Argonne NATIONAL LABORATORY

Future Machines

J. Taylor Childers (ATLAS Appointed Soothsayer)

Some guidelines for this overview

- For this talk I limited my inclusion to machines in the top 25 of the Top500 that I am aware that ATLAS could access (if not already has access).
- There were some machines that made me wonder if they should be included or targeted, but given 4 days to prepare this talk, I skipped them.
- Don't be offended if your favorite machine was not included.

Rank	Site	System
1	DOE/SC/Oak Ridge National Laboratory United States	**Summit** - IBM Power S AC922, IBM POWER9 2 NVIDIA Volta GV100, Du Mellanox EDR Infiniban IBM
2	National Supercomputing Center in Wuxi China	**Sunway TaihuLight** - Su Sunway SW26010 260C Sunway NRCPC
3	DOE/NNSA/LLNL United States	**Sierra** - IBM Power Sys S922LC, IBM POWER9 NVIDIA Volta GV100, Du Mellanox EDR Infiniban IBM
4	National Super Computer Center in Guangzhou China	**Tianhe-2A** - TH-IVB-FE Intel Xeon E5-2692v2 1 TH Express-2, Matrix-2 NUDT
5	National Institute of Advanced Industrial Science and Technology (AIST) Japan	**AI Bridging Cloud Infra (ABCI)** - PRIMERGY CX Xeon Gold 6148 20C 2.4 NVIDIA Tesla V100 SXM Infiniband EDR Fujitsu
6	Swiss National Supercomputing Centre (CSCS) Switzerland	**Piz Daint** - Cray XC50, X 2690v3 12C 2.6GHz, Ari interconnect , NVIDIA T Cray Inc.
7	DOE/SC/Oak Ridge National Laboratory United States	**Titan** - Cray XK7, Opter 16C 2.200GHz, Cray Ge interconnect, NVIDIA K Cray Inc.

Current HPC State (in Top25)

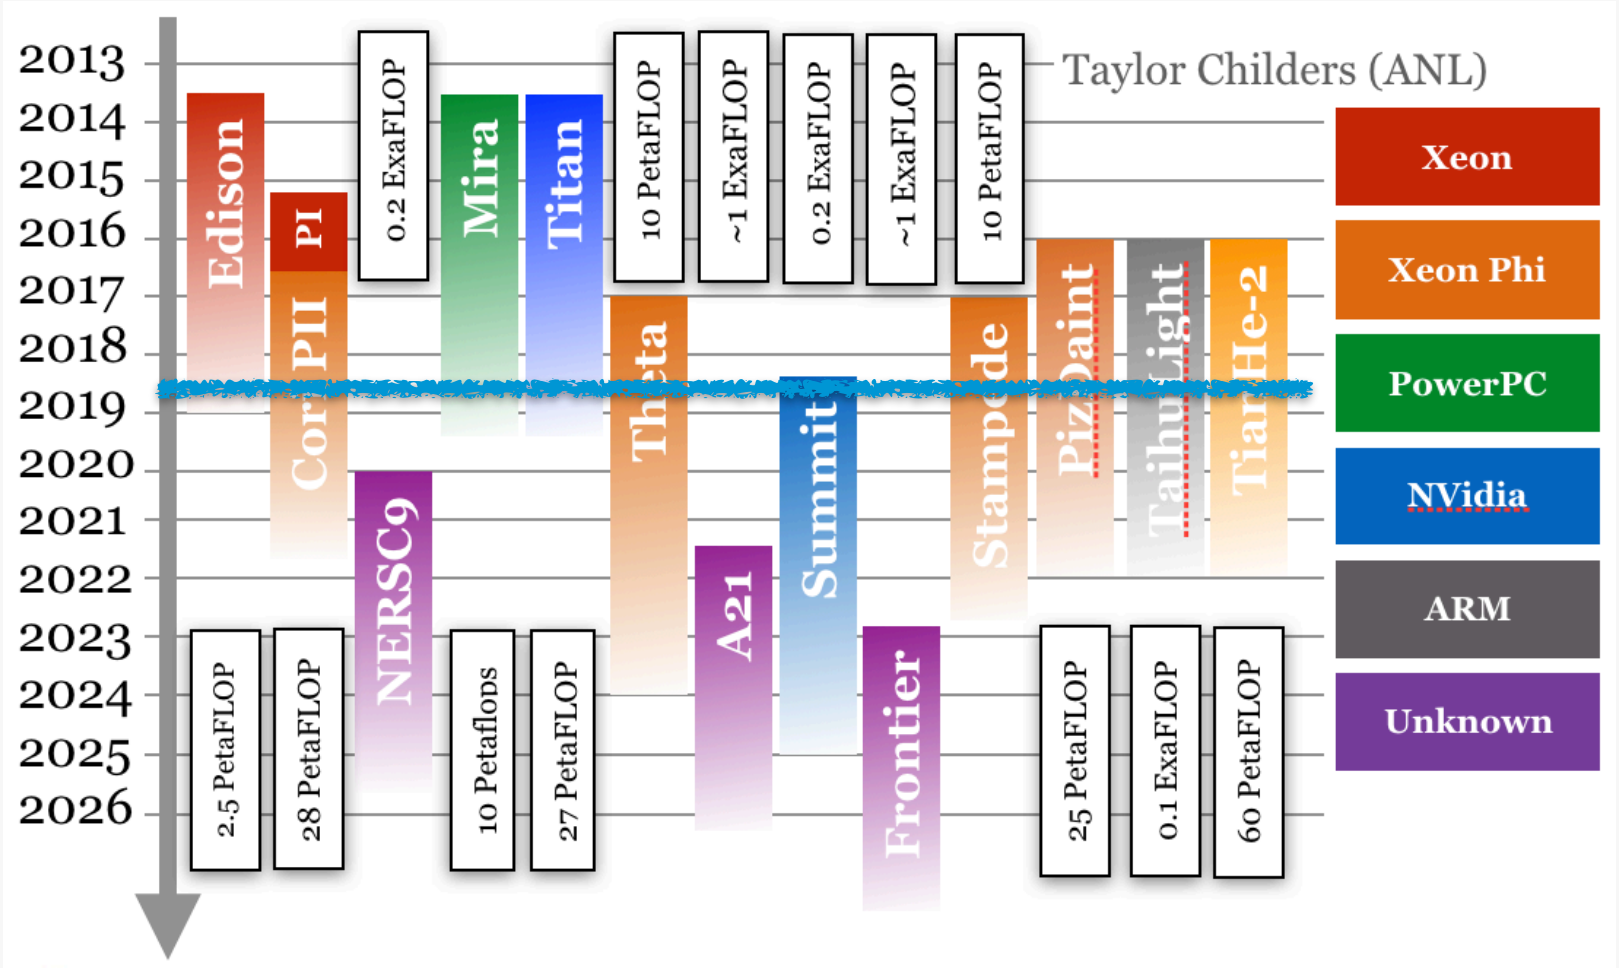

- Current HPC landscape contains a myriad of architectures
- Chinese resources have burst on to the Top500
- Swiss Supercomputing also
- Largely consumer product machines or at least they tried to be.
- HPC architectures in high flux and simulations no longer the singular driving force.

Future of HPC Around the World

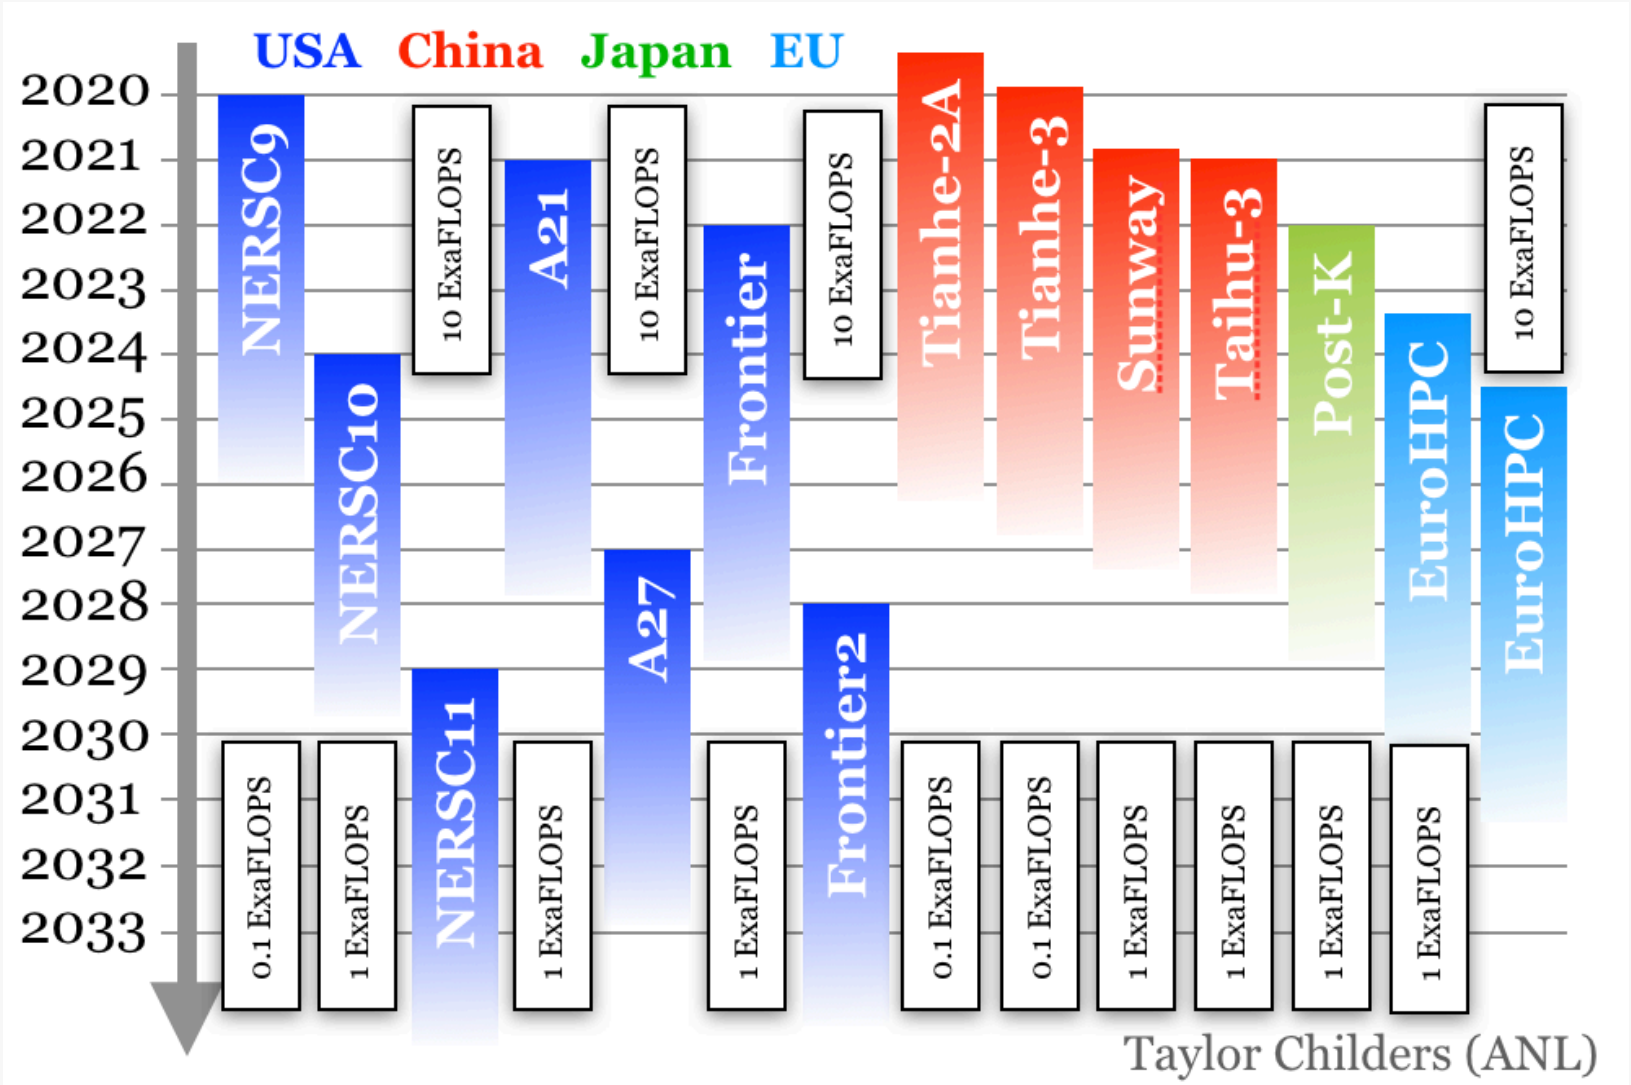

- US:
  - ~$600M/system + R&D
  - Intel & NVidia (ARM?)
- China:
  - $300-500M/system + R&D
  - ARM & x86 derivatives with accelerators
- Europe:
  - $350M/system + R&D
  - Aiming for 2 in Top5 and 4 in the Top25
  - New in-house Arch.
- Japan:
  - $800M/system
  - New in-house Arch.

https://www.youtube.com/watch?v=scTNGXyzdsg

Argonne NATIONAL LABORATORY

Future of HPC Around the World

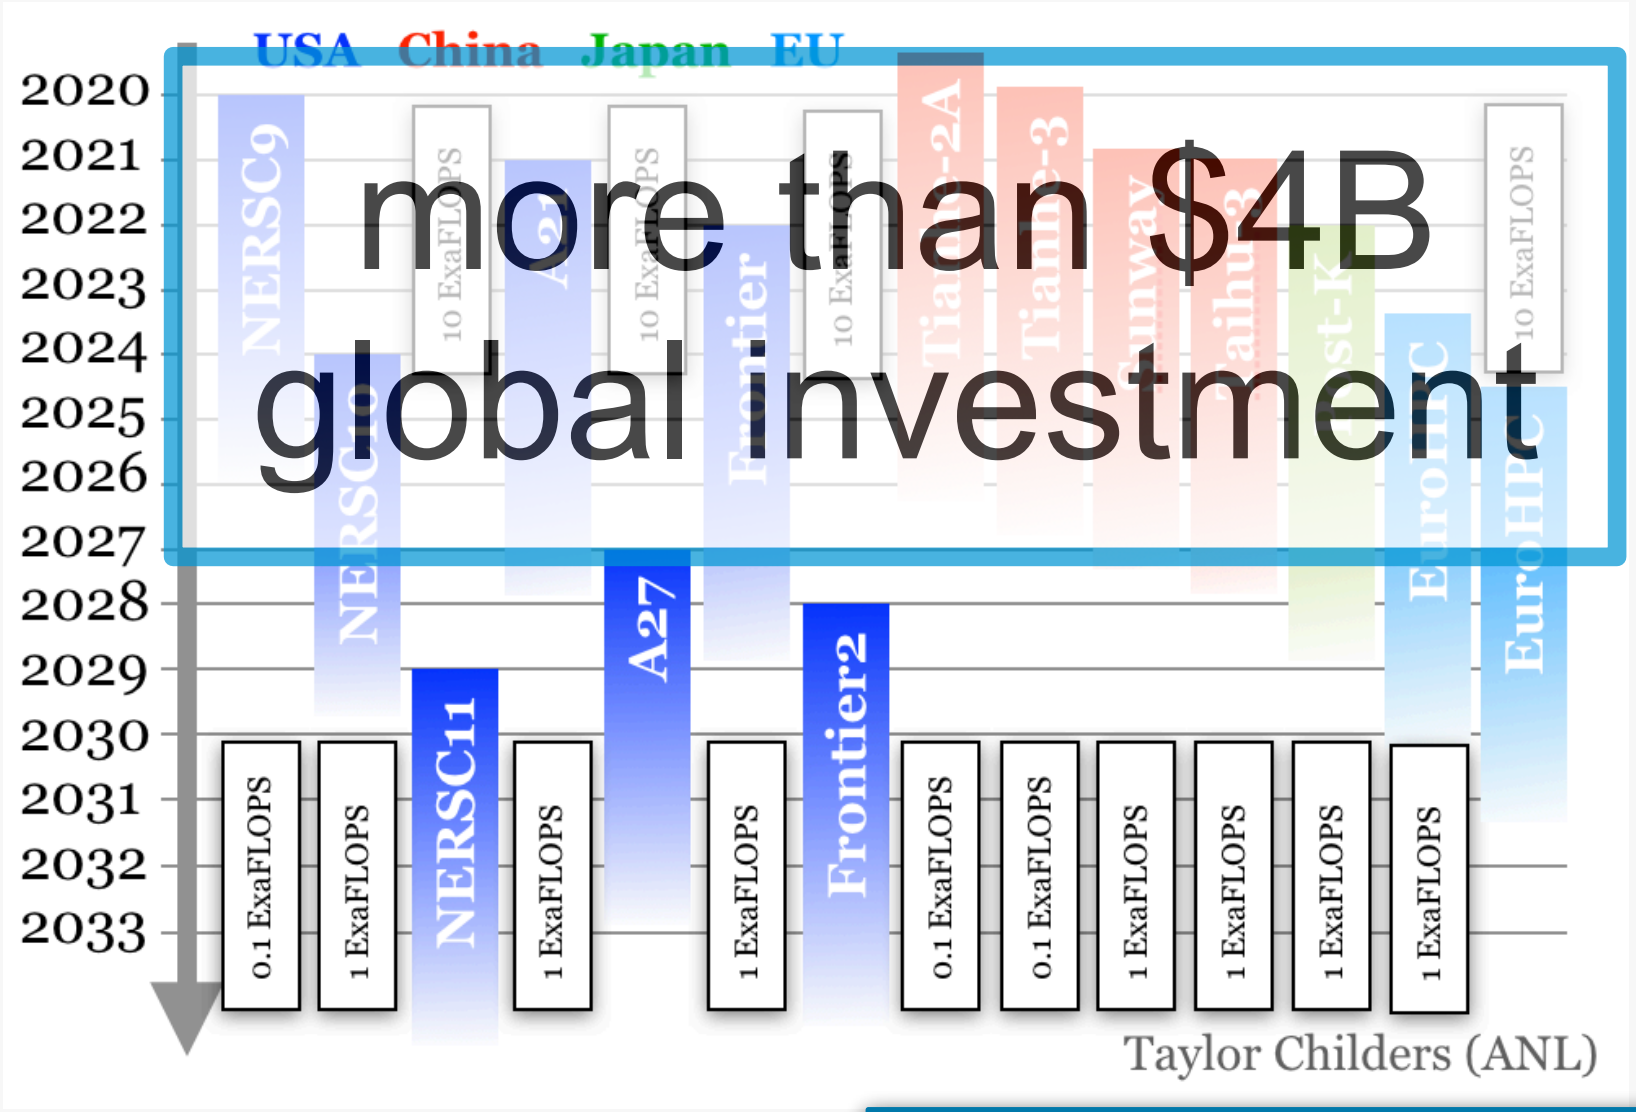

- US:
  - ~$600M/system + R&D
  - Intel & NVidia (ARM?)
- China:
  - $300-500M/system + R&D
  - ARM & x86 derivatives with accelerators
- Europe:
  - $350M/system + R&D
  - Aiming for 2 in Top5 and 4 in the Top25
  - New in-house Arch.
- Japan:
  - $800M/system
  - New in-house Arch.

https://www.youtube.com/watch?v=scTNGXyzdsg

Argonne NATIONAL LABORATORY

To Each Their Own

- Only thing certain about the future is the uncertainty.
- Only other thing certain about the future is every region will have their own home grown chips.
- While they may be x86 or ARM derivatives, it likely means no more x86 compatible binary execution so we must build our software for each system.
- These new architectures are being driven by the Machine Learning revolution, which is being driven by the data revolution that came before it.
- This means machines will continue to be made up of accelerators, like GPUs, many-core style, like KNLs, or some hybrid.
- Those who can represent their code as processing graphs, with small kernels operating at each step, will do well.

My Conservative Predictions for US

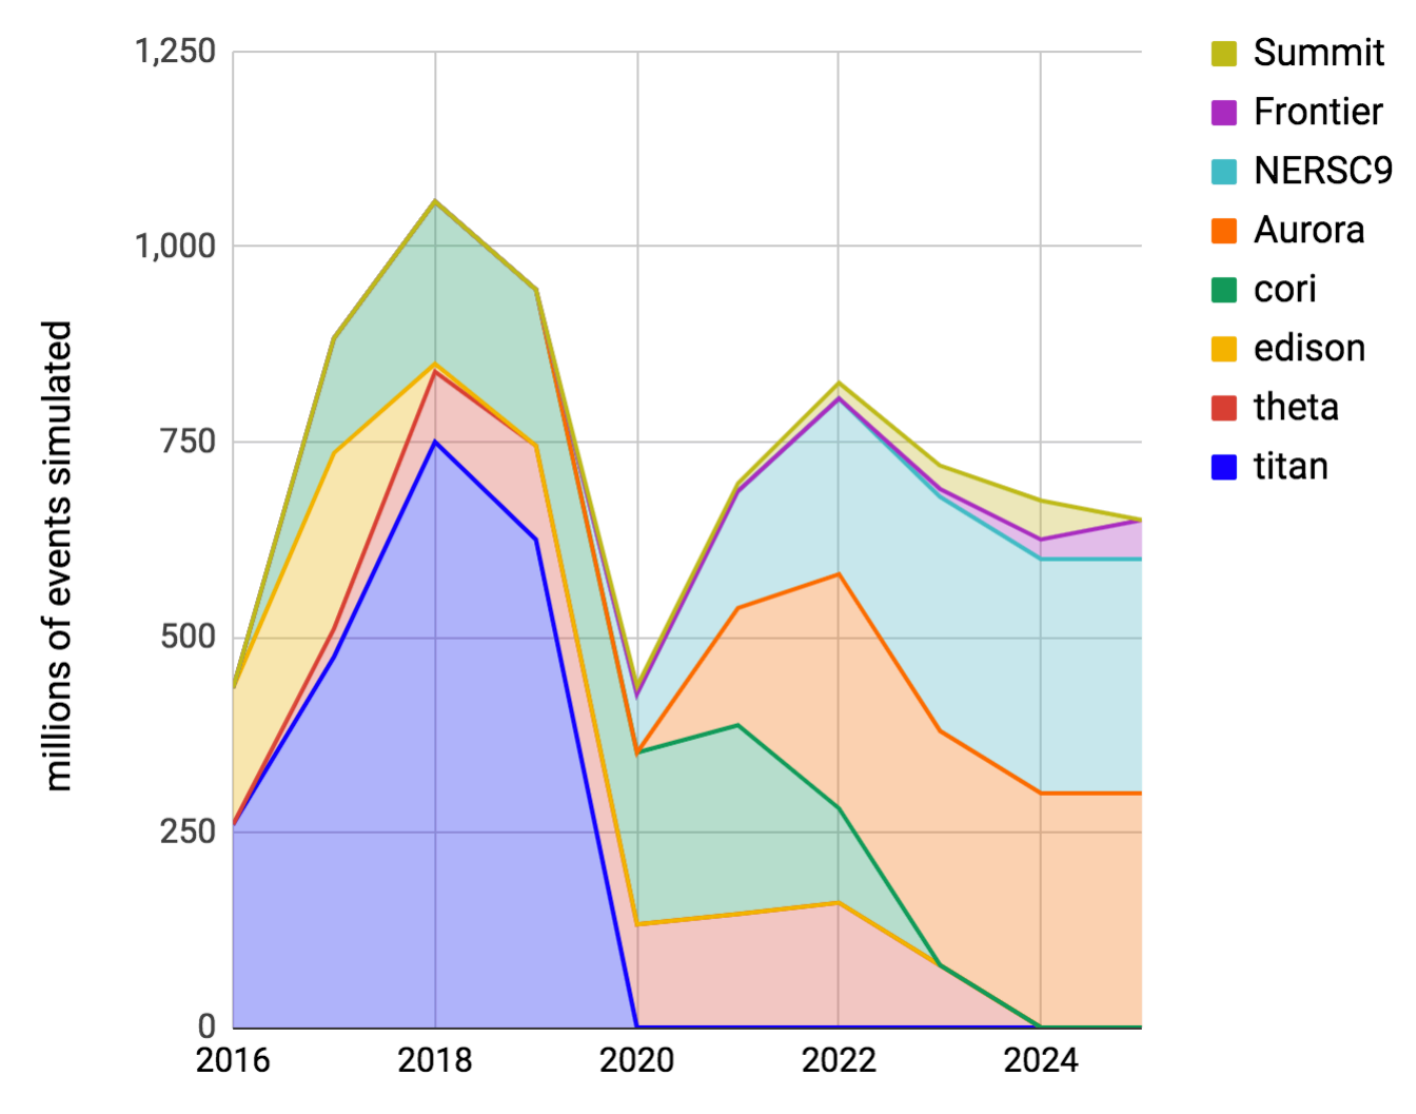

- Did not include TACC since we seem to be using small amounts and it was unclear to me if we will be getting a more permanent allocation.
- Did not include PizDaint for the same reason. I could not find it on Dashb or Grafana
- Left with the US machines.
- Perhaps people can send me more info and I can incorporate it into my projections.
- So lets go through this.

Argonne NATIONAL LABORATORY

My Conservative Predictions for US

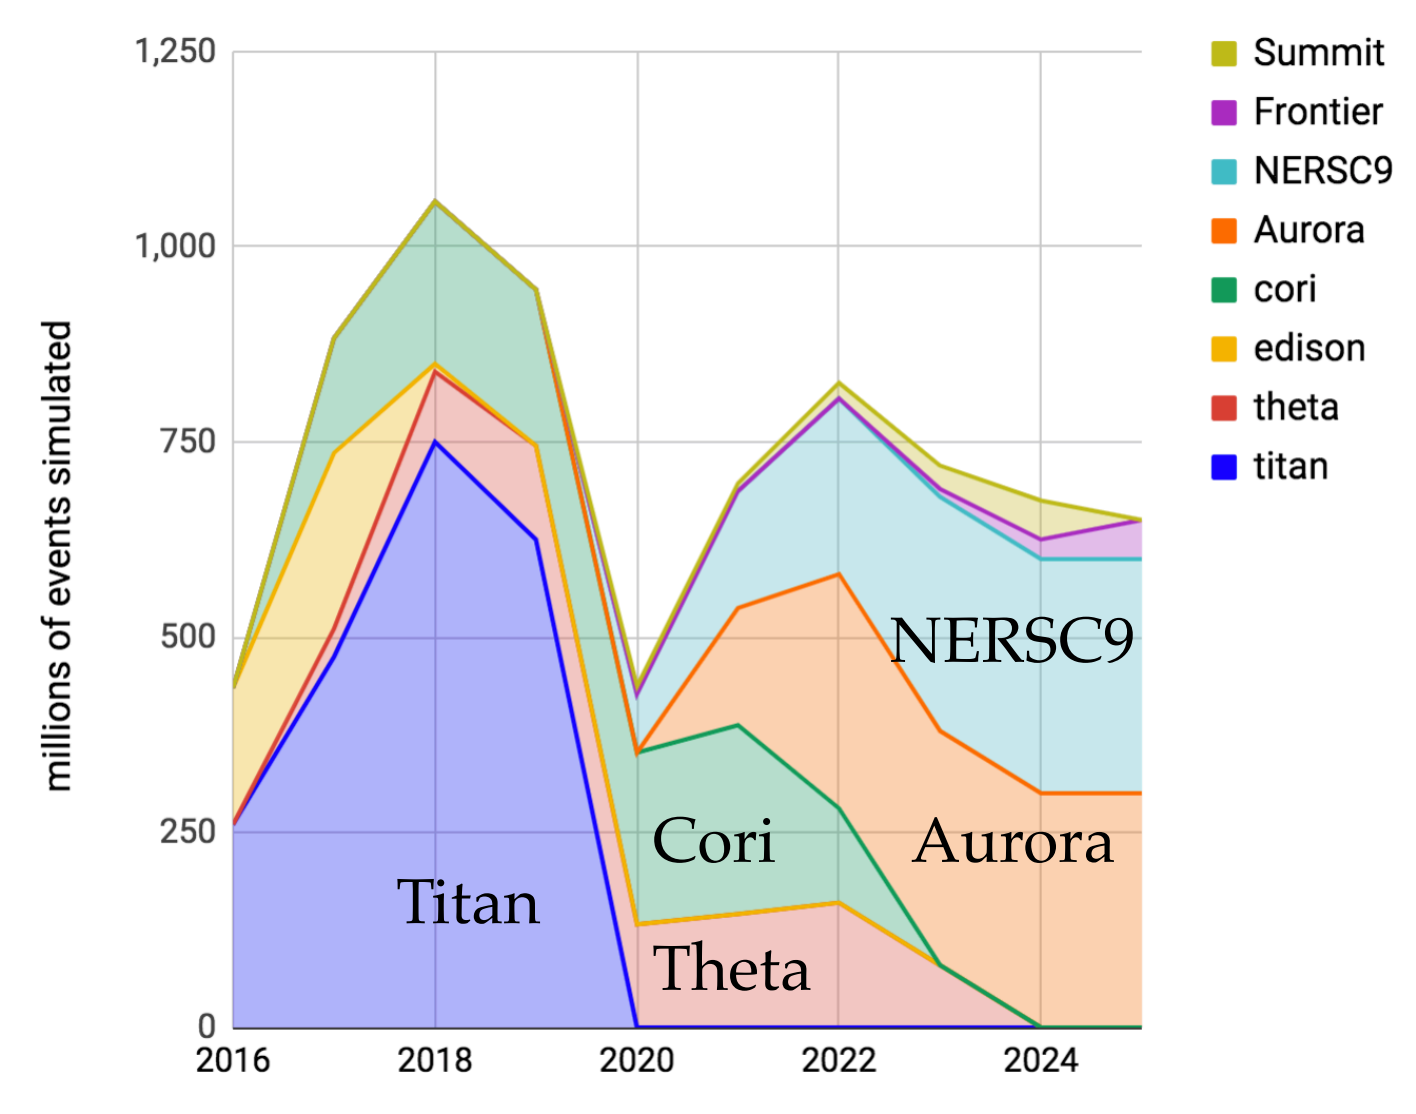

- I calculate the fraction of the total time available on a machine that we have been given historically and use that to calculate these numbers.
- Titan has been very productive, but now that Summit is online, will be decommissioned in the next year.
- Theta/Cori will have similar lifetimes, but I predict in this plot that Cori will still provide more hours for ATLAS. (not clear)

Argonne NATIONAL LABORATORY

My Conservative Predictions for US

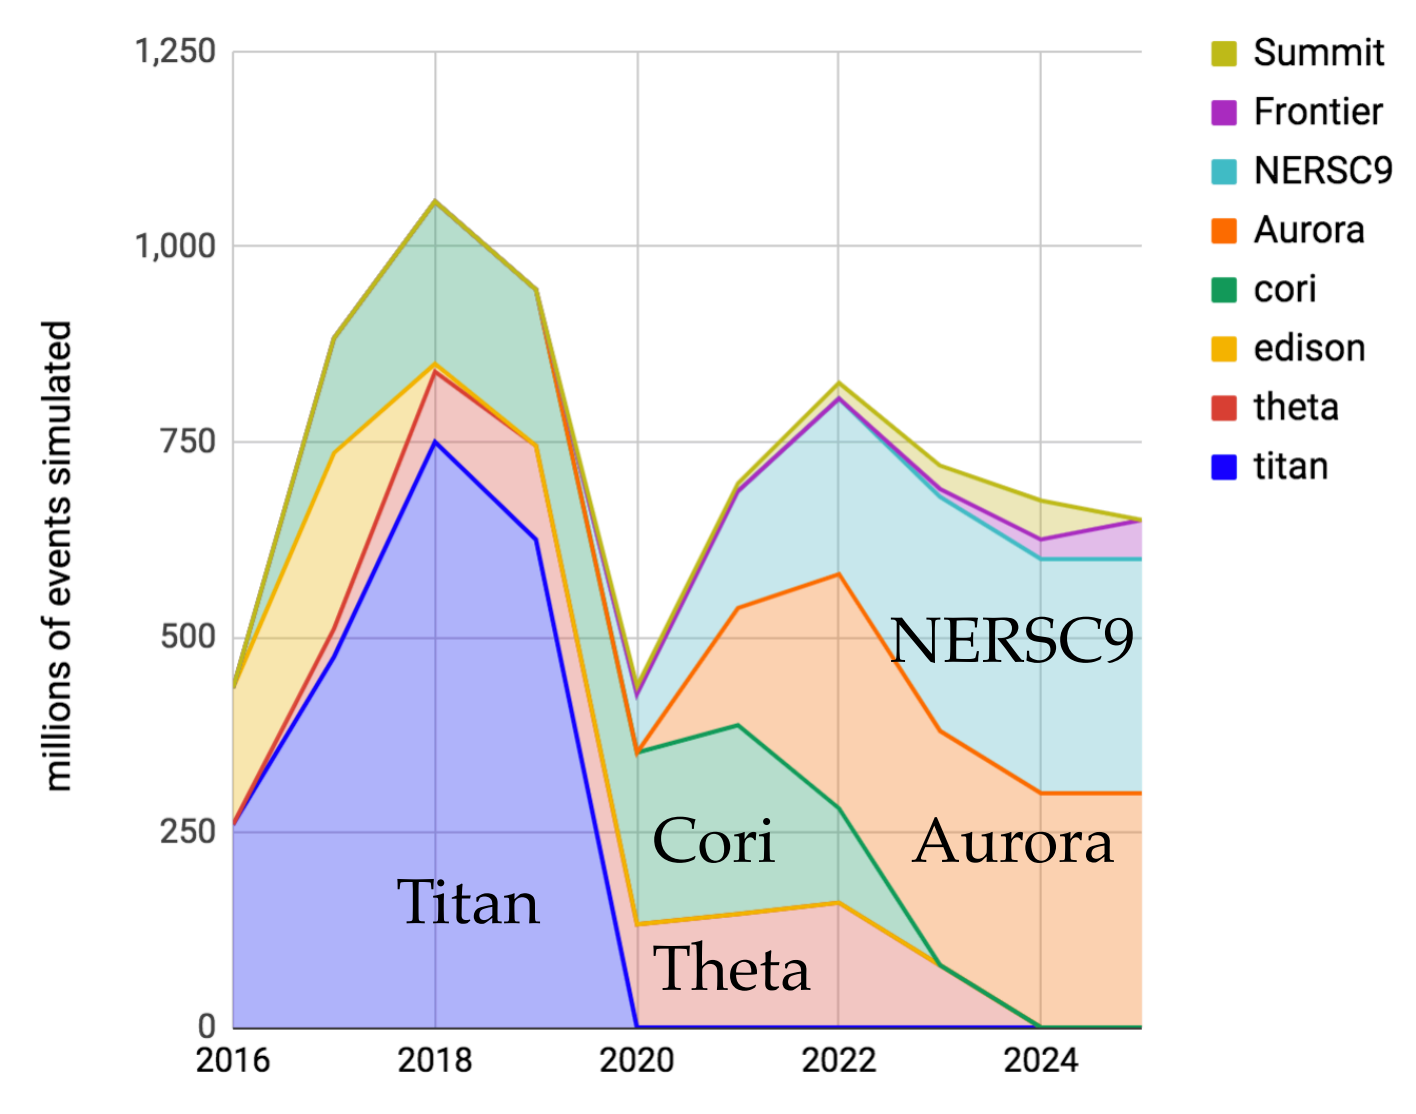

- With Cori/Theta I presume we will increase by 10% year over year either due to either more efficient running or more time.
- With NERSC9, I predict it will be a GPU machine which means we only have access to about 20% (or less) of the total capability. It's NERSC so we'd be given some time even if we don't use the GPUs.

Argonne NATIONAL LABORATORY

My Conservative Predictions for US

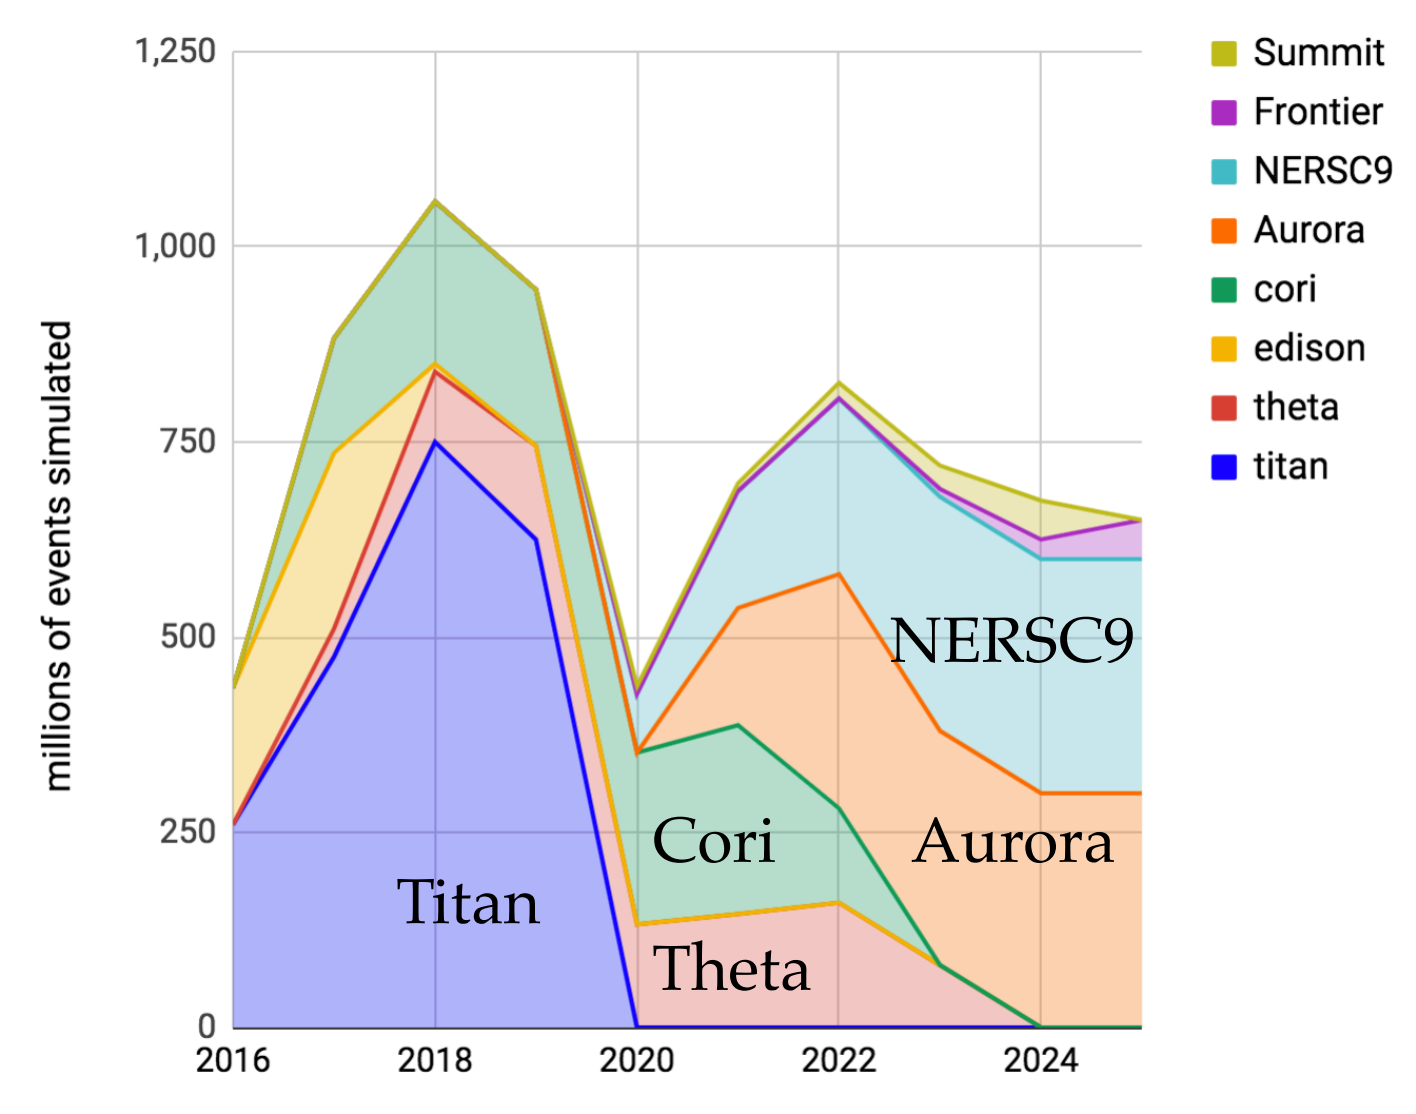

- For Aurora & NERSC9 I presume we get an event rate that is better than Xeon Phi but not as good as Xeon.
- For Summit & Frontier, I predict we will get very little time if any at all since 90% of their power will be in GPUs.

Argonne NATIONAL LABORATORY

Silver Lining

- Predicting the future is like a box of chocolates.
- Lots of R&D work ongoing for ML in ATLAS, offloading any significant portion of the Reco onto ML algs will open the door to accelerators.
- This plot does not include EuroHPC, CSCS, Chinese, or Japanese. If we can deploy our tools at these sites, and install software we could more than double the capacity on that plot.