Journal of Real-Time Image Processing (2018) 14:513–515
https://doi.org/10.1007/s11554-018-0763-2

GUEST EDITORIAL

Special issue on heterogeneous real-time image processing

Dietmar Fey[1] · Frank Hannig[2]

Published online: 14 March 2018

1 Introduction

Mobile devices, such as smartphones and tablets, offer a plethora of media-rich applications such as photograph and video recording and editing, natural user interfaces, and computer vision. Other areas of embedded image systems are characterized by close-to-sensor processing, such as advanced driver assistance systems (ADAS), mobile scanners, and smart devices used in medical and industrial imaging. Such applications demand highest computing capabilities at stringent resource and power budgets as well as hard real-time constraints.

Future scaling of computing performance mandates dramatically improving the energy efficiency of image systems. One rapidly rising trend is to use heterogeneous MPSoCs (multiprocessor system-on-chip), consisting of multiple processors, maybe of a different type, as well as accelerators such as digital signal processors (DSPs), embedded graphics processing units (GPUs), field-programmable gate arrays (FPGAs), or dedicated hardware. Another trend is to use new 3D integrated circuit technologies that allow for tighter integration of processor cores, memory, and sensors to reduce communication latency and improve bandwidth, leading to lower energy consumption, see, for example, the work by Dudek et al. [3].

This calls for novel methodologies for designing heterogeneous hardware architectures, and shielding software developers from growing complexity and allowing them to concentrate on algorithm development rather than on low-level implementation details. Thus, in the last years, several model-based design methods [4, 12, 17] and approaches based on domain-specific programming languages for programming heterogeneous image systems and corresponding compilers have been proposed [2]—prominent examples include Halide [13], HIPAcc [10, 14], and Darkroom [8]. These approaches mainly aim at improving productivity as well as optimizing utilization and performance, but hardly consider real-time aspects. Yet, this special issue covers heterogeneous processor architectures such as mentioned above with an emphasis on real-time image processing.

✉ Frank Hannig
frank.hannig@fau.de

Dietmar Fey
dietmar.fey@fau.de

[1] Department of Computer Science – Computer Architecture, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Martensstr. 3, 91058 Erlangen, Germany

[2] Department of Computer Science – Hardware/Software Co-Design, Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Cauerstr. 11, 91058 Erlangen, Germany

2 Heterogeneous real-time image processing

The special issue with the title "Heterogeneous Real-Time Image Processing" in Springer's Journal of Real-Time Image Processing was motivated by the *DATE Friday Workshop on Heterogeneous Architectures and Design Methods for Embedded Image Systems (HIS)* [6] in Grenoble, France, 2015. The special issue at hand embraces the entire technology stack of heterogeneous real-time imaging systems, namely close-to-sensor processing using 3D chip stacking and image processing in the analog domain, dedicated hardware/software architectures, synthesis and mapping methods for image processing applications targeting heterogeneous embedded architectures, as well as real-time applications implemented on CPUs, GPUs, and FPGAs. In reaction to an open call for papers, several submissions were received. After a careful peer-review process, seven manuscripts were accepted for inclusion in this special issue. It is our pleasure to introduce these articles in the following briefly.

The first article with the title "A flexible mixed-signal image processing pipeline using 3D chip stacks" by Shi et al. [16] proposes three-dimensional chip stacking suitable for smart camera applications. Here, an entire image processing pipeline consisting of a CMOS image sensor,

analog preprocessing, analog/digital conversion, and digital processing can be combined in one chip stack.

The second article entitled "Bio-inspired heterogeneous architecture for real-time pedestrian detection applications" by Maggiani et al. [9] presents a custom-tailored hardware/software architecture for real-time pedestrian detection as can be used as ADAS. More specifically, the authors tightly couple a computationally intensive histogram of oriented gradient (HOG) pipeline implemented in an FPGA with a bio-inspired spatiotemporal filter running on a CPU.

The next two articles propose systematic analysis, synthesis, and mapping methods for heterogeneous system-on-chip (SoC) architectures.

The article with the title "IPAS: A design framework for analysis, synthesis and optimization of image processing applications for heterogenous computing architectures" by Hartmann et al. [7] presents a framework for UML/SysML-based modeling, analysis, and synthesis of image processing applications. The proposed approach can target FPGAs by using a generic IP block library as well as embedded CPUs (ARM processors). The framework is evaluated for several edge detection algorithms and circle detection based on the Hough transform.

The article "A novel global methodology to analyze the embeddability of real-time image processing algorithms" by Saussard et al. [15] evaluates several heterogenous SoCs for ADAS and proposes three different mapping methods for applications with real-time constraints. Here, parallel applications are modeled in the form of a processing graph and performance models allow for pruning the vast search space of possible mappings to such parallel architectures.

This special issue is rounded off with three image processing applications parallelized and implemented on GPUs.

Garcia-Garcia et al. [5] present in their article "Interactive 3D object recognition pipeline on mobile GPGPU computing platforms using low-cost RGB-D sensors" a 3D object recognition application based on a Kinect sensor in combination with an NVIDIA Tegra K1 SoC. Here, several sub-algorithms, such as the point cloud creation, normal estimation, and bilateral filtering, are parallelized in CUDA for the embedded GPU.

The article with the title "GLSC: LSC superpixels at over 130 FPS" by Ban et al. [1] proposes a very fast superpixel segmentation method, which has many applications in computer vision. More specifically, the so-called *linear spectral clustering (LSC)* method is parallelized using CUDA and evaluated on several NVIDIA high-end GPUs.

Mújica-Vargas et al. [11] provide the last article entitled "An efficient nonlinear approach for removing fixed-value impulse noise from grayscale images." The authors propose an efficient nonlinear filter to suppress high-density fixed-value impulse noise—also known as salt-and-pepper noise—in large-size grayscale images. Two parallel implementations of the filter are presented and evaluated, one for shared-memory systems using OpenMP and the other one for GPUs using CUDA.

3 Conclusions

Finally, we thank the Editors-in-Chief, Nasser Kehtarnavaz and Matthias F. Carlsohn, of Springer's Journal of Real-Time Image Processing as well as the administrative staff for their valuable support throughout the preparation and publication of this special issue. Furthermore, we thank all authors for their contributions to this special issue and their excellent efforts. We are very grateful also to the reviewers for their careful work and valuable suggestions that helped to improve the quality of the articles. We hope you will enjoy reading this special issue.

References

1. Ban, Z., Liu, J., Fouriaux, J.: GLSC: LSC superpixels at over 130 FPS. J. Real Time Image Process. (2016). (ISSN: 1861-8219)
2. Bell, S., Pu, J., Hegarty, J., Horowitz, M.: Compiling algorithms for heterogeneous systems. Synth. Lect. Comput. Arch. **13**(1), 1–105 (2018). https://doi.org/10.2200/S00816ED1V01Y201711CAC043
3. Dudek, P., Lopich, A., Gruev, V.: A pixel-parallel cellular processor array in a stacked three-layer 3D silicon-on-insulator technology. In: Proceedings of the European Conference on Circuit Theory and Design (ECCTD), pp. 193–196. (2009). https://doi.org/10.1109/ECCTD.2009.5274946
4. Duran, A., Ayguadé, E., Badia, R.M., Labarta, J., Martinell, L., Martorell, X., Planas, J.: OmpSs: A proposal for programming heterogeneous multi-core architectures. Parall. Process. Lett. **21**(02), 173–193 (2011). https://doi.org/10.1142/S0129626411000151
5. Garcia-Garcia, A., Orts-Escolano, S., Garcia-Rodriguez, J., Cazorla, M.: Interactive 3D object recognition pipeline on mobile GPGPU computing platforms using low-cost RGBD sensors. J. Real Time Image Process (2016). (**ISSN: 1861–8219**).
6. Hannig, F., Fey, D., Lokhmotov, A.: Proceedings of the DATE Friday Workshop on Heterogeneous Architectures and Design Methods for Embedded Image Systems (HIS 2015). In: CoRR (2015). arXiv: 1502.07241
7. Hartmann, C., Häublein, K., Reichenbach, M., Fey, D.: IPAS: A design framework for analysis, synthesis and optimization of image processing applications for heterogenous computing architectures. J. Real Time Image Process (2016). https://doi.org/10.1007/s11554-016-0587-x
8. Hegarty, J., Brunhaver, J., DeVito, Z., Ragan-Kelley, J., Cohen, N., Bell, S., Vasilyev, A., Horowitz, M., Hanrahan, P.: Darkroom: Compiling High-level Image Processing Code into Hardware Pipelines. ACM Trans. Graphics (TOG) 33.4 (2014),

144:1–144:11.https://doi.org/10.1145/2601097.2601174. ISSN: 0730-0301

9. Maggiani, L., Bourrasset, C., Quinton, J.C., Berry, F., Sérot, J.: Bio-inspired heterogeneous architecture for real-time pedestrian detection applications. J. Real Time Image Process (2016). https://doi.org/10.1007/s11554-016-0581-3. ISSN: 1861–8219
10. Membarth, R., Reiche, O., Hannig, F., Teich, J., Körner, M., Eckert, W.: HIPAcc: HIPAcc: A domain- specific language and compiler for image processing. IEEE Trans. Parallel Distributed Syst. **27**(1), 210–224 (2016). ISSN: 1045-9219
11. Mújica-Vargas, D., de Jesús Rubio, J., Kinani, J.M.V., Gallegos-Funes, F.J.: An efficient nonlinear approach for removing fixed-value impulse noise from grayscale images. J. Real Time Image Process (2017). https://doi.org/10.1007/s11554-017-0746-8. ISSN: 1861-8219
12. Ngo, T.D., Sepulveda, D., Martin, K. J. M., Diguet, J.-P.: Communication-model based embedded mapping of dataflow actors on heterogeneous MPSoC. In: Proceedings of the Conference on Design and Architectures for Signal and Image Processing (DASIP), pp. 1–8 (2014). https://doi.org/10.1109/DASIP.2014.7115629
13. Ragan-Kelley, J., Barnes, C., Adams, A., Paris, S., Durand, F., Amarasinghe, S.: Halide: A Language and Compiler for Optimizing Parallelism, Locality, and Recomputation in Image Processing Pipelines. In: ACM SIGPLAN Notices 48.6 (2013), pp. 519–530. https://doi.org/10.1145/2499370.2462176. ISSN: 0362-1340
14. Reiche, Oliver., Akif Özkan, M., Membarth, Richard., Teich, Jürgen., Hannig, Frank.: Generating FPGA-based Image Processing Accelerators with Hipacc. In: Proceedings of the International Conference on Computer Aided Design (ICCAD). Irvine, CA, USA: IEEE, Nov. 2017, pp. 1026–1033 https://doi.org/10.1109/ICCAD.2017.8203894. ISBN: 978- 1-5386-3094-5
15. Saussard, R., Bouzid, B., Vasiliu, M., Reynaud, R.: A novel global methodology to analyze the embeddability of real-time image processing algorithms. In: J. Real Time Image Process. (2017). https://doi.org/10.1007/s11554-017-0686-3. ISSN: 1861-8219
16. Shi, L., Soell, C., Pfundt, B., Baenisch, A., Reichenbach, M., Seiler, J., Ussmueller, T., Weigel, R.: A flexible mixed-signal image processing pipeline using 3D chip stacks. J. Real Time Image Process. (2016). https://doi.org/10.1007/s11554-016-0628-5. ISSN: 1861-8219
17. Singh, A.K., Shafique, M., Kumar, A., Henkel, J.: Mapping on multi/many-core systems: Survey of current and emerging trends. In: Proceedings of the 50th Annual Design Automation Conference (DAC). Austin, TX, USA: ACM, pp. 1:1–1:10 (2013). https://doi.org/10.1145/2463209.2488734