主　题:	Re: Fw: ICCD 2024 notification for paper 99
发件人:	"Guojie Luo" <gluo@pku.edu.cn> 2024-8-2 14:45:53
收件人:	邹苏南 <zousunan@pku.edu.cn>
抄　送:	石弼钊 <shi_bizhao@pku.edu.cn>, 2000012970 <2000012970@stu.pku.edu.cn>

Good job!

---------- Original ----------

From: "Sunan Zou" <zousunan@pku.edu.cn>;
Date: Fri, Aug 2, 2024 10:36 AM
To: "gluo"<gluo@pku.edu.cn>;
Subject: Fw: ICCD 2024 notification for paper 99

> -----Original Messages-----

> From: "ICCD 2024" <iccd2024@easychair.org>
> Send time:Friday, 08/02/2024 09:17:39
> To: "Sunan Zou" <zousunan@pku.edu.cn>
> Subject: ICCD 2024 notification for paper 99
>
> Dear Sunan Zou,
>
> On behalf of the ICCD 2024 technical program committee, we are pleased to inform you that the following submission has been accepted as a regular paper for publication:
>
> 99 MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
>
> The selection process was very competitive, with a 28% acceptance ratio.
>
> Please see the comments below from our reviewers on your submission. We hope that they enable you to improve your work further before the final version and gain insights for your future research. The instructions for preparing your camera-ready version will be forwarded to you in the next few days.
>
> Registration & Presentation Requirement: At least one author of your paper must register and present your paper to be included in the conference proceedings. Please note that each paper must have a unique associated registration. If you are the author of more than one paper, a co-author must register for any additional papers, even if you present them.
>
> We look forward to meeting you and your co-authors at ICCD 2024 in Milan.
>
> If you have any additional questions, please feel free to contact us.
>
> Benjamin Carrion Schaefer, The University of Texas at Dallas, USA
> Sara Vinco, Politecnico di Torino, Italy
>
> ICCD 2024 Program Co-Chairs
>
> SUBMISSION: 99
> TITLE: MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
>
>
> ---------- REVIEW 1 ----------
> SUBMISSION: 99
> TITLE: MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
> AUTHORS: Sunan Zou, Bizhao Shi, Ziyun Zhang and Guojie Luo
>
> ----------- Overall evaluation -----------
> SCORE: 2 (accept)
> ----- TEXT:
> This paper proposes an FPGA-based multi-sketch accelerator, MUSA, considering the FPGA memory subsystem and using it effectively. The authors explained the sketch algorithm well, including its computation and memory needs. In addition to that, authors have introduced available memories and their efficient usage to accelerate and meet multi-sketch algorithm needs. MUSA employs hybrid parallelism, i.e., combining pessimistic (replica) and optimistic (banking) models to increase parallelism. MUSA has coalesced memory organization by storing sketch states of different sketches with the same index on the memory block. This helps reduce the number of memory accesses, and merges distributed memory into a structured one that solves routing congestion by shrinking the fan out of hash and updating modules connected to BRAM. The proposed architecture shows promising performance and accuracy. The paper could explore interleaved memory mapping to BRAM banks to further distribute traffic b
> etween BRAMs.
>
>
>
> ---------- REVIEW 2 ----------
> SUBMISSION: 99
> TITLE: MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
> AUTHORS: Sunan Zou, Bizhao Shi, Ziyun Zhang and Guojie Luo
>
> ----------- Overall evaluation -----------
> SCORE: 1 (weak accept)
> ----- TEXT:
> Due to the below reasons the paper should be accepted,
> The paper presents a novel hybrid parallelism and coalesced memory organization method for multi-sketch acceleration, addressing significant challenges in the field.
>
> The implementation on an AMD U280 acceleration card validates the practical applicability of MuSA, showing substantial improvements in real-world scenarios
>
> Extensive evaluations using both artificial and real-world datasets highlight MuSA's robustness and efficiency across different data streams
>
>
>
> ---------- REVIEW 3 ----------
> SUBMISSION: 99
> TITLE: MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
> AUTHORS: Sunan Zou, Bizhao Shi, Ziyun Zhang and Guojie Luo
>
> ----------- Overall evaluation -----------
> SCORE: 0 (borderline paper)
> ----- TEXT:
> * Paper summary
>
> : MuSA is a multi-sketch accelerator that proposes a hybrid (pessimistic & optimistic) parallelism and a coalesced memory organization. This paper is one of a few early research to support multiple sketch algorithms running on the device. The evaluation results show that MuSA achieves 15X of performance (throughput) improvements over the naive multi-skech accelerator.
>
>
> * Comments
>

> : The importance of supporting multi sketch algorithms seems interesting and reasonable. And there have been little research efforts to optimize multi sketch algorithms on FPGA. Therefore, it would expect a vast optimization space for this topic.
>
> The authors clearly identified the most critical factor for throughput; memory accesses. The proposed schemes are straight forward while there are a few things are not clearly explained. First, hybrid parallelism applies both pessimistic (e.g., replicate sketch states) and optimistic (e.g., banking). Therefore, it sounds like it needs more resources (area and bits). It is not clearly understood why hybrid parallelism is more resource efficient.
>
> Second, the coalesced memory organization does concatenate each state of different sketches and (it seems) does not reduce the total number of bits. The benefits of the coalesced memory organization is purely from the memory access reduction?
>
> In addition, from the proposed two schemes, the throughput improvement is naturally expected, but power consumption is also a critical metric for accelerator. It would be much greater if the paper could have discussed power consumption of each approach.
>
> Lastly, can the coalesced memory organization dynamically adapt different input sizes? or does it require software (library) that split the large input into the fixed size?
>
>
>
> ---------- REVIEW 4 ----------
> SUBMISSION: 99
> TITLE: MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
> AUTHORS: Sunan Zou, Bizhao Shi, Ziyun Zhang and Guojie Luo
>
> ---------- Overall evaluation ----------
> SCORE: 1 (weak accept)
> ----- TEXT:
> Summary:
> This paper proposes a multi-sketch accelerator on FPGA, called MuSA to optimize multi-sketch with hybrid parallelism. MuSA reduces memory resource consumption by a hybrid parallelism pattern which combines the strengths of pessimistic and optimistic parallelism. In addition, the author proposes a merged memory organization scheme for multiple sketches constructing data summaries in different sketches. Such organization which enables coalesced memory access patterns leads to better performance for multi-sketch deployment. To evaluate MuSA, the author compares it with state of the art accelerators on the dimensions of kernel thought, accuracy, resource utilization, and layout area and shows the strength of MuSA.
>
> Strengths:
> + Propose pessimistic and optimistic hybrid parallelism to achieve better trade-off for storage and throughput
> + Coalesced memory organization of mult-sketch with the same index to reducde memory access
>
> Weaknesses:
> - Advantages over prior work are not clearly presented
>
> Questions and Comments:
> Overall, this paper tries to solve an interesting problem of sketch algorithms on accelerators. It leverages the memory access pattern and hybrid parallelism to improve the efficiency of FPGA architecture.
>
> 1. How does MuSA compare to other approaches? This paper used SKT as the baseline, while there are other approaches proposes similar ideas, such as hybrid parallelism in [7]. How does MuSA compared to other studies?
>
> 2. What are the reasons of choosing (CM) sketch, Fast AGMS, and HyperLogLog(HLL) in this paper.
>
> 3. Introducing memory access earlier in the introduction section could help emphasize the importance and novelty of MuSA.
>
> 4. Incomplete sentence in the 3rd paragraph of Section II-B: "Therefore, routing con- gestion will likely appear around memory blocks and impede a" is an incomplete sentence.
>
>
>
> ---------- REVIEW 5 ----------
> SUBMISSION: 99
> TITLE: MuSA: Multi-Sketch Accelerator with Hybrid Parallelism and Coalesced Memory Organization
> AUTHORS: Sunan Zou, Bizhao Shi, Ziyun Zhang and Guojie Luo
>
> ---------- Overall evaluation ----------
> SCORE: 1 (weak accept)
> ----- TEXT:
> In the paper, the authors discuss how they implement an FPGA sketch
> accelerator leveraging a neat memory layout. They authors show
> respectable performance when compared to the state-of-the-are.
>
> The paper is well laid out and easy to read. I enjoy the discussion of
> the memory layout and the design trade-offs.
>
> That said, the the work is incremental.
>