

Two years ago, ARM decided to investigate the way in which its teams undertook hardware design. Using tools developed for cloud computing, the company now sifts through every result from the thousands of tests and simulations it runs on a daily basis in order to look for patterns.

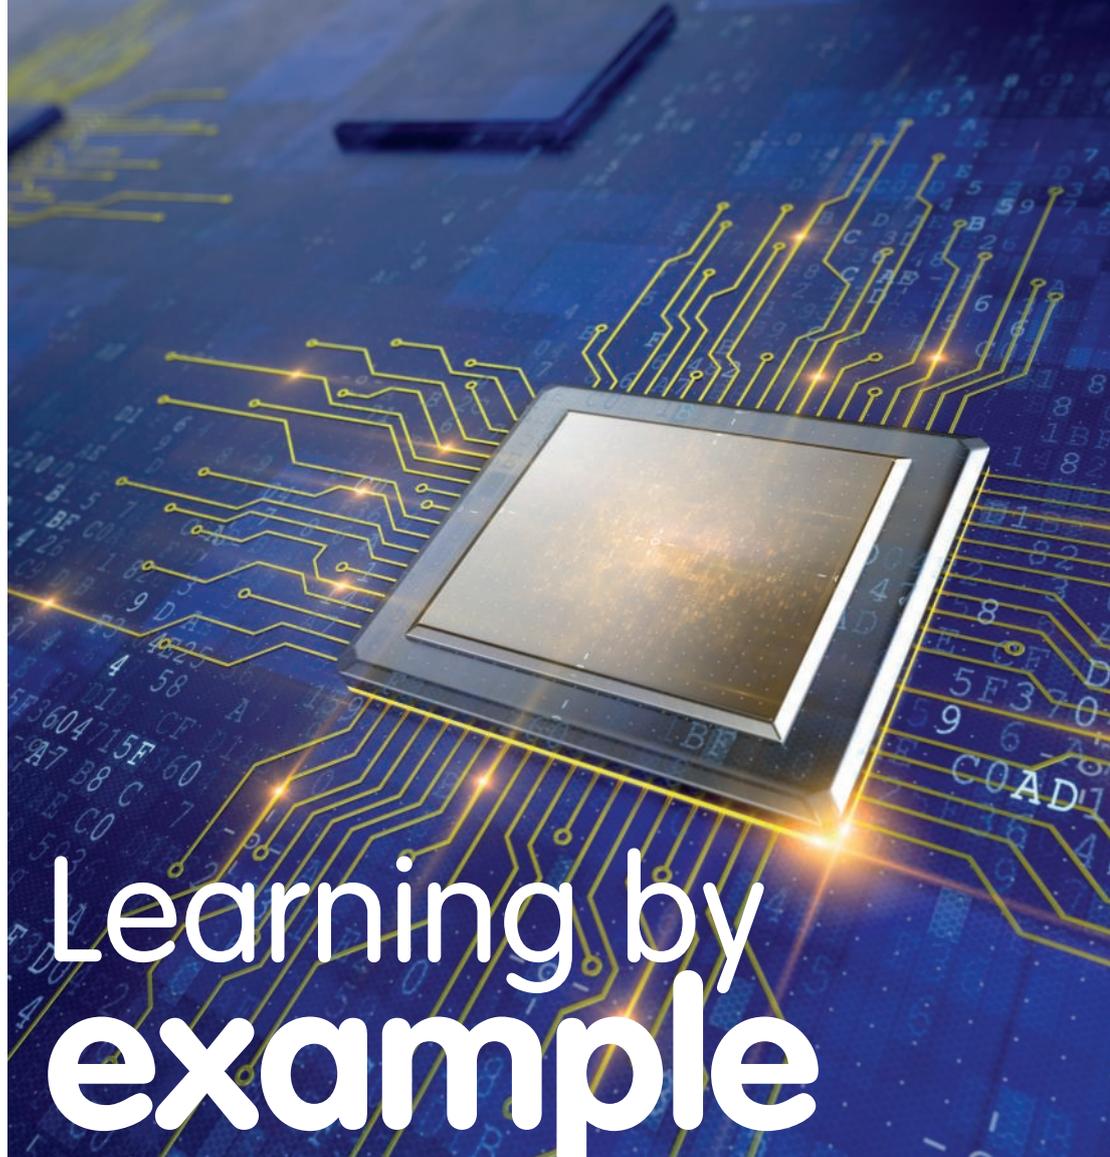
“RTL verification generates a lot of data,” Bryan Dickman, director of engineering analytics in ARM’s technology services group, told the TV&S Verification Futures seminar earlier in 2017. “It’s becoming a big-data problem.” One of the questions which ARM wants to answer is ‘are we running dumb cycles?’.

Dickman noted: “We do a lot of analytics around bugs. We drill into bugs by methodology and areas of the design: what methodologies find bugs and when. It is interesting to see how you are finding bugs against the amount of verification you are running. You can start asking questions like ‘I have been running lots of cycles, but no bugs were discovered: were they dumb cycles?’.”

So far, the emphasis has been on building visualisation tools to help engineers analyse the patterns. The next step is to build tools that can learn from the data. “We are looking at how we can take machine-learning algorithms and apply them to the data, then design predictive workflows that improve productivity,” Dickman said.

Machine learning (ML) has started to take hold across the EDA world in general and more than 20% of papers at the 2017 Design Automation Conference (DAC) dealt with ML techniques. But the applications are not going to be distributed evenly. Access to data is one of the biggest issues.

Speaking at a panel session during DAC, Ting Ku, senior director of engineering at nVidia, said ML techniques need to be tuned to the problem. “When people talk about



Machine learning techniques could play a leading role in getting bug free chips to market more quickly. By **Chris Edwards**.

ML, most think it’s neural network-related. But the data doesn’t have to be modelled by a neural network.”

Singapore-based Plunify combines domain knowledge with some forms of ML in its InTime tool, which tunes the parameters used to compile FPGA designs for performance and area.

Plunify co-founder Kirvy Teo noted: “The nature of ML in this area is different from the things we see at Google. Neural networks work very well for lots of data points, but we are dealing with thousands of data points, rather than billions.”

Fellow co-founder and CEO HarnHua Ng added: “We can’t tell customers ‘First you’ve got to run ten thousand compilations’.”

EDA tool developers can apply techniques ranging from fairly simple regression algorithms, that fit a curve

to a collection of data points, to more advanced techniques, such as support vector machines that, before the advent of deep learning, proved successful in image recognition.

Earlier this year, tool supplier Solido Design Automation launched its ML Labs initiative. The company built a toolkit of techniques for use in its own tools, such as one that analyses models of the low-level behaviour of the cell libraries used in IC design on advanced processes. In these processes, statistical variation at the nanometre scale strongly affects performance and the effects of temperature and voltage are not always linear. Traditionally, this demands many simulation runs to characterise the cells at all possible operating points. The tools analyse how the variables in existing

“We are looking at how we can take machine learning algorithms and apply them to the data, then design predictive workflows that improve productivity.”
Bryan Dickman

libraries interact and build models to predict the parameters for a subset of simulations that will create an effective library.

Amit Gupta, Solido CEO, explained: “We have been focused on ML for variation-aware design for the past seven years and are expanding that into other areas. The first area is characterisation, but we are getting demand from customers to expand into other areas within EDA.”

Sorin Dobre, senior director of technology at Qualcomm, points to an expansion of ML across more of the tasks needed at the physical-design level. “At 10nm and 7nm, we see a lot of process variation. You have to validate across multiple process corners.”

The problem that faces physical-verification teams is the number of combinations of temperature, transistor speeds and voltage that require checking under brute-force methods. “The question is to how to get good quality results without an explosion in process corners. ML can bring a 10x improvement in productivity.”

One area where RTL verification teams have seen success with ML methods is in eliminating some of the dumb verification cycles that concern ARM’s designers. Harry Foster, chief scientist at Mentor, now a subsidiary of Siemens, says. “Can I minimise the number of tests and still achieve

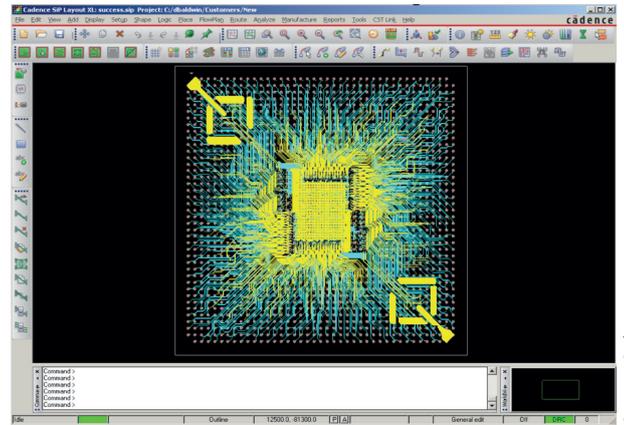
a reasonable level of coverage? There are papers on that and it is a technique now being applied in the industry.”

Yield analysis is another area where ML is having an impact. It can predict the types of geometry in a dense design that are likely to lead to lithographic or other process failures – and thus dud chips. “Improving yield by two or three points has a significant economic benefit,” Dobre notes.

Machine learning should make it easier to build simplified models of complex interactions that defy rule-based design methods. For example, researchers at Georgia Institute of Technology proposed using models generated automatically from electromagnetic measurements and simulations to help organise devices and I/O pins on multi chip modules so that signals do not suffer from crosstalk or IR drop.

“Where ML doesn’t work well is where you have to guarantee finding something, like manufacturing test. With verification, I can get away with less than 100% coverage, because my intent is mitigating risk. I can’t with manufacturing test, because it leads to returns,” Foster said.

The next step is to try to reuse the learning across projects. This is one of the obstacles that slows down adoption of data mining and ML in RTL verification.



Courtesy Cadence

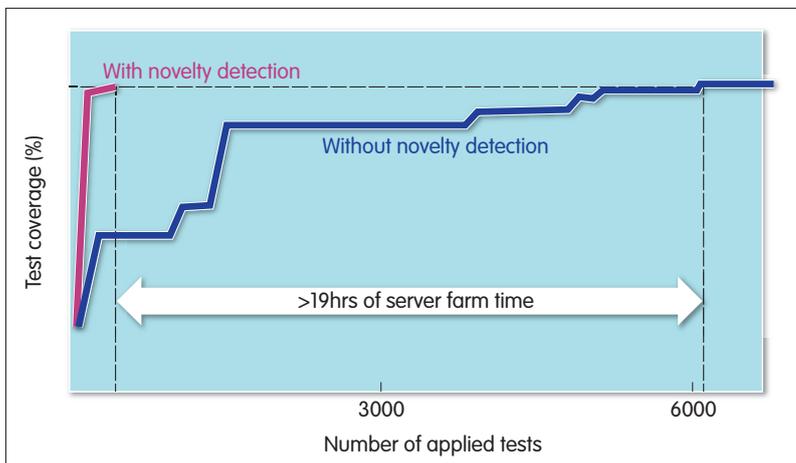
Researchers at Georgia Tech, amongst others, are looking at whether machine learning can be used to build better heat-aware models that will improve chip-package codesign

Foster continued: “The issue with [RTL] verification is that although people say we have lots of data, the reality is we don’t. If I’m starting a project, I have no data. How can we be successful in applying ML techniques if our data is limited? It can’t be based just on pure data mining; it has to incorporate some domain knowledge.”

Dobre says the issue of lack of data is reflected in physical implementation. “The question is how can you maintain cycle times when you go to the next process technology? We have highly automated flows from RTL to gates, but they are not repetitive – they are based on experience. If you bring in a new design team, you start from a pretty low level of experience.”

One possibility is to use data mining to look at repetitive tasks, such as timing closure, and to use ML to automate some of them, while learning about the new processes as each project continues. “We need to capture designer knowledge during implementation and reflect it back to the design community and the foundry side to fix issues,” Dobre argued, adding this would need data from many sources.

“There is a need for flexibility from the major EDA companies to work with these solutions,” he concluded. “If you have a closed system and don’t allow anyone to interact with your tools, we can’t do anything.”



A machine-learning algorithm developed by the University of California at Santa Barbara and Freescale filters achieved the same coverage in 310 tests as a normal testbench did in 6010.