# An Investigation of Square Waves for Evolution in Carbon Nanotubes Material

Odd Rune Lykkebø[1], Stefano Nichele[1] and Gunnar Tufte[1]
[1]Norwegian University of Science and Technology, Trondheim, Norway
{lykkebo, nichele, gunnart}@idi.ntnu.no

## Abstract

Materials suitable to perform computation make use of evolved configuration signals which specify how the material samples are to operate. The choice of which input and configuration parameters to manipulate obviously impacts the potential of the computational device that emerges. As such, a key challenge is to understand which parameters are better suited to exploit the underlying physical properties of the chosen material. In this paper we focus on the usage of square voltage waves as such manipulation parameters for carbon nanotubes/polymer nanocomposites. The choice of input parameters influence the reachable search space, which may be critical for any kind of evolved computational task. We provide common measurements such as power spectrum and phase plots, taken with the the Mecobo platform, a custom-built board for evolution-in-materio. In addition, an initial investigation is carried out, which links the frequency of square waves to comparability of the output from the material, while also showing differences in the material's physical parameters. Observing the behaviour of materials under varying inputs allows macroscopic modelling of pin-to-pin characteristics with simple RC circuits. Finally, SPICE is used to provide a rudamentary simulation of the observed properties of the material. This simulation models the per-pin behaviours, and also shows that an instance of the traveling-salesman-problem can be solved with a simple randomly generated cloud of resistors.

## Introduction and Background

Evolution-in-Materio (EIM) (Miller et al., 2014), (Miller and Downing, 2002) is a bottom-up approach where the intrinsic underlying physics of materials is exploited as computational medium. In contrast to a traditional design process where a computational substrate, e.g. silicon, is precisely engineered, EIM uses a bottom-up approach to manipulate materials with the aim of producing computation. Such manipulation is done with computer controlled evolution (CCE) (Harding and Miller, 2007), (Harding et al., 2008). CCE may program the materials with different kinds of stimuli, e.g. voltages and currents, temperature, and magnetic fields.

In the NASCENCE project (Broersma et al., 2012), novel nano-scale materials are being used as a substrate in which computation is attempted. In particular carbon nanotubes (CNTs) / polymer have shown promising for the solution of Travelling Salesman (Clegg et al., 2014), logic gates (Kotsialos et al., 2014), and function optimization problems (Mohid et al., 2014). To solve problems, the material is required to hold physical richness (Miller et al., 2014) under a certain manipulation scheme. The method used to manipulate the material into doing computation has until now largely consisted of setting up input signals of different kinds, e.g. static voltages (Clegg et al., 2014), square wave voltages (Lykkebø et al., 2014), a mix of both (Lykkebø and Tufte, 2014), across gold plated connectors that are exposed to the material via a glass plate. Square waves of different frequencies demonstrated potential to achieve a computationally rich behaviour (Nichele et al., 2015). As such, they are the main subject of investigation in this paper.

The hypothesis is that this electrical current exploits properties in the material that gives rise to a potentially useful output, i.e. there is an emergent behaviour emanating from interactions between current travelling through the provably non-linear material (in terms of the current/voltage relation (Massey et al., 2011) and jumps in conductivity under certain temperatures and geometric variation (Ebbesen et al., 1996)) and that measurements of this behaviour can be used for computation.

Is the material performing this computation? It can be very hard to pinpoint what computation actually is and where it happens. One key requirement is Ashby's requisite variety (Ashby, 1956), in which the importance of how many states the system can be in is underlined. In the context of computation in materials, the number of possible input states needs at least to have correspondence in physical states the material can be in. It is worth noting the difference between observable states and unobservable states. The number of observable states can be much lower than the number of states that the material can be in. When manipulating the material with an input, e.g. square waves, the material can iterate through a number of such internal states before eventually settling on the final emergent observable state, which can be in many forms, such as varying voltage peaks or phase offsets. Intuitively, these properties lead to an output of varying complexity, and an intuitively attractive idea is that the measured complexity of the output gives an indication about the internal states of the material as well; if the measured output is far more complex than the inputs, there must necessarily be a mechanism (computation) that gives rise to this complexity.

What is the best way of exploiting such computational properties? Our hypothesis is that materials with high CNT densities may act more as a conductive layer. With lower CNT densities, conductive paths may be closer to the electric percolation threshold of the CNT network (around 1% nanotubes). Another aspect that impact on computational properties is the selected range of frequencies for input signals. Specific frequencies may better penetrate the material as a result of the exploited CNT paths and signal feedbacks. Our aim is to gather basic knowledge of different material's computational properties as to be able to create a simple material model based on RC electrical circuits.

In the experiments herein, material samples with different concentrations of SWCNTs are investigated using square waves at different frequencies. Signal outputs are measured in terms of power spectrum and phase plots. In addition, compressibility of the output signals is investigated for different input frequencies and number of input pins, i.e. number of input frequencies. Such compressibility measure is close to Kolmogorov approximations (Kolmogorov, 1965), (Nichele and Tufte, 2013) of signal complexity.

Observing the response of materials with different concentrations of SWCNTs to given inputs indicate that a simple model of materials can be based on RC circuits. Such a model is developed using SPICE. The presented results provide initial thoughts regarding computational properties of used materials, together with aspects that need to be taken into account when tackling computational problems using evolution-in-materio, i.e. stability, repeatability, and noise.

The article is laid out as follows: Section II describes setup and experimental method. Section III presents the experimental results. In Section IV the modelling using SPICE is detailed and Section V provides a discussion, conclusion and ideas for further work.

## Setup and Methodology

The material used for this work, i.e. carbon nanotubes/polymer nanocomposite, is placed on a microelectrode array on a glass slide which is inserted into the Mecobo board. Each sample has 16 electrodes which are physically connected to the board and can be stimulated by electrical signals, e.g. static voltages, square waves. A sketch of the experimental setup is shown in Figure 1. See (Lykkebø et al., 2014) for more details. To investigate parts of the characteristics of the material in our system under the influence of various frequencies, we run a simple frequency analysis of each material, doing a Fourier transform and find the power spectrum and phase spectrums for 3 different material samples, and at 4 different frequencies: 1KHz, 10KHz, 50KHz, 1MHz (for space reasons and to try a broad range, relative to what the board can output). The tested materials are shown in Table 1.

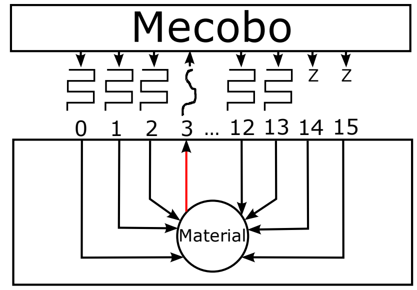


Figure 1: Experimental setup

| Material | SWCNT Concentration | DC resistance |
|---|---|---|
| B15S01 | 0.75% | 20k Ω |
| B15S02 | 1.00% | 5k Ω |
| B15S05 | 1.75% | 1k Ω |

Table 1: Parameters of tested materials

All measurements were taken with the Mecobo board as described in (Lykkebø et al., 2014). Pin 3 is used as sampling pin on all measurements, and the samples were taken at 500KHz for 50 ms, thus collecting 20,000 samples.

For the compression tests (see next chapter for details), pin 3 was used as a sampling pin and input pins *to* the material were added incrementally starting at pin 0 and increasing to pin 15, skipping pin 3. We also increased the number of tested frequencies since these plots were easier to visually compress and present in this paper. The number of samples collected is deterministic. After collecting the samples, they were converted to a binary string and appended to a growing string of bits of total length 800,000 (32 bits per float), which was then passed to the built-in Python 2.7 compress()-function. The length of the compressed buffer was then measured and plotted.

## Results and discussion

This Section outlines two sets of experiments. The first part deals with investigating how materials with different concentrations of SWCNTs behave when exposed to square waves at different frequencies. Output voltages, Fourier transform and phase plots are compared. In the second part, the material samples with different nanotube concentrations are analyzed again in terms of compressibility of output. As such, we may be able to relate input frequencies and output complexity in terms of compressibility for different materials.

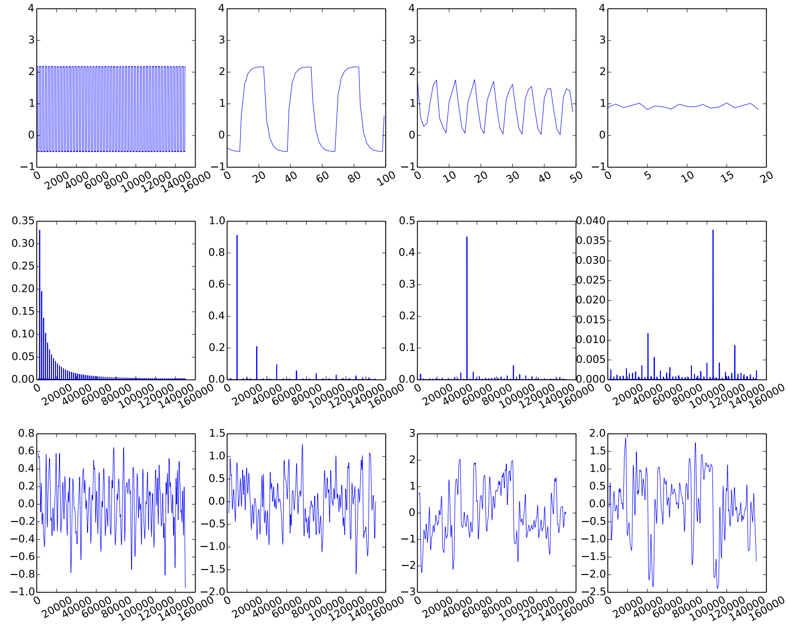

Figure 2: Material B15S01(0.75%) - Each column is a different frequency (1KHz, 10KHz, 50KHz, 1MHz). Row 1; raw voltage/time, row 2; power spectrum, row 3; phase spectrum.

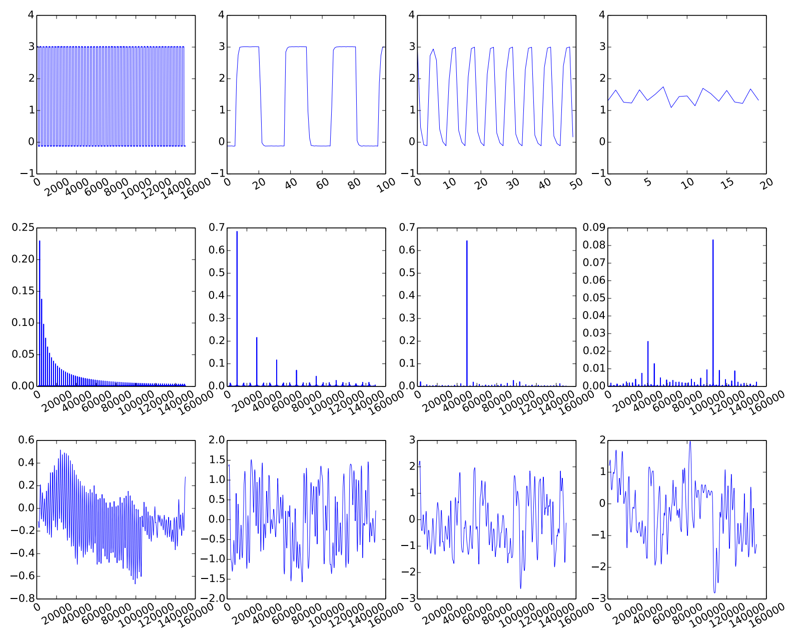

Figure 3: Material B15S02(1.00%) - Each column is a different frequency (1KHz, 10KHz, 50KHz, 1MHz). Row 1; raw voltage/time, row 2; power spectrum, row 3; phase spectrum.

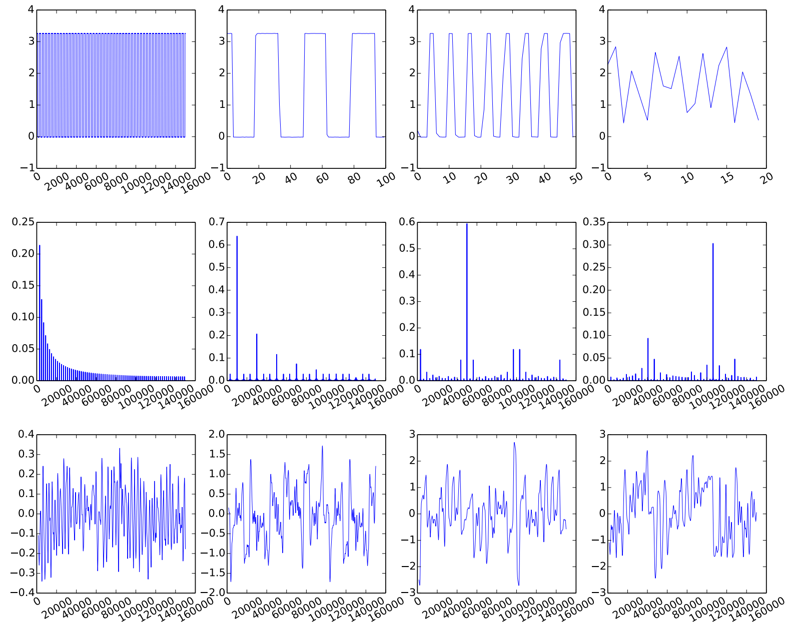

Figure 4: Material B15S05(1.75%) - Each column is a different frequency (1KHz, 10KHz, 50KHz, 1MHz). Row 1; raw voltage/time, row 2; power spectrum, row 3; phase spectrum.

## Basic frequency measurements

We will first discuss Figures 2, 3 and 4. Each column corresponds to a different frequency. As can be seen from the first row in the figures, the peak-to-peak amplitude goes down as we apply higher frequencies. We believe this is the effect of capacitance in the material.

When sweeping square wave frequencies from 1Hz to 1MHz, the material initially passes the signal through (low pass), and eventually starts to let the high frequency through as well (high pass). Applying Occam's razor, the easiest way to explain this behaviour is to assume that the material has a certain amount of capacitance in addition to the easy-to-measure DC resistance. This capacitance makes the material exhibit a charge/discharge cycle when exposed to a time-varying signal such as a square wave. Another possibility is that the inductance of the material is responsible for the filtering effects. Inductance, however, is due to winding of a conductor and the distribution of the CNTs in the PMMA does not immediately appear to be laid out in a such a way as to give rise to this phenomena. Since inductive impedance increases with frequency by $j\omega L$, where L is inductance, a relatively high inductance value is required to follow the curve of the signal, as seen in Figure 12. This is also the reason why the model described below does not include any inductors; they do not seem to contribute any major factor at this observation level and at the chosen range of frequencies to investigate.

Moving on to row 2 in Figures 2, 3, 4, which shows the power spectra for each material reveals a couple of things. First, it should be noted that we have removed the DC Fourier frequency bin from the plots. Secondly, the plots show that the platform is relatively free of obvious noise. The biggest frequency responses come from the applied signals, as expected. Notice by studying the y-axis of the plots that the power goes down significantly as we increase the frequency, as is expected since the peak-to-peak voltage is also rapidly decreasing. The high-frequency plots do seem to have several more frequencies present, but this is simply an artifact of the axis scaling.

The last row in Figures 2, 3 and 4, shows the phase information from the Fourier transform. Although harder to interpret intuitively, the relatively noisy response shows that there are no obvious trends in the phase of the signals across the frequency range, which means that it is little use in trying to infer any computational properties that make use of the phase information.

## Compression tests

The results from the compression test can be seen in Figures 5, 6 and 7. Note that the standard error is quite low. The main observation from these graphs is that as we add more pins that output current *into* the material, the length of the

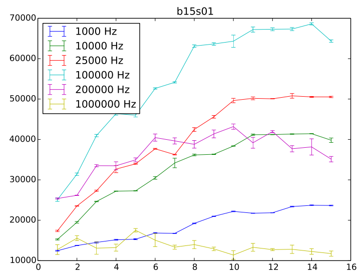


Figure 5: Material B15S01, length of compressed buffer vs. increasing number of input pins enabled.

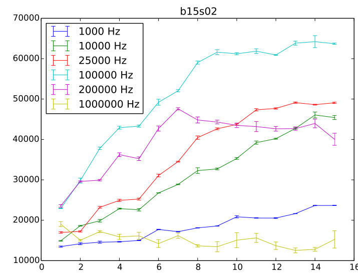


Figure 6: Material B15S02, length of compressed buffer vs. increasing number of input pins enabled.

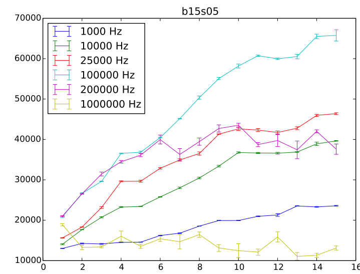


Figure 7: Material B15S05, length of compressed buffer vs. increasing number of input pins enabled.

compressed string rises sharply up to around 8-10 pins for all the materials. A threshold of saturation is further seen after this point. A possible explanation for the shape is that we add more pins, there are phase-offsets between input signals introduced due to imperfect scheduling on the Mecobo platform and potentially small interactions between the signals that give harder-to-compress output. As the number of pins increases, the amount of energy flowing through the system gives a more or less uniformly noisy measured signal. Manual inspection of the measured signal from the material confirms this.

Note that the 100KHz signal is higher than the rest. This is a combination of the measurement rate and bandwidth of the material. There is still relatively low damping seen on the signal at this rate, and 5 samples per cycle are sufficient to capture the main shape of the measured signal, and at the same time there is variation due to the frequency of output. At this frequency there is also ringing due to Gibbs phenomenon visible at high sample rates taken with an oscilloscope, giving rise to even more unpredictability.

The 1000Hz signal is almost undamped in the material, and it is therefore sharp and with a low amount of noise, making it easy to predict. The 1MHz signal is the exact opposite; very high damping ensures that the measured signal is mostly noise which should be relatively constant throughout the measurements, in particular at the sample rate we are limited to.

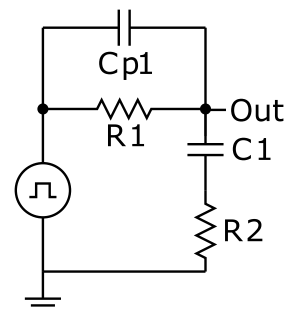


Figure 8: Simulation circuit, the output is measured from GND to the node between R1 and C1.

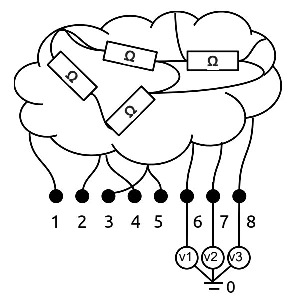


Figure 9: SPICE model used for the TSP experiments. Numbers correspond to SPICE nodes, v1,v2 and v3 are voltage sources. The cloud consists of randomly connected resistors of various values.

## SPICE modelling

SPICE is a circuit simulation tool. By observing the behaviour of the material under certain inputs, it is possible to create a macroscopic 'pin-to-pin' model of a material slide. For each pin pair, we measure the DC resistance and voltage response under various frequencies, and observe the response via an oscilloscope or via the Mecobo platform. This enables us to produce models based on common circuit elements such as capacitors, inductors and resistors that mimic what we observe at this level. A practical use for SPICE models is often to tune circuits in such a way that one obtains so-called *impedance matching*, in which the impedance seen by the observing element matches its own internal impedance. This allows for maximum power in the transferred signal. In the context of evolution-in-materials, the model can be used for such a purpose as well; since the materials provide a broad range of 'loads', we can use a model of the macroscopic properties to tune our measurement circuitry in such a way as to increase our chances of

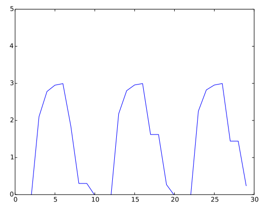

Figure 10: Mecobo capture of voltages for a 50KHz signal at 500KHz sample rate.

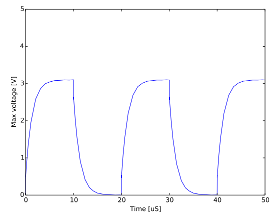


Figure 11: Spice output voltages for a 50KHz signal, simulated at 10ns timesteps.

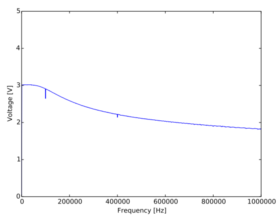


Figure 12: Mecobo capture of maximum voltage per frequency.

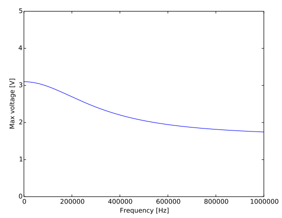


Figure 13: Spice output of maximum voltage per frequency.

finding emergent behaviours resulting from microscopic interactions in the material.

The model required to capture the effects of square waves seen at the level we are observing on can be simple. As discussed in the previous section, we believe that the main contributor to the filtering effects can be modelled by a single capacitor, and this is reflected in Figure 8. By doing these measurements on all pin pairs and inserting one of these models between each pin pair, a model of the complete material slide can be constructed.

The model for this can be seen in Figure 8. This circuit is an RC low pass filter (R1 and C1) connected together with an RC high-pass filter (R2 and C2), along with a parasitic capacitance in parallel with the 'main' resistor (Cp1). The DC load of the system is captured by R1. The parasitic capacitance is added to model ringing, either due to Gibbs phenomenon (Hewitt and Hewitt, 1979) or simply parasitic capacitances and inductances resonating at their characteristic frequency). We observe these effects when applying certain frequencies to the material, which also can be barely seen in Figure 12 for the captured signal and Figure 10 for the model, albeit not to a large extent. In both cases a signal of 50KHz has been applied to the material. The signal output from the circuit is measured over C1 and R2, and the input is modelled by a pulse train voltage source. Looking at Figures 10 and 11, it is obvious that the signal captured with Mecobo (which used a sample rate of 500KHz), does not reproduce all the effects seen at simulation level, such as the ringing present due to Gibbs phenomenon and possibly other small parasitic inductances and voltages, both in the material and the measurement apparatus. We therefore used an oscilloscope to capture a more detailed version of the signal, in which more of the effects mentioned above can be observed. This begs a deeper question: are there events in the material occurring that we fail to capture with our measurements? Take inductance, as mentioned above for instance. Since there is current flowing through the material (the conductor), there must also be a magnetic field associated with this moving current. Isolated it is likely very small, but the sum of the fields could be measurable with Hall effect detectors (Ramsden, 2006), which could lead us to further improve the model by adding the correct inductances as to mirror the effects seen.

## Travelling salesman

A further experiment was set up to attempt to solve the travelling salesman solved by Clegg in (Clegg et al., 2014) using a relatively simple SPICE model. In essence the problem as presented condenses into finding a way to configure the material into settling on a number voltages of different values. Each output pin is tagged with a city and the goal is to have the sorted values of the voltages on the output pins correspond to the shortest path between these cities, e.g. if pins 1, 2, 3 corresponds to the cities, and the shortest path between them is 2, 3, 1, the requirement is that $V(pin2) < V(pin3) < V(pin1)$. For details regarding the problem definition, we defer the reader to (Clegg et al., 2014).

Figure 9 shows our SPICE-model of this problem. We generate a 'cloud' consisting of nothing but resistors and connections (or 'nodes' in SPICE-terms) between them. An initial study of the problem revealed that we did not need to include capcitors to solve the problem. The first $k$ SPICE-nodes (1 to 5 in the figure) are used as 'cities' and are only used as voltage measurement points. The remaining pins are connected to voltage sources whose voltages are adjusted with a (1+4) evolutionary strategy. The resistor-cloud-generating procedure uses an adjustable number of internal nodes and an adjustable number of generated resistors, giving us the ability to adjust the density of the cloud with these two parameters. In addition, the values of the resistors are drawn from a uniform distribution in a range also passed to the procedure, which enables tuning of the 'point-to-point' resistance of the network.

The instance we are seeking to provide a solution for is one of the presented solutions in (Clegg et al., 2014); a circular arrangement of 8 cities in which the shortest path is simply that– a circle. The cities are arranged on a grid, and the path length of the suggested solutions (as described above) is measured as a real number and taken as the fitness.

The vector being optimized is a 4-tuple consisting of 4 floating point values, $(V_1, V_2, V_3, V_4)$, which maps directly to the voltages set on the voltage sources. Mutation is done by adding a number drawn from a standard normal distribution with $\mu = 0.0, \sigma = 2.0$ to one of the elements of the tuple. The SPICE simulation is run for 20ms using transient analysis with a resolution of 0.01ms. The final value of the measurements is the arithmetic average of the time series as measured across the nodes labeled as cities (i.e. nodes 1 to 8) and node 0 (which corresponds to ground).

*We find that solutions are relatively quickly obtained (within 100 generations) in some of the randomly generated networks, while in others the search settles into a local optima and is unable to escape before the termination of the search which happens at 2000 generations.*

This led us to investigating a few different parameter settings of the resistor cloud, the results of which can be seen in table 2. We generated 1000 networks for each parameter setting and counted the number of perfect solutions within 100 generations. As can be seen, there are relatively few solutions in all cases. Since we spent no significant time evaluating the different combinations, this is to be expected– the only conclusion that can be drawn from these numbers is that the composition of the resistor cloud matters when solving this type of problem. The somewhat unsurprising conclusion is that the composition of the materials used in the 'real' evolution in materials matters.

| **Nodes** | **Resistors** | **Solutions** |
|---|---|---|
| 35 | 150 | 30 |
| 35 | 300 | 19 |
| 35 | 500 | 32 |
| 60 | 150 | 69 |
| 60 | 300 | 11 |
| 60 | 500 | 34 |

Table 2: Results from running a search for a TSP solution in 1000 randomly generated resistor cloud networks.

## Discussion and conclusions

How should we view the results in (Lykkebø et al., 2014), (Mohid et al., 2014) in which square waves are used as one of the configuration parameters then? The only observable property of the material pinpointed thus far is the fact that it will act as low pass/high pass signal filter. The different pin pairs thus act like 'frequency selectors' and it is possible to think of ways to do a form of computation with these (a frequency discriminator (Thompson, 1997) for instance), which could provide a mechanism that an evolutionary algorithm could potentially find and make use of. Another explanation that needs to be investigated further is the very real possibility that the fitness functions have been defined in such a way that they can use *noise* to solve the problem, and that lack of re-evaluations in the experiments lead to solutions that worked once, but only by chance; by encoding the solution to the problem in the input and simply evolving a way to reconstruct the solution from the input and the noise added by the material.

As for configuration of the material using static DC voltages, we have shown that the material in this case behaves mostly as a resistor network. As we have shown in our simple SPICE model shown in Section IV it is indeed possible to solve the problem instance solved by Clegg in (Clegg et al., 2014) using a network of resistors. One important point to note is that the material is rich and more flexible than one single purpose-built resistor network; it contains a large amount of various such networks, each of which can be exploited by evolution to solve a number of problems.

An issue regarding this way of formulating the problem is that we are limited by the resolution of the measurements. A common ADC such as the one used in the Mecobo daughterboard, has a resolution of 12 bits, giving a maximum of $2^{12}$ different voltages, effectively limiting the number of cities to 4096. This is the best case, often times one can only hope to achieve 10-11 bits of resolution, depending on how well-matched the impedance of the load is to the measurement apparatus, which again is affected by the characteristics of the material which can vary by large amounts, as we have seen previously in this paper. One can of course invest in even higher-resolution measurement apparatus to achieve well over 20 bit resolution but the point of diminishing returns in terms of practicality of building a computing system seems likely to be reached quite fast by making further advancements in this.

For the results in (Kotsialos et al., 2014) we suggest that it might be hard to reproduce these results as well. Making a XOR gate out of a resistor network is not possible since a way of inverting a signal is needed; but it is not unlikely that it is possible to achieve with a noisy system, as there are an abundant amount of transistors, diodes and let's not forget, *environmental* noise readily available in a physical system. This of course leads us into the question of where to draw the line between the computational entity, the measurement apparatus and the input. For more discussion around this, see (Lykkebø et al., 2014).

As such, we conclude that future work using the SWCNTs thus needs to take particular care with proving that it is 1) better than a random search and 2) that the results obtained are repeatable with a high degree of confidence. One way of achieving this would be to operate *within the bandwidth* of the material, such that we can be sure we are measuring actual signal response and not noise. The second, obvious way, is to always discard unstable solutions if they fail to reproduce their behaviour a number of times. A third point to note is that we must construct the fitness functions in such a way as to minimize the evolutionary process' natural tendency to exploit unwanted effects such as noise.

Much work remains before we can draw final conclusions regarding the computational properties of random carbon nanotubes. One dimension that is immediately interesting in terms of richness is movability of the material, since geometric properties play an important role in solid state physics. Introducing a more viscous environment for the CNTs to move around in could prove fruitful, since there is evidence that the tubes are capable of self-organizing (Belkin et al., 2015), and that geometry matters (Ebbesen et al., 1996) and just as a human designer is free to exploit the spatial properties of electro-material interactions, so is it possible for an evolutionary process to do the same.

**Acknowledgements**

The research leading to these results has received funding from the [European Community's] Seventh Framework Programme ([FP7/2007-2013] [FP7/2007-2011]) under grant agreement no [317662].

## References

Ashby, W. R. (1956). *An introduction to cybernetics*. Chapman & Hall, London.

Belkin, A., Hubler, A., and Bezryadin, A. (2015). Self-assembled wiggling nano-structures and the principle of maximum entropy production. *Sci. Rep.*, 5.

Broersma, H., Gomez, F., Miller, J. F., Petty, M., and Tufte, G. (2012). Nascence project: Nanoscale engineering for novel computation using evolution. *International Journal of Unconventional Computing*, 8(4):313–317.

Cariani, P. (1993). To evolve an ear: epistemological implications of gordon pask's electrochemical devices. *System Research*, 10(3):19–33.

Clegg, K. D., Miller, J. F., Massey, K., and Petty, M. (2014). Travelling salesman problem solved 'in materio' by evolved carbon nanotube device. In Bartz-Beielstein, T., Branke, J., Filipič, B., and Smith, J., editors, *Parallel Problem Solving from Nature – PPSN XIII*, volume 8672 of *Lecture Notes in Computer Science*, pages 692–701. Springer International Publishing.

Ebbesen, T. W., Lezec, H. J., Hiura, H., Bennett, J. W., Ghaemi, H. F., and Thio, T. (1996). Electrical conductivity of individual carbon nanotubes. *Nature*, 382(6586):54–56.

Harding, S. L. and Miller, J. F. (2004). A tone discriminator in liquid crystal. In *Congress on Evolutionary Computation(CEC2004)*, pages 1800–1807. IEEE.

Harding, S. L. and Miller, J. F. (2007). Evolution in materio: Computing with liquid crystal. *Journal of Unconventional Computing*, 3(4):243–257.

Harding, S. L., Miller, J. F., and Rietman, E. (2008). Evolution in materio: Exploiting the physics of materials for computing. *Journal of Unconventional Computing*, 3:155–194.

Hewitt, E. and Hewitt, R. (1979). The gibbs-wilbraham phenomenon: An episode in fourier analysis. *Archive for History of Exact Sciences*, 21(2):129–160.

Kolmogorov, A. N. (1965). Three approaches to the quantitative definition of information. *Problems of Information Transmission*, 1:1–7.

Kotsialos, A., Massey, M., Qaiser, F., Zeze, D., Pearson, C., and Petty, M. (2014). Logic gate and circuit training on randomly dispersed carbon nanotubes. *International journal of unconventional computing.*, 10(5-6):473–497.

Lykkebø, O. R., Harding, S., Tufte, G., and Miller, J. F. (2014). Mecobo: A hardware and software platform for in materio evolution. In Ibarra, O. H., Kari, L., and Kopecki, S., editors, *Unconventional Computation and Natural Computation*, Lecture Notes in Computer Science, pages 267–279. Springer International Publishing.

Lykkebø, O. R. and Tufte, G. (2014). Comparison and evaluation of signal representations for a carbon nanotube compuational device. In *Evolvable Systems (ICES), 2014 IEEE International Conference on*, pages 54–60.

Massey, M. K., Pearson, C., Zeze, D. A., Mendis, B. G., and Petty, M. C. (2011). The electrical and optical properties of oriented langmuir-blodgett films of single-walled carbon nanotubes. *Carbon*, 49:2424.

Miller, J. F. and Downing, K. (2002). Evolution in materio: Looking beyond the silicon box. In *2002 NASA/DOD Conference on Evolvable Hardware*, pages 167–176. IEEE Computer Society Press.

Miller, J. F., Harding, S., and Tufte, G. (2014). Evolution-in-materio: evolving computation in materials. *Evolutionary Intelligence*, 7(1):49–67.

Mohid, M., Miller, J. F., Harding, S. L., Tufte, G., Lykkebo, O. R., Massey, M. K., and Petty, M. C. (2014). Evolution-in-materio: Solving function optimization problems using materials. In *Computational Intelligence (UKCI), 2014 14th UK Workshop on*, pages 1–8.

Nichele, S., Laketic, D., Lykkebø, O. R., and Tufte, G. (2015). Is there chaos in blobs of carbon nanotubes used to perform computation? In *7th International Conference on Future Comp. Tech. and Applications, IN PRESS*. XPS Press.

Nichele, S. and Tufte, G. (2013). Measuring phenotypic structural complexity of artificial cellular organisms - approximation of kolmogorov complexity with lempel-ziv compression. In *Innovations in Bio-inspired Computing and Applications - Proceedings of IBICA 2013, August 22 -24, 2013 - Ostrava, Czech Republic*.

Pask, G. (1959). Physical analogues to growth of a concept. *Mechanisation of Thought Processes*, pages 877–922.

Ramsden, E. (2006). *Hall-Effect Sensors - Theory and Application (2nd Edition)*. Elsevier.

Thompson, A. (1997). An evolved circuit, intrinsic in silicon, entwined with physics. In *1st International Conference on Evolvable Systems (ICES96)*, Lecture Notes in Computer Science, pages 390–405. Springer.