Veriware – model checking in a world of ubiquitous computing

We are moving into a new era of what's been called 'ubiquitous computing'. Processing power is no longer confined to what we've grown to think of as computers; it is now appearing in everything from phones and cars to household appliances.

Veriware, a new five-year project at Oxford University, aims to lay the foundations for the coming world of what futurist Adam Greenfield has termed 'everyware' – foundations that will let us be sure that the devices around us, on which our financial or physical safety may depend, work as we expect them to at all times.

Veriware recently received an award of just over €2m from the European Research Council. It aims to extend the discipline of 'model checking', which is used to ensure software programmes behave as they should, into the uncharted waters of the software embedded in everyware devices. It is led by Marta Kwiatkowska, a Professor of Computing Systems at Oxford University Computing Laboratory who specialises in the field.

The technology of everyware is still embryonic, but already many examples are on the market. Examples include Bluetooth phones, which automatically sense each others' presence and exchange information, or the electronics in cars that help the user control everything from engine function to entertainment systems. There's even a fridge that senses when the food stored in it is running low, and automatically orders more over the internet.

But many think the concept will be pushed much further. It won't be too many years before we spend our lives surrounded by tiny, barely-visible computers embedded in the objects around us, monitoring the world around them with sensors and sharing information with each other over wireless connections. Everyware will offer undreamed-of possibilities, making our lives far safer and more convenient – if it works properly. But there are major risks too; when the everyware we rely on breaks down, the consequences could be alarming.

Already, manufacturers have had to recall products at great expense after discovering flaws in their embedded software. If your internet fridge goes haywire and starts ordering large amounts of the wrong kinds of food, the results will be expensive. And if software that controls a car's braking system fails while the user is driving on the motorway, for example, the error could endanger more than just your bank balance. Products are tested, but they can never be exposed to every situation that could arise, so the risk remains that an unknown problem is lurking somewhere.

Kwiatkowska believes model checking could provide the answer. At present, model-checking software is provided with the source code of a programme written in languages like C or Java, and instructed to inspect it with a particular state in mind. It then systematically analyses every possible path through the programme, and shows whether or not that state will ever occur.

Unlike simply testing how a device works in practice, the results of model verification amount to a mathematical proof. The technique has been used for purposes ranging from automatically detecting errors in Windows device drivers to proving that certain problems can't arise in the software controlling Airbus passenger jets.

Kwiatkowska specialises in a more complex area of model checking, in which the software isn't expected to produce a simple yes/no answer about whether a programme can reach a specified desirable or undesirable state.

Instead, this 'quantitative model verification' uses probabilistic reasoning to determine the likelihood of a given situation occurring, or how long it can be expected to take. This lets investigators take account of the ways in which a system's smooth functioning can be affected by external factors outside its control. It's been used for tasks like finding the worst-case scenario for how long a specified amount of data could take to be transferred between two wireless devices.

The challenge Kwiatkowska is taking on with Veriware is to make the theoretical breakthroughs that will allow similar methods to be applied to the software behind everyware. It's a hugely ambitious goal. The software that model checking is presently used on has a well-defined beginning and end. It is given an input, and it does information processing and produces an output. This makes it relatively easy to check.

The programmes underlying everyware will be in a far more complex situation. Instead of finite, well-defined processing tasks they will be engaged in a continuous, ever-changing series of negotiations with devices around them, as they sense the world and wirelessly exchange the information they learn. They'll also have to operate in an uncertain environment in which resources like bandwidth may become scarce.

It won't be enough to model-check this software before release, although that would be a good start. What's also needed is 'online' model-checking technology which can continuously verify that everyware is processing its sensory input correctly and making the right decisions, and can take appropriate steps if a problem arises.

People will expect everyware to autonomously adjust to deal with any situation it faces – a key selling-point is the fact that the computers around us won't need frequent human intervention the way present-day computers tend to. Significant theoretical breakthroughs are needed to allow the kind of continuous, adaptable model-checking this will require.

The project's potential benefits outweigh its difficulty, though. Everyware will transform many aspects of our lives, but we have to be able to trust it to take care of important things on its own. If it succeeds, Veriware will provide a sound basis for that trust.