11<sup>th</sup> Workshop on General Purpose Processing using GPUs (GPGPU 11)

Proceedings

David Kaeli and John Cavazos February 25, 2018 Vosendorf, Austria The Association for Computing Machinery 2 Penn Plaza, Suite 701 New York New York 10121-0701

COPYRIGHT NOTICE. Copyright © 2018 by the Association for Computing Machinery, Inc. Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Publications Dept., ACM, Inc., fax +1 (212) 869-0481, or permissions@acm.org.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, +1-978-750-8400, +1-978-750-4470 (fax).

## Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that was previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG Newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform permissions@acm.org, stating the title of the work, the author(s), and where and when published.

ACM-ISBN: 978-1-4503-5647-3

## Message from the Workshop Organizers

We would like to welcome you to the proceedings of the 11<sup>th</sup> Annual Workshop on General Purpose Processing using Graphics Processing Unit. We have another strong program that includes a keynote, and invited paper, and presentations of 6 out of the 11 submitted papers. We want to thank the Program Committee for their efforts.

This year's Committee included:	
Jose-Luis Abellan - UCAM	Michela Becchi, University of Missouri
Brad Beckmann - AMD	Martin Burtscher, Texas State University
Trevor Carlson, Uppsala University	Christian Fensch, Heriot-Watt University
Magnus Ekman - NVIDIA	Xin Fu – University of Houston
Isaac Gelado – NVIDIA	Michael Gerndt – Tech. Univ. of Munich
Byunghyun Jang - University of Mississippi	Adwait Jog – William and Mary
John Kim – KAIST	Shin-Ying Lee – Samsung Research
Yun Liang – Peking University	Shelley Lin – Northeastern University
Sonia Lopez Alarcon – RIT	Simon McIntosh-Smith – Univ. of Bristol
Avi Mendelsohn - Technion	Fanny Nina-Paravecino - Intel
Lena Oden – Argonne National Labs	Antonio Pena – Barcelona SC
Lu Peng – LSU	Oscar Plata - University of Malaga
Xuehai Qian – USC	Tim Rogers – Purdue University
Hamed Tabkhi – UNCC	Yash Ukidave – AMD
Alejandro Valero – UPV	Jun Yang – University of Pittsburgh
Ayse Yilmazer – Istanbul Kemerburgaz U.	Antonia Zhai – University of Minnesota

Huiyang Zhou - North Carolina State University

We are also pleased to have two quality keynote talks:

- "Generating High Performance GPU Code using Rewrite Rules with Lift," Christophe Dubach, University of Edinburgh.
- "Initial Steps toward Making GPU a First-Class Computing Resource: Sharing and Resource Management," Jun Yang, University of Pittsburgh.

We would also like to acknowledge the efforts of our webmaster and proceedings co-editor, Trinayan Baruah. We hope you will enjoy the proceedings, and that the papers provide you with new ideas to further motivate your research on graphics processors. We would also like to thank the ACM ICPS and Conference Publishing service for enabling us to publish this proceedings.

Regards,

John Cavazos and Dave Kaeli

**GPGPU 11 Co-organizers** 

## **Table of Contents**

## Session: Persistent Data Structures

Keynote: Initial Steps toward Making GPU a First-Class Computing Resource: Sharing and Resource Management Jun Yang	1
<b>A Case for Persist Barriers in GPUs</b> Dibakar Gope, Arkaprava Basu, Sooraj Puthoor and Mitesh Meswani	2
Session: Applications / Frameworks	
Overcoming the Difficulty of Large-scale CGH Generation on multi- GPU Cluster Takanobu Baba, Shinpei Watanabe, Boaz Jessie Jackin, Takeshi Ohkawa, Kanemitsu Ootsu, Takashi Yokota, Yoshio Hayasaki and Toyohiko Yatagai	13
<b>Transparent Avoidance of Redundant Data Transfer on GPU-enabled</b> <b>Apache Spark</b> <i>Ryo Asai, Masao Okita, Fumihiko Ino and Kenichi</i> <i>Hagihara</i>	22
GPU-based Acceleration of Detailed Tissue-Scale Cardiac Simulations Neringa Altanaite and Johannes Langguth	31
Session: Concurrent Kernels	
Keynote: Generating High Performance GPU Code using Rewrite Rules with Lift Christophe Dubach	39
MaxPair: Enhance OpenCL Concurrent Kernel Execution by Weighted Maximum Matching Yuan Wen, Michael O'Boyle and Christian Fensch	40
<b>Oversubscribed Command Queues in GPUs</b> Sooraj Puthoor, Xulong Tang, Joseph Gross and Bradford Beckmann	50