SAND2020-5327PE

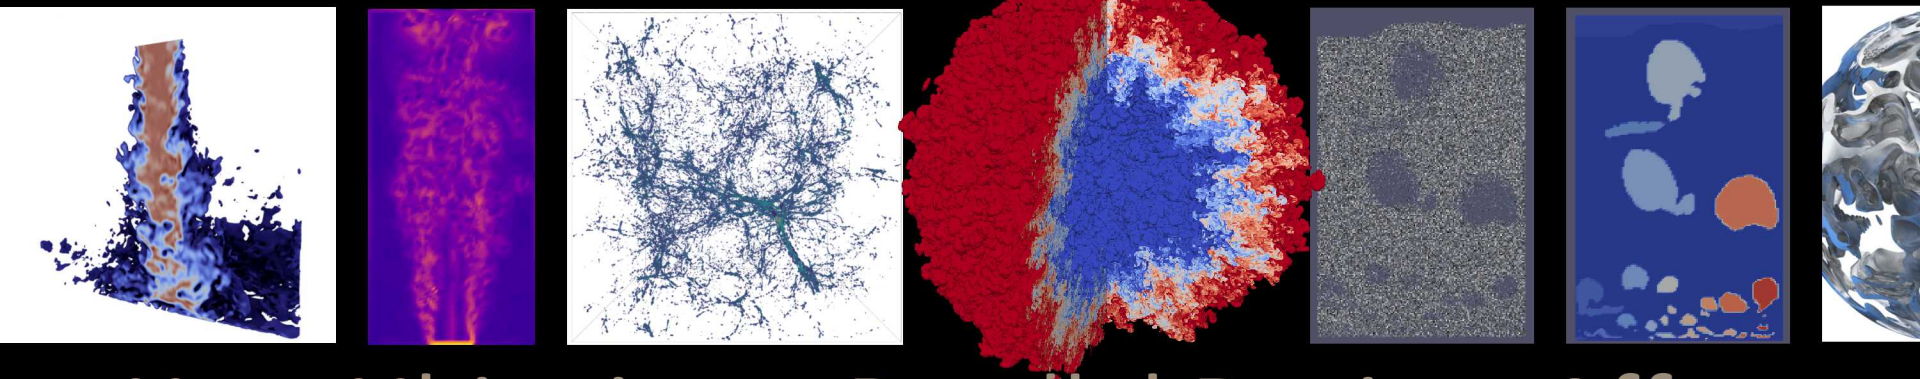

How Ubiquitous Parallel Devices Affect Visualization

EGPGV 2020 Panel

Monday, May 25, 2020

Kenneth Moreland Sandia National Laboratories

CCR Center for Computing Research

Panel Topic

Sandia National Laboratories

How Ubiquitous Parallel Devices Affect Visualization

HPC facilities chasing Exascale computing, desktops providing better productivity or gaming, and mobile devices cramming more features in tiny devices have very different concerns and goals, but all these diverse requirements converge on a common solution of increased parallel computing. Thus, **parallel computing is now relevant on everything from the worlds most powerful supercomputers to the phones in our pockets**. The topic for our panel discussion is how these ubiquitous parallel devices, from multicore CPUs to manycore GPUs, **affect the visualization community**. How is our research and development changing? How does EGPGV's role change? Is EGPGV becoming more or less relevant with ubiquitous parallel devices?

Parallel is Here

My new computer's got the clocks, it rocks
But it was obsolete before I opened the box

– “Weird” Al Yankovic, *It’s All About the Pentiums*, circa 1999

Moore’s Law is dead.

– Gordon Moore, circa 2005

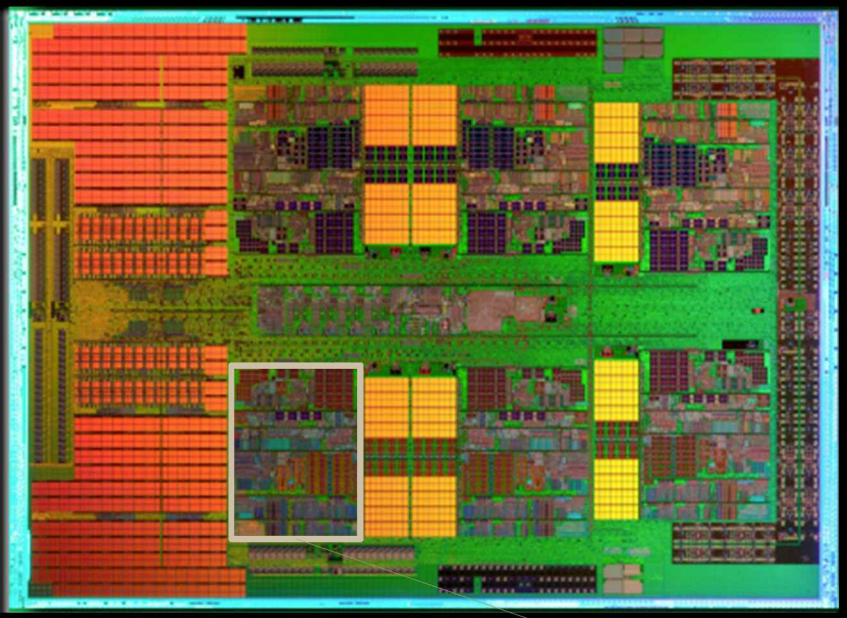

AMD x86

Full x86 Core
+ Associated Cache
6 cores per die
MPI-Only feasible

1 x86 core

1 Kepler "core"

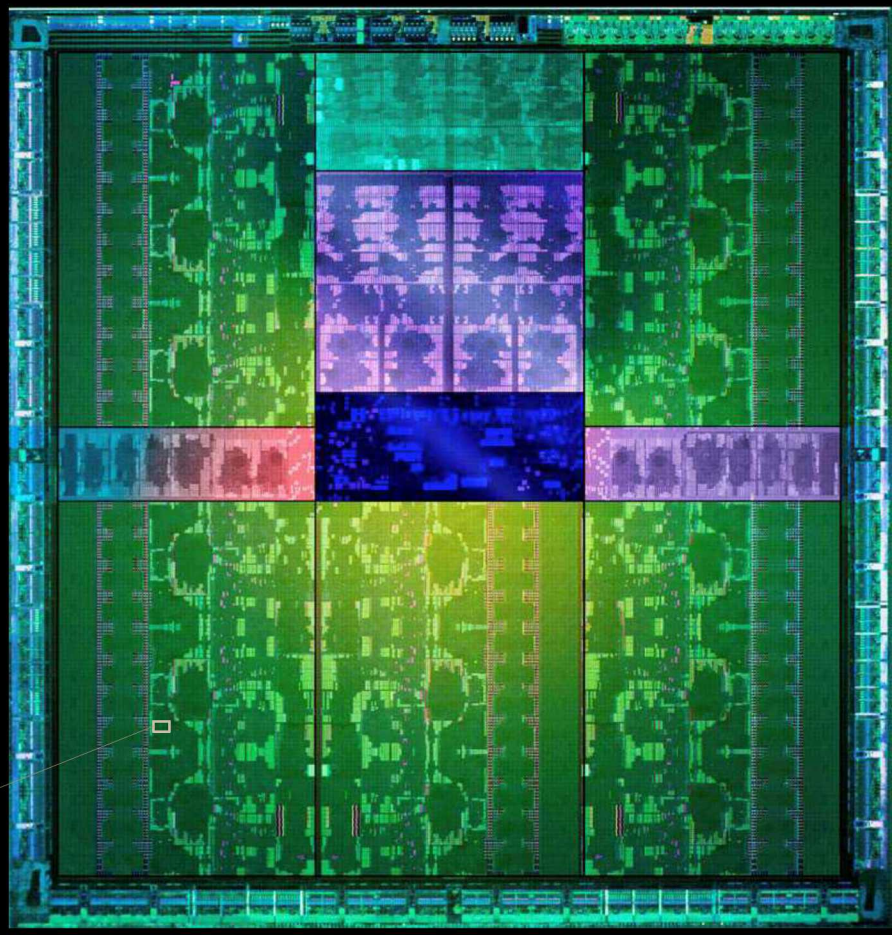

NVIDIA GPU

2,880 cores collected in 15 SMX
Shared PC, Cache, Mem Fetches
Reduced control logic
MPI-Only not feasible

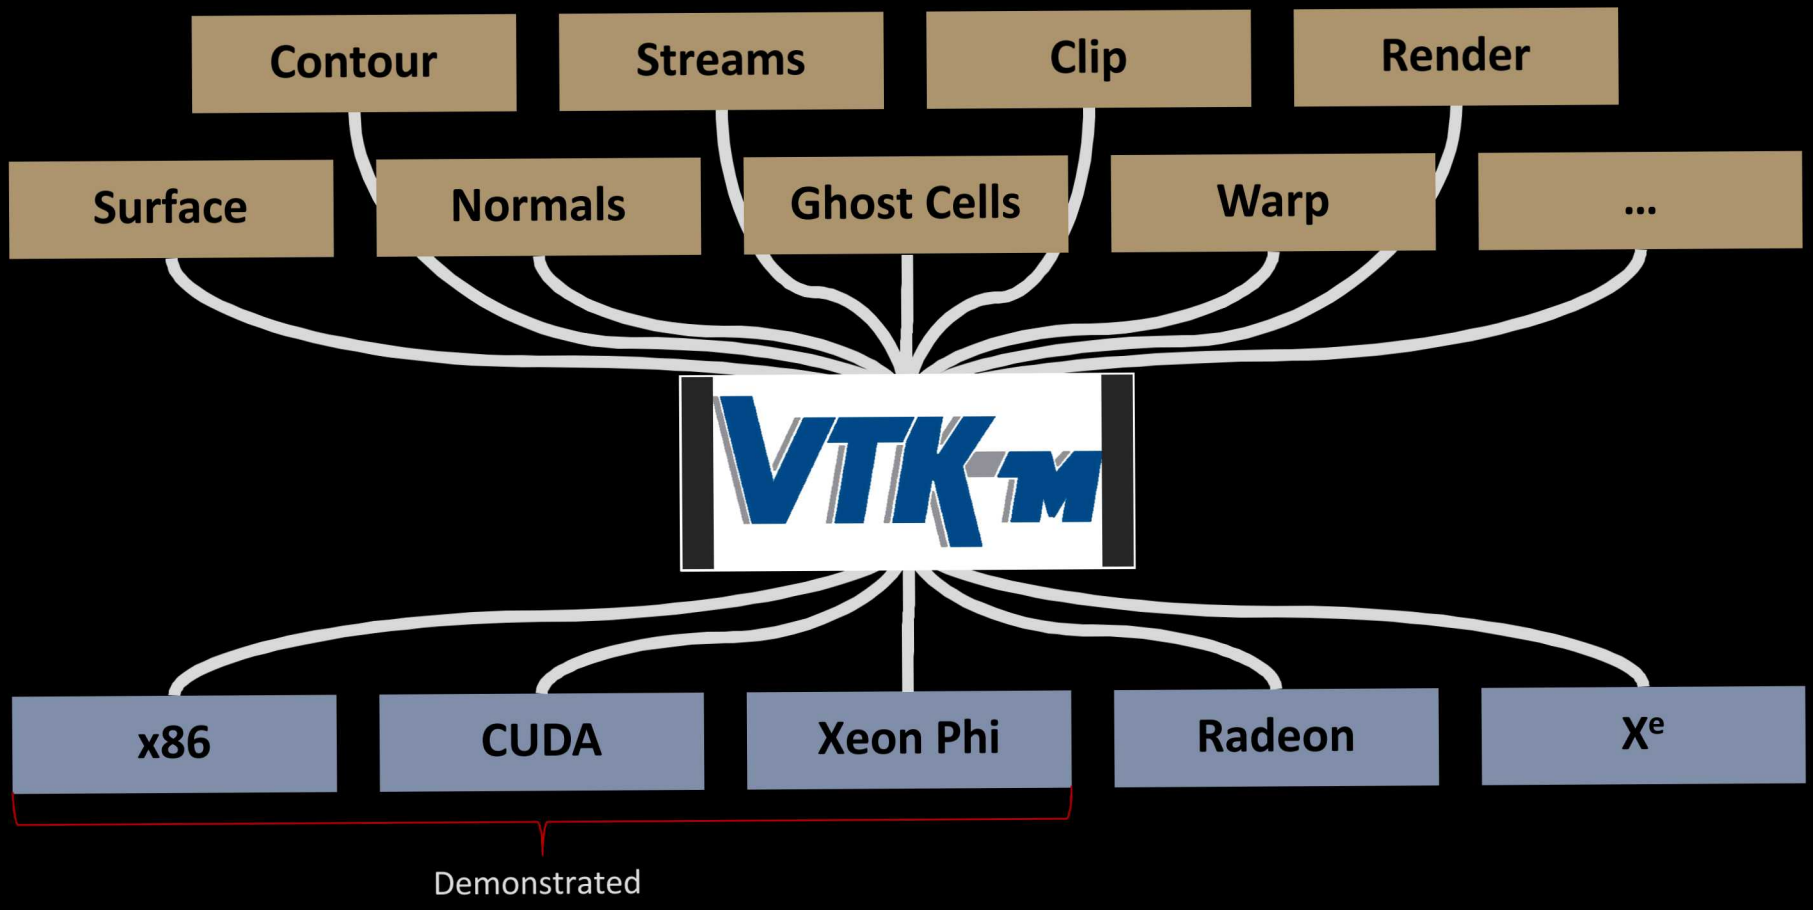
Contour
Streams
Clip
Render
Surface
Normals
Ghost Cells
Warp
...
VTK-m
x86
CUDA
Xeon Phi
Radeon
Xe
Demonstrated

Rendering (bigger is better) Moreland, et al., "VTK-m: Accelerating the Visualization Toolkit for Massively Threaded Applications, 2016

Dataset	Algorithm	Millions of rays per second
LT_350K	OptiX Prime	357.6
	EAVL	150.8
	VTK-m	164.5
LT_372K	OptiX Prime	322.4
	EAVL	124.7
	VTK-m	140.8
RM_350K	OptiX Prime	436.5
	EAVL	197.5
	VTK-m	200.8
RM_650K	OptiX Prime	420.4
	EAVL	172.9
	VTK-m	166.0
RM_970K	OptiX Prime	347.1
	EAVL	152.8
	VTK-m	163.5

Dataset	Algorithm	Millions of rays per second
LT_350K	Embree	51.9
	EAVL	27.7
	VTK-m	38.5
LT_372K	Embree	56.5
	EAVL	26.1
	VTK-m	36.0
RM_350K	Embree	64.8
	EAVL	33.3
	VTK-m	47.8
RM_650K	Embree	65.9
	EAVL	35.6
	VTK-m	49.1
RM_970K	Embree	59.1
	EAVL	29.3
	VTK-m	41.0

Particle Advection Pugmire, et al., "Performance-Portable Particle Advection with VTK-m," 2018

		GPU with data transfer			GPU without data transfer			CPU		
	File	K20X	K80	P100	K20X	K80	P100	$Intel_{16}$	$Intel_{28}$	$IBM\ P8_{20}$
W_1	Astro	0.627s	0.521s	0.389s	0.000s	0.011s	0.014s	0.001s	0.001s	0.001s
	Fusion	0.627s	0.521s	0.387s	0.001s	0.011s	0.015s	0.001s	0.001s	0.001s
	Thermal	0.627s	0.521s	0.392s	0.001s	0.011s	0.024s	0.001s	0.001s	0.001s
W_2	Astro	0.648s	0.543s	0.404s	0.021s	0.033s	0.029s	0.071s	0.046s	0.053s
	Fusion	0.649s	0.543s	0.400s	0.023s	0.033s	0.028s	0.071s	0.051s	0.052s
	Thermal	0.648s	0.541s	0.395s	0.021s	0.031s	0.027s	0.074s	0.048s	0.051s
W_3	Astro	1.511s	0.946s	0.577s	0.884s	0.436s	0.202s	3.003s	1.257s	2.327s
	Fusion	1.509s	0.961s	0.582s	0.883s	0.451s	0.210s	2.948s	1.208s	2.609s
	Thermal	1.508s	0.945s	0.583s	0.881s	0.435s	0.215s	2.801s	1.179s	2.691s
W_4	Astro	5.193s	2.851s	1.765s	4.566s	2.341s	1.390s	28.702s	10.688s	20.708s
	Fusion	5.327s	2.795s	1.776s	4.701s	2.285s	1.404s	26.295s	10.785s	19.949s
	Thermal	5.099s	2.785s	1.777s	4.472s	2.275s	1.409s	26.641s	11.266s	19.365s
W_5	Astro	38.660s	23.322s	13.338s	38.033s	22.812s	12.963s	256.900s	107.806s	185.852s
	Fusion	41.116s	24.450s	13.648s	40.490s	23.940s	13.276s	272.165s	107.113s	186.455s
	Thermal	39.444s	24.153s	13.626s	38.817s	23.643s	13.258s	260.740s	106.881s	193.110s

External Surface (smaller is better) Lessley, et al., "Techniques for Data-Parallel Searching for Duplicate Elements," 2017

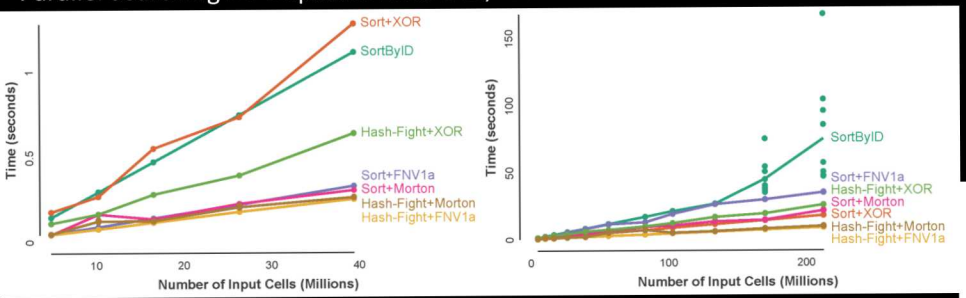

Point Merging (bigger is better) Yenpure, et al., "Efficient Point Merging Using Data Parallel Techniques," 2019

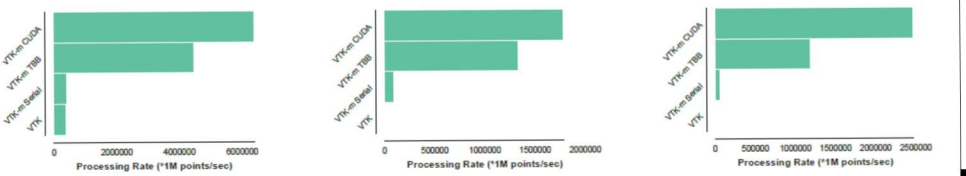

(a) Processing rates for fusion data (b) Processing rates for thermal hydraulics data (c) Processing rates for supernova data

Contour Tree (smaller is better) Carr, et al., "Parallel Peak Pruning for Scalable SMP Contour Tree Computation," 2016

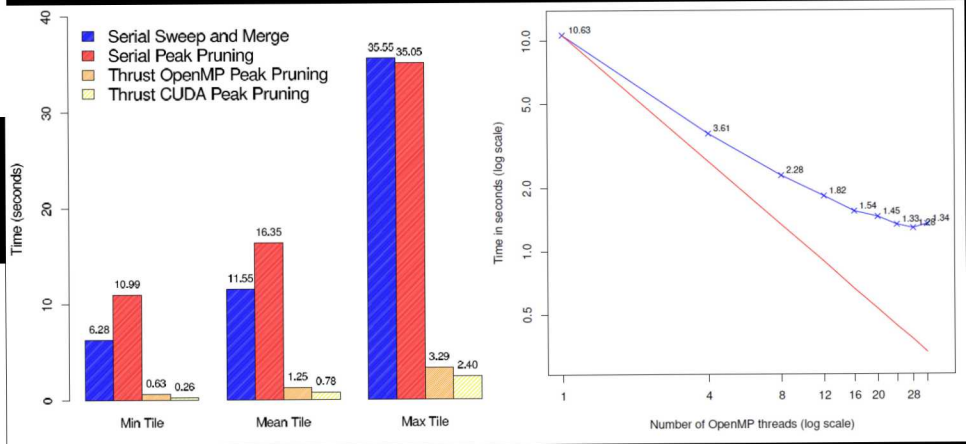

Are Serial Algorithms Still Relevant?

Sandia National Laboratories

- Yes
  - CPUs as we know them are not going away any time soon.
- But, on even cheap, old equipment you have extra processing that is sitting idle
- So, if performance is a concern, this is becoming less viable.
  - (You all know I'm being cavalier about important issues like parallel overhead, natural dependencies, data bandwidths, etc.)

Is EGPGV more or less relevant?

Sandia National Laboratories

- Obviously more relevant than ever:
  - Now that every computing platform is a parallel computing platform, a symposium on **parallel** visualization and graphics is vital to our community.
  - EGPGV leads the way to the future!
- Obviously completely irrelevant:
  - Now that every computing platform is a parallel computing platform, parallel visualization and graphics is no longer a niche community that needs its own symposium.
  - We take over EuroVis!

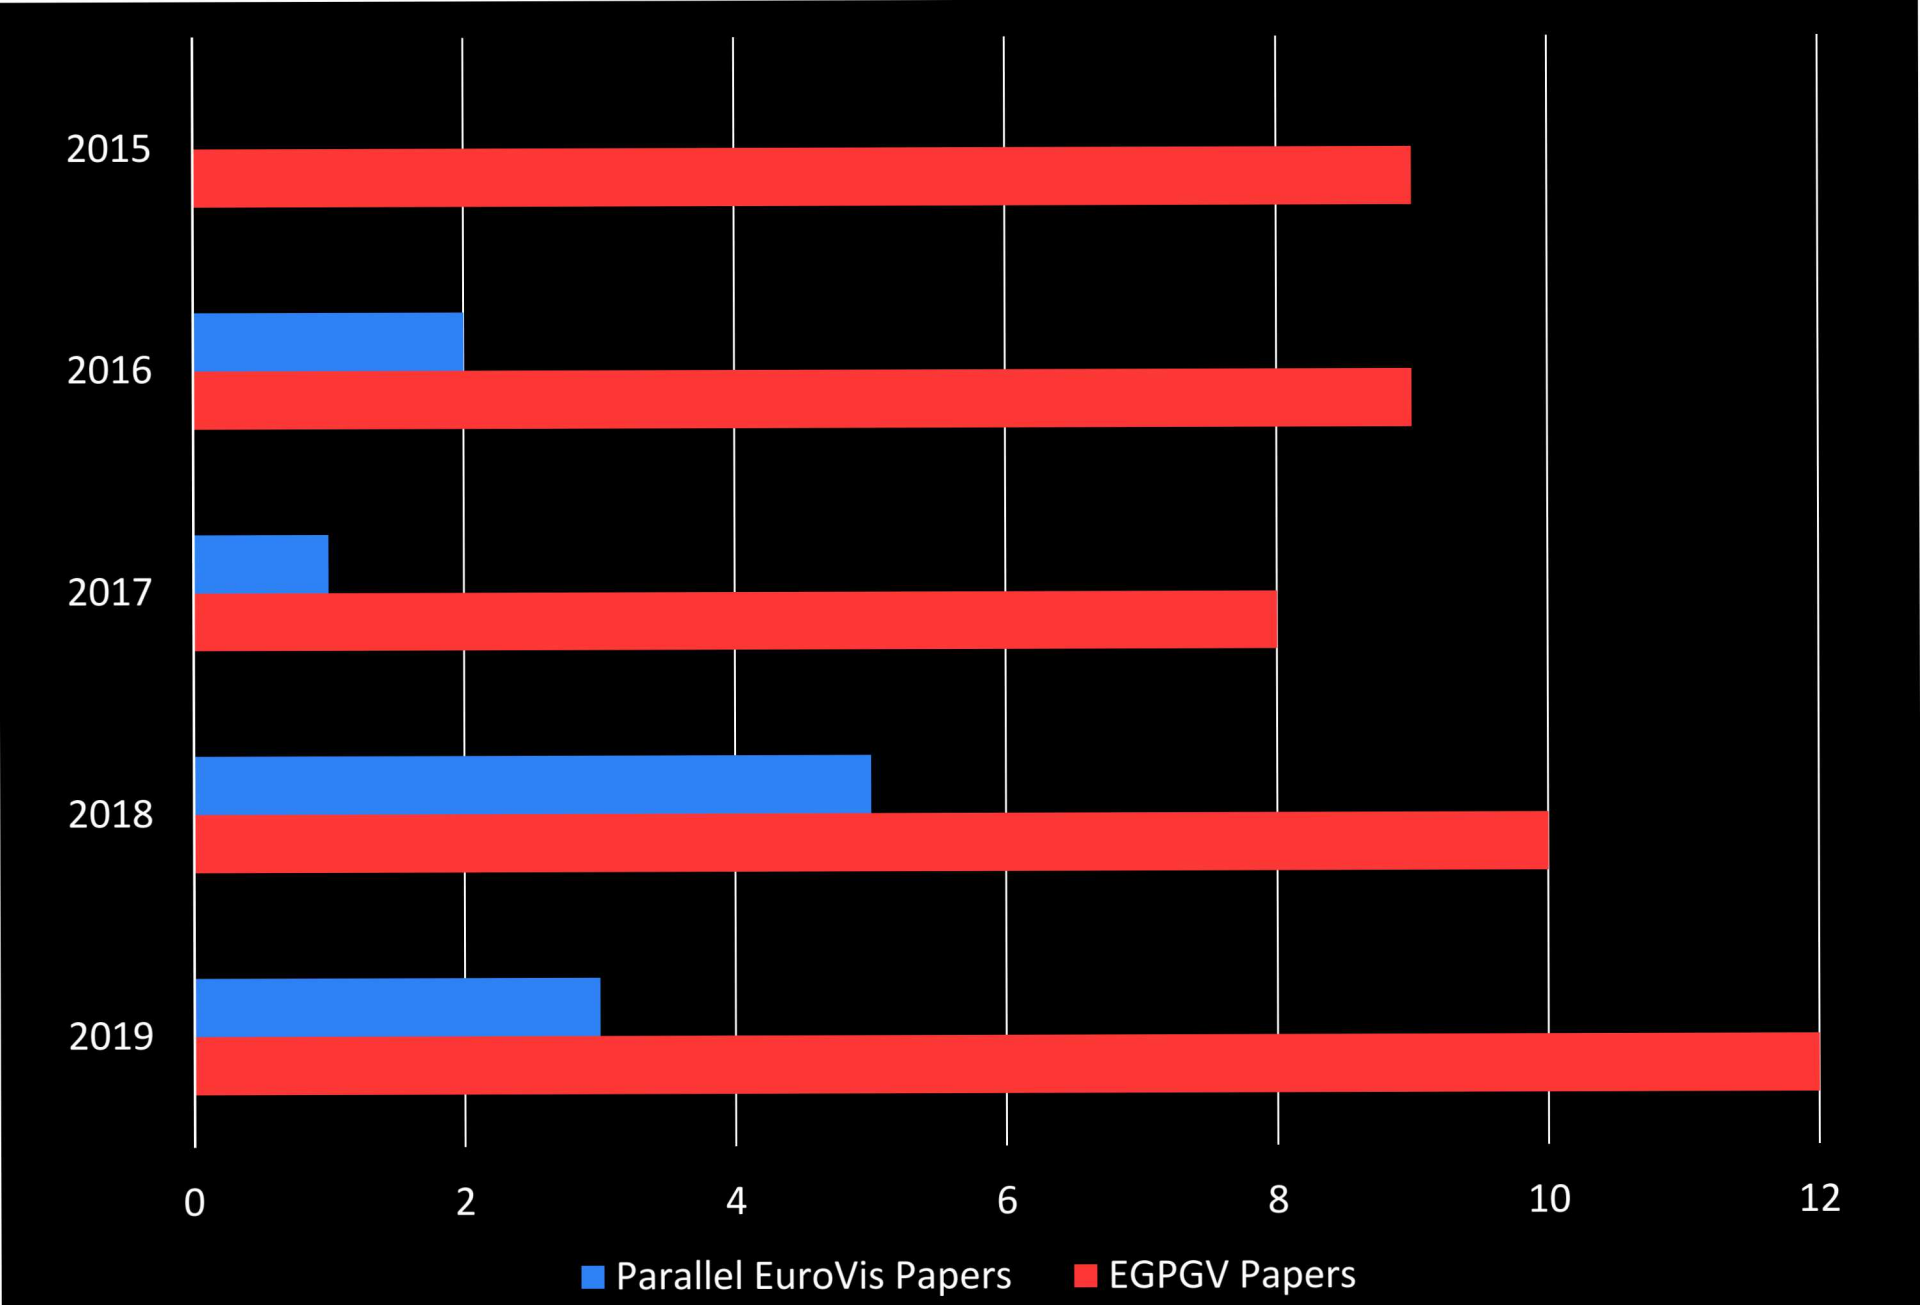

* A EuroVis paper is included if it has the word “parallel” in its abstract and it does not refer to something other than computing (e.g. parallel coordinates). All EGPGV papers are considered parallel.

In my Opinion…

- EGPGV exists because the general vis community does not appreciate the importance and difficulty of parallel algorithms.
  - Demonstrating an existing algorithm in parallel is both important and novel in its own right.
  - Showing known parallel algorithms at scale is both important and novel in its own right.
- Parallel algorithms might be a generational thing
  - Today's students have never known a world where parallel computing was not on every computer they used.