



Guest Editorial: JSPS Special Issue on 2018 IEEE Signal Processing Systems (SiPS) Workshop

Tokunbo Ogunfunmi¹ · John McAllister² · Bevan Baas³ · Mrityunjoy Chakraborty⁴

Received: 11 July 2020 / Revised: 13 July 2020 / Accepted: 14 July 2020 / Published online: 29 July 2020
© Springer Science+Business Media, LLC, part of Springer Nature 2020, corrected publication 2020

This Special Issue of the *Journal of Signal Processing Systems* focuses on papers selected and invited from all the papers accepted and presented at the **2018 IEEE SiPS Workshop**.

The special issue was preceded by an IEEE International Workshop on Signal Processing Systems (SiPS) held at Cape Town, South Africa, October 21–24, 2018

The IEEE Workshop on Signal Processing Systems (SiPS) is a conference dedicated to providing a premier international forum for researchers and practitioners from industry and academia to exchange the latest scientific and technical advances in the area of design and implementation of signal processing systems. It addresses current and future challenges and new directions in research and development of these systems.

SiPS 2018 workshop kicked off with seven tutorials in cutting-edge areas. Apart from regular technical sessions, the workshop had three interesting special sessions, namely, “Computing for Radio Astronomy”, “Signal Processing for Signal/Video and IoT Energy”, and “Recent Advances in Architectures for Machine Learning”. The workshop also organized a panel discussion on “The MeerKAT Radio Telescope”.

SiPS 2018 had a total of 80 regular paper submissions in addition to the 13 special session papers for a total of 93 papers. All the papers were assigned to three or more reviewers. A total of 61 papers were accepted – 43 were selected for lecture presentations and 18 were used to build the poster sessions.

At the end of the workshop, a committee selected 12 papers and invited the authors to expand on their work in SiPS 2018 papers and submit journal papers to our 2018 SiPS JSPS Special Issue.

The articles have undergone rigorous peer-review according to the journal’s high standards.

Out of the 12 papers selected, 7 papers made it through the rigorous peer review process.

These contributions encompass a wide range of research topics thereby appealing to both the experts in the field and those who want a snapshot of the current breadth of research results in current and future challenges and new directions in research and development of signal processing systems.

Collectively, these 7 papers illustrate the diverse range of issues that were presented at the workshop. This Special Issue presents these 7 papers. Here are the details:

The first article, by Ning Lyu, Bin Dai, Hongfei Wang and Zhiyuan Yan titled “Optimization and Hardware Implementation of Learning Assisted Min-Sum Decoders for Polar Codes” proposes a novel scaling offset min-sum (SOMS) algorithm and adapts the offset min-sum (OMS) algorithm for polar codes, and both algorithms are improved via learning. For all message updates, conventional min-sum decoding algorithms use the same scaling factor or offset, which is usually obtained by numerical simulations.

By modeling the data flow of min-sum algorithms as a deep neural network, the parameters used in the message passing updates of min-sum decoders can be different for each message update, and are obtained by training and optimizing the corresponding deep neural network.

The simulation results show that the proposed SOMS algorithm based on deep learning performs better than all existing belief-propagation (BP)-based algorithms. The

✉ Tokunbo Ogunfunmi
togunfunmi@scu.edu

John McAllister
jp.mcallister@qub.ac.uk

Bevan Baas
bbaas@ucdavis.edu

Mrityunjoy Chakraborty
mrityun@ece.iitkgp.ac.in

¹ Santa Clara University, Santa Clara, CA 95053, USA

² Queens University of Belfast, Belfast, N., Ireland BT7 1NN, UK

³ University of California, Davis, CA 95616, USA

⁴ Indian Institute of Technology, Kharagpur, W.B. 721,302, India

authors also present an efficient hardware architecture of the proposed SOMS algorithm. The proposed architecture of the SOMS algorithm for a (256,128) polar code is implemented and validated on the Xilinx Artix-7 field-programmable gate array.

The second article, by Joonas Iisakki Multanen, Heikki Kultala, Kati Tervo and Pekka Jääskeläinen on “Energy-Efficient Low Latency Multi-Issue Cores for Intelligent Always-On IoT Applications” proposed three cores targeting mixed control flow and data processing applications in internet-of-things always-on devices featuring an exposed datapath architecture with high performance, while retaining energy-efficiency. These features are achieved with exploitation of instruction-level parallelism, fast branching and the use of an instruction register file. The designs are evaluated targeting maximum clock frequency for high throughput tasks and energy-delay product for energy and delay critical tasks and compared against two RISC cores, LatticeMico32 and zero-riscy.

The third article, by Narges Mohammadi Sarband, Oscar Gustafsson and Mario Garrido on “Using Transposition to Efficiently Solve Constant Matrix-Vector Multiplication and Sum of Product Problems” presents an approach to alleviate the potential benefit of adder graph algorithms by solving the transposed form of the problem and then transposing the solution. A systematic way to obtain the transposed realization with a minimum number of cascaded adders subject to the input realization is described. In this way, wide and low constant matrix multiplication problems, with sum of products as a special case, which are normally exceptionally time consuming to solve using adder graph algorithms, can be solved by first transposing the matrix and then transposing the solution. Examples show that while the relation between the adder depth of the solution to the transposed problem and the original problem is not straightforward, there are many cases where the reduction in adder cost will more than compensate for the potential increase in adder depth and result in implementations with reduced power consumption compared to using sub-expression sharing algorithms, which can both solve the original problem directly in reasonable time and guarantee a minimum adder depth.

The fourth article, by Anatoly Prihozhy, Simone Casale-Brunet, Endri Bezati and Marco Mattavelli titled “Pipeline Synthesis and Optimization from Branched Feedback Dataflow Programs” develops an accurate algorithm and introduces fast dynamic and mixed static / dynamic heuristics that minimizes the number of pipeline stages for a given pipeline-stage time-period, and also minimizes the overall pipeline registers size by means of appropriate assignment of feedbacks and instructions to pipeline stages. The authors also propose a genetic algorithm for tuning the heuristics for a particular design. The experimental results show the algorithms give quick solutions that are very close to accurate

solutions and overcome the earlier developed algorithms regarding computing time and pipeline parameters.

The fifth article, by Jian Zhou, Antonia Papandreou-Suppappola and Chaitali Chakrabarti.

titled “Parallel Gibbs Sampler for Wavelet-based Bayesian Compressive Sensing with High Reconstruction Accuracy” proposes a two-stage parallel coefficient update scheme for wavelet-based Bayesian compressive sensing (BCS) which helps address ill-posed signal recovery problems where the first stage approximates the real distributions of the wavelet coefficients and the second stage computes the final estimate of the coefficients. While in the first stage, the parallel computing units share information with each other, in the second stage, the parallel units work independently. Even when the computing units share information, when the number of computing units is large, the process deviates from the sequential Gibbs sampler resulting in large reconstruction error. The authors propose two new coefficient re-computation schemes to reduce the reconstruction error at the cost of longer computation time and a new coefficient update scheme that updates coefficients in both stages based on data generated a few rounds ago. Finally, the authors design the corresponding parallel architecture and synthesize it in 7 nm technology node. For the system with 8 computing units, the proposed algorithm reduces the execution time up to 6.8x at maximum compared to the sequential implementation.

The sixth article, by Pavel Arnaudov and Tokunbo Ogunfunmi titled “Dynamically Adaptive Fast Motion Estimation Algorithm for HD Video” presents an adaptive Fast Motion Estimation (FME) algorithm, which reduces the number of search points, reduces computational complexity and therefore lowering power consumption, while providing improved quality per watt in FME. The pervasiveness of mobile devices with HD video capabilities demands such low power hardware accelerators. The proposed algorithm can identify the best search pattern within a given region of the HD video frame, based on its motion dynamics to achieve lower power video encoding. The goal is to achieve the best quality with the minimum number of search iterations. The reduced number of checks translates into power savings. The results show that the proposed algorithm can reduce computations by about 4 times compared to other fixed search patterns algorithms, while staying within 1 dB of PSNR results. That equates to about 75% power savings at the expense of not more than 1 dB of PSNR quality loss.

The seventh and final article, by Yaesop Lee, Yanzhou Liu, Karol Desnos, Lee Barford and Shuvra Bhattacharyya titled “Passive-Active Flowgraphs for Efficient Modeling and Design of Signal Processing Systems” develops a flowgraph for efficient modelling and design of signal processing systems by formulating the vertices and edges into concepts called “active blocks” and “passive blocks”, respectively in the graph representation. Computation in the dataflow graph is

represented as “active blocks”, while the concept of dataflow buffers is represented as “passive blocks”. Like dataflow edges, passive blocks are used to store data during the intervals between its production and consumption by actors. However, passive blocks can have multiple inputs and multiple outputs, and can incorporate operations on and rearrangements of the stored data subject to certain constraints. The authors define a form of flowgraph representation that is based on replacing dataflow edges with the proposed concept of passive blocks and present a structured design methodology for utilizing this new form of signal processing flowgraph, and demonstrate its application to improving memory management efficiency and execution time performance.

We thank all the authors, the reviewers, the JSPS journal administrative staff and the JSPS Editor-in-Chief for all their contributions to making the high quality of this JSPS Special Issue possible.

We hope you enjoy reading the articles.

Guest Editors:

[1] Tokunbo Ogunfunmi, Santa Clara University, Santa Clara, CA 95053, USA.

Email: TOgunfunmi@scu.edu (corresponding author for receiving proofs).

[2] John McAllister, Queens University of Belfast, Belfast, N. Ireland, BT7 1NN, UK.

Email: jp.mcallister@qub.ac.uk

[3] Bevan Baas, University of California, Davis, CA 95616, USA.

Email: bbaas@ucdavis.edu

[4] Mrityunjy Chakraborty, Indian Institute of Technology, Kharagpur, W.B., 721,302, India.

Email: mrityun@ece.iitkgp.ac.in

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.