# Soft Time in hARd Space (STARS), or mapping designs with blocks cut loose

A. Yakovlev and A. Bystrov, University of Newcastle

Microelectronic systems design is a creative process involving the following triad: *problem* description, *implementation technology* and a *human designer* equipped with design tools. The ways in which this triad develops are ruled by many cost functions and quality criteria. Research in microelectronic design should approach this triad from different angles and in step with the evolution of all three components. As novel problems arise and processes and tools advance, new ways of optimising system designs according to the required criteria must be found. It is important that *academic research addresses the most fundamental aspects* of this evolution rather than deals with specific engineering tasks. One such a fundamental issue is the cost of human effort. Should the designers be mostly involved in embedding the problem into the device and process constraints? Or should they be mainly focused on finding a good solution at the algorithmic level? With the rapid advances of the fabrication technology, the former option is less and less plausible due to increasing NRE costs, first of all the cost of system validation and testing. In this respect, more research in design should be targeted at methods of "technology-independent" problem solving.

In what way should we aim for *technology-independence*? A fundamental aspect of technology-independence is concerned with the notion of *timing or event coordination*, as it has a strong effect on the complexity of system behaviour and on the task of mapping the system specification into the implementation architecture. On one hand (a problem-oriented view), every algorithm and computation takes place in time, where there are clear boundaries between past, present and future, or between computation steps, or between cause and sequence etc. On the other hand (an implementation view), physical hardware that implements computations accords with physical law of inertia. The input and output values are not changed instantaneously in circuits, and therefore the idea of a value being valid or invalid at particular moments in time is crucial.

Designs of (predominantly) digital systems have traditionally been seen as either clocked (synchronous design) or self-timed (asynchronous design). In the former case, computations in circuits occur under the control of a global clock, whose signals attach a *validity tag* to the values on the circuit's inputs and outputs. In the latter case, the validity tagging is distributed between blocks of the circuit in the sense that the values are transmitted between blocks together with the local validity signals. Both approaches have their own advantages, which help reduce the complexity of the design process through regularising the operation of the system either by means of simple timing assumptions, or by imposing simple handshake (request-acknowledgement) rules. However, both these approaches are insufficiently powerful to sustain the challenges of the microelectronic technology at the nanometre level and the requirements for system functionality arising from new applications.

Both approaches lead to systems that are too rigid in terms of timing, i.e. systems that are essentially synchronous! The global clocking approach has a difficulty in coping with the clock distribution (without losing performance or functional correctness), timing closure, power consumption, EMI etc. The handshaking approach struggles to provide for (hard) real-time conditions, testability, tolerance to soft-errors and transient faults, all because the skeleton of the system is causally interlocked. There may be a need to have parts or all of the system being driven by clocks rather than being governed by handshakes, and at the same time, there may be a need for parts of systems be controlled by the availability of the resources and causal interactions rather than triggered by events. In general, a complex system can be a combination of open-loop and closed-loop temporal links, and the design methodology should be able to adjust to this requirement. Regardless of the synchronisation type, an *ideal design process* should meet a number of generic requirements. Here are just a few examples. Implementation should be adequate to specification. It should be correct by mapping. The issue of transparency, both structural and functional, versus optimality should be addressed. Predictability is guaranteed by transparency with clearly underwritten resources. Multi-level design is crucial, and at each level only the level-specific tasks are addressed. The refinement process should keep the details of all levels so as to facilitate low-level ad-hoc solutions in a systematic way.

In this vision statement we invite the design research community to look in the direction of ***problem-oriented event coordination***, particularly focusing on the description languages and associated analysis and synthesis methods for flexible and diversified paradigms of dynamic behaviour. This behaviour is coordinated by the requirements associated with the problem, both in functional and non-functional sense. Functionally, the level of mutual synchronisation between components can be determined by the real-time conditions or causal relations. Non-functionally, it may be associated with power, performance and area resources, as well as application-specific qualities of the design such as dependability and security.

Possible areas for research projects can be around the following issues.

**Design and test flow aspects:**
Front-end languages for loosely timed systems
Intermediate representations, token-based models
Provably correct and direct mappings
Inter-level and intra-level verification and model-checking
Human interfaces and visualisation
Inter-level design change monitoring
Fault-diagnosis, on-line and off-line testing, self-repair

**Novel circuit methods and IP solutions:**
Irregular and aperiodic signal sampling
Burst event handling
Synchronisation between clock domains
Fine-grained time measurement

**Case studies:**
Hardware event-handlers and resource schedulers
On-line error monitors and testers
Sensor network nodes
Adaptive gain controller

*Significance:*

- System timing is a core aspect of any system design. Our proposal may potentially have serious impact on academic and industrial research because it has a long-term objective of providing an adequate timing-related methodology which adjusts to the evolution of design problems (system drivers) and device technology (process integration).
- The criteria of success can be the overall effect on design technology, e.g. use of new specification tools for loosely timed systems, efficient token-based intermediate representations capable of coping with increasing system size, improvements in yield ramp, reduced time-to-market, ability to measure smaller delays in testing, increased predictability of systems (latency and performance) and better use of power resources.
- Better ways of handling timing in design should lead to the emergence of totally new devices in the areas of medical appliances, consumer electronics, car industry, smart cards, security and defense, which will appeal to general public.
- General public will see the merits of the new approach through the emergence of millions of novel products, quickly to their demands and even beyond their expectations. A good example of that is the sustainability of the Philips' Tangram design technology for handshake circuits!

*Scale*

- This research has clear international scope, with already existing collaborative links between the UK and EU, USA and Japan.
- Possible research projects can be set along the above-mentioned topics. Their goals can be identified in a more specific and short-term way, e.g. power-aware design for processing events in sensor networks, or low latency error handling for on-chip communications.
- Timing infrastructure is "glue matter" to any design. Therefore this research can easily adapt to various problem areas whether they are in cognitive science, biometrics, security, dependable systems, mobile communications, defence to name but a few.

*Timeliness*

- The objectives can be achieved within 10 year timescale. In the last 5-7 years, a good baseline has been created through research into GALS, desynchronisation, systems with strong and weak precedence, priority arbitration, Butler technology. It is thus feasible to address the generic objectives sketched in this proposal immediately.
- The risks are minimal because any advancement along the ways of system timing would bring useful contribution to the design flow.