

Taking its name from a dubious pun on the words 'rushing-dragging', the Russian Dragon was released on an unsuspecting market about five years ago. After receiving a number of very reasonable reviews and endorsements from major producers, it then appeared to vanish without trace and is now a rare sight even on the secondhand market.

The basic concept of the unit is to provide an accurate method of determining which of two different audio signals came first -- a perennial problem in the recording industry since the first drummer played along to a click track -- and one which has been further compounded since the arrival of frame accurate sequencers and MIDI triggered samplers.

The Russian Dragon offers a finite LED display capable of graphically showing the timed relationship of two individual sounds, all within a claimed accuracy of a fraction of a millisecond.

DESIGN

The rear of this extremely compact 1U unit is simplicity itself, housing only three sockets: two mono jacks for the audio inputs that feed channels 1 and 2, and a connection for the power supply (a simple external unit supplied as standard).

The front panel is dominated by a large, horizontal row of 25 square LEDs, consisting of a central display, surrounded by 12 smaller versions on either side. The comparative timing between the two main audio signals is then shown as a position on this scale, offering a visual indication of the current synchronisation state.

Another set of jack inputs for the two channels are fitted to the front of the unit, together with separate rotary input level controls. Simple LED VU meters are available for monitoring source levels, and each channel is also equipped with a Polarity Check button -- more of which later.

At the right-hand end of the panel are two Mask rotary controls, together with their associated trigger LEDs, which are used to set the speed at which an individual channel can re-trigger.

2.

Lastly, a nine position rotary switch, next to the power button, adjusts the scale of the main display from one to nine milliseconds in single units Channel 1 is designed to be the 'reference' or 'click' input, and channel 2 takes the sound or instrument that is to be checked. Almost any type of sound will do, but obviously, the more rhythmic the input, the greater accuracy achieved.

Once a reference source has been connected, the large central LED flashes whenever audio is received on this channel. Any signal detected on channel 2 then lights one of the other 24 LEDs, indicating its relative timing with reference to the original signal.

A lit LED to the left of centre indicates that the signal is 'dragging' (i.e. later than the reference), and one to the right indicates that the signal is 'rushing' (i.e. earlier). The further the outer LED is from the centre of the display, the greater the time differential detected between the two signals. To check a drummer's synchronisation with a click track, a wide scale would be used; a narrower one to test the timing of a sequencer against a recording.

One exceptionally good feature is the inclusion of the Mask controls, which determine a channel's re-trigger speed. If a channel appears to 'double trigger' because of extra noise between transients, or extra beats in the rhythm, then adjusting the Mask time can silence the input to the Russian Dragon when a trigger is not required. With careful use of the Mask function, it is possible to analyse even the most complex of drum tracks against a simple click.

ALL FIRED UP

Set up is quick and easy -- feed the two separate signals into the appropriate sockets, adjust the input levels, then sit back and watch for any timing discrepancies. For the really hi-tech approach, Jeanius have included a polarity check button on each channel, which displays whether a trigger sound begins with a positive or negative transient -- the Russian Dragon analyses the time between positive transients only, so this is a good way of testing for possible errors, especially when synchronisation times are very small.

3.

Feeding the same pulse into both channels immediately showed the 'snake eyes' display -- the two inner LEDs light up to confirm perfect synchronisation. Gradually introducing a delay into channel 2 caused the unit to track the timing difference with a moving LED. The owner's manual is well thought out, with a number of in-depth examples describing how to use the unit in typical situations.

The only caveat is that some care should be taken when setting the input levels, as overloading the unit (or alternatively feeding it a signal that is too low) can cause inaccurate readings. This is only to be expected with a unit such as this, and in general the system performed flawlessly.

CONCLUSION

I can think of many different uses for the Russian Dragon in today's technology-based studios. As the producer Keith Cohen states in the manual, the system is probably best left connected directly across a Mix-2 buss, so that if any sounds need to be checked for timing accuracy, all the engineer has to do is solo both sources and pan them hard left and hard right.

Programmers who resort to locking up kick drums because the synchronisation details have been left off the track sheet need worry no longer, and drummers playing along to click tracks can now also have the benefit of a visual display, rather than just a loud ticking noise in their cans.

Other uses include synchronising delayed loudspeaker systems for PA use or locking variably spaced microphones in phase. The overall accuracy of the system is very high, and as such, I am sure that the unit would be a worthwhile addition to any studio rack.

RUSSIAN DRAGON £445

PROS

- Simple to use.
- Mask function helps analysis of complex drum tracks.
- Extremely accurate timing reference.

CONS• None! **SUMMARY** Essential visual timing aid for drummers who play to click tracks or anyone wishing to check timing sources against an accurate reference.

The Russian Dragon RD-3/R

This Page Has been Intentionally Left Blank:

1.

We've all heard it said that being in the right place at the right time can make or break a career in music. Nowhere is this more important than in your musical performances. The difference between a great player and a good one can often be reduced to timing. Timing is what gives music that special funky flow that people want to groove to. The error might be as small as a few milliseconds, but those milliseconds determine whether people want to dance or even listen to your music.

When music is rhythmically right, it just feels good. It's always been difficult to judge where your rhythms are placed relative to a mechanical click. Or at least for most of us.... Stories of Donald Fagen's ears perceiving millisecond differences and voicing his disillusionment with MIDI gear are well documented and are what legends are made of.

Most producers and engineers I know would love a tool that would show players where they fall relative to the click. MIDI's Murky Pathways The desire for accuracy and the ability to measure it are also crucial in MIDI production.

In a studio, there are countless causes of timing delays and errors that affect the music in small and large ways. We are falsely led to believe that when we play something and record it, we will hear that performance played back without rhythmic alteration.

Admittedly, MIDI gear is a major perpetrator of these delays.

At the smallest level, MIDI itself can only communicate about one Note On per millisecond. "How Simultaneous Can Two MIDI Notes Be?" A good sequencer can read the differences in your performance as closely as one millisecond per note

The Math: For example, if your tempo is 120 bpm, and your sequencer has a resolution of 480 ticks per beat, you have a real time resolution of 960 ticks per second: $120/\text{beats per minute} = 120 \text{ bpm}/60 \text{ seconds} = 2 \text{ beats per second (beats per second)} \times (\text{ticks per beat}) = \text{ticks per second}$ $2 \text{ beats per second} \times 480 \text{ ticks per beat} = 960 \text{ ticks per second}$ = approximately 1000th of a second.

2.

Practically speaking, however, a serial cable is incapable of sending two notes at the exact same time: one note must follow the other by at least one millisecond. (On a Macintosh, serial cables are the modem and printer cables.)

More often than not, you are sending multiple events in the hope that they will play simultaneously, but they don't.

Plus, events scheduled for simultaneous playing will appear one tick apart, in random order. One way to solve the random-order problem is actually to spread simultaneous events by one tick in your sequencer.

So far, we know that the accuracy is hampered by the speed of the serial communication and the quantity of notes simultaneously played. This problem is compounded by other MIDI data communicated at the same time, further slowing this process and making playback even less consistent.

Many of you have seen your computer slow down when burdened with many complicated activities. This same problem plagues some sequencers' metronomes when reproducing a complicated MIDI performance.

This can be true for both computer sequencers and drum machines. The problems don't stop with sequencers, however. Once our MIDI data has left the serial cables, they arrive at a MIDI interface. These interfaces can create more delays as they translate information from serial to MIDI.

Finally, the data arrives at the MIDI modules. Guess what? More delays! Not only do different modules differ in the speed at which they produce audio from your MIDI data, sounds within a module vary in the time it takes them to be produced. What's a Few Milliseconds Between Brothers?

So what do we do about these millisecond differences? And how can we possibly guarantee that two MIDI notes will occur at the exact same time? Here's the product pitch: a tool that sent me scurrying to find all the faults with my own production studio.

3.

Jeanius, a company in San Antonio, has provided the recording world with a tool to simplify the judgment and analysis of all these time errors.

It is called the Russian Dragon (RD-R and RD-T), named for its ability to accurately perceive if your audio is "rushin'" ahead or "draggin'" behind a reference source you supply. The idea for the unit came from Audio Engineer/Designer Marius Perron and his drumming brother, who had always wished they had a mechanism to judge how closely he was playing to the sequenced tracks.

Marius made the prototype to remove both guesswork from the sessions and tension from not knowing who was right. The unit gave honest feedback to the drummer, the engineer, and the producer. Description: The Anatomy of the Russian Dragon. The Russian Dragon is available in a small rack-mount version (the RD-R) and a less expensive tabletop version (RDT).

I opted for the rack-mount model. The RD-R is an uncomplicated single-space rack-mount unit with 1/4-inch inputs on the back and the front. The top input (channel 1) is used for the reference and the bottom input (channel 2) is used for the other audio event you will judge to be dragging (slow), rushing (fast), or dead on (what Jeanius calls "snake eyes") relative to the reference. The unit features an input level control and four LEDs for visual confirmation of audio and amplitude adjustment.

A large row of 25 coloured LEDs indicates the timing differences from a tenth of a millisecond to 99 milliseconds. The LEDs are adjustable in 1-9 millisecond increments as set by a selector knob marked "ms per LED." A trigger LED indicates the presence and duration of the two signals. The duration is adjustable by a "Mask Control."

By increasing the length of the Mask Control, you can eliminate accidental re-triggering of the timing LEDs. This allows the unit to ignore delay characteristics or extraneous sounds that would make analysis difficult.

4.

A polarity check button for each input ensures that the shape of the transient wave forms start with positive (above the zero crossing point) sections. Polarity is checked for two reasons. First, the sensor that reads the beginning of the audio event is triggered when the signal goes above the zero threshold.

So, a signal with a rarefaction at its entrance would not trigger until later in the wave form, when it moved above the zero crossing. Second, two out-of-phase drums, with similar transients, combined and aligned by this unit would produce a cancellation or thinning phase effect, which is probably contrary to the Mother-of-God drum sound you were looking for when you aligned six snare samples.

The unit is rounded out by a wall wart power cord on the rear, a power switch on the front, and an extra pair of inputs for the convenience of having them on the front as well as the back. Figure 2: Picture of two out-of-phase waves Testing: Let's See Who's Been Naughty or Nice The minute I took the RD-R out of its box, I wanted to test the delays present in my studio system.

It's as if your ears have been bionically upgraded to distinguish millisecond differences.

Not all of the equipment in my eight-track MIDI studio is spanking new. But since good used gear is widely available, the Russian Dragon is especially valuable because it can measure the accuracy of used as well as new equipment. There is little documentation of MIDI delays in my older pieces (some of which create classic sounds which I still consider vital).

Adaptable and precise, the Russian Dragon measures reaction times, consistency, and sync ability in software and hardware sequencers, metronomes, sync boxes, digital delays, and MIDI modules. By testing the three delay units in my studio I could establish if they were consistent enough to aid me in performing other tests. I hooked up a click directly out the back of my Macintosh speaker output (using Performer software), through a Yamaha SPX90, into reference channel 1 in the RD-R. The click track also went to a Lexicon PCM70, then on to reference channel 2 in the RD-R.

5.

After setting the single-delay programs in both units to the same amount of time, the Russian Dragon showed the delayed clicks to be arriving together -- at "snake eyes," meaning that the units were accurate to a tenth of a millisecond. This consistency was wonderful to confirm, although "snake eyes" actually occurred when the units were set at slightly differing delay settings.

Thus, while the units retain the delay settings they claim to have (to a tenth of a millisecond), that delay setting is not necessarily what is displayed in the LCD. I also tested a Yamaha SPX90 against these units, with similar results.

These results are displayed in Figures 3 and 4. Next I was curious about the accuracy of the LED indicators on the front of the RD-R itself. By setting different delay times, I could check the unit's sensitivity to gradual increases in timing delays.

For instance, at the 1 millisecond setting, if the light lit up one unit to the left of the "snake eyes" setting, this meant that the two inputs were approximately 1 millisecond apart. At the 2 millisecond setting, the same light meant that the distance was approximately 2 milliseconds, and so forth (see Figure 5).

The most sensitive setting portrayed millisecond differences easily and a tenth of a millisecond accuracy only when the events were simultaneous (at "snake eyes"). The nine settings were measured with each of the three delays as well. The inconsistencies were more revealing of the discrepancies in the digital readout of the digital delay units.

The resulting numbers found in the various settings might look like the settings are far from the times they claim to measure.

In fact, they are halfway between the amount they are supposed to indicate. For instance, the first LED lights up when the SPX90 is set at 7.5 milliseconds and the "ms per LED" knob is set at 5 milliseconds because the delay is halfway between 5 and 10 milliseconds (see Figures 6, 7, and 8). Figure 6

6.

--- LEDs Lit --- snake KNOB

11 10 9 8 7 6 5 4 3 2 1 Snake eyes SETTING

10.9 9.9 9.0 8.0 7.1 6.1 5.2 4.2 3.2 2.1 1.3 --- (1)

2.9 --- (2)

4.4 --- (3)

6.0 --- (4)

7.5 --- (5)

8.6 --- (6)

10.4 --- (7)

12.0 --- (8)

13.6 --- (9)

Yamaha SPX90 II

--- LEDs Lit --- snake KNOB

11 10 9 8 7 6 5 4 3 2 1 eyes SETTING

10.3 9.3 8.4 7.4 6.5 5.5 4.6 3.6 2.6 1.7 .7 --- (1)

2.3 --- (2)

3.8 --- (3)

5.4 --- (4)

6.9 --- (5)

8.2 --- (6)

9.8 --- (7)

11.4 --- (8)

13.0 --- (9)

Lexicon PCM 70 --- LEDs Lit --- snake KNOB

11 10 9 8 7 6 5 4 3 2 1 eyes SETTING

10.4 9.68 8.82 7.61 6.84 5.77 4.80 3.64 2.87 1.81 .83 --- (1)

2.40 --- (2)

3.91 --- (3)

5.45 --- (4)

7.22 --- (5)

8.41 --- (6)

10.1 --- (7)

11.5 --- (8)

13.3 --- (9)

7.

What's the MIDI Module Delay Time? I wanted to know how long it takes for audio to come out of my MIDI modules when they were programmed to play fully quantised quarter notes in sync with the Macintosh click (see Figure 9).

You'll see that the delays range from 1 to 5 milliseconds. Keep in mind these are ideal circumstances - the units are only playing monophonic quarter notes and nothing further is being asked of the sequencer or the MIDI unit.

Hardly a realistic challenge, but still there are measurable differences. The differences themselves are inconsistent. I've indicated where it varied from note to note by as much as 2 milliseconds. Again, these may seem like small amounts of time. But when added together, they are perceivable and they change randomly.

Figure 9

Device vs. Device Mode Output Drag (ms)

Mac spkr Akai S1000 sample mode mix out 4-5

Mac spkr Akai S1000 program mode mix out 5-6

Mac spkr Akai S900 sample mode mix out 3

Mac spkr Akai S900 program mode mix out 2-3

Mac spkr Yamaha RX-5 pattern mode mono out 5

Mac spkr Yamaha RX-5 pattern mode rim out 4-5

Mac spkr Roland TR808 NA main out 2

Mac spkr Roland TR808 NA rim out 1-2

Mac spkr Emu SP-12 Turbo sequence mode mix out 2-4

Mac spkr Emu SP-12 Turbo sequence mode rim out 3-4

Hardware- and software-based sequencers were the next subject for analysis. All I really tested was the consistency of the metronomic pulse put out by these units. Given more time I'd test start inconsistencies and sync influences.

Once again, I used a digital delay as a constant unit of time. I delayed one quarter note to be in sync with the next. Include if we get AES permission: permission: I have also included findings from Marius Perron's AES presentation.

8.

Sequencer Delay Resolution: Complicated at 160 bpm Setting
Resolution drum-part with eight parts Mac Performer* 375.5 ms 1 ms 2 ms
(Macintosh output spkr click)

Emu SP12 Turbo 376.1 ms 0.1 ms 2 ms
(rimshot)

Roland TR808 NA 0.1 ms 0.1 ms but slower tempo
(rimshot)

Yamaha RX5 367.7 ms 0.1 ms 3 ms variation
(rimshot)

Boss Dr. Beat DB-66 376.1 ms 0.1 ms NA
(rimshot)

*Mark of the Unicorn, Version 4.2;

Macintosh SE/30 system 7.1 Marius Perron's Findings,
Presented at the 91st AES Convention. Sequencer Resolution Wittner
Taktell piccolo metronome
Korg M-1, quarter note side-stick 3

Boss Dr. Rhythm DR550 drum machine 4
Akai MPC-60 sequencer (metronome out) 1
Akai MPC-60 sequencer, quarter note pattern (individual out) 5
Boss DB music conductor no errors
Macintosh 2CX with Vision 1.2u2 (spkr out) 2
Korg DTM12 digital tuner/metronome no errors
Garfield Electronis Dr. Click no errors
Alesis HR16B drum machine 0.6
UREI model 964 digital metronome no errors
Emu III sequencer (metronome out) 3
Roland TR909 drum machine 5
Alesis MMT-8 sequencer, MIDI'd to drum machine 2
Roland SBX-80 (metronome out) 1

9.

These units also exhibited inconsistencies in their start-up time. We found that the tempo that each called 160 bpm varied -- as well as the consistency between quarter notes.

As a drummer, I've always wondered why some metronomes seemed easier to follow than others. I often attributed it to how I felt on a given day or the sound of these units. With the Russian Dragon,

I learned that metronomes can be as moody as I am. To this end, I tested a more complicated sequence as well as just a quarter note pulse. The findings are validating for the drummer in me . . . but upsetting for the producer in me. There was an 8.4 millisecond difference between the RX5's 160 bpm and the SP12's 160 bpm.

The quarter note test showed fairly stable tenths of a millisecond resolutions among the SP12, TR808, and RX5 sequencers. The surprises were the David and Goliath performances of the Macintosh SE/30 with Performer 4.2 and the Boss Dr. Beat metronome.

The Macintosh was inconsistent 1 millisecond, and the Dr. Beat held "snake eyes" at better than tenths of a millisecond stability for minutes on end.

The Mac and all of the drum machines suffered when many drum parts were added to the playback. The TR808, an analog drum machine, could not be measured for tempo accuracy because it has no digital readout.

Strangely, it was not less consistent when parts were added. But its tempo did slow down. Its determination to be consistent might be attributable to the 16-note resolution capacity of its sequencer.

Conclusion:

Everything is Not on the One The Russian Dragon helped illuminate many timing idiosyncrasies Like a tuner or oscilloscope, the Russian Dragon is a crucial tool for measuring timing in the studio.

10.

There are many ways to use the RD-R that I haven't had time to try yet. Some of the possibilities include finding offsets for slightly out-of-sync tapes and/or sequences, fixing the delay times between distant speakers, and providing feedback for drummers playing prerecorded and sequenced parts.

I'm looking forward to testing these options. One interesting solution to the MIDI serial problem is actually to split your simultaneous events by 1 tick to solve random-order problems.

In general, I'm a little less comfortable with the resolution of MIDI devices than I was before. I hope the system protocol will be improved.

That's a lot to ask for an industry that is slow to change designs involving cooperative schemes. So companies must agree and share information to make this happen.

If we take, for example, the battle between Opcode's OMS and Motu's Free MIDI, we may wait a long time. Company cooperation got us into MIDI. Let's hope the same spirit carries us to a new, more musical resolution. And let's hope our timing improves.

Thanks Notable Productions' timing is kept by Daniel C. Cantor and Mark Weltner, whose mission is to seek out new timing errors and remove them.

We'd like to thank Marius Perron at Jeanius for creating the Russian Dragon and providing us with the unit and information. [Include if we get AES permission: Thanks to AES for letting us reprint Marius's findings.] Thanks also to fellow time delay sleuth Robert Poor at Opcode R&D Systems for his box of info and his enthusiasm.

Finally Notable owes Kathy Wolff a heap of thanks for her split-second editing. How Simultaneous Can Two MIDI Notes Be? or Time Is Nature's Way of Keeping Everything from Happening At Once MIDI data sent on a MIDI cable are transmitted as serial data.

11.

A sequence of ones and zeros tells your drum machine to "play the side-stick on channel 3 mezzo forte." It takes a bit of time to transmit those ones and zeros. More important, all the information for one note must be transmitted before the next note can begin.

What is the delay between two "simultaneous" MIDI notes on a single cable? For those who just want the facts, the answer is 960 microseconds for ordinary Note

On events or 640 microseconds for Running Status Note On events. How do we get these numbers? Get out your calculator and follow along. The first thing to know is that MIDI data are transmitted at 31250 bits per second, also called 31.25 KiloBaud.

(This peculiar number happens to be 1/16th of 1,000,000 cycles per second. Many early computers had built-in crystal oscillators running at this frequency; this was an easy frequency to derive.)

Each byte of MIDI data requires 10 bits of serial data. The byte is transmitted with one start bit, eight bits of real data, and one stop bit.

The start and stop bits are needed to tell the receiving circuitry "here comes some data," and "that was the end of the data." At a rate of 31,250 bits per second, we can send 3,125 (31,250/10) of these "MIDI bytes" per second.

To send one Note On message requires three MIDI bytes. The first byte says that this is a Note On event and specifies which channel (1-16). The second byte says which key number is to be played. The third byte specifies the velocity.

We can send 1042 (3,125/3) such Note On messages per second. But we promised to tell you the duration of each Note On message, not the number of messages per second.

This is easy: $1/1042$ messages per second = .00096 seconds per message. In technospeak, this is .96 milliseconds, or 960 microseconds per Note On message.

12.

MIDI does offer a slight speed improvement through Running Status messages. After sending a normal three-byte Note On message, a sequencer can send additional key number/velocity messages, omitting the first byte of an ordinary Note On message.

For a string of notes that all occur on the same channel, this reduces the number of MIDI bytes per Note On message from three to two.

We can send 1563 ($3,125/2$) Running Status Note On messages per second, which corresponds to .00064 ($1/1563$) seconds, .64 milliseconds, or 640 microseconds per message.

The bottom line is that for MIDI data running on a single MIDI cable imposes an absolute "speed limit" of about 1000 Note On messages per second, or 1500 Note On messages when Running Status is in effect.

This number gets worse with other MIDI data, such as aftertouch and controller information. Your actual mileage may vary.

Robert Poor
Director of Research and Development
Opcode Systems, Inc.
By Daniel C. Cantor and Mark Weltner

Notable Productions@1994