On the Path to Energy-Efficient Exascale: A Perspective from the Green500

Wu FENG, wfeng@vt.edu

Dept. of Computer Science

Dept. of Electrical & Computer Engineering

Health Sciences

Virginia Bioinformatics Institute

THE GREEN 500

SyNeRG

synergy.cs.vt.edu

The Ultimate Goal of "The Green500 List"

- Raise awareness of energy efficiency in the supercomputing community
  - Drive energy efficiency as a first-order design constraint (on par with performance)
  - Encourage fair use of the list rankings to promote energy efficiency in high-performance computing systems.

On the Importance of the Green500

- K Computer
  - Power & Cooling: 12.66 MW → $12M/year

High-Speed Train
10 MW

- Google in *The New York Times*, June 14, 2006

Hiding in Plain Sight, Google Seeks More Power

Melanie Conner for The New York Times

Google is building two computing centers, top and left, each the size of a football field, in The Dalles, Ore.

POWER
First-Order Design Constraint in Data Centers

Google Details and Defends Its Use of Electricity

"Google disclosed Thursday that it continuously uses enough electricity to power **200,000** homes."

The New York Times, September 18, 2011

Agenda

- The Green500 List: A Brief History
  - From Green Destiny to the Green500 List
- Analysis of Green500 Lists
  - Energy Efficiency
- Projections for Energy-Efficient Exascale
  - Tracking Koomey's Law for High-Performance Computing (HPC)
  - Quantifying *Distance* from Exascale Goal: The Exascalar Metric

Brief History:

From Green Destiny to the Green500 List

- 02/2002: Green Destiny (http://sss.lanl.gov → http://sss.cs.vt.edu)
  - "Honey, I Shrunk the Beowulf!"
    31st Int'l Conf. on Parallel Processing, August 2002

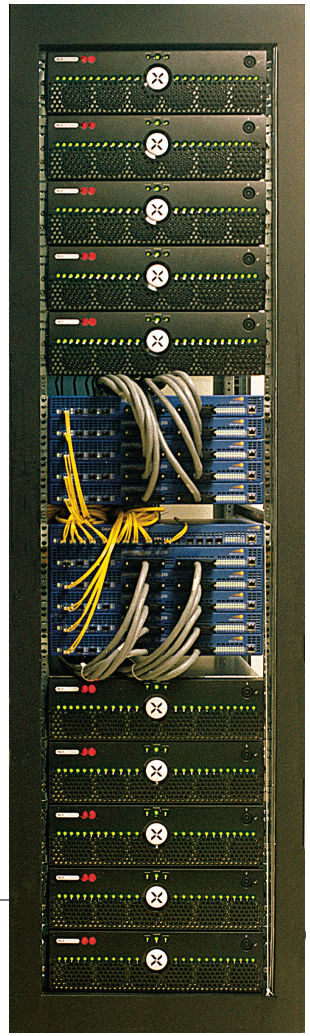

THE GREEN 500

Brief History:

From Green Destiny to the Green500 List

- 02/2002: Green Destiny (http://sss.lanl.gov → http://sss.cs.vt.edu)
  - "Honey, I Shrunk the Beowulf!"
    31st Int'l Conf. on Parallel Processing, August 2002
- 04/2005: Keynote Talk by W. Feng at the *IEEE Workshop on High-Performance, Power-Aware Computing*
  - Generates initial discussion for Green500 List
- 04/2006 and 09/2006: Making a Case for a Green500 List
  - Workshop on High-Performance, Power-Aware Computing
  - Jack Dongarra's CCGSC Workshop "The Final Push"

Brief History:

From Green Destiny to the Green500 List

- 02/2002: Green Destiny (http://sss.lanl.gov → http://sss.cs.vt.edu)
  - "Honey, I Shrunk the Beowulf!"
    31st Int'l Conf. on Parallel Processing, August 2002
- 04/2005: Keynote Talk by W. Feng at the *IEEE Workshop on High-Performance, Power-Aware Computing*
  - Generates initial discussion for Green500 List
- 04/2006 and 09/2006: Making a Case for a Green500 List
  - Workshop on High-Performance, Power-Aware Computing
  - Jack Dongarra's CCGSC Workshop "The Final Push"
- 09/2006: Launch of Green500 Web Site and RFC
  - http://www.green500.org

Brief History:

From Green Destiny to the Green500 List

- 11/2007: First official Green500 list released

Brief History:

From Green Destiny to the Green500 List

- 11/2007: First official Green500 list released
- 11/2010: First official Green500 run rules released

Brief History:

From Green Destiny to the Green500 List

- 11/2007: First official Green500 list released
- 11/2010: First official Green500 run rules released
- 06/2011: Collaborations begin on standardizing metrics, methodologies, and workloads for energy-efficient parallel computing
  - Energy-Efficient High-Performance Computing Working Group (EE HPC WG)
  - The Green Grid
  - TOP500
  - Green500

Agenda

- The Green500 List: A Brief History
  - From Green Destiny to the Green500 List
- Analysis of Green500 Lists
  - Energy Efficiency
- Projections for Energy-Efficient Exascale
  - Tracking Koomey's Law for High-Performance Computing (HPC)
  - Quantifying *Distance* from Exascale Goal: The Exascalar Metric

Analysis of Green500 Lists: Energy Efficiency

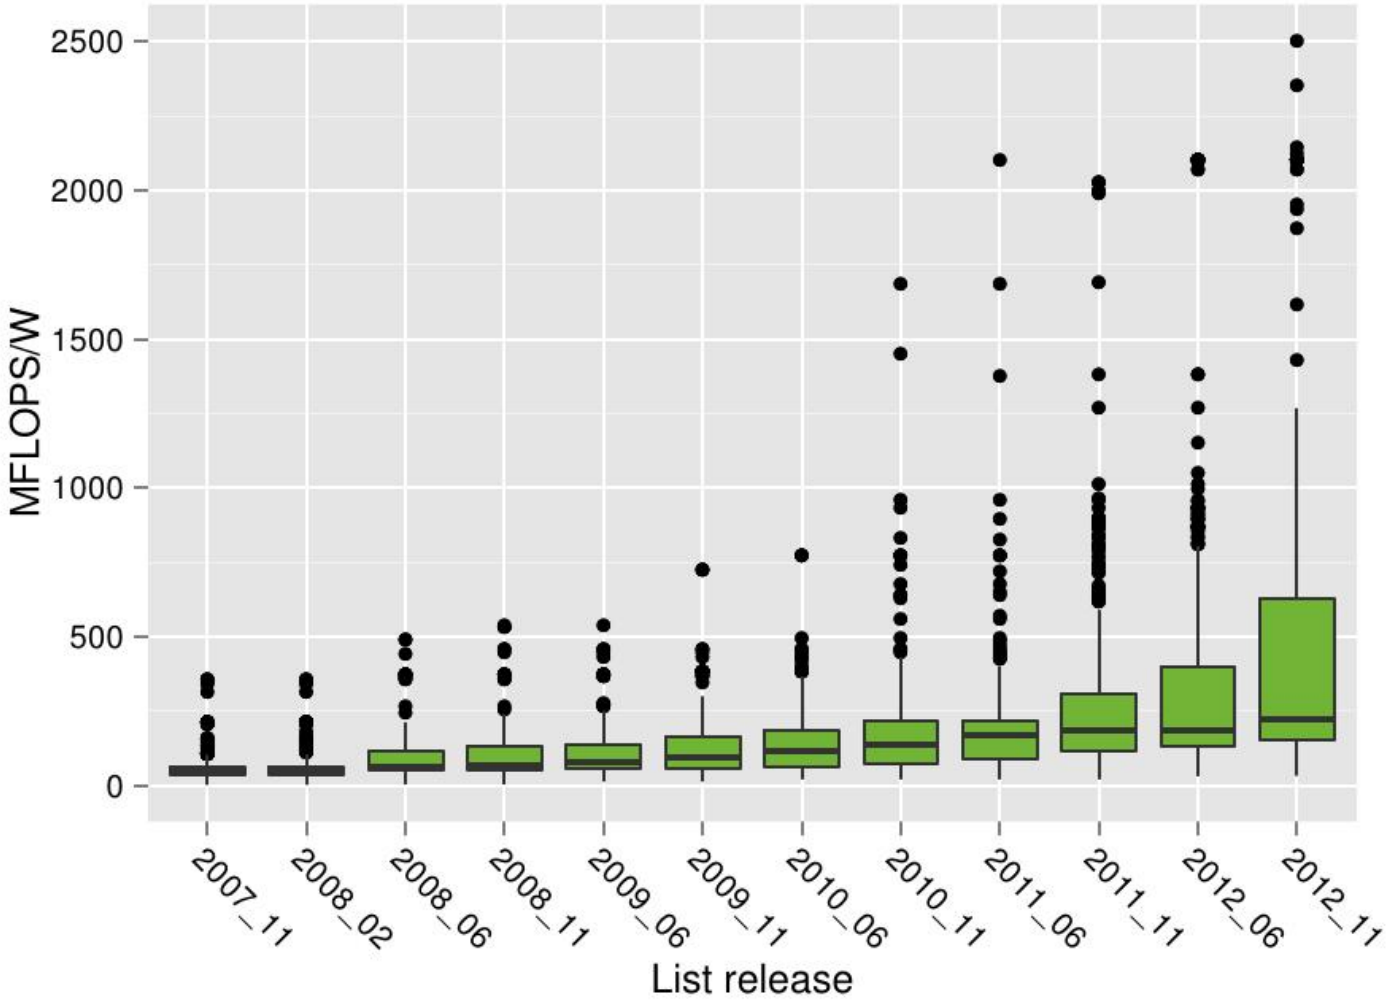

Analysis of Green500 Lists: Energy Efficiency

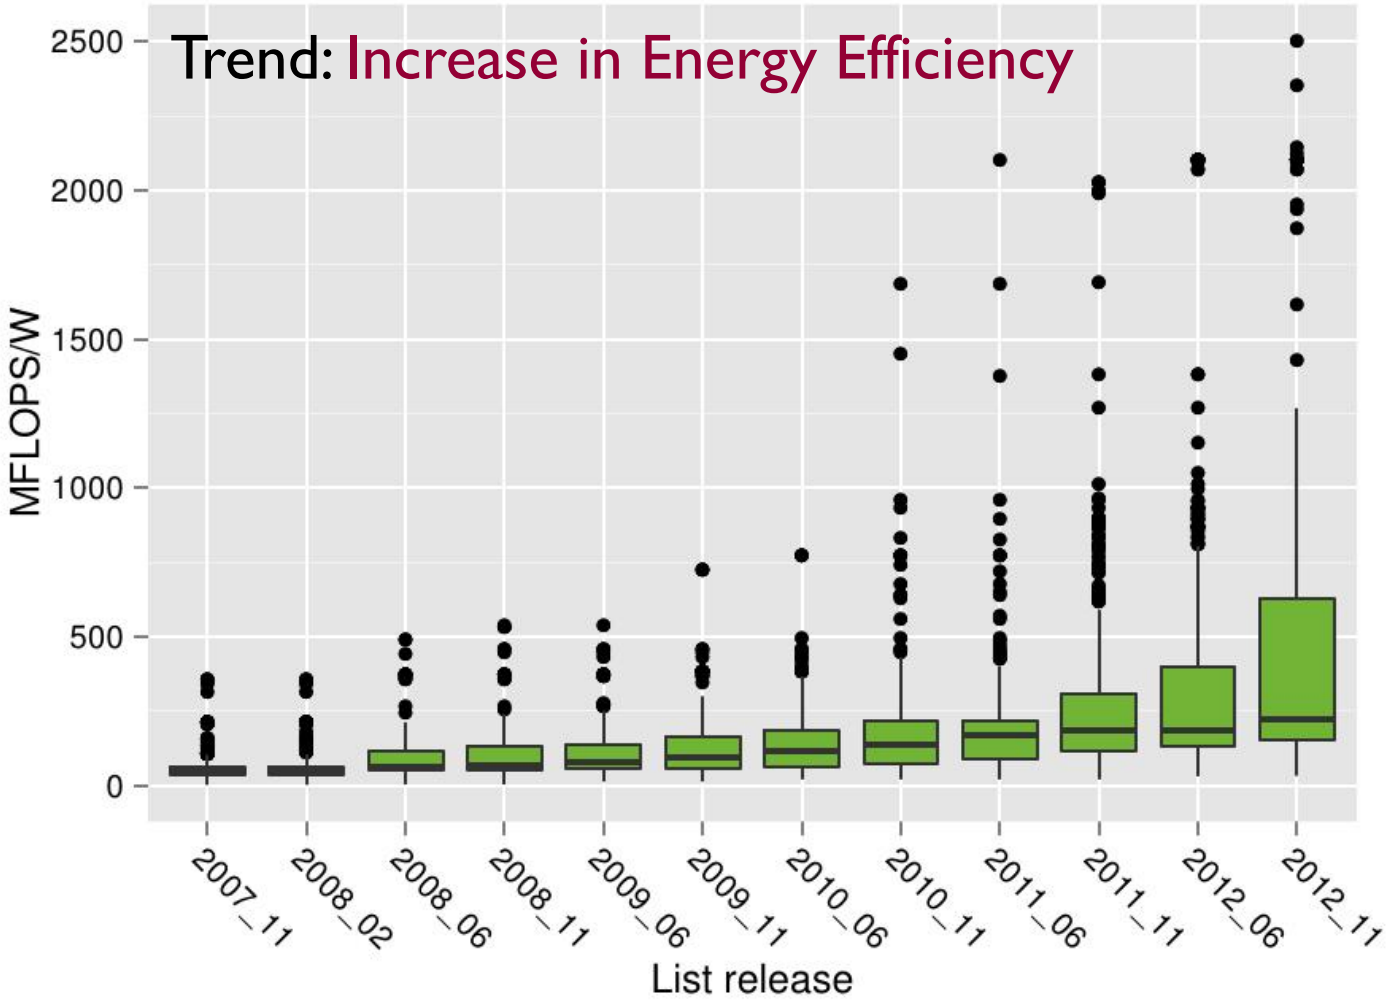

Analysis of Green500 Lists: Energy Efficiency

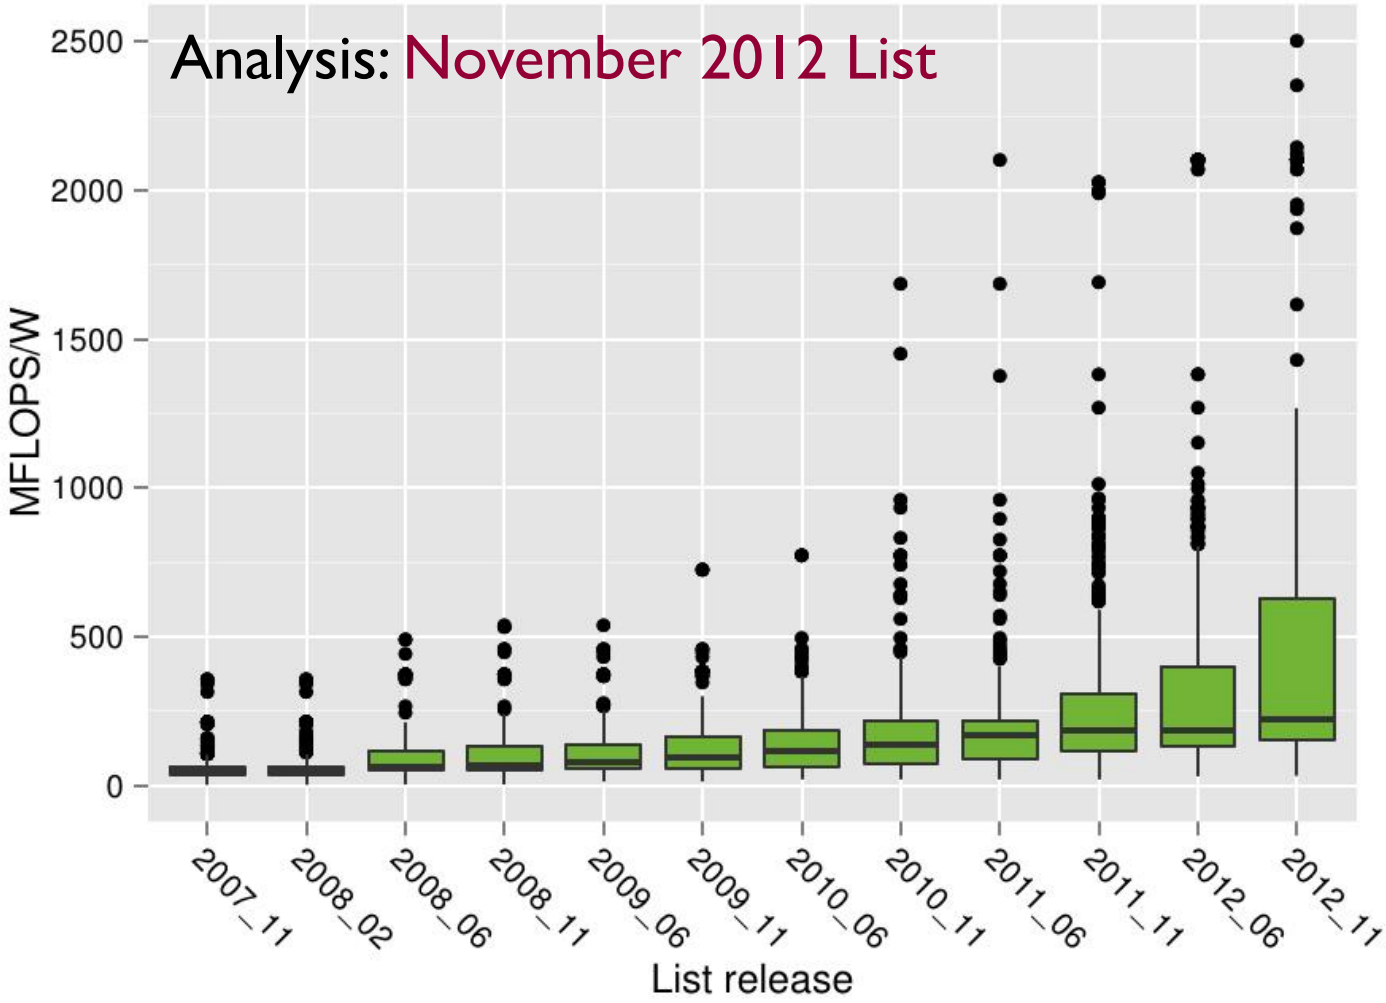

Analysis of Green500 Lists: Energy Efficiency

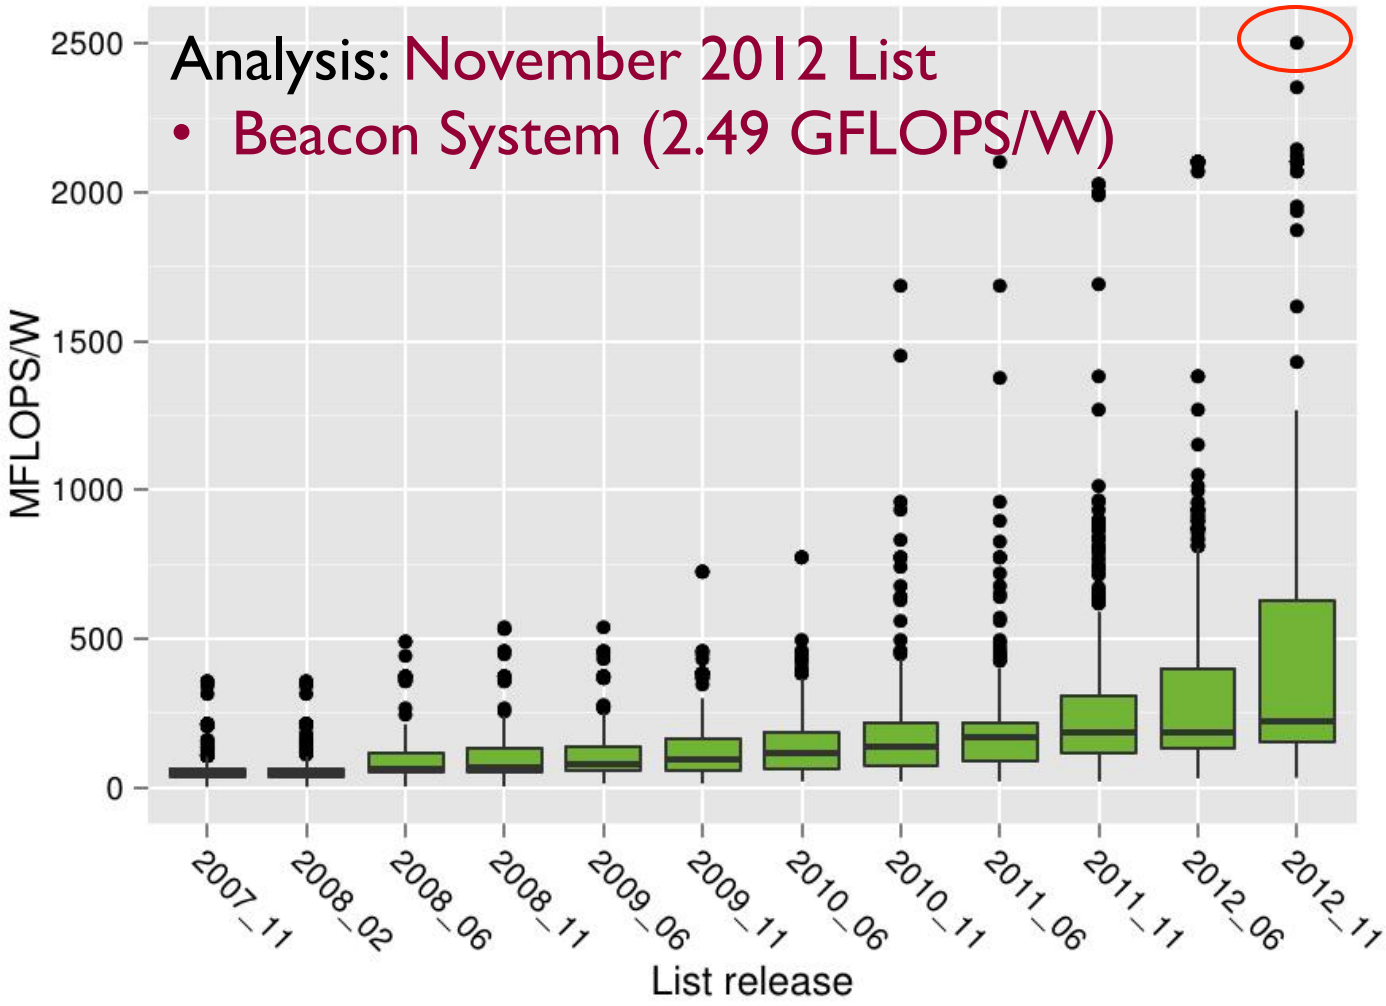

Analysis of Green500 Lists: Energy Efficiency

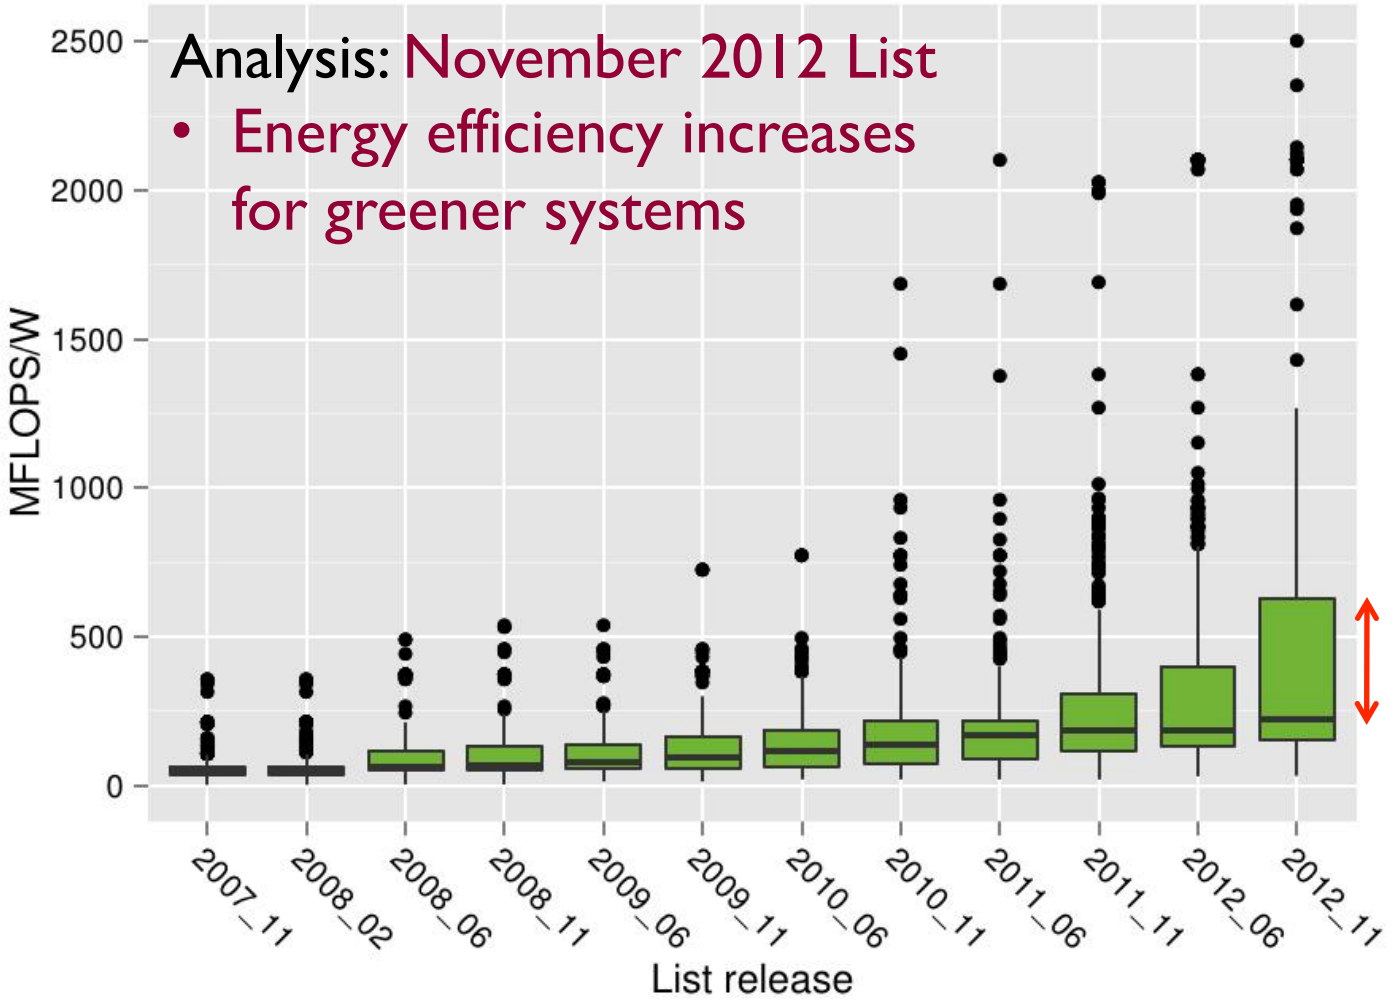

Analysis of Green500 Lists: Energy Efficiency

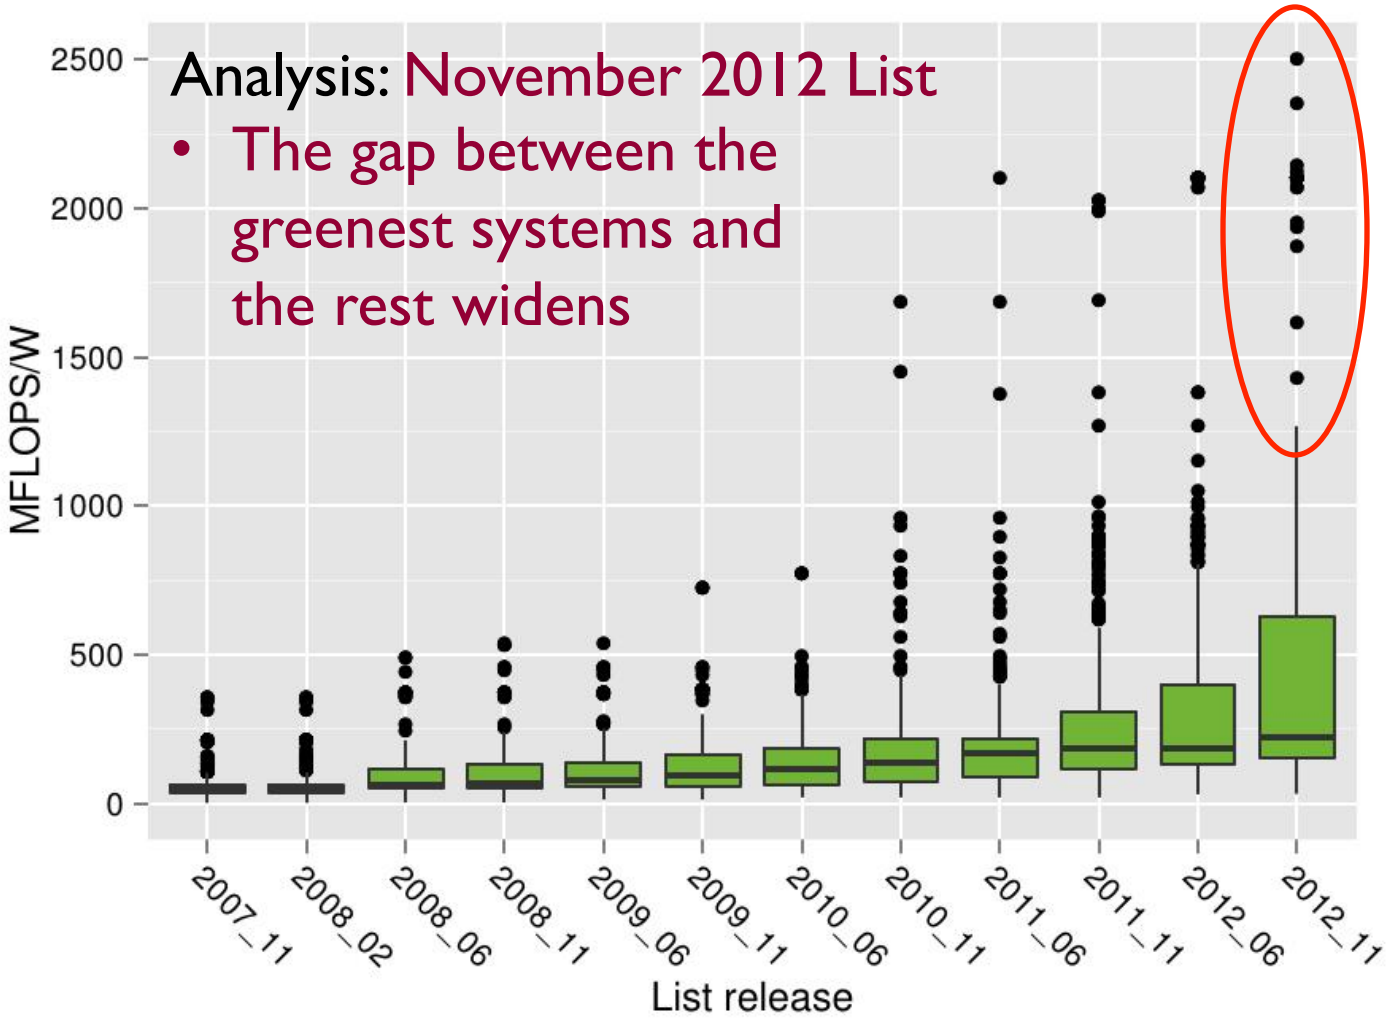

Analysis of Green500 Lists:
Efficiency by Machine Type

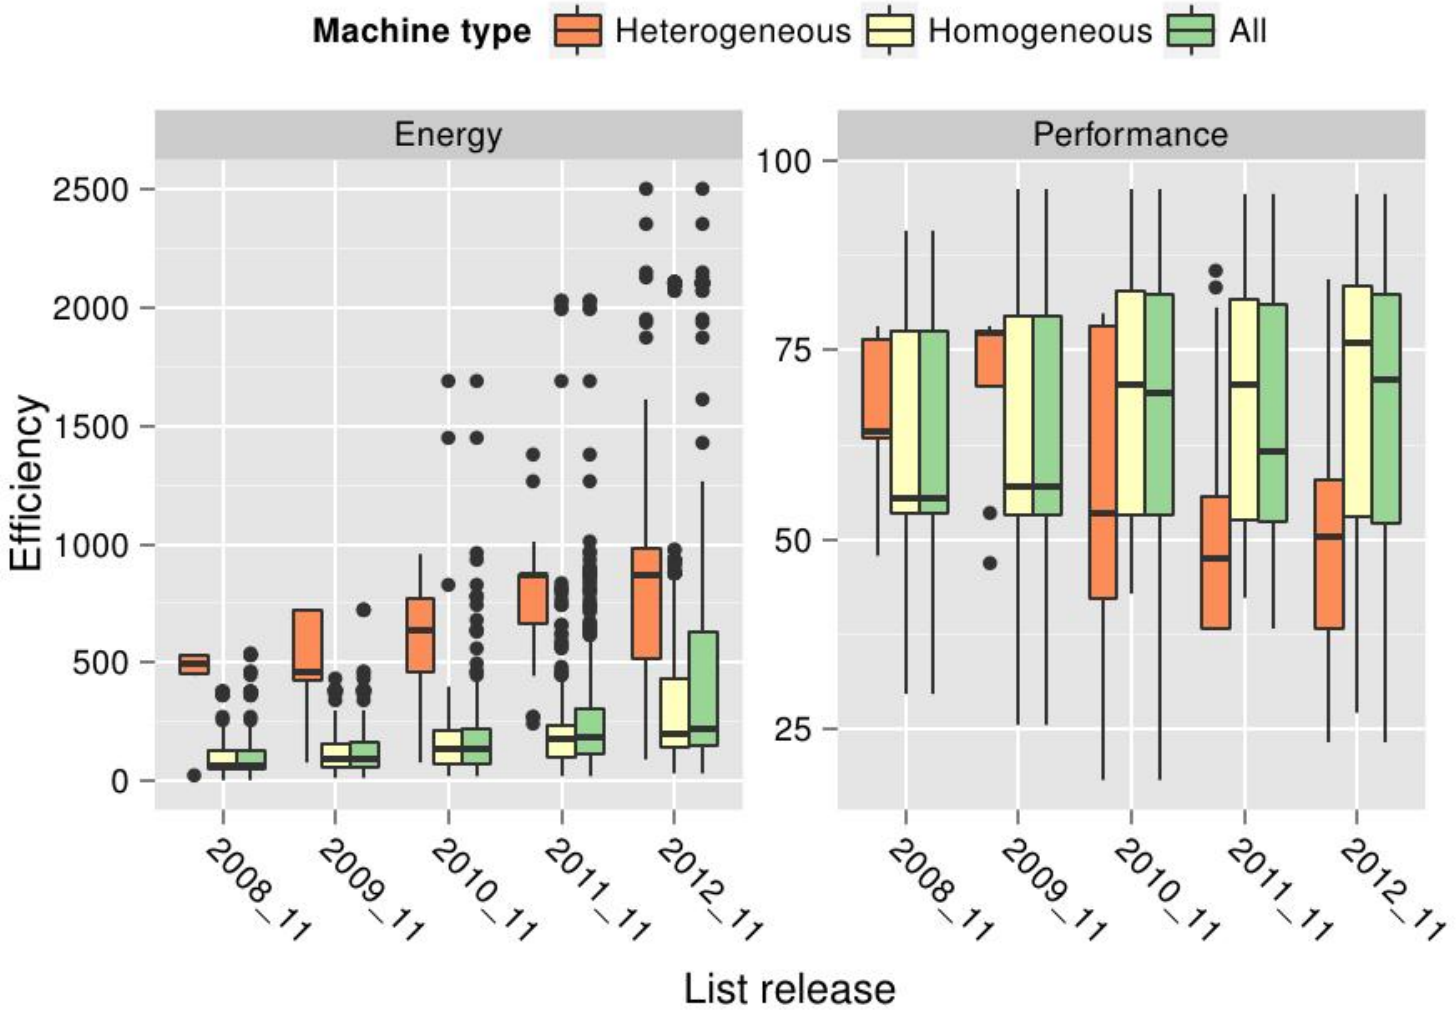

Analysis of Green500 Lists:
Efficiency by Machine Type

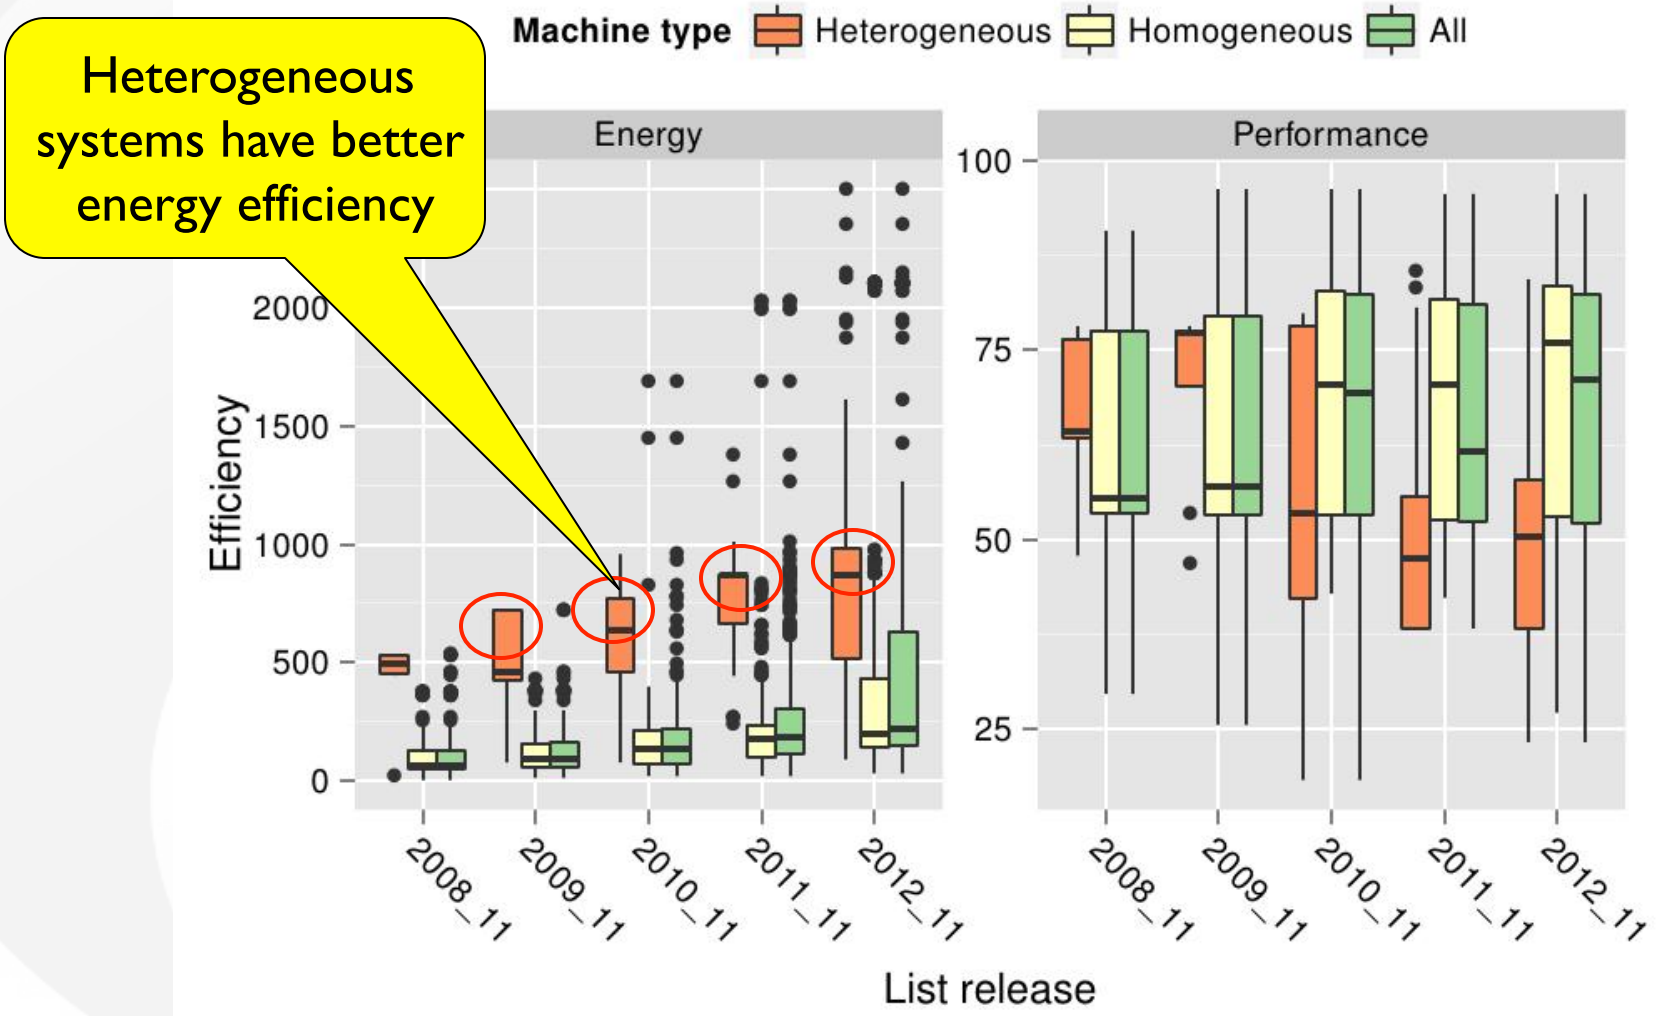

Analysis of Green500 Lists:
Efficiency by Machine Type

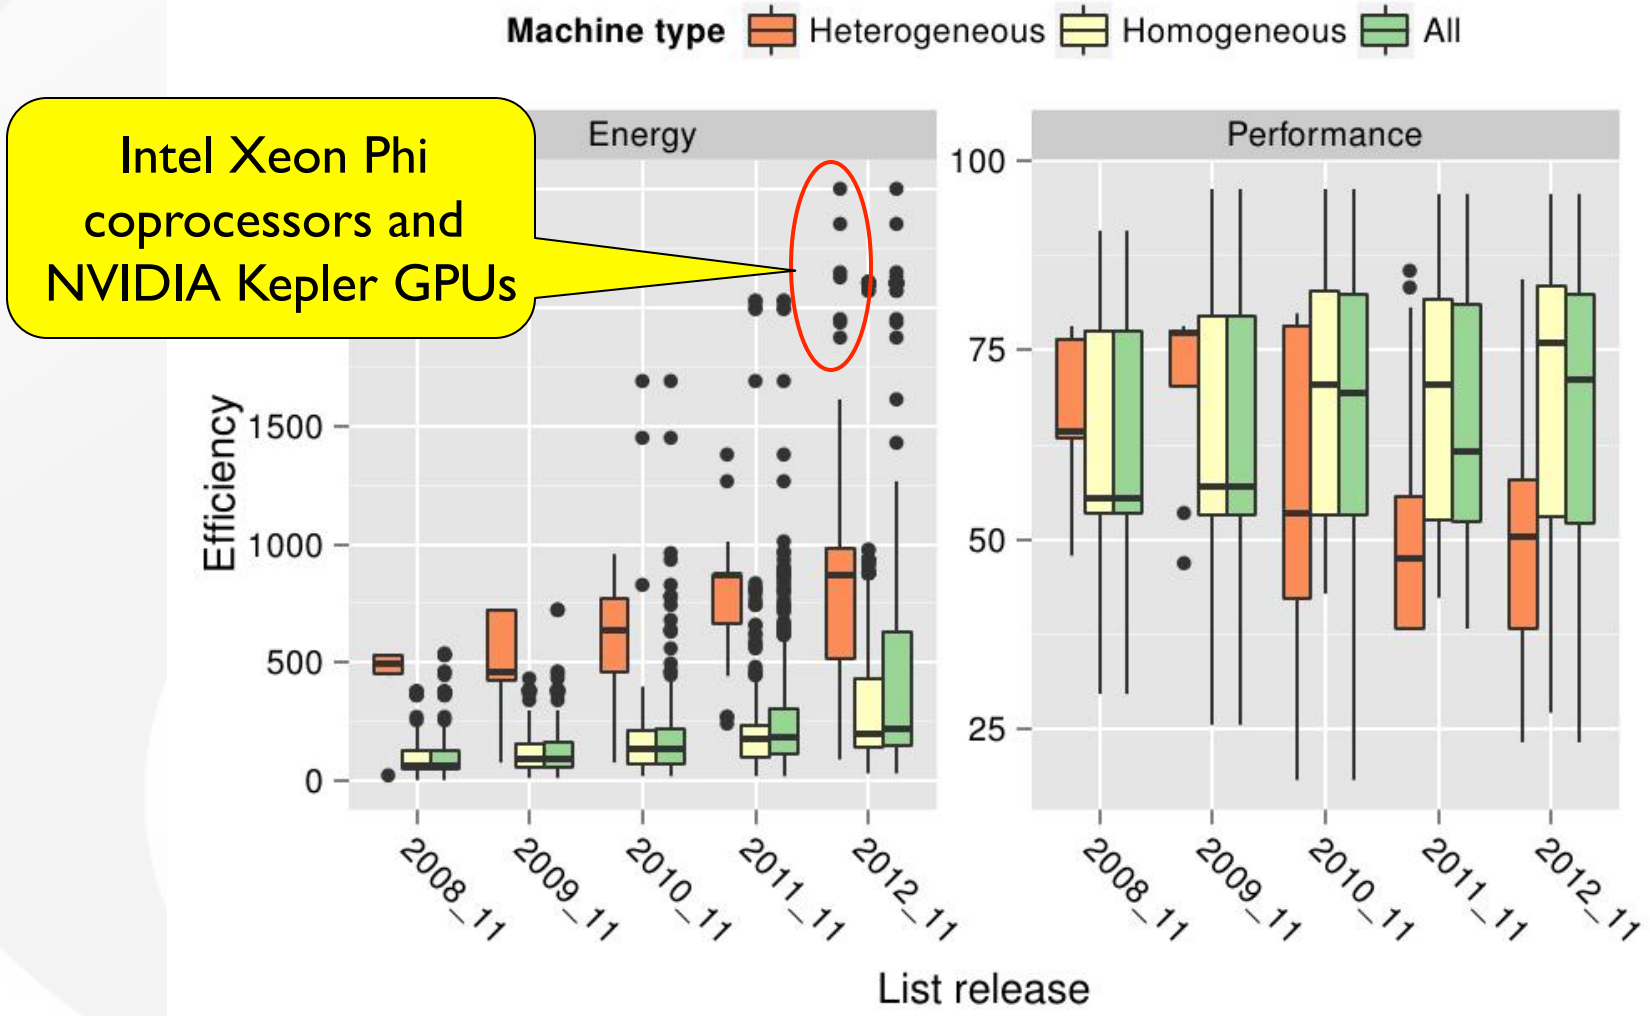

Analysis of Green500 Lists:
Efficiency by Machine Type

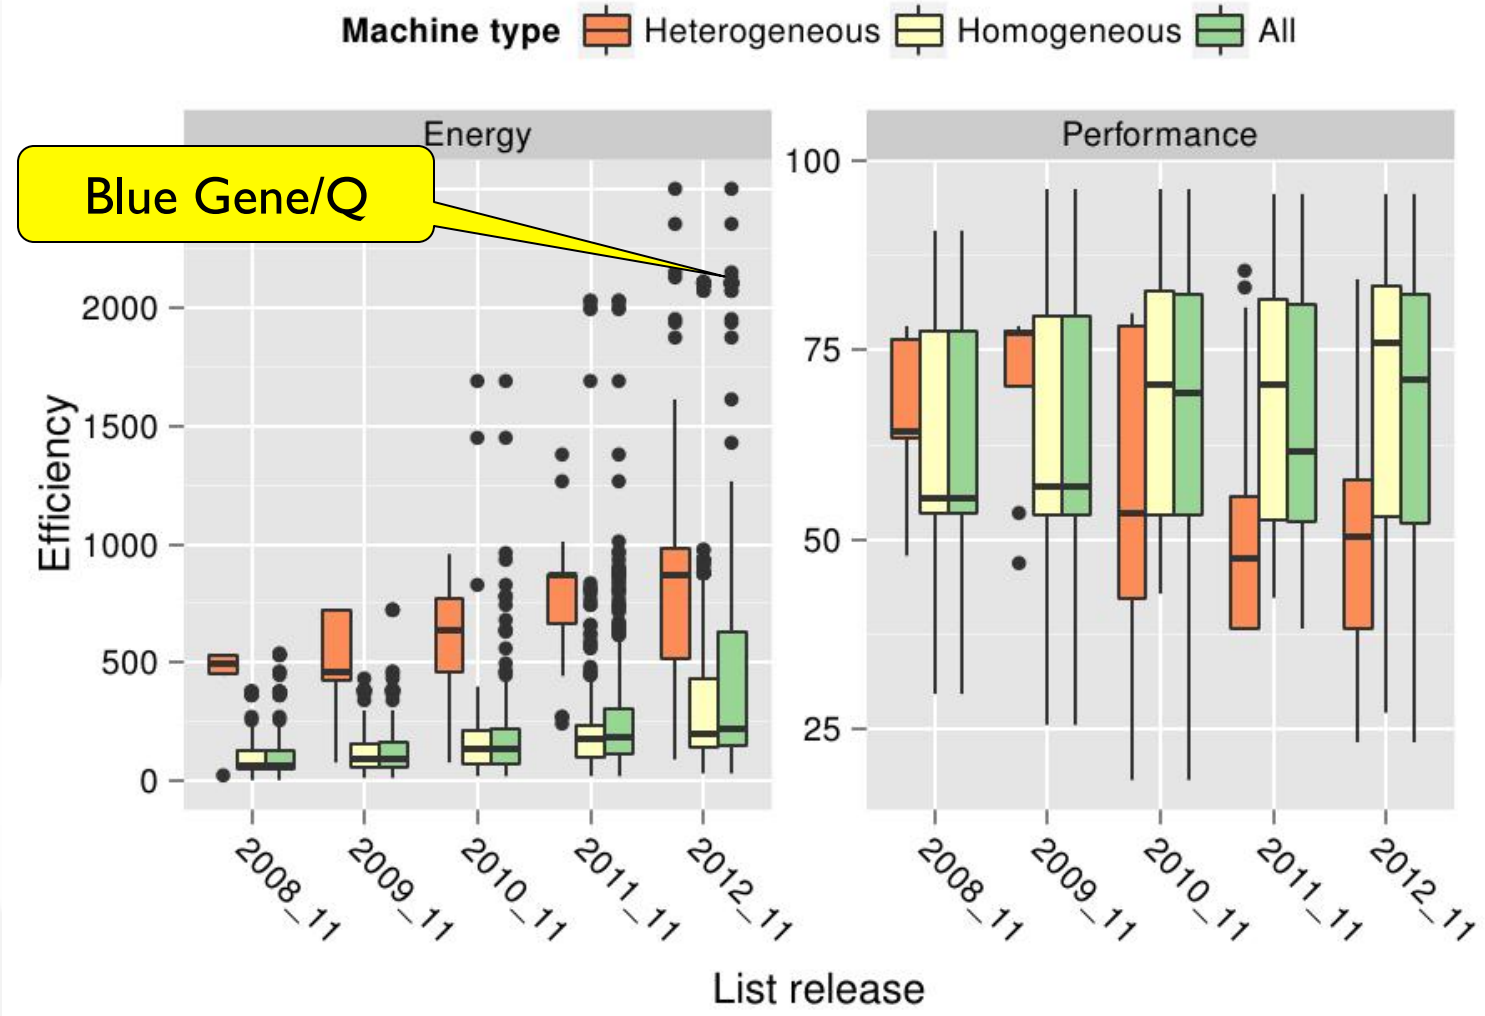

Analysis of Green500 Lists:
Efficiency by Machine Type

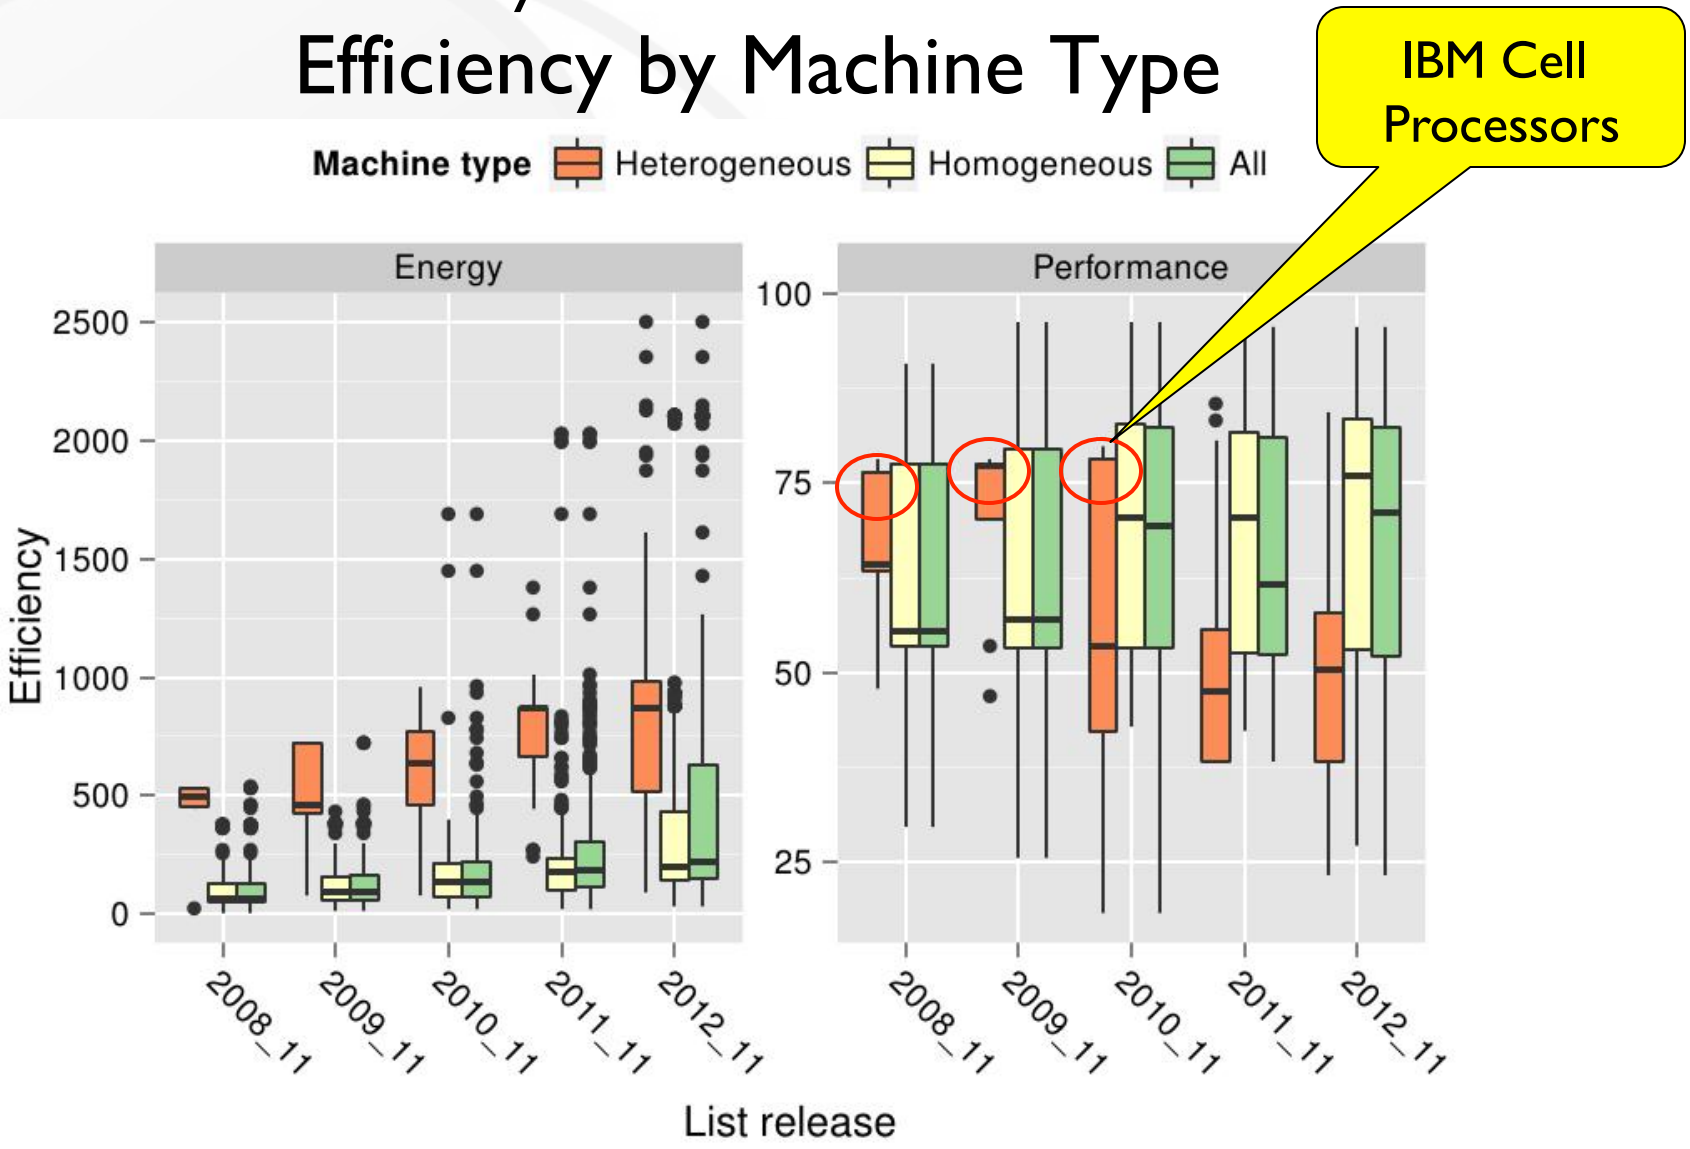

Analysis of Green500 Lists:
Efficiency by Machine Type

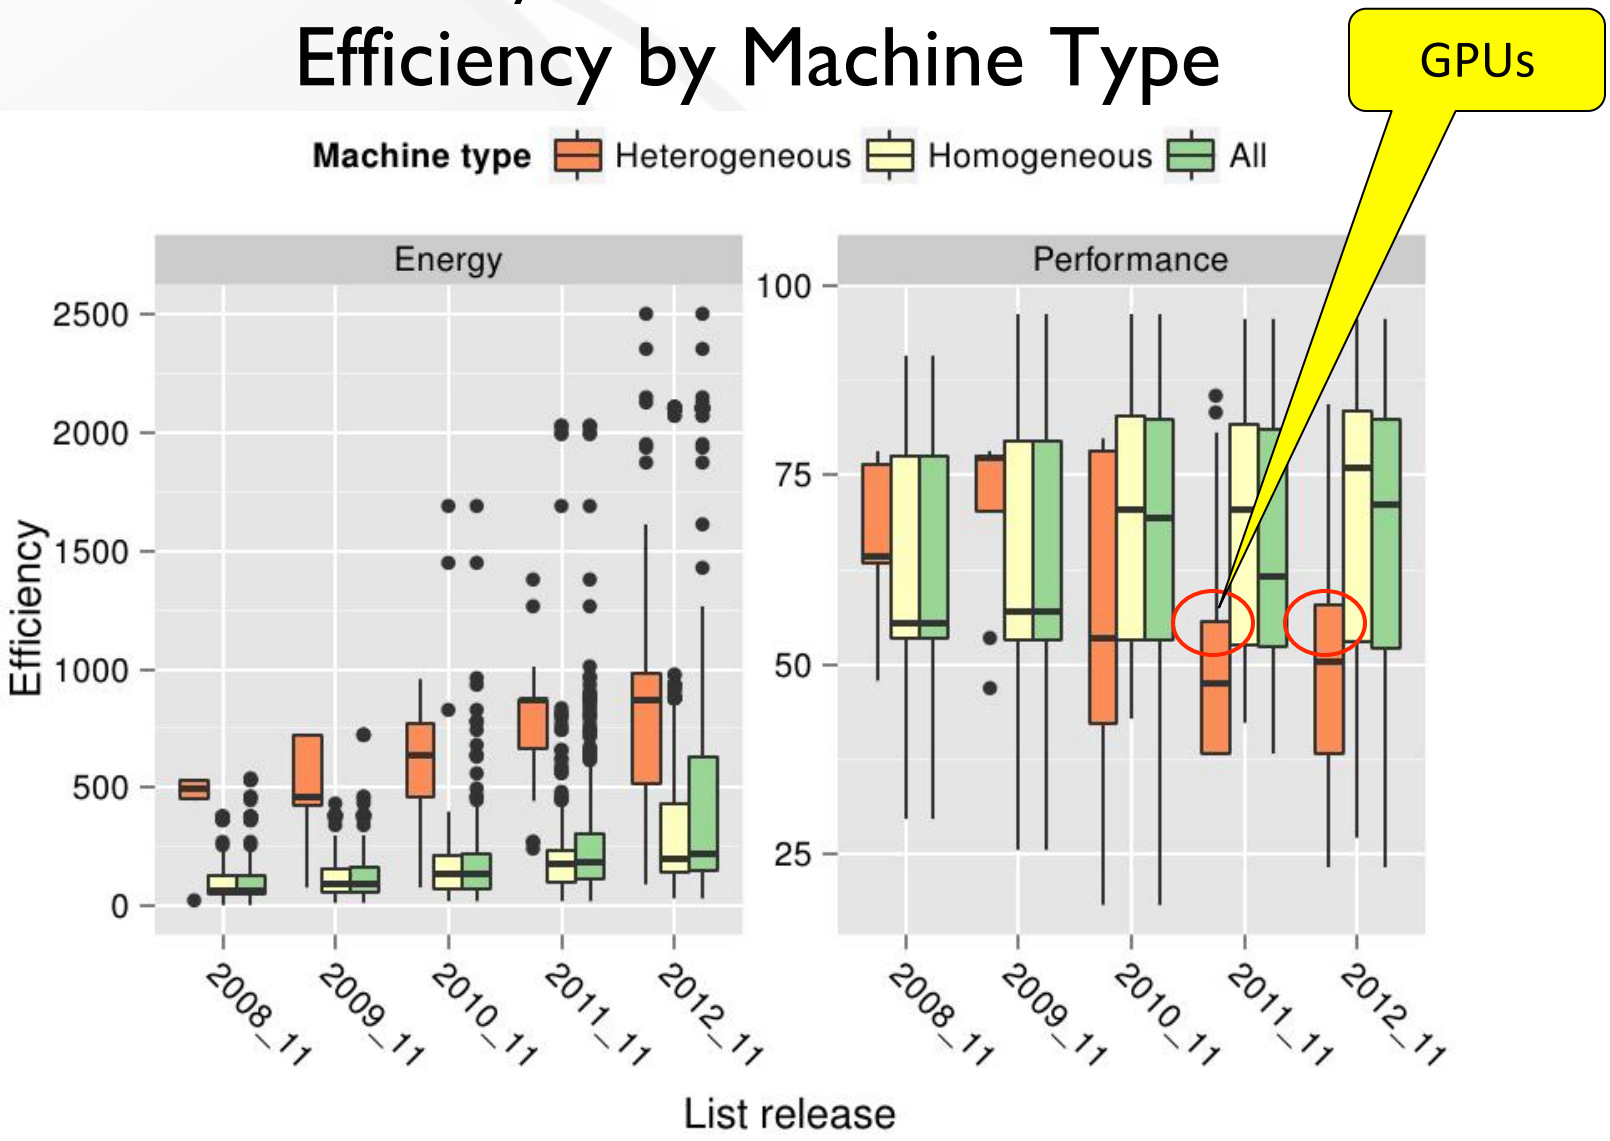

Analysis of Green500 Lists:
Efficiency by Machine Type

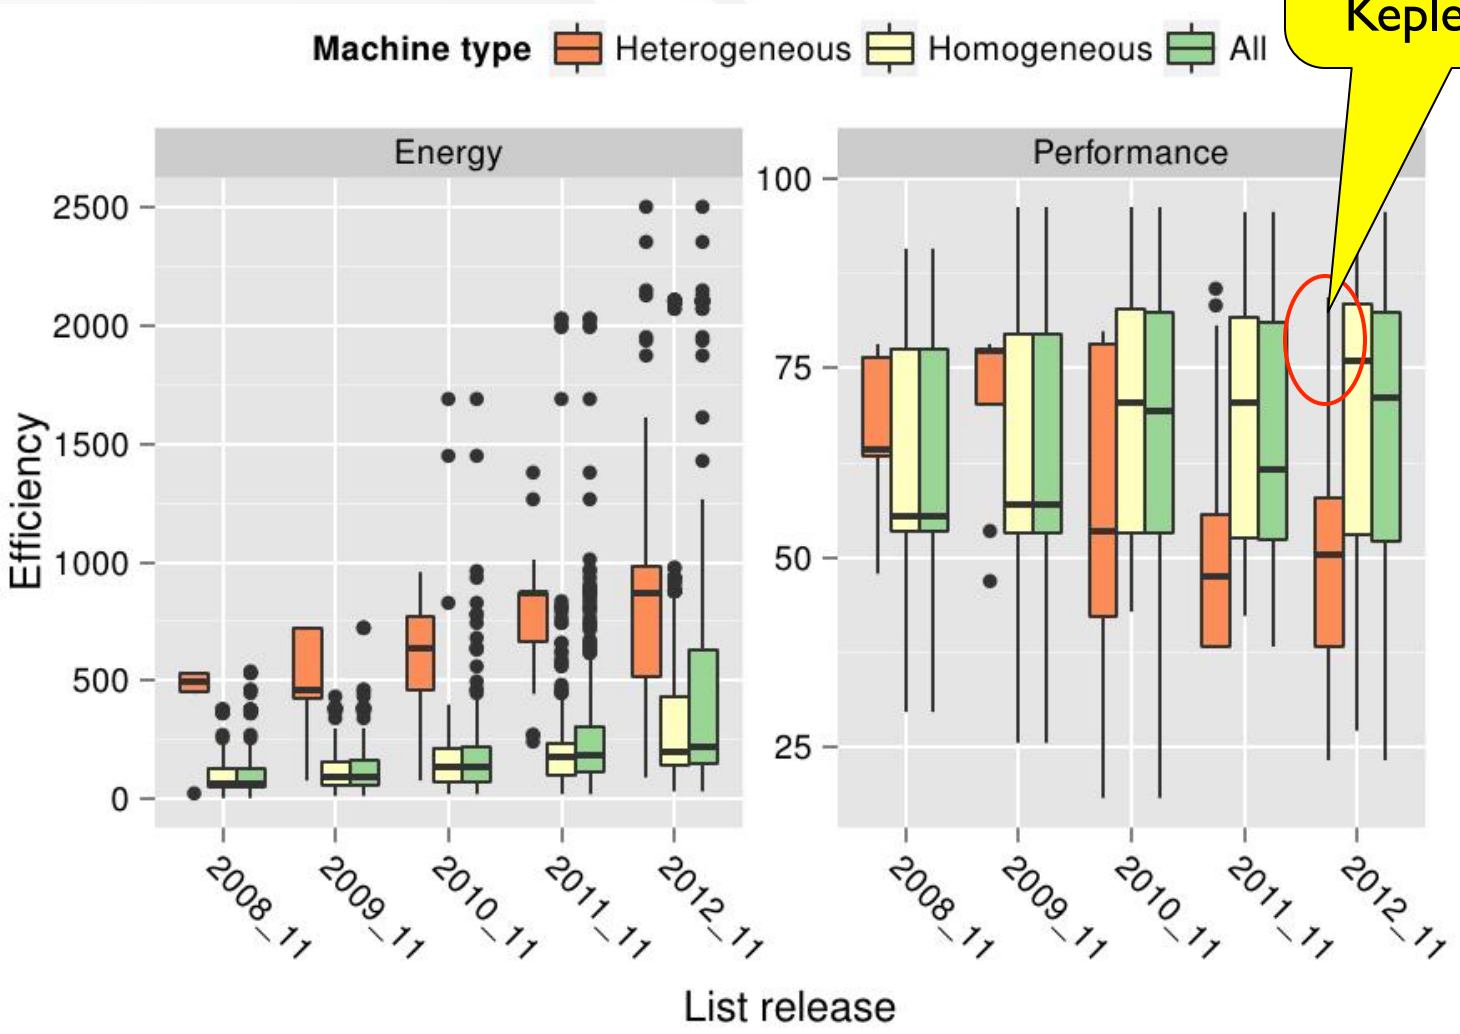

Intel Xeon Phis and NVIDIA Kepler GPUs

Agenda

- The Green500 List: A Brief History
  - From Green Destiny to the Green500 List
- Analysis of Green500 Lists
  - Energy Efficiency
- Projections for Energy-Efficient Exascale
  - Tracking Koomey's Law for High-Performance Computing (HPC)
  - Quantifying *Distance* from Exascale Goal: The Exascalar Metric

Projections for Energy-Efficient Exascale Systems

- How to achieve 20 MW exascale system?

Projections for Energy-Efficient Exascale Systems

- How to achieve 20 MW exascale system?
  - Current fastest supercomputer, Tianhe-2
    - 33.86 PFLOPS consuming 17.80 MW power.

Projections for Energy-Efficient Exascale Systems

- How to achieve 20 MW exascale system?
  - Current fastest supercomputer, Tianhe-2
    - 33.86 PFLOPS consuming 17.80 MW power.
  - Required: *30-fold increase in performance with only 1.12-fold increase in power.*

Projections for Energy-Efficient Exascale Systems

- How to achieve 20 MW exascale system?
  - Current fastest supercomputer, Tianhe-2
    - 33.86 PFLOPS consuming 17.80 MW power.
  - Required: *30-fold increase in performance with only 1.12-fold increase in power.*
- Is this target realistic?

Projections for Energy-Efficient Exascale Systems

- How to achieve 20 MW exascale system?
  - Current fastest supercomputer, Tianhe-2
    - 33.86 PFLOPS consuming 17.80 MW power.
  - Required: *30-fold increase in performance with only 1.12-fold increase in power.*
- Is this target realistic?
  - Track energy efficiency in the past
  - Project energy efficiency for the future

Tracking Koomey's Law for HPC

- Koomey's Law
  - Efficiency (computations/kilowatt-hour) doubles every 1.57 years
  - Potential Issue
    - Applied to *peak power and performance*.
    - What about systems running at less than peak power?

Source: Koomey et al., Implications of Historical Trends in the Electrical Efficiency of Computing, *IEEE Annals of the History of Computing*, 2011

Tracking Koomey's Law for HPC

Y-axis
- operations/energy
- computations/KWh

X-axis
- year

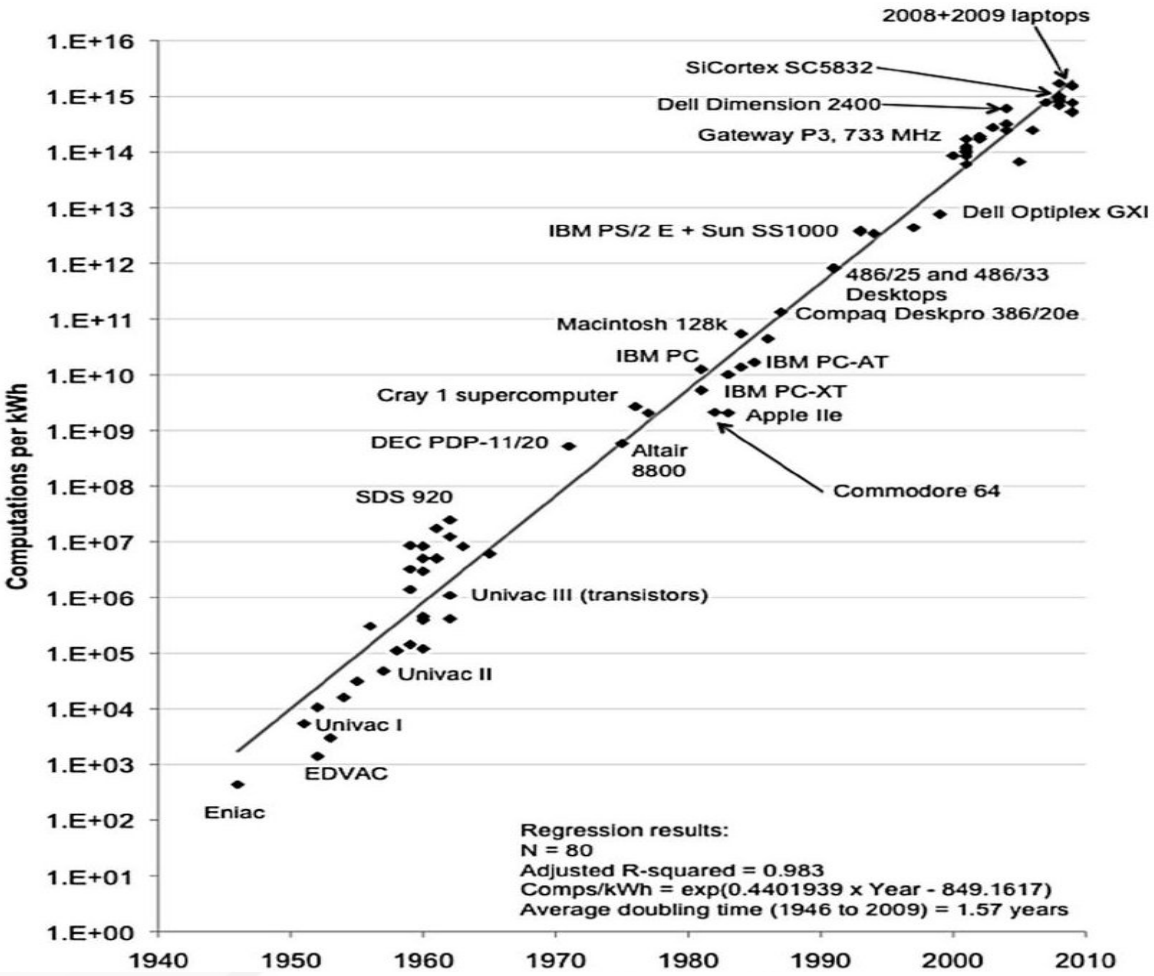

Source: Koomey et al., Implications of Historical Trends in the Electrical Efficiency of Computing, *IEEE Annals of the History of Computing*, 2011

Tracking Koomey's Law for HPC

- Analysis of Green500 Lists
  - Track energy efficiency for last 8 lists

Tracking Koomey's Law for HPC

- Analysis of Green500 Lists
  - Track energy efficiency for last 8 lists
    - Green5 (Greenest five systems on Green500)
    - Green10 (Greenest ten systems on Green500)
    - Green50 (Greenest fifty systems on Green500)
    - Green100 (Greenest hundred systems on Green500)

Tracking Koomey's Law for HPC

- Analysis of Green500 Lists
  - Track energy efficiency for last 8 lists
    - Green5 (Greenest five systems on Green500)
    - Green10 (Greenest ten systems on Green500)
    - Green50 (Greenest fifty systems on Green500)
    - Green100 (Greenest hundred systems on Green500)
  - Use linear regression to model the energy efficiency trend

Tracking Koomey's Law for HPC

- Analysis of Green500 Lists
  - Track energy efficiency for last 8 lists
    - Green5 (Greenest five systems on Green500)
    - Green10 (Greenest ten systems on Green500)
    - Green50 (Greenest fifty systems on Green500)
    - Green100 (Greenest hundred systems on Green500)
  - Use linear regression to model the energy efficiency trend
    - Regression input: List index
    - Regression output: Average energy efficiency at list index
    - Regression quality of fit: R^2 metric

Tracking Koomey's Law for HPC: Green5 and Green10

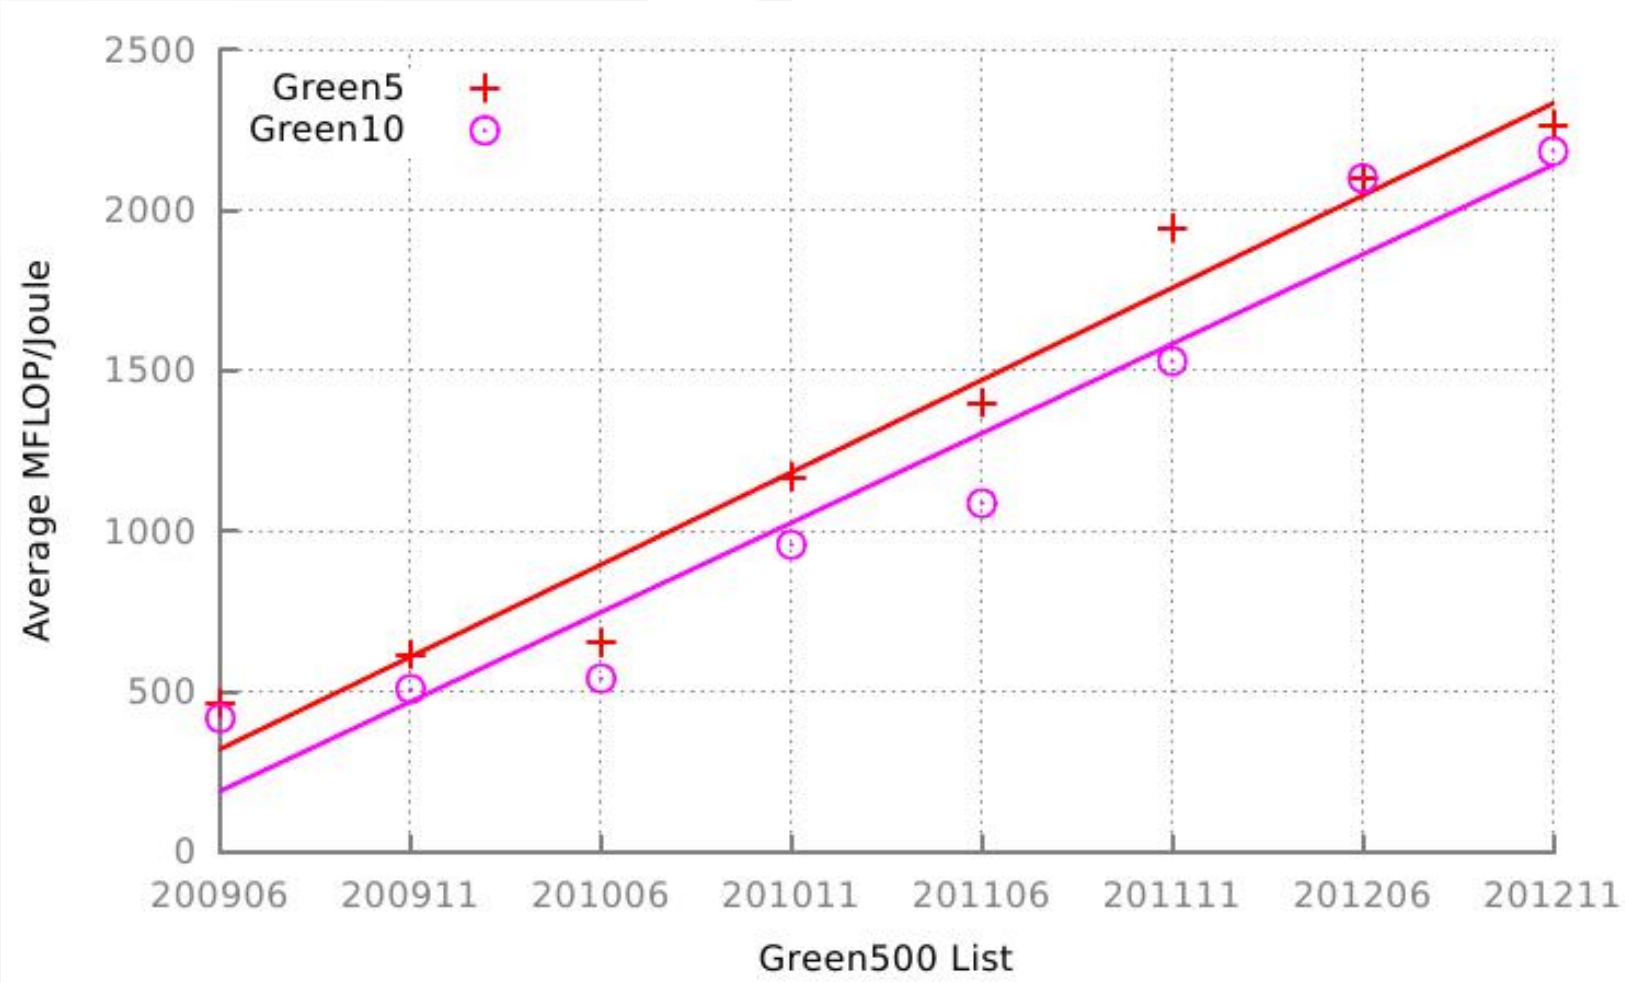

Tracking Koomey's Law for HPC: Green5 and Green10

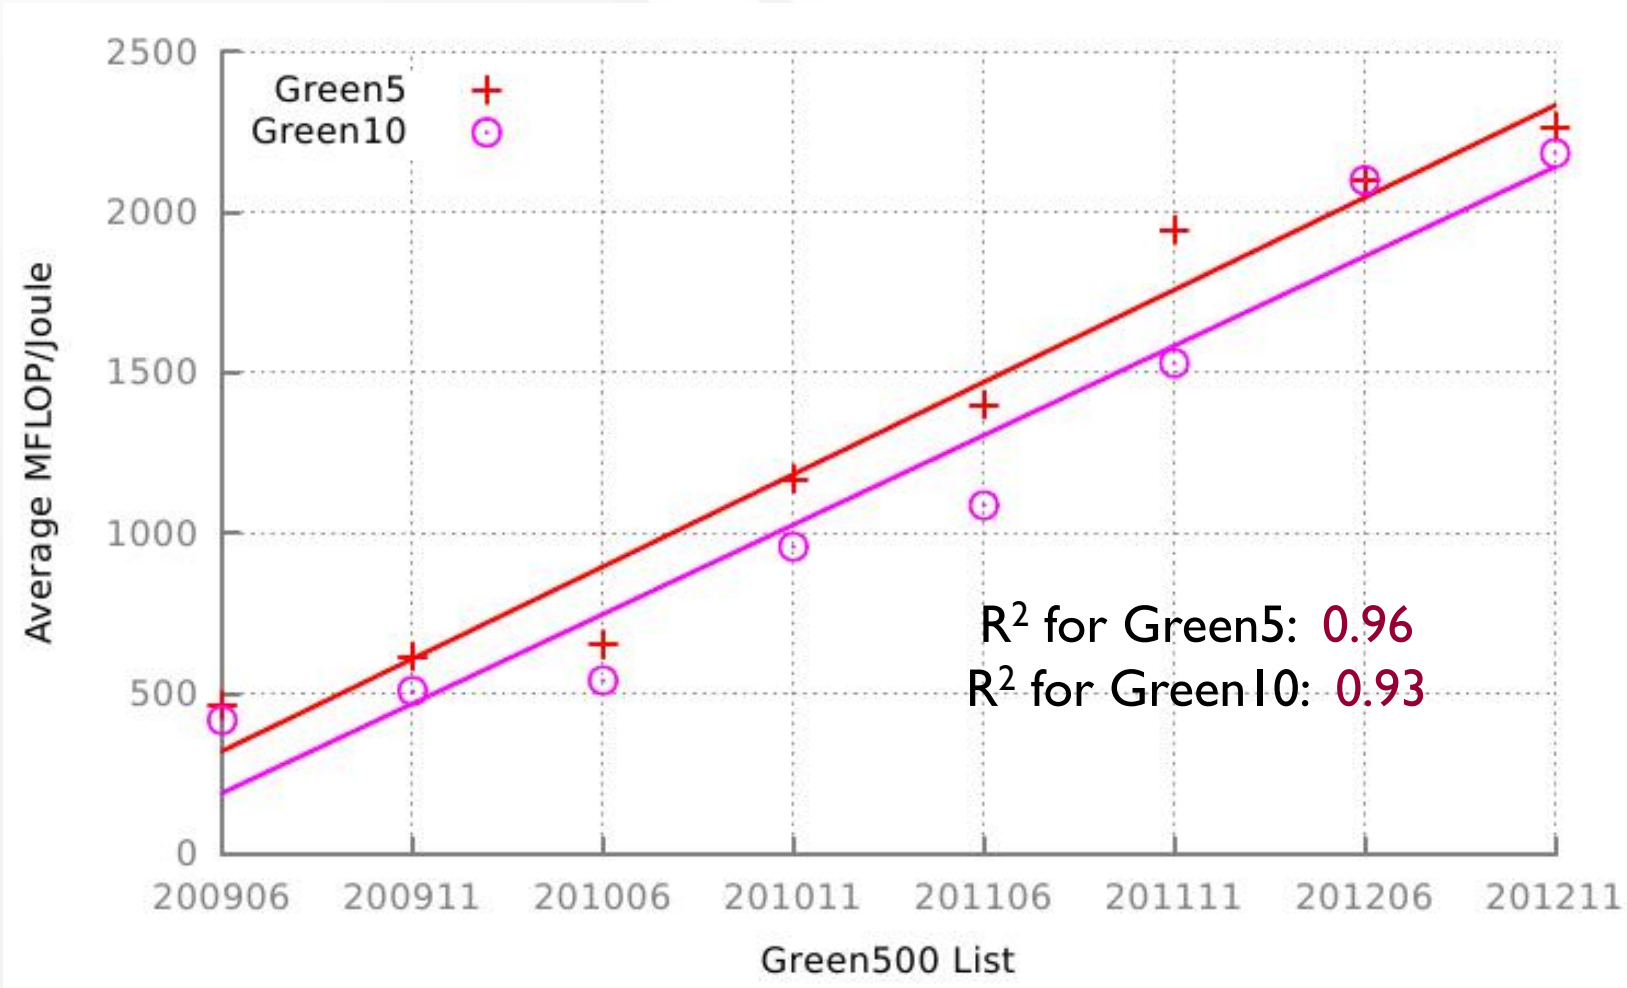

Tracking Koomey's Law for HPC: Green5 and Green10

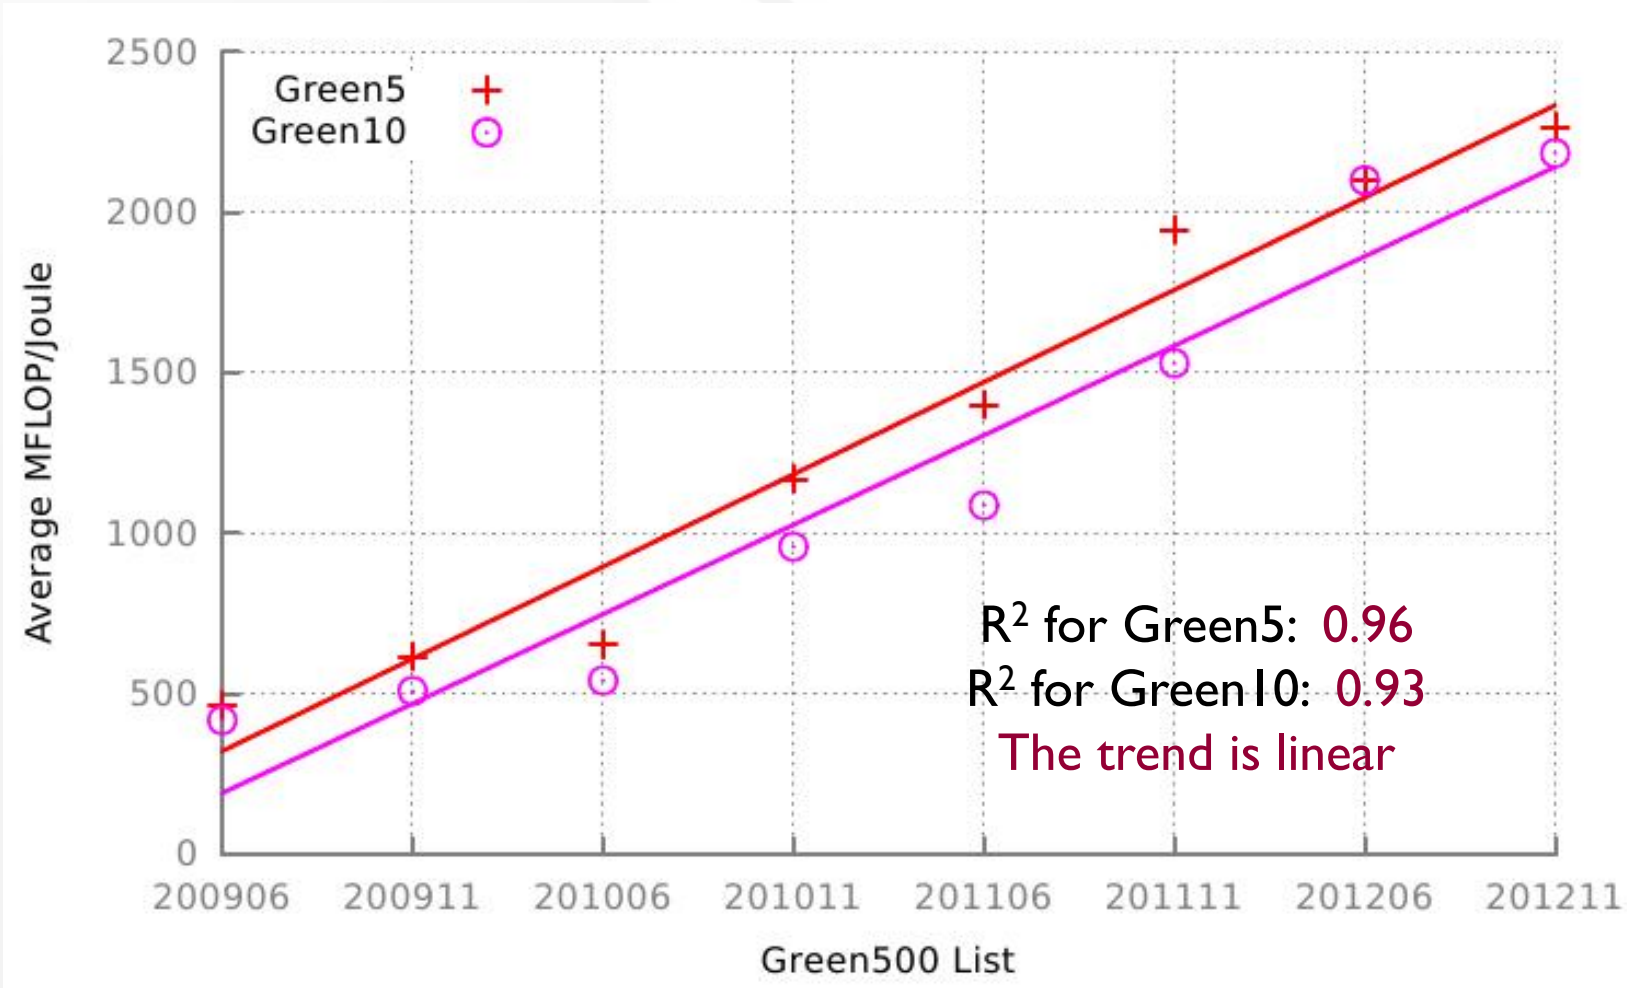

Tracking Koomey's Law for HPC: Green50 and Green100

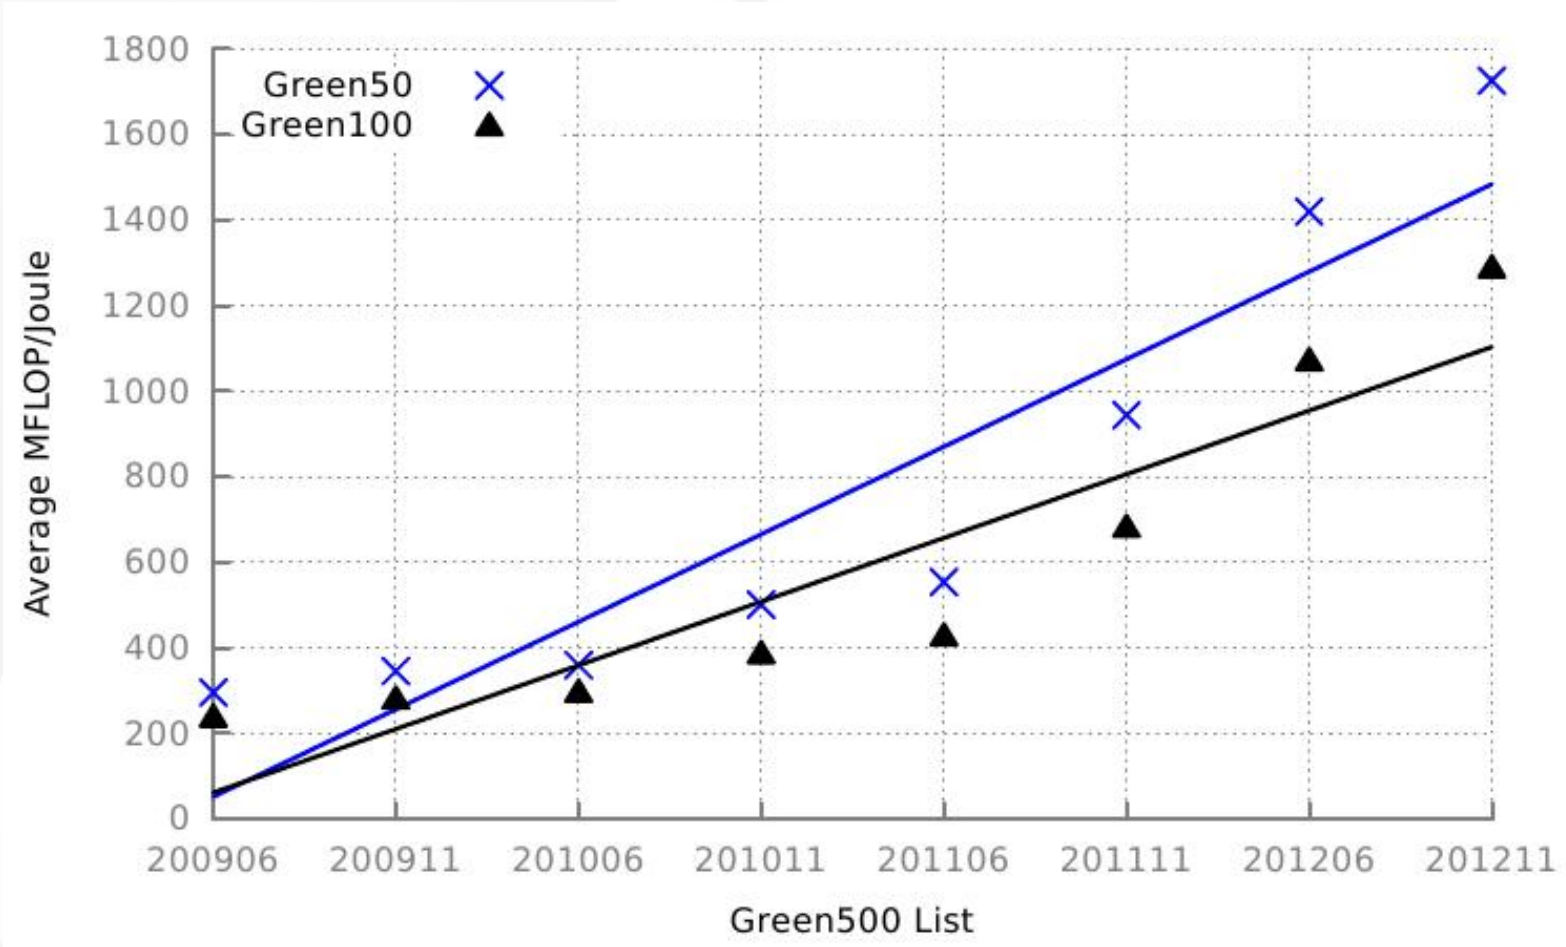

Tracking Koomey's Law for HPC: Green50 and Green100

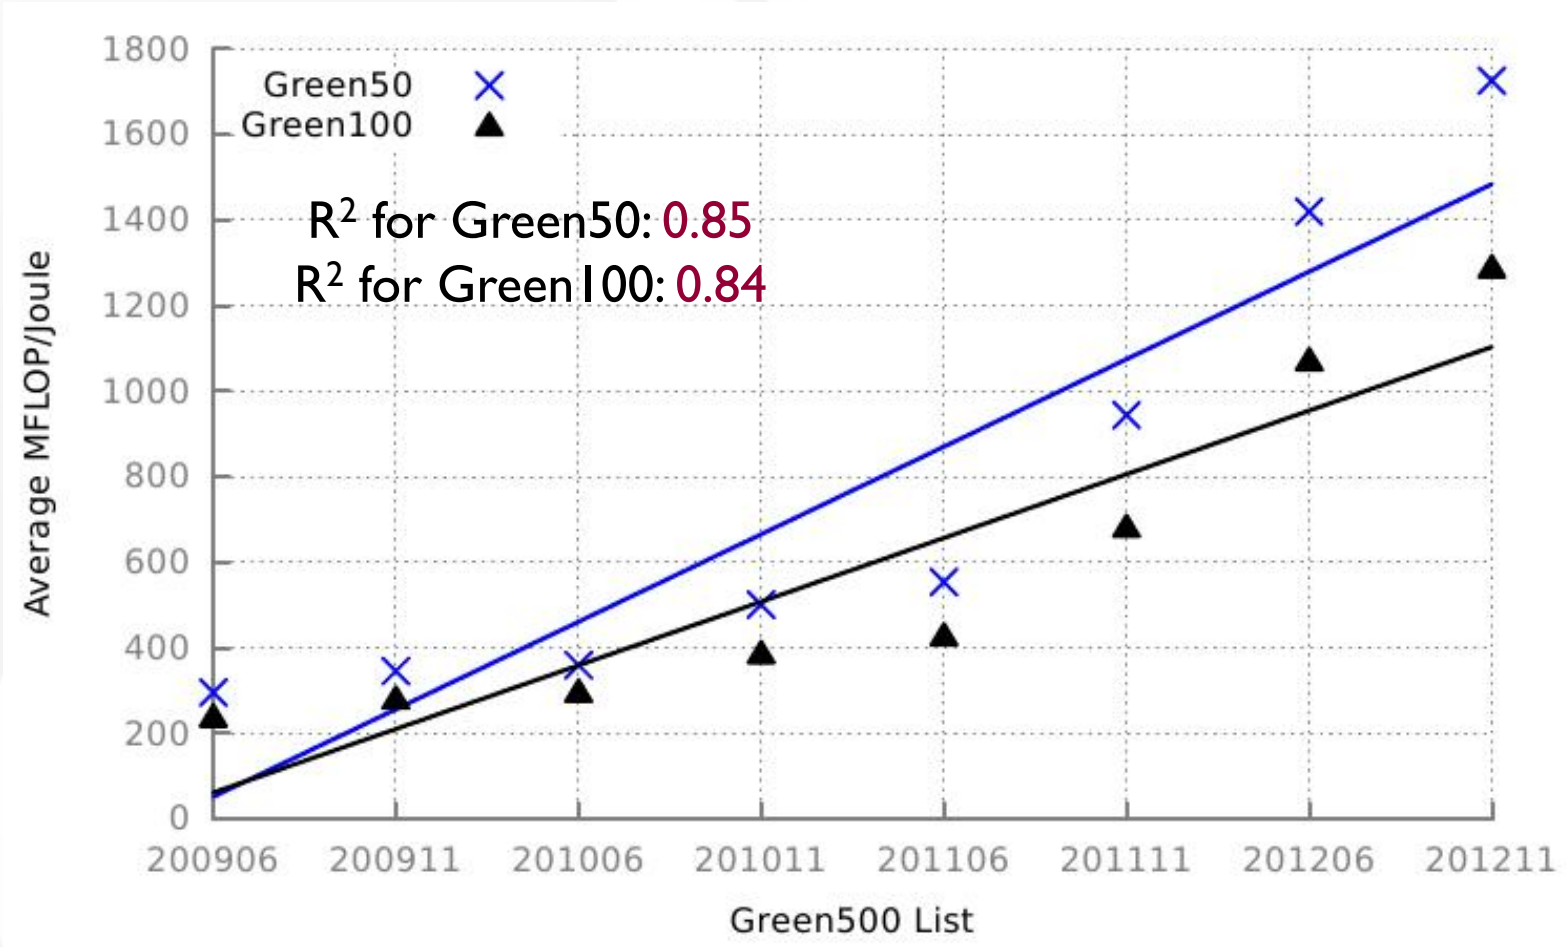

Tracking Koomey's Law for HPC: Green50 and Green100

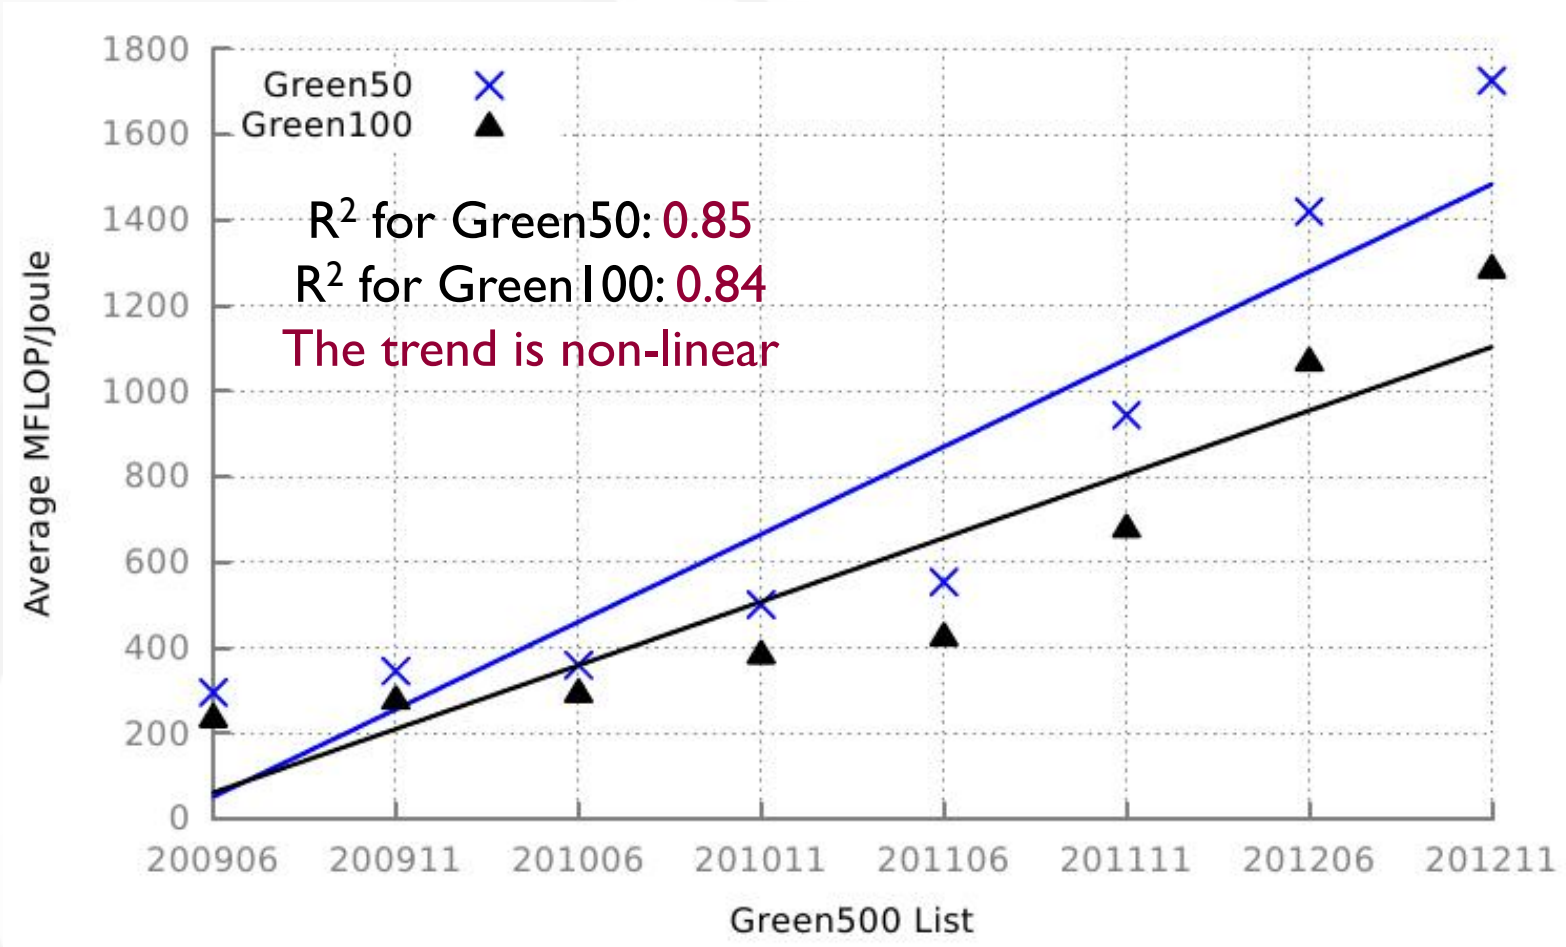

Tracking Koomey's Law for HPC: Comparison

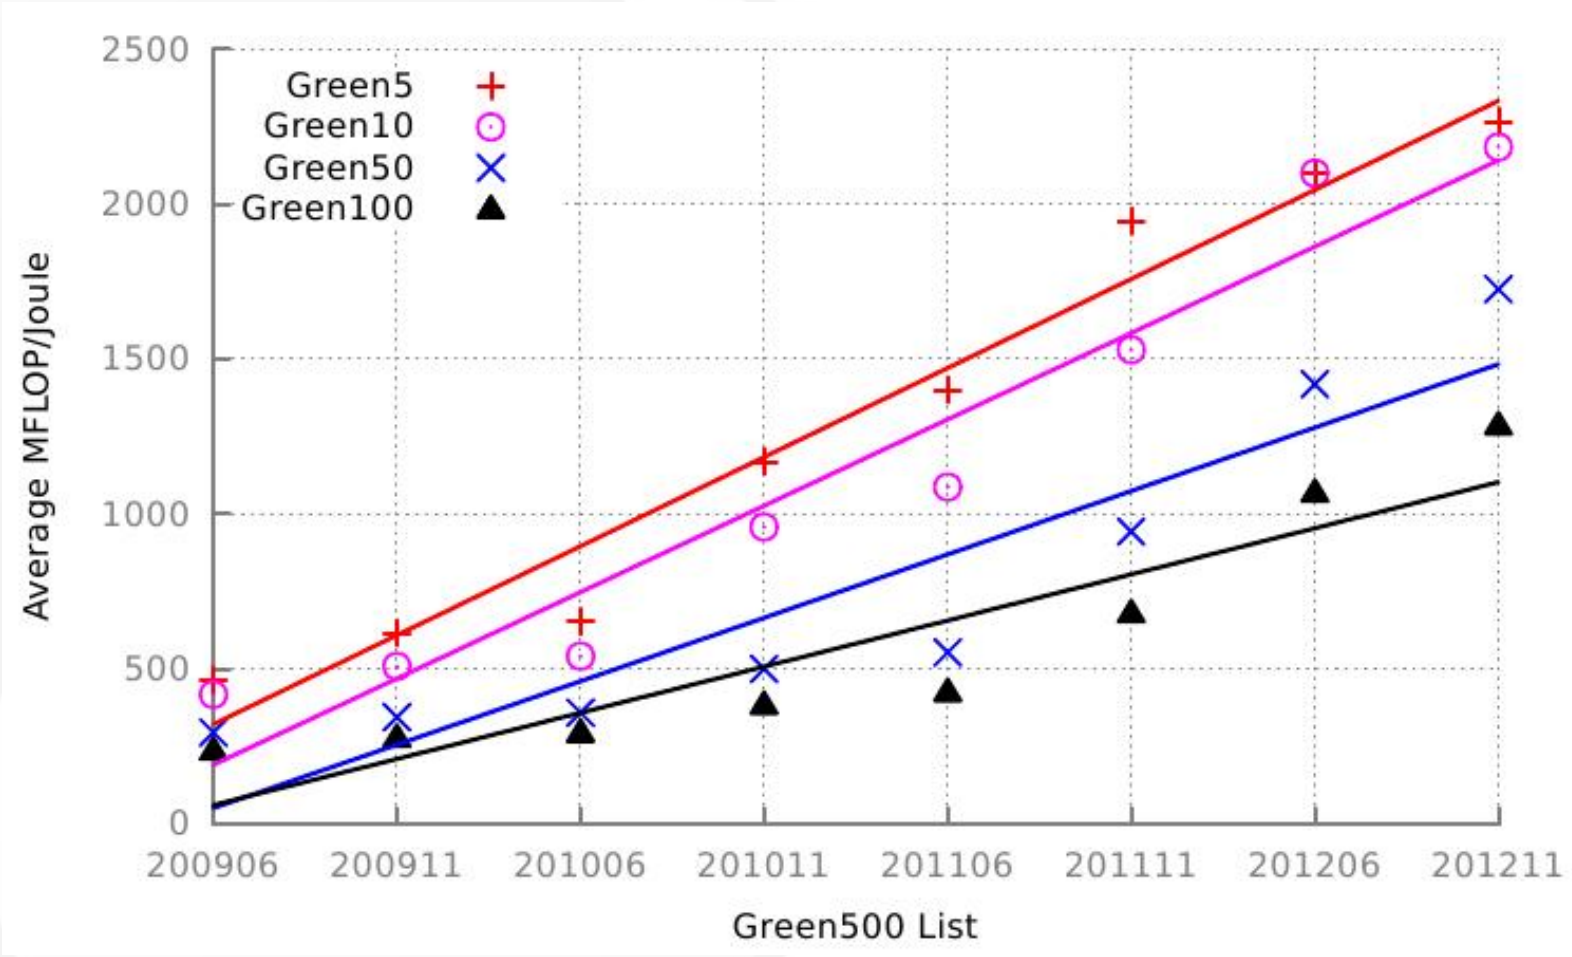

Projection to Exascale

- Efficiency of an exaflop system
  - 20-MW power envelope: 50 GFLOPS/watt
  - 100-MW power envelope: 10 GFLOPS/watt

Projection to Exascale

- Efficiency of an exaflop system
  - 20-MW power envelope: 50 GFLOPS/watt
  - 100-MW power envelope: 10 GFLOPS/watt
- State-of-the-art
  - Greenest system: 2.49 GFLOPS/watt
  - Fastest system: 2.14 GFLOPS/watt

Projection to Exascale

- Efficiency of an exaflop system
  - 20-MW power envelope: 50 GFLOPS/watt
  - 100-MW power envelope: 10 GFLOPS/watt
- State-of-the-art
  - Greenest system: 2.49 GFLOPS/watt
  - Fastest system: 2.14 GFLOPS/watt
- Projections: Efficiency (GFLOPS/watt) in 2018 and 2020

Class	In 2018	In 2020
Green5	5.78	6.93
Green10	5.48	6.60
Green50	3.93	4.75
Green100	2.88	3.48

Agenda

- The Green500 List: A Brief History
  - From Green Destiny to the Green500 List
- Analysis of Green500 Lists
  - Energy Efficiency
- Projections for Energy-Efficient Exascale
  - Tracking Koomey's Law for High-Performance Computing (HPC)
  - Quantifying *Distance* from Exascale Goal: The Exascalar Metric

Distance from Exaflop Goal

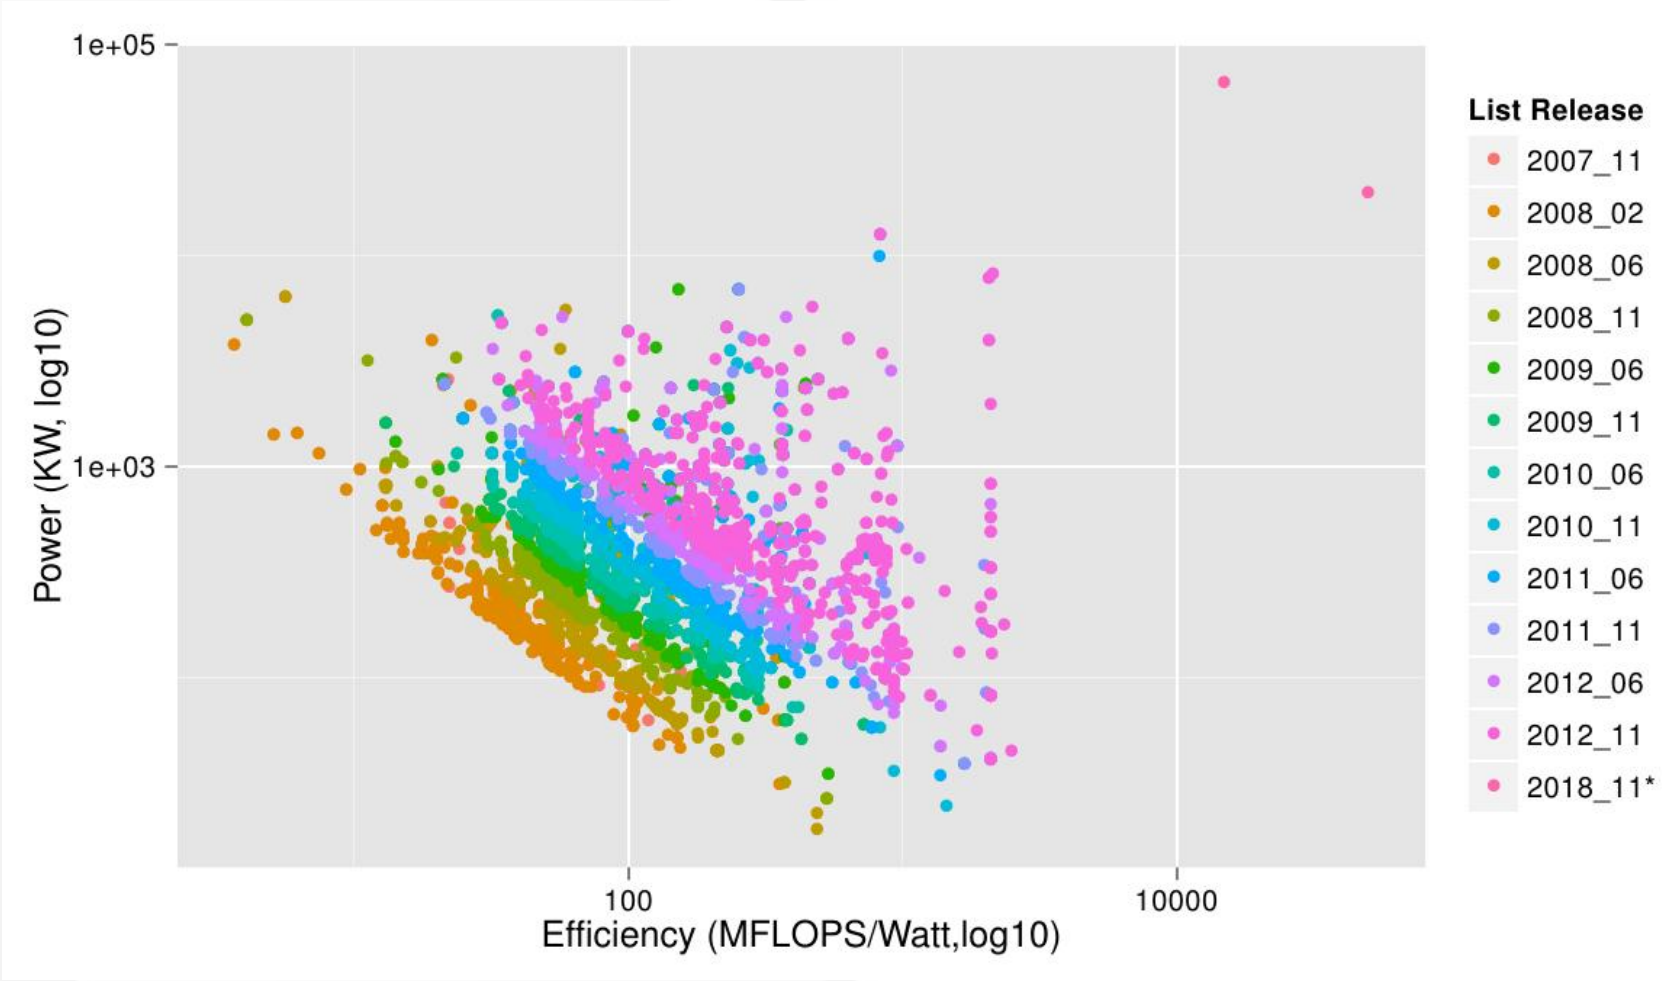

Distance from Exaflop Goal

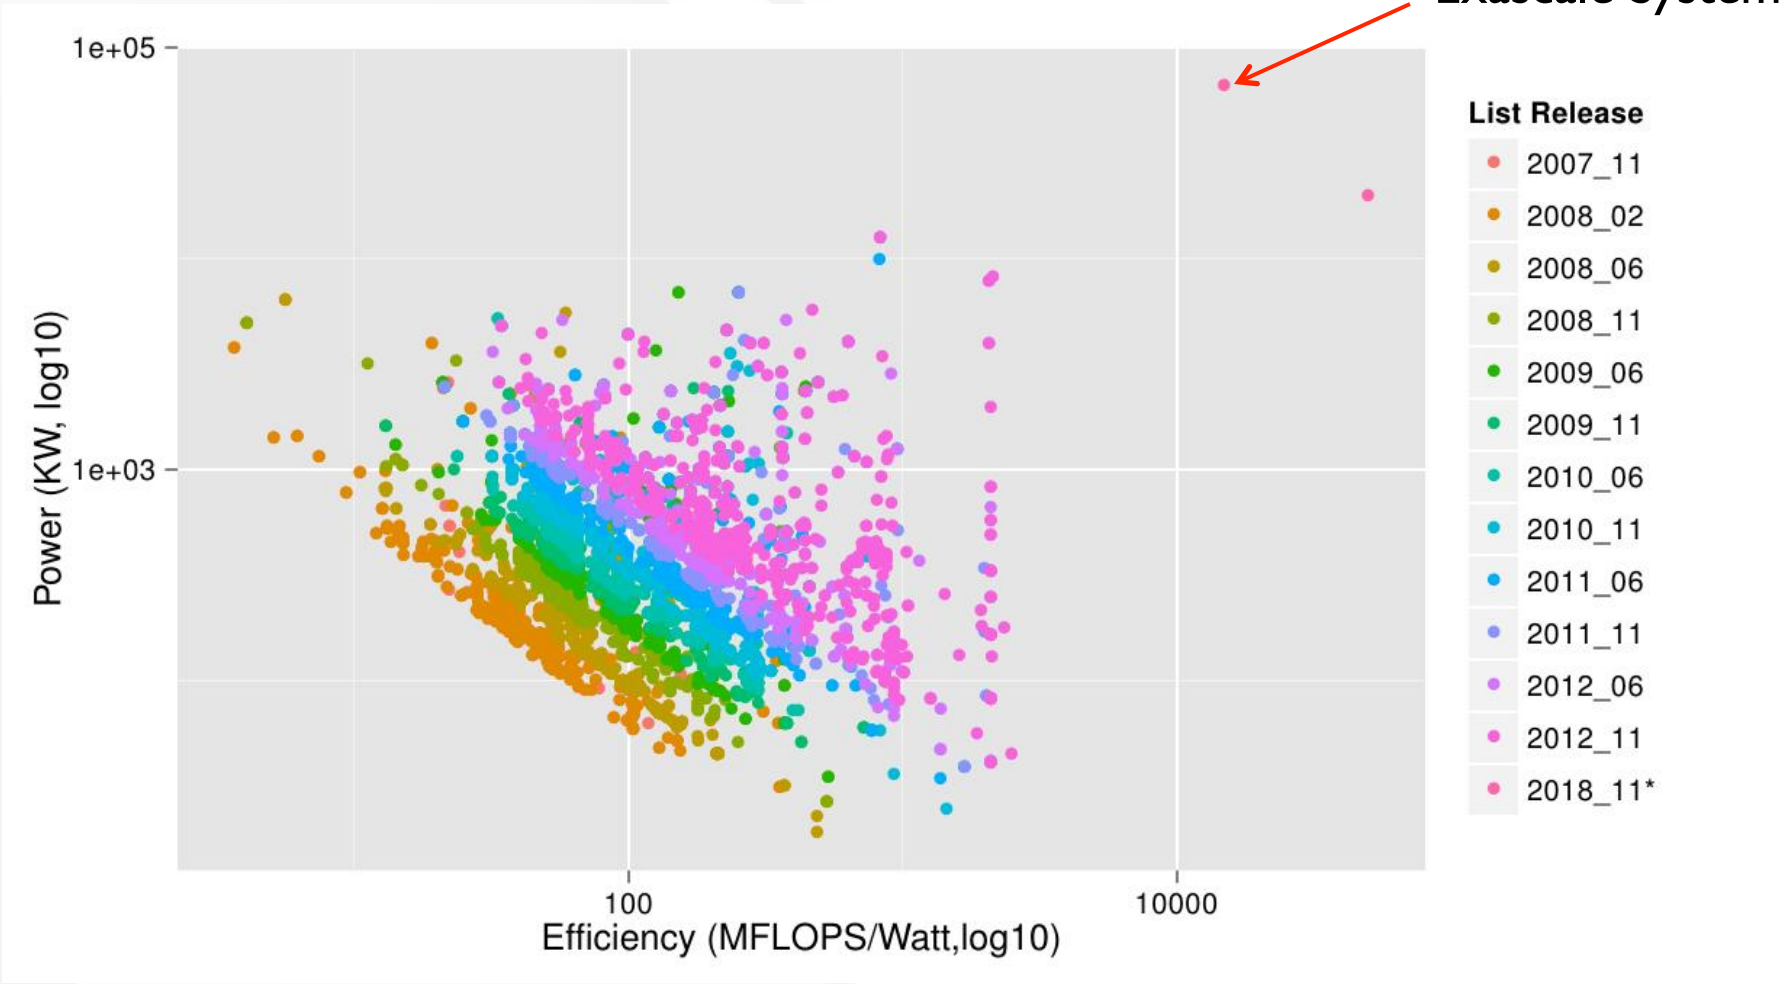

Distance from Exaflop Goal

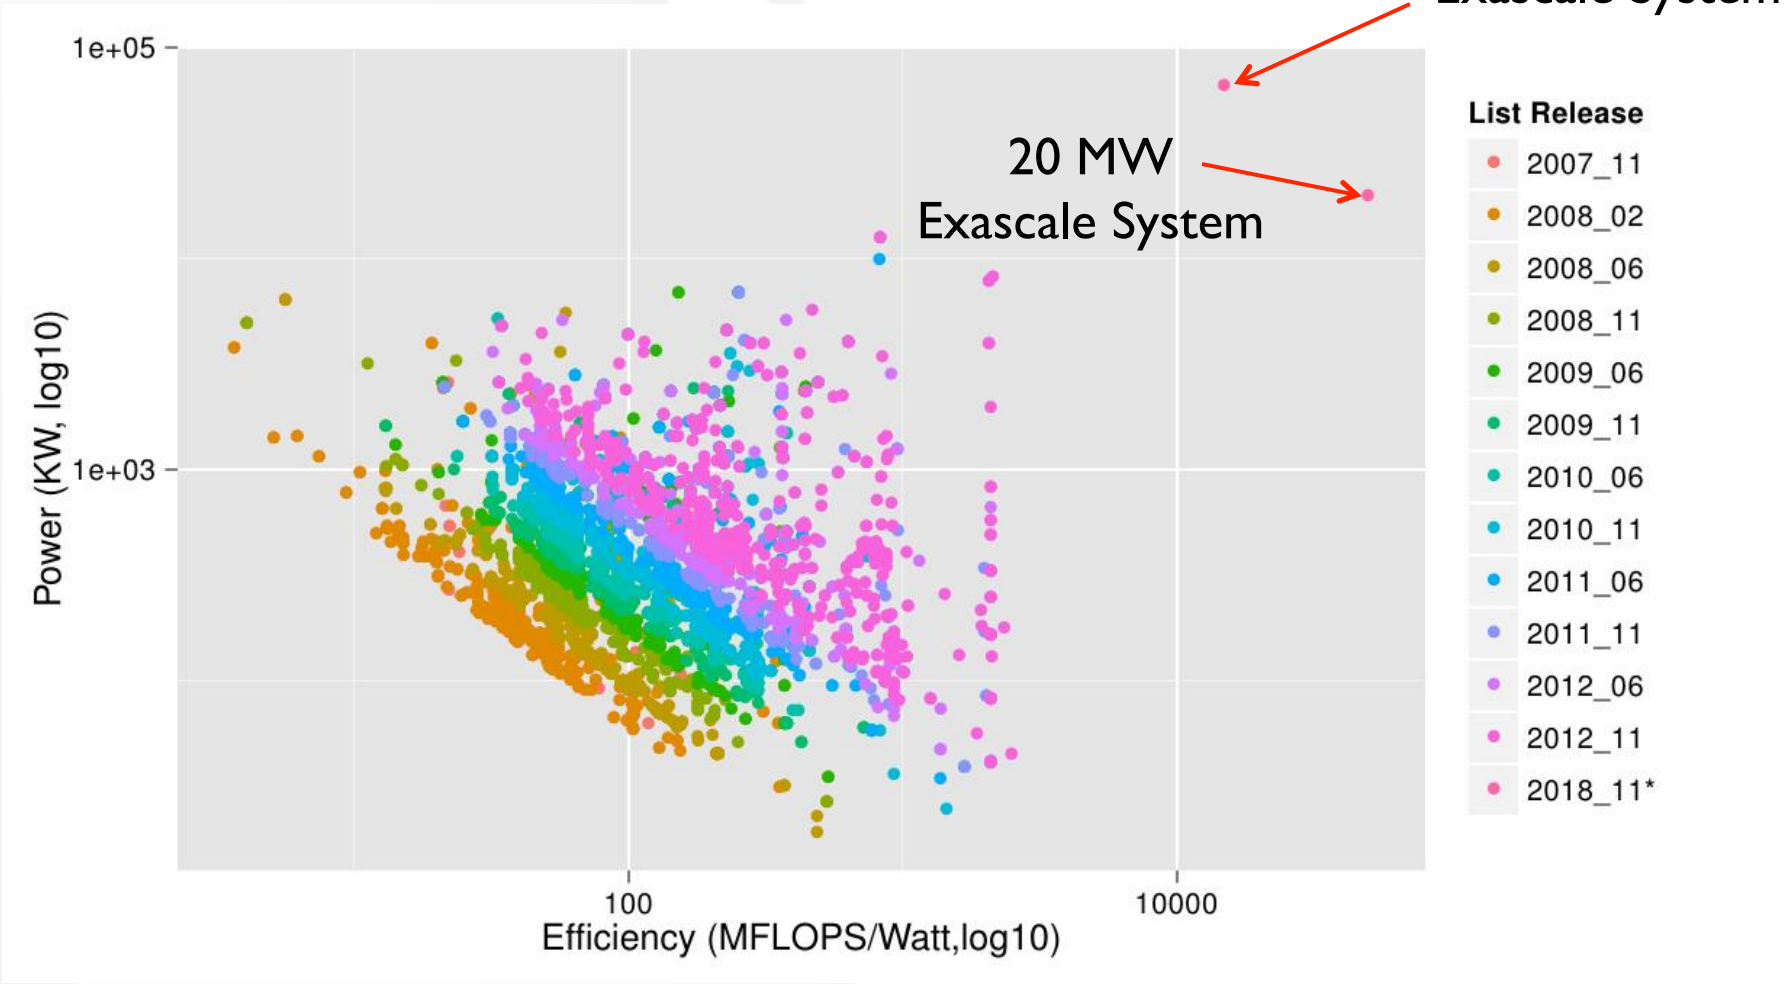

Distance from Exaflop Goal

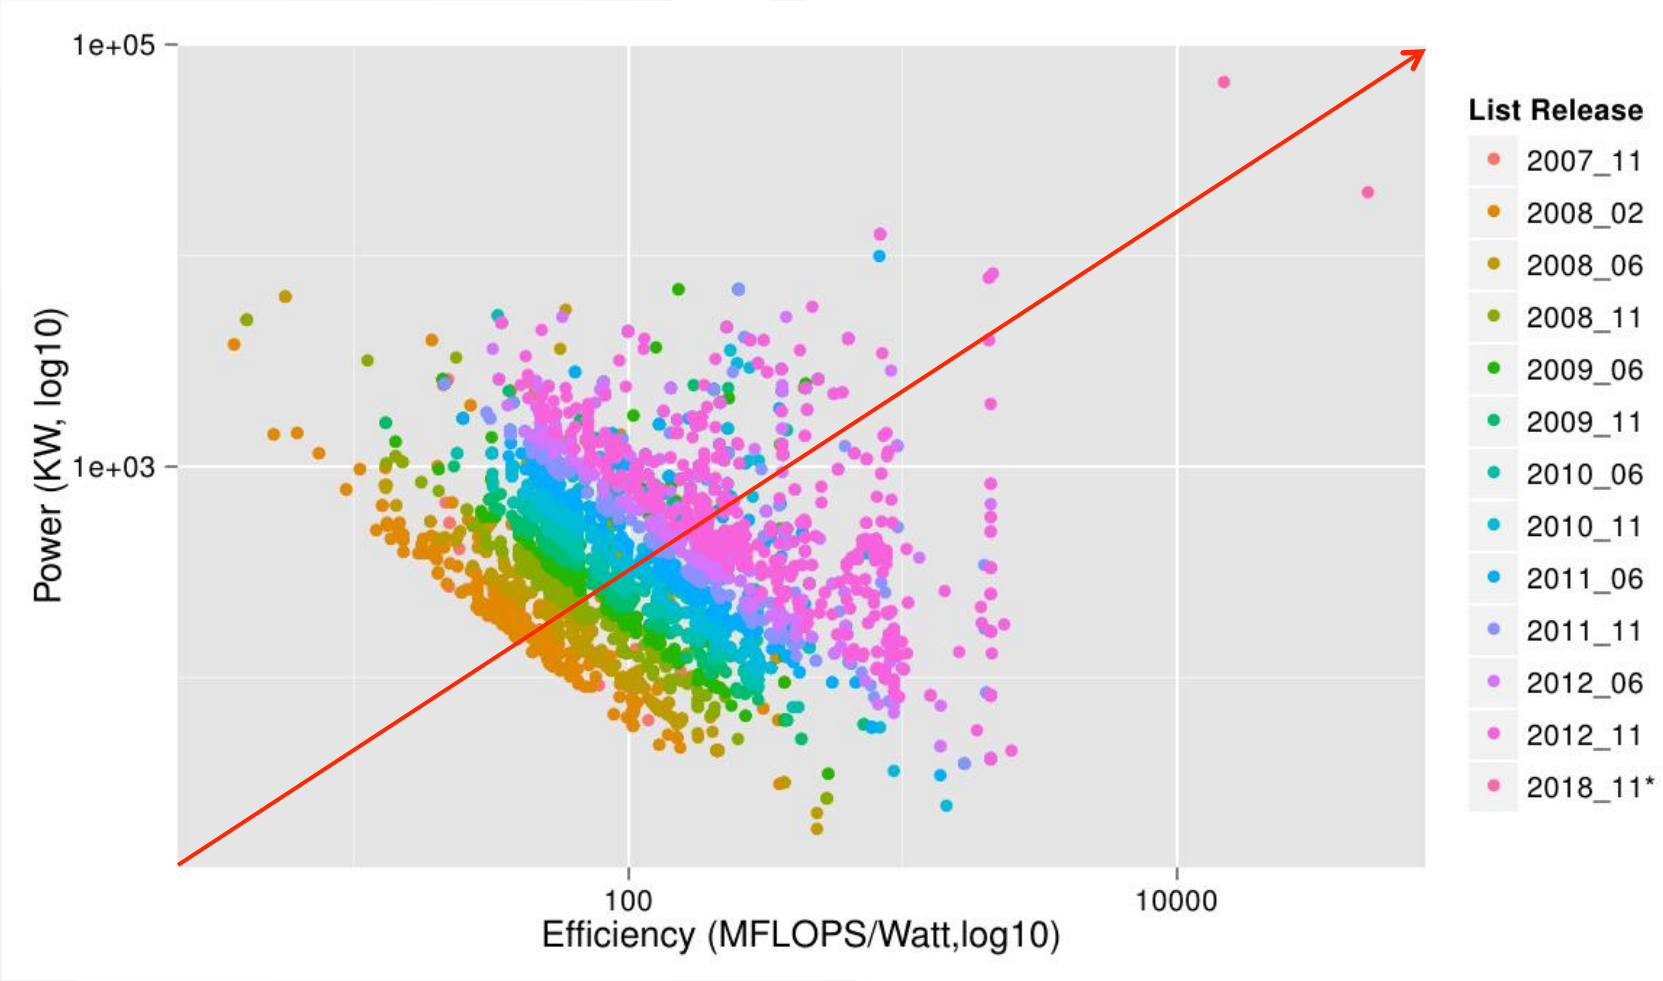

Quantifying Distance from Exaflop Goal:

The Exascalar Metric

- Exascalar
  - Holistic measure of distance from the exaflop goal
  - Simultaneous consideration of performance and efficiency

Quantifying Distance from Exaflop Goal:

The Exascalar Metric

- Exascalar
  - Holistic measure of distance from the exaflop goal
  - Simultaneous consideration of performance and efficiency
- Distance in X-Y coordinates

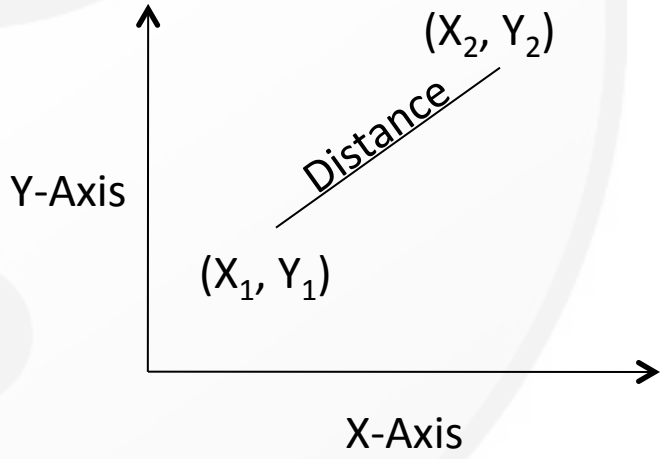

Quantifying Distance from Exaflop Goal:
The Exascalar Metric

- Exascalar
  - Holistic measure of distance from the exaflop goal
  - Simultaneous consideration of performance and efficiency
- Distance in X-Y coordinates

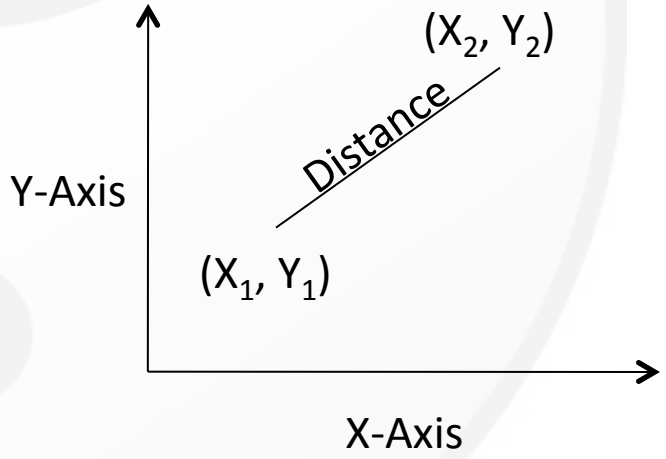

$$\text{Distance} = \text{Sqrt}((X_2-X_1)^2 + (Y_2-Y_1)^2)$$

Quantifying Distance from Exaflop Goal:

The Exascalar Metric

- Exascalar
  - Holistic measure of distance from the exaflop goal
  - Simultaneous consideration of performance and efficiency
- Distance in Energy Efficiency (EE) – Performance (Perf) coordinates

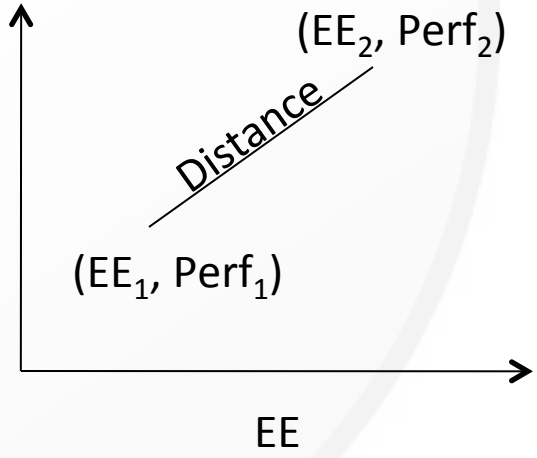

$$Distance = Sqrt((EE_2-EE_1)^2 + (Perf_2-Perf_1)^2)$$

Quantifying Distance from Exaflop Goal:

The Exascalar Metric

- Exascalar
  - Holistic measure of distance from exaflop goal
  - Takes both performance and power into account

$$\text{Exascalar} = \text{Sqrt}\left[\left(\log(\text{Perf}_{\text{System}}) - \log(\text{Perf}_{@20\text{MW\_Exaflop}})\right)^2 + \left(\log(\text{EE}_{\text{System}}) - \log(\text{EE}_{@20\text{MW\_Exaflop}})\right)^2 \right]$$

where Perf is Performance and EE is Energy Efficiency

Quantifying Distance from Exaflop Goal:
The Exascalar Metric

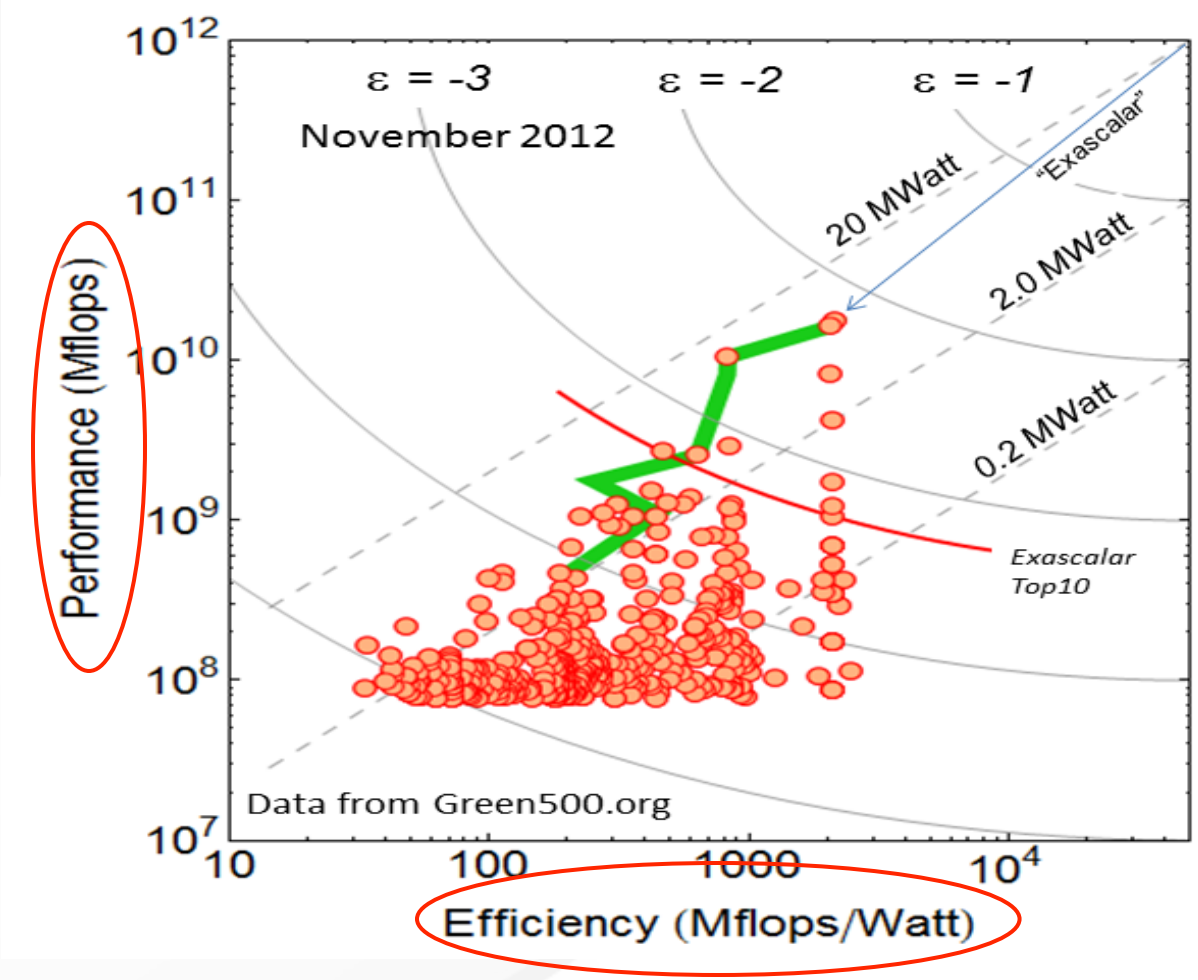

Quantifying Distance from Exaflop Goal:
The Exascalar Metric

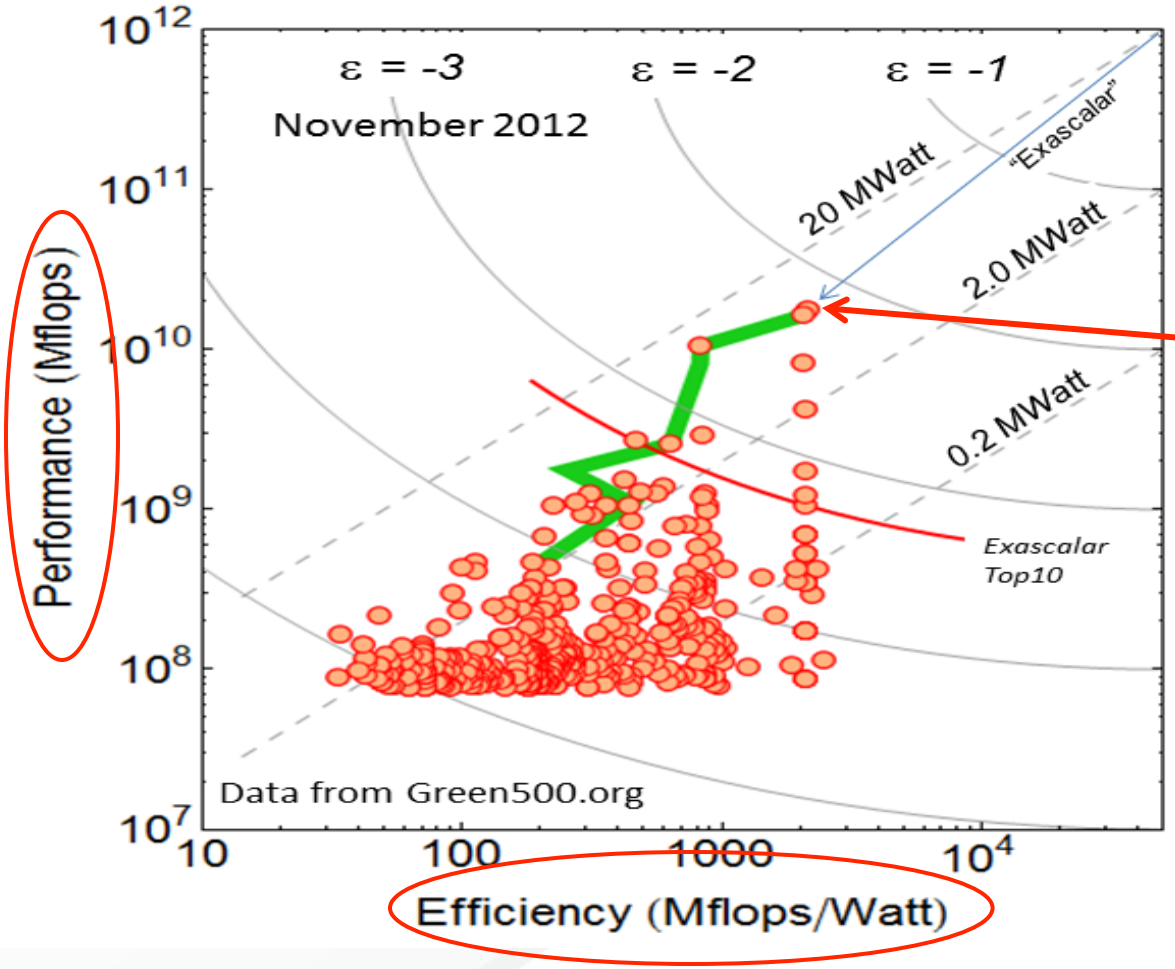

Top Exascalar Nov. 2012

Quantifying Distance from Exaflop Goal:
The Exascalar Metric

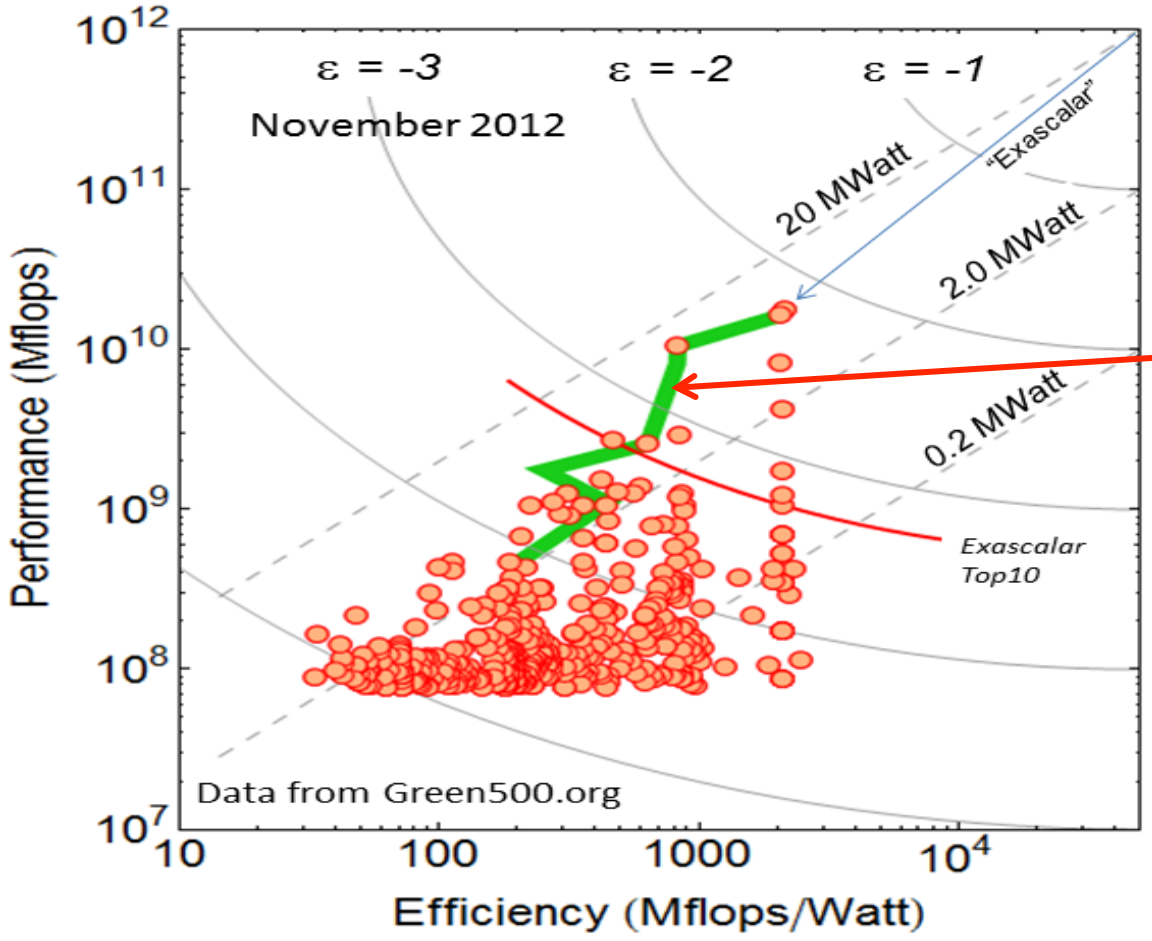

Top Exascalar Trend

Quantifying Distance from Exaflop Goal:

The Exascalar Trend

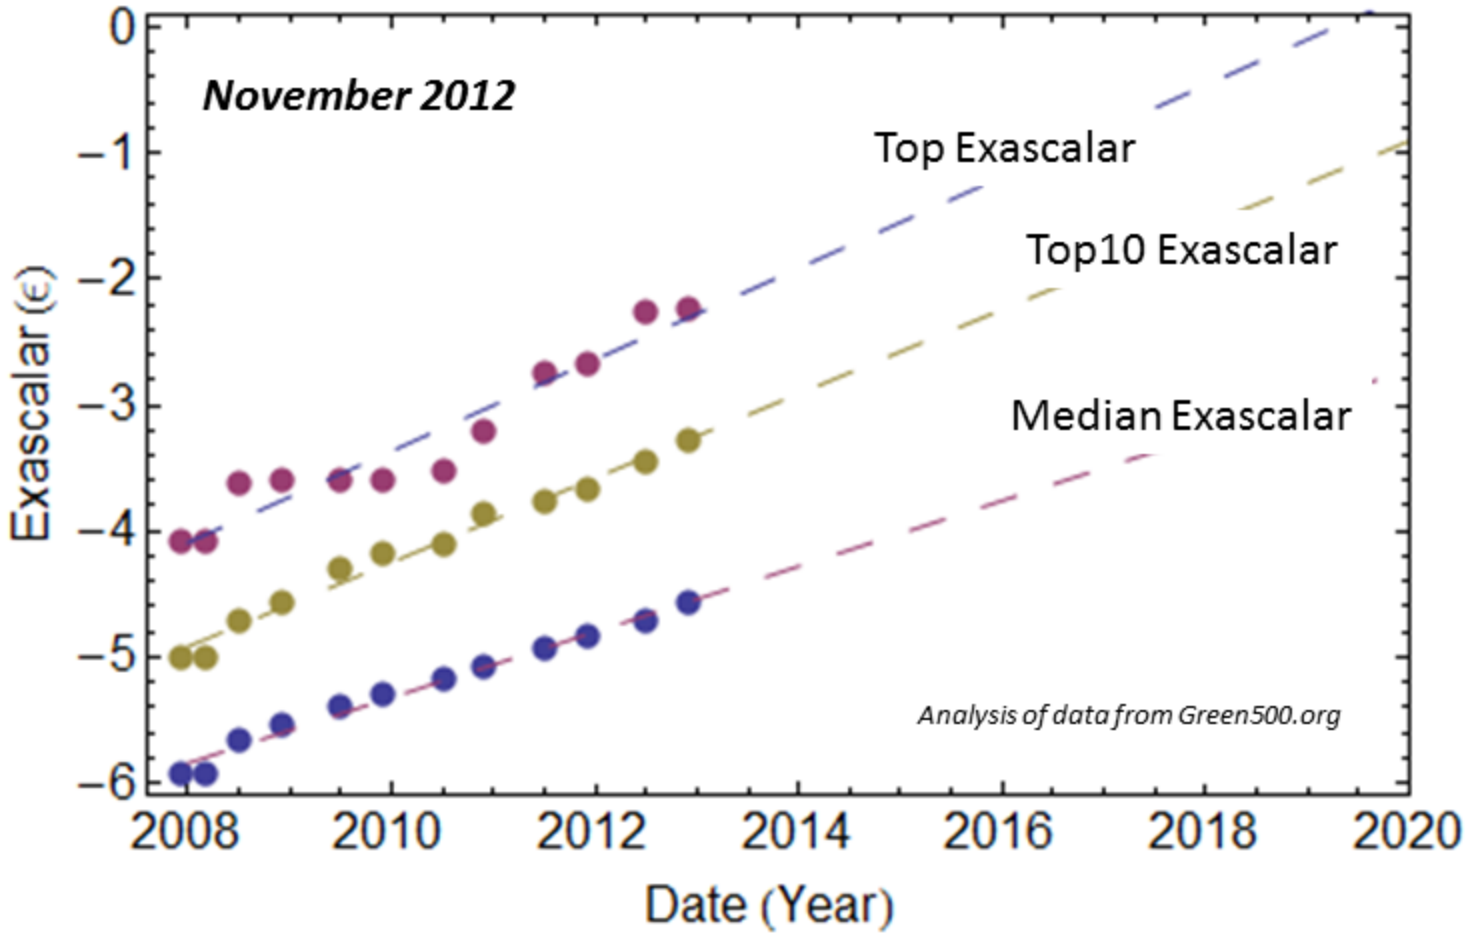

Quantifying Distance from Exaflop Goal:
The Exascalar Triangle

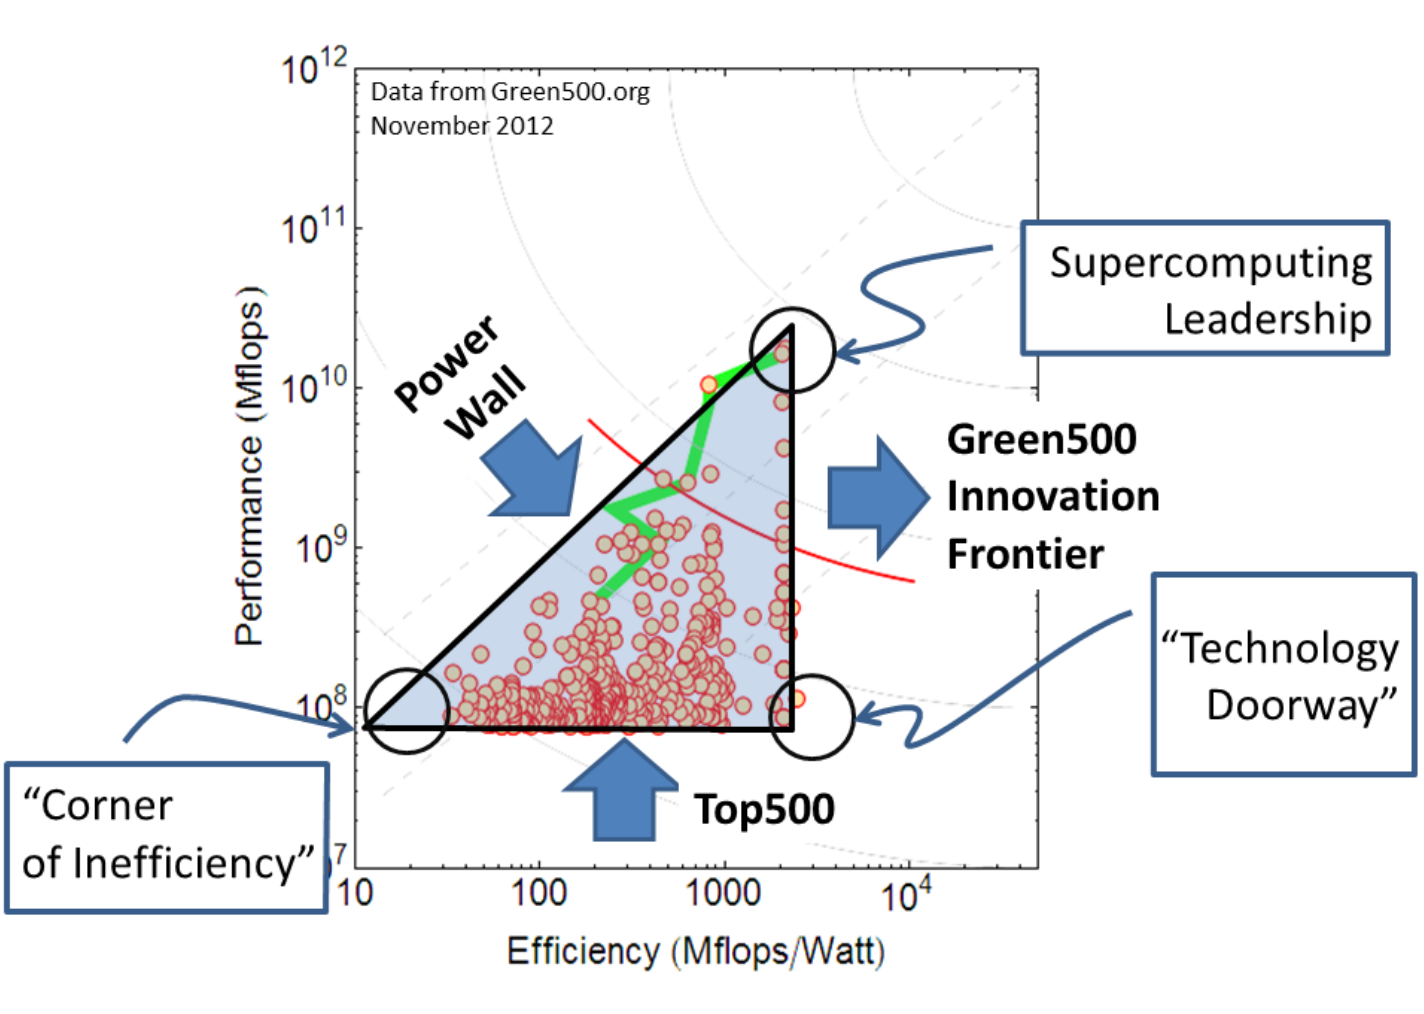

Conclusion

- Analysis of Green500
  - Energy efficiency continues to improve
  - Heterogeneous and custom-built systems are more energy efficient

Conclusion

- Analysis of Green500
  - Energy efficiency continues to improve
  - Heterogeneous and custom-built systems are more energy efficient
- Tracking Koomey's law for HPC
  - Systems would reach 5.78 and 6.93 GFLOPS/watt in 2018 and 2020

Conclusion

- Analysis of Green500
  - Energy efficiency continues to improve
  - Heterogeneous and custom-built systems are more energy efficient
- Tracking Koomey's Law for HPC
  - Systems would reach 5.78 and 6.93 GFLOPS/watt in 2018 and 2020
- Exascalar: Quantifying Distance fom Exaflop Goal
  - Current systems are 2.2 orders of magnitude away from exaflop goal

Conclusion

- Analysis of Green500
  - Energy efficiency continues to improve
  - Heterogeneous and custom-built systems are more energy efficient
- Tracking Koomey's Law for HPC
  - Systems would reach 5.78 and 6.93 GFLOPS/watt in 2018 and 2020
- Exascalar: Quantifying Distance fom Exaflop Goal
  - Current systems are 2.2 orders of magnitude away from exaflop goal

Acknowledgements

The Green500 Team: Balaji Subramaniam and Thomas Scogland

Contact: info@green500.org

Winston Saunders, Intel → The Exascalar Metric