

Granular Sampling Using a Pulse-Coupled Network of Spiking Neurons

Eduardo Reck Miranda and John Matthias

Computer Music Research Faculty of Technology and Faculty of Arts,
University of Plymouth, Smeaton 206, Drake Circus, Plymouth PL4 8AA, UK
{john.matthias, eduardo.miranda}@plymouth.ac.uk
<http://cmr.soc.plymouth.ac.uk/>

Abstract. We present a new technique for granular sampling using a pulse-coupled network of spiking artificial neurons to generate grain events. The system plays randomly selected sound grains from a given sound sample when any one of a weakly coupled network of up to 1000 neurons fires. The network can exhibit loosely correlated temporal solutions and also collective synchronised behaviour. This leads to very interesting sonic results, particularly with regard to rhythmic textures which can be controlled with various parameters within the model.

1 Brief Introduction to Granular Synthesis

Granular synthesis works by generating a rapid succession of very short sound bursts called grains that together form larger sound events. The notion behind it is largely inspired by a sound representation method published in a paper by Dennis Gabor back in the 1940s [1]. Gabor's point of departure was to acknowledge the fact that the ear has a time threshold for discerning sound properties. Below this threshold, different sounds are heard as clicks, no matter how different their spectra might be. The length and shape of a wavecycle define frequency and spectrum properties, but the ear needs several cycles to discern these properties. Gabor referred to this minimum sound length as an acoustic quantum and estimated that it usually falls between 10 and 30 milliseconds, according to the nature of both the sound and the subject.

2 Approaches to Granular Synthesis

As far as the idea of sound grains is concerned, any synthesiser capable of producing rapid sequences of short sounds may be considered as a granular synthesiser. Three general approaches to granular synthesis can be identified as follows [2]: sequential, scattering and granular sampling approaches. The sequential approach works by synthesising sequential grain streams. The length of the grains and the intervals between them are controllable, but the grains must not overlap. The scattering approach uses more than one generator simultaneously to scatter a fair amount of grains, not necessarily in synchrony, as if they were the 'dots' of a 'sonic spray'. The expression 'sound clouds' is usually employed by musicians to describe the outcome of the scattering approach.

Granular sampling employs a granulator mechanism that extracts small portions of a sampled sound and applies an envelope to them. The granulator may produce the grains in a number of ways. The simplest method is to extract only a single grain and replicate it many times. More complex methods involve the extraction of grains from various portions of the sample. In this case, the position of the extraction can be either randomly defined or controlled by an algorithm.

Thus far, most granular synthesis systems have used stochastic methods to control the production of the grains; for example, a probability table holding waveform parameters can be called to provide synthesis values for each grain during the synthesis process. As an alternative method, we have devised Chaosynth, a granular synthesiser that uses cellular automata to manage the spectrum of the sound grains [3]. Chaosynth explored the emergent behaviour of cellular automata to produce coherent grain sequences with highly dynamic spectra.

The challenge with granular sampling to find interesting and new controllable ways of playing back the grains, which are taken from the input sound sample. Grain events which are triggered independently will produce randomised signals which can have very interesting flow textural properties [4]. However in order to go beyond this one needs to look at introducing some kind of correlation in grain parameters whilst maintaining the inherent stochastic element which has been so effective in granular synthesis algorithms thus far. Chaosynth utilised the emergent behaviour of a cellular automata model in order to do this. Other attempts have looked at the collective properties of a large number of interacting particles, or swarms, to generate grain events [5]. In this work we have used the correlated firing properties of a large collection of pulse-coupled artificial neurons.

Spiking neural networks have a very rich dynamics and the relevant timescales are of the same order as those relevant to granular synthesis. (i.e., on the level of milliseconds and tens of milliseconds). This makes them very suitable to use as a triggering mechanism for a granular sampler. There is great variety in the dynamics at the level of the single neuron and this becomes even more interesting when we look at networked systems. The single neurons can show regular spiking, bursting (very fast spiking) so-called chattering and resonant behaviour. When connected they exhibit collective excitations on timescales larger than the inherent responses of single neurons (Fig. 1). Such collective excitations include synchronization of the firing times of large numbers of neurons in groups and repetition of signals over very large time scales (of the order of seconds) which have become known as Cortical Songs [6].

3 Spiking Neural Networks

Essentially one can visualise a neuron as an object that fires a spike signal when its input voltage exceeds a certain threshold [7]. The amplitude of the spikes of real neurons is of the order of 100mV (millivolts) and the duration of the spikes are of the order of 1-2ms. The spikes then travel to all the other (post-synaptic) neurons to which this (pre-synaptic) neuron is connected. The time taken for these spike signals to reach the post-synaptic neurons is also of the order of milliseconds. When one of the post-

synaptic neurons receives a spike it will fire a spike in turn if its current voltage state plus the signal of the spike are above its threshold and so on. Most of this processing takes place in the cortex. Each neuron in a mammalian cortex is typically connected to up to 10 000 other neurons so one can see quickly how complicated the resulting dynamics would be from simply sending out a single spike to just one neuron.

The firing patterns of individual cortical neurons are known to be very varied. In Fig. 1 we can see twelve of the most common forms of mammalian cortical neuron firing. As far as the interaction of the neurons with other neurons is concerned, the class of spiking neuron models in which we are interested are called Pulse Coupled Neural Networks (PCCN). Essentially, this means that when a neuron receives a spike it updates its connection with all the neurons with which it is connected. In such models, the connections between neurons are modelled by a matrix of synaptic connections $S = (s_{ij})$, and these synaptic connections are used inherently in the dynamics. In this paper, we have used the simple condition that when the j th neuron fires, the membrane potential, v_i of the all the connected neurons immediately increases by s_{ij} [8].

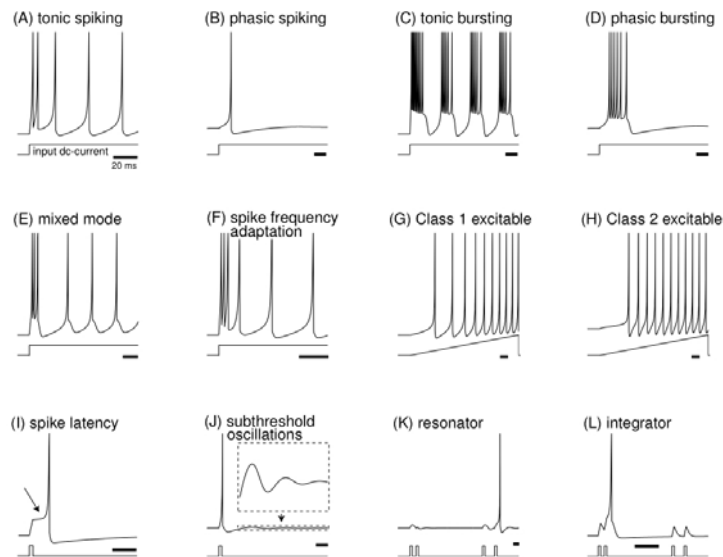


Fig. 1. Twelve of the different types of firing patterns exhibited by single neurons in the mammalian cortex. (This figure is reproduced with permission from Eugene Izhikevich.)

4 Izhikevich's Pulse-Coupled Neural Model

It has been recently discovered that surprisingly simple mathematical models of spiking neurons with random connections can produce realistic organised collective behaviour. The model of Eugene Izhikevich [8] [10] contains enough detail to produce the rich firing patterns found in cortical neurons (Fig 1), yet is also computationally very efficient.

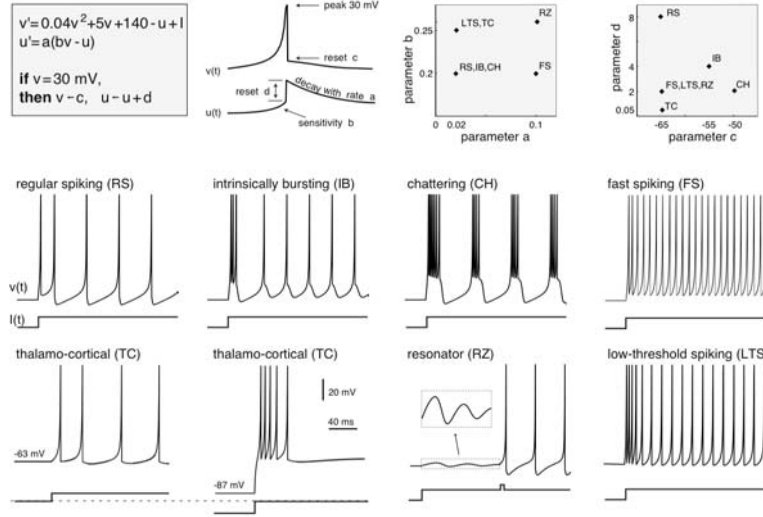


Fig. 2. Known types of neurons correspond to different values of the parameters a , b , c and d in Izhikevich’s model. Each inset shows a voltage response of the mode neuron to a step of dc-current $I=10$ (bottom). (This figure is reproduced with permission from Eugene Izhikevich.)

The temporal firing patterns of the network show both stochastic and synchronised behaviour depending on the values of various parameters, the number of neurons, the matrix of synaptic connections and the history of the behaviour of the network. The frequencies of collective modes of the system are between 1 and 40Hz and present a very interesting case for controlling a granular sampler, particularly in terms of rhythmic structure. We therefore use Izhikevich’s model along with a granular sampler such that grains of sounds (taken from a recording) are triggered when any of the neurons fire.

The model contains N neurons, each of which are described by two dimensionless variables v_i and u_i where v_i represents the membrane potential of the i th neuron and u_i represents a membrane recovery variable, which provides negative feedback to v_i . The system is then described by the following coupled ordinary (nonlinear) differential equations:

$$\frac{dv_i}{dt} = 0.04v_i^2 + 5v_i + (140 - u_i) + I_i \tag{1}$$

$$\frac{du_i}{dt} = a(bv_i - u_i) \tag{2}$$

with the following auxilliary after spike resetting; if $v_i \geq 30$ millivolts then $v_i \rightarrow c$ and $u_i \rightarrow (u_i + d)$. Essentially, the first of these conditions means that when a neuron receives a spike input then its membrane potential is immediately reset.

The neurons are coupled to one another through a matrix of synaptic connection weights. These synaptic connection weights are given by the matrix $S = (s_{ij})$, such that the firing of the j th neuron instantaneously changes variable v_i by s_{ij} . We have used a version of Izhikevich's model where the matrix S is a random matrix. However, in other versions of the model, S can updated itself according to various learning algorithms such as 'Spike Timing Dependent Plasticity' in which connections between neurons are reinforced according to temporal correlations. Synaptic currents or injected dc-currents (currents which come from either other neurons or from sensory information) are encompassed within variable I (which in our version is also a random variable) and, a , b , c and d are parameters whose effects are summarised in Fig 2. Essentially, different values of these parameters produce different individual intrinsic neuron firing patterns such that complex spiking, bursting or chattering of cortical and thalamic neurons can be simulated.

5 Controlling a Granular Sampler

The algorithm by which the granular sampler works is straightforward: When a neuron in the network fires at time t , a sound grain of random length (between 10-50ms) and random amplitude is taken from a random place in a recorded sample of sound and played back. The sound grain is convoluted within a Hanning envelope [11]. Effectively, the neural network plays a granular sampler. Synchronized firing of neurons sound like a pulse, whilst networks containing only a few neurons have a very interesting sparse rhythmic quality (between completely random and correlated). The system therefore has a very wide variety of temporal patterns and behaviours, which can be controlled according to the parameters in the mathematical model. One can control the parameters a , b , c and d , which determine the intrinsic properties of the neurons and one can control the number and type of neurons. In the current version, the connections are completely noisy in the sense that the matrix S is a random matrix and all current inputs are noisy. However it would be straightforward to extend the model by varying the connections and the input ('thalamic') current. Generally speaking, increasing the number of neurons in the model means more firing and therefore more sonic texture, although when the solutions exhibit synchronous behaviour increasing the number of neurons tends to lower the frequency of the collective response. It is interesting in itself that such random (noisy) inputs can produce synchronous pulses of sound.

Generally speaking, in this version of the model (without any temporal correlation such as Spike Timing Dependent Plasticity [7]) one gets interesting sounds if we have either rather few (up to 10) or very many (over 500) neurons. The result with up to 10 neurons sounds very sparse but one can hear rhythms, which appear and then transiently die away. They do not repeat exactly; the network is effectively isolated from any sensory input (unlike real neurons in a mammalian cortex) and therefore not stimulated by correlated information. The synchronous solution appears in the dynamics if all the neurons selected are the same and if there are more than 500 of them. This sounds like a very gritty pulse, especially if the selected grain size is short.

6 Concluding Remarks

The technique we have introduced successfully fulfills the object of our enquiry in that its domain lies right in between the completely random and completely correlated in its temporal behaviour. Given a set of initial parameters, outputs are not predictable fully due to the large number of noisy elements in the model, but do follow discernable dynamical patterns especially when the system is in a dynamically synchronised state. The output is also controllable to a large extent. There would seem to be much profitable study from looking at this method further.

Acknowledgements

We would like to thank Stephen Coombes, University of Nottingham's Institute of Neuroscience, for stimulating discussions.

References

1. Gabor, D. (1947). "Acoustical quanta and the theory of hearing", *Nature*, 159(4044):591-594.
2. Miranda, E. R. (2002). *Computer Sound Design: Synthesis Techniques and Programming*. Oxford (UK): Elsevier/Focal Press.
3. Miranda, E. R. (1995). "Granular Synthesis of Sounds by Means of a Cellular Automaton", *Leonardo*, 8(4):297-300.
4. Truax, B. (1988). "Real Time Granular Synthesis with a Digital Signal Processor", *Computer Music Journal*, 12(2):14-26.
5. Blackwell T. M. and Young M. (2004). "Swarm Granulator", *Proceedings of Applications of Evolutionary Computing EuroWorkshops 2004*, LNCS 3005, Springer-Verlag, pp. 399-408.
6. Ikegaya Y., Aaron G., Cossart R., Aronov D., Lampl L., Ferster D. and Yuste R. (2004), "Synfire chains and cortical songs: Temporal modules of cortical activity", *Science*, 304:559-564.
7. Gerstner, W and Kistler, W. M. (2002). *Spiking Neuron Models*. Cambridge (UK): Cambridge University Press.
8. Izhikevich, E. M. (2003). "A simple model of a spiking neurons", *IEEE Transactions on Neural Networks*, 14:1469
9. Hoppenstadt, F. C. and Izhikevich, E. M. (1997). *Weakly connected neural networks*. New York (NY): Springer-Verlag.
10. Izhikevich E. M. (in print). "A Simple model of a spiking network' *IEEE Transactions on Neural Networks*".
11. Abramovitz M. and Stegun I. (1972) *Handbook of Mathematical Functions*. New York (USA): Dover.