

A reprint from

# American Scientist

the magazine of Sigma Xi, The Scientific Research Society

This reprint is provided for personal and noncommercial use. For any other use, please send a request Brian Hayes by electronic mail to [bhayes@amsci.org](mailto:bhayes@amsci.org).

# The Memristor

Brian Hayes

WHEN BELL TELEPHONE Laboratories announced the invention of the transistor in 1948, the press release boasted that “more than a hundred of them can easily be held in the palm of the hand.” Today, you can hold more than 100 billion transistors in your hand. What’s more, those transistors cost less than a dollar per billion, making them the cheapest and most abundant manufactured commodity in human history. Semiconductor fabrication lines churn out far more transistors than the world’s farmers grow grains of wheat or rice.

In this thriving transistor monoculture, can a new circuit element find a place to take root and grow? That’s the question posed by the memristor, a device first discussed theoretically 40 years ago and finally implemented in hardware in 2008. The name is a contraction of “memory resistor,” which offers a good clue to how it works.

Memristor enthusiasts hope the device will bring a new wave of innovation in electronics, packing even more bits into smaller volumes. Memristors would not totally supplant transistors but would supplement them in computer memories and logic circuits, and might also bring some form of analog information processing back into the world of computing. Farther out on the horizon is a vision of “neuromorphic” computers, modeled on animal nervous systems, where the memristor would play the role of the synapse.

Whether the memristor will ultimately fulfill all these hopes remains to be seen. The history of invention is littered with promising novelties that failed to dislodge an incumbent tech-

*The first new  
passive circuit element  
since the 1830s  
might transform  
computer hardware*

nology. On the other hand, there is now widespread agreement that *some* fundamental shift in circuit design will be needed if computer hardware is to remain a growth industry. The memristor looks like a strong candidate.

## Titania

The device that has sparked all the recent excitement over memristors was created in 2008 by R. Stanley Williams and several colleagues at Hewlett-Packard Laboratories. The Williams memristor consists of two metal electrodes separated by a thin film of titanium dioxide, or TiO<sub>2</sub>. This substance, also known as titania, is familiar to artists as a white pigment and to beachgoers as an ingredient of sunscreen.

In its natural form titania is an electrical insulator, presenting very high resistance to the flow of electric current. In the memristor, part of the titania layer retains this natural insulating character, but the rest is altered during deposition by restricting the amount of oxygen available. The resulting oxygen vacancies in the crystal lattice reduce the resistance of the material by supplying mobile electrons that can carry a current. The oxygen-starved layer is said to be “doped.” (This term usually refers to *added* impurity atoms, but the effect of the oxygen deficiency is the same.)

An electric current passing through the memristor has to cross both the

doped and the undoped regions, and so the total resistance is the sum of contributions from the two layers. The total depends on the relative thickness of the layers, or in other words on the position of the boundary between them. What gives the memristor its special traits is that this boundary can move.

Consider what happens inside the titania film when a voltage is applied to the terminals of the memristor, so that a current flows through it. The current is carried by conduction electrons—mostly electrons liberated by the oxygen vacancies. Electrons have a negative charge, and so they are repelled by the negative terminal and attracted to the positive one. In the background, meanwhile, another process is going on. The oxygen vacancies also have an electric charge; they act as positive ions, which drift toward the negative electrode. Movement of the vacancies requires physical rearrangement of the crystal lattice, and so it is much slower than the flow of electrons.

The relatively slow drift of the oxygen vacancies makes no significant contribution to the electric current flowing through the memristor, but by shifting the boundary line between doped and undoped layers, it alters the overall resistance of the device. Depending on the polarity of the applied voltage, the resistance can either increase (if the doped region is squeezed into a narrower layer) or decrease (if the doped region expands to include more of the total thickness). When the external voltage is removed, the boundary line stays put in its new position.

It is the migrating boundary between doped and undoped regions that gives the memristor its memory. And it’s not hard to see how this property can be put to work for information storage. One simple scheme defines a low-resistance state as a binary 0 and high-resistance state as a binary 1. To write a bit into the memory cell,

Brian Hayes is senior writer for American Scientist. Additional material related to the Computing Science column appears at <http://bit-player.org>. Address: 11 Chandler St. #2, Somerville, MA 02144. E-mail: [brian@bit-player.org](mailto:brian@bit-player.org)

apply a strong voltage pulse of the appropriate polarity, thereby setting the resistance either high or low. To read the stored state of the cell, use a lower voltage or a briefer pulse, which can measure the resistance without appreciably altering it.

A notable advantage of the memristor is that it can be made very small. As a matter of fact, it *must* be small, at least along one dimension—the thickness of the TiO<sub>2</sub> film. The ratio of maximum to minimum resistance varies inversely as the square of this thickness. In practical devices the film might be as thin as 10 nanometers, which is just 25 or 30 atomic diameters.

It's also notable that the memristor offers *nonvolatile* storage: The device retains its memory even when the power is turned off.

### Lost and Found

There is a long tradition of explaining electric circuits by hydraulic analogies. Thus a conductor is compared to a pipe; electric current is analogous to the flow of water through the pipe; and voltage is like the pressure difference that drives the flow. In this imaginary world of electrical plumbing, a resistor is a small orifice that restricts the flow of water through a pipe. Similarly, a diode (or rectifier) might be likened to a one-way check valve, with a flap that the water pushes open when flowing in the right direction; pressure in the opposite direction closes the flap, which prevents any flow.

What is the hydraulic equivalent of a memristor? The closest analogy I can think of is a sand filter, an item of apparatus used in water-purification plants. As contaminated water flows through a bed of sand and gravel, sediment gradually clogs the pores of the filter and thereby increases resistance. Reversing the flow flushes out the sediment and reduces resistance. Note that this behavior differs from that of a check valve. Although in both cases the direction of flow is what controls the state of the device, at any given instant the resistance of the sand filter is the same in both directions. The memristor, too, is symmetric in this sense.


Plumbing analogies offer intuition about how a component works, but engineers need more—they need a predictive mathematical theory. The memristor has such a theory. It was formulated by Leon O. Chua of the University of California, Berkeley, in the

early 1970s—when he had no physical device to which the theory applied. Chua's 1971 paper on the subject was titled "Memristor—the missing circuit element." Williams and his colleagues titled their 2008 announcement "The missing memristor found."

Chua's theory has nothing to say about oxygen vacancies or other details of materials and structures. It is framed in terms of the basic equations of electric circuits. Those equations link four quantities: voltage ( $v$ ), current ( $i$ ), charge ( $q$ ) and magnetic flux ( $\phi$ ). Each equation establishes a relation between two of these variables. For example, the best-known equation is Ohm's Law,  $v = Ri$ , which says that voltage is proportional to current, with the constant of proportionality given by the resistance  $R$ . If a current of  $i$  amperes is flowing through a resistance of  $R$  ohms, then the voltage measured across the resistance will be  $v$  volts. A graph of current versus voltage for an ideal resistor is a straight line whose slope is  $R$ .

Equations of the same form but with different pairs of variables describe two more basic electrical properties, capacitance and inductance. And two more equations define current and voltage in terms of charge and flux. That makes a total of five equations, which bring together various pairings of the four variables  $v$ ,  $i$ ,  $q$  and  $\phi$ . Chua observed that four things taken two at a time yield *six* possible combinations, and so a sixth equation could be formulated. The missing equation would

The memristor is a new building block for electrical circuits, an addition to the family of "passive" devices that also includes the resistor, the capacitor and the inductor. A version of the memristor invented at Hewlett-Packard in 2008 has a layer of insulating titanium dioxide, TiO<sub>2</sub>, sandwiched between two platinum (Pt) electrodes. Part of the TiO<sub>2</sub> layer is "doped": It has oxygen vacancies (orange disks) that act as positive ions and liberate electrons (purple dots) that can carry an electric current. The overall resistance of the device depends on the position of the boundary between the doped and the undoped regions; moreover, this boundary can be moved by an applied electric field. The diagrams at right show several possible states of the memristor. Proceeding upward from the middle of the sequence, a forward bias voltage expands the doped region and thereby lowers the resistance; proceeding downward from the middle, a reverse bias voltage leaves most of the layer undoped, raising the resistance.



connect charge  $q$  and magnetic flux  $\varphi$  and would describe a new circuit element, joining the resistor, the capacitor and the inductor. Those three devices had all been known since the 1830s, so the new element would be a very late and unexpected addition to the family. Chua named it the memristor.

No law of physics demanded that such a device exist, but no law forbade it either; the existing theory of circuits with resistance, capacitance and inductance could be augmented in a straightforward way to include memristance as well. Chua argued for the plausibility of the memristor on grounds of symmetry and completeness, suggesting an analogy with Dmitri Mendeleev's construction of the periodic table. Nature is not required to fill every square of this table, but a blank spot is certainly a good place to look for a new chemical element—or a new circuit element.

What would a device linking charge and flux look like? Framing the question in this way may be part of the reason it took so long to identify a physical memristor. The variables  $q$  and  $\varphi$  invite visions of electric and magnetic fields interacting in some conspicuous way. But the memristor invented at Hewlett-Packard has no obvious con-

nection with magnetic phenomena. Instead it works as a special kind of variable resistor. How can this device be described in terms of  $q$  and  $\varphi$ ?

Chua's answer is that  $q$  and  $\varphi$  are more important as mathematical variables than as physical quantities. The charge  $q$  is the time integral of an electric current: The current is a rate of flow—the number of electrons per second passing some point in the circuit—whereas the charge is the total number of electrons passing that point. A similar relation defines voltage in terms of magnetic flux. By making use of these definitions, we can describe the action of the memristor in terms of voltage and current instead of charge and flux.

The simplest form of the memristor equation is just a variant of Ohm's Law:  $v = M(q)i$ . Where Ohm's Law has a constant,  $R$ , representing resistance, the memristor equation has a function,  $M(q)$ .  $M$  is not a constant; instead it varies as a function of the quantity of charge that has passed through the device. This functional dependence allows memristance to be controlled in ways that ordinary resistance cannot. (Nevertheless, memristance is expressed in the same unit of measure as resistance, namely the ohm.)

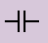
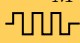

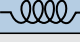
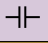


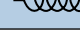
Long before Williams announced the  $\text{TiO}_2$  memristor, there were reports of "anomalous" resistance effects that can now be understood in terms of memristance. Chua has compiled a list of examples going back to 1976, and Williams himself had been exploring such phenomena since 1997. What changed in 2008 was the recognition that Chua's memristor theory could be applied to these devices. The connection between theory and experiment is more than a formality; it allows memristors to be modeled in circuit-simulation software, an essential in the design of large-scale systems.

### Hysteresis and Memory

The resistor, the capacitor, the inductor and the memristor are all described as "passive" circuit elements, to distinguish them from "active" devices such as transistors, which can amplify signals and inject power into circuits. But the memristor differs from the other passive components in a crucial way: It is necessarily a nonlinear device. In an ideal resistor, as mentioned above, the relation between current and voltage is one of simple proportionality, and so the graph of this relation is a straight line of slope  $R$ . The equivalent graph for an ideal memristor is not a line but a curve, where the slope varies from place to place.

In the  $\text{TiO}_2$  memristor it's easy to see where the nonlinearity comes from. Suppose the device is connected to a source of constant voltage. As current passes through the memristor in the "forward" direction—enlarging the conductive, doped, layer—the memristance decreases; this allows more current to pass, which further reduces the memristance, and so on. Reversing the polarity of the voltage source leads to the opposite kind of feedback loop, where increasing memristance causes still further increases.

The nature of the nonlinearity can be seen clearly by tracing the response of the device to a sinusoidal signal—a smoothly alternating voltage. The plot starts at zero volts and zero amperes. As the voltage steadily increases, so does the current, at an accelerating rate reflecting the nonlinear memristance. Then, after the voltage reaches its maximum and starts to fall again, the current continues to rise briefly because the resistance of the  $\text{TiO}_2$  film is still diminishing. When the current finally does retreat, the descending branch of

	charge $q$	current $i$	voltage $v$	magnetic flux $\varphi$
charge $q$		$q = \int i dt$	capacitance $q = Cv$ 	memristance $q = \frac{\varphi}{M}$ 
current $i$	$i = \frac{dq}{dt}$		resistance $i = \frac{v}{R}$ 	inductance $i = \frac{\varphi}{L}$ 
voltage $v$	capacitance $v = \frac{q}{C}$ 	resistance $v = Ri$ 		$v = \frac{d\varphi}{dt}$
magnetic flux $\varphi$	memristance $\varphi = Mq$ 	inductance $\varphi = Li$ 	$\varphi = \int v dt$	

Equations of circuit theory led to a prediction of the memristor almost 40 years before the device was discovered. The equations state relations among four variables: charge ( $q$ ), current ( $i$ ), voltage ( $v$ ) and magnetic flux ( $\varphi$ ). Taking these variables in pairs, there are six possible combinations, but only five equations were known. Leon O. Chua of the University of California, Berkeley, showed that the missing sixth equation, which links  $q$  and  $\varphi$ , defines the property he named memristance. In this matrix each equation appears in two forms, one the inverse of the other. Four paired equations (colored squares) are associated with basic circuit elements: resistance, capacitance, inductance and memristance. The remaining equations (gray squares) define charge as the time integral of current and voltage as the time derivative of flux.


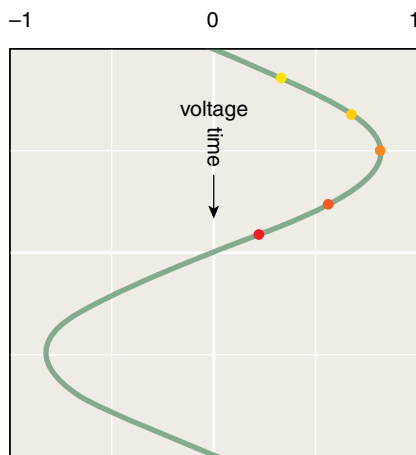
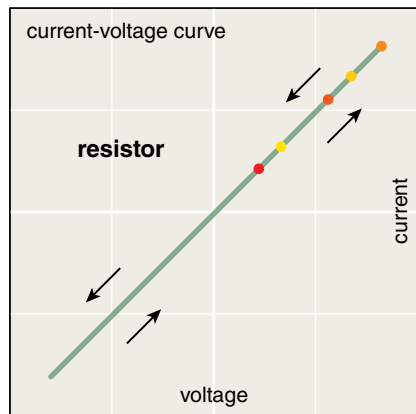
the curve does not retrace the path of the ascending branch. Instead it forms a loop, called a hysteresis loop (a term borrowed from the study of magnetic systems). Specifically, the memristor's curve is a "pinched" hysteresis loop, because the two branches cross at the origin. It's a characteristic of the memristor that whenever the voltage is zero, so is the current, and vice versa. This fact implies that the memristor stores no energy, not even briefly. (The same is true of resistors, but not of capacitors or inductors.)

Hysteresis creates a fundamental distinction between resistors and memristors. In a resistor, current is a simple, single-valued function of voltage; the same voltage always elicits the same current. The hysteresis curve of a memristor driven by a sinusoidal input signal implies that the same voltage can yield two different currents. More generally, when we consider inputs other than simple sine waves, a given voltage can correspond to *many* different values of current. What value is observed depends on the internal state of the memristor, which in turn depends on its history. This is just another way of saying that the memristor retains a memory of its own past.

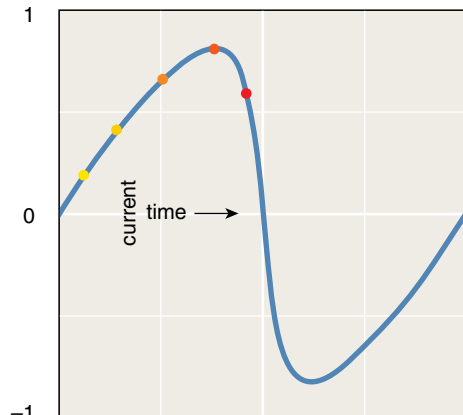
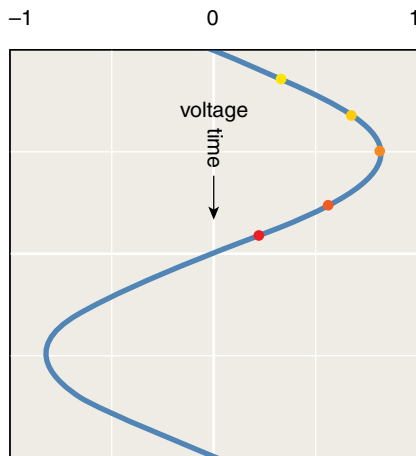
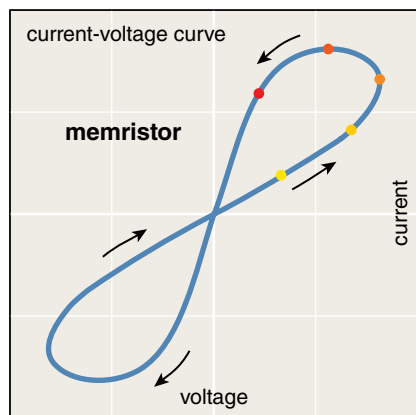
### Switching in Time

The transistor is a three-terminal device, with three connections to the rest of the circuit. It acts as a switch or amplifier, with a voltage applied to one terminal controlling a current flowing between the other two terminals. No such design is possible with memristors, which have just two terminals. But memristors can nonetheless be used to build both memory and digital logic. The key is to exploit the memristor's built-in sense of history: A signal applied at one instant can affect another signal that later travels the same path. The first signal exerts this control by setting the internal state of the memristor to either high or low resistance. The first signal exerts this control by setting the internal state of the memristor to either high or low resistance.

The favored layout for memristor memory is a crossbar structure, where perpendicular rows and columns of fine metal conductors are separated by a thin, partially doped layer of TiO<sub>2</sub>. In this way a memristor is formed at every point where a column crosses a row. Each bit in the memory is individually addressable by selecting the correct combination of column and row conductors. A signal pulse applied to



Plotting the relation of current to voltage highlights the difference between a resistor (left and above) and a memristor (below). In these graphs, voltage and current are each plotted separately as a function of time, and then combined in a current-voltage curve that shows the evolving state of the system. (The separate voltage and current graphs are turned 90 degrees to each other so that their axes match those of the current-voltage graph.) The input voltage is a sine wave. For an ideal resistor, current is proportional to voltage, and so the current is also sinusoidal. The current-voltage curve is a straight line whose slope is the resistance, a constant.



The memristor current-voltage curve takes the form of a "pinched hysteresis loop." Because of the shifting boundary between doped and undoped regions, memristance is not a constant. Current initially grows slower than voltage, then speeds up and continues increasing even after voltage has reached its peak. The loop closes again at the origin: Whenever the voltage is zero, so is the current. The colored dots are meant to help in tracing the trajectory of the system along the current-voltage curve. They identify five corresponding points in the first half-cycle of the sine wave; the sequence progresses from yellow through orange to red.

these conductors can write information by setting the resistive state of the TiO<sub>2</sub> junction. A later pulse on the same pair of conductors reads the recorded information by measuring the resistance.

A near-term role for memristor crossbar arrays might be as a competitor for “flash” memory, the nonvolatile storage technology used in cell phones, cameras and many other devices. Each cell in a flash memory is a single transistor, modified for long-term storage of electric charge. The memristor structure is simpler and requires only two connections, so it might be made smaller than the flash-memory transistor. Thus there’s the possibility of higher density and lower cost.

Building logic circuits out of memristors would be a somewhat greater departure from current practice. In the early years of solid-state electronics, a technology called resistor-transistor logic had a brief vogue; the idea was to minimize the number of expensive transistors and maximize the number of cheap resistors. But with the coming of integrated circuits the economic incentives changed. With potential advantages in size and power consumption, memristors might shift the balance back toward a technology that combines active and passive devices. Williams and his colleagues have demonstrated a set of memristor logic gates that are computationally complete—they can implement any Boolean logic function.

Active components such as transistors would still be needed even if most information processing were done by memristors. The reason is that signals are reduced in amplitude by every passive circuit element, and at some point they must be restored to full strength. This requires a transistor or some other active device. Methods of fabricating hybrid circuits that combine transistors and memristors on the same substrate are an active area of investigation.

In binary digital circuits, memristors would operate as switches, toggling between maximum and minimum resistance. In this mode, the state of a memristor encodes one bit of information. If several intermediate resistances could be distinguished reliably, then the information density could be raised to two or three bits per device. The writing and reading processes would have to be calibrated to resolve four or eight levels of resistance. (Some flash memory chips already achieve this.) The end point of this evolution is to let the resis-

tance vary continuously, and operate the memristor as an analog device.

One intriguing way to exploit analog memristors would be to build a machine modeled on the nervous system. In biological neural networks, each nerve cell communicates with other cells through thousands of synapses; adjustments to the strength of the synaptic connections is thought to be one mechanism of learning. In an artificial neural network, synapses must be small, simple structures if they are to be provided in realistic numbers. The memristor meets those requirements. Moreover, its native mode of operation—changing its resistance in response to the currents that flow through it—suggests a direct way of modeling the adjustment of synaptic strength.

### Empire Building

Will the memristor turn out to be a transformative technology, the key to putting hundreds of *trillions* of devices in the palm of your hand? Or will we be asking, a few years from now, “Whatever happened to the memristor?”

The empire of the transistor has fended off many other rivals and would-be invaders. A memory technology based on magnetic “bubbles” floating on a garnet crystal once held great promise, but you can read about it now on the website of the Vintage Technology Association. The charge-coupled device, or CCD, was another candidate for main memory and mass storage; it failed to gain a foothold in that role, although it did find another niche, as the image sensor of digital cameras. And there were wilder flights of fancy, such as superconducting computers and photonic data processing.

This roster of defeated challengers might lead one to conclude that no innovation has a chance of displacing an entrenched technology. However, the transistor itself offers the obvious refutation. At its debut in 1948 it had to compete with the vacuum tube, which had dominated the electronics industry for 30 years. Although the transistor took more than a decade to establish itself, in the end the vacuum tube became a quaint collector’s item.

Today the TiO<sub>2</sub> memristor is just one of many contending new technologies. Considering only the realm of switched-resistance memory elements, there are several other candidates, including devices based on phase chang-

es, on magnetic fields and on electron spin. (Chua argues that all these devices should be classified as memristors.) To evaluate the long-term prospects of such technologies, one would have to go beyond basic principles of operation to questions of reliability, longevity, uniformity, cost of manufacturing and dozens of other details.

In a telephone conversation I asked Williams why he believes the memristor will be the technology that prevails. He offered several substantive arguments, but he also added, candidly: “It’s the one I’m working on. I have to believe in it.” In a sense this is the strongest endorsement anyone can give. As a bystander, I have the luxury of waiting on the sidelines to see how the contest comes out. But *someone* has to make choices, take risks and commit resources, or nothing new will ever be created.

### Bibliography

- Borghetti, Julien, Gregory S. Snider, Philip J. Kuekes, J. Joshua Yang, Duncan R. Stewart and R. Stanley Williams. 2010. ‘Memristive’ switches enable ‘stateful’ logic operations via material implication. *Nature* 464:873–876.
- Chua, Leon O. 1971. Memristor—the missing circuit element. *IEEE Transactions on Circuit Theory* 18:507–519.
- Chua, Leon O., and Sung Mo Kang. 1976. Memristive devices and systems. *Proceedings of the IEEE* 64(2):209–223.
- Chua, Leon. 2011. Resistance switching memories are memristors. *Applied Physics A* 102(4) (In press).
- Joglekar, Yogesh N., and Stephen J. Wolf. 2009. The elusive memristor: properties of basic electrical circuits. *European Journal of Physics* 30:661–675.
- Keyes, Robert W. 2009. The long-lived transistor. *American Scientist* 97:134–141.
- Li, Hai, and Yiran Chen. 2010. Emerging non-volatile memory technologies. In *Proceedings of the 53rd Midwest Symposium on Circuits and Systems*. doi:10.1109/MWSCAS.2010.5548590.
- Rose, Garrett S. 2010. Overview: Memristive devices, circuits and systems. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems (ISCAS 2010)*, pp. 1955–1958.
- Strukov, Dmitri B., Gregory S. Snider, Duncan R. Stewart and R. Stanley Williams. 2008. The missing memristor found. *Nature* 453:80–83.
- Strukov, D. B., D. R. Stewart, J. Borghetti, X. Li, M. Pickett, G. Medeiros Ribeiro, W. Robinson, G. Snider, J. P. Strachan, W. Wu, Q. Xia, J. Joshua Yang and R. S. Williams. 2010. Hybrid CMOS/memristor circuits. In *Proceedings of 2010 IEEE International Symposium on Circuits and Systems (ISCAS 2010)*, pp. 1967–1970.
- Versace, Massimiliano, and Ben Chandler. 2010. The brain of a new machine. *IEEE Spectrum* 47(12):30–37.
- Williams, R. Stanley. 2008. How we found the missing memristor. *IEEE Spectrum* 45(12):28–35.