Heterogeneous Systems on Chip and Systems in Package

I. O'Connor, Ecole Centrale de Lyon, France; B. Courtois, TIMA Laboratory, France;
K. Chakrabarty, Duke University, USA; N. Delorme, CEA-LETI, France;
M. Hampton, Certess Inc., USA; J. Hartung, Cadence Europe, Germany

Abstract

This paper discusses several forms of heterogeneity in systems on chip and systems in package. A means to distinguish the various forms of heterogeneity is given, with an estimation of the maturity of design and modeling techniques with respect to various physical domains. Industry-level MEMS integration, and more prospective microfluidic biochip systems are considered at both technological and EDA levels. Finally, specific flows for signal abstraction heterogeneity in RF SiP and for functional co-verification are discussed.

1 Introduction

From the miniaturization of existing systems (position sensors, labs on chip ...) to the creation of specific integrated functions (memory, RF tuning, energy ...), MEMS and non-electronic devices are being integrated to create heterogeneous systems in package (SiP) and even systems on chip (SoC). However, are the expected advantages worth the investment and effort? In addition, should heterogeneous SoC/SiP be technology or application-driven? This embedded tutorial and panel aims to confront several points of view on this subject from experts in the field.

It is necessary first to capture here the many meanings of the term "heterogeneous" itself. The first and most obvious meaning is that *more than one physical domain* (electrical, mechanical, optical, chemical ...) is involved in the functionality of the system. Another meaning of "heterogeneous" concerns the technological fabrication process and the use of *more than one basic material* (silicon, III-V, organic ...) for the functional devices, whether by co-integration techniques (planar or stacked SoC) or bonding (SiP). When applied to SoC/SiP, it is implicit that one of the domains is electrical, and that one of the materials is silicon. From the technological viewpoint, many choices, issues and tradeoffs exist between the various packaging and integration techniques (above IC, SiP, heterogeneous integration, bulk ...).

Other meanings of "heterogeneous" relate to the design process involved before the object exists physically (specification, synthesis, simulation, verification). In fact, the term was first commonly used to describe systems based on (digital) hardware and software, which can be more generally defined as system description *using more than one level of abstraction* - both for hardware and signal description. This is essentially driven by the need to handle the massive complexity of SoC/SiP by simplifying assumptions, and which in turn drives many of the requirements for modeling, design and simulation techniques of non-digital hardware. Specific branches of heterogeneity can also be identified, concerning *more than one type of signal description* (continuous and discrete time and value, cycle and bit accuracy) and *more than one model of computation* (dataflow, sequential processes, discrete event ...).

The concept of abstraction levels is one that must be addressed for heterogeneous SoC/SiP. Valid abstraction is difficult to achieve when tightly coupled physical phenomena are present in the system – this is the case even for digital electronics at nanometric technology nodes, and the rise in analogue, mixed-signal (AMS), RF and heterogeneous content to address future application requirements compounds this problem. Efficient ways must be found to incorporate non-digital objects into design flows in order to ultimately achieve AMS / RF / heterogeneous / digital hardware / software co-design.

The main objective of such an evolution is to reduce the design time in order to meet time to volume constraints. It is widely recognized that for complex systems at advanced technology nodes, a radical evolution in design tools and methods is required to reduce the "design productivity gap". Production capacity increases annually by around 50%, while design capacity increases annually by a rate of only 20-25%. The 2003 and 2005 ITRS Roadmaps both clearly state that "cost [of design] is the greatest threat to continuation of the semiconductor roadmap". Without the introduction of new design technology, design cost becomes prohibitive and leads to weak integration of high added value devices (such as RF circuits) for the various application sectors (automotive/transport, biomedical, telecommunications ...).

978-3-9810801-2-4/DATE07 © 2007 EDAA

A high-level vision of the maturity of these abstraction levels for various physical domains is given in Tab. 1, with examples of adequate modeling languages or simulation engines where solutions exist. To achieve design technology capable of fully exploiting the potential of heterogeneous SoC/SiP in a potentially holistic approach, high-level modeling techniques should be explored to cover more physical domains, and co-simulation/co-design tools should efficiently cover more abstraction levels. Hence, for EDA, it is clear that the impact of heterogeneity on design flows is or will be high, and necessary to facilitate heterogeneous device integration in SoC/SiP.

physical domain / level of abstraction	software	digital	analog	radiofrequency	mechanical	optical	fluidic	chemical
service	CORBA							
transaction	SystemC / UML		SystemC-AMS					
macro-architecture	SystemC		Ptolemy			SystemC		
micro-architecture	SystemC / VHDL							
block		VHDL	RF Simulation / VHDL-AMS		VHDL-AMS	VHDL-AMS	VHDL-AMS	
circuit		Electrical Simulation		RF Simulation			FEMLab	
physical		Finite Element Methods				Finite Difference		analytical

industrial solution laboratory solution no known solution

Tab. 1: Abstraction levels for various physical domains

The scope of this paper is not intended to hold the answers to these issues. Instead, it covers several aspects of heterogeneity in systems on chip and systems in package. Sections 2 and 3 cover physically heterogeneous systems, where electronics is integrated with components from a different physical domain. One is a relatively mature approach (MEMS for automotive applications), the other more prospective (microfluidics-based biochips). Sections 4 and 5 cover two ways in which heterogeneity is handled within EDA flows, concerning signal abstraction and integration heterogeneity in RF systems, and the validation of hardware/software IP components.

2 Design for MEMS integration

MEMS-based systems are increasingly used in our everyday life. One representative example is cars, where more than 40 MEMS sensors can typically be found today, including pressure, acceleration and angular rate sensors. Typical advantages are low-cost, low-power, small size, and sufficient reliability and performance.

The design of microsystems embedding MEMS cannot typically be considered to be a trivial task since it involves multiple engineering disciplines, multiple description levels and multiple design tools. Among the disciplines involved are MEMS-related sciences such as electrical, mechanical, thermal and fluidics engineering. Circuit design knowledge is also required, including all specific branches of the field (analog, digital, radio and power management). Finally, much of the performance in mid- to high-end microsystems comes from proper architectural design at the system level, which usually means control theory expertise as well as skills in mathematical, digital signal processing and software science.

Through the design details of a 15-bit SNDR, 50Hz BW closed-loop accelerometer [1], this section proposes to picture the actual steps usually required for this kind of "closely-coupled" heterogeneous microsystem. Fig. 1 depicts the architecture used in this design, which consists of a transducer, closely embedded in its readout electronics and the closed loop controller.

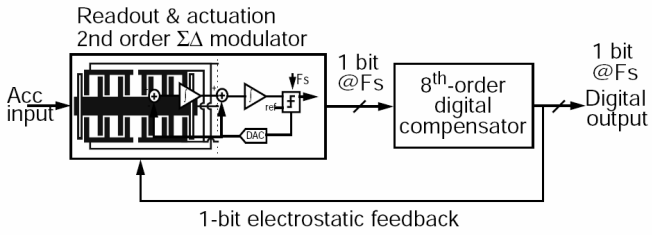

Fig. 1: Closed-loop accelerometer architecture

These three main components must be designed concurrently, not just for the performance of the system (noise, linearity, power), but also for its raw functionality (loop stability, controller robustness). However, no single tool enables the microsystem designer to achieve this.

Several commercial packages allow the finite-element simulation of the transducer to extract its main characteristics. Most of the time however, these are only used for auxiliary verifications, and the system designer relies on analytical models which allow for joint simulation using system-level (e.g. Matlab/Simulink) or lower-level (e.g. VHDL-AMS or Verilog-A) languages.

At the system level, one has to model all the various components with a reasonable level of non-idealities. In our case, the transducer is modeled with its non-linear effects as well as thermal noise. The readout ΣΔ modulator requires the modeling of effects such as switch and opamp noise, integrator leaks due to opamp finite gain, slew-rate, correlated double sampling effects, etc. Finally, the digital controller coefficients and finite-precision data and arithmetic can be taken into account. Modeling at this level allows for fast simulation of the whole system with a reasonable set of non-ideal effects and this simulation speed enables a certain degree of optimization through iterations. Simulation times can be reduced 100- to 1000-fold using this approach compared to transistor-level modeling. These figures are especially accurate in the case of stiff and in particular oversampled systems where the

clock frequency is much faster than the physical measurand (100-1000x) and where many mesurand cycles (10-100) need to be simulated to estimate performance.

In closed-loop systems an adaptive controller has to be inserted in the feedback path to ensure loop stability and performance. This task is usually done at the system level, for the same simulation speed reasons. However, while optimization tools do exist to tune the controller coefficients, there is usually a gap between control theory experts and systems developers / circuit designers. This gap typically needs to be filled using a proper re-formulation of the loop using standard control theory (which allows calculation of the coefficients) and a subsequent return to the electrical domain (to map the coefficients to the controller architecture hardware).

Once the system study is complete, actual circuit design can start. For efficient simulation purposes, the transducer is modeled using behavioral modeling languages compatible with transistor-level simulation, such as VHDL-AMS or Verilog-A. Analog circuits are simulated using Spice-like simulators, while digital parts are efficiently described with VHDL or Verilog netlists. Commercial packages allow for the joint simulation of these heterogeneously described subparts. However, no performance information usually comes out of these simulations. Only the basic functionality can reasonably be validated because of the long simulations times required by these multi-time-constant systems (low frequency transducer and measurand, oversampled A/D converter, higher frequency digital part for controller calculations).

This example design of a closed-loop accelerometer shows the diversity of the tools, technologies and disciplines involved. A coherent design framework to handle these interactions at various levels of abstraction is clearly lacking. Further, the question of knowledge management and IP re-use arises in order to ensure the preservation of know-how and reduce time-to-fab.

3 Integration of biofluidics with electronics

In this section, we discuss how the integration of biofluidics with electronics is enabling many exciting applications, while also illustrating many of the challenges of heterogeneity, both from a technological and EDA point of view. Microfluidics-based biochips (electronic chips for biofluidics), also referred to as lab-on-a-chip, are revolutionizing laboratory procedures involving molecular biology [2]. These devices automate highly repetitive laboratory tasks by replacing cumbersome equipment with miniaturized and integrated systems, and enable the handling of small amounts, e.g. micro- and nano-liters, of fluids.

3.1 Technology Platforms

Early biochips were based on the concept of a DNA micro-array, which is a piece of glass, plastic or silicon substrate on which segments of DNA (probes) have been affixed in a microscopic array. A drawback of DNA arrays is that they are neither reconfigurable nor scalable after manufacture. The basic idea of microfluidic biochips is to integrate assay operations, detection, as well as sample pre-treatment and preparation onto one chip using microfluidics technology. Continuous-flow microfluidic technologies are based on the manipulation of continuous liquid flow through microfabricated channels. Actuation of liquid flow is implemented either by external pressure sources, integrated mechanical micropumps, or by electrokinetic mechanisms [2]. Alternatives to the above closed-channel continuous-flow systems include novel open structures, where the liquid is divided into discrete, independently controllable droplets, and these droplets can be manipulated to move on a substrate [3].

The basic unit cell of an electrowetting-based "digital" microfluidic biochip consists of two parallel glass plates, as shown in Fig. 2(a). The bottom plate contains a patterned array of individually controllable electrodes, and the top plate is coated with a continuous ground electrode. A dielectric insulator, e.g., parylene C, coated with a hydrophobic film of Teflon AF, is added to the plates to decrease the wettability of the surface and to add capacitance between the droplet and the control electrode. The droplet containing biochemical samples and the filler medium, such as the silicone oil, are sandwiched between the plates; the droplets travel inside the filler medium. In order to move a droplet, a control voltage is applied to an electrode adjacent to the droplet, and at the same time, the electrode just under the droplet is deactivated. By varying the electrical potential along a linear array of electrodes, electrowetting can be used to move nanoliter volume liquid droplets along this line of electrodes [4]. The velocity of the droplet can be controlled by adjusting the control voltage (0~90 V), and droplets can be moved at speeds of over 20 cm/s. Droplets can also be transported, in user-defined patterns and under clocked-voltage control, over a two-dimensional array of electrodes shown in Fig. 2(b) without the need of micropumps and microvalves.

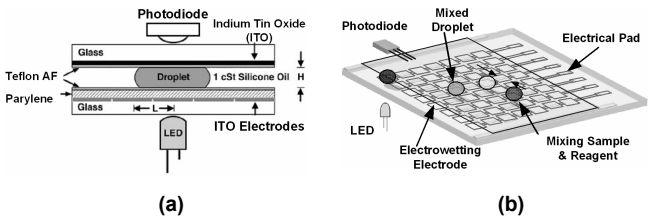

Fig. 2: Schematic of a digital microfluidic biochip used for colorimetric assays: (a) basic unit cell; (b) top view of microfluidic array

On-chip multiplexed assays for determining the concentrations of target analytes is a natural application for digital microfluidics [3] [5]. The on-chip processing steps include the following: 1) pre-diluted sample and reagent loading into on-chip reservoirs; 2) droplet dispensing of analyte solutions and reagents; 3) droplet transport; 4) mixing of analyte solution and reagent droplets; and 5) reaction product detection. On-chip detection of commercial-grade 2,4,6 trinitrotoluene (TNT) and pure 2,4-dinitrotoluene has been demonstrated [6]. In the realm of environmental monitoring, automated on-chip measurement of airborne particulate matter has been proposed using a scanning droplet method [7]. On-chip gene sequencing by synthesis is another emerging application [8]. Finally, protein crystallization is another emerging application area for microfluidic biochips [9].

3.2 Fabrication issues

While early fabricated digital microfluidics devices were based on a glass substrate and a top plate, recently, an open (i.e. without a top plate) structure has been designed on PCB (using a 3/3 mil linewidth/spacing) [10]. Electrodes (1.5 × 1.5 mm^2) are patterned on ¼ oz (~8.5 μm) copper with a final thickness of ~25 μm due to electroplating. A number of 150 μm via holes are drilled into each electrode to provide electrical contacts to the backside of the board. Ground electrode rails are patterned alongside all the drive electrodes to provide a continuous ground connection to the droplets, and a liquid photoimageable (LPI) soldermask (~17 μm) is patterned to act as an insulator, exposing only the ground rails. As the only post-processing step, Teflon AF is brush-coated to render the entire surface hydrophobic. A prototype device is shown in Fig. 3.

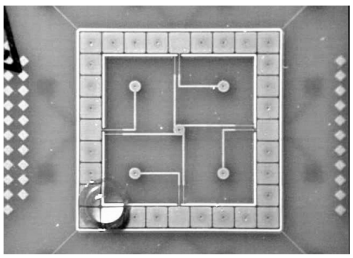

Fig. 3: A PCB implementation of "topless" coplanar microfluidics [10]

3.3 Design Automation

The "digital" nature of fluid flow in droplet-based microfluidics facilitates a system design methodology that is similar to electronic system design. *Architectural-level synthesis* [11] for microfluidic biochips can be viewed as the problem of scheduling assay functions and binding them to a given number of resources so as to maximize parallelism, thereby decreasing response time. On the other hand, *geometry-level synthesis* addresses the placement of resources and the routing of droplets to satisfy objectives such as area or throughput. Resource binding refers to the mapping from bioassay operations to available functional resources. Scheduling determines the start times and stop times of all assay operations, subject to the precedence and resource-sharing constraints.

A key problem in the geometry-level synthesis of biochips is the placement of microfluidic modules such as different types of mixers and storage units [12]. Since digital microfluidics-based biochips enable dynamic reconfiguration of the microfluidic array during run-time, they allow the placement of different modules on the same location during different time intervals. A synthesis methodology has recently been developed to unify operation scheduling, resource binding, and module placement [13]. Exact placement information, instead of a crude area estimate, is used to judge the quality of architectural-level synthesis. This method allows architectural design and physical design decisions to be made simultaneously. Finally, droplet routing techniques to determine droplet transportation paths have also been developed recently [14].

4 RF SoC/SiP flows

In this section, we examine the requirements for EDA to address the challenges of next generation wireless systems, which must be targeted to multi-standard and reconfigurability. Market demand has grown explosively over the last years for more functionality, higher performance, smaller size and lower cost. The choice of integration strategy is key to concurrently achieving these demands, and therefore it is necessary to evaluate the feasibility of any proposed integration approach. Several integration strategies exist for these systems, but SiP (rather than SoC or discrete) is emerging as a strong contender as the solution to facilitate these market applications. It allows a high degree of flexibility in the package architecture, particularly for RF applications.

RF SiP is an enabling packaging platform for wireless communication, which allows the integration of digital ICs, logic IC and RFICs plus passive components, SAW filters and mechanical parts (Fig. 4). The package is no longer just a connectivity interposer between an IC and a board and has become the system integration vehicle. What is still impractical to put into a single SoC can now be integrated into a single package.

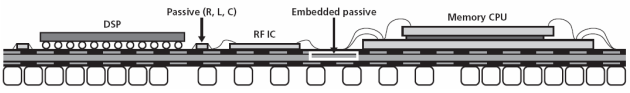

Fig. 4: SiP example

Aggressive time-to-market demands, complex cost structures and tradeoffs in system architecture add further complexity, in that this new highly complicated package must be designed concurrently with the system going into the package. However, SiP design will span a complex design chain of system, SoC, circuit, package and board designers. This also means it requires expert engineering talent in widely divergent fields. These domains have traditionally been designed, simulated, implemented and verified separately due to the different mindsets of the engineers and the underlying tools, methods and flows.

For SiP, a new "co-design" approach capable of supporting complete modeling and analysis of the system interconnect across the IC, its package, and the PCB is required. A single "mega-environment" cannot be the answer; it needs to integrate with IC design and physical package implementation. For any tool to be used effectively, it must be a natural part of the environment that a particular engineer is using. At integration (a key aspect of a SiP design enabling solution), special attention needs to be paid to *who* will be running the top level simulations and performing top level physical design, and where design collateral (netlists, models, databases, etc.) is coming from.

Recent product developments built around proven methodologies now enable a complete front-to-back SiP design and implementation (Fig. 5). This approach allows companies to adopt what were once expert engineering SiP design capabilities for mainstream product development.

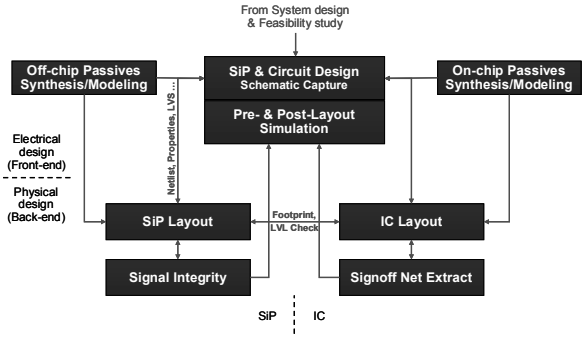

Fig. 5: Front-to-back SiP design and implementation flow

A similar approach now needs to be taken to include MEMS and non-electronic devices into an overall co-design methodology, where "design" is a convergent process, starting with an abstract model at specification and maturing into manufacturing. Co-design brings the model of the entire design domains (IC, package, board, MEMS, LCD displays …) into a common design environment that allows global optimization and characterization of the design in development, starting at the top-level electrical simulation of the entire system or sub-system. This represents a design process that manages the physical, electrical and manufacturing interfaces between design components across all of the associated design domains.

The development of the EDA flow and dedicated product extensions towards a fully integrated solution needs to go hand-in-hand with a methodology shift within the industry and enable a co-design mentality across business units and companies.

5 Verification of heterogeneous IP

This section deals with a point tool that addresses one of the long-standing issues in functional verification: the measurement of the quality of the functional verification process of IP blocks. Since the building blocks for heterogeneous systems are becoming increasingly disparate, digital SoC can provide some lessons regarding the challenges of integrating heterogeneous IP.

The design of today's industrial SoC is, in most cases, evolutionary and "middle out". Current practices within the organization, combined with individual and collective engineering experience, provide a context for system specification and design. A high quality of components in a system allows for assumptions that are more aggressive and therefore a simplification of the design process.

Producing functionally correct silicon is a primary concern for SoC designers, but is becoming increasingly difficult and expensive to achieve. It is alarming to see functional verification consuming an increasing percentage of project resources. Functional verification is a process and like any process, we must be able to measure it accurately in order to improve it effectively.

Since current management practices rely on measurable objectives, it is essential that management has objective measurements for quality throughout the design process. The alternative is to isolate design teams from quality concerns by measuring productivity and constraining available resources. Quality is a holistic notion and heterogeneous systems demand a more holistic approach. A key challenge is then the measurement of the quality of the functional verification of each IP block.

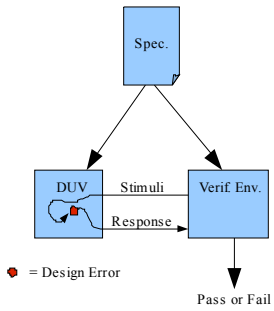

Fig. 6: Typical verification strategy

Fig. 6 shows a DUV (Design Under Verification) and a VE (Verification Environment), both derived from a

specification. The VE generates stimuli to exercise the specified behavior and checks the DUV's responses are compliant with the specification. The objective of Functional Verification is to find design errors (represented by the circle in the DUV of Fig. 6).

Existing tools are able to measure the ability of the verification environment to generate sufficient stimuli, often through the use of *Code Coverage* and *Functional Coverage*. However, these techniques do not measure the ability of the VE to check for correct responses - for example, if there are missing checkers or errors in checkers, this is not reflected in coverage metrics.

Mutation analysis [15] is a technology that can measure both the ability of the verification environment to generate sufficient stimuli and the ability of the VE to check for correct responses. Functional errors or "mutations" are injected into the design under verification. If the user's existing VE is not able to detect these errors then this highlights weaknesses in the verification. This technique is applicable to both hardware and software verification. Mutation analysis can be used to (i) improve, and (ii) measure the quality of Functional Verification.

Mutation analysis has a long history of research in the software community [16]. Recent patented innovations now allow this technology to be applied to today's leading industrial projects. In particular, Certitude [17] is an EDA tool that takes as input:

1. an existing functional verification non-regression test suite developed using any language or tool, and
2. the DUV developed in VHDL and/or Verilog.

The output of the tool is a display of the original DUV's source code. This source code is color highlighted to indicate which areas of the code are not verified, these are the verification weaknesses. The verification weaknesses are classified, by the tool, either as missing stimuli or missing/erroneous response checking. The tool also provides a metric, which measures the percentage of mutations which can be detected or "killed" by the current verification environment.

6 Conclusion

This paper has given a short overview of some of the issues at stake in heterogeneous systems on chip and systems in package. Four cases were considered: design for MEMS integration, microfluidic biochips, RF SiP and functional IP co-verification. While the establishment of a single environment to handle complex heterogeneous systems seems unrealistic, it is clear that an overall co-design methodology is required at all phases of design, from abstract high-level specifications to mature physical manufacturing. Ways forward include the extension of the concept of synthesizable IP to heterogeneous content, advances in hierarchical modeling techniques, and clear identification of point-tool collaboration and roles within a complete flow managed by transversal platforms.

References

[1] C. Condemine et. al., "A 0.8mA 50Hz 15b SNDR ΔΣ Closed-Loop 10g Accelerometer Using an 8th-order Digital Compensator", *Proc. ISSCC*, pp. 248-249, 2005.

[2] E. Verpoorte and N. F. De Rooij, "Microfluidics meets MEMS", *Proc. IEEE*, vol. 91, pp. 930-953, 2003.

[3] V. Srinivasan, V. K. Pamula, and R. B. Fair, "An integrated digital microfluidic lab-on-a-chip for clinical diagnostics on human physiological fluids", *Lab on a Chip*, pp. 310-315, 2004.

[4] M. G. Pollack, A. D. Shenderov and R. B. Fair, "Electrowetting-based actuation of droplets for integrated microfluidics", *Lab on a Chip*, vol. 2, pp. 96-101, 2002.

[5] T. H. Schulte, R. L. Bardell and B. H. Weigl "Microfluidic technologies in clinical diagnostics", *Clinica Chimica Acta*, vol. 321, pp. 1-10, 2002.

[6] K. Pamula et al., "A droplet-based lab-on-a-chip for colorimetric detection of nitroaromatic explosives", *Proc. IEEE MEMS Conf.*, 2005.

[7] Y. Zhao and S.K. Cho, "Microparticle sampling by electrowetting-actuated droplet sweeping", *Lab on a Chip*, vol. 6, 137-144, 2006.

[8] C.A. Emrich, "Microfabricated 384-lane capillary array electrophoresis bioanalyzer for ultrahigh-throughput genetic analysis", *Anal. Chem*, 74, pp. 5076-5083, 2002.

[9] B. Zheng, L.S. Roach, R.F. Ismagilov, "Screening of protein crystallization conditions on a microfluidic chip using nanoliter-size droplets", *Journal of American Chemical Society,* 125, 11170-11171, 2003.

[10] P. Y. Paik, V. K. Pamula, M. G. Pollack and K. Chakrabarty, "Coplanar digital microfluidics using standard printed circuit board processes", *Proc. Int. Conf. Miniaturized Chemical and Biochemical Analysis Systems (MicroTAS)*, 2005.

[11] F. Su and K. Chakrabarty, "Architectural-level synthesis of digital microfluidics-based biochips", *Proc. IEEE Int. Conf. CAD*, pp. 223-228, 2004.

[12] F. Su and K. Chakrabarty, "Module placement for fault-tolerant microfluidics-based biochips", *ACM Transactions on Design Automation of Electronic Systems*, vol. 11, pp. 682-710, 2006.

[13] F. Su and K. Chakrabarty, "Unified high-level synthesis and module placement for defect-tolerant microfluidic biochips", *Proc. IEEE/ACM Design Automation Conference*, pp. 825-830, 2005.

[14] F. Su, W. Hwang and K. Chakrabarty, "Droplet routing in the synthesis of digital microfluidic biochips", *Proc. Design Automation and Test in Europe*, pp. 323-328, 2006.

[15] R.A. DeMillo, R.J. Lipton and F.J. Sayward. "Hints on Test Data Selection: Help for the Practicing Programmer", *IEEE Computer*, 11(4): p. 3443, 1978.

[16] Jeff Offutt and Roland H. Untch, "Mutation 2000: Uniting the Orthogonal", *J. Software Testing, Verification, and Reliability*, pp. 45-55, San Jose, CA, October 2000.

[17] Certitude is a registered trademark of Certess Inc. www.certess.com