

A Granular Toolkit for Cycling74's Max/MSP

Nathan Wolek
School of Music, Northwestern University
711 Elgin Rd., Evanston, IL 60208-1200, USA
Email: n-wolek@northwestern.edu

ABSTRACT

Since the generation of granular textures was first automated using a computer (Roads 1978), granular synthesis has grown to become a popular tool for creating new sounds in electro-acoustic music. Many effects can be achieved through the granulation of sampled sound including time compression (Jones and Parks 1988) and expansion (Truax 1990) independent of pitch alterations. Such effects can be created using Cycling74's Max/MSP software, allowing them to be utilized in real-time. However, the software does not include sufficient externals to meet the need for maximum efficiency and flexibility in creating such effects. This paper details a toolkit for Max/MSP that the author has created with the aim of meeting these needs.

BACKGROUND

The basic concept behind granular synthesis is often traced back to the ideas presented by Dennis Gabor, who made the assertion that "it is our most elementary experience that sound has a time pattern as well as a frequency pattern (Gabor 1947, p. 591)." Gabor developed a model of sound that linked these to two attributes in fundamental particles that came to be known as "grains." These short bursts of sound are only a few milliseconds in length with variable frequency content and combine with other grains to form larger sounds.

Research into the musical applications of Gabor's granular sound model was limited until computers were first used to generate the hundreds of the short sounds per second that this model necessitates. Since the first computer implementation of

granular synthesis (Roads 1978), emphasis has shifted away from synthesizing grains based on mathematical models toward the sampling of grains from a pre-recorded sound source. These short bits of sound are then re-synthesized to create a desired effect that builds upon the original sound source. Effects such as time-contraction (Jones and Parks 1988) and time-expansion (Truax 1990) without any variation to the pitch of a source recording are two examples of resulting effects one might obtain from the granular method of sampling and re-synthesis. Depending on the control parameter settings, varied sounds can be created using a single source sample, ranging from rhythmic chains of pitched particles to dense clouds of broadband noise.

Software performing granular synthesis has taken different forms in the past, with each implementation having its own methods of control and available parameters. However, there has not been a platform for users to freely experiment with the concepts of this process and build customized structures that create unique effects using basic building blocks that work with individual grains. This type of tool for free exploration of granular processes seems like a logical method for working with the technique since it is in reality a concept of sound generation and not a single effect.

MAX/MSP

Cycling74's Max/MSP is "a high-level, graphical programming language (Cycling74 2001, p. 3)" that allows the user to easily put together signal networks and control structures for the production of audio, MIDI and multimedia projects, all of which can be controlled in real-time. The current version of the software runs on Apple PowerPC-based computers under versions 8.1 to 9.x of the Mac OS. It is an ideal environment to experiment with ideas because one can quickly connect (using virtual "patch cords") smaller blocks of code (known as "objects") together to form larger audio processing networks ("patches"). The environment offers a semi-open architecture in

which anyone capable of programming in C can use the software developers' kit to develop custom objects ("externals") that can be freely distributed and used by others working with the software. In addition to the development of externals using the C programming language, Max/MSP users can encapsulate patches so that they can be re-used in other patches (most often referred to as "abstractions"). Once in this form, abstractions can be called in the same manner as objects or externals.

This level of customization makes Max/MSP an ideal environment for exploring granular concepts of sound and developing methods of employing them in processing effects. An implementation of granular sampling using an earlier incarnation of Max that ran on the IRCAM Signal Processing Workstation has been previously reported (Lippe 1994). However, externals that deal specifically with sound according to granular concepts are just beginning to appear, including those detailed in this paper. These new externals which work specifically on the granular level should allow Max/MSP users to more easily create new and interesting patches built on the concepts of granular synthesis.

NEW GRANULAR EXTERNALS

This toolkit began with the idea that new externals were needed for the Max/MSP environment that worked at the granular level. These externals needed to manage the various control parameters, define each grain internally and ensure that updates only occurred when a new grain was generated. The process of development was somewhat unscripted, with revelations in the development of one leading to the concept for the next and techniques discovered in the development of another leading to revisions in previously completed externals. In the end, four externals were created to meet the different control needs that the various forms of granular synthesis require. The name of

each begins with the prefix of "grain" combined with the suffix that describes the control method each employs: 'grain.bang~', 'grain.pulse~', 'grain.stream~', and 'grain.phase~'.

The externals are setup to operate in a consistent manner with most control parameters the same between these four externals. All four of the externals require two arguments in order to be used. The first argument names the buffer~ holding the sound file to be sampled and the second names the buffer~ holding the window shape to be used. The use of a buffer for the sampling source allows these granular objects to sample in a manner similar to other objects included in the Max/MSP distribution with which most users should be familiar (such as 'groove~' or 'wave~'). The second buffer gives users the flexibility to either use the window files included in the toolkit or to design their own. Interpolation can be applied to either of the buffers being sampled in the form of an all-pass interpolation filter (Dattorro 1997). By default interpolation is turned on for sampling from the sound source buffer and off for sampling from the window buffer. These decisions on the default states were made based on the improvement of audio quality compared with the added processing load that interpolation on each buffer would have. Each can be turned on or off by simply using the respective 'sndInterp' and 'winInterp' messages with an argument of 1 or 0, denoting on or off respectively.

The leftmost inlet on each of the externals is dedicated to the control input in each of the four externals. Each has only one outlet that is used to output the granular signal. The remaining inlets are capable of receiving their data in the form of either Max messages or MSP signals. In either case, changes sent to these inlets are always deferred to the beginning of the next grain generated by the object. This ensures that their implementation is consistent with the concept of grains as being fundamental units with constant parameters throughout their short durations (Xenakis 1971). The second inlet on each is labeled "sound begin" and it determines the offset from the beginning of the sound buffer to begin sampling for the grain. It is specified in milliseconds and

values that would be beyond the boundaries of the buffer's length are simply wrapped around in a manner similar to the modulo operator. The last inlet is described as the "grain pitch multiplier" value and affects the sampling increment (how many samples are skipped for each sample in the audio output) for the sound buffer. It is labeled in terms of pitch because this is the aural effect that changes in sampling increment have upon output. A pitch multiplier of 1.0 causes the samples to sound at their normal recorded pitch. A multiplier of 2.0 effectively causes the samples to sound an octave higher, while a multiplier of 0.5 causes the sample to sound an octave lower. A floating-point value may be specified making more minute pitch transpositions possible. Both the 'grain.bang~' and 'grain.pulse~' require the additional specification of the grain length. This is necessary due to the methods with which their outputs are controlled as is described next.

The first external is named 'grain.bang~' and, as the name suggests, is controlled with the common Max control message of 'bang'. When it receives a 'bang' in its leftmost inlet, a single grain is output from the signal outlet. While it is in the process of outputting that grain, all bangs that might be received by the control inlet are ignored. This protects from the unwanted effect of interrupting a grain that is in the processes of sounding. Once the grain has completely sounded, the external will once again begin listening to the inlet for bangs. However, the use of bangs as a control signal means that grains can only occur at the interrupts of the Max scheduler and never in between. One consequence of this is that the interval between successive grains can never be smaller than the interrupt for the scheduler. For example, if the Max scheduler is set for an interval of 20 milliseconds, this means that grains can never be closer than 20 milliseconds apart or repeat at a frequency higher than 50 Hertz. This may not seem to be a problem upon first glance since grains are typically of a length on the order of 40 to

50 milliseconds, however there is an additional consequence of the granular frequency's dependence on the scheduler.

Dependence on the scheduler means that the granular frequency will always be sub-harmonically related to the scheduler's interrupt frequency. If grains can only be generated at the scheduler interrupts, it means that intervals in between successive grains will always be multiples of the scheduler interrupt. Sub-harmonics can be calculated by dividing the interrupt frequency by integers greater than one. Using the 20 millisecond example again, it would mean that the only granular frequencies possible would be 50, 33.3, 25, 20, 16.6, etc. Lowering the scheduler interval helps to alleviate the problem, since a smaller interrupt increases the possible frequency limit, but the sub-harmonic relation persists. An interrupt interval of one millisecond would allow for a theoretical maximum granular frequency of 1000 Hertz and sub-harmonics of 500, 333.3, 250, 200, 166.6, etc. Additionally, lowering the interval of the scheduler works against the computational efficiency that the lower frequency is meant to create. In order to create a better resolution of granular frequencies a control signal that operates at the sampling rate is needed and is exactly what the 'grain.pulse~' external is meant to provide.

The 'grain.pulse~' external monitors the signal connected to its leftmost inlet for zero-to-one transitions to which reacts by outputting a grain. This type of control signal (in which all samples have a value of zero except for a single sample that has a value of one) is more commonly referred to as an impulse. In the Max/MSP environment, this signal can be generated by the 'train~' object. The fact the control signal allows a resolution of grain placement equivalent to the sampling rate allows the sub-harmonics to be much closer together at the lower end of the frequency spectrum and therefore provides a much better resolution within the range of typical granular frequencies. The additional computational load of a control signal that operated at the audio rate within

Max/MSP was found to be not very significant. Like the manner in which bangs are handled in its counterpart 'grain.bang~', the 'grain.pulse~' external will ignore any pulse received while a grain is being produced. Additional parameters are again set through the other inlets, including the length of the grain to be produced upon receipt of an impulse.

With both of the previous externals, a single event triggered the output of one grain with a predetermined length. This one-to-one correlation between control and output meant that in order to create a continuous stream of grains, the control signal had to be repetitious. In order to eliminate this requirement, 'grain.stream~' external was created. This external requires that the user only specify the granular frequency in order to create a continuous stream of grains. With this external, grains are created consecutively at the given frequency and changes in frequency are only applied at the start of a new grain. Therefore changes in the granular frequency may not proceed in direct correspondence to the changes in the frequency that the user makes. This may seem undesirable at first, however it is a better fit with the concepts behind granular synthesis.

A single voice of continuous grains works very much like amplitude modulation applied to the sampled sound source, but it can be combined with other voices to fill out the sound and obtain a more even re-synthesis. However, the 'grain.stream~' object provides no manner for the synchronization of phases between multiple instances of itself. The multiple streams can move in and out of phase with one another in an indeterminate manner. However with the 'grain.phase~' external, the need for a control signal is re-introduced in order to more precisely maintain phase relationships. The control input takes the form of a signal modulating from zero to one which controls the phase of the window applied to the sampled source. This type of output is generated by the standard Max/MSP object called 'phasor~', multiple instances of which can maintain

constant phase relationships. It is important to note that the phase signal controls only the window and not the sampling increment for the sound buffer. The sampling increment is still computed in relation to the pitch multiplier specified and remains fixed for the duration of a single grain.

ADDITIONAL EXTERNALS

Two additional externals were developed in order to more easily manage multiple instances of the 'grain.pulse~' and 'grain.phase~' objects. When using several instances of either object, it is often desirable to have their control inputs be out of phase with one another. Most often they needed to be consistently even in their phase relationships and this meant individually setting the initial phase for multiple instances of either the 'train~' or 'phasor~' objects. This prompted the development of custom versions of these objects that have multiple outlets evenly out of phase with one another. These externals are called 'train.shift~' and 'phasor.shift~'.

They each require a single integer argument to specify the number of outlets the object should have. Once instantiated, the outlets will produce outputs similar to the objects from which they derive their names with each outlet being an additional 1-over-the-initial-argument cycle out of phase with the outlet to its left. For example, if the 'phasor.shift~' is given an argument of 2, it will have two outlets with the first having an initial phase of 0.0 and the second being 0.5; if given an argument of 3, three outlets will be produced with phases of 0.0, 0.33, and 0.66; and so on up to a maximum argument of 32. These greatly reduce clutter in patches where multiple voices of grains are needed. Although they have been developed for use with the granular externals described here, perhaps others will find additional uses for them.

EFFECT ABSTRACTIONS

In addition to the new externals that have been created for this toolkit, several abstractions were developed that provide ready-made granular effects built with the externals described above. These abstractions were built with a consistent style and interface in order to make incorporating them into other patches as easy as possible for users. Most of the effects have two versions: one that operates on static sound files and a second that processes audio live as it streams into the abstraction. These two parallel methods of granular sound sampling are analogous to those previously implemented by Barry Truax for the DMX-1000 (Truax 1988). All of the effects operate in real-time with very reasonable loads on the current generation of G3 and G4 PowerPC processors.

Changes are made to the sound input through the leftmost inlet. For those effects working with a static buffer this inlet handles messages used to load sound files into that buffer. A bang message will bring up the appropriate dialogue window so that the user can select the file from the hard drive, while a message containing the name of a file itself will load that file into the buffer. If you would like to use a 'buffer~' object that is within your own patch, a message with the text 'setBuffer' followed by the name associated with an instance of the 'buffer~' object will allow the abstraction to use it instead of the internal buffer.

For those abstractions with a streaming input buffer, the user simply needs to connect the signal that is to be processed to the leftmost input. Any signal may be used including one from a live audio input. The internal buffer of the abstraction records the input and simply loops itself while the processing is being performed. In addition to this, the streaming abstractions have a second inlet from the left that allows the user to switch buffer recording on and off as needed using values of 1 or 0. A short, 10-millisecond fade is applied when switching on or off in order to prevent clicks due to signal discontinuity that this switching may cause.

All of the toolkit's abstractions have a common prefix of 'gran' to help the user identify them. Those with a static buffer are given the suffix of 'file', while those with a continuous streaming buffer are given a suffix of 'live'. These naming conventions should help the user to identify common effects, while separating the different forms of buffer sampling provided for each processing effect.

Each effect abstraction has several control parameters that can be changed in real-time to affect the processing output. By double-clicking on an instance of a specific abstraction, a window will be opened that can be used as a control interface for the effect. It uses standard Max/MSP interface items, such as number boxes and drop down menus, to provide the user with basic abilities to affect the processing output. Within the interface, effort has been made to place similar parameters next to each other and in common locations between the different interfaces. This should help the user quickly become familiar with the interface elements in a single abstraction and translate that knowledge easily to the other abstractions.

In addition to the provided interface, the abstractions offer a fairly simple method for creating custom interfaces and control mechanisms. The rightmost inlet and outlet are used to manage changes in the numerous control parameters through messages exchanged with the parent patch. All changes to the control parameters are output from the outlet in a consistent message format with the name of the parameter followed by its new value. For example, if a message were output reading "pitch_mult 1.2" would mean that the pitch multiplier had been changed to a value of 1.2. Similar messages can also be sent to the corresponding inlet in order to change the parameter setting. Again in the previous example, if a message were sent reading "pitch_mult 1.0" to the rightmost inlet it would change the value for the pitch multiplier to 1.0. Using the 'route' and 'prepend' objects from the standard Max distribution, it is quite simple to create a custom interface that communicates with the abstraction. It is important to note that changes made

through the inlet are also echoed to the rightmost outlet and so the user must protect against creating infinite loops when creating custom interfaces.

Any additional outlet found on a given abstraction is used to output a granular signal for either playback or further processing. Most of the abstractions produce only a single, mono signal. All abstractions accept only mono signals as input, either in as single connection for the streaming buffer effects or a mono sound file for the static buffer effects. A few of the effect abstractions produce stereo output and one type produces eight individual streams at varied phases. The names of these abstractions and a brief description of their processing effect follows. Included in the descriptions are any reference to readings on the specific type of granular synthesis used and the processor load of the abstraction. Processor utilization was monitored using the DSP status window in Max/MSP running on an iBook G3 clocked at 500 MHz. A change in the control parameters can affect the amount of processing the abstraction uses and so the figures stated here are only approximates using mostly default settings.

gran.ASstream.file~ & *gran.ASstream.live~* (5-6% CPU utilization) – These are perhaps the simplest of the abstractions. Using their respective form of sampling input, both produce a single stream of grains based on the concept of asynchronous granular synthesis (Roads 1991). These abstractions use very little processor time but also have minimal control parameters affecting the sound.

gran.pitch.file~ & *gran.pitch.live~* (16-17% CPU utilization) – The pitch abstractions use eight voices of granular synthesis with evenly spaced phases to maintain an even re-synthesis of the sampled sound. Using a single pitch multiplier, the user can transpose the pitch height of an input, which in the case of live sources can be used to produce harmonies.

gran.chord.file~ & *gran.chord.live~* (16-18% CPU utilization) – Similar construction to the above "pitch" abstractions with the added ability to specify a list of

pitch multipliers. Individual grains are then transposed to selected intervals within the given list. The end result sounds like a chord based on the list of multipliers.

gran.groove.file~ (15-16% CPU utilization) – This abstraction is meant to mimic the behavior of the 'groove~' object which is included with Max/MSP. However, since this abstraction uses granular processing to loop the sound file, the speed and pitch can be manipulated independently. This effect does not have a streaming buffer version because its time compression and expansion capabilities require that the input sample be static. It uses the same phase-related grain streams as previously discussed in order to create an even re-synthesis.

gran.space2.file~ & *gran.space2.live~* (19-20% CPU utilization) – Based on the "pitch" effect, this abstraction places the individual streams at specific panning locations across the stereo spectrum. These streams are numbered one to eight from left to right and can be turned on or off in order to create different spatial effects.

gran.space8.file~ & *gran.space8.live~* (18-19% CPU utilization) – The same concept as the above "space2" effect with the difference that the individual streams are output from different outlets. This gives the user the ability to process the streams separately or distribute them within a multi-channel audio output system.

gran.cloud.file~ & *gran.cloud.live~* (19-20% CPU utilization) – These use the model of asynchronous granular synthesis with an element of random deviation from basic control parameters to produce thick "cloud" textures (Roads 1991). These abstractions provides more parameters than the others described in this paper and are far more processor intensive at higher granular frequencies.

ABSTRACTION CONTROL PARAMETERS

Many of the effect abstractions have common names for parameters that are the same or operate in a highly similar fashion. Rather than explain them in the context of

the individual abstractions, the following section provides a common reference for the different parameter names and a brief explanation of what is controlled by each. The shortened names are the labels used in the control interface that is displayed when the user double-clicks on a given abstraction instance. These are also the names used by the inlet and outlet that handle messages for custom control of the parameters.

samp_offset – Sample offset. In abstractions that use a static buffer, this parameter is used to define the number of milliseconds from the beginning of the buffer to begin sampling from for grain production. It is not available in the streaming buffer abstractions because the sampling position always follows the recording position within the buffer.

buf_size – Buffer size. Used by streaming buffer abstractions in order to define the length of the buffer used for sampling the input signal in number of milliseconds.

speed – Only available in the "gran.groove.file~" abstraction. Used to control the speed of playback for the sampled sound upon re-synthesis. For example, a value of 1.0 equivalent is to normal playback speed, 0.5 equals half speed, and 2.0 equals double speed.

offset_bw – Offset bandwidth. Defines the number of milliseconds that the sampling position may randomly deviate from itself. The bandwidth represents the total possible deviation in both the positive and negative direction combined. The maximum deviation possible in either single direction is one half of the value specified.

freq – Frequency. For some abstractions, specifies the number of grains per seconds produced in Hertz. This is equivalent to the inverse of the time between grain onsets expressed in seconds. It is preferred over a simple length parameter when the length is being varied in some manner that results in it not being equivalent to the time of the next grain onset, as is the case with asynchronous granular synthesis.

freq_bw – Frequency bandwidth. Defines as a percentage the maximum amount of random deviation allowed for the frequency parameter. The percentage is used to derive the actual deviation in Hertz using the frequency parameter value. The use of a percentage ensures that the bandwidth is perceived consistently across the frequency spectrum. Again the bandwidth is the total deviation possible with both the positive and negative directions combined.

length – The length of grains to be produced in milliseconds. Lengths for granular processing are typically limited to a range of 10 to 100 milliseconds and are most often in the 40 to 50 millisecond range.

length_max – Maximum length. Expressed as a percentage of the inverse of the frequency, this parameter allows the grains to overlap one another producing a thicker texture. At 100 percent, no overlap should occur because the length would equal the inverse frequency.

length_min – Minimum length. Another percentage of the inverse frequency that allows the grains to be shortened, creating short delays between grains. The actual length of individual grains fluctuates randomly between the maximum and minimum length.

winShape – Grain windowing function. Allows the user to choose from a selection of common windowing functions to be applied to the samples retrieved from the buffer. Different shapes have different effects on the frequency spectrum produced by the grains. For a more thorough discussion of window functions see Roads' *The Computer Music Tutorial* (Roads 1996).

pitch_mult – Pitch multiplier. Used to raise or lower the perceived pitch of the grains. It is actually used to control the sampling increment of the sound buffer, but the aural result is in the form of a pitch manipulation. For most samples, a setting of 1.0 would be normal playback, 2.0 an octave higher and 0.5 an octave lower.

pitch_bw – Pitch multiplier bandwidth. Provides the bandwidth for a random deviation from the specified pitch multiplier expressed as a percentage of that value. Again the use of a percentage gives the bandwidth a consistent aural width across the frequency spectrum.

stream_config – Stream configuration. Used in the "space" effects, this parameter allows the user to turn the individual streams of granular sampling on or off. The related message has a somewhat different format from those for other parameters. It must have the parameter name followed by the stream number and a 1 or 0 to change the stream's state. For example, a message reading "stream_config 4 0" would turn stream four off, while "stream_config 4 1" would turn it on again.

gain – Provides a basic gain control over the granular output. Expressed in decibels, a value of 0.0 dB would multiply the output by 1.0 while a value of -76.0 dB would multiply the signal by 0.0.

gain_bw – Gain bandwidth. Controls the bandwidth of the random deviation from the specified gain parameter. The value is expressed in decibels and allows the individual grains to have varied gains.

pan_min – Minimum panning position. It is used only in the "cloud" abstractions in which individual grains are randomly panned across the stereo spectrum. This control gives the user the ability to limit the width of the random panning. The value is expressed as a floating point number from 0.0 to 1.0, with 0.0 representing hard left and 1.0 representing hard right.

pan_max – Maximum panning position. Used in connection with "pan_min" in order to limit the area of the stereo spectrum in which the random panning can occur.

CONCLUSION

With its externals and abstractions, this granular toolkit should appeal to both the novice and expert user of Cycling74's Max/MSP. Novices will find useful effects built to be easily incorporated into their own work. Experts will find useful tools that allow them to freely experiment with the possibilities of granular synthesis at various levels of precision. Both should be pleased with the efficiency and flexibility that the granular toolkit offers to all.

This toolkit is freely available for download from the author's web site at the following address: "<http://www.nathanwolek.com>" Follow the menu link to the software section in order to download a compressed version of the files. The user must follow the included instructions on where to place the files on the hard drive. A working copy of Cycling74's Max/MSP must be present in order to use the toolkit.

ACKNOWLEDGEMENTS

This toolkit began as part of the development for several projects on which I worked as programmer. They include Amnon Wolman's *Thomas and Beulah*, Paul Hertz's *Fool's Paradise* and the Northwestern Center for Art and Technology's *Portal Project*. I would like to sincerely thank the artists involved in those projects for pushing me to develop the ideas that evolved into this toolkit. Additional support this research was provided by a grant from the Northwestern University Graduate School.

REFERENCES

- Cycling74. 2001. *Max 4: Getting Started*. Online. Available: <http://www.synthesisters.com/download/Max4GettingStarted.pdf>. 1 March 2002.
- Dattorro, Jon. 1997. "Effect Design, Part 2: Delay-Line Modulation and Chorus." *Journal of the Audio Engineering Society* 45 (10): 764-788.
- Gabor, Dennis. 1947. "Acoustical Quanta and the Theory of Hearing." *Nature* 159 (4044): 591-594.
- Jones, Douglas L. and Parks, Thomas W. 1988. "Generation and Combination of Grains of Sound." *Computer Music Journal* 12 (2): 27-33.
- Lippe, Cort. 1994. "Real-Time Granular Sampling Using the IRCAM Signal Processing Workstation." *Contemporary Music Review*, Vol. 10, United Kingdom, Harwood Academic Publishers, pp. 149-155.
- Roads, Curtis. 1978. "Automated Granular Synthesis of Sound." *Computer Music Journal* 2 (2): 61-62.
- Roads, Curtis. 1991. "Asynchronous granular synthesis." In G. DePoli, A, Piccialli, and C. Roads, eds. *Representations of Musical Signals*. Cambridge, Massachusetts: MIT Press, pp. 143-186.
- Roads, Curtis. 1996. "Granular Synthesis." In *The Computer Music Tutorial*. Cambridge, Massachusetts: MIT Press, pp. 168-184.
- Truax, Barry. 1988. "Real-Time Granular Synthesis with a Digital Signal Processing Computer." *Computer Music Journal* 12 (2): 14-26.
- Truax, Barry. 1990. "Time-Shifting of Sampled Sound with a Real-Time Granulation Technique." *Proceedings of the 1990 International Computer Music Conference*. San Francisco: International Computer Music Association, pp. 104-107.
- Xenakis, Iannis. 1971. *Formalized Music*. Bloomington: Indiana University Press.