

Panel: Saudi TopXY

- Context of the Top500 international list
 - courtesy of Erich Strohmeier, who curates the Top500 list, along with Horst Simon and Jack Dongarra
- Lesson from the China Top100
- International context of TopXY lists
- Motivations for creating and curating a recognized, consensual Saudi TopXY list


TOP500 Roots – End of 1980s

- Used Linpack as measure
 - Simplistic, but widely available
- Self-adapting over time
- Chose 500 to be “comfortably inclusive”
 - Easy to keep filled
 - Not leaving any major players out


TOP500

- After first list, settled on publishing twice a year
 - ISC'xy in June in Germany
 - SCxy in November in the U.S.
- Absolute data is less important than *trends over time*
- All information available from the TOP500 web site at: www.top500.org

History of HPL performance



No. 1's – HPL R_max



China TOP100

- In 2002, China started its first national “Top 50”
- Fifteen years later, China has more supercomputers in the international Top500 than any other country
 - the sum of China’s top two supercomputers outranks the sum of the top ten of every other country (not claiming cause 😊)
- China's national list instantly attracted attention internationally
 - including mine as a regular collaborator

Other national TOP lists

- India has maintained such a list since 2008
- Ireland has done so since 2013
- These lists currently show, respectively, 13 and 9 systems with an R_{\max} for HPL above 100 Tflop/s
 - Both of these countries have less to track than Saudi Arabia
- Encouraging: we would start off impressively
- Discouraging: patiently tracking the national growth in supercomputing does not "automatically" generate growth

Why a national Saudi list?

- Give visibility to the importance of supercomputing
 - Kingdom can accelerate its economic transformation
 - “To outcompute is to outcompete”
 - Supercomputing leads to workforce sophistication in the form of data analytics and simulation
- Track national progress in Saudi supercomputing
 - regular attention to such a list, at annual intervals such as every meeting of HPC Saudi, will inevitably help progress.
- Each new system for *any* player buoys *all* players
- You cannot easily improve what you cannot reliably measure

Why a national Saudi list?

- Saudi Arabia is in an unusually opportunistic position
- Major national projects emphasize transition to an information economy
- Large number of industries and agencies are not *yet* very fully exploiting the data that they are sitting on
 - but know that they should be

Saudi TOPXY

- What is “XY”
 - 100? 25? 10?
- The number has to be large enough that many organizations find it within range to climb on near the bottom now and then aspire to climb up over time. It must be small enough that the work to verify and update it is manageable
- We can start small and grow it
- Let’s take advantage of coming later and be more interesting

What it would look like today

Rank	Name	Manufacturer	Year	Segment	Rmax [TFlop/s]	Rpeak [TFlop/s]	Specification	Site
1	Shaheen II	Cray Inc.	2015	Academic	5536.99	7235.17	Cray XC40, Xeon E5-2698v3 16C 2.3GHz, Aries interconnect	King Abdullah University of Science and Technology
2	Makman-2	Dell EMC	2015	Industry	2249.68	3041.28	Dell PowerEdge R630, Xeon E5-2680v3 12C 2.5GHz, Infiniband QDR	Saudi Aramco
3	Faris	HPE	2014	Industry	816.58	917.50	Cluster Platform SL230s Gen8, Intel Xeon E5-2680v2 10C 2.8GHz, Infiniband QDR	Saudi Aramco
4	Dammam	Dell EMC	2017	Industry	694.33	1074.13	Dell PowerEdge R630, Xeon E5-2695v4 18C 2.1GHz, 10G Ethernet	Saudi Aramco
5	SANAM	Adtech	2012	Research	532.60	1098.00	Adtech, ASUS ESC4000/FDR G2, Xeon E5-2650 8C 2.000GHz, Infiniband FDR, AMD FirePro	King Abdulaziz City for Science and Technology
6	Makman	IBM/Lenovo	2013	Industry	441.82	506.19	iDataPlex DX360M4, Xeon E5-2670 8C 2.600GHz, Infiniband QDR	Saudi Aramco
7	Aziz	Fujitsu	2015	Academic	211.30	228.56	Fujitsu PRIMERGY CX400, Intel Xeon E5-2695v2 12C 2.4GHz, Intel TrueScale QDR	King Abdulaziz University
8	??							
9	??							
10	??							

What benchmark?

- Top500 is based on HPL: High Performance LINPACK
 - A dense matrix LU-factorization and backsolve
 - Great to show off flop/s: does N^3 flops in N^2 memory
 - Trend correlates with Gordon Bell Prize over time
 - But, not representative of mainstream simulation-based supercomputing
 - Further, getting the highest number takes a long time to run
 - can be several days on the world's highest memory systems
- There are many other benchmarks possible

Other benchmarks

- Candidates
 - HPL
 - HPCG
 - HPGMG
 - Green
 - Graph
 - I/O
- Could attempt some composite on one list
 - Reservoir simulation, seismic inversion, some analytics
 - Exhibit “scientific taste” and intellectual leadership!