Microelectronics System Design Research Group

School of Electrical and Electronic Engineering

Enabling Survival Instincts in Electronic Systems: An Energy Perspective

A.Yakovlev

Technical Report Series

NCL-EEE-MSD-TR-2013-181

May 2013

Contact: alex.yakovlev@ncl.ac.uk

Supported by EPSRC grant EP/K012908/1 and EP/G066728/1

NCL-EEE-MSD-TR-2013-181

Microelectronics System Design Research Group
School of Electrical and Electronic Engineering
Merz Court
Newcastle University
Newcastle upon Tyne, NE1 7RU, UK

http://async.org.uk/

ENABLING SURVIVAL INSTINCTS IN ELECTRONIC SYSTEMS: AN ENERGY PERSPECTIVE

Written as a chapter to Peter Cheung's 60 Festschrift (3rd May 2013)

Alex Yakovlev

School of Electrical and Electronic Engineering, Newcastle University
Newcastle upon Tyne, NE1 7RU, England
E-mail: alex.yakovlev@ncl.ac.uk

The writing of this paper has been inspired by the motivating ideas of incorporating self-awareness into systems that have been studied by Prof Cheung in connection to dealing with variability and ageing in nano-scale electronics. We attempt here to exploit the opportunities for making systems self-aware, and taking it further, see them in a biological perspective of survival under harsh operating conditions. Survivability is developed here in the context of the availability of energy and power, where the notion of power-modulation will navigate us towards the incorporation into system design of the mechanisms analogous to instincts in human brain. These mechanisms are considered here through a set of novel techniques for reference-free sensing and elastic memory for data retention. This is only a beginning in the exploration of system design for survival, and many other developments such as design of self-aware communication fabric are further on the way.

1. Introduction

Complex information and communication systems have been studied for a long time. Many approaches and methodologies for their modelling, analysis and design exist to date. Amongst the properties of interest in

those studies a prominent place is occupied by the property of systems to stay alive and function in spite of harsh environmental conditions that may surround them. Typically such conditions are assumed to generate higher rates of errors such as those that are for example caused by radiation. They are considered mostly in the scope of information processing, and to a lesser extent in the domain of resource availability, for example, the availability of energy, the mother of all resources. While the system may remain fully functional under the nominal conditions of energy supply, its behaviour may be highly unpredictable when the energy flow to the system is impaired for one or another reason. Design of systems with varying power modes is a rapidly emerging area of research, and it comes from many different directions; for example, intelligent autonomous systems, systems with energy harvesting, green computing etc. Much of this research is about systems that are sufficiently complex that even their most energy-frugal mode of action still requires a certain stable level of energy flow. What about systems that have to 'live on the poverty line', the conditions in which power levels drop to zero and systems that have to self-recover upon the arrival of the 'first beam of sunlight'?

In this paper we shall look at the first glimpses of, perhaps, still naive, approaches to building electronic computer systems whose power sources can be defined in a wide band of modes. Such systems will effectively need survival instincts as part of their intrinsic characteristics. An important element of this new design discipline is a close link between the design methods required for power conditioning and those necessary for computational blocks as the latter form the load in the overall power chain. This proximity and even interplay of energy and information flows, and associated holistic nature of system development activities, is what drives us towards a new type of co-design, which involves new methods for modelling, simulation, synthesis and hardware and software implementation. This paper will address a number of paradigms for such designs, such as power-modulated computing and elastic system design. It will present examples of problems formulated and solutions obtained in the context of research on the new generation of systems with higher self-awareness for survivability. A prominent place in this exploration is taken by what we call reference-free sensing, which allows the system to check its power conditions without relying on external references in voltage or clock.

On-chip sensing is generally a very important area of research in modern times due to the high variability of devices produced in nanometre technologies. Before any piece of fabricated silicon is put into action, it has to be measured and tuned to help its performance best meet its individual characteristics. Ageing is another factor that requires adaptation of functional settings, voltage and frequency scaling, throughout the lifetime of the system. This has been realised by Prof. Peter Cheung and his co-workers at Imperial College who investigate methods for health monitoring of chips, exploring their individual character and looking for ways of run-time performance optimisation (e.g. [1]). In many respects the various built-in self-awareness facilities for adaptation to variations and ageing are similar to those for survival. This interesting relationship and long term professional friendship with Prof Cheung has inspired the author in writing this paper for such a wonderful occasion!
Before we start our journey into the subject of this work, it would be pertinent to bring two important quotations:

"The very essence of an instinct is that it is followed **independently of reason**." 1871: C. Darwin Descent of Man I. iii. 100

"The operation of instinct is **more sure and simple than that of reason**." 1781: E. Gibbon Decline & Fall (1869) II. xxvi. 10

We bring these quotations with one purpose. For our study of certain basic functionalities in electronic systems that are retained in the conditions of austerity, we need an analogy with biology. Biological world is the realm where survival is a key property of organisms, whether it concerns each organism individually or organisms as a species. As we postulated above, instincts are seen as something which is adherent to survival. So is the importance of these quotations - they define the place and role of instinct along and in comparison with reason, something that is regarded as highest form of biological activity. Armed with this analogy, we will start looking at the ways of how electronic systems can be built where their 'reason' parts operate along with their 'instinct' parts. The outline of topics discussed in this paper is as follows:

- Bio-inspiration: survival instincts in real life.
- "Survival instincts" in ICT systems.
- Energy-Power modulation and layers of functionality.

- Mechanisms in energy and data processing:
    - Reference-free sensing,
    - Elastic memory for data retention,
    - Elastic power supply for survival
- Future developments.

2. Survival and Instincts in Real Life

So, what are survival and instinct in general terms? Among the many definitions of survival and instinct that can be found in OED, perhaps the following serve our needs best: **"Survival**: The continuing to live after some event; remaining alive, living on". "**Instinct**: (a) An innate propensity in organized beings (esp. in the lower animals), varying with the species, and manifesting itself in acts which appear to be rational, but are performed without conscious design or intentional adaptation of means to ends. Also, the faculty supposed to be involved in this operation (formerly often regarded as a kind of intuitive knowledge). (b) Any faculty acting like animal instinct; intuition; unconscious dexterity or skill".

If we were looking at instincts from biological or even psychological perspective, we would have distinguished between instinct and intuition. In our present analysis, we will also do that, and see intuition as, perhaps, the highest form of instinct that is close to reasoning. It is akin to prediction in information systems, which often connects higher forms such as reasoning with sensory-signalling forms. In our analysis we will not go to the level of intuition analogy, but rather stay at the level of basic instincts. What's more we will mostly approach instincts from the perspective of energy in the system, and see how energy or power levels determine the role of instincts, particularly focusing on their manifestation under the low energy conditions.

To get better sense of how instincts may reveal themselves both structurally and behaviourally, we illustrate them in the following way. Firstly, we bring an example of a 'case study' which shows the energetic aspect of instincts quite vividly. A few years ago, the world had heard a story about a French cave explorer Jean-Luc Josuat, who got lost in a cave and spent there five weeks without food and water before he was found by his rescuers. During this ordeal his first (conscious) reaction was to actively search for food - due to orexin, a hormone produced in the hypothalamus; orexin is normally generated to trigger alertness and

all parts of the body to work faster. But at a later stage, some 'more hardwired' instincts (inherited by humans from more primitive species through evolution) started to prevail in the brain and everything slowed down to ensure survival when energy sources became short. There is a video about this case on YouTube that can be accessed from this website: http://videos.howstuffworks.com/discovery/6835-human-body-built-for-survival-video.htm

Secondly, a good illustration of where instincts rest in humans is provided by Paul McLean's triune model. The model states that the human brain has three independent (and behaviourally concurrent!) brains, which were developed successively in response to evolutionary needs. They are reptilian (responsible survival), paleomammalian or limbic (responsible for emotions) and neomammalian or neocortex (responsible for higher-order thinking). The lowest one, reptilian brain (or R-complex), is the one which is inherited from reptiles. This is where our instincts rest. This brain is active all the time even in deep sleep. We do not sense this reptilian brain in our consciousness under normal conditions. But in the conditions like those of Jean-Luc Josuat's ordeal the R-complex takes control of our bodies to help them survive.
So in this paper, we strongly hypothesize that the manifestation of these different brains is driven by the energy levels in the body, and with this hypothesis we enter the cyber-world and think of electronic systems of the future – with the idea of Darwinian evolution also being transferred to the cyber-world.

3. Survival and Survivability in Electronic Systems

Let's now turn our attention to artificial systems, like information systems, and raise two key questions about survival: "survival from what?" and "survival of what?" First of all, let's see what sort of 'disasters' we should imagine that the systems would need to survive from. We can roughly categorise them into the following three groups:

(1) Faults and degradation inside the system: defects, ageing, transients (inside gates, crosstalk on signal lines, IR drops).
(2) Upsets outside the system: radiation, power supply drops, signal distortions.
(3) Miscellaneous physical effects (both internal and external): temperature fluctuations, electro-magnetic interference.

Now, what aspects of the systems can we consider for survival? They are mainly, but not exclusively: structure, behaviour, and specific (or purposeful) functionality (defined by the system's user for example).
Combining the sources of impairments and their effects on the system, one would conventionally consider ways of how the system would react to them. Here, the reader might see some relationships if not similarities between the property of survivability and following properties, sufficiently well explored in the ICT domain: tolerance, resilience, recoverability, longevity etc. (It is very tempting to start thinking about such even 'more biological' properties such as reproducibility, especially if our notion of survival may one day stretch to thinking about genetics and preservation of species – well, in a few years with the developments in DNA computing we may have a chance!). Let's, at least, briefly contrast survivability with two fairly common properties:

- Dependability (Fault-tolerance …):

Dependable systems typically want to restore their full functionalities, hence large costs for redundancy; survivability is supposed to be less resource-demanding, or in other words the system may continue to work even with incomplete power levels.

- Graceful degradation:

Gracefully degrading systems typically have a smooth (often quantitative) reduction in their performance (cf. today people talk about approximate computations and trade-offs between accuracy and quality of service), rather than "qualitative" transitions to a more restricted (more critical) set of functionalities as needed for survival.
From these two brief comparisons we can see that the key difference between survivability and other seemingly similar properties lies in the way how we approach the energy aspect. We start to talk about survivability when the system's power is variable, intermittent, sporadic etc. Of course, the scale and range of power and energy disruptions would matter here as well, but in our simple approximation, the notion of survivability, similar to biology, refers first of all to the power conditions. For years, ICT systems have been designed to be fault-tolerant, robust and resilient to faults, ageing etc, but they have always been assumed to be fully powered. Of course, otherwise, how can one activate the fault-detection and correction procedures and engage recovery mechanisms.
At this point, however, the reader might actually stop us by saying that survivability has been studied in ICT. Indeed, it has – but conventional

survivability in ICT is more about software systems (cf. [2]) that make transitions between different services depending on the operating environment.

What we are interested in here is different. It is what we call "Deep, or Instinct-based, Survival", as opposed to conventional survivability, where again, as it is about software, there is very little scope to think about serious power-related issues, such as power deficiency or interruptions.

So, conventional survivability does not consider deep, embedded layers of hardware/software that work in proportion to the level of available energy/power resources. Thus, **Deep Survival** is a new concept, inspired by nature, which maintains operation in several structural and behavioural layers, with mechanisms ("instincts") developed and accumulated in bodies due to biological evolution. So, we end this section by postulating that survivability cannot be achieved in the system without providing it with sufficient back-up in the form of instinct. And, as we can see it from our quotations of Darwin and Gibbon, we must really talk about an independent layer of activity in the system's structure, so independent that even the ways of its powering are independent of those of the 'reasoning' layers. We will therefore have to first look at how power may modulate the system's functionality, the subject of our next section.

4. Power-Modulated Computing and Functionality Layers

In this paper we postulate that the principle of **power (energy)-modulated computing** [3] is fundamental for deep survival. In other words, until and unless we start designing systems in such a way that the incoming power is actually the driver of the functional behaviour we will not be able to build systems that can survive. Yet, putting it even stronger, until we only limit our design approaches to power-efficiency rather than power-modulation, our systems will not be fully survivable. Here are some further arguments in favour of this view.

Any piece of electronics becomes active and performs to a certain level of its delivered quality in response to some level of energy and power. A quantum of energy when applied to a computational device can be converted into a corresponding amount of computation activity. Depending on their design and implementation systems can produce meaningful activity at different power levels. As power levels become

uncertain we cannot always guarantee completely certain computational activity. Good characterisation of power profiles for the system in space and time is important for designing systems for survival. Fig. 2 illustrates this idea.

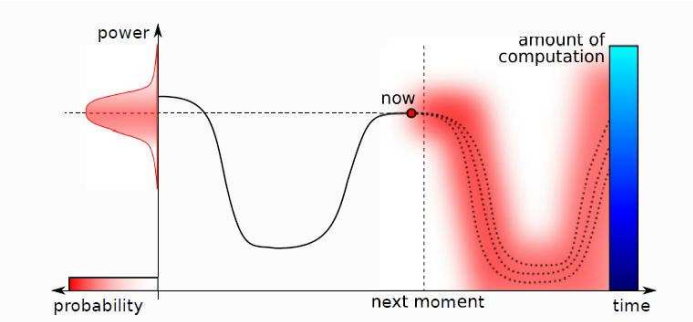

Fig. 2. Power profile in time, its uncertainty and illustration of power-modulated computing.

At any moment in time we have a determinate trace of power supply in the past but the future is indeterminate. The system, thanks to its sensing abilities, and initial forms of intuition, can make some localised prediction, from every moment at present and its ability to compute (say in terms of the rate of activity) will be determined by the actual power levels. This brings us to the link with the recently published ideas of power-proportional computing [3, 4]. Power proportionality however has two forms. One, more conventional, form concerns the fact that the system is power-proportional when its power consumption is proportional to its service demand. When systems are driven by the service demand they tend to follow the principle of multi-modality, where the system "consciously" switches between a full functionality mode to a hibernating mode primarily depending on the data processing requirements. Survival aspects here are limited to the ability of mode management.

But what if the power level drops? Here we face with the second form of power-proportionality, which in our view lends itself to a more general form of survivability. To extend the frontier of survivability, system design should also follow the **power-modulation approach**, and this leads to structuring the system design along partially or fully independent layers (cf. Darwin's "The very essence of an instinct is that it is followed **independently of reason**.")

Multiple layers of the system architecture **can turn on/off at different power levels** (cf. analogies with living organisms' nervous systems or underwater life, or layers of expensive/cheap labour in most of the resilient economies). As power goes lower higher layers turn off, while **the lower layers ("back up") remain active** – this is where instincts become more in charge!

The more active layers the system has, the more resourceful and capable of surviving it is. This layered view is reflected in Fig.3, which puts it in analogy with the sea layers and ability of different forms of life to survive in different conditions of sunlight penetration.

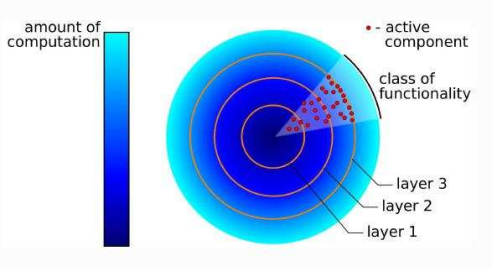

Fig. 3. Layered computational activity in response to power levels.

Fig. 4 illustrates the difference between traditional and energy-modulated system design. In the next section we will attempt to present our list of most basic instincts that the system needs to maintain for survivability.

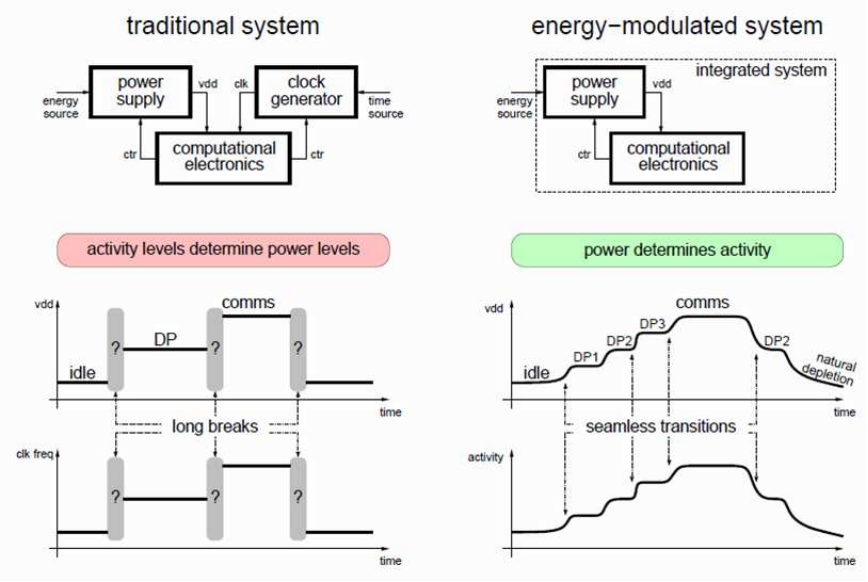

Fig.4. Traditional versus energy-modulated design.

5. Basic Instincts: Self-Awareness and New Sensing

The following categories of instincts can be identified in electronic computer systems that can help them to be better equipped for survival. The most important is probably energy/power-awareness, i.e. sensing, detection and prediction of power failures. The next one is the ability of storing energy "for the rainy day". Other instincts involve mechanisms for retaining key data, reactive and optimising mechanisms, and layers of power-driven functionality.

These instincts cannot work without the following basic abilities and associated actions:

- ability to accumulate some energy, initially and at any time after long interruption, say by charging a passive element;
- ability to switch, e.g. generate events;
- ability to decide, e.g. whether there is an event or not.

These actions underpin two major categories of instinct-supporting mechanisms:

- Mechanisms in energy and data processing domains:
  - Reference-free self-sensing and monitoring [5-8],
  - Retention memory for survival [9],

  - Elastic power-management for survival [10].
- Mechanisms in communication fabric:
  - Monitoring progress in transactions (link level failures, deadlock detection) [11,12],
  - Power noise and thermal monitoring [13],
  - Non-blocking communications [14].

In this paper we restrict ourselves by discussing only the first category of mechanisms. An interested reader may find the description of mechanisms in the second category in our papers [11-15]. Our main focus here in on self-awareness, hence sensing is our priority. Sensors must work in changing environments with uncertainty, where constant and reliable references are not available. Traditionally, sensors used in electronic systems are quite heavy – their purpose is to convert some physical form of information into digital form so that it can be processed in the computing system. Normally, this is done with the purpose of digital signal processing with fairly high requirements for fidelity and signal-to-noise ratio. This leads to having sensors with fully-fledged A-to-D converters involving accurate voltage or time references supplied from outside. In systems that are autonomous and in the conditions where the aim is to survive this is not possible. Hence our target is design an entirely different sort of sensors. In this paper we focus on the so-called **reference-free** sensors, where we will consider the following options:

- Sensing by charge-to-digital conversion;
- Sensing by differentiators in delays;
- Sensing by crossing characteristic mode boundaries such as oscillations;
- Sensing by measuring metastability rates;

All of these sensors have some digital parts whose behaviour is modulated by the voltage that they sense, and this voltage is connected to the power terminal of the digital part. In this way these sensors are inherently power-modulated. We shall now describe some of such sensors.

5.1 **Sensing by Charge-to-Digital Conversion**

This method involves sampling the input signal into a capacitor in the form of its electric charge and then discharging the capacitor in such a

way that its charge is converted to digital code. Basically it is inspired by the challenge of building a sensor that is powered by the energy of the sensed signal itself. So, the principle of operation of such a sensor is that energy sampled in the capacitor as charge is proportional to the sensed voltage. It is then discharged through some load registering the quantity of energy (just like in a waterwheel!). As such a load we can use a self-timed counter as shown in Fig. 5. The bottom part of the figure shows voltage on the capacitor as a function of time. We have investigated this relationship and found that it is subject to a complex behaviour of the switching gates in the counter, which are defined by the characteristics of their constituent transistors in different modes and mechanisms, including superthreshold, subthreshold, leakage etc. Under reasonable approximations the analytical characteristic of voltage versus time is a hyperbola rather than exponential while the transistors operate in superthreshold mode [15].

Let's now discuss the reference-free issue in this method. In the absence of external voltage and time references, we still need to control time in order to decide when to stop the discharging process while the level of voltage in the counter is sufficiently high, so the code stored in the counter can be recorded before the counter stops counting. We should stop counting irrespective of Vin – constant sensing/conversion delay.

However, this "same time" implies timing reference or some clock. Hence we need to produce a voltage level Vd such that is a constant reference. Vd could be based on some internal constant such as the threshold of a transistor (similar to the idea of bandgap).

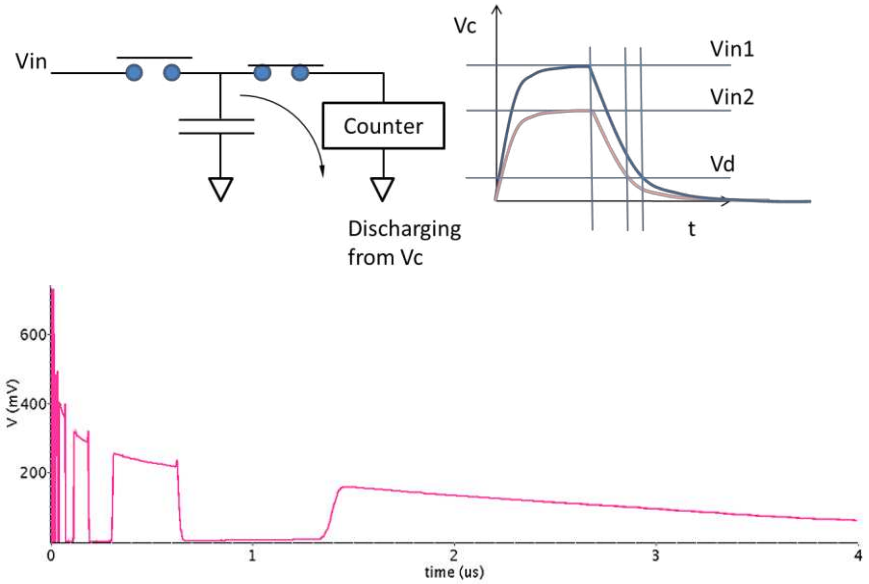

Fig. 5. Charge-to-digital conversion principle.

The circuit shown in Fig.6 illustrates how the control circuit and internal reference generator can be built. The waveform in this figure shows that the event of crossing the second threshold corresponds to stopping the counting and latching the code from the counter. We have designed and fabricated a sensor chip in 180nm TSMC via Europractice. We connect the chip to a 10nF sampling capacitor and tested the sensor – the results are plotted in the above figure.

This experiment has shown the feasibility of building a sensor that is powered by the signal it senses and that is reference-free. In the following sections we will show ideas for building sensors that can be used in the highly variable conditions. We have not yet brought them to the same level of experimental implementation as the above sensor, but there are plans to do that.

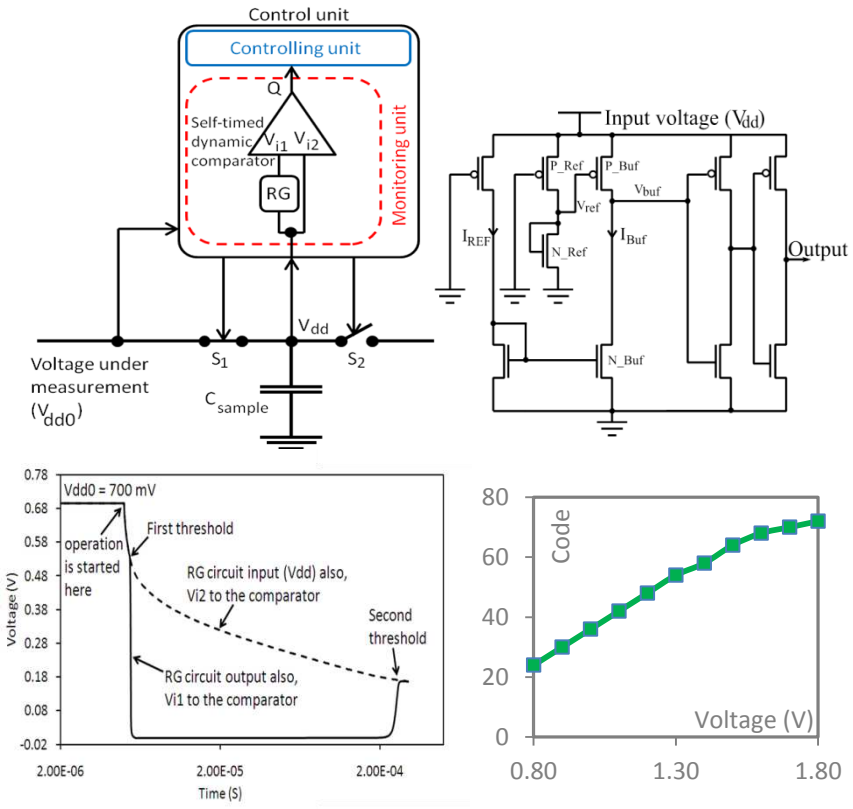

Fig. 6. Sensor control and its internal reference generator, the timing controlled by RG and the measured code vs input voltage (data from the fabricated 180nm chip).

5.2 **Sensing by Delay Differentiators**

The idea of sensing using delay differentiators is as follows. We need to design two circuits, which can operate in a range of voltages of our interest. The circuits, however, must have their delays scaled differently with the supplied voltage, as shown in Fig.7 (left hand side). If this is the case, then the difference between these delays will represent some characteristic form of (ideally, proportional in some critical range of interest) dependence on the supply voltage. The right hand side of Fig.7 shows that the digital value of the measured voltage can be obtained by

measuring the time when Circuit 1 finishes against Circuit 2. We thus need a mechanism of registering the position of where the signal is in Circuit 2 when Circuit 1 is finished.

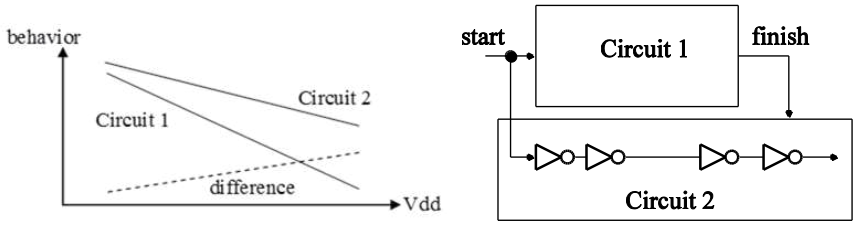

Fig. 7. Principle of delay difference based sensing.

For example, we have observed the difference (mismatch) in delay scaling between SRAM cells and logic gates, as shown in Fig.8. This mismatch rapidly increases (in terms of the number of inverters that need to match the delay of the SRAM cell, which acts as Circuit 1) when Vdd drops below 0.7V (for 90nm technology). Now, replacing the line of inverters with a self-timed counter (similar to the one used in the charge-to-code converter), to act as Circuit 2, which is started together with the SRAM cells and stopped when the reading (or writing) the cell finishes (we used a self-timed SRAM with explicit completion detection), allows us to register the binary code for the delay difference. This is shown in Fig.9 on the basis of spice simulations for a 90nm technology node. Although the linearity of this sensor is quite limited, it can still be used for the purposes of condition monitoring we are interested, and what's really important it is completely reference-free.

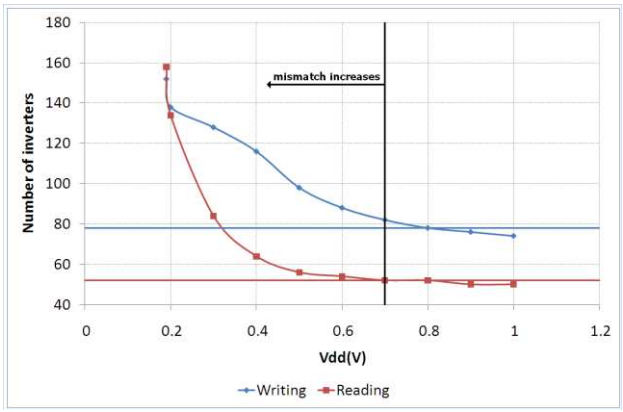

Fig. 8. Mismatch between inverter chain and memory cell delay (90nm technology)

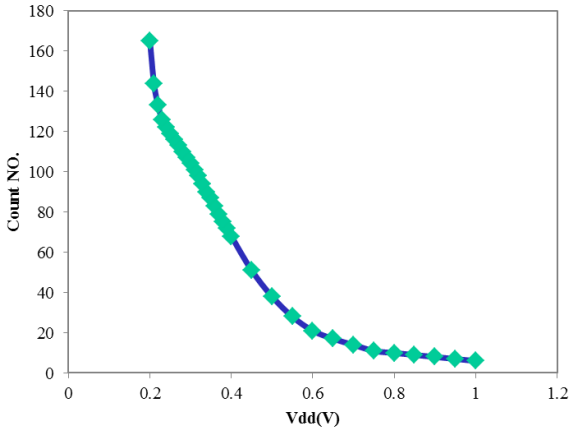

Fig. 9. Voltage sensing result using memory-logic mismatch.

5.3 **Sensing by Oscillation Detection**

It is often the case that we need to sense voltage, say power supply, only to the point where it crosses certain level, for example, the level at which some 'reasoning' parts of the system can no longer be trusted. This kind of sensing can be done with a circuit which changes its operating mode, for example, from stable to oscillatory. An example of such a threshold-crossing oscillator is shown in Fig. 10. It consists of two stages, each containing a pair of forward (F) inverters and a pair of cross-coupled

(CC) inverters. The circuit has two operating modes: oscillation and latching/locking. When the supply voltage Vdd drops below the certain Vthr level the circuits oscillates, as shown in Fig.11.

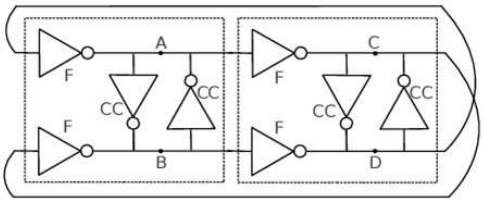

Fig. 10. Voltage modulated oscillator.

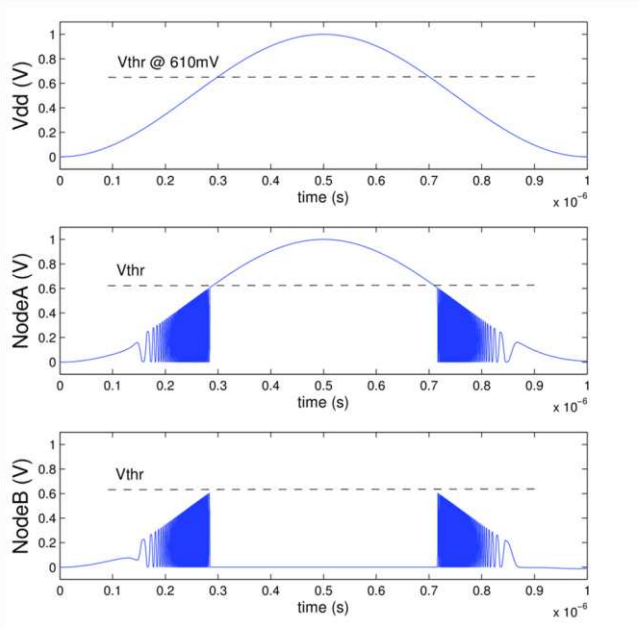

Fig. 11. Voltage modulated oscillation.

The specific value of Vthr can be defined at design time by setting the ratio between transistor sizes:

$$r = \frac{width\ of\ CC}{width\ of\ F}$$

The effect of the ratio on the Vthr is shown in Fig.12. The overall setup for detecting an event of Vthr crossing via oscillation is shown in Fig. 13. It uses a self-timed counter, initially reset to zero but introduces some delay of counting until the most significant bit is set to 1, to guarantee that the oscillations are stable. As before, it is easy to see that this

method of sensing is free from external references. The behavior is completely determined by the internal characteristics of the devices.

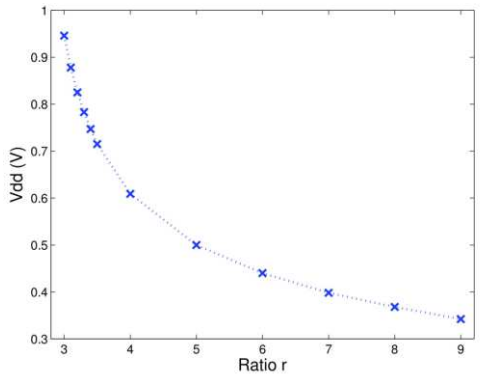

Fig. 12. Transistor size ratio vs Vdd at which the circuit oscillates.

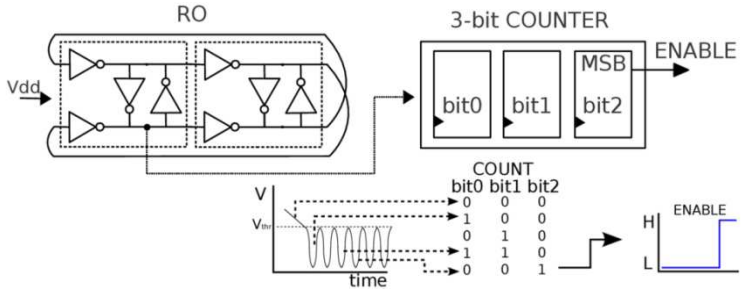

Fig. 13. Setup for oscillation-based sensing.

5.4 **Sensing by Measuring Metastability Rates**

Finally we present another technique for voltage sensing (it is also applicable to temperature sensing). It is based on the use of metastability in bistable devices. Metastability offers a nice way of removing external references in Voltage and Temperature sensor. When the setup and hold time conditions of a flip-flop are not met, the flip-flop may become metastable. A metastable flip-flop will take extra time to decide whether to go logic high or low (decision time = clock-to-q delay). The "decision making" time constant (τ) is a function of Vdd. So, the idea of the method is to use the time constant (τ) to quantify Vdd. What we need to do is to count the rate at which the flip-flop fails to decide!

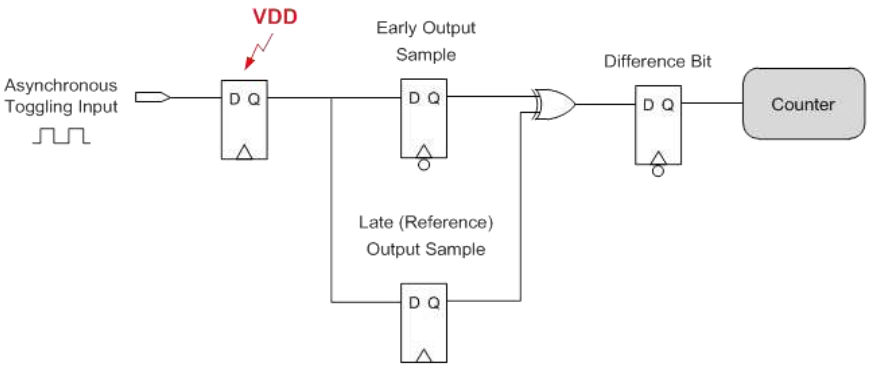

Fig. 14. Circuit for metastability rate measurement

The sensor circuit shown in Fig. 14 works is as follows. Firstly, the left-most flip-flop (call it FF1) often becomes metastable because its input is asynchronous. Secondly, when FF1's output is delayed, the early and late samples of FF1's output (captured at the following falling and rising edges respectively) will be different. Finally, the counter counts these instances. Its output after a fixed period of time is an exponential function of the time constant τ, which is determined by the sensed parameter. The advantages of this method are that it is purely digital, very compact and offers sufficiently high precision. We proved this concept in FPGA (Altera Cyclone II), and the results are shown in Fig. 15 (in a semi-logarithmic plot).

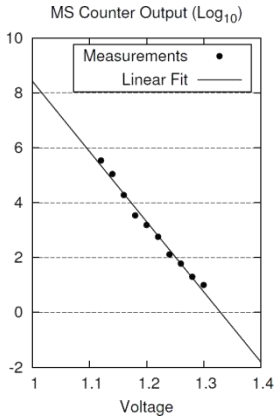

Fig. 15. Voltage sensing results for the FPGA prototype

6. Elastic Memory for Data-Retention in Instincts

We now illustrate a way of designing retention storage (SRAM) for survival. We call it elastic because it is completely self-timed and operates correctly in a wide range of supply voltages, both stable and time-varying. This SRAM can be built around different types of cells; for example, we have designs for 6T and 10T cells. One can use a 6T solution for energy-efficiency and 10T for core-function survivability. One can build control for such an SRAM array with different types of completion detection, again depending on the need to mitigate variation between columns. For example, a version with more economic completion detection (data bundling) is shown in Fig. 16.

A **speed-independent control** circuit for the SRAM is shown in Fig. 17. The timing diagram shown in Fig. 18 shows the simulation trace for the SRAM as it works for Data Write with the time varying Vdd supply. It is easy to notice the response time with which the memory sends the Wack signal is modulated by the Vdd (for smaller Vdd the delay between Wreq and Wack is longer).

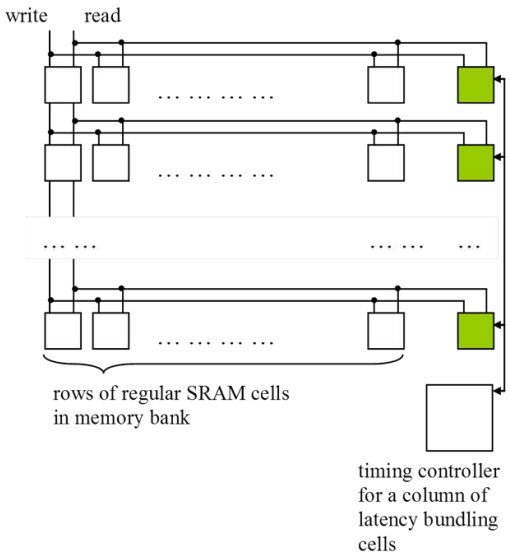

Fig.16. SRAM array with data bundling

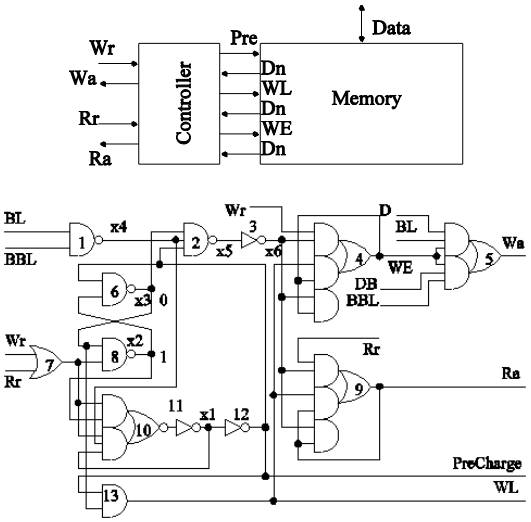

Fig. 17. Speed-independent control circuit for SRAM

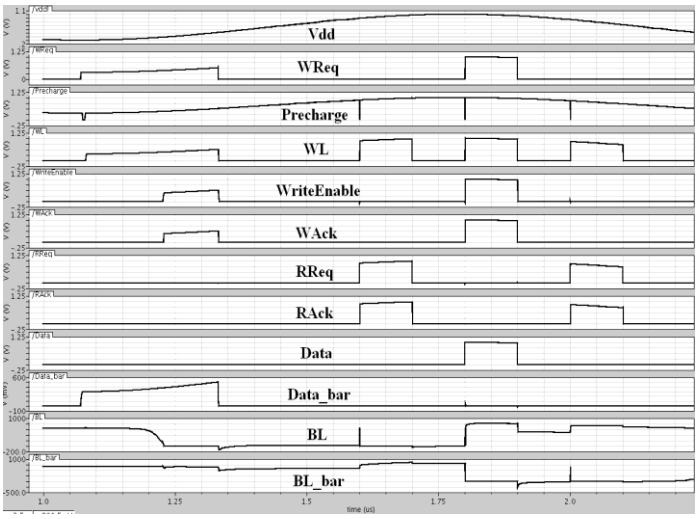

Fig. 18. Write simulation for time-varying Vdd

For a 6T case we have built an ASIC prototype to prove the concept. The layout of the die is shown in Fig.19. The chip was successfully tested and one of the traces of the Wack signals captured by oscilloscope clearly

shows the effect of the switching behaviour modulated by Vdd (one can observe the Vdd changing in a quick-charge-slow-discharge shape).
During testing the chip we discovered interesting effects of self-timed SRAM which confirm its time elasticity and useful properties for survival. Despite the fact that in the simulation we saw the circuit working down to the level of 190 mV, the real silicon showed that the SRAM worked steadily for Vdd above 0.75V, after which its control logic 'froze' in either its setting or resetting phases. This can be observed in the trace of Fig.19 where the Wack signal gets 'stuck' either in the low or high state. Interestingly, that due to the speed-independent nature of the circuit, the circuit smoothly recovers from the 'frozen' state as soon as Vdd goes back to the level above 0.75V. What's important is that when Vdd is below 0.75 the data is safely retained in the SRAM (this was checked during the testing process). The data is retained while Vdd is greater than 0.4V.
The above behaviour shows that a fully speed-independent SRAM is excellent as retention storage for survival in power-deficient regimes. It provides self-detection of the power condition by 'freezing', an early warning, well before the system starts to lose its data.

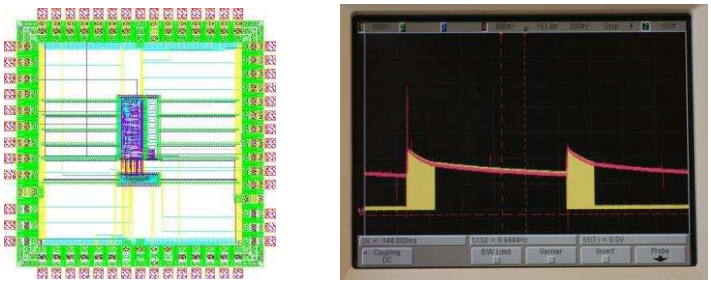

Fig. 19. SRAM chip layout (UMC 90nm Europractice) and control signal trace for varying Vdd

7. Retaining Energy: Elastic Power Management for Instincts

The design of ICT systems destined for survival will increasingly be more holistic and will have to take care of not only their data processing parts, such as sensing and computing electronics, but also their power supply electronics. We are exploring new ideas in this direction. They involve more active use of switched capacitor circuits for DC/DC

conversion. Conventionally there are switched capacitor DC/DC converters (SCCs). They convert constant input Vdd to constant output Vdd according to a set of ratios. However, SCCs usually rely on the availability of stable sources of time and voltage references. Instead, under harsh operating conditions such references may not be available. Hence, we develop a different type of switched capacitor circuits that are aware of the presence of self-timed circuits as their load. We call them capacitor bank blocks (CBBs). We have also designed hybrid CBBs that can work as SCCs and CBBs depending on the conditions and whether the load electronics is synchronous or asynchronous. Details of this method can be found in [10].

8. Conclusions and Outlook

As stated in the abstract and introduction, this paper was inspired by the ideas of incorporating self-awareness into systems that have been studied by Prof Cheung in the context of improving the performance of electronic systems under process variations and ageing. We take self-awareness further, and with the help of biological analogy, consider survival instincts here. The paper has focused almost exclusively on the techniques and examples of circuits for survivability that support an 'instinct layer', which is supposed to remain alive and operational under the conditions of power instabilities and lack of power.

We are currently involved in an EPSRC-funded project "Staying alive in variable, intermittent, low-power environments" (SAVVIE), in collaboration with Dr Bernard Stark of University of Bristol. The project's main aim is to develop techniques for enabling systems to survive in the top left corner of the energy-power state space depicted in Fig. 20. While there exist methods that support trajectories like T1 and T2 in this state space, approaches to cater for trajectories such as T3 and T4 are in their infancy. We hope that the ideas that have been described in this paper will contribute to this aim.

A list of outgoing research directions we are currently pursing:

- More diversification – power and data processing paths intertwined, mixed digital and analogue fabrics, synchronous and asynchronous fabrics, multiple technology fabrics.
- New modelling and design approaches – models that capture multi-modal and multi-layer architectures; combining structure and behaviour in models, capturing overlay in functionality.

There is plenty to investigate on this path, and the research is already under way at Newcastle and in collaboration with our partners from Southampton, Imperial and Manchester under the programme grant PRiME that will explore energy-reliability tradeoffs in designing future many-core embedded systems.

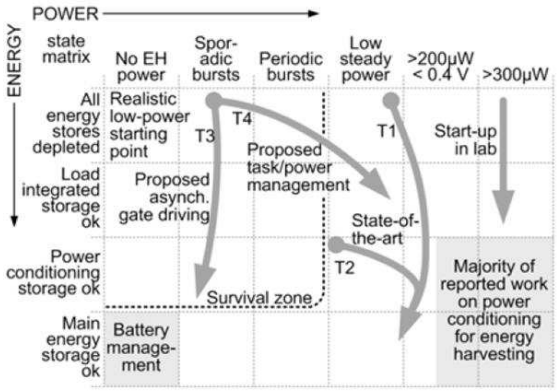

Fig.20. Energy-power state space in the SAVVIE project (courtesy of B. Stark).

Acknowledgments

This paper would not have been written without the combined effort of the author's research team, the Microelectronics Systems Design group at Newcastle. The list of people actively working in this area can be found on the group's webpage: http://async.org.uk.
Our collaboration with Dr Bernard Stark's team at Bristol in the SAVVIE project and our recent collaboration with Universities Southampton, Bristol and Imperial College London under the Holistic project (http://www.holistic.ecs.soton.ac.uk/) are acknowledged with deep gratitude.

References

1. J.M. Levine, E.A. Stott, G.A. Constantinides, and P.Y.K. Cheung, Online Measurement of Timing in Circuits: For Health Monitoring and Dynamic Voltage & Frequency Scaling, FCCM 2012, pp. 109-116.

2. J.C. Knight and E.A. Strunk, Achieving Critical System Survivability through Software Architectures, Proc. Architecting Dependable Systems II, LNCS 3069, 2004, pp. 51-78.
3. A. Yakovlev, Energy-Modulated Computing, Proc. DATE 2011, pp. 1340-1345.
4. R. Ramezani, D. Sokolov, F. Xia and A. Yakovlev, Energy-modulated quality of service: New scheduling approach, Proc. Faible Tension Faible Consommation (FTFC), 2012, pp.1-4.
5. R. Ramezani, A. Yakovlev, F. Xia, J.P. Murphy and D. Shang, Voltage Sensing Using an Asynchronous Charge-to-Digital Converter for Energy-Autonomous Environments, IEEE J. of Emerging and Selected Topics in Circuits and Systems, Vol. 3, No.1, March 2013, pp. 35-44.
6. D. Shang, F. Xia and A. Yakovlev, Wide-Range, Reference free, On-chip Voltage Sensor for Variable Vdd Operations, Proc. ISCAS 2013, Beijing, May 2013.
7. G. Tarawneh, T. Mak and A. Yakovlev, Intra-chip Physical Parameter Sensor for FPGAs using Flip-Flop Metastability, Proc. FPL2012, pp. 112-119.
8. I. Syranidis, F. Xia, A. Yakovlev. A reference-free voltage sensing method based on transient mode switching, Proceedings of the 2012 8th Conference on Ph.D. Research in Microelectronics and Electronics (PRIME), pp.143-146.
9. A. Baz, D. Shang, F. Xia and A. Yakovlev. Self-timed SRAM for Energy Harvesting Systems, Journal of Low Power Electronics, vol. 7, no. 2, April 2011, pp. 274-284, American Scientific Publishers.
10. X. Zhang, H. S. Low, D. Shang, F. Xia, and A. Yakovlev, "A Hybrid Power Delivery Method for Asynchronous Loads in Energy Harvesting Systems", Proc. 10th IEEE International NEWCAS Conference (NEWCAS2012), pp. 413-416
11. L. Dai, D. Shang, F. Xia and A. Yakovlev, Monitoring circuit based on threshold for fault-tolerant NoC, Electronics Letters, Vol. 46, Issue 14, pp. 984-985.
12. R. Al-Dujaily, T. Mak, F. Xia, A. Yakovlev, and M. Palesi, Embedded Transitive Closure Network for Runtime Deadlock Detection in Networks-on-Chip, IEEE Transactions on Parallel and Distributed Systems, vol.23, no.7, pp.1205-1215
13. N. Dahir, T. Mak, F. Xia and A. Yakovlev, Minimizing Power Supply Noise through Harmonic Mapping in Networks-on-Chip, Proc. CODES+ISSS 2012, ACM, NY, pp. 113-122.
14. F. Xia, A.V. Yakovlev, I.G. Clark and D. Shang. Data communication in systems with heterogeneous timing, IEEE Micro, vol. 22, No. 6, pp. 58-69, Nov./Dec. 2002.
15. R. Ramezani and A. Yakovlev. Capacitor Discharging through Asynchronous Circuit Switching, Proc. ASYNC 2013, Santa Monica, CA, May 2013.