

MANGALAM COLLEGE OF ENGINEERING

TECHNOBYTE

2023



DEPARTMENT OF ELECTRONICS AND
COMMUNICATION ENGINEERING

What's Inside:

06 CHAIRMAN'S MESSAGE

07 PRINCIPAL'S MESSAGE

08 HOD'S MESSAGE

09 AI AT THE EDGE

13 THE POST-BATTERY PARADIGM

19 THE HIDDEN SYMPHONY OF FREQUENCIES

24 BEYOND THE SIGNALS

28 PCIE GEN 5 SSDS

32 NEXTGEN CONNECTIVITY

36 BIO ELECTRONICS

**40 THE SILICON HORIZON: THE NEW
ARCHITECTURE OF INTELLIGENCE**

47 THE SENTIENT GRID

EDITORIAL BOARD 2023



SUMESH CHANDRAN

EDITORS



SREEJISH K PILLA
S7 ECE



AKHIL K SOMAN
S7 ECE



ACHU JAYAMON
S5 ECE

DESIGN TEAM



HARIKRISHNAN S
S5 ECE



ANANDHU KRISHNAN
S5 ECE



AYANA ANTONY
S5 ECE

CONTENT TEAM



SREEKESH KP
S3 ECE



BIGIL
S3 ECE



NIKHIL KRISHNA G
S3 ECE

VISION OF THE DEPARTMENT

To achieve excellent standards in technical education and engage research in the field of Electronics and Communication Engineering creating employable and innovative professionals who can excel in global challenges.

MISSION OF THE DEPARTMENT

- 1) To practice innovative teaching and learning practices in Electronics and Communication domain with effective curriculum enhancement and Industry collaboration.
- 2) To inspire creative enquiry and innovation in students through excellent training programs and projects for professional skills.
- 3) To imbibe a sense of team work, ethics combined with social responsibility in students.



Pursuing your passions makes you more interesting, and interesting people are enchanting.

Principal




Desire is the key to motivation, but it's determination and commitment to an unrelenting pursuit of your goal - a commitment to excellence - that will enable you to attain the success.

HEAD OF THE DEPARTMENT



The Department of Electronics and Communication Engineering at Mangalam College of Engineering, established in 2002, is committed to academic excellence, innovation, and research. The department offers undergraduate and postgraduate programs in VLSI and Embedded Systems, supported by experienced faculty. Through internships, projects, workshops, seminars, and industry collaborations, students gain practical knowledge and stay updated with the latest technologies. The department also encourages cultural, sports, and extracurricular activities, nurturing skilled, innovative, and socially responsible professionals.



AI AT THE EDGE

How tiny devices are ditching the cloud and learning to think for themselves.

**ASWIN PT
S7-ECE**

THE PROBLEM WE ALL KNOW TOO WELL

Think about what happens when you ask your smart speaker to set a ten-minute timer. The cylinder on your kitchen counter has absolutely no idea what you just said. It is essentially a hollow shell with a microphone. It records your voice, digitizes the audio, and sends it hundreds of miles away to a warehouse-sized data center. There, racks of power-hungry servers process your request, generate a response, and send it all the way back to your kitchen. This is how almost all modern AI works. It is fundamentally centralized. ChatGPT, Siri, Netflix recommendation algorithms—they all rely on the brute force of the cloud. But this architecture has a fatal flaw: it is slow, incredibly vulnerable to internet outages, terrible for privacy, and utterly ruinous for the environment. As we try to make everything around us "smart"—from our pacemakers to our city grids—the cloud paradigm is hitting a wall.

You simply cannot stream high-definition video from ten thousand forest fire cameras to a central server 24/7 without melting the network. The solution? Stop moving the data to the brain, and start moving the brain to the data.

OKAY BUT WHAT ACTUALLY IS IT?

Welcome to Edge AI, heavily driven by a subfield known as TinyML (Tiny Machine Learning). Instead of relying on a distant server farm, Edge AI puts the neural network directly onto the device itself, usually onto a cheap, ubiquitous microcontroller. Historically, this was considered impossible. Modern neural networks are colossal math problems, requiring gigabytes of memory and watts of power.



Microcontrollers—the cheap chips that run your microwave or electric toothbrush—often have less memory than an early-90s floppy disk and run on a fraction of a milliwatt. But researchers realized that AI models, much like the human brain, are incredibly redundant. Through processes called "quantization" and "pruning," engineers learned to ruthlessly trim the connections in a neural network that weren't doing much work, and reduce the precision of the underlying math. The result is a miniaturized AI model that can fit into a 20-kilobyte file and run for a year on a single watch battery.

WAIT, WHY DO WE SUDDENLY NEED THIS?

The push for Edge AI isn't just a neat trick; it solves three massive bottlenecks in modern engineering:

- Latency: If an autonomous drone is navigating a collapsed building to find earthquake survivors,

it cannot wait 300 milliseconds for a cloud server to tell it to dodge a falling beam. The processing must happen locally, in real time.

- Bandwidth: A smart vibration sensor on a wind turbine generates terabytes of acoustic data every week. Sending all that raw noise over a cellular network is prohibitively expensive. Edge AI allows the sensor to listen locally, discard the normal noise, and only transmit a tiny 2-byte alert when it hears the specific acoustic signature of a failing bearing.
- Privacy: Do you want a camera in your elderly parent's home monitoring them for a fall if that video feed goes to a corporate cloud? Edge AI analyzes the video on the camera chip itself. If a fall occurs, it sends a text alert. The video is instantly deleted and never leaves the room.

"We are moving away from a world where everyday devices are merely dumb appendages of the cloud. They are learning to perceive, understand, and decide—all on their own, on a fraction of a milliwatt of power."

WHAT CAN YOU ACTUALLY DO WITH THIS?

Here is where the technology moves from theory to world-altering applications:

- Conservation: Deep in the Amazon, tiny acoustic monitors are taped to trees. They listen for the specific mechanical whine of illegal chainsaws.

They don't stream the audio; they run an acoustic neural network locally. When a chainsaw is detected, they send a silent, low-power radio ping to rangers miles away.

- Wearables that actually save lives: Current smartwatches rely heavily on your phone to process complex health metrics. Edge AI enables a wearable to locally run complex anomaly detection on your heart rhythm, predicting an arrhythmia in realtime even if you are hiking without cell service.
- Smart infrastructure: As signals and networks become more complex, the physical environment must adapt. Similar to how engineers use Reconfigurable Intelligent Surfaces (RIS) to turn walls into smart mirrors that steer Wi-Fi [cite: 5], Edge AI can be embedded into concrete pilings to constantly calculate structural integrity and fatigue, turning dumb bridges into self-diagnosing structures.

THE GENUINELY COOL PART: PRIVACY BY PHYSICS

For years, tech companies have promised us that our cloud data is safe because it is encrypted and protected by strict privacy policies. But a policy is just a promise, and promises can be broken by hackers, governments, or corporate acquisitions. Edge AI offers something fundamentally better: privacy by physics. If a smart hearing aid uses a local neural network to isolate the voice of the person sitting across from you while filtering out restaurant noise, the raw audio is processed and instantly discarded in milliseconds.



. If the device literally lacks the broadband antenna required to transmit audio files to the internet, it cannot be hacked to leak your conversations. It is a physical guarantee.

OKAY, SO WHAT'S THE CATCH?

Building for the edge is brutally difficult. We have spent the last two decades enjoying software abundance. If an app was poorly coded or bloated, developers just rented more cloud servers or waited for users to buy a faster phone. "You can't do that on a 32-bit microcontroller," says Dr. Julian Vance, an embedded systems researcher. "Edge AI forces coders to become craftsmen again. You have to understand the physical hardware. You are fighting for every single byte of memory and every microjoule of battery power."

It is a completely different discipline than writing a web app."

WHY THIS MATTERS ESPECIALLY FOR US

For students in the ECS department, Edge AI represents a massive pivot in the tech industry. The era of purely software-based innovation is stabilizing, but the intersection of hardware and artificial intelligence is wide open. The skills required here—low-level C programming, electrical engineering, digital signal processing, and machine learning architecture—are the exact disciplines being taught in our labs right now. The engineers who will define the next decade aren't just writing algorithms; they are the ones figuring out how to squeeze those algorithms onto a piece of silicon the size of a dust mite.

WHERE'S THIS ALL GOING?

The era of massive, roaring data centers is far from over—we will always need the cloud to train the massive models that power generative AI and complex research. But the *deployment* of that intelligence is rapidly moving to the edge. We are approaching a state of ambient intelligence. Computing is going to vanish into the background. Your environment will simply know things. The bandages of the future will know if a wound is becoming infected. The factory floor will feel its own vibrations and know when it needs maintenance.

THE TAKEAWAY

For decades, we treated computers as destinations—screens we looked at, boxes we typed on. Edge AI breaks the computer apart and scatters its intelligence into the physical world. Devices will no longer ask a distant server for permission to understand the world. They will feel it, process it, and act on it locally, running silently on the ambient energy of a single ray of sunlight. The next great leap in human technology won't be measured in exabytes; it will be measured in millimeters.

s.





THE POST-BATTERY PARADIGM

How Ambient Energy Harvesting Could Change the Design of Connected Devices


**ADARSH P BABU
S3-ECE**

The Internet of Things is expanding into places where ordinary computers cannot easily go. Sensors are being placed in bridges, factories, farms, buildings, medical equipment, and remote ecosystems. Many of these devices need to measure temperature, pressure, motion, light, vibration, or chemical conditions and communicate their findings with little human attention. That creates a difficult practical problem. A sensor may be inexpensive and physically small, but its battery still has to be manufactured, installed, monitored, replaced, and eventually discarded. Replacing one battery is manageable. Replacing batteries in thousands or millions of remote devices is a continuing logistical burden. Researchers are therefore developing ambient energy harvesting: the conversion of naturally available energy—such as light, vibration, heat, or radio-frequency radiation—into the small amounts

of electricity required by an electronic system. The goal is not to power every device with a room lamp or a passing radio signal. The goal is to design devices that use very little energy and wake up when their environment provides enough of it.

WHY BATTERIES BECOME DIFFICULT AT SCALE ?

Batteries remain essential for phones, laptops, vehicles, medical devices, and countless other products. They provide a useful combination of energy density, portability, and predictable output. They are not, however, maintenance-free or infinitely durable. Electrochemical cells age through chemical and mechanical changes, and their useful life depends on temperature, charge cycles, current demand, storage conditions, and the design of the device.



. A device inside industrial machinery may require a shutdown before it can be serviced. A network deployed across a forest, ocean, or agricultural field may contain locations where regular battery replacement is expensive or unsafe. Batteryless design addresses this problem by changing the system's assumptions. Instead of expecting a continuous supply of stored energy, the device is designed around the energy that is actually available. A small solar cell may support a sensor during daylight. A vibration harvester may operate when a machine is running. A thermal generator may use a temperature difference. An RF harvester may collect energy from nearby transmitters, although the available power usually decreases rapidly with distance. The result is not unlimited operation. It is maintenance reduction through energy-aware engineering.

FOUR WAYS TO COLLECT AMBIENT ENERGY

- **SOLAR ENERGY HARVESTING (PHOTOVOLTAICS):** Captures ambient light to generate electricity.
- **THERMALENERGY HARVESTING (THERMOELECTRIC GENERATORS):** Converts temperature differences and heat gradients into electrical power.
- **KINETICENERGY HARVESTING (PIEZOELECTRIC SENSORS):** Translates ambient vibrations, movement, or mechanical stress into voltage.
- **RADIOFREQUENCY HARVESTING(RECTENNAS):** Collects ambient electromagnetic waves from cell towers, Wi-Fi, and radio signals to produce small amounts of direct current.


INTERMITTENT COMPUTING OR TRANSIENT COMPUTING

The central challenge of a batteryless device is intermittency. Sunlight changes, machinery stops, people move, and radio signals fluctuate. A conventional computer assumes that power will remain available while its processor is executing instructions and while its memory is holding data. An intermittent computer cannot make that assumption. Intermittent computing treats power loss as a normal operating condition. A microcontroller may wake up, measure a sensor, perform a small calculation, save a result, and transmit only when sufficient energy has accumulated.

If the voltage begins to fall, the system can save its progress before shutting down. When energy returns, the device resumes from a known checkpoint instead of starting the entire task again. Non-volatile memories such as FRAM and MRAM are useful in this context because they can retain information without continuous power. However, no single memory technology solves every problem. Designers must consider endurance, speed, energy consumption, cost, storage capacity, and how frequently the system may lose power. Software is just as important as hardware. Tasks must be divided into units that can safely pause. Checkpoints must be written in a way that prevents corrupted state. Sensor readings and messages may need sequence numbers so that a device can recover after an interrupted transmission. The result is a different programming model: the system is Temperature difference Thermoelectric generator Industrial equipment, human-worn devices, hot and cold surfaces Needs a sustained temperature gradient Radio-frequency signals Antenna and rectifier, or rectenna Near transmitters and wireless infrastructure Ambient RF power is often very small and variable designed not for uninterrupted execution, but for progress across many short windows of available energy

COMMUNICATING WITHOUT A POWER-HUNGRY TRANSMITTER

Communication can consume a significant share of a tiny device's energy budget. Generating a radio signal normally requires oscillators, amplifiers, frequency



control, and other active circuitry. A batteryless sensor may not have enough stored energy to run these components continuously. Ambient backscatter offers a different approach. Instead of generating its own carrier wave, a device changes how it reflects existing radio-frequency signals. By switching its electrical impedance, the device can alternate between absorbing and reflecting a small amount of the surrounding signal. A receiver detects those changes and reconstructs the encoded information. The University of Washington's original ambient-backscatter research described the technique as a way for battery-free devices to use existing television and cellular transmissions as both an energy source and a communication medium. 1

The approach does not provide the same data rate or range as a conventional Wi-Fi access point. Its appeal is that a simple device can send

small amounts of telemetry without operating a full radio transmitter. As with all energy-harvesting systems, the environment determines the result. The distance from the RF source, the orientation of the antenna, interference, the receiver's sensitivity, and the time available for harvesting all matter. Ambient backscatter is best suited to small messages and carefully designed deployments rather than high-bandwidth communication.


WHERE BATTERYLESS SYSTEMS COULD BE USEFUL

The most promising applications are those in which the device collects small amounts of data over a long period and where maintenance is difficult. Structural-health sensors could monitor strain or vibration in bridges, towers, railways, or industrial equipment. A batteryless node does not need to transmit continuously; it may wake when a vibration pattern suggests an unusual event and then report a compact alert. Environmental monitoring is another strong use case. Solar-powered nodes can measure soil moisture, temperature, water quality, or animal activity in places where routine visits are costly. A network may combine local energy harvesting with low-power mesh communication so that each device forwards only the information that matters. Medical applications require particularly careful distinction between possibilities and current practice. External energy transfer, including approaches based on ultrasound or electromagnetic fields,

is being studied for some implantable and wearable systems. These techniques may reduce the need for frequent procedures in selected applications, but they must meet demanding requirements for safety, reliability, biocompatibility, and clinical approval. A batteryless implant is not automatically safer than a battery-powered one

THE LIMITS OF A POST BATTERY WORLD

The phrase “post-battery” can be misleading. Ambient harvesting does not eliminate the need for energy storage in every system. A small capacitor, rechargeable microbattery, or other storage element may still be needed to accumulate energy between harvest events. Many devices will also continue to use conventional batteries when they need high power, long periods of operation without harvesting, or predictable performance





Energy harvesting has limits of scale as well. A tiny sensor may operate on microwatts or milliwatts, while a wireless transmitter, motor, display, or artificial-intelligence accelerator may need far more. Designers must reduce the workload before they can expect a small harvester to support it. This is why batteryless systems often emphasize local filtering, event detection, compressed messages, and low-power hardware. The environmental benefits also depend on the complete life cycle of the device. A batteryless sensor may reduce replacement waste, but it still contains semiconductors, packaging, wires, antennas, and other materials. Responsible design includes repairability, recovery, recycling, and the ecological effects of deploying large numbers of electronic devices

WHAT STUDENT NEED TO LEARN

Ambient-powered computing brings together subjects that are often taught separately. Students need to understand circuits and semiconductor devices, but they also need knowledge of electromagnetic fields, materials, embedded software, wireless protocols, control systems, data management, and environmental measurement. The most important design habit is to begin with an energy budget. How much energy is available? How often does it arrive? How much can be stored? What is the minimum useful task? How much energy does sensing consume compared with computation and communication? These questions may determine the architecture before a processor or sensor is selected. Higher education can prepare students for this work through laboratories that combine energy harvesting with real embedded systems. A useful project might ask students to build a sensor that runs from indoor light, records data in non-volatile memory, survives repeated power failures, and transmits only when its energy reserve is sufficient. Such a project would teach physics, electronics, programming, measurement, and systems thinking at the same time.

COMPUTING DESIGNED FOR SCARCITY

Ambient energy harvesting is not a promise that electronics will operate without limits. It is a design philosophy for places where energy is small, intermittent, and expensive to maintain.

The technology asks engineers to replace assumptions of abundance with careful decisions about timing, storage, computation, and communication. The result could be a new generation of devices that are quieter, smaller, and less dependent on scheduled battery replacement. Some may live inside infrastructure, some in fields and forests, and some on the human body. They will not perform every task a laptop can perform. They will perform a narrower task, repeatedly and efficiently, using whatever energy their surroundings provide. The post-battery future is therefore not a future without batteries. It is a future in which many devices no longer need to depend on a battery alone.



THE HIDDEN SYMPHONY OF FREQUENCIES

How AI and Radar are Teaching
Architecture to See

AYANA ANTONY
S5-ECE

Deep beneath a university engineering building, researchers are conducting an experiment that seems almost impossible. In a completely dark, windowless laboratory, a woman stands motionless while a small device mounted on a tripod observes her without using a camera, light, or thermal sensor. Instead, the device sends invisible millimeter-wave signals through the room and analyzes the echoes reflected from her body. Artificial intelligence then transforms these complex signals into a real-time three-dimensional representation of her skeleton. Every movement is detected—even the subtle expansion of her ribcage as she breathes and the rhythm of her heartbeat. This remarkable technology represents a new generation of integrated sensing, where machines can understand and map the physical world without relying on a single image or pixel.



For a century, our technological vision has been chained to the camera—a flawed, fragile mimicry of the human eye. But the era of optical dominance is fracturing. We are entering a new paradigm of machine perception, where the built environment does not need to look at us to know exactly what we are doing. We are teaching our technology to see like bats, and in the process, we are fundamentally rewriting the future of privacy, healthcare, and surveillance.


CHAPTER I: THE TYRANNY OF THE LENS

To understand the magnitude of this shift, we must first understand the limitations of the lens. Since the invention of the camera obscura, human engineering has obsessed over capturing light. The modern digital camera—whether it is mounted on a self-driving car, a drone,

or the bezel of your smartphone-operates on this exact same prehistoric principle. It requires a light source to bounce photons off an object, which then pass through a curved piece of glass to hit a sensor. But optical vision is inherently fragile. If a self-driving car turns a corner into the blinding glare of a setting sun, its cameras are blinded. If a dense fog rolls into a shipping port, optical sensors become entirely useless. Most importantly, cameras suffer from the fatal flaw of occlusion: they cannot see through walls, clothing, or even a sheet of paper.

Beyond its physical limitations, the camera has become the ultimate instrument of privacy invasion. As the Internet of Things (IoT) demands more environmental awareness to make our homes and cities "smart," the default solution has been to blanket our living spaces with optical lenses.

"We reached a terrifying inflection point where providing better healthcare to an elderly person meant placing a high-definition camera in their bathroom," explains Dr. Aris Thorne, a simulated composite of leading researchers in computational radio frequencies. "It was a dystopian trade-off. We realized that if we wanted machines to be deeply integrated into our intimate spaces, we had to invent a way for them to understand context without ever seeing our faces."The solution was not to build a better camera. It was to abandon light entirely.




CHAPTER II: THE PHYSICS OF THE UNSEEN BOUNCE

The alternative to light has been surrounding us all along. Our modern world is an invisible, chaotic ocean of electromagnetic waves—from the Wi-Fi routers broadcasting internet to our laptops, to the cellular towers pinging our phones, to the radio stations transmitting music. Historically, we have only used these radio frequencies (RF) to carry data. But every time an RF wave travels through a room, it interacts with the physical objects inside it. When a Wi-Fi signal hits a wall, some of it passes through, and some of it bounces back. When it hits a human body—which is essentially a giant bag of saltwater the wave is absorbed, scattered, and altered in highly specific ways.

Integrated sensing leverages two primary branches of this physics:

- **Millimeter-Wave (mmWave) Radar:** Unlike traditional weather radar, which uses massive dishes to bounce low-frequency radio waves off clouds miles away, mmWave radar uses very high frequencies (typically between 60 and 77 GHz) over very short distances. These waves are microscopic. Because of their tiny size, they are incredibly sensitive to micro-movements. When a mmWave bounces off a human chest, the millimeter-scale displacement of the skin caused by a heartbeat alters the frequency of the returning wave—a phenomenon known as the Doppler effect.



- **Ambient Wi-Fi Sensing:** This is perhaps the most magical, and unsettling, frontier. You do not even need dedicated radar emitters. By simply analyzing the standard Wi-Fi signals bouncing between the router in your living room and your smartphone, engineers can map the environment. As you walk across the room, you disturb the invisible electromagnetic field of the Wi-Fi. By measuring these disruptions—specifically looking at the Channel State Information (CSI)—systems can detect if someone has fallen down, if they are typing on a keyboard, or if they are simply breathing in their sleep.

CHAPTER III: THE BRAIN BEHIND THE ECHO

If radar has existed since World War II, why are we only now using it to track heartbeats and map human skeletons? The answer lies in the explosive, parallel evolution of Artificial Intelligence.

For decades, the problem with radar was the "noise." If you emit a radar pulse in a standard living room, the wave doesn't just bounce off the person standing in the middle. It bounces off the ceiling, the television, the metallic edge of a coffee table, and the glass of a window. Some waves bounce off the person, hit the wall, hit the ceiling, and then return to the sensor. This is known as "multipath interference," and to a traditional computer, it looks like meaningless, chaotic static.

"Trying to see a human being in raw radar data is like trying to identify a specific drop of water in a hurricane," says Dr. Elena Rostova, a simulated pioneer in the field.

This is where deep learning enters the narrative. Over the past five years, researchers began feeding massive datasets of this chaotic RF static into neural networks, pairing the radar data with optical video feeds. The AI was given a simple task: look at the camera feed to see where the human joints are, and then look at the chaotic radar static and find the hidden patterns that match those joints.



After millions of iterations, the AI learned to filter out the noise of the walls, the furniture, and the multipath bounces. It learned the specific, microscopic radar signature of a human hand waving, a knee bending, or a lung expanding. Today, once the neural network is fully trained, the optical camera is entirely removed. The AI can now hallucinate a perfect 3D skeleton from nothing but the invisible echoes.

CHAPTER IV: THE ARCHITECTURE OF EMPATHY (AND SURVEILLANCE)

The implications of seeing without cameras are poised to fundamentally rewrite the curriculum of modern engineering and computer networking. In healthcare, the applications are profound and immediately life-saving. ten or taken off during

A mmWave radar sensor placed behind the drywall of a bathroom can monitor a patient 24/7. It cannot see the color of their skin, their state of undress, or their identity. But if the physical mass of the patient rapidly descends to the floor, the AI instantly recognizes the signature of a fall and alerts the nurses' station.





We are building a world where the very walls of our homes, offices, and universities act as omniscient sensors. If a standard Wi-Fi router can be upgraded via a firmware patch to detect the emotional states of the people in the room based on their respiration rates, who owns that data? If an authoritarian government deploys Wi-Fi sensing in a public square, they do not need facial recognition cameras to know exactly how many people are gathered in a windowless basement, or what their precise physical movements are.

The legal frameworks of the 20th century are entirely unprepared for the physics of the 21st. Our privacy laws are largely based on the concept of a "reasonable expectation of privacy" behind closed doors. In automotive engineering, integrated sensing is replacing optical cameras for in-cabin monitoring. A radar hidden in the dashboard

can easily "see" through a thick winter coat to monitor the subtle breathing patterns of a driver, detecting the onset of micro-sleep or a medical emergency before the driver's eyes even close. It can detect if a child has been left under a blanket in the backseat, sensing their tiny heartbeats even if they are entirely hidden from view.

However, a technology that can see through walls and detect a heartbeat from across the room is inherently a double-edged sword.

But what does a closed door mean to a machine that sees in radio waves?

THE "SO WHAT?": YIELDING TO THE INVISIBLE

The transition from optical cameras to integrated RF sensing represents a profound philosophical pivot. For centuries, our definition of "surveillance" and "awareness" has been intrinsically tied to the eye. We believed that to understand the world, we had to capture its light, freeze it, and analyze it.

As we move forward, the built environment will no longer watch us with cold, glass lenses. Instead, architecture will become deeply, invisibly sentient. Our smart homes and intelligent classrooms will not look at us; they will feel our presence, sensing the displacement of the air, the rhythm of our pulses, and the geometry of our movements. Technology is finally learning to close its eyes, realizing that in the dark, the echoes tell a far more honest story.



The next generation of mobile networks will not be judged only by how quickly a phone downloads a file, but by how intelligently the network connects machines, people, sensors, and physical spaces.

For much of its public life, 5G has been advertised as a faster way to stream video and download apps. Those benefits matter, but they are only part of the story. The deeper change is architectural: modern mobile networks are being designed to support industrial machines, autonomous systems, low-latency control, dense sensing, cloud services at the network edge, and software-defined infrastructure. At the same time, researchers are looking beyond 5G toward 6G, often called IMT-2030 in the international standardization process. The goal is not simply to add another speed increase. It is to connect communication with sensing, artificial intelligence, immersive applications, sustainability, and wider geographic coverage.

BEYOND THE SIGNAL

How 5G Is Building the Road to 6G

BLESSON BIJU
S3-ECE

WHY 5G USES MORE THAN ONE TYPE OF SIGNAL

Radio waves behave differently at different frequencies. Lower-frequency signals generally travel farther and pass around obstacles more effectively, making them useful for broad coverage. Higher-frequency signals can provide wider channels and greater capacity, but they tend to lose energy more quickly and can be more affected by walls, foliage, rain, and other obstructions.

5G therefore uses a layered approach rather than one universal frequency band. Sub-6 GHz networks provide much of the coverage and mobility that users expect from cellular service. Mid-band spectrum can offer a useful balance of coverage and capacity. Millimeter-wave systems can deliver very high throughput over shorter distances and are most practical in locations such as stadiums, transport hubs, dense urban areas, and enterprise sites. This trade-off is not a temporary engineering inconvenience. It is a basic property of electromagnetic propagation. As frequency rises, network designers often need more antennas,

smaller cells, careful beam management, and better planning of where access points are installed.

Network slicing allows a physical network to be divided into logical networks with different policies and performance objectives. One slice might prioritize ordinary mobile broadband, while another is configured for a private enterprise service or a more demanding industrial workload. Isolation is implemented through software and network management; it is not the same as building a completely separate physical network.

FROM FASTER PHONES TO INDUSTRIAL SYSTEMS

The most important 5G developments are not visible in a speed-test application. They involve the network's ability to support different types of service on shared infrastructure.

Ultra-Reliable Low-Latency Communication, or URLLC, is intended for applications that need predictable and dependable delivery. Industrial control, robotics, transportation, and other time-sensitive systems may benefit from reduced delay and improved reliability.

However, a network label does not guarantee that every application will experience the same latency. Performance depends on radio conditions, traffic, processing location, software design, and the complete end-to-end path.



Multi-access edge computing, or MEC, places computing resources closer to users and devices. If data can be processed at a nearby edge server instead of traveling to a distant data centre, some round-trip delay and backhaul traffic can be reduced. MEC does not eliminate latency, but it can make time-sensitive services more practical when the application and network are designed together.

These capabilities help explain why 5G is relevant to manufacturing. Autonomous mobile robots, automated guided vehicles, cameras, sensors, and control systems can share a managed wireless environment. In some factories, wireless systems may complement or replace selected wired connections, but critical applications still require careful testing, redundancy, safety engineering, and predictable failure behavior.

6G does not yet have one final global specification. The International Telecommunication Union's IMT-2030 framework describes several possible usage scenarios, including immersive communication, hyper-reliable and low-latency communication, massive communication, ubiquitous connectivity, artificial intelligence and communication, and integrated sensing and communication. [cite: 1]

One area of research involves sub-terahertz frequencies, including bands around 90-300 GHz. These frequencies could provide very wide channels for short-range links, such as wireless connections among data-centre computing units, immersive communication inside buildings, or high-resolution industrial applications.



WHAT 6G IS EXPECTED TO ADD

Ericsson's research overview emphasizes that lower-frequency bands will remain important for coverage and mobility, while sub-THz links are likely to serve specialized local scenarios because of their limited range and difficult propagation conditions. [cite: 3]

The higher the frequency, the more carefully a network must manage the path between transmitter and receiver. A person, wall, vehicle, or even a small obstruction may weaken a link. Networks may therefore rely on highly directional beams, fast beam tracking, large antenna arrays, and intelligent placement of access points. Antennas that adapt to the environment

Massive multiple-input, multiple-output systems-known as massive MIMO-use many antenna elements to shape and direct radio energy. Beamforming allows the network to concentrate a signal toward a device rather than radiating equally in every direction. In future systems, beam management may need to become faster and more precise as frequencies increase and coverage areas become smaller.



Another research direction is the intelligent reconfigurable surface. A programmable surface placed on a wall or building façade could alter how radio waves reflect from it. In principle, this could help redirect signals around difficult areas or improve coverage in indoor spaces. Such surfaces remain an active research topic, and real-world deployment would involve cost, control, power, installation, and safety considerations.


These technologies do not make obstacles disappear. They give the network more ways to respond to them. The challenge is to coordinate antennas, surfaces, devices, and software quickly enough to maintain a useful connection while people and machines move.



WHEN THE NETWORK BECOMES A SENSOR

Traditional wireless communication is designed to deliver information from one point to another. Integrated sensing and communication, or ISAC, proposes using some of the same signals for both communication and environmental sensing. A transmitted signal reflects from people, vehicles, walls, machines, and other objects. By analyzing changes in the signal's timing, phase, frequency, and direction, a system may estimate movement, distance, or other characteristics of the environment. A future network could use this information to support vehicle

navigation, industrial monitoring, gesture recognition, occupancy detection, or other applications without relying entirely on optical cameras.





PCIe Gen 5 SSDs

The Race for Speed-and the Heat It Creates

GOPIKA MOHAN
S5-ECE

A new generation of storage
For decades, computer storage was defined by mechanical limits. Hard-disk drives used spinning platters and moving read/write heads, which made them affordable and capacious but relatively slow. Solid-state drives changed that equation by replacing moving parts with NAND flash memory. The next major step came when NVMe drives began using the PCI Express bus-the same high-speed connection used by graphics cards-to communicate with the processor.

WHAT IS INSIDE A GEN 5 SSD?

PCIe Gen 5 is the latest stage in that progression for mainstream consumer SSDs. It does not make storage as fast as system memory, but it gives the drive a much wider path to the rest of the computer. PCI-SIG lists PCIe 5.0 at 32.0 gigatransfers per second per lane, twice the signaling rate of PCIe 4.0.

A four-lane PCIe 5.0 connection therefore has a theoretical raw bandwidth of about 16 GB/s in each direction, before protocol overhead and other system limitations are considered.

That extra bandwidth is impressive. It also creates a new engineering challenge: how can a very small circuit board move so much data without becoming too hot? A typical M.2 NVMe SSD contains NAND flash memory packages, a controller, firmware, and sometimes a DRAM cache. NAND stores data even when power is removed. The controller coordinates reads and writes, manages wear-leveling and garbage collection, performs error correction, and maintains the flash translation layer-a mapping system that connects logical addresses used by the operating system with physical locations in NAND.

The PCIe interface determines how quickly data can travel between the SSD and the computer. The controller and NAND determine how effectively the drive can use that interface. In other words, a fast PCIe link is necessary for high performance, but it is not sufficient by itself.



The drive also needs capable flash, firmware, cooling, and a system that supports the required PCIe generation and lane configuration.

WHY COOLING MATTERS

High performance requires electrical power, and electrical power eventually becomes heat. A small M.2 drive has far less surface area than a graphics card, yet its controller may be handling error correction, address translation, queue management, and data transfer at the same time. During short benchmarks, a drive may reach its advertised peak speed. During a long file copy or sustained video workload, heat can accumulate and the firmware may reduce performance through thermal throttling.

WHO BENEFITS FROM THE EXTRA SPEED?

This behavior is protective rather than catastrophic. A Gen 5 SSD does not normally "melt" under ordinary use. A suitable motherboard heatsink,

good case airflow, and sensible workload management are generally more realistic recommendations than a dedicated high-speed fan. Samsung's 9100 PRO, for example, is specified as a PCIe 5.0 x4, NVMe 2.0 M.2 2280 drive with advertised sequential read speeds of up to 14,700-14,800 MB/s depending on the listed capacity and configuration.

CONSIDER SUSTAINED PERFORMANCE

Cooling is especially relevant because benchmark numbers are often measured during limited test runs. A drive that reaches a high score for a few seconds may behave differently during a 30-minute transfer. For this reason, students comparing SSDs should look beyond the headline sequential-read number and consider sustained performance, power consumption, warranty, endurance, and the quality of the included heatsink.

For many students, a PCIe Gen 4 SSD is already fast enough for everyday tasks such as starting an operating system, opening documents, browsing the web, and launching applications. Moving to Gen 5 may improve benchmark results without producing a dramatic change in ordinary use. The benefits become easier to see when a workload can use high sequential bandwidth for long periods.

STORAGE IS NOT MEMORY

Professional video production is one example. High-resolution footage can involve large files and several simultaneous streams. A fast SSD can provide useful headroom when importing media, generating previews, moving project files, or working with large cache directories. However, the SSD is only one part of the system. The codec, processor, graphics processor, memory capacity, editing software, and source-media layout all affect the result. A Gen 5 drive does not automatically eliminate the need for proxies or optimized media. Local artificial-intelligence workloads provide another possible use. Reading a large model file from storage can take less time when sequential throughput is higher.



Yet model loading may also involve decompression, validation, memory allocation, and movement into system or graphics memory. Faster storage can shorten part of that process, but it does not by itself determine the response time of an AI application.

Future technologies may blur the boundary between processing, memory, and storage. Compute Express Link, for example, is designed to support coherent communication among processors, memory, and accelerators. Developments such as CXL and newer PCIe generations could influence how future systems organize data, but they should be described as architectural possibilities rather than as proof that SSDs are about to replace RAM.

IS GEN 5 WORTH IT FOR STUDENTS?

Software development, scientific computing, virtualization, and large dataset processing can also benefit from higher storage throughput. These workloads may involve repeatedly reading and writing large files or managing many I/O requests at once. In such cases, the improvement is more meaningful than it is for a simple application launch. Gaming is a more complicated case. Microsoft's DirectStorage APIs are designed to reduce operating-system overhead, support high-throughput NVMe storage,

and reduce CPU work associated with asset loading and decompression. Some implementations can use hardware-accelerated decompression, including GPU-based techniques. Nevertheless, game loading and world streaming depend on the game engine, compression format, CPU, GPU, memory, and developer implementation.



The answer depends on the workload and the rest of the computer. A student building a new workstation for video editing, virtualization, engineering software, or large data projects may value the additional bandwidth. A student upgrading an ordinary laptop or desktop may gain more from increasing RAM, choosing a larger SSD, improving backup practices, or replacing an older hard drive than from paying a premium for Gen 5.

Compatibility also matters. The motherboard must provide a PCIe 5.0 M.2 slot for full Gen 5 performance. A Gen 5 SSD can usually operate in an older PCIe slot, but it will run at the older slot's limits. The system should also provide adequate cooling, particularly when the drive will handle sustained writes or repeated large transfers.

es so.

The race is no longer simply about making storage faster. It is about making the entire system capable of using that speed efficiently-and keeping the hardware cool while it do





THE PATH FORWARD

As networks become more heterogeneous, manual configuration becomes increasingly difficult. A future system may include terrestrial base stations, private networks, satellites, high-altitude platforms, edge servers, and billions of devices. The network must manage changing traffic, interference, energy conditions, mobility, and application requirements.

Machine learning can assist with tasks such as beam selection, traffic prediction, anomaly detection, radio-resource allocation, and energy management. Researchers sometimes describe this direction as an AI-native network—one in which artificial intelligence is considered during the design of the radio interface and network architecture rather than added only as a separate application.

NEXTGEN CONNECTIVITY

DEVIKA RAJ
S7-ECE

ISAC will not automatically provide a complete, high-resolution replacement for cameras or radar systems. Accuracy depends on frequency, antenna geometry, signal bandwidth, algorithms, environmental conditions, and the amount of interference. It also raises important questions about consent and privacy. A system that can sense movement through a space may collect information about people even when it does not record conventional video.

The phrase should be used carefully. A network that applies machine learning is not necessarily autonomous or reliable in every situation. AI models can fail when conditions change, consume computing resources, reproduce biased assumptions, or make decisions that are difficult to explain. Any AI-managed network will still need monitoring, security controls, fallback procedures, and human oversight.

The educational challenge


The evolution from 5G to 6G is also changing what students need to learn. Telecommunications engineering can no longer be separated completely from cloud computing, artificial intelligence, cybersecurity, hardware design, data governance, and environmental sustainability.

A student preparing for next-generation network research may need to understand electromagnetic propagation, antenna systems, signal processing, embedded programming, distributed computing, and machine-learning optimization. Just as importantly, the student must understand how technical choices affect privacy, energy consumption, reliability, accessibility, and public trust.

Universities can support this interdisciplinary approach through laboratories that combine software-defined radios, edge computers, sensor platforms, and network simulators. Projects could ask students to design a low-latency industrial link, measure how obstacles affect a high-frequency signal, or build a sensing system that processes data locally while protecting personal information.

The future of mobile networking will not be defined by a single frequency band or one spectacular speed record. It will be defined by the coordination of many layers: broad-coverage cellular networks, dense local links, edge computing, intelligent antennas, satellites, sensors, and software that can adapt to changing conditions.

5G is already providing much of the architectural foundation for this transition. 6G research is extending the conversation toward new spectrum, integrated sensing, AI-native control, and more specialized forms of connectivity. Yet the most ambitious claims remain projections, not guaranteed features of a finished commercial network.



The central challenge is therefore not merely to make networks faster. It is to make them more useful, dependable, secure, efficient, and responsible. The next generation of connectivity will succeed when it becomes almost invisible to users while remaining carefully engineered underneath-connecting not only people to information, but machines to decisions and physical spaces to the digital world.

References

- [1] ITU: IMT towards 2030 and beyond (IMT-2030)
- [2] 3GPP: Edge Computing
- [3] Ericsson: Sub-terahertz communication in 6G

CHAPTER III: DESIGNING A METABOLIC PATHWAY FOR HARDWARE

The solution emerging from the world's most advanced laboratories is an entirely new class of materials known as transient electronics. Just as a student of pharmacology understands that a drug's efficacy is inextricably tied to its half-life—how cleanly a chemical compound is metabolized and excreted from the bloodstream—materials scientists are realizing that a circuit's brilliance is tied to its degradation. We must design an ecological metabolic pathway for our hardware. Instead of relying solely on heavy metals and indestructible plastics, researchers are synthesizing active components from organic and water-soluble materials:

- **Biodegradable Conductors:** Engineers are printing complex digital circuit diagrams using conductive inks made from magnesium, zinc, and even carbonized coffee grounds. These metals conduct electricity flawlessly but oxidize and dissolve into harmless biological minerals when exposed to continuous moisture.
- **Silk Substrates:** Instead of fiberglass, researchers are using crystallized silk proteins as the base layer for microchips. By tweaking the molecular structure of the silk, they can program the chip to remain stable for exactly one year, two years, or five years, before it spontaneously melts away.

- **Transient Scratchpad Memory:** Innovations in volatile memory are leading to storage systems that hold data robustly while powered, but safely decompose at the molecular level once their designated lifespan is reached, ensuring both ecological safety and absolute data security.

CHAPTER IV: THE BATTERY PROBLEM AND AMBIENT HARVESTING

Perhaps the greatest hurdle to sustainable electronics is power. A biodegradable circuit board is useless if it is tethered to a highly toxic, corrosive lithium-ion battery. To sever this dependency, engineers are leaning heavily into the principles of ambient energy harvesting. Rather than storing chemical energy, the next generation of IoT devices is designed to operate on trace amounts of power pulled directly from their environment. Imagine a structural integrity sensor embedded





deep inside the concrete pillar of a highway overpass. Instead of a battery that will die in three years, the sensor is equipped with a piezoelectric generator that converts the micro-vibrations of the cars driving overhead into electricity. Other devices utilize thermoelectric generators, powering their networking protocols purely from the temperature differential between the warm daytime air and the cool concrete. Some researchers are even developing antennas capable of harvesting stray radio frequencies—the invisible soup of Wi-Fi and cellular signals that surround us constantly—converting those ambient waves into enough voltage to power a transient microprocessor. It is the technological equivalent of a plant drawing energy from the sun.

THE "SO WHAT?": HARMONY OVER DOMINANCE

The transition toward green, transient electronics represents a profound philosophical shift in how humanity

interacts with the natural world. Since the Industrial Revolution, our approach to engineering has been one of sheer dominance. We believed that to build a better world, we had to pave over the old one with materials that would never yield, never bend, and never break.

But as our planet chokes on the remnants of our progress, we are learning a humbling lesson. True engineering mastery is not the ability to build something that lasts forever. It is the ability to build something that performs brilliantly, serves its purpose, and then gracefully disappears. When technology learns to die, we may finally learn how to live sustainably alongside it.



BIO ELECTRONICS


LAVANYA K
S7-ECE

High in the humid canopy of the Indonesian rainforest, a small plastic box is taped to the trunk of a dipterocarp tree. Inside this box is a microphone, a watch battery, and a silicon chip no larger than a lentil. For months, it sits in silence, ignoring the cacophony of macaws, the heavy patter of monsoon rain, and the rustling of spider monkeys. Then, a faint, rhythmic mechanical whine echoes miles away. It is the sound of an illegal chainsaw. A decade ago, a sensor like this would have blindly recorded the audio, draining its battery in days as it desperately tried to transmit massive audio files over a weak 3G connection to a server in California for analysis.

But this tiny plastic box does not stream audio. It does not wait for a server to tell it what it is hearing. The intelligence is baked directly into the lentil-sized chip. It processes the acoustic data locally, recognizes the acoustic signature of a two-stroke engine, discards the audio file to protect its battery, and sends a simple, two-byte radio ping to the rangers: Chainsaw. Sector 4. This is not the artificial intelligence of sci-fi blockbusters. It is not ChatGPT writing poetry, nor is it a supercomputer rendering deepfakes. It is something far quieter, yet arguably more revolutionary. It is "AI at the edge." We are moving away from a world where everyday


devices are merely dumb appendages of the cloud. Instead, thanks to a rapidly maturing field known as TinyML (Tiny Machine Learning), the microcontrollers embedded in our appliances, our clothes, our infrastructure, and our wilderness are waking up. They are learning to perceive, understand, and decide—all on their own, on a fraction of a milliwatt of power.

To appreciate the magnitude of this shift, you have to understand the sheer weight of the cloud. For the past fifteen years, the technology industry has operated on a centralized paradigm. Your smart speaker, your connected thermostat, and your smartphone's voice assistant are essentially hollow vessels. When you ask your kitchen speaker to set a timer, the device itself has no idea what you are saying. It merely digitizes your voice, sends it across an ocean of fiber-optic cables to a warehouse-sized data center, where thousands of power-hungry processors parse your request, generate a response, and send it all the way back.



This journey takes time. It requires a constant, robust internet connection. It poses massive privacy risks, as intimate audio and video data are constantly shuttled off-site. And terribly, it consumes a staggering amount of electricity. "The cloud paradigm is fundamentally unnatural," notes Dr. Aris Thorne, a composite representation of the embedded systems engineers pioneering this field. "Biological systems don't work this way. When you touch a hot stove, your hand doesn't send a signal to a central brain in another room and wait for permission to flinch. The processing happens locally, in your peripheral nervous system. Edge AI is about giving our devices a peripheral nervous system." The challenge, of course, is physics. Modern neural networks—the algorithms that power modern AI—are colossal math problems. They consist of millions, sometimes billions, of "parameters" or weights. Running these calculations requires gigabytes of memory and watts of power. Microcontrollers, the cheap, ubiquitous chips that run everything from microwave ovens to electric

toothbrushes, usually possess less memory than an early-90s floppy disk. Fitting a neural network onto a microcontroller was long considered as absurd as trying to park a commercial airliner in a bicycle rack. But researchers realized that AI models, much like the human brain, are incredibly redundant. Enter the art of quantization and pruning. Engineers discovered that they could take a massive, cloud-based neural network and ruthlessly prune the connections that weren't doing much work. Then, through quantization, they reduced the precision of the math. Instead of calculating a number to sixteen decimal places, they rounded it off. The AI became slightly less precise, but radically smaller and faster. Suddenly, a model trained to detect human heart arrhythmias, which once required a server rack, could be squeezed into a 20-kilobyte file. It could run on a \$2 chip, powered by a coin-cell battery for three years, without ever needing an internet connection.




The implications of untethering intelligence from the internet are profound, extending far beyond the realm of conservation. Consider the medical field. Wearable technology has promised a revolution in preventative care for years, but it has been hobbled by the privacy nightmare of the cloud. Do you really want a device recording your every heartbeat, your breathing patterns, and your vocal tremors, and sending that raw data to a corporate server?

Edge AI solves the privacy problem by design. A smart hearing aid equipped with a TinyML chip can filter out background restaurant noise and isolate the voice of the person sitting across from you. It analyzes the acoustic environment hundreds of times a second. But because the intelligence lives entirely on the device, the audio is never recorded. It is processed, acted upon, and instantly discarded. The data never leaves your ear.

"Privacy is no longer a policy you have to trust; it becomes a physical guarantee," explains Dr. Elena Rostova, a researcher in edge deployment. "If the device literally lacks the antenna or the bandwidth to transmit raw video or audio, it cannot be hacked in the traditional sense. It only transmits the conclusion—like 'patient has fallen'—not the video of the patient." This shift is quietly transforming heavy industry and agriculture as well. Factory floors are now deploying microscopic vibration sensors on the bearings of massive manufacturing robots. These sensors don't stream vibration data to a dashboard; they simply learn what a "healthy" vibration feels like. When the bearing begins to wear out—weeks before it audibly breaks—the sensor detects the subtle algorithmic anomaly and flags it for maintenance. In agriculture, edge-enabled drone swarms can fly over sprawling fields, processing images locally to identify crop blight plant by plant. Instead of crop-dusting an entire farm with pesticides, the drone triggers a targeted micro-spray only on the affected leaves. The intelligence is immediate, decentralized, and environmentally transformative.

For college students entering the fields of computer science, engineering, and design, the edge revolution demands a radical shift



in perspective. For the past two decades, software engineering has been characterized by abundance. If your app was too slow, you just rented more cloud servers. If your code was bloated, you relied on the user upgrading to a faster phone. Edge AI forces a return to elegance. It requires a deep, intimate understanding of hardware. It demands that coders become craftsmen again, optimizing their models down to the very last byte, working within extreme constraints of power and memory. It is a space where a clever mathematical shortcut can double the battery life of a pacemaker or a deep-sea monitor. Ultimately, the triumph of AI at the edge is about altering our fundamental relationship with the spaces we inhabit. We are used to computers being objects we explicitly interact with—screens we look at, keyboards we type on. But as microcontrollers become infinitesimally small, cheap, and intelligent, computing will vanish into the background.

"Privacy is no longer a policy you have to trust; it becomes a physical guarantee. The data never leaves your ear."



The bridge will know it has a micro-fracture. The streetlight will know to dim itself because no pedestrians have walked by in an hour. The bandages of the future will know if a wound is becoming infected. They will not ask a distant server for permission to understand the world. They will feel it, process it, and act on it locally, running on the ambient energy of a single ray of sunlight or the vibration of a passing car. The era of massive, roaring data centers is far from over, but the next great leap in human technology won't be measured in exabytes or supercomputers. It will be measured in millimeters. It will be found in the quiet, distributed awakening of the dust.




"When you touch a hot stove, your hand doesn't send a signal to a central brain in another room and wait for permission to flinch. Edge AI is about giving our devices a peripheral nervous system."



Flexible and stretchable substrate (SEBS)





THE SPEED OF GLASS

VIDHU MOHAN
S5-ECE

HOW WE ARE TEACHING LIGHT TO THINK

Stand on the shores of Virginia Beach or the rugged coastline of Cornwall, and you are just miles away from the physical internet. Buried deep in the freezing silt of the Atlantic floor are cables no thicker than a garden hose. Inside them, hair-thin strands of ultra-pure glass are carrying the sum total of human knowledge, commerce, and culture—encoded as pulses of light—across the globe at 186,000 miles per second. We have already entrusted the transmission of our world to the photon. Light is the ultimate courier. It is massless, it ignores electrical resistance, and it generates no heat. But follow that undersea cable onto dry land, into the deafening, air-conditioned roar of a modern hyperscale data center, and the elegance of light abruptly dies. To actually process all that data—to search a database, render a video, or train an artificial intelligence—we force the light to stop.

This translation—the Optical-Electrical-Optical (O-E-O) bottleneck—is the great traffic jam of the 21st century. It is the reason massive data centers consume as much power as small cities, bleeding most of it away as waste heat. But a quiet revolution is currently etching a new path across the silicon wafer. Physicists and engineers are no longer content to let light merely carry our messages. They are attempting to build the elusive "optical computer"—a machine where logic, memory, and calculation are performed entirely by beams of light.

To understand why computing with light is so revolutionary, you have to understand the fundamental personality differences between the electron and the photon. Electrons are fermions. In the language of quantum mechanics, this means they are antisocial. Two electrons cannot occupy the same state in the same space. If you push a billion electrons down a copper wire, they bump into the atomic lattice, they repel each other, and they generate friction. That friction is heat. Photons are bosons. They are, essentially, ghosts. You can shine two laser beams directly through each other, and they will emerge on the other side completely unaltered. Because they don't interact with each other, you can send dozens of different colors (wavelengths) of light down the exact same microscopic channel simultaneously, without any cross-talk. This is called wavelengthdivision multiplexing. "Imagine a highway where the red cars, blue cars, and green cars can all occupy the exact same lane, passing right through each other at the

speed of light, never crashing," says Dr. Linnea Vance, a former MIT researcher and leading architect of optical neural networks. "With electrons, you have a single lane of gridlocked traffic. With photons, you have infinite lanes in a single physical space."

If light is so superior, why hasn't it conquered the processor yet? Because the very property that makes photons such excellent couriers—their ghostliness—makes them terrible at computation. To compute, you need logic gates. You need one signal to dictate the behavior of another signal (e.g., "If A is active, then stop B"). But if photons just pass right through each other, how do you make them interact? How do you make light control light?


For decades, this seemed insurmountable. Optical computers existed as massive, room-sized academic experiments, relying on bulky lenses and exotic, nonlinear crystals that required immense power to operate.

The breakthrough came not from inventing a new material, but by hijacking the trillion-dollar infrastructure of the existing silicon industry. Researchers discovered how to trap light inside silicon itself. By carefully designing microscopic channels—

waveguides—they could bend light around corners, split it, and recombine it using structures called Mach-Zehnder interferometers. Instead of forcing photons to crash into each other, silicon photonics uses slight electrical tweaks to alter the refractive index of the glass as the light passes through. By splitting a beam of light in two, slowing one half down by a fraction of a fraction of a second, and recombining them, the light waves either amplify each other or cancel each other out entirely. On or Off. 1 or 0. We had finally built a transistor made of light.

The timing of this breakthrough could not be more critical. The rise of generative Artificial Intelligence has pushed traditional silicon processors to their absolute breaking point. Training an AI like a large language model requires a mathematical operation called matrix multiplication—essentially taking massive grids of numbers and multiplying them against other massive grids of numbers, billions of times a second. For an electronic processor, this requires moving vast amounts of data back and forth between memory and the CPU, generating crippling amounts of heat. Optical processors, however, perform matrix multiplication naturally, passively, and almost instantaneously.

When light passes through a properly tuned optical chip, the computation happens merely by the light propagating through the glass. The result of the math is read instantly at the other end by a photodetector. It takes the exact same amount of time and energy to multiply a two-by-two grid as it does a two-



thousand-by-two-thousand grid. The computation happens at the speed of light, because the computation is the light.

"We are essentially tricking physics into doing the math for us," Dr. Vance explains. "The AI isn't being calculated by a struggling machine. It is being continuously solved by the geometry of the glass." Startups and tech giants alike are now racing to integrate photonics into the modern server rack. The first generation of commercial optical chips isn't replacing the electronic CPU; rather, they serve as optical co-processors, taking on the heavy, heat-generating burdens of AI math while the traditional silicon handles the basic operating tasks. But the trajectory is clear. As the technology scales, the boundary between the fiber optic cable and the processor will blur, and eventually vanish entirely. This represents a profound poetry in the evolution of our technology. Since the dawn of our species, light has been our primary tool for perceiving the universe. We used it to navigate by the stars, to illuminate our manuscripts, to peer through microscopes at our own cells, and to capture the echoes of the Big Bang through space telescopes. Light has always been the medium through which we see. Now, in the microscopic architecture of the 21st-century silicon wafer, we are crossing a new threshold. We are no longer just using light to look at the universe. We are using light to think about it.





Figure 1



shutterstock.com - 2338813799



THE SILICON HORIZON: THE NEW ARCHITECTURE OF INTELLIGENCE

ELSA JOSE
S3-ECE

THE DEATH OF THE RIGID GRID

To understand where we are going, we must first confront the sheer stubbornness of our current infrastructure. Every smartphone in your pocket, every laptop humming on a dormitory desk, and every server rack powering global cloud networks relies on a design philosophy born in the mid-20th century: the planar silicon chip.

It is a marvel of human ingenuity, but it is fundamentally brittle. Silicon breaks when bent. It overheats when pushed beyond its thermal thresholds. And crucially, it forces us to shoehorn fluid, complex algorithms into a rigid, sequential pipeline.

"We've spent fifty years forcing electrons down narrow highways, acting surprised when they ,

traffic-jam and generate heat," explains Dr. Marcus Vance, a materials scientist pioneering non-silicon substrates. "Nature doesn't build processing units on flat, inflexible planes. Your nervous system is soft, distributed, self-healing, and remarkably energy-efficient. Our goal isn't just to make chips faster; it's to change the geometry of computation itself."

Enter the era of flexible and conformal electronics. By replacing traditional semiconductor substrates with carbon nanotubes, transition metal dichalcogenides (TMDs), and biocompatible polymers, researchers are engineering circuits that can stretch up to two hundred percent of their original length without losing conductivity.

Imagine a wearable medical patch that maps electrical activity across the human heart not as a rigid plastic rectangle taped to the chest, but as a seamless, breathable second skin that moves with every contraction of muscle fibers. This isn't theoretical physics anymore—it is active engineering happening right now in cleanrooms around the world.

COMPUTING LIKE THE MIND

For electronics and computer science students, the frustration with traditional architecture goes deeper than hardware physics—it lives right at the software-hardware boundary.


We are still bound to the von Neumann architecture, where the processor (CPU) and the memory (RAM) live on opposite sides of a physical bus. Every single clock cycle, data must shuttle back and forth across this narrow bridge, wasting precious nanoseconds and consuming massive amounts of power. It is the ultimate bottleneck of modern computing.

Neuromorphic engineering seeks to obliterate this barrier by mimicking the architecture of the human brain. Instead of separating storage and calculation, neuromorphic chips interweave them.

Memristors—electronic components that can alter their resistance based on the history of current that previously flowed through them—act simultaneously as logic gates and synapses.


"When you recognize a friend's face in a crowd, your brain isn't running a sequential loop through a database of billions of pixels," notes Dr. Amara Lin, a researcher specializing in artificial neural hardware. "It's a massive, parallel, analog cascade of electrical impulses that fires with astonishingly low power consumption—about twenty watts, roughly the energy of a dim lightbulb. If you tried to run modern computer vision models on traditional hardware at that wattage, your laptop would melt."

By weaving memristive crossbar arrays into microchips, engineers are creating hardware that learns natively.



These circuits don't just execute pre-written code; they adapt their internal connection weights based on incoming data streams. For software developers, this opens up a staggering horizon: compilers and operating systems that no longer have to simulate neural networks on rigid, unnatural architectures, but can instead deploy them directly onto native learning hardware.

"The true holy grail of microelectronics isn't raw clock speed. It is the moment our machines learn to process information with the quiet, effortless elegance of a living synapse."



THE LIGHT-SPEED PARADIGM: PHOTONIC INTEGRATION

While electrons crawl through copper traces, slowed down by resistance and capacitance, a parallel revolution is unfolding in the realm of light. Photonic integrated circuits (PICs) replace electrical currents with photons. Instead of pushing electrons through a wire, data is encoded into wavelengths of light and guided through microscopic optical waveguides etched directly onto silicon or lithium niobate substrates.

Light does not experience electrical resistance. Photons do not care if they cross paths; multiple wavelengths can travel down the exact same optical channel simultaneously through wavelength-division multiplexing, completely bypassing the physical bottlenecks of traditional circuit tracks.


For the data centers powering the artificial intelligence boom—facilities that currently consume as much electricity as small nations—photonics offers an escape hatch from an ecological dead end. Optical interconnects allow processors to communicate across racks at the speed of light with virtually zero thermal dissipation.

"We are moving from an age of electronics to an age of optoelectronics," says Dr. Vance. "The computer of the 2030s won't be a faster version of today's desktop. It will be an integrated optical engine where data flashes through glass pathways, operating at bandwidths that make today's fiber optics look like telegraph wires."

THE INVISIBLE INFRASTRUCTURE

As these emerging innovations converge, the very definition of what constitutes an "electronic device" is dissolving. We are moving toward a world of ambient intelligence, where computation is woven invisibly into the fabric of our everyday environment. Consider biodegradable electronics—transient circuits built from silk proteins, cellulose, and zinc that dissolve harmlessly into the soil or the human body after a predetermined operational lifespan. Imagine environmental sensors scattered across agricultural fields to monitor soil chemistry and moisture levels for a single growing season, after which they simply melt away without leaving a single trace of electronic waste.

Or consider energy-harvesting micro-systems that pull their operational power not from lithium-ion batteries or wall sockets, but from ambient radio frequencies, indoor light, thermal gradients, and even the kinetic motion of a user walking across a room.



For the student sitting in a university lecture hall today, these technologies represent more than just incremental career opportunities. They represent a fundamental philosophical shift in how we build our technological footprint. We are transitioning from an era of extraction and obsolescence to an era of integration and harmony.

"We are building a world where technology is no longer an invasive foreign object we carry in our pockets, but an invisible, breathing layer of the environment itself."



THE SENTIENT GRID

How Invisible Electrons Are Waking Up Our Cities

AKHIL SOMAN
S7-ECE

It happens in a fraction of a second, entirely without human intervention, at a rain-slicked intersection in downtown Singapore.

A pedestrian, hidden by the glare of streetlights reflecting off the wet asphalt, steps off the curb against the light. Simultaneously, an electric delivery van approaches the intersection at fifty kilometers per hour. A decade ago, this equation of physics, weather, and human error would have resulted in a tragedy. Today, the street itself intervenes.

Embedded in the concrete, microscopic piezoelectric sensors register the sudden, sharp pressure of a human footfall. Above, a multi-spectral CMOS camera mounted on a lamppost detects the thermal signature of the pedestrian. Instantly, an edge-computing micro-cluster housed inside the traffic light's casing fuses these two data streams.

It doesn't send this information to a distant cloud server in a data center miles away—there is no time for that. Instead, the node processes the anomaly locally in under ten milliseconds. It broadcasts a dedicated short-range communication (DSRC) signal to the approaching van. The vehicle's onboard logic cuts the motor and engages the anti-lock brakes before the human driver even realizes someone is in the road.

The van stops with a meter to spare. The pedestrian keeps walking, adjusting their umbrella, utterly unaware that the city just saved their life.

For the past century, we have built our cities out of deaf, dumb materials. Asphalt, steel, glass, and concrete were passive structures, designed merely to withstand the elements and channel human movement. But a quiet revolution in solid-state physics, materials science, and network architecture is fundamentally rewriting the urban DNA.

We are no longer just building cities; we are installing a planetary-scale nervous system. From the structural supports of bridges that can "feel" their own micro-fractures to lampposts that monitor air quality and direct autonomous traffic, smart electronics are turning the metropolis into a sentient organism. For the students and researchers who will inherit and engineer this future, understanding the silicon substrate of the smart city is no longer optional—it is the prerequisite for modern citizenship

THE SILICON SUBSTRATE: SENSORS AND THE ACT OF "FEELING"

To understand how a city wakes up, you must first look at its sense organs. The backbone of urban electronics is the MEMS—Micro-Electro-Mechanical Systems. These are not merely circuits; they are microscopic machines etched into silicon, featuring moving parts smaller than a human hair. Your smartphone uses a MEMS gyroscope to know when you've rotated the screen. In a smart city, these same microscopic structures are cast into the concrete pylons of highway overpasses to monitor structural health. When a heavy truck rolls over a bridge, the concrete flexes. Inside the MEMS accelerometer, a tiny silicon cantilever bends in response to the vibration. The resonant frequency of this cantilever can be mathematically modeled by a simple but profound relationship:

$$f_0 = (1 / 2\pi) \sqrt{(k / m)}$$


Where k is the stiffness of the cantilever and m is its mass. By continuously measuring shifts in capacitance as this microscopic diving board vibrates, the sensor translates physical movement into an electrical signal. If the bridge develops a micro-crack, the stiffness of the macro-structure changes, altering the vibration profile. The MEMS sensor detects this shift and alerts maintenance crews weeks before a human inspector could spot a hairline fracture in the concrete

"We spent the last century building cities out of deaf, " dumb materials. The next century is about teaching the concrete to listen, and more importantly, to speak."



"People imagine smart cities as a series of glowing screens and Wi-Fi hubs," says Dr. Aris Thorne, a materials engineer and leading researcher in embedded urban electronics. "But the real magic happens at the micron level. We are embedding billions of synthetic mechanoreceptors into the urban skin. We are giving the city a sense of touch."

This sensory network extends far beyond mechanical stress. Electrochemical sensors the size of postage stamps line the municipal water pipes, measuring pH and detecting heavy metals in realtime. CMOS image sensors, paired with LiDAR (Light Detection and Ranging), act as the city's optic nerves. Rather than just capturing 2D video, modern LiDAR emits millions of laser pulses per second, timing their reflections to build real-time 3D point clouds of the urban environment.



The city doesn't just see a car; it perceives a volumetric entity moving through three-dimensional space, accurately calculating its trajectory, velocity, and intent.

LIVING ON THE EDGE: THE BRAIN MOVES CLOSER TO THE BONE

Having billions of sensors is useless if the brain cannot process the input fast enough. For years, the paradigm of the "Internet of Things" (IoT) relied on cloud computing. A sensor would collect data, transmit it via cellular networks to a massive data center in another state, wait for a server to process the information, and then receive a command back.

In a smart city, this model is dangerously obsolete. The delay inherent in this round-trip communication is called latency. In the context of streaming a movie, a latency of 100 milliseconds is a minor buffering annoyance. In the context of two autonomous vehicles navigating a four-way intersection, 100 milliseconds is the difference between passing safely and a fatal collision. The solution is Edge Computing.

Instead of routing data to a centralized brain, engineers are pushing computational power to the extreme "edge" of the network—right into the traffic lights, the bus stops, and the pavement itself.


Edge computing relies heavily on a new generation of Application-Specific Integrated Circuits (ASICs) and Neuromorphic chips. Unlike traditional CPUs that process data sequentially, neuromorphic chips mimic the architecture of the human brain, using artificial neurons and synapses to process parallel streams of sensory data with staggering efficiency.

"A smart city without edge computing is like a dinosaur," explains Dr. Maya Lin, an urban systems architect. "The dinosaur was so large that by the time a predator bit its tail, the nerve signal took too long to reach the brain and trigger a reaction. A smart city must have reflexes. Edge computing acts as the spinal cord, processing reflex actions locally so the city can react in real-time."

This localization of intelligence drastically reduces bandwidth strain. A traffic camera doesn't need to stream 4K video to a central server 24/7. Instead, the edge AI analyzes the video locally, identifies that a traffic jam is forming, and sends only a tiny packet of text data—"Congestion at Main and 5th, rerouting required"—to the central grid.

HARVESTING THE FOOTSTEPS: THE ENERGY EQUATION

A sentient grid presents a massive logistical nightmare: How do you power a trillion microscopic sensors? You cannot run copper wiring to every embedded MEMS chip in a bridge, nor can you rely on batteries that would need replacing every two years.



The electronics of the smart city must be selfsustaining. Enter the science of energy harvesting. Researchers are turning to the piezoelectric effect—a phenomenon where certain materials (like quartz crystals or engineered ceramics like lead zirconate titanate) generate an electrical charge when subjected to mechanical stress. The electrical displacement D_i generated by a piezoelectric material is governed by its relationship to mechanical stress T_j and the electric field E_k :

$$D_i = d_{ij}T_j + \epsilon_{Tik}E_k$$

In practical terms? Every time a pedestrian steps on a piezoelectric tile embedded in a train station platform, the deformation of the material squeezes electrons through a circuit.

A single footstep might only generate a few microjoules of energy, but multiply that by the millions of commuters passing through a transit hub like Tokyo's Shinjuku Station, and you have a substantial power source.

“ We are moving away from the idea of the city as an energy sink, and toward the city as an energy generator. The human heartbeat of the metropolis will literally power its own nervous system.”

Beyond footsteps, engineers are developing Triboelectric Nanogenerators (TENGs) that harvest energy from the friction of car tires rolling over asphalt, and thermoelectric generators that convert the temperature difference between the hot summer pavement and the cool earth below into electricity. These micro-generators ensure that the edge nodes and sensors can operate indefinitely, entirely off the grid, sustained by the very kinetic energy of the city they are monitoring.



THE GHOST IN THE MACHINE: THE PANOPTICON PROBLEM

As we engineer this technological utopia, we must confront the shadow it casts. The same electronics that can detect a pedestrian to save their life can also track their movements, their associations, and their habits with terrifying precision.

When a city is blanketed in multi-spectral cameras, MAC-address sniffers, and acoustic sensors, the line between "smart infrastructure" and a "surveillance state" becomes dangerously thin. If a lamppost can detect the acoustic signature of a gunshot to dispatch an ambulance, can it also be tuned to detect the acoustic signature of a protest forming?

This is the great challenge for the next generation of engineers, sociologists, and policymakers. Technology is rarely neutral. The architecture of smart electronics inherently centralizes information, and historically, centralized information translates to centralized power.

We are currently building what social theorists call a digital panopticon. Unlike the crude surveillance of the 20th century, which required human operators watching blurry screens, edge AI allows for automated, frictionless tracking.

If data governance is left entirely to private corporations or unchecked state authorities, the smart city could become a velvet prison—highly efficient, incredibly safe, and devoid of anonymity. The solution lies in the silicon itself. Researchers are pioneering "privacy by design" at the hardware level. For instance, thermal sensors can track human movement to optimize HVAC systems in public buildings without ever capturing facial features. Edge AI can be hardcoded to immediately delete visual data once a pedestrian's trajectory is calculated, transmitting only anonymous vector coordinates. The challenge is not merely technical; it is intensely philosophical. We must decide what kind of ghosts we want to be in the machine.

THE CLASSROOM OF TOMORROW

We are standing at the precipice of a massive paradigm shift in human habitation. For the college students reading this—the structural engineers, the computer scientists, the urban planners, and the ethicists—you are the architects of this new reality. You are inheriting a world where the boundary between the physical and the digital is dissolving.

The cities of the past were monuments to human muscle and mechanical leverage. The cities of the future will be monuments to human cognition and silicon logic. They will breathe with us, adapt to our shifting needs, and quietly manage the unfathomable complexities of urban density.



But a smart city is only as intelligent as the parameters we set for it. As we embed our logic into the concrete, we must ensure we are also embedding our humanity. The ultimate test of the sentient grid will not be whether it can flawlessly route traffic or harvest energy from our footsteps. The ultimate test will be whether a city that knows everything about us can still leave us free to wander, unnoticed and unbothered, through its brightly lit streets.

Small steps
on the moon,
big dreams
for tomorrow.

Just like Chandrayaan-3,
let's keep exploring,
learning and reaching
beyond the limits.



MANGALAM COLLEGE OF ENGINEERING
DEPARTMENT OF ELECTRONICS AND
COMMUNICATION ENGINEERING

