Dennard, Carbon & Moore

Scenarios for the Future of NSF Advanced Computational Infrastructure

Andrew A. Chien, University of Chicago, Argonne National Laboratory
Mark J. Dietrich, Bloodstone Solutions Inc.

How can we satisfy continued high growth in scientific computing demand?

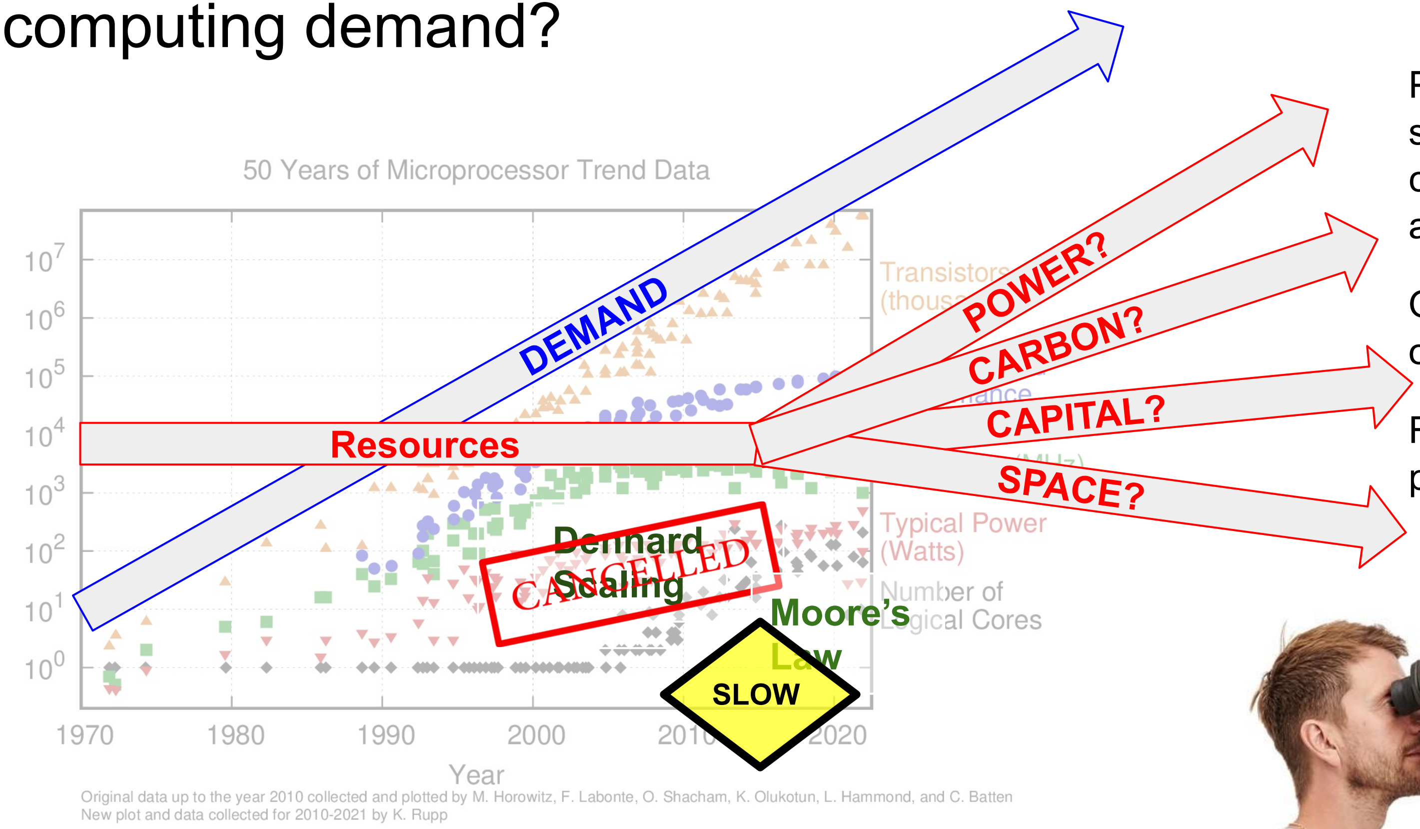

Perspectives of system operators, as consortia, as funding agencies?

Quantitative change or disruption?

Replacement cycles, partial upgrades?

Approach and Data

Model performance and total cost of ownership:

- Capital; Rpeak; Power/power costs/CO_2 (Scope 2); Space, space costs

Calibrate with a range of real systems

- Wide scale (400:1 Rpeak, 100:1 kW)
- Representative of Top100

Construct scenarios around investment objectives:

- **I0**: Constant (annualized) capital investment
- **I1**: Achieve 40% annual growth in capacity

System	Site	In Service (Mo/Yr)	Rpeak (Gflops/s)	Compute Power (kW)
Jetstream	Indiana	1/2016	516,096	140.0
Big Red II+	Indiana	8/2016	286,157	90.0
Big Red 3	Indiana	8/2019	928,512	179.0
Comet	SDSC	7/2015	2,831,699	550.9
Expanse	SDSC	7/2020	5,078,656	407.4
Stampede 2	TACC	9/2017	18,394,522	2,200.0
Frontera	TACC	7/2019	40,977,504	3,407.1
Wrangler	TACC	7/2015	92,160	28.8
Chameleon	TACC	6/2018	370,944	100.8
Bridges	PSC	1/2016	2,131,341	283.4
Bridges 2	PSC	10/2020	2,957,069	257.0

Average Site	Power Cost/kWyr	kg Carbon/kWyr
Campus (Site 9)	$764.34	3,829
Green Data Center (Site 8)	$477.12	-0-

Trends and Scenarios

I* Investment Objective, CD* Compute Density, PD* Power Density, DC Data Center

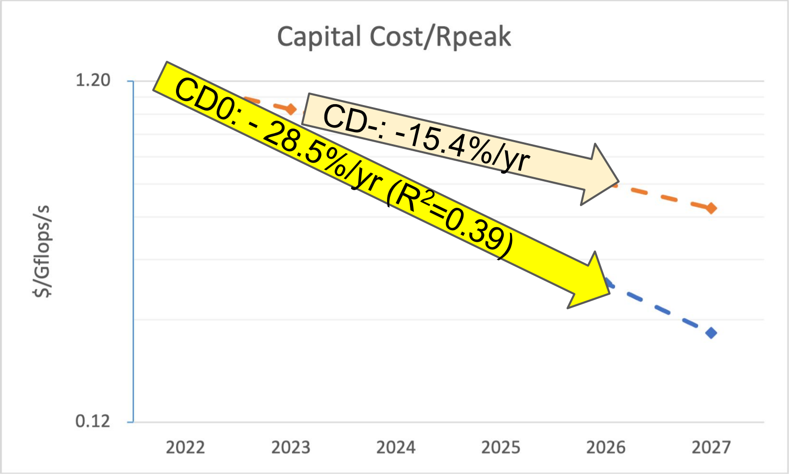

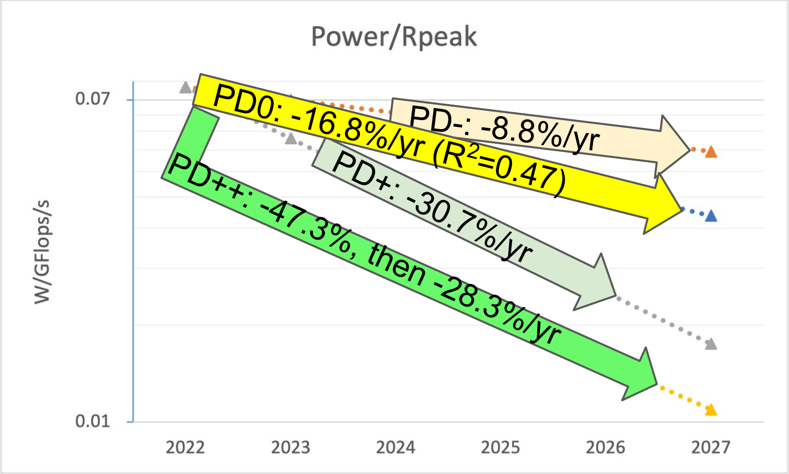

Trend from Sample → Conservative → Better → Optimistic

24* Scenarios = { I0, I1 } × { CD-, CD0 } × { PD-, PD0, PD+, PD++ } × { Campus DC, Green DC }

* I0|CD0 = I1|CD0

Baseline Scenario: Power Overwhelms Capital by 2037

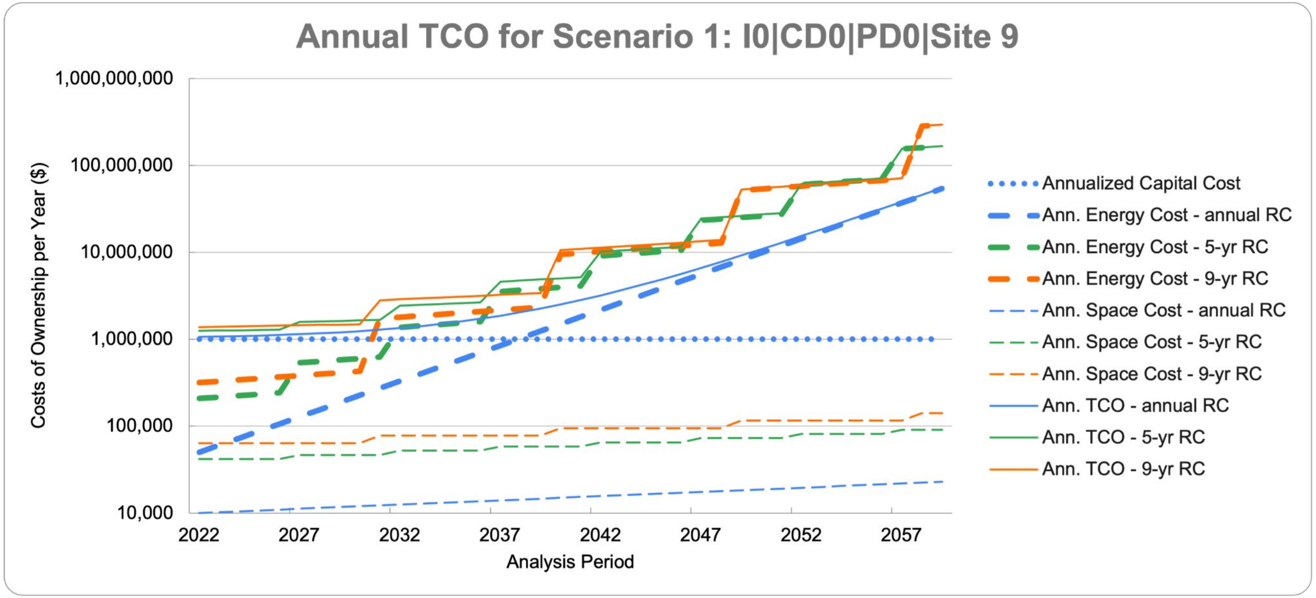

- **I0**: Constant capital investment of $1M/year (annualized for replacement cycles (RC) >1 year)
- Replacements fix power/space density at each commissioning – penalizing older systems
- Annual power costs > annualized capital costs, dominating Annual TCO, by 2037.
- 2027 number
- Space costs << capital costs
- NB: Log scale for all costs

24 Scenarios: Significant Differences

Power density:	Baseline (PD0)	Conservative (PD-)	Better (PD+)	Optimistic (PD++)
Measured Compute Density trend continues (CD0); investment stays constant (I0)	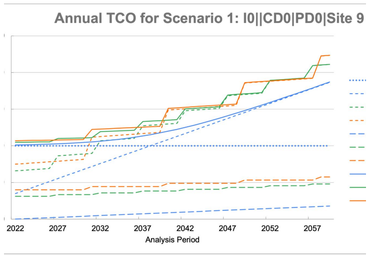	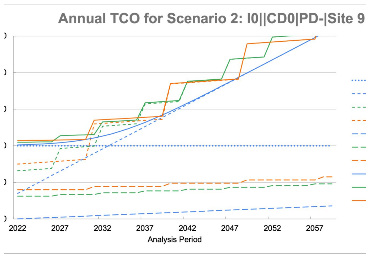	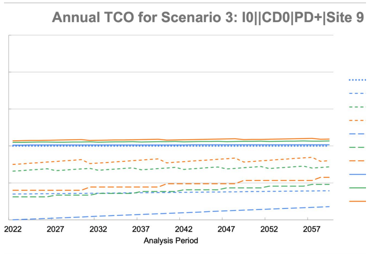	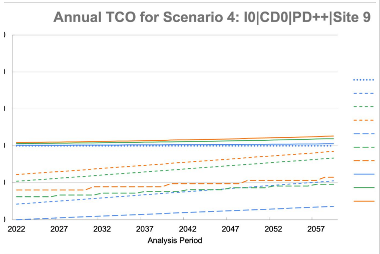
	Not shown: Compute density slows and investment stays constant			
Compute density trend slows 2x (CD-), investment increased to maintain growth of flops/s (I1)	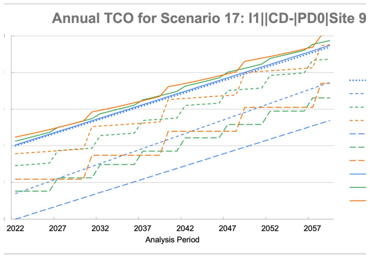	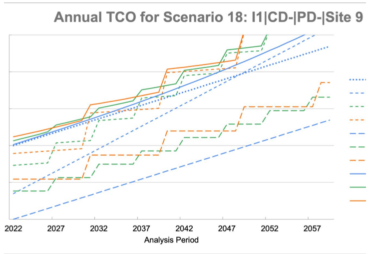	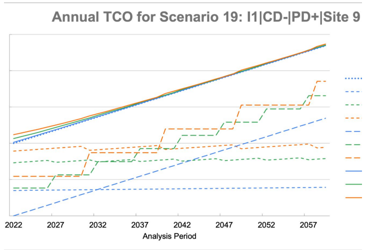	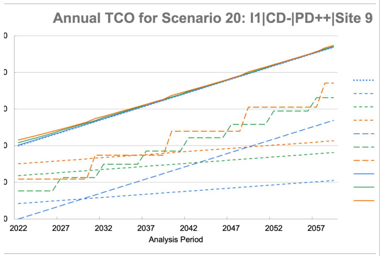

- Legend: same as prior chart, log scale for costs
- Not shown: Scenarios for green data centres

Two Key Patterns

Power density:	Baseline (PD0)	Conservative (PD-)	Better (PD+)	Optimistic (PD++)
Measured Compute Density trend continues (CD0); investment stays constant (I0)	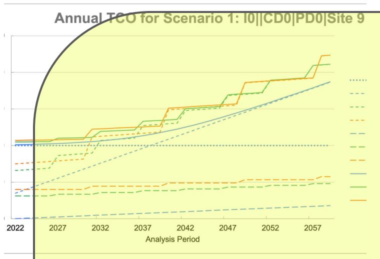	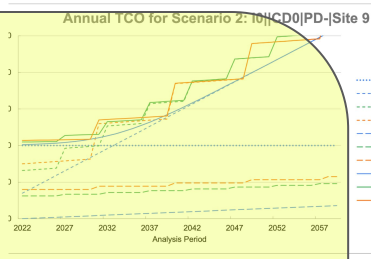	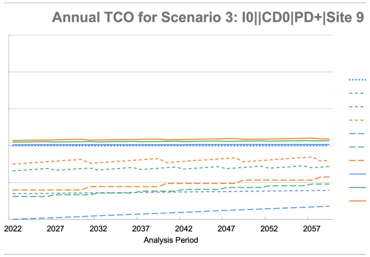	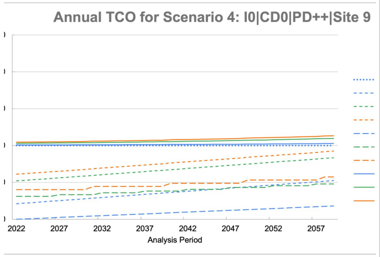
	Not shown: Compute density slows and investment stays constant			
Compute density trend slows 2x (CD-), investment increased to maintain growth of flops/s (I1)	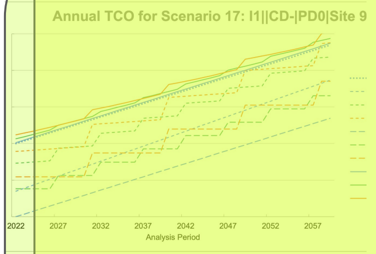	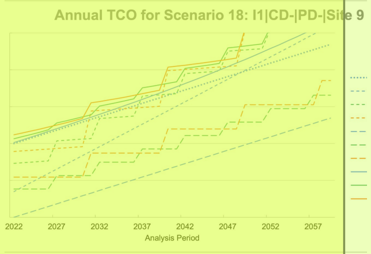	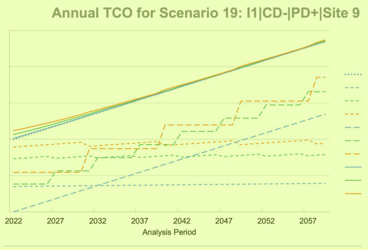	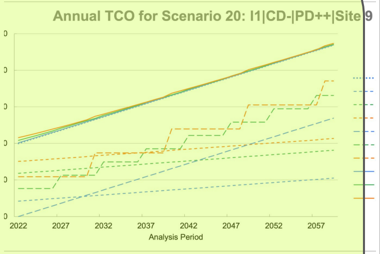

Capital costs grow 12x by 2037

Power costs grow 10-40x by 2037

Capital/Space Explode 12x by 2037 if Compute Density Slows

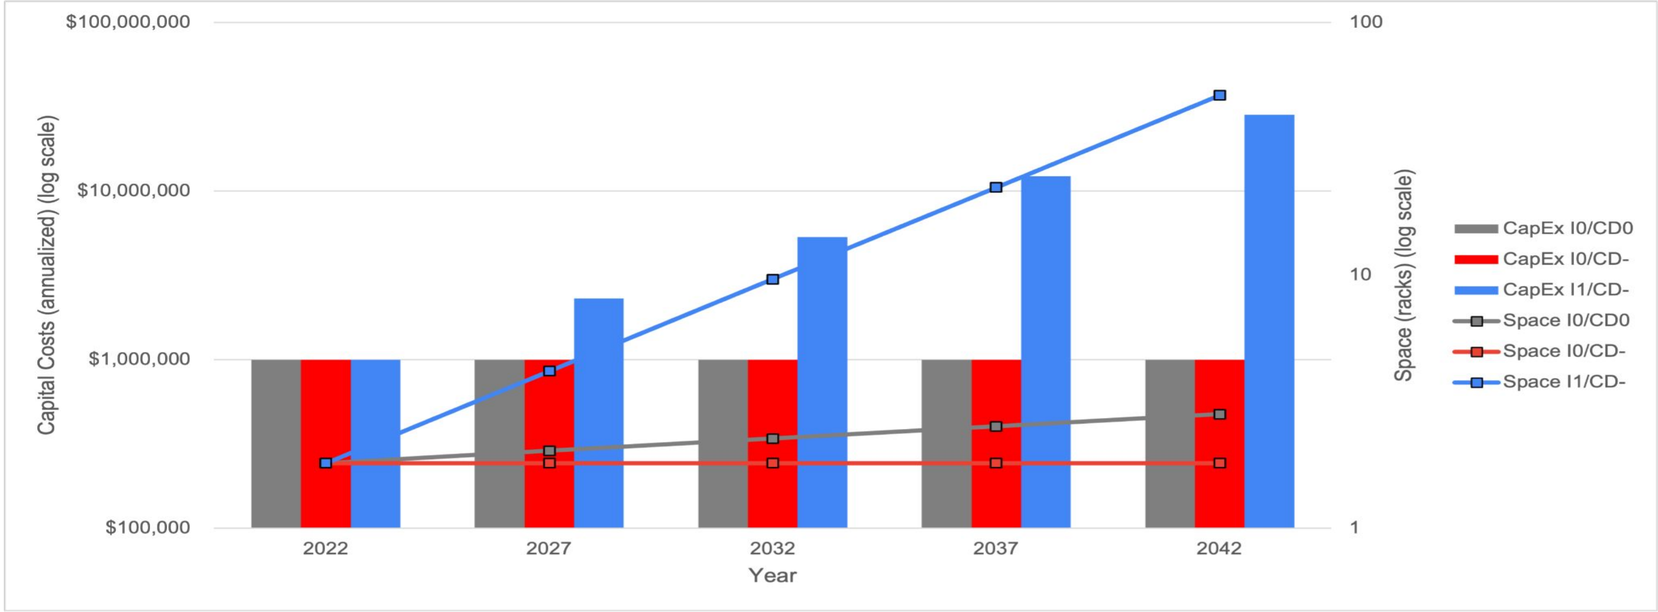

Increase 5x by 2027

…if we increase investment to grow capacity to meet demand → Consider new hosting options?

Otherwise capital & space requirements stay under control.

Power Overwhelms Capital in 10-15 Yrs in Many Scenarios

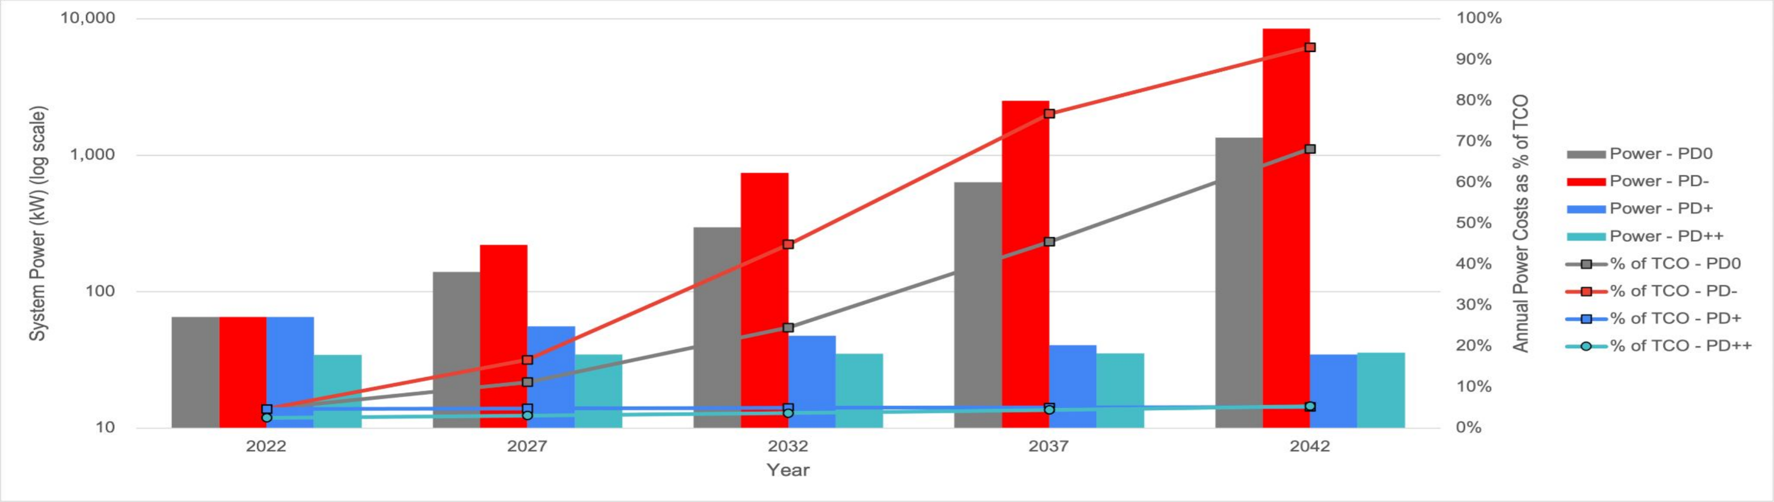

Increase 2-3x, accounting for 11-17% of TCO, by 2027

More kW → Higher Power costs (absolute and as % of TCO), More CO_2

→ New data center designs, new cooling tech

→ Need to consider green options

Slower Compute Density Trend Extends Replacement Cycles

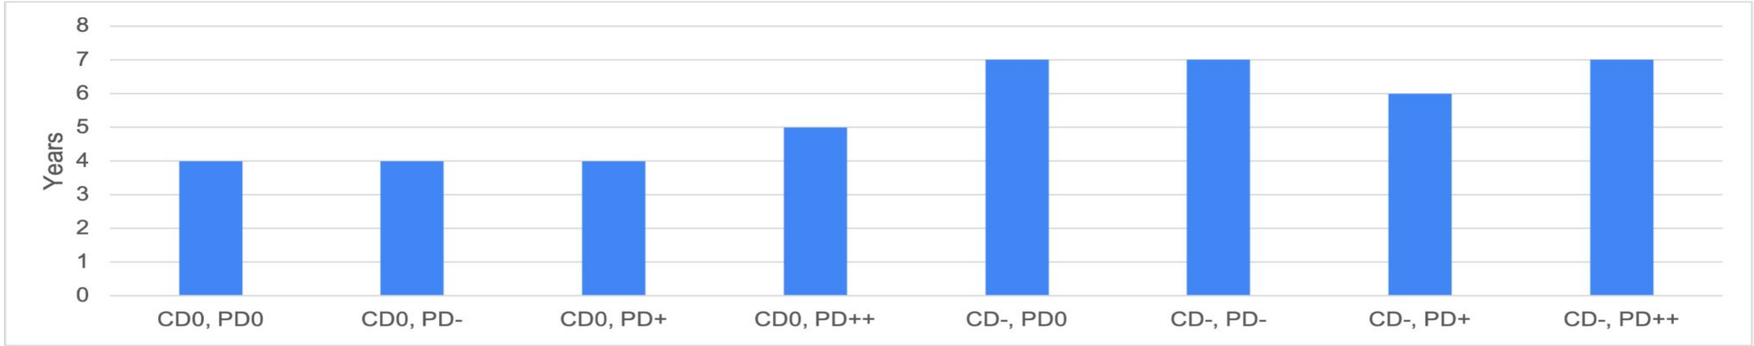

- Using systems through their full “economic life” optimizes TCO ($) per unit capacity
- Slower compute density/performance improvement trends extend replacement cycles - optimizing TCO/capacity

Power and Carbon Mitigation Need to Start Soon: Green Data Centers Offer Solutions

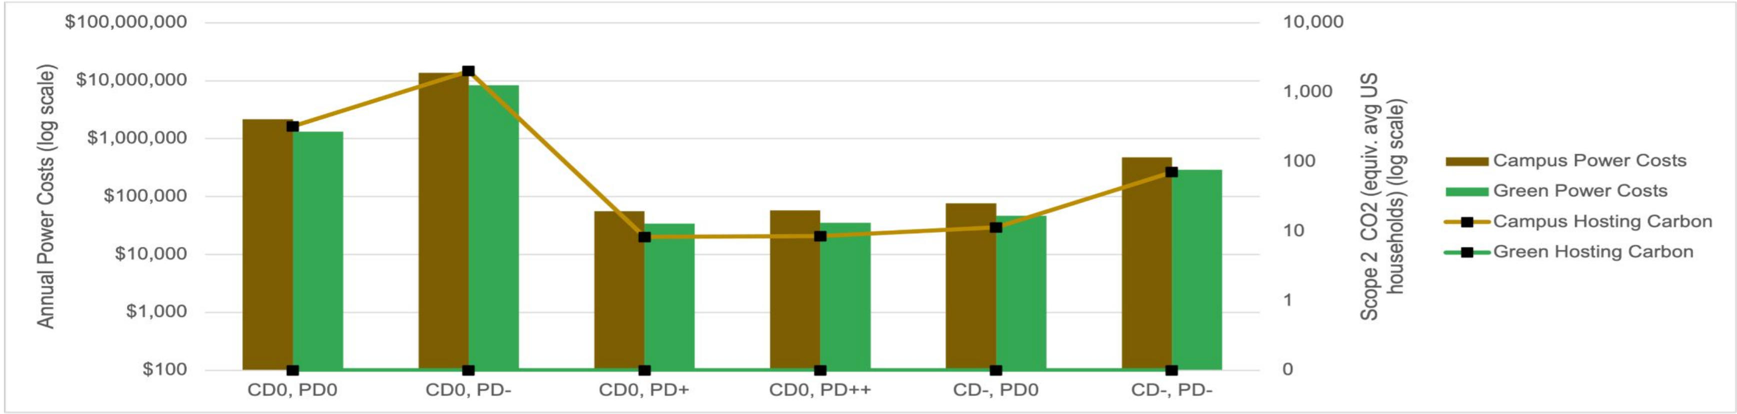

- <u>Eliminate</u> Scope 2 GHG emissions
- Reduce costs by 37% exploiting stranded power

Key Conclusions

Capital & Hosting Space Explode 28x in 20 Yrs if Compute Density Trend Slows

Power Overwhelms Capital in 10-15 Yrs if Power Density Trend Continues (or slows)

Slowing Compute Density Trend Lengthens Optimum Replacement Cycle

Green Data Centers Mitigate Extreme Impacts of Carbon and Power Costs

Link to short paper: https://dl.acm.org/doi/10.1145/3491418.3535156

Link to more detailed analysis: https://newtraell.cs.uchicago.edu/research/publications/techreports/TR-2022-04

BOF at PEARC'20: http://people.cs.uchicago.edu/~aachien/lssg/research/zccloud/lifetime/bof-index.html

CASC survey: https://newtraell.cs.uchicago.edu/research/publications/techreports/TR-2020-05

Thank you!
Questions?

Andrew: achien@cs.uchicago.edu
Mark: mjdietrich@bloodstonesolutions.org

Thanks to our research sponsors! NSF CMMI-1832230 and OAC-2019506, as well as VMWare, Google, and Intel.

Thanks to the Zero-carbon Cloud group!

Baseline Scenario: Projecting Key Parameters

Year	2022	2027	2032	2037	2042	20-year Growth
Capital Investment	$1,000,000	$1,000,000	$1,000,000	$1,000,000	$1,000,000	1.00x
Petaflops/s (PF)	0.9	4.5	24.2	129.1	688.3	809x
System Power (kW)	65	139	297	634	1,352	20.7x
Cost of power ($/kWyear, 1st year of operation, average campus site)	$764	$921	$1,110	$1,337	$ 1,612	2.11x
Power costs (annual, average campus site)	$49,907	$128,272	$329,689	$847,378	$2,177,955	43.6x
Scope 2 CO2 Emissions (MT, annual, average campus site)	250	533	1,137	2,426	5,175	20.7x
System Size (Racks)	1.8	2.0	2.3	2.5	2.8	1.56x
Space costs (annual, average campus site)	$10,026	$11,208	$12,529	$14,005	$15,656	1.56x
TCO (annual)	$1,059,933	$1,139,480	$1,342,218	$1,861,383	$3,193,611	3.01x
TCO/PF	$1,245,840	$251,129	$55,465	$14,422	$4,640	0.00372x

Power density in the Top500: GPU-heavy systems significantly better on W/Gflops/s

N=2,235 unique systems

Trend: -26-27%/yr, R^2>.078

- except for "hiCPU" (CPUs account for 14-99.9% of total Rpeak): -19%/yr

W/Rpeak parallels W/Rmax