A new spin on magnetic memories

Andrew D. Kent and Daniel C. Worledge

Solid-state memory devices with all-electrical read and write operations might lead to faster, cheaper information storage.

Spin-transfer-torque magnetic random access memory (STT-MRAM) devices store information in the orientation of the magnetization of nanometre-scale ferromagnetic elements. As such, they are like hard disk drives, which use magnetic states to store information. In contrast to hard disk drives, however, STT-MRAM is written and read electrically, and does not have moving parts. This is a key difference that enables the integration of magnetic devices with semiconductor chips. Such devices might fulfil the speed requirements of a computer's working memory while having the inherent advantages of using magnetic states — that no energy is needed to retain information.

STT-MRAM is the result of important advances in physics and materials science made over the past 20 years. The first key finding was the theoretical prediction of spin-transfer torque between conduction electrons and magnetization: spin-polarized electrical currents can transfer spin angular momentum to the magnetic moments of a ferromagnet, thus reorienting them[1–3]. In a ferromagnet, the majority and minority electron spin states are shifted in energy. Thus, if the spin polarization of electrons incident on a ferromagnetic layer is not aligned with its magnetization (that is, the electron is not in a definite majority or minority spin-state), the electron spin precesses rapidly around a momentum-dependent internal field of the ferromagnet. Electron spins dephase because of the distribution of electron momenta associated with current flow[4]. As a result, the component of spin-polarization transverse to the magnetization decays, transferring spin angular momentum to the ferromagnet. In transition metal ferromagnets, this dephasing typically occurs at the interface of the ferromagnet, on a length scale of several atomic layers. However, the entire ferromagnetic layer responds to the torques because of the strong exchange coupling of moments throughout its thickness.

Spin-transfer torque provides a mechanism to write information. On the other hand, information can be read by measuring the device resistance. The magnetoresistance refers to the percentage change in resistance between parallel and antiparallel magnetization alignment of the electrodes in a magnetic tunnel junction, which is made of a ferromagnetic metal/insulator/ferromagnetic metal stack[5]. Until 2004, the maximum magnetoresistance reported[6–8] at room temperature was 70%. Magnetoresistance greater than 100% had, however, been predicted in crystalline Fe/MgO/Fe tunnel junctions[9] and was then observed experimentally[10,11]. Subsequent rapid advances in the growth of thin-film materials have led to junctions with large magnetoresistance of several hundred per cent, through the use of transition metal electrodes (typically CoFeB).

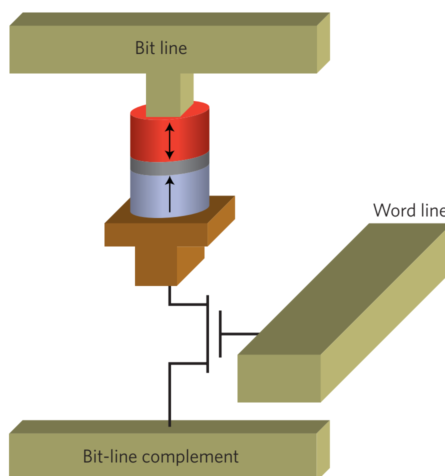

Figure 1 | STT-MRAM bit cell. A magnetic tunnel junction is formed by a fixed reference layer (purple), a tunnel barrier (grey) and a free-layer element (red), with both layers magnetized perpendicular to the plane of the junction (black arrows). The bit is selected by a word line and transistor, and operated by applying biases to the bit lines.

Spin-polarized currents in magnetic tunnel junctions provide a source for spin-transfer torques. The orientation of the magnetization of one electrode of the junction is fixed (by any of a variety of means) and serves as a reference layer that sets the spin-polarization direction. The other electrode acts as a 'free layer' in which the information can be written. Figure 1 shows a 1-bit STT-MRAM cell with a patterned free layer and reference layer, both magnetized perpendicular to the plane of the junction. The cell is accessed through a transistor using a word line; one transistor is required for each cell. Voltage biases on the bit lines operate the cell. The read bias voltage that is used to measure the cell resistance and determine the bit state is low, around 100 mV. The write bias voltage is higher and allows the magnetic moment of the free layer to be reversed by spin-transfer torque.

Device attributes and applications

STT-MRAMs are potentially suitable for a variety of uses, including as replacement of battery-backed static random access memory (SRAM) and as a fast-write buffer in a hard disk or solid-state drive. Table 1 lists the key features of existing and emerging memory technologies. STT-MRAM is the only non-volatile memory expected to have unlimited endurance. This is because there is no inherent magnetic wear-out mechanism for switching magnetic moments back and forth. No atoms are moved during writing operations, as is the case in phase change memory (PCM) or resistive random access memory (RRAM); only the magnetization is rotated. There is, however, an electrical wear-out mechanism — the dielectric breakdown of the MgO tunnel barrier. To avoid this, the write voltage must be kept sufficiently low (roughly 400 mV across the tunnel barrier)[12]. STT-MRAM can be read and written in 10 ns, making it a much faster memory than Flash. This combination

© 2015 Macmillan Publishers Limited. All rights reserved

Table 1 | Comparison of key features of existing and emerging memories.

	SRAM	eDRAM	DRAM	eFlash (NOR)	Flash (NAND)	FeRAM	PCM	STT-MRAM	RRAM
Endurance (cycles)	Unlimited	Unlimited	Unlimited	**10^5**	**10^5**	10^{14}	**10^9**	Unlimited	**10^9**
Read/write access time (ns)	<1	1–2	30	10/**10^3**	100/**10^6**	30	10/100	2–30	1–100
Density	**Low (six transistors)**	Medium	Medium	Medium	High (multiple bits per cell)	**Low (limited scalability)**	High (multiple bits per cell)	Medium	High (multiple bits per cell)
Write power	Medium	Medium	Medium	**High**	**High**	Medium	Medium	Medium	Medium
Standby power	**High**	Medium	Medium	Low	Low	Low	Low	Low	Low
Other	**Volatile**	**Volatile. Refresh power and time needed**	**Volatile. Refresh power and time needed**	**High voltage required**	**High voltage required**	**Destructive readout**	**Operating *T*<125°C**	**Low read signal**	**Complex mechanism**

Significant disadvantages are marked in bold. Estimates for emerging memories are based on expectations for functioning chips, not demonstrations of individual bits. See text for abbreviations.

of speed and endurance, together with the potential for scaling well below 20 nm, makes STT-MRAM promising also as a volatile memory, as the only candidate to replace standalone dynamic random access memory (DRAM), which is expected to be difficult to scale down much below 20 nm (ref. 13). With further improvements in writing (lower switching current) and reading (higher magnetoresistance and narrower resistance distributions), it may be possible to develop embedded STT-MRAM that operates reliably at reading/writing speeds of 2 ns, which would allow replacement of embedded DRAM (eDRAM) and use as cache memory.

Even without further improvements in speed, the non-volatility of STT-MRAM makes it attractive as an embedded memory (a memory built on the same chip as logic and other functions) to store code and data in microcontrollers and digital signal processors, for example for smartcard, automotive and mobile applications[14]. Embedded Flash, the incumbent embedded non-volatile memory, requires the expensive and time-consuming development of specialized high-voltage transistors and large-area charge pumps to enable high-voltage operation; in contrast, STT-MRAM uses standard low-voltage transistors. STT-MRAM also requires a much simpler fabrication process than embedded Flash, and does not require processing above 400 °C, as ferroelectric RAM does when depositing the ferroelectric oxide. For widespread adoption as an embedded technology, STT-MRAM must be processed at the industry-standard temperature of 400 °C used for dielectric materials in interconnections, and not above or below. Most current STT-MRAM is processed around 300 °C to protect the magnetic materials, but recent results[15] show that it is possible to improve thermal endurance to 400 °C. There is also substantial interest in high-operating-temperature embedded memory for automotive and industrial applications that require operation at 180 °C or above. In principle, following materials optimization to increase retention, and a concomitant increase in the write current, STT-MRAM could operate at these high temperatures, whereas PCM and RRAM may suffer retention problems due to atomic motion in the storage material. Non-volatility and the ease with which it can be embedded with logic circuits make STT-MRAM a leading candidate to replace Flash for embedded non-volatile storage applications.

In addition to use as storage memory, STT-MRAM can also potentially be used as an embedded working memory in some mobile applications, even if it does not achieve the speed of embedded DRAM or SRAM. For example, STT-MRAM could replace SRAM in mobile co-processors owing to its low standby power, high density, high endurance and 10-ns speed[14]. Such co-processors run in the background at low speed and low power while the main processor sleeps, providing monitoring and processing of sensor data and low-power wireless connectivity. Because this application does not require high-speed operation, STT-MRAM is a natural solution, significantly reducing both chip area and leakage power. The Internet-of-Things, smartwatches and other wearable electronics may drive a substantial increase in this type of low-power, low-performance, always-on application.

Device physics and materials challenges

Figure 2 shows the electrical characteristics of an STT-MRAM device. The device is bistable near zero applied field, providing memory states (Fig. 2a), with positive field favouring parallel (P) magnetization alignment and a low-resistance state, and negative field favouring antiparallel (AP) magnetization alignment and a high-resistance state. The magnetic field reorients the magnetization of the free layer while that of the reference layer is fixed; these field-dependent measurements give access to device characteristics such as its resistance, magnetoresistance and the free-layer switching fields. Device operation, however, does not require applied magnetic fields. A voltage bias alone generates a current that switches the device between P and AP states by the STT mechanism (Fig. 2b), where positive voltage leads to AP-to-P switching and negative voltage the reverse transition, P to AP.

Data retention is associated with the lifetimes of the P and AP states under device operating conditions, in terms of temperature and read-out bias. The lifetimes are set by thermally activated transitions between states, and are given, to a good approximation, by an Arrhenius law

© 2015 Macmillan Publishers Limited. All rights reserved

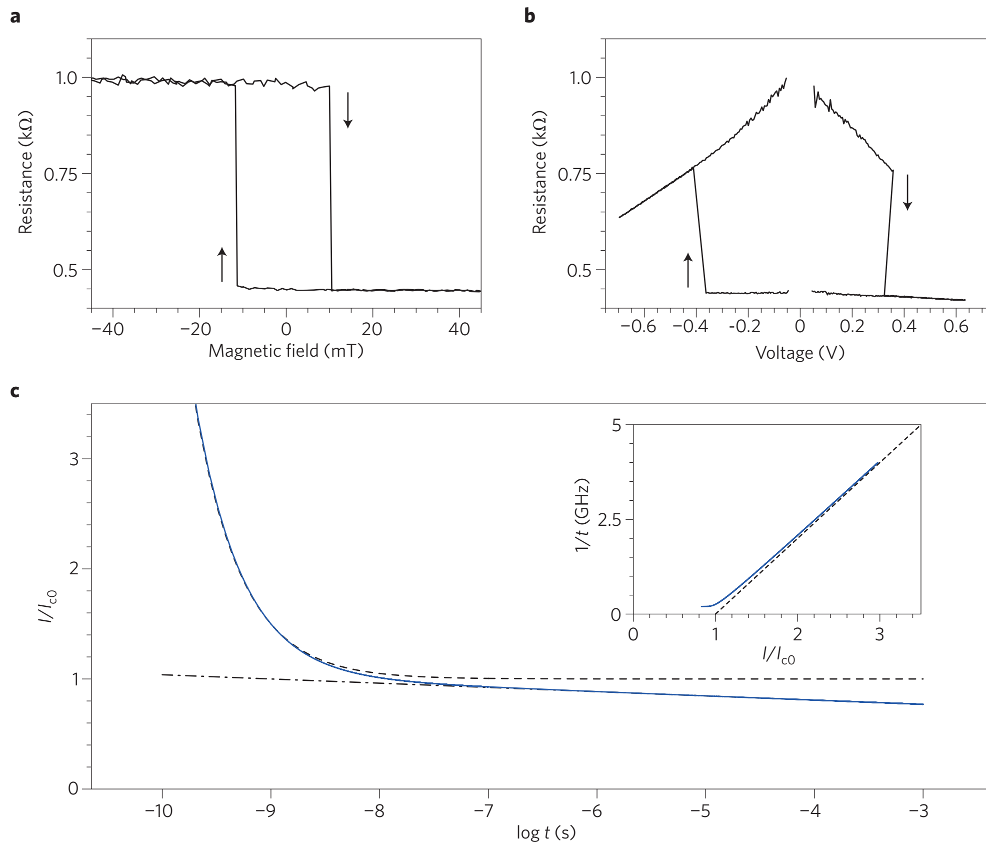

Figure 2 | STT-MRAM electrical characteristics. **a**, Resistance versus applied magnetic field, showing bistable resistance states near zero-field associated with parallel (P) and antiparallel (AP) magnetized bits. **b**, Resistance versus voltage, showing switching between AP and P states (positive bias) and vice versa (negative bias). **c**, Pulse switching amplitude versus pulse duration, on a logarithmic scale for fixed switching probability. The dashed line shows the inverse of the pulse duration proportional to pulse amplitude, characteristic of the ballistic switching limit at short times, while the dashed-dotted line is characteristic of the long-time behaviour, thermally activated transitions assisted by STT. The slope of the dashed-dotted line is inversely related to the energy barrier to magnetization reversal $\Delta = U/(kT)$ (ref. 20); measurements of these device characteristics can thus be used to estimate Δ. The inset shows the inverse pulse duration versus pulse amplitude in the short time limit. The slope of this curve is the STT dynamic parameter A, and the intercept with the x axis occurs at I_{c0}, the threshold current for STT switching, permitting determination of key device parameters from short-time pulse switching data.

$\tau = \tau_0 e^{U/(kT)}$, where U is the energy barrier between P and AP states, T is the device operating temperature, k is Boltzmann's constant, and τ_0 is a characteristic attempt time, of the order of 1 ns. Retention times of 10 years (the industry standard) thus require bit stability factors $\Delta = U/(kT) = 40$. To minimize data loss in large memory arrays, however, larger bit stability factors ($\Delta > 60$) are needed. U is proportional to the free layer's magnetic anisotropy and its volume; thus, reducing the bit size requires a corresponding increase in its magnetic anisotropy. Asymmetric bit shapes — such as with an elliptically shaped, in-plane magnetized thin-film element — lead to anisotropy barriers associated with magnetic dipole interactions (shape anisotropy). But the maximum shape anisotropy is limited by a material's magnetization density, and sets a minimum element width of about 40 nm. Scaling to smaller feature sizes requires using large perpendicular bulk magnetocrystalline and interface anisotropies, which are magnetic anisotropies associated with spin–orbit interactions. Thin CoFeB films with interfaces to MgO are perpendicularly magnetized, and have a sufficiently large perpendicular, interface-induced magnetic anisotropy to lead to stable bits with dimensions smaller than 20 nm (refs 16,17).

The more stable a bit, the larger the torques needed to reverse its magnetic moment. In a perpendicularly magnetized bit cell (in which the free layer has a uniaxial magnetic anisotropy pointing out of the plane) the current I_{c0} required to destabilize a state (P or AP) is directly proportional to the energy barrier between the P and AP states. Within a simplified macrospin model, $I_{c0} = 4e\alpha U/(\hbar P)$, where α is the Gilbert damping of the free layer, P is the spin polarization of the current, e is the electron's charge and $\hbar$ is the reduced Planck constant[18]. Values of $P \approx 1$, $\alpha = 0.01$ with $\Delta = 60$ give $I_{c0} \approx 15$ μA, showing that low write currents are possible, within this model of the magnetic moment dynamics. Fast switching of the magnetization in less than 10 ns requires currents larger than I_{c0}, that is, a current overdrive $i = I/I_{c0}$. Figure 2c shows a schematic plot of the amplitude of the switching current versus its pulse duration. For short pulse durations t, $1/t = A(i - 1)$ (dashed line in Fig. 2c, main figure and in the inset): that is, the switching pulse duration is inversely proportional to the current overdrive, reflecting conservation of angular momentum, where A is the STT dynamic parameter[19,20]. This relation characterizes the limit for ballistic switching, in which STT is larger than torques associated with thermal fluctuations; I_{c0} and A can be determined from short pulse switching data, as shown in the inset of Fig. 2c. Instead, for $i < 1$, the physics of device operation is based on thermally activated transitions assisted by STT; STT effectively lowers the energy barrier U between states. The result is a logarithmic dependence of switching

© 2015 Macmillan Publishers Limited. All rights reserved

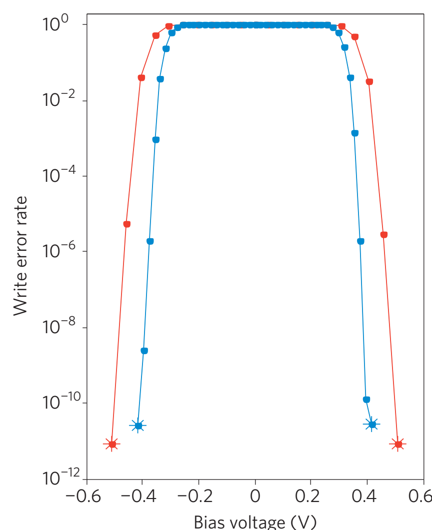

Figure 3 | Write error rate versus the junction bias voltage. The error rate is shown for two pulse durations, 10 ns (red line) and 50 ns (blue line). The stars represent an upper limit on the write error rate, that is, error rates below the experiment's detection limit. Figure adapted with permission from ref. 26, IEEE.

pulse amplitude on pulse duration (dotted-dashed line in Fig. 2c), with a slope that is inversely proportional to the energy barrier $\Delta = U/(kT)$. It is important to note that thermal fluctuations render the switching process stochastic at all times. Thus i versus t in Fig. 2c reflects a particular switching probability; the form of the curve, however, is not very sensitive to the switching probability.

Thermal fluctuations lead to write errors — write pulses that do not change the bit state — and read disturbs, that is read biases that change the bit state. Figure 3 shows the write error rate versus bias voltage for two different pulse durations. Write errors decrease rapidly with overdrive and are predicted to be an exponential function of the overdrive[20]. Errors in the limit of short time pulses are associated with magnetization states that are collinear or nearly collinear with the spin-polarization axis, where there is no transverse component of spin polarization to create a STT. This has led to STT-MRAM device concepts that incorporate spin-polarizing filters orthogonal to the magnetization of the free layer to ensure the onset of STT the moment a device is biased[21]; this can reduce both the switching time and write error rate in certain cases.

Advances in materials engineering are needed to improve the performance of STT-MRAM devices. Chief among them is to achieve even larger magnetic anisotropy to enable stable magnetic bits with lateral sizes smaller than 10 nm. The interface magnetic anisotropy (for example in CoFeB/MgO) alone may not be sufficient, and may need to be augmented with bulk magnetocrystalline magnetic anisotropy, while, of course, maintaining large magnetoresistance and low damping. Faster device operation (<10 ns) requires further reduction in the write current and a careful engineering of STT to reduce the write errors at a given voltage or current overdrive. These goals, along with reducing the distribution of device parameters in a memory array, remain the central challenges to advancing STT-MRAM technology.

Commercial challenges and status

A number of technological challenges must be met before STT-MRAM can be widely adopted in the most advanced applications. Most importantly, the write current must be further reduced to enable the use of smaller-area transistors. Reducing the write current below 20 μA (which means I_{c0} is less than 10 μA) would enable multi-gigabit memories. A distinct but related requirement is to keep the write voltage low, around 400 mV across the magnetic tunnel junction, to avoid breakdown of the MgO barrier. Increasing the magnetoresistance to over 300% is needed to enable large memories and faster reading. For embedded applications in particular, magnetic materials must be developed that can withstand processing at the industry standard temperature of 400 °C, instead of the ~300 °C that is now commonly used to process MRAM. An on-pitch etch must be developed that does not cause shorting or damage the magnetic properties of the magnetic tunnel junction, as is the case with current etching methods based on methanol reactive ion etching and ion beam etching. Improved etching to give lower spreads in resistance from bit to bit is also required, both to provide sufficient read margin for larger memories and to increase read speed. The height of the magnetic tunnel junction must be reduced to enable easier etching and to allow the device to be placed in between standard levels of metallization in embedded architectures. The layer composition of the magnetic tunnel junction stack may need to be simplified and deposition tooling improved to increase throughput from roughly one wafer per hour to more than 20 wafers per hour. Tool vendors, including Anelva, TEL/AMAT and Singulus, are currently developing 300-mm production-level tools for deposition and etching of magnetic tunnel junctions to help to meet these challenges.

STT-MRAM has come a long way since the switching mechanism was proposed[1–3] and the first demonstration using a magnetic tunnel junction was made[22]. Sony made the first integrated demonstration chip in 2005, building a 4-kbit chip using in-plane magnetized materials[23]. Toshiba developed the first magnetic tunnel junctions with perpendicular magnetization, and demonstrated STT switching in individual devices[24]. Tohoku and IBM independently developed magnetic tunnel junctions with perpendicular magnetization using interface perpendicular magnetic anisotropy[16,17], and IBM demonstrated sufficient margins in write bias and reliable writing down to error rates of 10^{-12} (refs 25,26). Samsung has realized the smallest functioning devices[27], with lateral dimensions of 17 nm. TDK demonstrated a 400 °C process and reliable writing at 2 ns on individual devices, and functionality of an 8-Mbit chip designed by IBM[15]. Everspin is the closest to commercialization, now sampling an in-plane 64-Mbit chip[28].

Perspectives

Along with this rapid progress in STT-MRAM have come discoveries that may make the mechanism of STT even more energy efficient, as well as enabling new device geometries. Recently, it has been demonstrated that spin–orbit interactions in non-magnetic conducting layers can also generate strong STTs on a proximal magnetic layer[29,30]. For example, spin–orbit interactions can lead to spin currents transverse to the charge currents, called the spin Hall effect. In this case, STT can be generated without charge currents traversing the tunnel barrier. This enables a three-terminal device with separate write and read paths, with the advantage that these characteristics (for example read and write impedances and magnetoresistance) can be optimized separately, and the disadvantage of a potentially larger memory cell (that is, a lower memory density). Moreover, the surface states on three-dimensional topological insulators have been shown to produce large STT[31] and magnetic switching[32]. It is too early to say what the impact of these discoveries will be on STT-MRAM, but these new means of generating spin torques are an active and exciting area of basic research. ❐

Andrew D. Kent is in the Department of Physics, New York University, New York, New York 10003, USA. Daniel C. Worledge is at the IBM T. J. Watson Research Center, Yorktown Heights, New York 10598, USA.
e-mail: andy.kent@nyu.edu; worledge@us.ibm.com

© 2015 Macmillan Publishers Limited. All rights reserved

References

1. Slonczewski, J. C. *Phys. Rev. B* **39**, 6995–7002 (1989).
2. Slonczewski, J. C. *J. Magn. Magn. Mater.* **159**, L1–L7 (1996).
3. Berger, L. *Phys. Rev. B* **54**, 9353–9358 (1996).
4. Stiles, M. D. & Zangwill, A. *Phys. Rev. B* **66**, 014407 (2002).
5. Jullière, M. *Phys. Lett.* **54A**, 225–226 (1975).
6. Moodera, J. S., Kinder, L. S., Wong, T. M. & Meservey, R. *Phys. Rev. Lett.* **74**, 3273–3276 (1995).
7. Miyazaki, T. & Tezuka, N. *J. Magn. Magn. Mater.* **139**, L231–L234 (1995).
8. Wang, D., Nordman, C., Daughton, J. M., Qian, Z. & Fink, J. *IEEE Trans. Magn.* **40**, 2269–2271 (2004).
9. Butler, W. H., Zhang, X-G., Schulthess, T. C. & MacLaren, J. M. *Phys. Rev. B* **63**, 054416 (2001).
10. Yuasa, S., Nagahama, T., Fukushima, A., Suzuki, Y. & Ando, K. *Nature Mater.* **3**, 868–871 (2004).
11. Parkin, S. S. P. *et al. Nature Mater.* **3**, 862–867 (2004).
12. Min, T. *et al. IEEE Trans. Magn.* **46**, 2322–2327 (2010).
13. Liu, J., Jaiyen, B., Veras, R. & Mutlu, O. in *Proc. 39th Annu. Int. Symp. Computer Architecture (ISCA '12)* 1–12 (IEEE Computer Society, 2012).
14. Lee, K., Kan, J. J. & Kang, S. H. in *Proc. 2014 Int. Symp. Low Power Electronics and Design (ISLPED '14)* 131–136 (ACM, 2014).
15. Thomas, L. *et al. J. Appl. Phys.* **115**, 172615 (2014).
16. Ikeda, S. *et al. Nature Mater.* **9**, 721–724 (2010).
17. Worledge, D. C. *et al. Appl. Phys. Lett.* **98**, 022501 (2011).
18. Sun, J. Z. *Phys. Rev. B* **62**, 570–578 (2000).
19. Bedau, D. *et al. Appl. Phys. Lett.* **97**, 262502 (2010).
20. Liu, H. *et al. J. Magn. Magn. Mater.* **358–359**, 233–258 (2014).
21. Kent, A. D., Özyilmaz, B. & del Barco, E. *Appl. Phys. Lett.* **84**, 3897–3899 (2004).
22. Huai, Y., Albert, F., Nguyen, P., Pakala, M. & Valet, T. *Appl. Phys. Lett.* **84**, 3118–3120 (2004).
23. Hosomi, M. *et al.* in *Proc. IEDM Tech. Dig.* 459–462 (2005).
24. Kishi, T. *et al.* in *Proc. IEEE Int. Electron Devices Meeting (IEDM)* 1–4 (2008).
25. Worledge, D. C. *et al. Proc. IEEE Int. Electron Devices Meeting (IEDM)* 296–299 (2010).
26. Nowak, J. J. *et al. IEEE Magn. Lett.* **2**, 3000204 (2011).
27. Kim, W. *et al. IEEE Int. Electron Devices Meeting (IEDM)* 24.1.1–24.1.4 (2011).
28. http://www.everspin.com/
29. Miron, I. M. *et al. Nature Mater.* **9**, 230–234 (2010).
30. Liu, L. *et al. Science* **336**, 555–558 (2012).
31. Mellnik, A. R. *et al. Nature* **511**, 449–451 (2014).
32. Fan, Y. *et al. Nature Mater.* **13**, 699–704 (2014).

Acknowledgements

A.D.K. thanks G. Wolf for comments on the manuscript and for preparing Fig. 2. He acknowledges support from the National Science Foundation, grant number NSF-DMR-1309,202. D.C.W. thanks J. DeBrosse for comments on the manuscript.

Competing financial interests

A.D.K. is the founder of Spin Transfer Technologies.

Memory leads the way to better computing

H.-S. Philip Wong and Sayeef Salahuddin

New non-volatile memory devices store information using different physical mechanisms from those employed in today's memories and could achieve substantial improvements in computing performance and energy efficiency.

Current memory devices store information in the charge state of a capacitor; the presence or absence of charges represents logic 1's or 0's. Several technologies are emerging to build memory devices in which other mechanisms are used for information storage. They may allow the monolithic integration of memories and computation units in three-dimensional chips for future computing systems[1]. Among those promising candidates are spin-transfer-torque magnetic random access memory (STT-MRAM) devices, which store information in the magnetization of a nanoscale magnet. Other candidates that are approaching commercialization include phase change memory (PCM), metal oxide resistive random access memory (RRAM) and conductive bridge random access memory (CBRAM).

Today's computing systems use a hierarchy of volatile and non-volatile data storage devices to achieve an optimal trade-off between cost and performance[2]. The portion of the memory that is the closest to the processor core is accessed frequently, and therefore it requires the fastest operation speed possible; it is also the most expensive memory because of the large chip area required. Other levels in the memory hierarchy are optimized for storage capacity and speed (Fig. 1). The main memory is often located in a separate chip because it is fabricated with a different technology from that of the microprocessor.

For over 30 years, static random access memory (SRAM)[3] and dynamic random access memory (DRAM)[3] have been the workhorses of this memory hierarchy[4]. Both SRAM and DRAM are volatile memories — that is, they lose the stored information once the power is cut off. For non-volatile data storage, magnetic hard disk drives (HDDs) have been in use for over five decades[5–7]. Since the advent of portable electronic devices such as music players and mobile phones, however, solid-state non-volatile memory known as Flash memory[8] has been introduced into the information storage hierarchy between the DRAM and the HDD. Flash has become the dominant data storage device for mobile electronics; increasingly, even enterprise-scale computing systems and cloud data storage systems are using Flash to complement the storage capabilities of HDD.

Resistive switching memory technologies

The design specifications for memory (volatile data storage, fast, expensive) and for storage (non-volatile data storage, slow, inexpensive) are different, and they often have different data access standards and protocols. Around 15 years ago, researchers started exploring the possibility of blurring the design boundary between memory and storage[9,10], and coming up with new data access modes and protocols that are neither 'memory' nor 'storage'. Indeed, the adoption of Flash in the memory hierarchy (albeit on a separate chip from the processor) inspired the exploration of computing architectures that capitalize on the salient features of Flash: non-volatility and high density[11]. At the same time, new types of non-volatile memory have emerged that can easily be integrated on-chip with the microprocessor cores because they use a different set of materials and require different device fabrication technologies from Flash[12]. Some of them can be programmed and read quickly; others can have very high data storage density. Importantly, all of these memories are free from the limitations of Flash — that is, low endurance, need for high voltage supply, slow write speed

© 2015 Macmillan Publishers Limited. All rights reserved