NICTA

EMBEDDED SYSTEMS VISION

Author	Chris Nicol
Subject	Embedded Systems Vision
Document Version #	1.0
Date Finalised	August 2010

NICTA Embedded Systems Research Theme

Vision

"Every person, every day, interacts with an embedded system developed by NICTA"

We are leveraging our strengths in operating systems, microelectronics, software engineering, communications and image processing to create mutually-reinforcing embedded systems research activities. Individually, we strive for global research excellence and together we drive competitive advantage for our customers in areas of great importance to the embedded systems industry.

Embedded systems

Embedded Systems are control systems embedded in products (from the very large, such as aeroplanes and traffic systems to the small, such as cell phones and cameras) and devices (including implantable devices). They are implemented using software and hardware (including digital, analog, mechanical, chemical and biological). They are effectively pervasive and in many cases provide major value creation opportunities in product roadmaps. Since 1999, the annual growth rate of embedded devices markets is over 10% and all indicators predict an increased growth rate in the deployment of embedded systems, making it a major growth area of ICT. Today, 16 billion embedded processors are controlling systems with only a minor proportion in personal computers, representing a significant opportunity for the Australian ICT sector to deliver embedded systems components, software, design methodologies and technology platforms into global supply chains over the next 10 to 20 years. Embedded systems are growing in complexity. No longer firmware on a microcontroller; today's embedded systems are complex systems of multi-processors, communications networks and distributed control. With this complexity comes challenges and these provide opportunities for NICTA to deliver research outcomes with lasting, valuable impact. NICTA will tackle aspects of the following 4 major research challenges where we have the ability to create critical mass, and partner to effect real impact. We will create research outcomes that grow the international reputation of our research teams and in doing so, will underpin NICTA's endeavor of becoming a top 10 ICT research institute.

1. The systems and software engineering challenge

Embedded systems are often mission-critical. If the embedded control system fails, the entire product is likely to fail. This creates challenges in design, development and verification, especially where safety of people and property is imperative. Furthermore, as the architecture of embedded systems grows in complexity, the challenge of ensuring reliable operation dominates development schedules and cost. Large, complex, feature-rich software systems are notoriously difficult and expensive to construct reliably, safely and securely. Software is 70% of the development effort in a 3rd Generation cellular mobile phone, and 70% of the development cost of an automobile. Development cost is dominated by verification and validation. The challenge is to find new ways to rapidly develop correct embedded systems with efficient architectures and implementations.

There exist many architectural solutions to a given problem. An overarching principle driving embedded systems architectural research lies in the way designers arrive at a particular solution. “Architectural fit” is the alignment of an architectural solution to a problem. The consequences of a poor architectural fit can include increased design time, verification, production cost and power consumption.

Model driven development involves selecting an architecture that best fits the domain (platform independent) model derived from the problem specification. These abstract models are transformed into platform specific models through a series of transformations. It is a principle of this approach to systems design that the verification process is an integral part of each step of the engineering process that results in the validation of a system with regard to its specification without additional work. The first challenge is deriving the domain model from the specification. The next challenge is a model transformation into a platform specific model that defines the architecture, functional processing, synchronisation, bandwidth, and memory subsystems for implementation. At this stage, the constraints of the platform specification are carried forward into the design. Success with this approach comes from an excellent requirement specification (one that is use-case driven), followed by a set of testable executable specifications derived from the requirements, quantitative assessment of candidate architectures, and mapping to implementation technologies such as software, electronics and mechanics.

In the NICTA Green Transport Green Car program we are modelling thousands of vehicles (including the real time modelling of engine & transmission plant, control subsystems and communications) as they interact with each other (driver models, radar, V2V) and with infrastructure control (V2I, traffic control & signals). This modelling technology will enable model driven design processes for unified automotive and infrastructure control system design. This requires a multi-disciplinary engineering approach that leverages many of NICTA’s strengths.

At NICTA we are also developing new reconfigurable and scalable development process models (focusing on early and continuous verification) as extensions or replacements of traditional reference development process models. We are also creating simulation and analysis techniques that can help create, evaluate and optimize a customized development process using a range of parameters including user-defined objective functions, system architectures, process-impacting technologies and traditional resource/scheduling/cost variables. The unique feature of this work is its ability to go beyond process and project management parameters and include system architecture and quality requirements as model and analysis parameters.

NICTA Project: Green Transport Green Car: “Empirical engineering platform.”

NICTA Project: Green Transport Green Car: “The process of distributed systems engineering.”

NICTA Project: Scalable Vision Machines. This project uses model transformations to create an efficient heterogeneous multi-processor implementation model from a high level platform independent model. This transformation and optimisation is performed at run-time.

2. The safety, security and trusted computing challenge

The global revolution of the “internet of things” will lead to embedded systems that are continually communicating with other systems. These systems may be subjected to intentional attacks, or plagued by software failure. As the complexity of systems grows, applications that require different levels of verification and priority will share the same computing environment. Ensuring temporal and spatial fault isolation on a secure, trustworthy software environment will be critical to ensure overall system operation. For example in a vehicle, a failure of the air conditioning subsystem should not interfere with the operation of the park brake. In vehicular networking, vehicles communicate with each other and infrastructure to co-operatively improve safety and resource utilisation. Vehicles are required to verify the validity of the messages they receive and ensure that the messages they send are secure. The cost of not getting this right ranges from loss of life, to loss of consumer trust. NICTA is a world leader in the development of embedded operating systems for safety and security applications.

The ERTOS group will develop embedded systems software platforms, and associated technologies, which enable the construction of secure, safe, and reliable complex embedded systems. Our approach uses trustworthy microkernel-based operating systems, tools for construction and analysis of component-based systems constructed upon microkernels, and the strong enforcement of functional and temporal isolation boundaries between components, enabling systems composed of both highly-trustworthy software and un-trusted software.

NICTA has developed a mathematically verified implementation of a microkernel to serve as a trustworthy foundation of embedded systems. This is complemented by techniques to identify the temporal behavior of the microkernel. We will develop operating-system components that provide system services to higher-level software layers, such as secure user-interface components. Some of these components will possess the same rigorous guarantees as the underlying kernel. We are also developing a system-composition framework which enables vendors to architect systems that take advantage of the mathematically proven isolation guarantees, in a provable way, without the vendors requiring formal methods or real-time analysis experts. While isolation in the presence of failure is critical to our approach, we also expect to develop analysis tools for our platform, which will aid in the development of correct systems and minimise the reliance on the strong isolation guarantees.

NICTA Project: “ERTOS” Secure Trustworthy Embedded Software.

3. The concurrency and power challenge

This challenge is likely to create an inflexion point across the industry. It is brought about by (a) processor clock ceiling at 4GHz, (b) power ceiling at 130W, and (c) increased variance at deep sub micron (DSM) <40nm CMOS technology nodes. The clock and power ceilings have forced processor architectures away from superscalar to multi-processors and yet the majority of legacy software is compiled to sequential threads that are difficult to map across these multi-processors efficiently. The performance available to execute a sequential thread has levelled off but customer data sets continue to grow. As a result, execution times will grow that in-turn will erode competitive advantage. This can only be fixed with a re-architecture of the software however, less than 2% of the world’s programmers are adequately trained in multi-threaded programming.

Researchers are considering new languages that enable the description of the concurrency of the problem. At NICTA, we are taking the approach of building domain-specific languages on top of a functional programming paradigm. Unlike imperative languages (e.g. C), a functional language requires the user to explicitly specify the relationship between computations. This enables the automatic generation of a parallel implementation by a compiler that maps concurrent threads onto multiple cores to achieve speed-up. Indeed, NICTA's Scalable Vision Machines project[1] is developing a domain specific language for describing computer vision and real time image processing algorithms; and aims to automatically deploy their implementation to typical heterogeneous processing architectures found in smart IP camera systems. Some approaches being considered by researchers and developers are Intel Concurrent Collections[2], F Sharp[3], CUDA[4], OpenCL[5] and Haskell[6].

As the numbers of cores on a chip increases to 1000, the cores will necessarily become smaller, simpler and more computationally efficient. Compilers and operating systems will evolve from mapping threads to a single core to mapping programs across arrays of cores. This represents a fundamental shift in systems thinking that could have dramatic impact on operating systems of the future (an area that NICTA is considering for future ERTOS activity). It also places so-called many-core systems on a path that will intersect with ASICs and FPGAs. In the ASIC industry, rising NRE costs are already resulting in fewer viable design starts. ASICs require substantial IP reuse which will move their "customisation" from hardware (RTL) to embedded software implementation. More applications will migrate away from ASIC platforms to programmable or reconfigurable solutions. To address this trend, the FPGA research community is investigating arrays of programmable processors as future field programmable architectures.

ASICs, field programmable gate arrays and multi-core processing could all evolve into *massively parallel arrays of simple programmable processing cores (MPP)*. In anticipation of this, NICTA is partnering with Wave Semiconductor to develop projects that map applications to MPP arrays.

Dynamic voltage and frequency scaling (DVFS) is currently combined with temperature and process sensors to match the power consumption to the expected computational workload. DVFS in the synchronous world is a (cute) kludge that will become increasingly difficult to implement reliably in DSM. The overhead in the clock period that is attributable to variance in circuit delay will grow as a proportion of the total clock period (to approximately 40%). We predict that synchronous systems will hit a ceiling and the industry will reconsider clock-less (asynchronous) logic families for increased performance, lower power consumption and EMI. The robust operation of delay and voltage insensitive circuits also enables efficient voltage scaling.

We have identified the following research challenges that are likely to leverage our strengths and deliver valuable technologies into the inflection point described above.

- Delay insensitive circuits with dynamic voltage scaling to sub-threshold operation.
- Mapping programs to heterogeneous multi-processor architectures.
- Compiling programs to data flow graphs and mapping to MPP arrays.
- Real Time Operating Systems architectures for spatial computation fabrics.

---

[1] http://www.nicta.com.au/research/projects/scalable_vision_machines

[2] http://software.intel.com/en-us/articles/intel-concurrent-collections-for-cc/

[3] http://research.microsoft.com/en-us/um/cambridge/projects/fsharp/

[4] http://www.nvidia.com/object/cuda_home_new.html

[5] http://www.khronos.org/opencl/

[6] http://haskell.org/

NICTA Project: Scalable Vision Machines

NICTA Project: Wave-Semi (Clock-less)

4. The dynamic scene understanding challenge

We envisage a future with millions of smart cameras deployed as hierarchies of dynamically reconfigurable visual sensor networks. These networks of cameras will work together to perform dynamic scene analysis with distributed algorithms. They will be able to track objects between scenes. Furthermore, sensors will themselves move about freely in space. Cameras will be placed in cars, and on pedestrians (like smart eyewear with real time augmented reality or visual sensors for bionic eye processing to assist people with vision impairments). The other trend that we acknowledge here is that “smarts” are migrating out to the sensors. Smart sensors will measure the environment and be able to store the measurements, process the data, store the results and transmit some or all of the data to other sensors and back through the network for processing.

It is with these trends in mind that we create a theme plan for embedded dynamic scene analysis. In doing so, we combine the unique mix of NICTA’s computer vision capability, with sensor networking and the making sense of data research theme, to create solutions that leverage these skills for competitive advantage.

The need to extract real-time information from dynamic scenes is a unifying requirement of a broad range of NICTA's projects that require computer vision. Rather than a single camera, or multiple cameras, we need visual sensor networks that can, in real time, give information about self-motion, view-points of other cameras that are linked in the scene, the structure of the scene, and the motion of any objects in the scene. At the same time, the sensor network should identify certain key objects in the scene and track and recognize people (the so-called “face in a crowd” problem). Moreover, these requirements are relevant to a broad range of possible sensing technologies in other areas including Intelligent Transport Systems (for both in-car devices and road infrastructure), video surveillance, and low-vision assistive devices.

To address these applications, we propose to develop embedded technologies for smart visual sensor networks for the analysis of dynamic scenes. These visual sensors make use of a mix of embedded parallel hardware and programmable processors to facilitate real-time analysis, and include NICTA algorithms for scene structure and object motion recovery. They will also identify a class of objects in the scene.

Visual sensors will have “trained” descriptors of a set of objects that are recognized and tracked whenever they appear in the visual scene. Segmented objects, their respective positions and motions, as well as that of the cameras, will be output via wireless network from the camera. If there are multiple cameras in a scene, these will automatically interface to each other to form a visual sensor network that enables the cameras to work together to improve rates of object recognition as well as the quality of positional and motion data.

Networks of cameras might be controlled via a network controller that performs operations equivalent to soft hand-over in a mobile communications system, enabling objects to be tracked across groups of sensor networks and creating a fully scalable sensor network.

We propose to tackle this problem using geometric and biologically-inspired approaches. Research in the Canberra Research Lab focuses on classic computer vision approaches to scene segmentation and understanding using embedded cameras. The Queensland Research team are exploring processing for smart sensor networks with massively parallel low-resolution processing that mimics the human visual processing system. Arrays of reconfigurable processing elements embedded in smart cameras can be used to implement different bio-inspired parallel algorithms and filters. A top-level algorithm reconfigures the array elements dynamically based on application and scene. This type of reconfigurable system enables the migration of hardware and software algorithm “agents” across camera boundaries.

To build such a system will require research into embedding computer vision algorithms in hardware, on real-time and parallel algorithms for detection/recognition/tracking of objects, and on scene structure and motion recovery.

NICTA Project: Visual Processing for the Bionic Eye (VIBE)

NICTA Project Advanced Surveillance: Advanced Surveillance using visual sensor networks for Safety and Security Systems.

NICTA Project: Automap

How it all fits together

The diagram below shows how the various embedded systems projects span the challenges outlined above. For example, the *scalable vision machines* project is relevant to systems engineering and model transformation as well as the automated exploitation of concurrency – specifically for computer vision applications. This project has applications in transport systems, mobile systems, safety and security and defence.

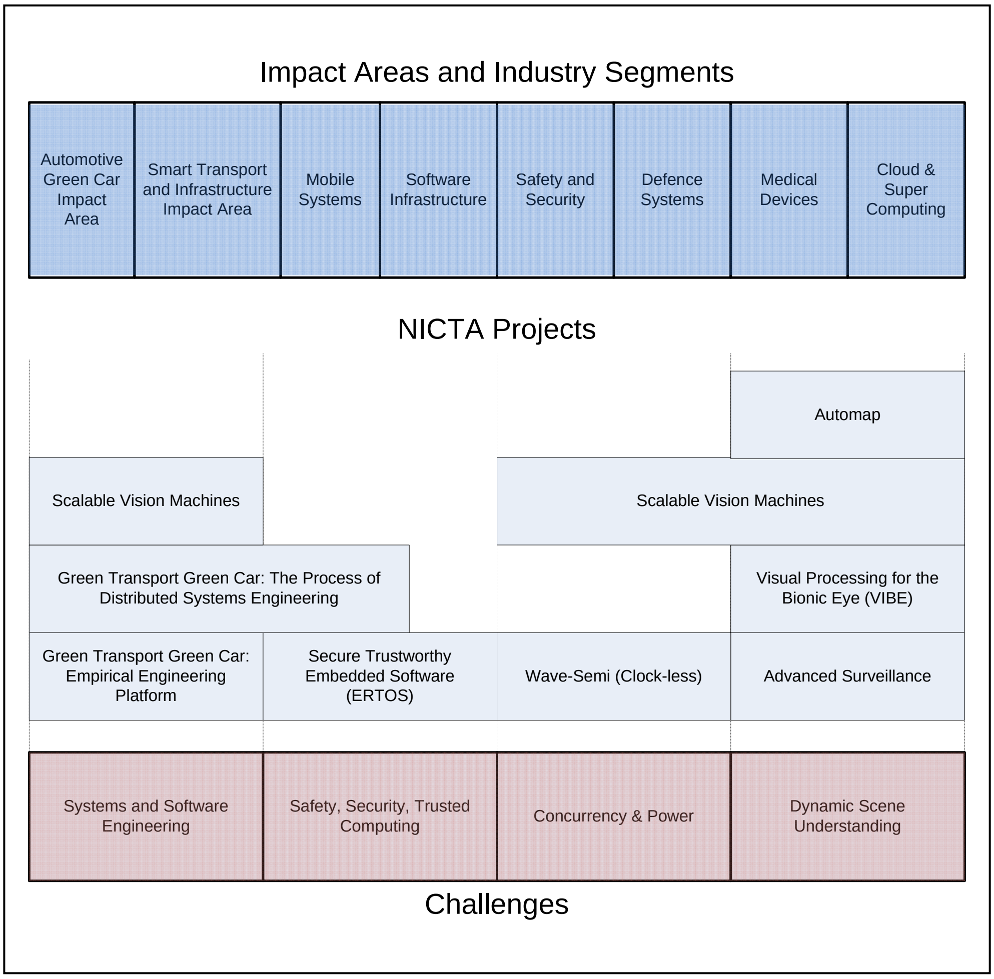

What is not covered ...

NICTA has a substantial implant technology research program that is part of the embedded systems theme. The development of the bionic eye in the Bionic Vision Australia and the Implant Systems projects are in this category. There is the Human Performance Improvement project in our Canberra lab which could be grouped with the implant technology projects under some kind of biomedical systems challenge but it would be just that, a grouping of projects, and not part of a co-ordinated plan. NICTA works in wireless sensor networks and develops systems technologies in the SmartBasin project. NICTA also has a 60GHz CMOS development that includes RFIC and Digital Baseband. While these projects are substantial in terms of innovation, resources and impact, they are not core to the theme vision. They are, nevertheless, highly innovative in their respective fields and great assets within the Embedded Systems theme.