GPU Acceleration of WCT

Kwangmin Yu, Meifeng Lin, Zhihua Dong

Computational Science Laboratory

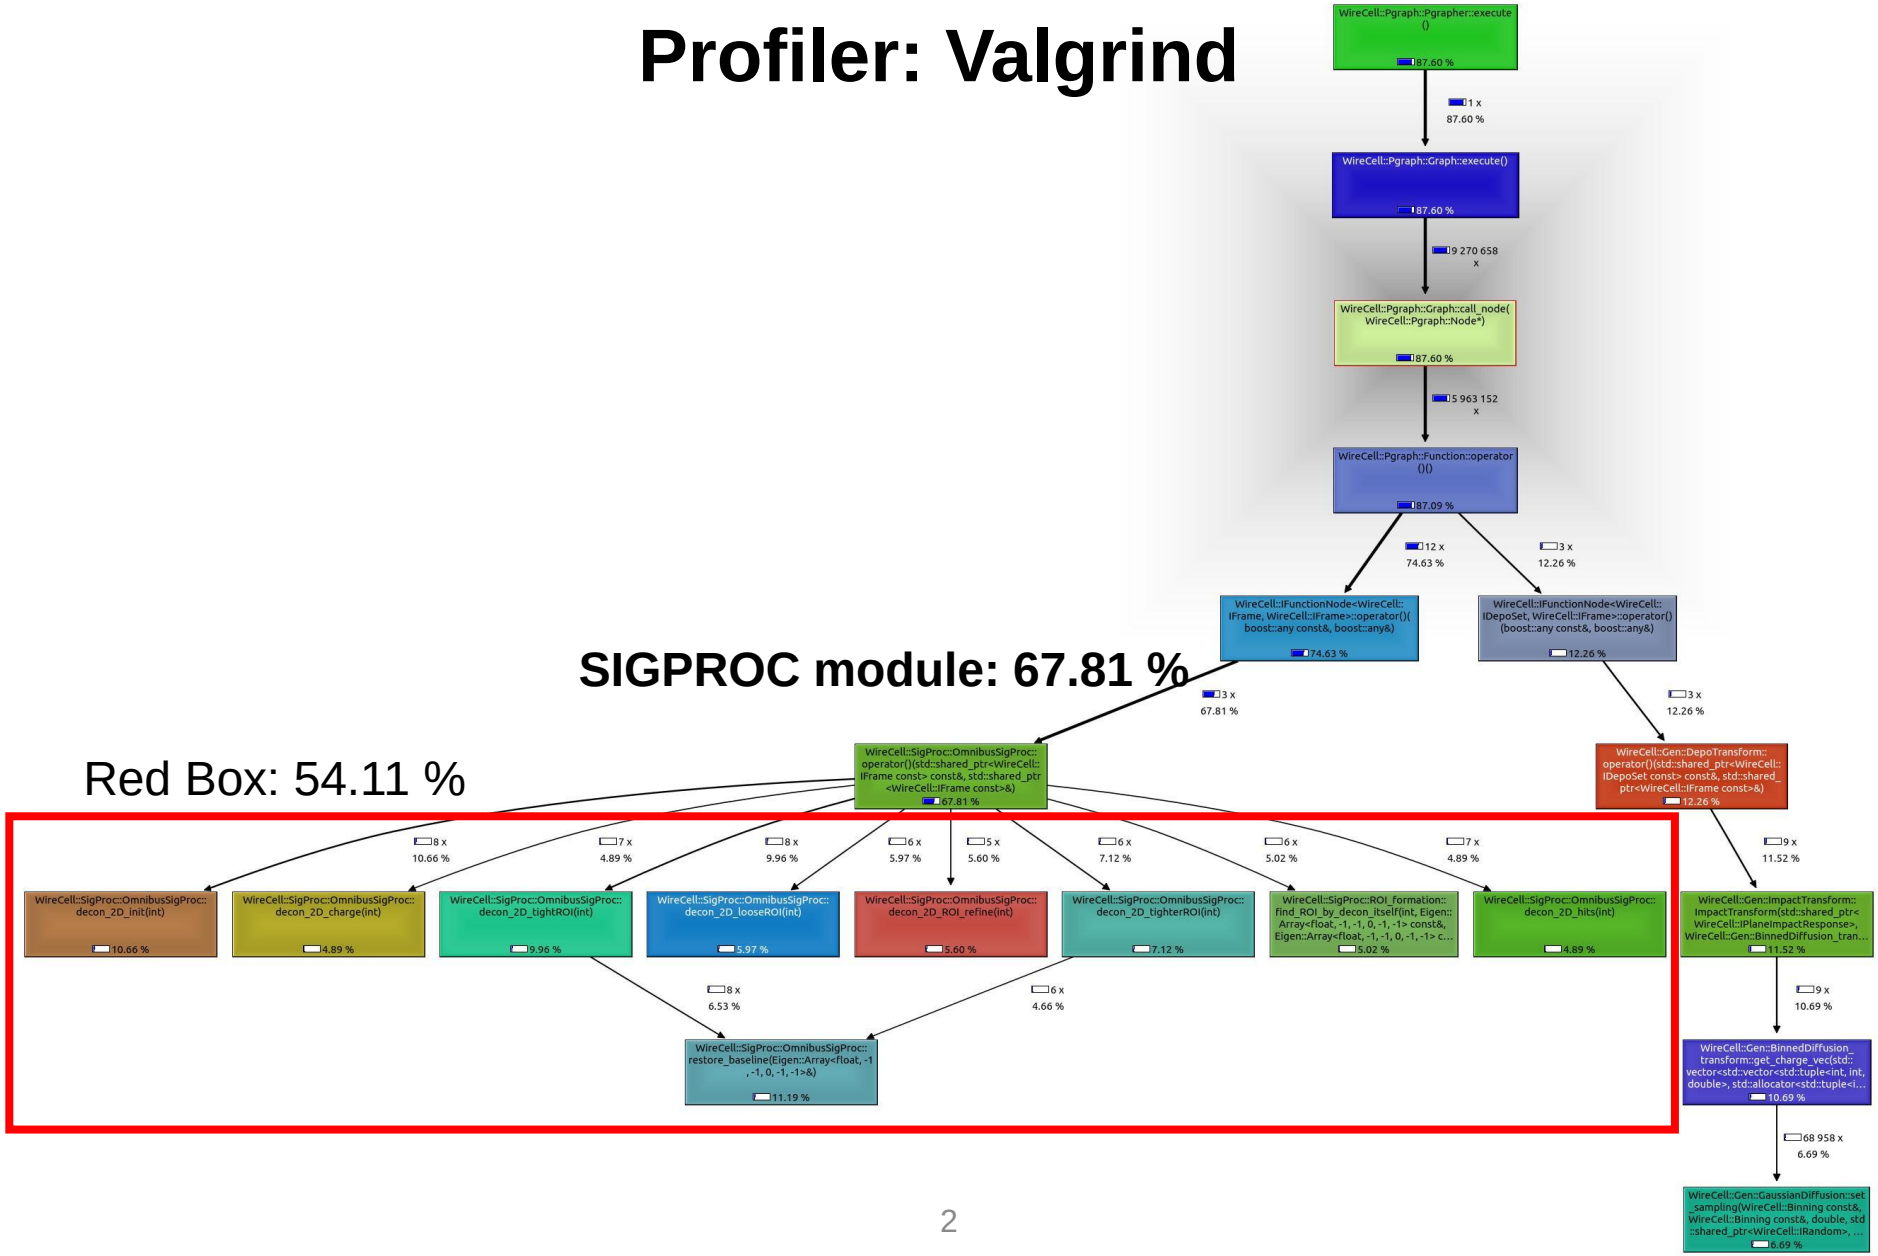
BROOKHAVEN NATIONAL LABORATORY

U.S. DEPARTMENT OF ENERGY

Jan. 27th, 2020

Profile

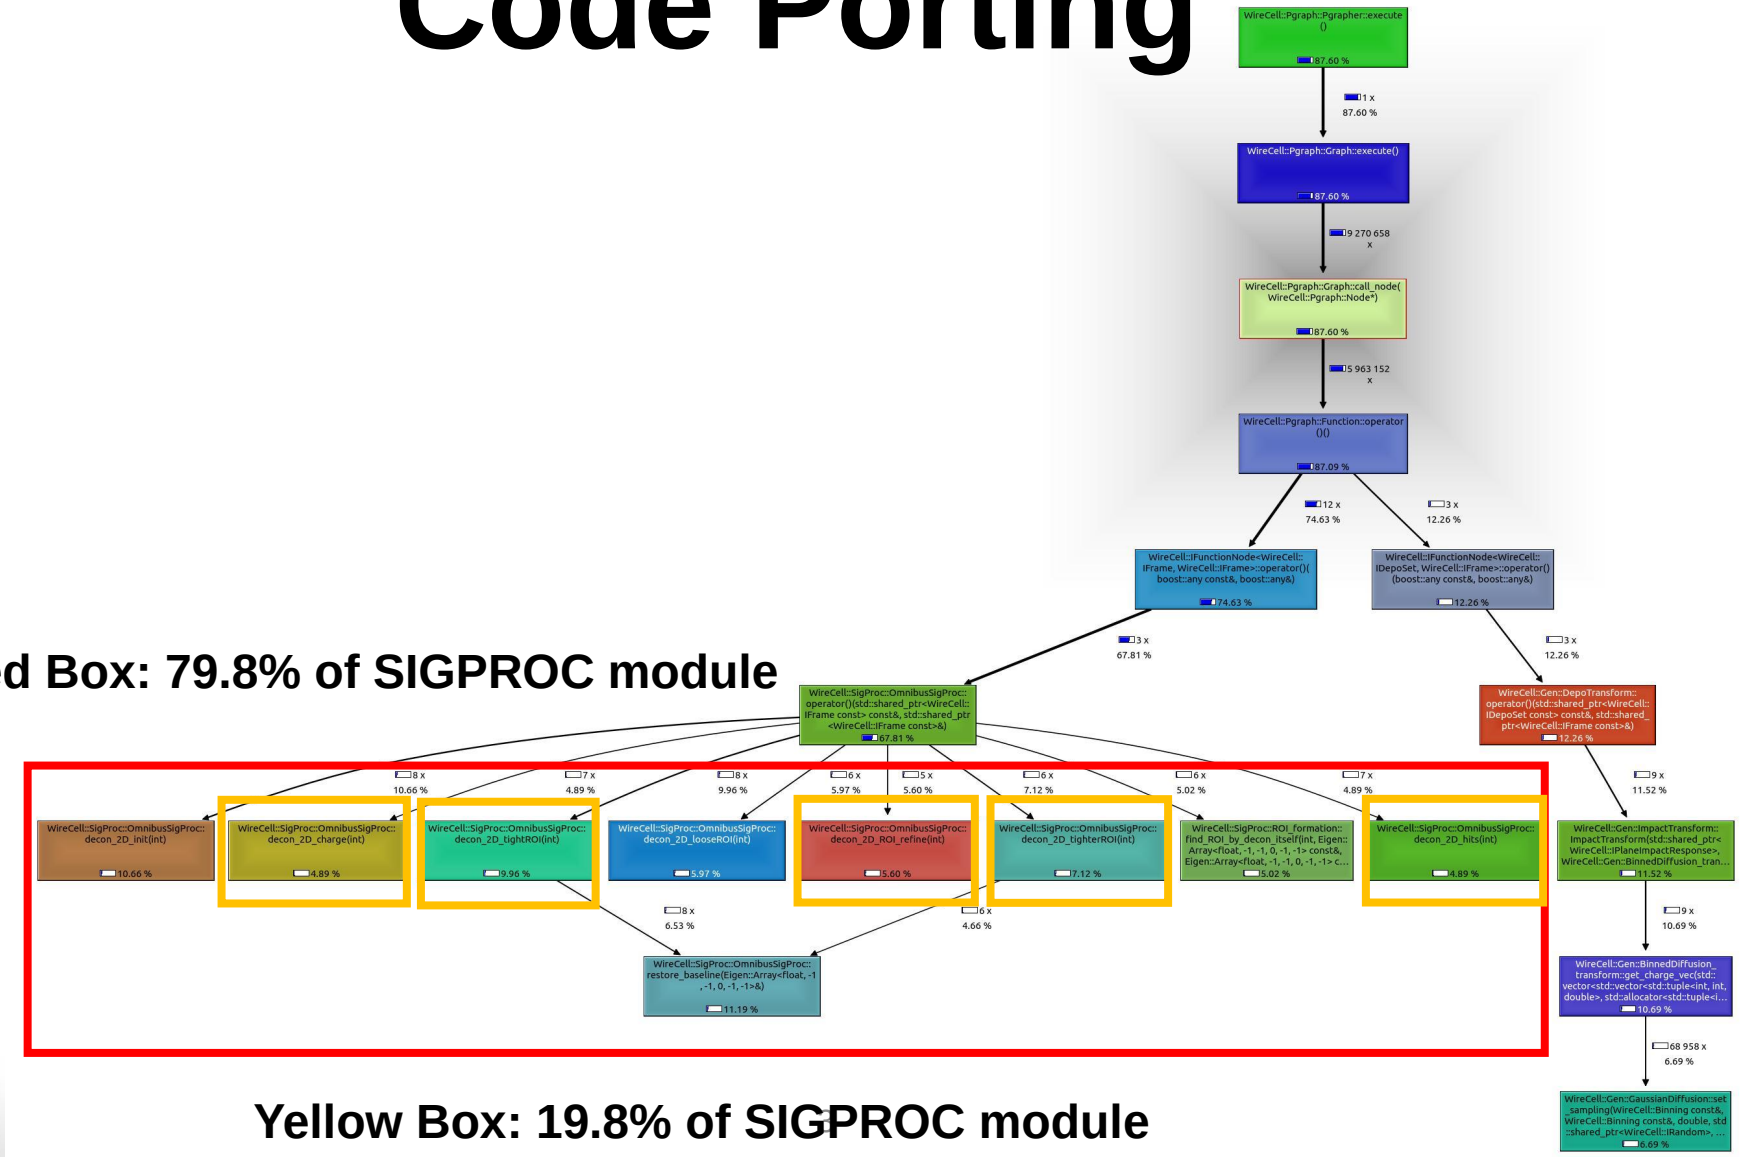
Profiler: Valgrind



SIGPROC module: 67.81 %

Red Box: 54.11 %

Code Porting



Red Box: 79.8% of SIGPROC module

Yellow Box: 19.8% of SIGPROC module

Code Porting

- GPU acceleration by CUDA was applied on the subroutines of the **yellow box** (**decon_2D_charge(), decon_2D_tightROI(), decon_2D_ROI_refine(), decon_2D_tighterROI(), decon_2D_hits()**).
- Base data format is a matrix (ex. 800 ~ 960 X 6000). Each row is one signal.
- Most significant part:
  - ✓ Inverse Fourier Transform (C -> R) & Applying filter
  - ✓ Several filters on signals
- Porting Strategy:
  - ✓ cuFFT library for Inverse Fourier Transform.
  - ✓ Several CUDA kernel functions are implemented.
- Testbed Hardware:
  - ✓ Intel(R) Xeon(R) CPU E5-2695 v4 @ 2.10GHz (36 Cores)
  - ✓ NVIDIA Tesla P100
- Test Result:

(sec)	Accel. Part	SIGPROC
CPU	77.14	389.53
GPU	4.02	316.41
Speed-up	**19.19**	**1.23**

Conclusion and Plan

- In the input setting, several functions in SigProc::OmnibusSigProc class takes about 80% of SIGPROC module.
- Problem size (800 ~ 960 signals of 6000 sampling data) is still small to utilize whole resource of GPUs.
- In the test CUDA implementation, we have 19 times speed-up.
- Data migration from CPU to GPU will reduce data transfer more.
- Supporting multiple GPU devices will be considered.