

STORAGE SOLUTIONS: CAN WE REIMAGINE THE TRADE-OFFS THAT LIMIT FLASH STORAGE DEVICES?

EXPERTS FROM THE INTERNATIONAL DATA GROUP HAVE PREDICTED THAT BY 2025, EVERY PERSON WHO IS CONNECTED TO THE INTERNET WILL ENGAGE WITH DIGITAL DATA ALMOST 5,000 TIMES EVERY DAY. THIS WILL RESULT IN AN EXTRAORDINARY AMOUNT OF DATA, AND A LOT OF IT WILL NEED TO BE STORED SOMEWHERE. MOST OF THE DATA THAT WE CURRENTLY CREATE, LIKE PHOTOS, VIDEOS AND WORD DOCUMENTS, ARE STORED ON FLASH DRIVES ON OUR PHONES, LAPTOPS OR PORTABLE USB DEVICES. **DR BRYAN S. KIM**, FROM **SYRACUSE UNIVERSITY** IN THE US, IS STUDYING SOME OF THE INNOVATIVE METHODS THAT ARE BEING DEVELOPED TO HELP IMPROVE INFORMATION STORAGE

TALK LIKE A COMPUTER ENGINEER

SOLID-STATE DRIVES (SSDs) –

a storage device that uses electronically programmable flash memory to store data. They are the *de facto* storage for mobile systems

HARD DISK DRIVES (HDDs) –

a storage device that uses magnetic and mechanical components to store data. They are bulkier, slower and more power-intensive than SSDs

CAPACITY – the amount of data that can be stored on a storage device

RELIABILITY – the ability to maintain the original data without errors or corruption

PERFORMANCE – the speed at which data in a storage device can be accessed (latency) or how much data can be accessed in a given time (throughput)

In today's technology-driven world, the amount of data that each of us generates is increasing rapidly. Every photo you take on your phone, every save you make in a video game and every word document you create for school needs to be stored somewhere. The chances are that most of your devices make use of something called flash storage to store all your data.


Flash storage, also known as solid-state drives (SSDs), is a type of electronic data storage that keeps its data persistently. This means that when you turn the storage device off and back on again, all the data will still be there. These days, SSDs are everywhere, from USB sticks and smartphones to PCs and large-scale cloud systems.

SSDs are so popular due to their advantages over the more traditional hard-disk drives (HDDs). HDDs store data using a mix of magnetic and mechanical parts, making them bulkier and susceptible to damage from vibrations and other physical disturbances like being dropped. SSDs,

on the other hand, are completely electronic, allow computers and devices faster access to information stored in them and consume less power than HDDs.

The one advantage that HDDs have over SSDs is their price. In terms of cost per byte, SSDs are more expensive than HDDs. However, there is a huge amount of effort being focused on how to increase the density of SSDs to reduce their cost. Many of these techniques involve making trade-offs between three aspects of data storage: capacity (cost per byte), reliability (error rate) and performance (how fast data can be accessed).

Finding the best way to balance these trade-offs is a complex problem that many scientists and engineers are trying to solve. One such researcher is Dr Bryan S. Kim, an assistant professor at the College of Engineering and Computer Science at Syracuse University in the US. Bryan has been researching flash storage systems to discover innovative methods that might improve their efficiency.



BRYAN S. KIM

Syracuse University, USA

.....

FIELD OF RESEARCH

Computer engineering

.....

RESEARCH PROJECT

Exploring better ways to build data storage systems to balance the trade-off between capacity, reliability and performance

.....

FUNDER

US National Science Foundation (NSF)

This material is based on work supported by the NSF under Grant CNS-2008453.

Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

WHAT TECHNIQUES ARE USED TO INCREASE SSD DENSITY AND REDUCE ERROR RATES?

Storage density is a measure of how many ‘bits’ of data can be stored in a device. A bit represents a single binary data point (1 or 0) and is the most basic unit of information in computing. “The techniques for increasing the density of SSDs include packing more bits into a smaller cell or stacking cells on top of each other vertically,” says Bryan. “However, these, in turn, increase the error rate of the storage.”

There are a few different techniques that SSDs use to reduce their error rate and increase their reliability. “The first line of defence against errors are error correction codes (ECCs),” explains Bryan. ECCs can detect errors in the original code and correct them, but they only have a limited number of errors that they can correct. If the ECC fails, then the SSD may attempt to re-read the data using a different set of parameters. Some SSDs even proactively read the data stored on them and make corrections before the data become corrupted.

Unfortunately, these techniques, which are designed to increase reliability, have a detrimental effect on the performance of the system. Here, we can see how the trade-off between capacity, reliability and performance works: engineers attempt to increase the capacity of SSDs, which reduces the reliability of the storage drives. And when engineers attempt to increase the drive’s reliability, its performance is reduced, which can lead to fail-slow symptoms.

WHAT ARE FAIL-SLOW SYMPTOMS AND HOW CAN THEY BE AVOIDED?

As the components of a storage system age, wear out and begin to fail, the performance of the storage system will get worse. This decreased level of performance can result in stalling, glitches and slower processing speeds, all of which are examples of fail-slow symptoms. The storage units within an SSD, known as cells, wear out over time and become more error prone. This deterioration causes the performance to slow down as the system puts more effort into reducing errors and maintaining its capacity.

Bryan’s research suggests that a concept known as capacity variance may provide a solution to the problem of fail-slow symptoms. Most storage drives currently have a fixed capacity. “For example, a new 1 terabyte (TB – 1,000 gigabytes) SSD will always be 1 TB in total size,” says Bryan. “To maintain the illusion of a 1 TB-sized healthy set of blocks, the device internally uses reliability enhancement techniques that sacrifice performance, leading to fail-slow symptoms.”

Bryan argues that a storage system with the ability to vary its capacity could be more effective. A capacity-variant storage system would be able to reduce its capacity as its cells begin to deteriorate and fail. “For example,” explains Bryan, “a 1 TB SSD may reduce its capacity to 900 GB when some of its cells degrade and exhibit higher error rates.” Removing worn-out cells would get rid of the need to perform reliability-enhancing techniques. This would allow the system to maintain its performance and avoid developing fail-slow symptoms. “Essentially,” says [➔](#)

Bryan, “the fixed-capacity model sacrifices performance to maintain capacity and reliability.” However, a capacity-variant model would allow a system to sacrifice capacity to maintain performance and reliability.

WHAT ARE THE CHALLENGES BRYAN FACES WHEN UNDERTAKING THIS RESEARCH?

The deterioration of SSDs occurs over an extended period as errors accumulate. The errors that cause this deterioration occur by chance and are therefore difficult to predict. This makes studying fail-slow symptoms challenging because it is hard to accurately re-create the scenarios in which they occur.

When studying SSDs by reproducing these states, Bryan must contend with another trade-off. “That is,” Bryan explains, “we can either age an SSD quickly while sacrificing how realistic the aging process would be, or we can age it slowly while being close to how it would really behave.” Bryan is hoping to bypass this problem by creating a data-driven model that learns how SSDs deteriorate based on past occurrences.

WHAT HAS BRYAN'S RESEARCH UNCOVERED SO FAR?

As a precursor to the idea of capacity variance, Bryan and his team have been arguing against the effectiveness of some reliability-enhancing techniques. One such technique is known as wear levelling. This involves spreading the deterioration of cells equally around the storage system. “It makes sense to do this if we would like the SSD to exhibit an all-or-nothing failure state,” says Bryan. “Either all of its cells are worn out and unusable, or none of them is.”

With a capacity-variant system, wear levelling would no longer be necessary. Bryan and his team have collected and analysed lots of data that highlight the problems with wear levelling, which they hope will strengthen their argument for capacity-variant systems.

WHAT ARE THE NEXT STEPS FOR BRYAN'S RESEARCH?

“Our research group is always thinking about exploring better ways to build storage systems,” says Bryan. A lot of the time, this involves making sure that a system is using the right kind of interface. Most storage systems

use a fixed capacity interface; however, as we have seen, these are not always the most effective. Aside from their capacity-variant concept, Bryan and his team are excited by two other kinds of interfaces.

The zoned namespace (ZNS) interface allows incoming data to be grouped based on their use and how often the information needs to be accessed. This allows the design of the SSD to be simplified and increases its efficiency and longevity. Bryan also emphasises his enthusiasm for another kind of interface known as a key-value (KV) interface. Key-value interfaces make SSDs more complex, but significantly increase the processing power of the system. “Again, there exists a set of trade-offs,” says Bryan. “We can’t have it all!

EXPLORE A CAREER IN COMPUTER ENGINEERING

- The Association for Computing Machinery (www.acm.org) has many resources and general material that may be of interest to anyone thinking about a career in computer engineering.
- USENIX (www.usenix.org) is another society that provides resources and support for people who are interested in computer engineering. They also have a discounted student membership.
- The Chartered Institute for IT (www.bcs.org) has lots of resources for people exploring a career in computer engineering. It also runs an apprenticeship programme to help young people start a career in the industry.
- According to Salary.com, the average salary for a computer engineer in the US is \$89,000 (www.salary.com/research/salary/listing/computer-engineer-salary)

PATHWAY FROM SCHOOL TO COMPUTER ENGINEERING

- Studying maths at school is fundamental and is an entry requirement for most computer engineering courses. Taking classes in physics and IT (if they are offered) will also provide you with a good foundation.
- Studying these subjects should help you get a place on a computer engineering degree at university, which is vital if you want to be a computer engineer. There are many different courses to choose from, so make sure to do your research.
- Join your school's computing club. This is a wonderful way to get hands-on experience with computers and make some new friends.
- Consider going to a computer science summer school (e.g., www.summerschoolsineurope.eu/search/discipline;comsc). This is an interesting way to gain experience and will look great on your university application.

ABOUT COMPUTER ENGINEERING

Computer engineering is a field that combines electrical engineering and computer science. Computer engineers design and develop hardware and software that is used in all kinds of computing contexts, from laptops and supercomputers to self-driving cars and robots.

"IDEAS ARE NEVER STUCK AT THE 'ON THEORY' OR 'ON PAPER' STAGE, BUT ARE ALWAYS IMPLEMENTED AND EVALUATED TO SEE HOW THE IDEAS WOULD FARE IN THE REAL WORLD."

The demand for computer engineers is expected to continue increasing over the next decade. Computers and computer-based systems play a huge role in our lives, and that does not look like it will be changing any time soon.

WHAT IS REWARDING ABOUT A CAREER IN COMPUTER ENGINEERING?

"Computer systems research tends to be practical and hands-on with a lot of real-world impact," says Bryan. "Ideas are never stuck at the 'on theory' or 'on paper' stage but are always implemented and evaluated to see how the ideas would fare in the real world."

Computer engineers are often involved in developing cutting-edge technology. "We are

at the forefront of the industrial revolution," says Bryan. Computer engineering is an exciting field to be a part of because new, innovative technologies are being created all the time. As we advance further into the 21st century, our need for technological innovation is only going to increase.

WHAT SKILLS AND PERSONAL QUALITIES ARE USEFUL FOR COMPUTER ENGINEERS TO HAVE?

Curiosity is an excellent quality to have as it drives you to learn more and understand the inner workings of many different devices and things. Abstract reasoning is also another excellent quality because it helps you put those tiny details from your curiosity into the big picture; the two qualities should go together.

HOW DID BRYAN BECOME A COMPUTER ENGINEER?

WHAT WERE YOUR INTERESTS WHEN YOU WERE YOUNGER? HAVE YOU ALWAYS BEEN FASCINATED BY ENGINEERING AND COMPUTER SCIENCE?

I enjoyed playing and building with Lego as a child and have always been interested in math. Up until when I was applying for college, I actually did not think that engineering and computer science would be the path that I would be taking – in fact, a lot of my high school friends thought I would major in math. I sent out a mix of college applications and chose to go in the direction of engineering and computer science.

I did not do very well in college, however. Most (if not all) other students had experience with programming and computer science coming into college while I had none, and everyone seemed smarter than me.

WHAT DO YOU ENJOY DOING WHEN NOT WORKING?

While my toddler is awake, I enjoy playing and spending time with him. When he is asleep (and finally peaceful), my wife and I like to play board games together.

BRYAN'S TOP TIPS

1. Tinker with stuff. Build stuff, find ways to improve, tear it down and make it better.
2. Manage your time well. Don't miss deadlines but spend time on what is important to you.
3. Develop your logical thinking. It is fundamental for computer engineering.