] and **WEAVE** [c.f. **EXTEND ZIGZAG**]).

However, others operate on the channel amplitudes (and have more in common with the **HILITE** group): **SUPPRESS** is the opposite of **HILITE TRACE** and **SCATTER** likewise thins the spectrum, but randomly; **AVRG** blurs by averaging the spectral energy across a number of adjacent windows, **SPREAD** reduces the prominence of spectral peaks and **NOISE** raises the gain in low-level (noisy) channels. **CHORUS** randomises amplitudes and frequencies to produce a spectral chorusing effect, though it can also introduce noise.

COMBINE – combine frequency analysis data from two or more files

This group combines analysis data from two spectral files, for example by interleaving them or producing the sum, difference or maximum value of the spectra. Several of these functions do not have as radical an effect as might be imagined (**MEAN** is an exception), but others are more powerful, such as **CROSS**, which goes further than **MORPH VOCODE** in *replacing* one spectral envelope with another (and there are also similar functions: **SPECROSS**, **SPECTWIN** and **SPECSPHINX**).

MAKE (and **MAKE2**) are key tools for combining **PITCH DATA** files with **FORMANT** files to produce spectra that can be re-synthesized back to sound.

DISTORT – distortion based on wavesets (pseudo-wavecycles)

The **DISTORT** functions are variants of the grain concept, but the segments are pseudo-wavecycles called **wavesets**, found in zero-crossings. The wavesets are processed by repetition, omission, replacing waveforms, adding "harmonics", and so on. Cycles can be processed singly or in groups. Many processes do produce distortion with small groups of cycles, because of rapid and uneven changes of waveform, but larger groups produce smoother and often unpredictable results.

ENVEL – alterations to the amplitude envelope

The 'envelope' is the overall loudness contour of a sound. You can extract an envelope (either as a binary or **breakpoint** text file) and edit or transform it in over 30 different ways. Envelope functions can alter the dynamic shape, creating fades and swells, warp it in many ways, or totally re-draw it, creating new envelopes. Most importantly, you can impose or replace the envelope of one sound with that of another.

Most envelope functions use a parameter called window-size. This is the time-period used to measure the amplitude successively and it determines how accurately the envelope is defined: a small window-size gives a precisely defined envelope; with a larger one, amplitude values are averaged over a longer time.

ENVNU – specialised operations on envelopes

A small group which extends the scope of the ENVEL functions. **PEAKCHOP 2** produces an envelope based around the peaks in the soundfile; when this replaces that of the original source, the result is a gated sound with silence-separated events, particularly relevant to the RETIME programs.

EXTEND – processes which extend and/or segment a sound in time

The EXTEND group is concerned with repetitions of various sorts, which prolong the sound. The iterations tend not to repeat the whole file, but only segments of it. Typical of this approach are:

- **LOOP**, which not only loops a segment but can move the looping start-point through the file to create a form of time-stretching.
- **ZIGZAG**, which reads segments forwards and backwards, using an ever-changing start-point, and
- **DRUNK**, which takes a random walk through the file, extracting and splicing segments selected by a certain randomness of movement and segment size.

SEQUENCE and **SEQUENCE 2** are quite versatile sample-playing functions, which can play a sequence of sounds at given times, with optional transposition, loudness and duration. (*Soundshaper* has a MIDI to SEQUENCE 2 conversion utility which enables the program to render standard MIDI files.)

Newer standalone functions attached to the group include **CERACU** and **SHIFTER**, which can execute complex polyrhythms.

FILTER – techniques to filter or focus on frequency bands

Filtering changes the tone-colour of a sound, attenuating some of the harmonics and emphasizing others. The filter's power is well known through the classic "subtractive synthesis" model in which a synthesised waveform, rich in harmonics, is shaped by a time-varying filter. The classic filter types are:

- **Low-pass**: cut off top frequencies
- **High-pass**: cut off bottom frequencies
- **Band-pass**: a mixture of Low and High which preserves a frequency band; and hence **Notch**, in which a band is rejected
- **All-pass**: all frequencies are let through but with changed phase relationships

CDP's filter functions include all the classic types plus particularly effective Filter Banks (c.f. a Graphic EQ), which can be used to "tune" the sound by strongly emphasising certain frequencies over others. There are also spectrally based filters (see **HILITE**).

FOCUS – focus or sustain spectral amplitude data

The FOCUS functions complement those in HILITE and BLUR by emphasising certain aspects of the spectral envelope, or operating across a number of time-windows. The time-based functions freeze, step-freeze or hold the sound over a period, while **SUPERACCU** (replacing the older ACCU) sustains energy in each spectral band (a little like a piano pedal). **EXAG** exaggerates the spectral envelope and **FOCUS** similarly emphasises the spectral peaks by filtering around them. **FOLD** radically alters the brightness of a sound by octave-transposing frequencies into a more concentrated range.

FORMANTS – extract and impose the spectral envelope (formants)

Formant files extract the time-varying spectral envelope. This can then be imposed on a different sound or combined with a pitch file to create a new sound. The group includes a spectral version of the classic Vocoder, in which a bank of filters controlled by the characteristics of one sound is imposed onto a second sound.

GRAIN – granulate and manipulate grains

The main granular program is **MODIFY BRASSAGE** and its variants **MODIFY SAUSAGE**, **WRAPPAGE** and **NEWTEX**. The last two of these are included in the GRAIN group: **WRAPPAGE** is a multi-channel version of BRASSAGE, while **NEWTEX** generates a time-varying texture of transposable segments cut from one or more sources.

The other Grain functions manipulate the grains in 'grainy' sounds, i.e. sounds with silences separating the segments/grains. **BRASSAGE** Mode 5 can create such sounds, as can **ENVNU PEAKCHOP** combined with **ENVEL REPLACE**. Grains can be treated in many ways, such as being duplicated, omitted, repitched, repositioned, shuffled and reversed in order.

HILITE – emphasise spectral amplitude, including spectral filtering

The spectral filtering functions are arguably the most powerful of the HILITE group: **FILTER** implements a number of classic filter types with the precision only possible in the frequency domain and **GREQ** (graphic EQ) is a versatile (though visually non-graphic) filter bank; **TRACE** retains the N most prominent bands in each time-window and **VOWELS** successfully imposes the formant characteristics of vowel-sounds onto any sound.

HOUSEKEEP – basic soundfile housekeeping utilities

Some HOUSEKEEP functions are in fact sound utilities, handling basic channel functions (**CHANS** and **CHANPHASE**), gating functions (**GATE** and **EXTRACT**), or tidying-up facilities (**TOPANTAIL2**, **EXTRACT 3 and 4**, **DEGLITCH** and **ENDCLICKS**). **RESPEC** modifies the soundfile header directly.

The rest of the group consists of file-handling utilities – **BAKUP** concatenates files for archival; **BUNDLE** creates a text list of files and **SORT** sorts such a list; **BATCHEXPAND** expands a batchfile to take more inputs; **COPY** makes multiple copies of a soundfile.

MODIFY – Wide variety of effects, including brassage, transposition, loudness, panning, delay and ring-modulation

The **MODIFY** group has most of the fundamental time-based functions such as **SPEED** (transposition), **LOUDNESS**, **SPACE** (panning) and **REVECHO** (delay and echo) as well as more radical traditional techniques such as Reverse, Ring-modulation and Shredding (in **MODIFY RADICAL**).

BRASSAGE implements granular reconstruction (also known as "granular synthesis"), typically creating thickly textured sounds out of overlaid segments (grains), and CDP's program **GRAINMILL** provides a standalone graphical interface for this. The technique is closely related to the separate group of **GRAIN** functions.

MORPH – Create smooth transitions between sounds, using amplitude and frequency analysis data

Morphing interpolates between the spectra of two frequency-analysis files, creating a gradual transition between them. The Release 7 functions **NEWMORPH** and **NEWMORPH 2** are particularly effective, including tuning options and morphing between the frequencies of spectral peaks. Another powerful function is **GLIDE**, which makes a frequency transition between two single spectra over any time; a spectral envelope from a suitable sound can then be imposed on the result.

MULTICHANNEL – multi-channel diffusion and processing

The functions in the MULTI-CHANNEL group are all standalone programs which have been brought together purely for reference purposes. Several are multi-channel equivalents of earlier mono/stereo functions.

A second sub-group manipulate channels and sounds in multi-channel space, involving spatial rotation and panning, re-ordering channels, distributing sounds across a spatial **PANORAMA**, or segmenting and dispersing the fragments spatially (**FRACTURE**). Mention should also be made here of the **MULTI-CHANNEL TOOLKIT**, which similarly includes functions for channel re-ordering and rotation.

The third sub-group is concerned with multi-channel mixing: the multi-channel mixing program **NEWMIX MULTICHAN** has its own multi-channel mixfile format and a support program **MULTIMIX** to help with creating such files.

TANGENT and TRANSIT

Multi-mix files are also the output of the **TANGENT** and **TRANSIT** groups of functions, which move a sound or a sequence of sounds along a tangent to (Tangent) or across (Transit) a circle of loudspeakers.

MULTI-CHANNEL TOOLKIT – multi-channel file handling functions

The Multi-Channel Toolkit is a group of versatile tools for handling multi-channel files, including a special emphasis on ambisonics and the **WAVE_EX** file format. The tools include functions to extract channels, change the channel order, mix and interleave multi-channel files, play them, obtain information about them, plus **COPYSFX** (CDP's "Swiss Army Knife") which can copy/convert between a very wide range of soundfile formats.

ONEFORM – operations on single formants

The ONEFORM group complements the **FORMANTS** functions, which extract and employ time-varying spectral envelope (formant) files. **GET** extracts a single spectral envelope from a formant file at a specific time. **PUT** and **COMBINE** produce sounds of static amplitude (which may still vary in frequency content), the first by imposing the single formant on a spectral file, the second by combining it with a pitchfile (c.f. **COMBINE MAKE**).

PITCH (&Harmony/Freq) – operations based on partials

The PITCH spectral group consists of functions which directly alter frequency content and includes two transposition functions (**TRANSP** and **OCTMOVE**). **CHORD(F)** does not produce a conventional chord, but rather a timbre enhanced by superimposing transposed spectra. The most powerful functions in the group are probably **PICK** and **TUNE**, which tune a spectrum onto the specified partial sets; **TUNEVARY** can do this in a time-varying manner.

PITCHINFO – information about partials or extracted pitch traces

PCHINFO functions are for binary pitch files (**.frq**), extracted from spectral files by **REPITCH GETPITCH**. They either report to the terminal or produce textfile reports. Basic file header information is provided by **SNDINFO PROPS**. **INFO** gives the maximum, minimum and mean pitch in the pitch file and the total pitch range. **ZEROS** reports whether the file contains unpitched windows. You can also hear the pitch the file has found, or view it (or transposition data **.trn**) as a pseudo-soundfile.

PSOW – manipulate pitch-synchronous grains (FOFs)

PSOW is a set of experimental grain processes for vocal sounds. It attempts to find and then manipulate FOFs, which are formant waveforms (grains), used to synthesise the singing voice. (The term comes from IRCAM's CHANT program.) PSOW processes require a pitch-trace taken from a spectral analysis file. The programs allow you to alter formants independently of pitch, often with unexpected results.

PVOC – FFT analysis and resynthesis

The public-domain PHASE VOCODER (**PVOC**) is used by CDP to convert between the waveform representation of sound in the time-domain (**.wav** / **.aif** files) and time-varying spectral analysis files (**.ana**) in the frequency domain – and back again.

While in the frequency domain, the different "bins", or frequency bands, are manipulated by Spectral and Pitch processes to create effects that could not be achieved in the time-domain. CDP has been especially innovative in this regard.

REPITCH – Extract and alter pitch-related frequency analysis data

Apart from **REPITCH TRANSPOSE(F)**, which transposes pitch without altering time, the REPITCH programs are concerned with extracting pitch from frequency analysis files and processing the resultant pitch files, which can be in binary or breakpoint-text format.

To re-synthesise sound, **COMBINE MAKE** combines the pitch file with a formant file (.for), which represents the time-varying spectral envelope. Typically, the formant file might come from another sound. When processing pitch, it's important to realise that the dynamic profile will be that of the formant file and that changes to pitch will only be convincing to the extent that the frequencies of the altered pitches are also present (i.e. strong or weak) in the formant file.

REVERB – functions to reverberate soundfiles

The REVERB group's two main programs are **REVERB** and **ROOMVERB**, comprehensive reverberation functions incorporating such features as reverberation time, dry/wet mix, absorption and early reflections. The programs use built-in sets for small, medium and large rooms, or they can use a datafile created by the separate **ROOMRESP** program. There is also optional filtering and pre-delay (giving the effect of a bigger room). ROOMRESP's data can also be used by **TAPDELAY**, a tapped delay line.

Also in the group is **FASTCONV**, an FFT-based convolution program, which can simulate a wide range of reverberation types using a sampled impulse response of a building or other responsive space. (Many suitable impulse-response files are available on the Internet.) More experimentally, the impulse-response input can be any soundfile. Convolution can also implement an FIR linear-phase filter.

SFEDIT – basic soundfile editing

One of the largest function groups, SFEDIT has the fundamental cutting and splicing functions found in visual editors, but CDP also extends these operations into relatively uncharted areas.

For example, **MASKS** covers portions of the sound with silence; applying different masks to copies of the sound effectively partitions it in different ways. Similarly, **ISOLATE**, and **PARTITION** partition sound into separate output files, each retaining the original segment times; like the MASK outputs, these files can be further processed independently prior to a possible remix. **RANDCHUNKS** and **RANDCUTS** also output multiple files, cut at random within a specified range. **TWIXT** and **SPHINX** switch between a number of sounds seen as being streamed simultaneously. **PACKET** outputs short enveloped sound grains that can be used with TEXTURE, BRASSAGE or PANORAMA.

JOIN is a standard splicing program, but there are also simple sequencing functions **JOINSEQ** and **JOINDYN** (c.f. **SEQUENCE / 2** in the **EXTEND** group.)

SNDINFO – basic soundfile information

SNDINFO functions either report to the terminal or produce textfile reports. The functions cover such aspects as soundfile properties, length, maximum level, lowest level, loudest channel and channel differences, zero-crossings, and start-times. (**RMSLEVEL**, in the Multi-Channel Toolkit, reports the overall power level of each channel.) **PEAKFIND** lists the times of loudness peaks found and is an important data file for other programs such as **RETIME**. **PRTSND** converts soundfile data (time-varying amplitude) to text. Some functions require two or more infiles for the purpose of comparing soundfile levels, lengths, and summing lengths.

Other reporting functions specific to particular Function Groups are found in that group. These include **cycle count** for **DISTORT**, **grain count** for **GRAIN**, **find pan** for **MODIFY SPACE** and **stretch factor** for **STRETCH TIME**.

UNITS is a general-purpose musical calculator, offering a number of useful unit conversions for pitch and frequency, transposition and interval, speed and time, and loudness as gain and dB. **TIMESMP** converts time in a soundfile to sample-count, for sample-accuracy editing.

SPEC – gain and editing utilities in the spectral domain

A collection of spectral utilities. **CUT** and **GAIN** are spectral counterparts of time-domain functions. **BARE** zeroes channels which are not harmonics, according to a pitchfile used for comparison. The reference manual lists some possible uses.

GRAB retains a single analysis time-window at the specified time. This is required by **MORPH GLIDE** (only). **MAGNIFY** expands a single window over a specified time, not unlike the **FOCUS** functions **HOLD** and **FREEZE**.

CLEAN has been superseded by **SPECNU CLEAN**; both require a sample of noise for comparison, so that the noise frequencies may be spectrally removed.

SPECFNU – Modify spectral shape in relation to formant peaks

SPECFNU is a single program of 23 modes, modifying the spectral shape while aiming to retain the existing formants. Many of the functions are equivalents of existing ones, but with the addition of preserving formants. Other modes manipulate the formants themselves.

SPECINFO – information about spectral data

SPECINFO functions either report to the terminal or produce textfile reports. Basic information includes the number of time-windows, and conversions between channel number and frequency. File header information is provided by **SNDINFO PROPS**. Various functions report on the time-varying spectral peaks in the sound or extract the amplitudes of partials; **NEWMORPH 2** (documented in the MORPH function group) lists the frequencies of the most prominent spectral peaks. You can also view the changing level as a pseudo-soundfile or print (part of) the analysis data to a textfile. (This produces huge files and is best used for tiny portions.)

SPECNU – various functions, mainly to clean up frequency analysis files

Three of the original functions in this group (**CLEAN**, **REMOVE** and **SUBTRACT**) are concerned with cleaning up files, though the latter two may have more general uses.

SLICE partitions the spectrum into frequency bands based on permutations of channel numbers and outputs these as separate files. This is an important resource for separate processing of the files, prior to a possible remix. For example, a mono source can be turned into pseudo-stereo by giving the slices different pan positions.

SLICE Mode 5 inverts the spectral frequencies around a given frequency value, which destroys the formant relationships in speech or singing, while retaining the rhythm.

Newer functions in the group include **RAND**, which randomises the order of spectral time-windows. (The Release 7 program **SPECGRIDS** is identical to **SLICE**, Mode 1.)

STRANGE – some unusual and unpredictable effects

The "strangeness" of this Function Group is perhaps subjective. **GLIS** attempts to create glissandi inside the changing spectral envelope, with a limited degree of success (**SUPERACCU** does it better). **WAVER** oscillates between a harmonic and inharmonic state. Its detuning can be subtle or extreme; if minimal, it can be used to give a chorus effect. **INVERT** inverts the spectral envelope in each time-window, transferring the energy in the lowest frequency bands to the highest ones and vice versa; the result is typically a much brighter timbre.

More useful is **SHIFT**, which performs a linear frequency shift of all or part of the spectrum, turning harmonic sounds into inharmonic ones. An upwards shift compresses the musical intervals between the partials and a downwards shift stretches them apart, not unlike **STRETCH SPECTRUM**.

STRETCH – stretch time or frequency data

A mini-group, consisting of two processing functions:

- **STRETCH SPECTRUM** warps the spectrum by stretching or compressing its frequency components, so that, for example, a harmonic sound may become inharmonic. This function is closely related to **STRANGE SHIFT**
- **STRETCH TIME** stretches or compresses the duration of the sound without changing the pitch. The degree of stretching or compression is itself time-variable.

Attached to the group is the utility program **STRETCHA**, which calculates the *stretch factor* needed to convert one duration to another, especially in relation to a given number of beats and tempo.

SUBMIX – mixing and related processes

The centrepiece of the SUBMIX group is **MIX**, which mixes the sounds listed in a text *mixfile*, the lines consisting of: **soundfile name, start-time, no_of_channels, ch1_level, ch1_pan [ch2_level, ch2_pan]**.

There are many support functions which manipulate the data in mixfiles, such as **PAN, ATTENUATE, SHUFFLE, SYNC, SYNCATTACK, SPACEWARP** and **TIMEWARP**.

Other mixing functions are also available, which do not use mixfiles. These include:

- **MERGE** – a quick mix of 2 soundfiles
- **MERGEMANY** – a quick mix of several soundfiles
- **BALANCE** – mix 2 soundfiles, using a balance function
- **CROSSFADE** – quick crossfade between 2 soundfiles
- **FADERS** – mix several mono or stereo files using a time-changing balance function

The two **INBETWEEN** functions produce multiple outputs, involving a transition in which the balance shifts from predominantly one sound to the other. Finally, an important workhorse – **INTERLEAVE** – combines mono signals to create a stereo or multi-channel file. (A similar function, **INTERLX**, is found in the Multi-Channel Toolkit.)

In general, CDP Mix functions are best for mixing small sounds to create larger units. Although larger mixes are perfectly possible, a visual editor is more intuitive and tends to have better controls.

SYNTH – functions to synthesise sounds

Only a small number of synthesis functions are provided within CDP, which is predominantly a sound-processing system. The older functions create basic waveforms, noise, silence, a chord for test purposes, a tone of variable pitch within a frequency band (**SPECTRA**), and a click-track.

The **NEWSYNTH** program is in a different league. Modes 1 and 2 (akin to CSOUND's ADSYN) generate a waveform with a time-varying spectrum and any number of harmonic or inharmonic partials. Mode 2 allows each wave-cycle envelope to be shaped, to a degree. In Mode 3, a multi-channel synth, pulsed selections of the spectrum are sent to a multi-channel output. Each output event is a random selection of partials from the overall spectrum and subject to octave-transposition over a specified pitch range.

Mention should also be made of undocumented functions developed for sonification and distributed with Release 7. (These will be documented in due course.) Sound can be created from time-series text data, or the data can be assumed to be a pitch-trace. The **SPECTRUM** programs can generate a time-varying spectrum from "frq-amp" pairs, and there are supporting spectral analysis functions which can generate such text data from input sound.

SYSUTILS – system utilities

A group of standalone support programs, including functions to play and record soundfiles (**PAPLAY**, **RECSF**) and play spectral analysis files (**PVPLAY**), list available audio devices (**LISTAUDEVS**) and (sound)files in a folder (**DIRSF**), and create a shortcut to a file (**ALIAS**, PC only). There is also a MAC utility (**CDPCONV**) to convert spectral files from PPC to Intel format. **COPYSFX**, which converts between soundfile formats, is also included in this group (as well as in the **Multi-Channel Toolkit**).

COLUMNS is a major data-messaging program which is implemented in *Soundloom's* Table Editor and *Soundshaper's* Data Editor. Around 100 functions alter or generate data textfiles, for possible use with CDP programs. The **COLUMNS** functions are not concerned with the appropriateness of the resulting data for any particular purpose: the user must ensure that ranges and values are suitable for the intended parameter.

TEXTURE – create textures of sounds

The TEXTURE suite is based on algorithmic repetition: the processes repeat and transpose the input sound or sounds in various ways to create a texture of events ('**notes**'). The 'note events' are whole sounds, which last for a specified time range, or may optionally be played to the end. However, you can easily edit a sound source to provide segmented events for TEXTURE: **ISOLATE**, **PACKET** and **PARTITION** are three processes which can produce suitable small segments.

Note events can be treated as simple repetitions, or repeated in groups, or with a timed rhythm, or as transposed ornaments or fully-defined motifs (timed and transposed). Key parameter values such as pitch, duration, gain (level), spatialisation and sound number are generally chosen randomly from within a specified time-variable range (which can also be a fixed value).

In each of the eight main functions, repetitions can be pitched at random within the defined pitch range, or restricted to a user-defined **pitch-set** or **harmonic-field** (which octave-transposes the pitches); the set/field can be time-varying. More than one input sound can optionally be used, and the range of sounds chosen can be time-varied.

It is not easy to grasp the differences between the Texture programs, and the suite is best explored at first using a simple note sample or similar short sound. CDP supplies two sets of tutorial examples for this purpose.