SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN

[RE]IMAGINE EDUCATION

A weekly Design·AI newsletter

CYBER TWINS FOR EDUCATION

Nagarajan Sivanadipatham, iTrust
In Conversation with Dr. Nachamma Sockalingam, OSP

Nagarajan Sivanadipatham, Research Associate

https://www.sutd.edu.sg/profile/nagarajan-sivanadipatham/

What if students could experiment with a water plant, power system or other critical infrastructure — test what happens when something fails, simulate an attack, redesign a response, and then reset the entire system to try again?

At iTrust, SUTD's centre for research in cyber security, this is increasingly possible through **cyber twins**: high-fidelity virtual replicas of real operational technology environments. Originally developed for cyber-security research, infrastructure operators and defence training, these virtual environments also open up interesting possibilities for teaching and learning across SUTD.

Because the environment is virtual, students can explore scenarios that would be difficult, costly or unsafe to reproduce using live infrastructure. They can observe how systems behave, test hypotheses, investigate failures and design possible solutions — turning complex infrastructure into an experiential learning environment.

In conversation with **Naga from iTrust**, we explore how cyber twins work, how AI is accelerating their development, and how they could become a shared platform for students and faculty to **learn by testing, simulating and designing**. How might we explore cyber twins in our areas of work?

Editor of Newsletter:
Dr. Nachamma Sockalingam, Office of Strategic Planning

Contact us at nachamma@sutd.edu.sg to share your Design·AI stories

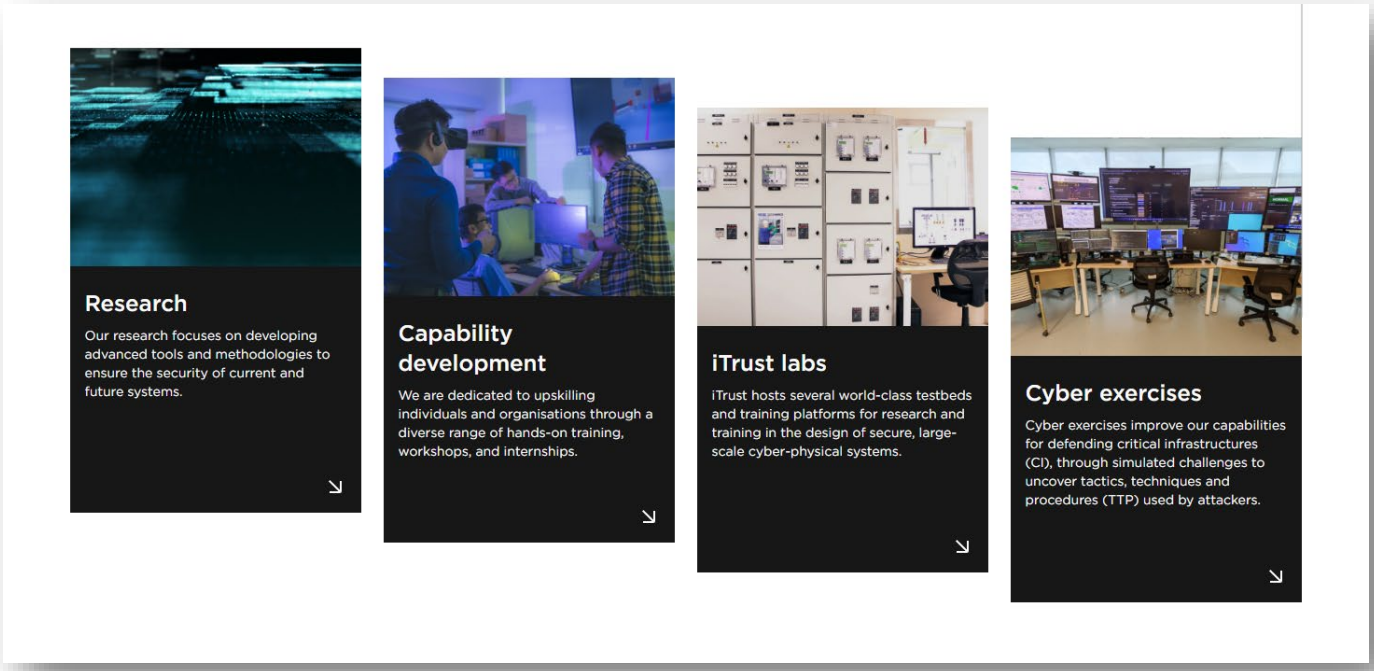

About iTrust

Naga, could you tell us a little about your role at iTrust and your work on cyber twins?

At iTrust, I lead work on digital twins, particularly what we call **cyber twins** — high-fidelity virtual replicas of operational technology environments. These can recreate everything from controllers and SCADA systems to industrial protocols and the underlying physical processes that keep systems such as water plants and power grids running.

This work grew quite naturally from iTrust's experience with physical testbeds such as SWaT, WADI and EPIC. These testbeds let us study cyber-physical systems and attacks in a safe environment, without putting real infrastructure at risk.

Cyber twins take that same idea into software. They give us much of the realism of a physical testbed, but in a form that is portable, replicable and easier to access. That means we can simulate failures or attacks, observe how the system responds, reset it and try again — which makes them useful not only for research and training, but potentially for teaching and learning as well.

You mentioned that cyber twins give us the realism of a physical testbed, but in a much more portable and accessible form. What makes them particularly powerful for experimentation and learning?

Anatomy of Cyber Twin

The real advantage is that you can do things with a cyber twin that you would never dare to do to a real facility. **You can launch a full-scale attack, watch a pump fail, flood a process with bad data, and observe how the system responds — and then reset everything and do it again.**

Because **the twin reproduces how the actual plant behaves** — from the sensors and controllers to the underlying process and industrial protocols — you are not just looking at a diagram or a dashboard. You can actually see how one action affects another part of the system and understand how failures propagate.

What makes this even more powerful today is the way it connects to **Design AI and physical AI**. A cyber twin is no longer just a simulation environment; it becomes a **design space for intelligent systems**. With AI agents embedded in the loop, you can start to design, test, and refine defensive strategies, control logic, and even attack scenarios in a much more automated and exploratory way. **Design AI** helps researchers and students rapidly generate new configurations, stress-test assumptions, and explore "what-if" scenarios at a scale that would be impossible manually.

That **ability to test, observe, design with AI, reset, and try again makes cyber twins a very powerful environment for a defender, an operator, or a student. It creates a safe, intelligent space to experiment with situations that would otherwise be too risky or impractical to reproduce on live infrastructure** — while also preparing us for a future where AI systems are not just analysed, but actively designed and trained in the physical world they will eventually operate in.

Building these environments used to take years. How is AI changing the way you develop and use cyber twins today?

Traditional bespoke OT twins take one to two years; AI-assisted builds at iTrust take weeks.

What has really changed in recent years is not just what cyber twins can do, but how quickly we can build them. I have seen the traditional way of building a bespoke OT twin, and it is slow. Industry guides put enterprise-scale, multi-asset digital twin builds at anywhere from twelve to twenty-four months, and much of that time is spent on painstaking, repetitive engineering — wiring up protocols, reconstructing control logic, building interfaces and validating process behaviour by hand.

In our work at iTrust, **we have compressed that timeline dramatically by putting AI at the centre of the build process. Rather than hand-coding every component, we run an in-house, AI-assisted build framework: a coordinated set of specialised AI agents, each tuned to a different part of the job** — one to plan the architecture, one to write and assemble the engineering, one to handle routine work — while human experts scope the problem, make the design decisions and validate every result.

The **AI does the heavy lifting** on the parts that used to eat our calendars — protocol plumbing, boilerplate control logic, HMI dashboards, the first pass of the process simulation. That frees the engineers to do what only they can do: decide what “realistic” means for this particular plant, and prove that the twin actually behaves like the real thing.

Crucially, **speed is not allowed to come at the cost of security.** Client environments carry sensitive operational detail, so we enforce a strict separation — a “sanitisation wall” — between the cloud-connected machines that do the building and the air-gapped data that describes a client’s real infrastructure. Nothing sensitive crosses that boundary.

How are cyber twins being used today in real-world research, training and cyber exercises?

Cyber twins are no longer just research prototypes — they are now becoming operational tools used in very practical and high-stakes settings.

In research, they allow us to study cyber-physical attacks in a controlled but realistic environment. Instead of theorising about how an attack might propagate through an industrial system, researchers can actually simulate it, observe the cascading effects, and test different detection or mitigation strategies. This is especially valuable in critical infrastructure domains where real-world experimentation is simply not possible.

In training, cyber twins are being used to prepare the next generation of engineers and cyber defenders. Trainees can interact with a live, behaving system that mirrors real industrial plants, learning how to recognise abnormal behaviour, respond to incidents, and understand the operational consequences of cyber events. Because the environment is safe and resettable, they can make mistakes, learn from them, and try again — something that is rarely possible in real operational settings.

We are also **seeing increasing use in cyber exercises and "live-fire" simulations, where teams from industry, government, and defence come together to test their response to coordinated attacks. In these scenarios, the cyber twin acts as a shared, realistic battlefield where defenders can practise decision-making under pressure without risking actual infrastructure.**

Across all these use cases, the **key value is realism without risk.** Cyber twins provide a space where we can stress-test systems, people, and processes in ways that closely mirror the real world, but with full control over safety, repeatability, and scale.

AI is helping your team build cyber twins much faster, but in a safety-critical environment, speed cannot come at the expense of accuracy. What do you see as the key benefits and risks of using AI, and where does human judgement come in?

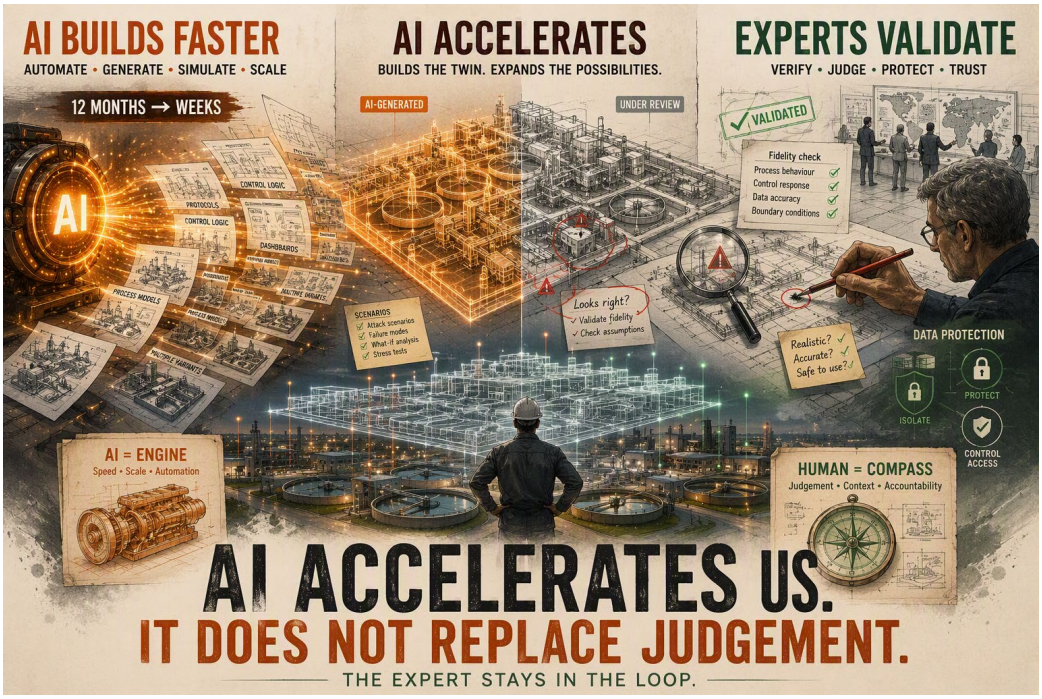

It is not AI versus Human; It is Human + AI

The benefits are significant. **AI allows us to automate many of the repetitive parts of building a cyber twin** — things like protocol plumbing, control logic, dashboards and the first pass of process simulation. That means builds that once took a year or more can potentially be completed in weeks. We can also create more variants, test many more attack and failure scenarios, and generate realistic synthetic data for training and research.

But there are risks too. **An AI-generated model can look very convincing and still be subtly wrong, and in an operational technology environment, even a small error can matter**. So fidelity cannot simply be assumed — it has to be validated. There are also important considerations around sensitive infrastructure data and making sure that information is properly isolated and protected.

This is where human judgement becomes essential. AI can accelerate the engineering work, but it cannot decide by itself what “realistic” means for a particular plant or whether a twin is behaving correctly. That requires people who understand the actual process and can validate the results.

For us, the principle is very clear: **AI accelerates us; it does not replace judgement. The expert stays in the loop.**

Looking ahead, how could students and faculty across SUTD use cyber twins for teaching, learning and research?

I think cyber twins can really become a **canvas for education and innovation**. Instead of learning cyber security only through slides or static labs**, students can actually interact with systems that behave like real industrial environments.** They can see how a small misconfiguration or a malicious action propagates through a system, and more importantly, they can experiment safely and repeatedly until they truly understand what is happening.

For teaching, **faculty can design very rich, hands-on modules where students are not just passive learners but active participants in a living system.** You can teach concepts like network security, control systems, anomaly detection, or even system resilience by letting students "break and fix" the system in a controlled way. That kind of experiential learning is very hard to achieve in traditional classroom settings.

For research, **cyber twins open up a completely new space. Instead of relying only on historical data or limited testbeds, researchers can generate realistic scenarios, simulate rare or dangerous attack conditions, and evaluate defence mechanisms at scale.** It becomes possible to test ideas much earlier in the research cycle, and iterate quickly without waiting for physical infrastructure or real-world incidents.

Ultimately, I see **cyber twins as something that can bring students, researchers, and even industry partners onto a common platform. It is not just a tool for simulation — it is a shared environment where we can learn, experiment, and design the future of secure cyber-physical systems together.**

Cyber twins may come from the world of cyber security and critical infrastructure, but the educational idea behind them is much broader.

They create spaces where learners can test, observe, fail safely, reset and try again. Instead of only learning about a system, students can learn through interacting with it and seeing the consequences of their decisions.

Just as importantly, Naga's work reminds us that AI is most powerful when it accelerates human capability without replacing human judgement.

For educators, the opportunity may not be to replicate cyber twins literally, but to borrow the principle behind them: create safe, realistic environments for experimentation, iteration and reflection. The question then becomes: what could we make possible for our learners if they had more opportunities to simulate, test, redesign and learn by doing?