Utilising jitter noise in the precise synchronisation of laser pulses

Róbert Mingesz[a], Zoltán Gingl[a], Gábor Almási[b] and Péter Makra[a, *]

[a]Department of Experimental Physics, University of Szeged, Dóm tér 9, Szeged, Hungary, H6720.

[b]Department of Informatics in Physics, University of Pécs, Ifjúság útja 6, Pécs, Hungary, H7624

ABSTRACT

We present a laser delay control system based on adaptive averaging which utilises the jitter noise of the laser to stabilise the delay more precisely. The system contains a delay line to measure the laser delay and a microcontroller that runs our control algorithm. The algorithm regulates the laser delay on the basis of the average of detected delay values, wherein the steps with which the delay is varied and the averaging length are chosen adaptively, depending on the distance from the target delay. Our complementary numerical simulations show that the jitter of the laser may play a beneficial role here: the error of the delay has a distinct minimum at a non-zero noise level. In a way similar to the dithering principle applied in analogue-to-digital conversion, averaging the noise-modulated detection instances yields a precision in setting the delay that is well beyond the resolution provided by detection time windows, and is close to the theoretical limit determined by the step size of the delay line we applied.

Keywords: Dithering, stochastic resonance, laser, thyratron, jitter, anode delay

1. INTRODUCTION

Hydrogen thyratrons play an important role in a wide range of applications from pulse gas lasers (excimer, carbon dioxide and copper vapour lasers) to fast kicker systems used in high-energy physics. It may be crucial in many of these applications to have the thyratron precisely synchronised with some external event. The uncertainty of the synchronisation is often required to be less than 1-2 ns.

Thyratrons switch with a delay of 0.1–1 µs after receiving the trigger signal. This amount of time—called the *anode delay*—is required for the area between the thyratron grid and the anode to become conductive. The anode delay is not constant: first, it has a more or less deterministic, slow drift of a few hundred nanoseconds, which is a result of changes in the external parameters (especially the temperature); second, it also displays a shot-to shot uncertainty called the *jitter,* which is significantly smaller and whose source is primarily the gas discharge in the thyratron.

The jitter is an unpredictable random process and we cannot fully eliminate it. Yet the question arises whether we can use it for our own purposes. The constructive potential of noise has long been exploited in a class of technical procedures collectively known as *dithering,* and starting with the problem of ice age cycles,[1, 2] the study of noise in beneficial role has also emerged as a new scientific discipline. A central notion of the latter, *stochastic resonance,* refers to phenomena in which noise contributes to the optimal performance of a system: improves signal transfer, facilitates the detection of weak signals, etc. Stochastic resonance has been observed in a broad range of systems from electronic circuits[3] through ring lasers[4] to sensory processes of certain animals[5] or even human circulation.[6]

In this work, we harness the constructive potential of the jitter noise inherent in the laser system in our algorithm to compensate the long-term drift of the anode delay. We present a control technique based on adaptive averaging and show that this control performs best at a non-zero jitter noise level.

* Corresponding author. E-mail: phil@titan.physx.u-szeged.hu

2. REALISATION OF THE CONTROL SYSTEM

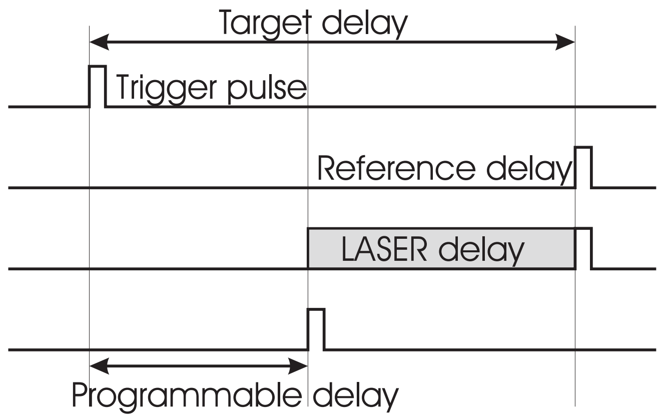

Figure 1. The control principle.

The aim of the control is to ensure that the laser pulses always follow the trigger signal with the same delay. Since the delay is not constant because of the drift (and the jitter), we can achieve this by introducing a programmable delay between the trigger pulse and the launch of the laser pulse. After a time, the laser emits the pulse. The control compares the time instant of the actual emission to a reference delay and decides upon the length of the programmable delay accordingly. The control principle is illustrated in Fig 1.

2.1 Hardware

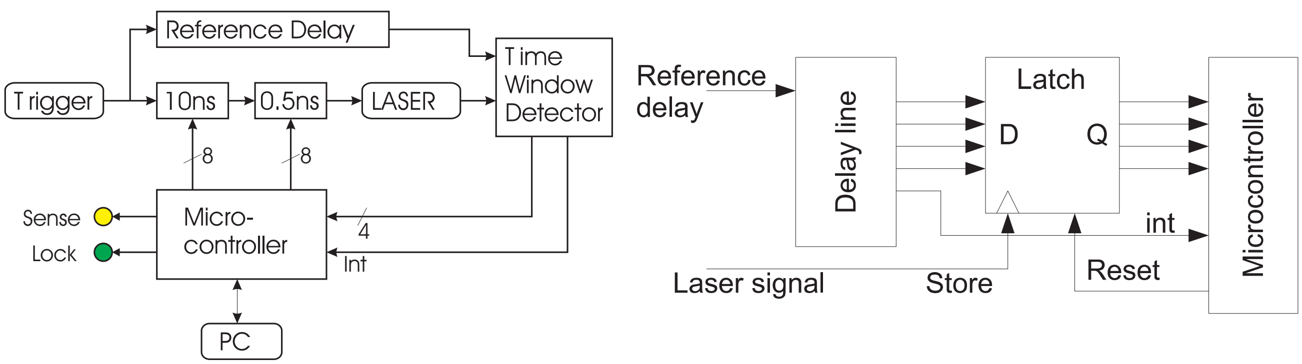

Figure 2. The block diagram of the hardware.

The block diagram of the hardware we built is given in Fig 2. The control circuit consists of two branches: the reference branch, which defines the desired time of the laser pulse, and the driving branch, which sets the programmable delay in order to ensure synchronicity. In both branches we use programmable delay units which can provide delays up to a few microseconds in sub-nanosecond steps. To determine the actual position of the laser pulse, we need several replicas of the reference pulse which are shifted in time with different delays as compared to the original (see Fig 3). We obtain these using a delay line provided by a Maxim DS1020-15 timing unit, which has five output channels and produces pulses positioned at a distance of 4–100 ns from each other. As we shall see later, the system can have to modes: one in which the time windows are arranged in a way that the target delay is located in the middle of the second time window (mid-positioned mode), and another in which the target delay is positioned on the boundary between the first and second time windows (edge-positioned mode).

The signal corresponding to the laser pulse writes the output of the delay line into a latch. The state of the latch will depend on when the laser pulse arrived. The delay line also activates the interrupt routine of the microcontroller, which then reads the state of the latch and executes the control algorithm, setting the new values of the programmable delay. If no laser pulse arrives until the interrupt is activated, the state of the latch will not be overwritten. We can use this fact to detect the absence of pulses: at the end of the interrupt, we write a value which cannot occur in normal operation into the latch, and if we read this forbidden value when the interrupt is activated next, we shall know that no laser pulse has arrived in time.

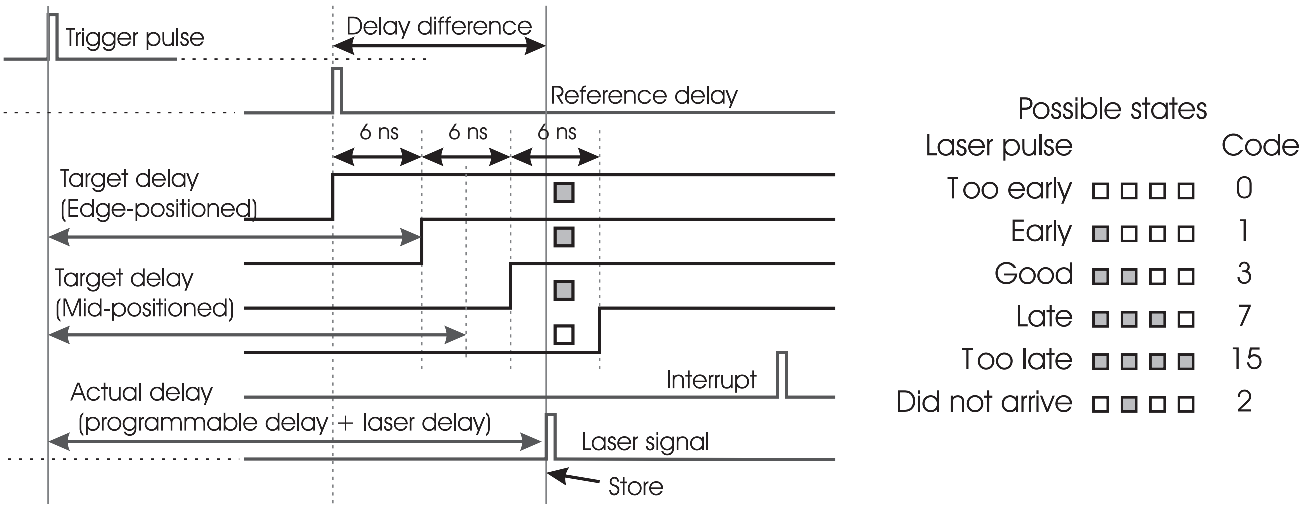

Figure 3. The arrangement of time windows provided by the delay line and the corresponding latch states.

The interrupt can find the latch in a number of states which are summarised in Fig 3. Consequently, we can realise a simple control scheme as follows:

- If we get code 0, the laser pulse arrived too early (but we do not know to what extent): we should increase the delay by a big step.
- If we get code 1, the laser pulse arrived early (in the first time window): we should increase the delay by a small step.
- If we get code 3, we are in the desired time window (the difference is within ± 3 ns): we should do nothing.
- If we get code 7, the laser pulse arrived late (in the third time window): we should decrease the delay by a small step.
- If we get code 15, the laser pulse arrived too late (yet still before the interrupt occurs): we should decrease the delay by a big step.
- If we get code 2 (that is, no pulse had arrived until the interrupt occurred), devoid of the necessary information, we perform no control operation, we only let the user know that problems have arisen (eg, the LED signalling the detection of a pulse will no longer glow).

This simple control scheme has a number of serious problems. It is very sensitive to disturbances: even a small jitter can force it to keep stepping continuously, and individual outlier pulses lead to big steps at once, seriously damaging system performance. In the laser system we set out to improve, a median filter had formerly been applied to circumvent these shortcomings, but it also had a tendency to oscillate about the equilibrium. Similarly, simple averaging cannot yield sufficient stability either.

2.2 The control principle

Though a simple averaging can reduce the effects of disturbances, it still cannot provide a reliable solution, so the idea lends itself that we should use an averaging whose length is adjusted to the current state of the control system (adaptive averaging).

Before averaging we convert the code output by the latch to a code that represents the index of the corresponding time window, so in the case of a valid pulse (ie, that had arrived before the control interrupt occurred) we can average the code values thus obtained directly. The block diagram of the control algorithm is given in Fig 4.

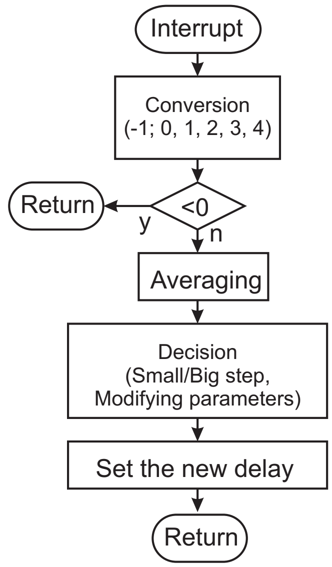

Figure 4. The block diagram of the control algorithm.

A distinct advantage of averaging is that it filters out the disturbances, yet it comes with a trade-off: fast changes are harder to follow. So in the software we adjusted the number of averages to the actual circumstances. We followed the logic outlined below in choosing the length of one averaging sequence:

- Whilst we are looking for the proper value of the programmable delay (that will produce the target delay) and we are yet far from the target delay, we apply a big step. In this case we must not use a lengthy averaging because if we did, we would likely overstep the proper value and an oscillation might ensue. Here we average one or two elements.
- If we are getting close to the target delay, we should choose small steps. In this case, a short averaging can provide a sufficient protection against disturbances whilst retaining the ability to react fast. Simulations show that averaging about four elements is adequate here.
- If neither the big nor the small step has proved necessary for a given time, we can increase the length of averaging, thus improving the precision and the resistance to disturbances.

Apart from yielding a better protection against disturbances, this method can also enhance the resolution of delay detection. If because of the jitter we get delay values in more than one time window, the average of these will be much closer to the real value than any one of the discrete delay values corresponding the windows (see Fig 5). This is very similar to the dithering principle applied in analogue-to-digital conversion.[7] Knowing the error of the delay with a higher resolution allows us to make better decisions in the choice of the programmable delay.

If the average is around the middle of the target interval, no stepping is necessary. In this case, extending the interval will make the system less sensitive to disturbances, whilst narrowing it can make the delay value more precise. We should opt for the latter if a sufficiently large number of averages is available, thus the system will not follow the effects of disturbances.

If the difference between the actual delay and the target delay is greater than a given value, we adjust the programmable delay to approach the target value. After stepping, we must revert to a shorter averaging because

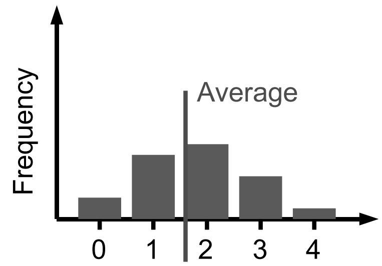

Figure 5. Averaging the noise-modulated detection instances can enhance the precision of delay detection.

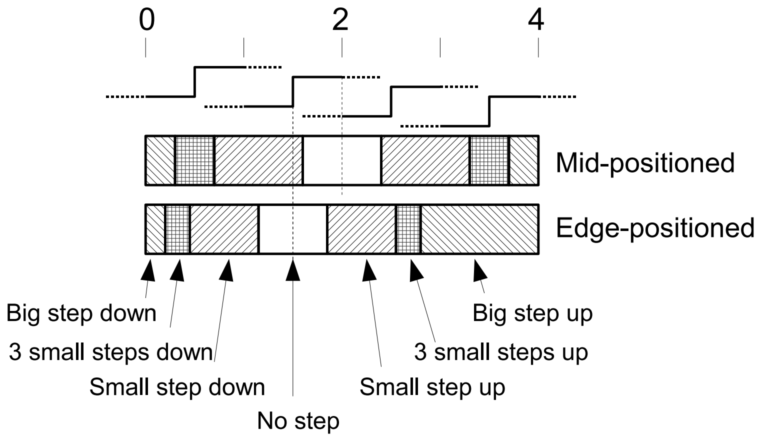

Figure 6. Averaged window code intervals in the refined control algorithm and the corresponding decisions in the mid-positioned and the edge-positioned modes.

the stored values of detection codes correspond to a previous set-up and including them in the average would corrupt it. (We could circumvent it by using the real delay values instead of window codes in the algorithm.)

Small steps are not necessarily sufficient to follow fast changes in the anode delay; on the other hand, big steps may lead to oscillations. To avoid this, we can refine the scheme above by introducing an intermediate adjustment level: when the difference between the actual delay and the target delay is large enough to trigger intervention, three small steps are taken instead of one big step. This way we can extend the drift range in which the control can follow the target delay without resorting to big steps.

If the detected difference is too large, either we are too far from the target delay or the laser system has started to creep. In this case, we can use big steps to correct the delay. We summed up the states of the system and the corresponding decisions in Fig 6.

3. THE PERFORMANCE OF THE CONTROL AS A FUNCTION OF JITTER LEVEL

The width of the time windows in the system we used is 6 ns. If the laser pulse always falls into the same window, we cannot tell its exact position within the window. Consequently, the difference between the target value and the set value can be as large as 3 ns, whilst the delay line could offer a resolution of 0.5 ns. But if a small amount of fluctuation (like the combined jitter of the laser and the measurement unit) is present in the system, the laser pulse may occasionally fall into the neighbouring time windows. If we calculate the average, we can infer the position of the laser pulse within the time window more precisely. This way, the noise inherent in the system can enhance the precision of the control (see Fig 5).

We carried out numerical simulations to assess the performance of the control as a function of the jitter level. In the simulations, we first set the standard deviation of the noise, then chose the value of the target delay randomly, waited until the control approached the target delay sufficiently close, then recorded the differences between the actual delay and the target delay through 5000 iterations. The error of the control was the root mean square of these differences (the square root of the average of the squared values). For a given standard deviation value of the noise, we repeated the process above 100 times (each time choosing a new random value of the target delay) and calculated the average of the errors.

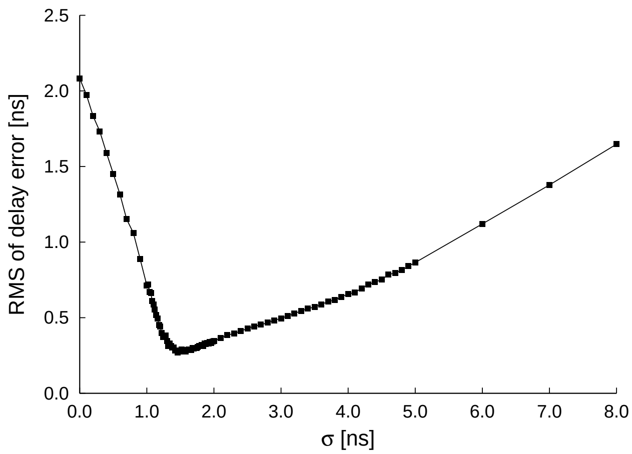

Figure 7. The error of the control as a function of the standard deviation of the noise (σ).

Fig 7 shows how the error of the control depends on the standard deviation of the noise. We can see that the error has a definite minimum around a standard deviation value of 1.5. This minimum is very close to the theoretical minimum that the step size of the control (0.5 ns in our case) allows. It is important to note that this minimum is located at a jitter level typical of the laser system we targeted. Plotting the 'quality factor' (the reciprocal of the error) of the control, we get a familiar profile that reminds one strongly of stochastic resonance (see Fig 8).

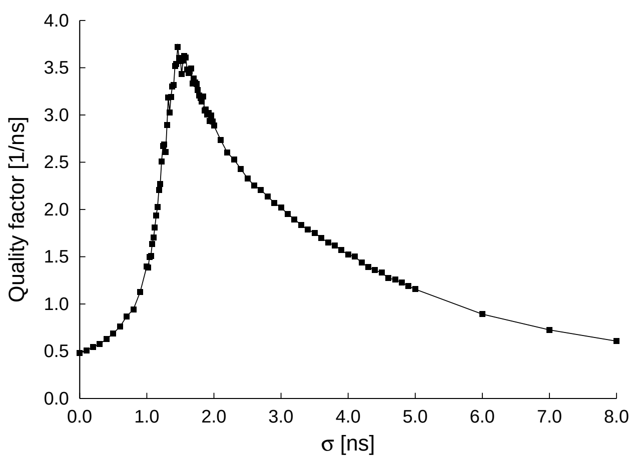

Figure 8. The 'quality factor' (the reciprocal of the error) of the control as a function of the standard deviation of the noise (σ).

The question may arise how we can improve the precision of the control in low-jitter systems. One option is to augment the inherent jitter with external noise fed into the control system so that together they provide the optimal noise level. Tuning the control so that the end of the target delay interval is positioned on the boundary between two time windows instead of the middle of a given time window (see Figs 3 and 6) is the other option. In the latter case, the smallest amount of noise would cause detection instances of laser pulses oscillate evenly between the neighbouring time windows.

From Fig 9 we can see that using this edge-positioned control a small amount of noise is sufficient to approach the target delay. However, there is no stochastic resonance-like effect in this case: the performance gets poorer with increasing noise levels, and from a given standard deviation value, it is worse than that of the mid-positioned setting. We also included the performance of the non-adaptive averaging control in Fig 9. The constructive role of noise is clearly visible even for this simple control principle. Though the non-adaptive averaging control works fairly well at high noise levels, it causes a significant error for small noise standard deviations. The latter is due to the oscillations that may occur in this case.

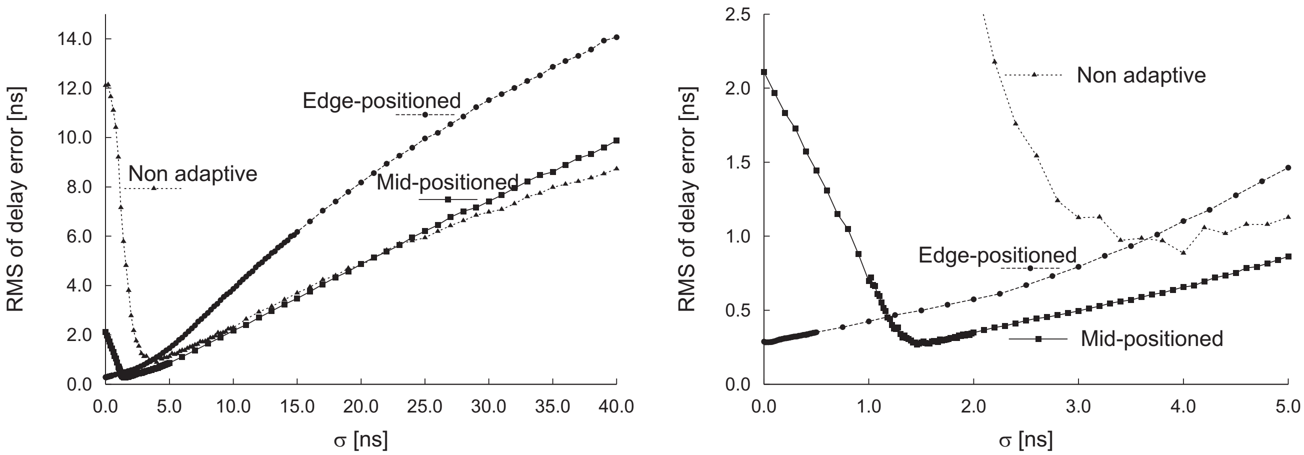

Figure 9. The error of different control modes as a function of the standard deviation of the noise (σ). The right panel shows the area around the optimum of the mid-positioned mode.

If the anode delay of the laser changes in time, the control must follow this. Noise can play a significant role also in how precisely the control is able to track changes. In the next simulation the control must adapt to a sinusoid drift in the anode delay whose amplitude is 40 ns and whose period is 1000 shots. The results are shown in Fig 10. We can see that the edge-positioned control performs much better than the mid-positioned one for low noise levels, whilst for strong jitter its performance is again poorer.

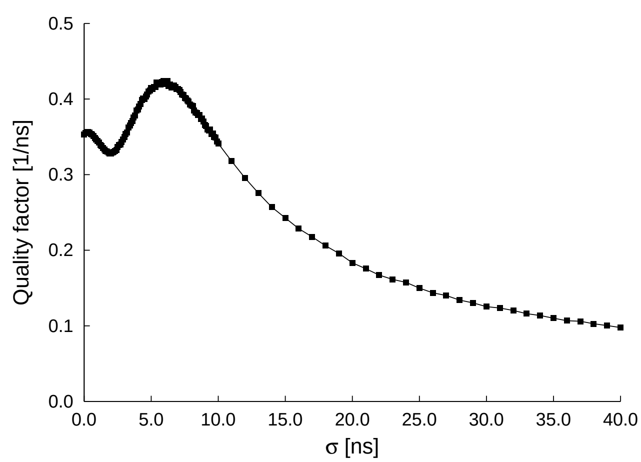

Figure 10. The error of the mid-positioned control when a sinusoid drift is to be tracked. The standard deviation of the noise is denoted by σ.

4. CONCLUSION

In the above, we have presented an electronic system capable of the precise synchronisation of laser pulses along with the corresponding control principle. The control is based on adaptive averaging and utilises the noise inherent in the system, enabling us to set the delay of the laser more precisely than the resolution of the hardware would allow: whilst the width of a detection window is 6 ns, the error of the control at an optimal noise level is around 0.25 ns, close to the theoretical limit determined by the 0.5 ns resolution of the programmable delay.

The control is tunable through a large number of parameters, thus we can tailor the behaviour of the laser system to meet the actual requirements. Depending on the settings, the noise inherent in the system can play different roles; in a mid-positioned setting (when the target delay is set in the middle of a time window), the system shows a stochastic resonance-like behaviour, that is, the error of the control has a minimum at a given non-zero noise level.

The system is also capable of adapting to relatively fast changes; the stochastic resonance-like behaviour can be observed here as well.

ACKNOWLEDGEMENTS

Zoltán Gingl acknowledges support from the Bolyai Fellowship of the Hungarian Academy of Sciences.

REFERENCES

1. R. Benzi, A. Sutera, and A. Vulpiani, "The mechanism of stochastic resonance," *Journal of Physics A* **14**, pp. L453–L457, 1981.
2. R. Benzi, G. Parisi, A. Sutera, and A. Vulpiani, "Stochastic resonance in climatic change," *Tellus* **34**, pp. 10–16, 1982.
3. S. Fauve and F. Heslot, "Stochastic resonance in a bistable system," *Physics Letters* **97A**, pp. 5–7, 1983.
4. B. McNamara, K. Wiesenfeld, and R. Roy, "Observation of stochastic resonance in a ring laser," *Physical Review Letters* **60**, pp. 2626–2629, 1988.
5. D. F. Russell, L. A. Wilkens, and F. Moss, "Use of behavioural stochastic resonance by paddle fish for feeding," *Nature* **402**, pp. 291–294, 1999.
6. I. Hidaka, D. Nozaki, and Y. Yamamoto, "Functional stochastic resonance in the human brain: noise induced sensitization of baroreflex system," *Physical Review Letters* **85**, pp. 3740–3743, 2000.
7. W. Kester, "ADC input noise: the good, the bad, and the ugly. Is no noise good noise?," *Analog Dialogue* **40**, pp. 13–17, 2006.