

Editorial: Opinion: the next infrastructure waves behind AI data centers



By the Editor of MIIJ, Dr. Chong Leong, Gan

Every technology era has its illusion of progress. Today's is the belief that faster accelerators alone will carry AI forward. What truly decides whether AI data centers scale or stall, are not models or GPUs, but three deep physical realities: how data moves, where it lives and how power arrives. Interconnect, memory and power are no longer supporting actors. They are the main plot.

Interconnect: When physics votes against copper

For decades, copper earned the benefit of the doubt. It was cheap, familiar and "good enough." AI changed the math. At 200+ Gbps per lane and rack level power densities pushing past 100 kW, copper faces a hard physical wall: loss, heat and reach that no amount of signal conditioning can fully fix. This is why the industry's shift toward silicon photonics and co-packaged optics is accelerating at a pace that surprises even insiders.

The debate is non-optical versus copper: it is how long partial copper architecture can survive. Co-packaged optics move photons closer to computers not because they are fashionable, but because power per bit has become existential. In future AI systems, wasting watts on interconnect is simply unaffordable. Optical links are becoming mandatory for plumbing, not exotic upgrades, as clusters scale toward hundreds of thousands and eventually millions of accelerators. However, some technical challenges are on the front for enabling optical data transfer (Source: Emerald MIIJ SiPh article, [semiengineering.com]).

Memory: Bandwidth became the bottleneck of ambition

If interconnect determines how fast data moves, memory determines how far AI can think. High bandwidth memory (HBM) has quietly become the constraint that shapes model size, training efficiency and even architectural creativity. Each new HBM

generation delivers staggering bandwidth, but at rapidly rising thermal and power costs. Memory stacks are no longer passive companions to compute; they are energy-dense devices demanding codesign with interposers, cooling and packaging strategies.

What troubles the AI datacenter application is how often memory is treated as a supply chain contest instead of a system level reckoning. As HBM cube's power per stack climbs, the industry must confront uncomfortable questions: how much bandwidth is truly useful, where near-memory compute makes sense and when smarter data movement beats brute force scaling. AI will not simply "out memory" physics forever.

Power: From afterthought to first order constraint

Power has become the loudest voice in AI infrastructure and the most neglected in technical narratives. Traditional 12 V delivery designs cannot support the current demands implied by modern AI racks. The shift to 48 V busbars and wide bandgap conversion is no longer optional; it is the minimum viable architecture for survival. Connection delays, grid congestion and permitting constraints are now shaping where data centers can even exist. This is where hydrogen fuel cell technology enters the conversation, not as a green side project, but as a reliability and deployment enabler.

Hydrogen fuel cells: Backup today, strategic power tomorrow

Hydrogen fuel cells are emerging as a credible replacement for diesel generators in AI data centers, offering zero onsite emissions, rapid ramp up and high reliability for backup power scenarios. More interesting, and more disruptive, is the shift from backup to strategic power. For AI data centers, uptime and deployment velocity often matter more than theoretical efficiency.

Closing thought

Interconnect, memory, power and now hydrogen are converging into a single system problem. Optimizing one in isolation no longer works. The next wave of AI infrastructure will be decided by engineers and architects willing to think across photons, electrons and molecules at once. From an editorial standpoint, this is both sobering and exciting. The future of AI data centers will not be defined by marginal GPU/TPU gains alone, but by how honestly the industry confronts physics and how boldly it redesigns the infrastructure beneath the algorithms!

Chong Leong Gan

Department of Package Development Engineering,
Micron Memory Taiwan Co Ltd, Taichung City, Taiwan

Further reading

Gan, C.L., Chung, M.H. and Huang, C.Y. (2026), "Silicon photonics packaging for AI datacenters: technical challenges, reliability requirements and future directions – a critical review", *Microelectronics International*, Vol. 43 No. 3, doi: [10.1108/MI-12-2025-0262](https://doi.org/10.1108/MI-12-2025-0262).

Website (2026), "All AI data center interconnects will Be optical within 5 years".

The current issue and full text archive of this journal is available on Emerald Insight at: <https://www.emerald.com/insight/1356-5362.htm>



Microelectronics International
43/4 (2026) 213
© Emerald Publishing Limited [ISSN 1356-5362]
[DOI [10.1108/MI-10-2026-310](https://doi.org/10.1108/MI-10-2026-310)]