▶**Samar Aseeri**
4700 KAUST – THUWAL/SA
Phone: +966(12)8080314
E-mail: samar.aseeri@kaust.edu.sa

Education
Ph.D., Master and Bachelor in Applied Math   (2009, 2005 and 2001 respectively)
▶   **Seven papers published in international journals and conferences.**

Experience
Postdoctoral Fellow (2009 –2010)

Computational Scientist (2010 –present)

▶   **Application User Support (Consultation + Research involvement)**
▶   **Provide Lab-Wide Coordination for Training Activities**

Skills
Installing and Benchmarking Computer Applications
Performance Tools such as Scalasca and Tau + OpenFOAM

Outreach
ECRC project involvement.
WEP Lectures to Reach-Out to Campus to Engage on Supercomputing
HPBench track Chair (2016 – 2018) part of HPCS/IEEE conference
Invited Speaker at the M3HPCST, India, Dec (27-29) 2015
ACM, SIAM membership

Research
1) **S. Aseeri, S. Chandrasekaran, L. Dalcin, Y. Feng, F. Franchetti, A. Gholami, J. H. Goebbert, M. Mortensen, B.K. Muite, D. Pekurovsky, S. Plimpton, T. Popovici, and D. Takahashi. Distributed memory fast fourier transforms in the exascale era. Submitted, 2018.**
2) **S. Aseeri, O. Batrasev, M. Icardi, B.Leu, A. Liu, N. Li, B.K. Muite, E. Müller, B. Palen, M. Quell, H. Servat, P. Sheth, R. Speck, M. Van More, J. Vienne, "Solving the Klein-Gordon equation using Fourier spectral methods: A benchmark test for computer performance", ACM DL - Proceedings of the 23rd High Performance Computing Symposium (HPC 2015), held in Conjunction with 2015 Spring Simulation Multi-Conference, April 2015.**

Benchmarking Distributed Memory Fast Fourier Transforms

Samar Aseeri, PhD
Computational Scientist
KAUST

Benson Muite, PhD
Research Fellow of Distributed Systems
University of Tartu

Motivation

1. FFT is an important component of many programs.

2. On many emerging high performance computing architectures, the FFT may not work well on the full parallel computer

3. A good benchmark will lead to adoption of the best FFT software technology.

4. Identification of alternative algorithms to the FFT along with comparisons of efficiency will lead to optimal use of high performance computers

5. A galvanized and involved benchmarking community is required to do this
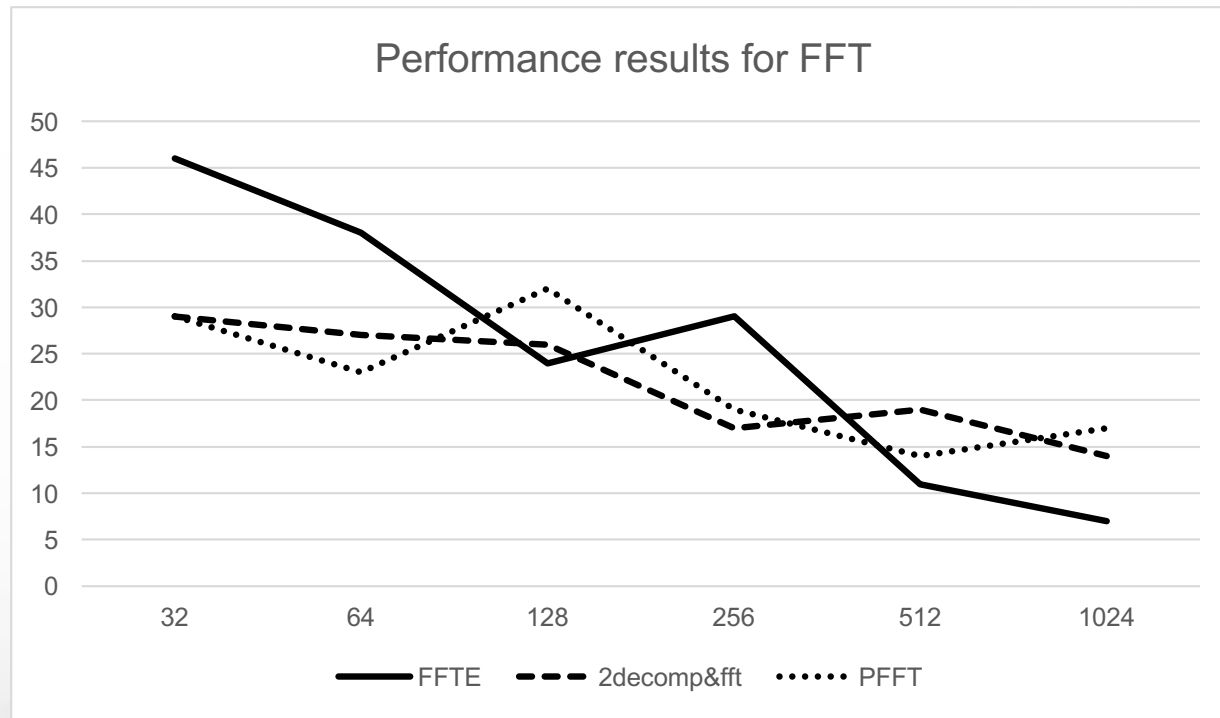
Previous Work

1.   FFTW benchmark webpage -> Not regularly updated

2.   1D FFT benchmark in HPCC -> Not many submissions since 2013

3.   3D FFT in NAS parallel benchmarks -> Limited uptake

4.   GearSHIFFT data comparison -> No distributed memory data

Comparing Parallel FFT libraries

Performance results for FFT

Comparing Parallel FFT libraries

- Even on one computer, many sources of variation including

  - topology mapping

  - compiler and compiler options

  - other jobs running

  - operating system factors

- Wide variety of possible results on same machine with same source code

- May need to also estimate likely variation between jobs

- May still want to compare effects of optimization and operating system

Speakers	Affiliation	Software
Daisuke Takahashi	University of Tsukuba, Japan	FFTE
Thom Popovoci	Carnegie Mellon University, USA	Spiral
Yu Feng	Berkeley Center for Cosmological Physics, USA	Pfft.py
Jens Henrik Göbbert	Jülich Supercomputing Center, Germany	nb3dfft
Sunita Chandrasekaran	University of Delaware, USA	PsFFT, cusFFT
Amir Gholami	University of California at Berkeley, USA	AccFFT
Dmitry Pekurovsky	San Diego Supercomputing Center, USA	P3DFFT
Steve Plimpton	Sandia National Laboratories, USA	Parallel FFT

Questions?

SC17
Denver, CO | hpc connects.

1.   Why did you write your own FFT?

2.   What considerations are important for you in an FFT implementation?

3.   What might you look for if there were to be a unified FFT interface (similar to BLAS, LAPACK and SCALAPACK interfaces)?

4.   How important are performance, portability, and scalability for you?

5.   Will FFT be needed in exascale computing and if so how will it be achieved?

6.   What would be a good FFT benchmark or a good way to include the FFT in a high-performance computer benchmark?

Links

- Webpage: http://www.fft.report/FFT_BoF.html

- Community forum: https://www.forum.fft.report

Room: 52–104 (Wednesday, March 7)

1:00 PM – 2:40 PM

MS2 State-of-the-Art FFT —— Algorithms, Implementations, and Applications – Part I of II

Organizer: Daisuke Takahashi *University of Tsukuba, Japan*
Franz Franchetti *Carnegie Mellon University, USA*
Samar A. Aseeri *King Abdullah University of Science & Technology (KAUST), Saudi Arabia*
Benson K. Muite *University of Tartu, Estonia*

1:00-1:20 Implementation of Parallel FFTs on Cluster of Intel Xeon Phi Processors
Daisuke Takahashi, University of Tsukuba, Japan
1:25-1:45 SPIRAL FFT
Franz Franchetti, Carnegie Mellon University, USA
1:50-2:10 Pipelining Fast Fourier Transform on the OpenPOWER Cluster
Jun Doi, IBM Research - Tokyo, Japan
2:15-2:35 Automatic FFT Kernel Generation for CUDA GPUs
Akira Nukada, Tokyo Institute of Technology, Japan

3:10 PM – 4:50 PM

MS13 State-of-the-Art FFT —— Algorithms, Implementations, and Applications – Part II of II

Organizer: Daisuke Takahashi *University of Tsukuba, Japan*
Franz Franchetti *Carnegie Mellon University, USA*
Samar A. Aseeri *King Abdullah University of Science & Technology (KAUST), Saudi Arabia*
Benson K. Muite *University of Tartu, Estonia*

3:10-3:30 Fast Fourier Transforms (fft)
Samar A. Aseeri, King Abdullah University of Science & Technology (KAUST), Saudi Arabia
3:35-3:55 Use of the Fast Fourier Transform in Solving Partial Differential Equations
Benson K. Muite, University of Tartu, Estonia
4:00-4:20 Parallel Fast Gauss Transform
Shravan Veerapaneni, University of Michigan, USA; *Hari Sundar*, University of Utah, USA
4:25-4:45 Implementation of OpenFFT and Its Application to Industrial Problems
Truong Vinh Truong Duy, Nissan ARC, Japan

Summary

- Several different proposed measures for benchmarking FFT have evolved with architectural trends
- A widely run long lasting set of benchmarks is required
- An FFT benchmark should allow for easy comparison with possible FFT alternatives

Objective

- Find a good way to benchmark the FFT
- Enable performance prediction for the FFT pattern in parallel computing
- Ensure uptake of the benchmark

Approach

- Community discussion
  - Online
  - In person
- Testing of proposed benchmarks

Outcome

- Data repository
- More productive supercomputers
- Fame and/or Fortune
- Community involvement