A Cyber-Physical Systems Approach to Personalized Medicine: Challenges and Opportunities for NoC-based Multicore Platforms

Paul Bogdan
Department of Electrical Engineering, University of Southern California, 3740 McClintock Avenue, Los Angeles, CA 90089-2560
pbogdan@usc.edu

ABSTRACT
This paper describes a few fundamental challenges concerning the design of Network-on-Chip (NoC) based multicores as the backbone of cyber-physical systems (CPS) for personalized medicine. One fundamental challenge in designing such CPS architectures is the need for a unifying mathematical description of the dynamical interactions between bio-physiological processes and cyber states. Another fundamental challenge is to build a rigorous mathematical optimization framework that allows the CPS to adapt to varying workloads and demands. To enable large-scale parallelism, we need a rigorous understanding of the CPS workloads that can guide the design and optimization of wired and wireless NoCs. We advocate for the development of *goal-oriented self-organization* algorithms that seek to both optimize specific design cost functions and maximize information about future system state. It is necessary to identify basic local rules of interaction not only for solving large scale optimization problems in a distributed fashion, but also for inducing an overall degree of autonomy and intelligence in the CPS architecture.

Keywords
Cyber-physical systems; networks-on-chip; multicore platforms; highly-variable workloads; scalability; adaptive autonomous systems; goal-oriented self-organization; personalized medicine; non-stationary fractal behavior; real-time performance guarantees.

1. INTRODUCTION
Stressful and sedentary lifestyles coupled with irregular sleep patterns and inappropriate diets contribute to a higher prevalence of chronic diseases (e.g., diabetes, heart disease, cancer) and increased healthcare costs [34]. Recent progress in genomics, proteomics, metabolic phenotyping and physiological sensing revealed that numerous biomarkers representing collections of malfunctioning molecular regulatory pathways can serve for early detection of abnormal pathophysiology [40]. A better understanding of the complex landscape of diseases can lead to more effective, early and reliable medical diagnosis and therapy. To reach this goal, we require holistic integration and accurate mining of biological processes (from DNA and protein levels to pathways, cells, tissues and organs).

To address these emerging clinical and healthcare challenges, we advocate for a cross-disciplinary approach to cyber-physical systems (CPS) design [3][4][21][22], aiming at seamlessly and safely integrate sensing, computation, communication, control and actuation for developing new technology for personalized and precise medicine. This envisioned CPS is conceptualized in Figure 1, which shows targeted product and tissue cells measured from a patient. Based on signal processing and data mining techniques, the sensed processes are communicated to computational platforms. Of note, the measured dynamical physical processes represent most of the inputs into a cyber-physical task graph we define in Figure 2. The role of the computing platform is to model and analyze the dynamics of various biomarkers, as well as the interactions between biological entities and therapeutic agents in order to assess the efficacy and success of various medical treatments. In addition, the outcome of the computational analysis will help us find control strategies for drug administration that eliminate infections or suppress diseases. This optimization process requires constant interactions with physicians and continuous monitoring of biophysical processes.

The heterogeneity and variability of data- and computation-intensive CPS tasks require massively parallel architectures [26] for zetta-scale computing. Networks-on-Chip (NoC) offer a promising approach for dealing with the resource-intensive CPS tasks (processing, communication storage). However, they also require built-in *intelligence* for adapting to CPS workloads [6] and meeting their performance constraints. Even though microfluidic technology [25] provides a rich CPS sensing modality, the stringent CPS timing constraints require mathematically founded strategies for identifying the loosely and tightly coupling parallelism and its exploitation in NoC design and optimization. Developing *intelligent run-time adaptation* strategies for NoCs requires not only fixed-cost (power/ throughput) optimization, but also system prediction capabilities. To accomplish this goal, we envision the need for designing *goal-oriented self-optimization algorithms*. While performing a set of tasks, the role of these algorithms is both to minimize design cost metrics and to monitor the platform resources in order to gain information about its future state. This calls for an understanding of how simple (distributed) rules contribute to optimizing the design cost and maximizing information about the system state.

Starting from these premises, this paper is organized as follows: Section 2 outlines several advances in CPS sensing for genomic, proteomic, metabolic and physiological processes. Section 3 describes a unified description of interactions between sensed physical processes and cyber-components. Section 4 highlights a few research problems related to NoC-based multicore design and outlines several unsolved issues. Section 5 concludes this paper.

2. CPS APPROACH TO PERSONALIZED MEDICINE: PREMISES & CHALLENGES

2.1 Advances in genomic sensing
Two sensing technologies have proven highly successful in attempts at analyzing the structure and dynamics of genomic processes: 1) Single-molecule level probing has enabled the detection of the structure and dynamics (transient events) of genomic processes within living cells [13]. For instance, single-molecule experiments allow us to monitor in real-time base-pair-by-base-pair the RNA polymerase progress and the replisome activity in living cells [11]. A major feature of in-vivo single-molecule approaches is that they provide dynamic information regarding genomic processes (e.g., transcription, translation, replication) taking place in living cells. Also, single-molecule

The author acknowledges the support by US National Science Foundation under Grant 1331610 and University of Southern California.

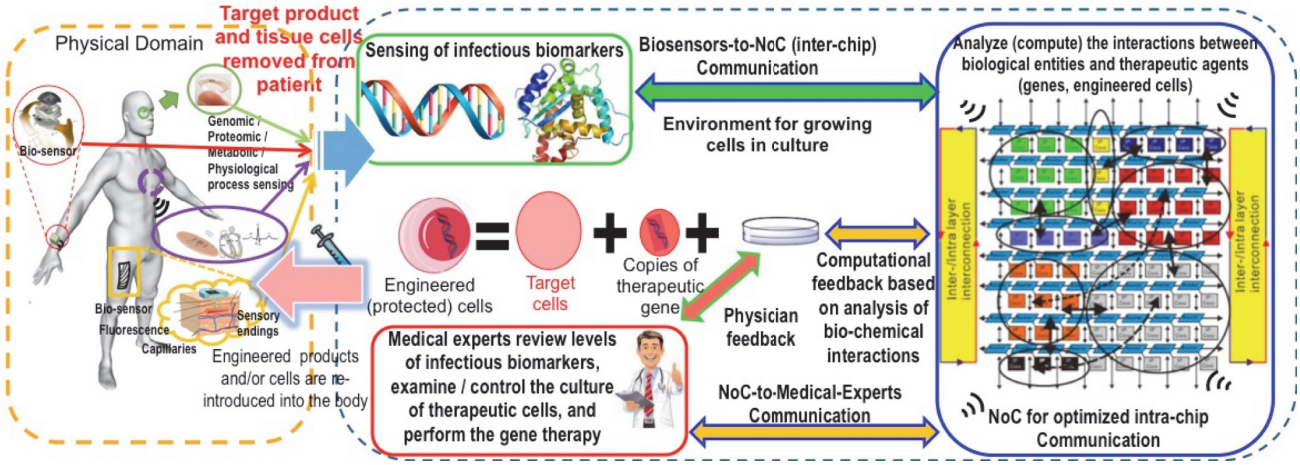

Figure 1. Envisioned multicore-based CPS entails the following strategy: Continuous monitoring of DNA, RNA, protein, metabolic biomarkers and physiological processes help at assessing the health state of a patient. This information together with specifics and gravity of the disease / infection are transmitted to a computational platform. The optimized NoC-based multicore performs an in-depth analysis of the bio-chemical interactions between native (unprotected) cells, engineered products and (protected) cells and the disease markers or virulence of the virions. Depending on the aggressiveness of the disease / spread of the virions, medical experts can select the best therapeutic strategy using the optimized computational platform and grow ex-vivo only the one solution that minimizes the presence of disease markers or virion spread with minimal toxicity. The last step of the envisioned CPS-enabled therapy consists of infusion of the amount of engineered products and cells and the process is repeated until no disease or virion biomarkers are detected.

fluorescence allows measurement of a large number of genome sequences in parallel. It enables single-molecule DNA sequencing and the evaluation of genomic characteristics (e.g., transcription dynamics and efficiency). 2) The dimethyl sulphate sequencing (DMS-seq) approach can be applied to human-derived samples to analyze the structure-function relationships for both informational and functional RNAs present in a cell. More precisely, the DMS-seq approach permits the analysis of long noncoding RNAs, the relationship between mRNA structure and microRNA/RNA interference targeting, and the functional identification and analysis of ribozymes, ribo-switches and thermal sensors [37]. In spite of several technological breakthroughs in real-time monitoring of genomic processes [11][12], we also need algorithms for genomic data processing and CPS platforms enabling mining of collected biological data [44]. Numerous algorithms (e.g., LASSO regression, Bayesian networks) have been proposed, but none of these provide either accurate information about the gene regulatory networks or near real-time information about the genomic patterns of diseases.

2.2 Advances in proteomic sensing

Proteomics complements genomics by providing a comprehensive perspective of biological processes at protein level and can be used for disease diagnosis [41]. In recent years, mass spectroscopy proteomics technologies have become a routine part of biological investigations and have enabled the mapping of cellular pathways and networks [23]. Surface plasmon resonance (SPR) technique requires 1pg of proteins bound per $1mm^2$ for detection purposes and measures the variation in the refractive index occurring on a metal surface, which in turn represents protein binding/interaction events. The single molecule fluorescence resonance energy transfer (FRET)[18] exploits the energy transfer between a pair of donor and acceptor and can provide detailed conformational information about protein-protein complexes. More recently, a real-time single-molecule fluorescence imaging within a coimmunoprecipitation approach was developed to study not only the structural properties of protein-protein interaction networks, but also their transient (dynamic) characteristics [23]. This real-time single-molecule processing has provided new insights into signaling kinetics of proteins in patients with cancer [24]. Besides inference of protein-protein networks, we need efficient computing platforms for protein folding analysis that facilitate early disease detection.

2.3 Advances in metabolomics sensing

Metabolic sensing and phenotyping refers to a comprehensive analysis of biological fluids (e.g., plasma, urine) or tissue samples. Metabolic sensing provides a complementary perspective to genomics and proteomics. Monitoring the dynamics of various metabotypes reveals macroscopic information about the continuous interactions among genes and factors such as the patient's lifestyle, gut microbiome and drug-efficiency or side-effects [32]. Since static snapshots of metabolic processes are currently extensively used in the medical diagnosis, many early signs of diseases remain unnoticed because time-based correlations are overlooked. Nuclear-magnetic resonance (NMR) and mass spectroscopy can provide time-varying information about various metabotypes, but CPS platforms capable of real-time mining (e.g., quantifying the degree of insulin resistance in diabetes, detecting the concentrations of cancer biomarkers) are missing. Thus, there is a need for CPS platforms that can integrate and efficiently process (e.g., principal component analysis, partial least squares discriminant analysis (PLSDA), hierarchical clustering (HC), self-organizing maps) measured data for mining metabolic processes and providing accurate information about early disease signs.

2.4 Advances in physiological sensing

Physiological processes encompass macroscopic signatures of the dynamic interactions among human organs and represent the basis for current medical diagnosis and therapy. Thus, numerous sensing technologies for physiological processes have been developed [19][33]. For instance, the possibility of continuously monitoring the blood glucose [38] has enabled the development of artificial pancreas devices [9][14]. Despite recent advances in sensing physiological processes, efficient CPS architectures capable of mining and determining the control strategies based on the characteristics [7] of physiological processes are still missing.

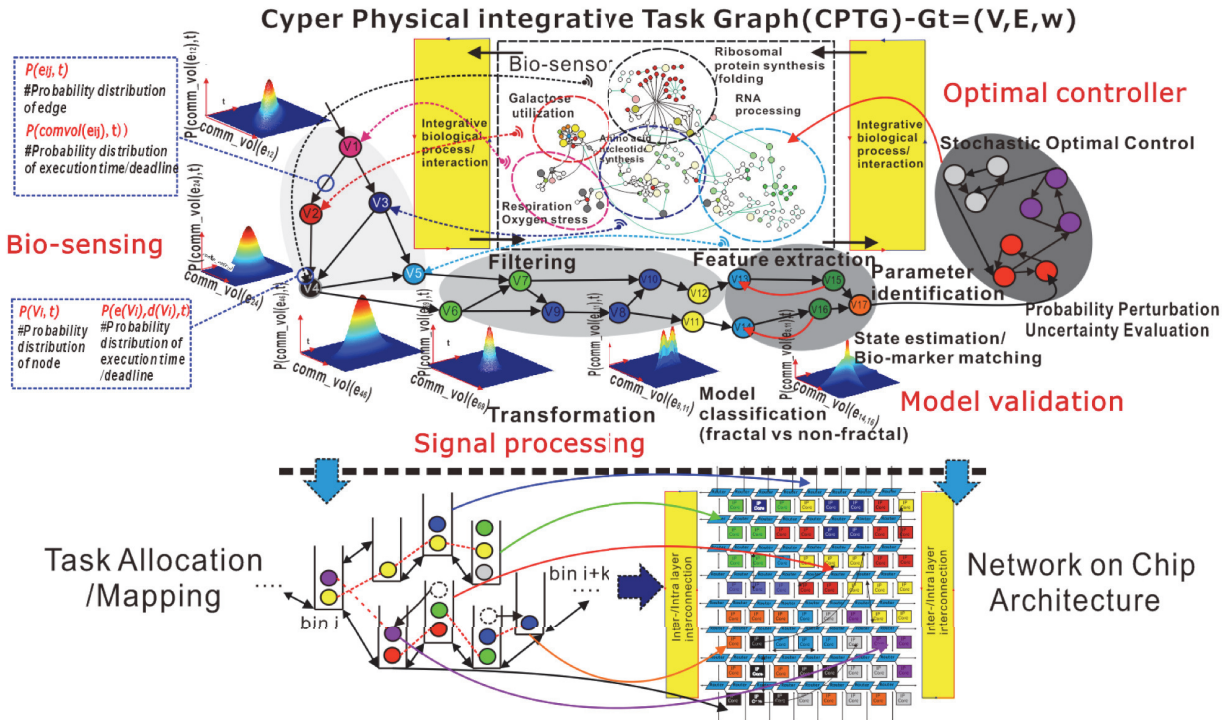

Figure 2. Unifying mathematical primitives and abstractions to describe the interactions and dependencies between the cyber and physical components of a system in a dynamic environment. Unlike traditional system-level approaches developed for embedded system design, we define the time-dependent cyber-physical task graph capable to describe the dynamic interactions between the system (cyber), physical processes and the environment. Both computation and communication events cease to be time independent; instead, they are both coupled to the interaction with the sensors measuring the physical processes and their dynamics is constrained and dictated by the actual sensed values. This calls for rethinking the theoretical foundations of CPS design and for proposing new mathematical strategies and algorithms for CPTG partitioning. Clustering and mapping on the multicore platforms, for time dependent and correlating computation and communication scheduling, self-organizing network traffic routing as a function of the emerging network load, hardware solutions for run-time reconfiguration of the multicore topology (wired vs. wireless, reactive vs. proactive routers), buffering and memory allocation.

2.5 Additional CPS Challenges

Ignoring the characteristics of sensed physical processes and working with data-oblivious formatting strategies for the sake of simplicity leads to inefficient solutions. This approach contributes not only to exacerbated workloads, but also to restricted choices in the optimization of multicores. The characteristics of sensed physical processes should influence not only the way data is formatted and stored, but also the way it is modeled and analyzed [4][6]. So far, most data mining and machine learning strategies raise computational complexity issues although they advance our understanding. Most of them ignore the time dependence and memory properties of physical processes since it is assumed that these lead to even greater complexity levels. Truly breakthrough solutions in this field will come from understanding and exploiting correlation structures [14].

Transferring massive amounts of formatted data through wide area networks is costly in terms of both performance and energy consumption. An improved understanding of the modalities and characteristics of sensed physical processes can help us exploit the heterogeneous dimensionalities of data by designing communication and storage strategies in a distributed fashion on networked multicores [29]. These strategies will improve accessibility and reduce the processing speed. The distributed organization and management of the sensed data also has significant implications for tasks like mining and prediction.

3. MATHEMATICAL FRAMEWORK FOR CPS DESIGN AND OPTIMIZATION

As shown in Figure 1, we propose to adopt a radically new approach to CPS design not driven by the reductionism and decoupling of sensing, computation and communication, but guided by the mathematical characteristics (e.g., time dependence, significant variability, heterogeneity) of physical components. Building on the fact that genomic, proteomic, metabolic and physiological signals are highly dynamic, we define the ***cyber-physical task graph*** (CPTG) (see Figure 2) as a time-dependent representation where nodes denote the set of physical processes (sensed processes or processed signals) and monitored cyber states, while the edges represent the interactions among physical and cyber entities (communicated events or processing updates). Of note, we do not see the cyber and physical components as distinct. Instead, we embed the dependencies between the computation (state processing), sensing (state observation) and communication (state update) into the CPTGs. In addition, due to intrinsic biological complexity, multi-scale spatio-temporal dynamics and significant heterogeneity / variability, the CPTGs possess stochastic features at both computation and communication levels, which in turn affect the optimization and control actions of the entire CPS. More precisely, we do not assume that computations are imprecise as in stochastic computing; instead the computation reflects the stochastic nature and evolution of bio-chemical kinetics that is becoming the de-facto method for analyzing and controlling biological processes.

The benefit of introducing the CPTG is that it allows us to build a *unifying mathematical framework* that supports composition by bridging the computational and physical aspects of time and space. It also copes with stochastic events and physical components' uncertainty.

Definition 1: A *cyber-physical task graph* (CPTG) is a directed time-dependent multi-valued hypergraph $G_t=(V,E,w)$, where each node $v_i \in V$ denotes a sensed physical process, a processed signal or a monitored cyber-variable and each directed edge $e_{ij} \in E$ connects one node to one or more nodes in graph G_t. Each of the edges can be either single-valued (meaning that is tagged with a single random variable) or multi-valued (meaning that is tagged with multiple random variables) encoding the nature of interactions between sets of pairs of nodes in G_t. The nodes and edges in the CPTG (Figure 2) represent the basic computations and communications (interdependencies) between bio-sensing, signal processing (proces filtering, process transformation, feature extraction), model classication and identification (elucidating when a process exhibits a fractal or non-fractal dynamics, estimating model parameters, performing a state estimation for physical processes, validating the model by testing its prediction capibilities), and control (optimizing the control variables to maintain the genomic, proteomic and physiological processes within a fractal domain). Thus, the CPTG nodes and edges are augmented with specific CPS design variables:

- $P_{ex}: V \times R \rightarrow [0,1]$ and $P_d: V \times R \rightarrow [0,1]$ are time-dependent probabilities associated with the execution time and deadline of a task $v_i \in V$;
- $P: V \times V \times E \times R^m \rightarrow [0,1]$ represents a multidomensional probability distribution describing the characteristics (e.g., communication volume, bandwidth requirements, communication delay, timing constraints) of the directed multi-valued interactions between sets of CPTG nodes.

Building on these grounds (see Figure 2), we advocate for a science of CPS, which conceptualizes the interdependencies among the engineered system and the environment and rethinks the mathematical methodologies, architectural models and algorithms for system design and optimization. The CPTG allows encoding within node-state equations both discrete and continuous time dependence of cyber and physical components and account for uncertainty via conditional multi-dimensional probabilities represented along the CPTG edges. From this perspective, numerous CPS problems such as CPTG partitioning, clustering and mapping onto the multicore platform, computation distribution and communication scheduling, network traffic flow routing and regulation, network topology synthesis and reconfiguration, communication buffering and memory allocation require new mathematical strategies and algorithms. Consequently, we advocate for a decentralization (democratization) of the CPS design and optimization so that such problems are solved in a distributed fashion and only sparse hierarchical feedback is allowed to meet high performance requirements with minimum overhead.

4. RESEARCH PROBLEMS & SOLUTIONS FOR MULTICORE-BASED CPS

In order to enable discoveries in the medical CPS domain and provide contextual information to support effective data analytics and mining, we will need to develop new solutions for improving the design and optimization of NoC architectures. In what follows, we review three key research problems, briefly outline current progress and highlight the need for developing a rigorous mathematical approach that allows goal-oriented self-organization and optimization of multicore-based CPS.

4.1 Goal-oriented self-organization inspired solutions for CPTG-based resource allocation

Resource allocation problems such as task partitioning/ clustering, mapping and scheduling have received significant attention from the research community [17][28] under both static and dynamic conditions. Most of these solutions assume that application task graphs are either deterministic and time-independent or slowly varying, with process variables characterized by Gaussian distributions. However, rich unconstrained sensing of medical CPS not only provides a large amount of biological data, but also highlights the importance of mathematical characteristics of physical processes (high variability, heterogeneity, non-Gaussianity, fractality and long range memory dynamics, higher order spatio-temporal correlations [5]). These biophysical process characteristics influence not only the model identification, state estimation and control algorithms [6], but also the computational features and requirements of the CPTG. Unlike classical embedded systems, time in medical CPS is crucial for performance, accuracy, dependability and reliability of the overall system. Traditional mathematical approaches for reducing such problems to classical quadratic assignment or integer linear programming formulations are not scalable in the context of a large number of variables with pronounced time dependent behavior and whose dynamics and interactions are characterized by a complex uncertainty structure.

The characteristics of biophysical processes imply not only challenges for CPTG description, but also opportunities for theoretical system design and optimization. For instance, one popular approach to dynamic resource allocation is represented by the *d-choice balls-into-bins* formulation [2][20][35]; it aims to solve the randomized allocation problem of m balls into n bins where each ball is allowed to choose the least loaded bin out of d bins independently and uniformly at random. In the CPS context, it is important to consider that the process in which a ball chooses a bin has some correlation structure. For instance, a ball is constrained to hop from one bin to another only by following the paths provided by the NoC architecture and any further migration incurs a cost. In other words, letting correlated balls (tasks in the CPTG in Figure 2) choose random bins out of d independent bins may lead to inefficient situations where the balls reside on distant cores on the NoC architecture. Along these lines, the authors in [8][10] studied maximum load and cover time for their proposed *local search allocation* for m balls choosing among n bins where bins are placed at the vertices of a graph and each ball moves only along the edges of the graph. Such research endeavors provide analytical tools for investigating distributed dynamic resource allocation rather than centralized allocation. The drawbacks lie in the assumption that balls are independently and uniformly generated at random across the network. As shown by the dotted lines linking two or more balls in Figure 2, each ball has intrinsic correlations / dependencies with other balls, which further imposes constraints on the allocation process. This implies that the birth of balls should be seen as a correlated process. Although the execution / computation time for individual tasks is minimized by allowing such balls to choose independently and uniformly at random from the NoC resources, this strategy can also increase the communication delay and exacerbate traffic requirements. Concomitantly, while the above situation may suggest that some first order linked balls should be mapped and scheduled to be

processed on nearby NoC resources, long range dependencies across the CPTG can require that these balls spread more evenly across the direction dependency sub-graph of the CPTG (i.e., when designing the local interactions and migration rules for individual balls, one needs to take into consideration not only the immediate nodes that require particular computation results, but also other CPTG nodes whose execution may be affected). Another drawback in CPTG mapping and scheduling is related to the NoC architecture: *i*) Given that computational tasks are highly dependent on sensed data, which can present high variability, how can we predict the task arrival and execution times on a particular NoC resource? *ii*) This uncertainty in task execution time also raises the question of how to efficiently determine the execution order of tasks *iii*) Given that tasks are highly heterogeneous running on highly variable data, it is harder to predict the task execution time and estimate the availability of NoC resources. To address this lack of complete system information and tackle the scalability issue, Xue et al. [42][43] explored the energy landscape-based algorithms for protein folding analysis and proposed a task mapping strategy inspired from *self-organization in natural systems* where balls move from one NoC node to another and check the availability of the processing elements (PEs) along a source-destination path. If an available PE is discovered, then the computational task gets processed at this particular PE and the result is communicated back to a master PE. This leads not only to a significant improvement in the utilization of the NoC architecture, but also considerably reduces the execution time (up to 200X improvement in computation time when compared to randomized mapping for an NoC with 256 cores). In addition, this self-organizing randomized strategy enables the NoC architecture to offer a higher performance with minimal impact on energy consumption. Of note, the master node was not required to solve a complete mapping and scheduling problem. We let the intrinsic characteristics of the computations dictate the availability. In turn, the mapping follows a self-optimization process. The only constraint imposed for the distributed computation was to migrate within a predefined region of the NoC architecture to minimize communication delay.

4.2 Goal-oriented self-organization inspired solutions for CPTG-based traffic routing

Since time is essential in CPS, we need to figure out how to route efficiently (communicating fast with minimum energy consumption) significant amount of data and processed results within the NoC architecture. Although reviewing the vast literature on NoC routing protocols is beyond the scope of this paper, it is important to note that there have been numerous adaptive routing algorithms proposed to tackle the traffic congestion [1][15][36][39][43]. For instance, Qian et al. [36] proposed a dynamic distributed routing algorithm inspired from statistical physics that learns a fitness metric in order to evaluate the neighboring router availability. To capture the traffic dynamics, the fitness metric of each candidate output channels depends on both the channels occupancy and the speed of the packets traversing that channel. Alternatively, a scalable self-organized task-aware routing algorithm is presented in [43], which balances the computational and communication loads by employing a randomized search for available NoC resources. Although some degree of self-organization can be achieved with such routing protocols, there remain numerous open problems: *i*) What are the best localized cost metrics (e.g., fitness, traffic pressure) that can adequately capture the traffic dynamics and enable a *goal-oriented self-organization of multicore traffic*? Of note, while current adaptive routing schemes improve NoC performance over deterministic routing, they fail to enforce and guarantee that particular traffic flows/corridors meet their timing requirements. *ii*) How should the traffic dynamics metrics be defined when considering various networking alternatives such as reconfigurable wired and wireless links or expressing virtual channels? *iii*) Given that traffic dynamics can dictate over time the appropriateness of choosing among different cost metrics, how can the CPS platform optimally select the right metric to route traffic and for how long should the rules be enforced? *iv*) Given that establishing global information about the network state is prohibitive, what is the best decomposition of the network into traffic decision islands such that the aggregated distributed/localized routing schemes meet the predefined timing requirements of the CPS?

4.3 Goal-oriented self-organization solutions for NoC infrastructure optimization

For multicore architectures, the communication infrastructure plays a very important role. Numerous solutions ranging from fixed topology customization to dynamic topology optimization have been proposed [16][26][30][31][43]. For instance, Hollis et al. [16] defined a set of micro-rules for inducing adaptiveness and demonstrated how these rules combined with the skip-link concept contribute to the self-optimization of NoC performance and energy consumption. Matsutani et al. [30][31] improved the performance of 3D NoCs by combining a traffic-aware approach with a randomized topology customization that exploits both wired and inductively coupling wireless capabilities. Aiming to describe the interactions between biological entities resulting either from a disease/infection progression or a drug effect on human genomic, proteomic and physiological processes, Majumder et al. [27] developed a simplified CPTG description of the multi-scale spatio-temporal stochastic interactions among bio-chemical reactions, identified the computation and communications workloads, and proposed an optimization algorithm for establishing wireless NoC connections among PEs. Although significant improvements in terms of both performance and energy consumption are achieved through the proposed solutions, there remain numerous unsolved problems: *i*) How to find the optimal number of wireless connections and where to place them such that the traffic flows meet their performance bounds (delay and throughput bounds)? *ii*) How to dynamically turn on / off connections to sustain the goal-oriented self-organization of the NoC infrastructure?

5. CONCLUSIONS

This paper has presented a few fundamental challenges faced by the design of NoC-based multicores as the backbones of CPS architectures. To cope with such challenges, the paper advocates for developing mathematical approaches and distributed algorithms for inducing a degree of autonomy, self-awareness and self-adaptation concerning not only specific design metrics, but also the prediction of the CPS future states.

ACKNOWLEDGMENT

We thank to Dr. A. Stauffer, Dr. H. Ghasemzadeh, Dr. P. Pande, Dr. R. Marculescu and Y. Xue for valuable feedback.

REFERENCES

[1] G. Ascia, V. Catania, M. Palesi, D. Patti, “Implementation and analysis of a new selection strategy for adaptive routing in networks-on-chip," *IEEE Trans. on Computers*, vol.57, no.6, pp.809-820,2008.

[2] Y. Azar, A.Z. Broder, A.R. Karlin, and E. Upfal, “Balanced allocations,” *SIAM J. Comput*., vol. 29(1), pp.180–200, 1999.

[3] R. Baheti and H. Gill, “Cyber-physical systems,” From: *The Impact of Control Technology*, T. Samad and A.M. Annaswamy (eds.) 2011.

[4] P. Bogdan and R. Marculescu, “Towards a Science of Cyber-Physical Systems Design,” *Proc. of the 2011 IEEE/ACM 2nd Intl. Conf. on Cyber-Physical Systems* (ICCPS), Chicago, USA, 2011.

[5] P. Bogdan and R. Marculescu, “Non-stationary traffic analysis and its implications on multicore platform design,” *IEEE Trans. on Computer-Aided Design*, vol. 30, issue 4, pp. 508-519, April 2011.

[6] P. Bogdan and R. Marculescu, “ Cyberphysical systems: workload modeling and design optimization,” *IEEE Design and Test of Computers*, vol. 28, no. 4, pp. 78-87, July/Aug. 2011.

[7] P. Bogdan, S. Jain, K. Goyal, and R. Marculescu, “Implantable pacemakers control and optimization via fractional calculus approaches: a cyber-physical systems perspective,” *Proc. of IEEE/ACM 3rd Intl. Conf. on Cyber-Physical Systems* (ICCPS), 2012.

[8] P. Bogdan, T. Sauerwald, A. Stauffer, H. Sun, “Balls into bins via local search,” *Proc. of the 24th Ann. ACM-SIAM Symposium on Discrete Algorithms* (SODA), 2013.

[9] M. Breton et al., “Fully integrated artificial pancreas in type 1 diabetes: modular closed-loop glucose control maintains near normoglycemia,” *Diabetes*, vol. 61, no. 9, pp.2230-2237, 2012.

[10] K. Bringmann, T. Sauerwald, A. Stauffer, H. Sun, “Balls into bins via local search: cover time and maximum load,” *Proc. of the ACM-SIAM Symp. on Theoretical Aspects of Computer Science*, 2014.

[11] J. Chen et al., “High-throughput platform for real-time monitoring of biological processes by multicolor single-molecule fluorescence,” *Proc. of National Academy of Sciences of the USA.*, vol. 111, 2014.

[12] X. Darzacq, Y. Shav-Tal, V. de Turris, Y. Brody, S.M. Shenoy, R.D. Phair, R.H. Singer, “In vivo dynamics of RNA polymerase II transcription,” *Nature Structural & Molecular Biology*, 14, 2007.

[13] D. Dulin, J. Lipfert, M.C. Moolman, and N.H. Dekker, “Studying genomic processes at the single-molecule level: introducing the tools and applications,” *Nature Reviews Genetics*, vol. 1, pp. 9-22, 2013.

[14] M. Ghorbani and P. Bogdan, “A cyber-physical system approach to artificial pancreas design,” *Proc. of 9th Intl. Conf. on Hardware/Software Codesign and System Synthesis* (CODES+ISSS), 2013.

[15] P. Gratz, B. Grot, S.W. Keckler, “Regional congestion awareness for load balance in networks-on-chip,” *IEEE 14th Intl. Symp. on High Performance Computer Architecture* (HPCA), pp.203-214, 2008.

[16] S.J. Hollis, C. Jackson, P. Bogdan, and R. Marculescu, “Exploiting emergence in on-chip interconnects," *IEEE Trans. on Computers*, vol.63, no.3, pp.570-582, March 2014.

[17] J. Hu and R. Marculescu, “Energy- and performance-aware mapping for regular NoC architectures,” *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, vol.24, no.4, April 2005.

[18] C. Joo, H. Balci, Y. Ishitsuka, C. Buranachai, and T. Ha, “Advances in single-molecule fluorescence methods for molecular biology,” *Annu. Rev. Biochem*. 77, pp. 51–76, 2008.

[19] E. Kaniusas, *Biomedical Signals and Sensors*, Springer, 2012.

[20] R.M. Karp, M. Luby, and F.M. auf der Heide, “Efficient PRAM simulation on a distributed memory machine.,” *Algorithmica*, 1996.

[21] E.A. Lee, “Cyber physical systems: design challenges,” *IEEE Intl. Symp. on Object Oriented Real-Time Distributed Computing* (ISORC), pp.363,369, 5-7 May 2008.

[22] I. Lee et al. “Challenges and research directions in medical cyber–physical systems," *Proc. of the IEEE* , vol.100, no.1, Jan. 2012.

[23] H.W. Lee et al., “Real-time single-molecule co-immunoprecipitation of weak protein-protein interactions,” *Nature Protocols*, vol. 8, 2013.

[24] H.W. Lee et al., “Real-time single-molecule co-immunoprecipitation analyses reveal cancer-specific ras signalling dynamics,” *Nature Communications*, vol. 4, article number 1505, 2013.

[25] Y. Luo, K. Chakrabarty, T.Y. Ho, “A cyberphysical synthesis approach for error recovery in digital microfluidic biochips,” *Design, Automation& Test in Europe Conf. & Exhibition* (DATE), 2012.

[26] T. Majumder, S. Sarkar, P.P. Pande, A. Kalyanaraman, “NoC-based hardware accelerator for breakpoint phylogeny," *IEEE Trans. on Computers*, vol.61, no.6, pp.857-869, June 2012

[27] T. Majumder, X. Li , P. Bogdan, P. Pande, “NoC-enabled multicore architecture for stochastic analysis of biomolecular reactions,” *Design, Automation & Test in Europe Conf. & Exhibition*, 2015.

[28] R. Marculescu, U.Y. Ogras, L.-S. Peh, N.E. Jerger, Y. Hoskote, “Outstanding research problems in NoC design: system, microarchitecture, and circuit perspectives,” *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 2009.

[29] R. Marculescu and P. Bogdan, “The chip is the network: toward a science of network-on-chip design,” *Foundations and Trends in Electronic Design Automation*, vol. 2, no. 4, pp. 371-461, 2009.

[30] H. Matsutani et al. “A case for wireless 3D NoCs for CMPs," *Proc. of the 18th Asia and South Pacific Design Automation Conference* (ASP-DAC), pp.23-28, 22-25 Jan. 2013.

[31] H. Matsutani et al., “Low-latency wireless 3D NoCs via randomized shortcut chips," *Proc. of the Design, Automation and Test in Europe Conference and Exhibition* (DATE), pp.1-6, 24-28 March 2014.

[32] J.K. Nicholson et al., “Metabolic phenotyping in clinical and surgical environments,” *Nature*, vol. 491, pp. 384–392, 15 November 2012.

[33] R.B. Northrop, *Noninvasive Instrumentation and Measurement in Medical Diagnosis*, CRC Press, Medical, 2010.

[34] Partnership to fight chronic diseases, “The growing crisis of chronic disease in the United States,” Available online at: http://www.fightchronicdisease.org/sites/fightchronicdisease.org/files/docs/GrowingCrisisofChronicDiseaseintheUSfactsheet_81009.pdf

[35] Y. Peres, K. Talwar, and U. Wieder, “The (1 + β)-choice process and weighted balls-into-bins,” *ACM-SIAM Symp. on Disc. Alg.*, 2010.

[36] Z. Qian, P. Bogdan, G. Wei, C.-Y. Tsui, R. Marculescu, “A traffic-aware adaptive routing algorithm on a highly reconfigurable network-on-chip architecture,” *IEEE/ACM intl. conf. on Hardware/software codesign and system synthesis* (CODES+ISSS), 2012.

[37] S. Rouskin, M. Zubradt, S. Washietl, M. Kellis and J.S. Weissman, “Genome-wide probing of RNA structure reveals active unfolding of mRNA structures in vivo,” *Nature*, 15 December 2013.

[38] G. Sparacino, A. Facchinetti, C. Cobelli, “Smart continuous glucose monitoring sensors: on- line signal processing issues,” *Sensors*,2010.

[39] M. Tang, X. Lin, and M. Palesi, “Routing pressure: a channel-related and traffic-aware metric of routing algorithm,” *IEEE Trans. on Parallel and Distributed Systems*, vol.PP, no.99, pp.1,1, 2013.

[40] R.L. Winslow, N. Trayanova, D. Geman and M.I. Miller, “Computational medicine: translating models to clinical care,” *Sci. Transl. Med.*, Vol. 4, Issue 158, p. 158rv11, 31 October 2012.

[41] J.D. Wulfkuhle, L.A. Liotta, E.F. Petricoin, “Proteomic applications for the early detection of cancer,” *Nature Reviews Cancer*, vol. 3, issue 4, pp. 267-275, 2003.

[42] Y. Xue, Z. Qian, P. Bogdan, F. Ye, and C.-Y. Tsui “Disease Diagnosis-on-a-Chip: Large Scale Networks-on-Chip based Multicore Platform for Protein Folding Analysis,” *Proc. of the Design Automation Conference* (DAC), June 2014.

[43] Y. Xue et al. “An Efficient Network-on-Chip (NoC) based Multicore Platform for Hierarchical Parallel Genetic Algorithms,” *Proc. of the IEEE/ ACM Intl. Symp. on Networks on Chip* (NoCS), 2014.

[44] M. Yandell and D. Ence, “A beginner's guide to eukaryotic genome annotation,” *Nature Reviews Genetics*, vol. 13, pp. 329-342, 2012.