# Ing. Martin Musil:
# Ghost-Free HDR Video Using FPGA
## a review by Jan Schmidt, CTU in Prague

The reviewed work proposes a new algorithm to deghost high dynamic range images acquired by multiple exposures. The proposed algorithm is explained in Section 4.1 in sufficient detail to be reproduced, and evaluated using available metrics in Section 4.3. The evaluation results are discussed and illustrated to hint at image processing detail. The complexity of the algorithm is not evaluated in an abstract way (e.g., MAC per pixel), rather, implementations of the algorithm are compared to implementations of other algorithms by their running time as thoroughly as possible. Although such comparisons tend to be less relevant, I believe that the proposed method is indeed faster while providing better or comparable results than other state-of-the art methods.

When an algorithm for a specific environment is proposed, that environment is usually characterized in an abstract way, so that the results can be applied by others. In this work, I have difficulty to derive what the target environment is:

- the title says *Using FPGAs;*
- the claim in Chapter 4 says *for an embedded hardware device,* the next paragraph is more specific with *based on FPGA technology with FullHD CMOS sensor onboard.*

The reader does not know, what are the important limitations to be considered in the construction of the proposed algorithm. In that chapter, a comparison is promised with *...and even state-of-the-art methods, that are too computationally demanding and even not feasible to implement and/or accelerate on FPGA.* Yet, what constitutes the infeasibility, is not analyzed.

Sections 3.1 and 3.2 inform about main classes of embedded platforms. Unfortunately, the sections are merely descriptive, not analytic, and do not arrive at any conclusion. Properties important just for engineering and even production are mentioned, as if the entire work was only an engineering exercise. Descriptions in such a limited scope are bound to be rather shallow. Besides some inaccuracies, they also get outdated quickly (Xilinx 7 Series SoC versus Ultrascale SoC devices). The sections contains almost no references.

In further sections of Chapter 3, existing solutions to problems in HDR acquisition and deghosting are reviewed. Again, the sections are not analytic, although it would be easy to find design patterns common in embedded platforms and important for the ongoing work, such as precomputed tabular functions.

I assume that the implementation in Section 4.2 serves two purposes:

- It is a case study, documenting the feasibility of the algorithms and hence at least one example confirming the claim in Chapter 4, and
- it is a vehicle for algorithm testing and comparison.

When considering it as a case study, the target parameters are not given, at least not in a succinct way. In the text, FullHD is mentioned, that is, 1920 × 1280 pixels. When studying this part of the work, the referenced paper by Nosko et al. is a required reading. It states 3 exposures per image and the target speed of 30fps, and internal 10.8 bit format. The amount of resources available for the deghosting hardware could be *comupted* form that paper and from the 7020 Zync data sheet by a kind reader. The target clock speed can also be calculated from those data.

The implementation is stated to be *fully pipelined*, and to consist of *two components, Certainty map creation and HDR merging*. I cannot estimate whether the two certainty maps (Figure 4.3) are computed in parallel, how do we obtain the scaling factor at the bottom of Figure 4.3, etc. No architecture has been given. It may be the case that the data flow in Figure 4.3 has been projected into hardware in a straightforward manner, but otherwise more than one operation in the flow could have been mapped to a single unit (scaling). If the measure of 96.45 FPS in Table 4.13 is the *output* rate, then one may suspect that, with ideal scaling, *one third* of the resources may suffice for the target 30 FPS.

Also, the design process is unknown. It may be the case of an RTL design using the Vivado system (briefly mention Section 4.4.1). I can only guess that high-level synthesis using Vivado HLS was not used.

So, the quality of the implementation cannot be estimated, and hence the relevance of quantitative details. Nevertheless, the implementation and its measurements prove the main claim of the work.

I can conclude that the work contributes to present knowledge, support its claims well and therefore I recommend it for defense.

Jan Schmidt
CTU in Prague
Apr. 16th, 2021