Using the cloud for high performance computing: my perspective

Sai Gautam Gopalakrishnan

Materials Engineering, Indian Institute of Science

saigautamg@iisc.ac.in; https://sai-mat-group.github.io

SSWR-IMechE Quarterly Lecture
Department of Aerospace Engineering, Indian Institute of Science
Dec 3, 2022

Acknowledgments


Sai Gautam Gopalakrishnan
Principal Investigator


Reshma Devi Parthasarathy
Ph. D. student


Debolina Deb
Ph. D. student


Dereje Bekele Tekliye
Ph. D. student


Ankur Srivastava
Ph. D. student (co-advised)


Vijay Choyal
Institute of Eminence Fellow


Adilakshmi Chirumamilla
Masters student


Sachin Kumar
Masters student


Javeed Ahmad Dar
Ph.D. student


Pritam Ghosh
Project Associate


Sougat Purohit
Ph.D. student


Ankit Kumar Gupta
Masters student

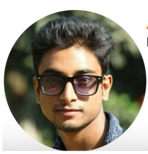
Arnab Pyne
Masters student


Hasna Sabreen
Masters student

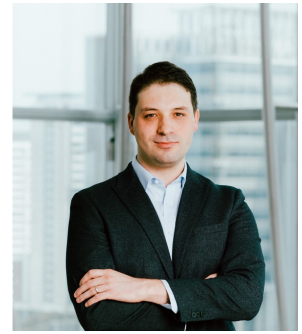
Piero Canepa
NUS
+CaRe group


Prabeer Barpanda
MRC, IISc
+FaMaL


Phani Motamarri
CDS, IISc


May 2022

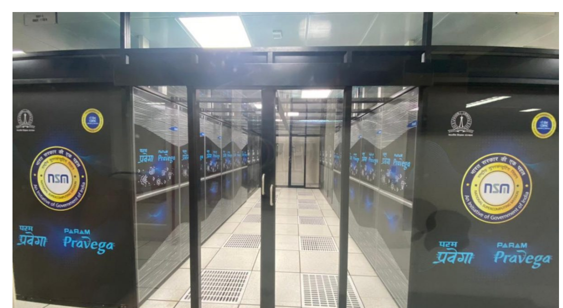

+other collaborators, computing resources, …

Climate change is here

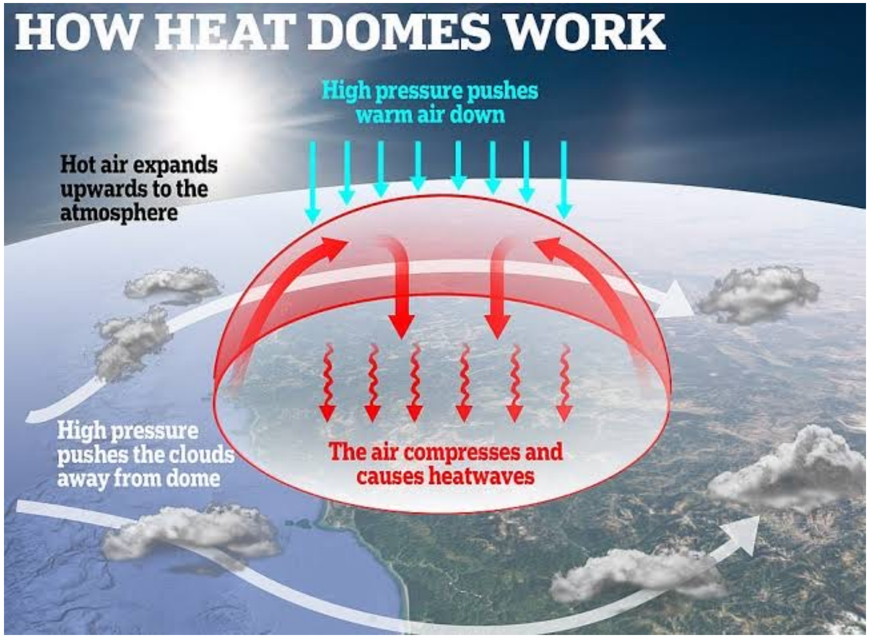

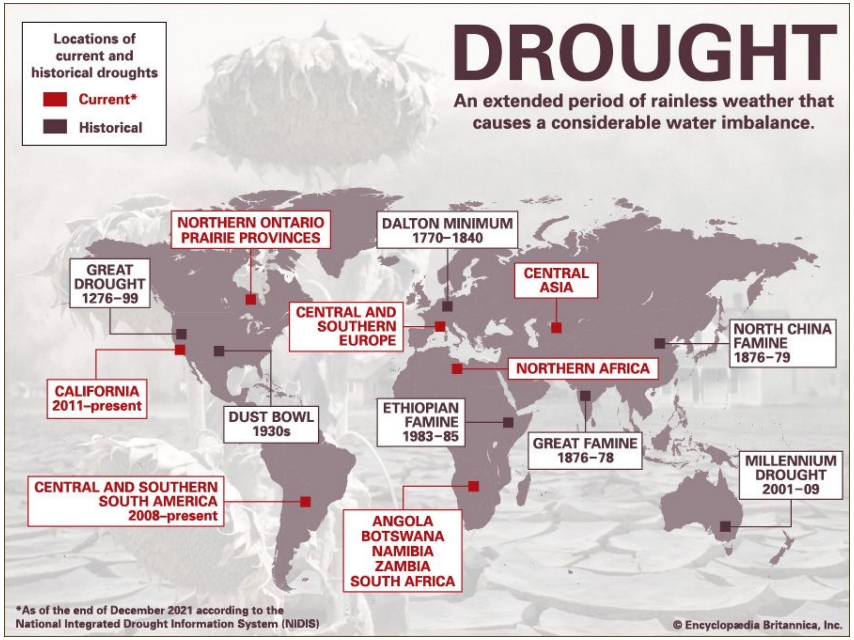

Heat waves and wildfires

Droughts and floods

Non-fossil-fuel options for mitigating climate change

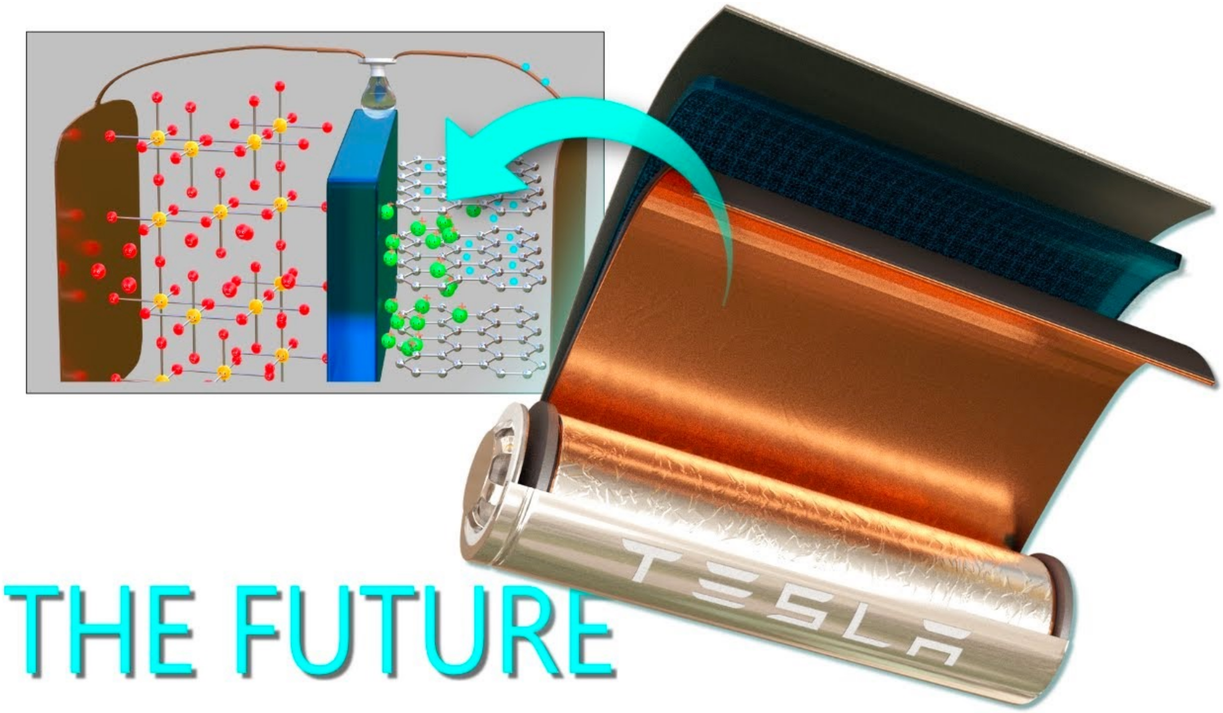

When the sun doesn’t shine or the wind doesn’t blow

Materials form the performance-bottlenecks of most renewable energy devices: **how do we understand and improve the material bottlenecks?**

- How can we improve the amount of energy stored (i.e., energy density) in a battery?
- Can we find materials better than Si as photovoltaics?
- Are there better thermochemical water splitters?

We work broadly on energy materials

Design better electrodes and solid electrolytes

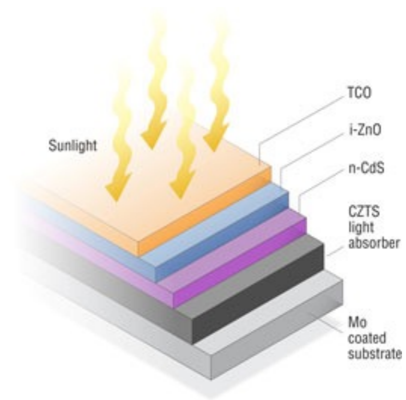

Develop better light-absorbing semiconductors

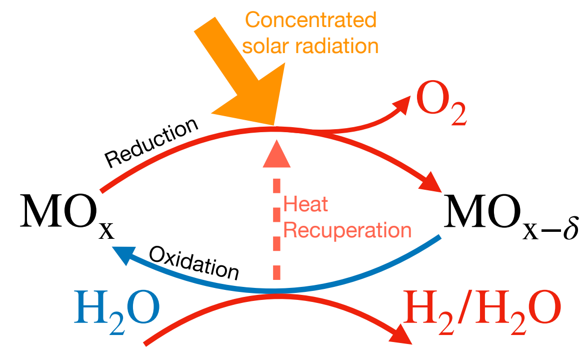

Identify better thermochemical H_2O-splitters

Identify novel materials for applications

- Use high-throughput screening +/- machine learning (**ML**) to generate key performance-determining descriptors
- Collaborate with experimental groups for validation of theoretical predictions

Understand underlying materials phenomena better

- In-depth studies focused on thermodynamic, kinetic or electronic behavior of a given (candidate) material
- Predict ”stable” configurations, mobility bottlenecks, suppress/enhance defect formation, etc.

Make theory better

- Benchmark existing theoretical models against experimental data to identify best ones
- Develop better models for simulating complex phenomena

We do theory, computations, & ML

Density functional theory (**DFT**): (Approximately) predict material properties

- Structural (lattice parameters)
- Thermodynamic (voltages, stabilities, phase diagrams)
- Electronic (band gaps)
- Magnetic (oxidation states, magnetic moments)
- High-throughput “screening”

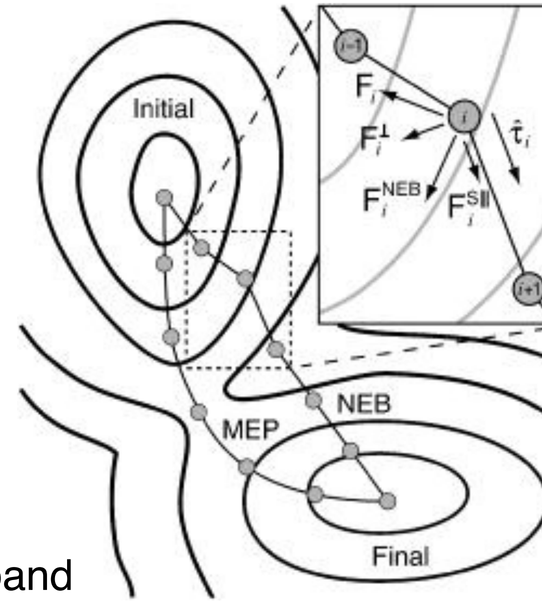

Nudged elastic band (**NEB**): migration barriers

ML: regressions and interatomic potentials

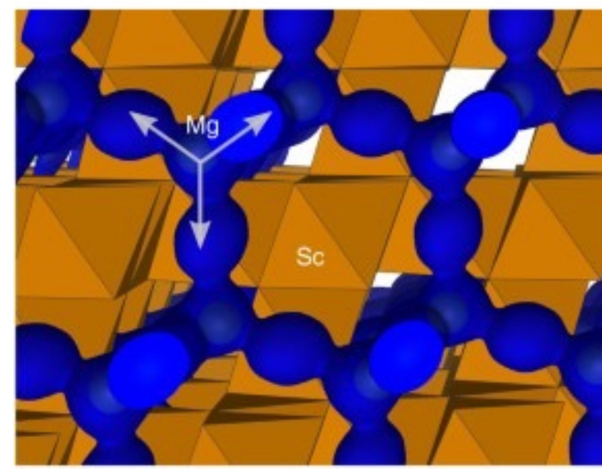

Ab initio and classical (ML) molecular dynamics: kinetic properties

Examples of recent works

Screening cathodes for Ca-batteries

Ca-batteries: beyond Li-ion system(s) that can mitigate challenges with current Li-ion batteries

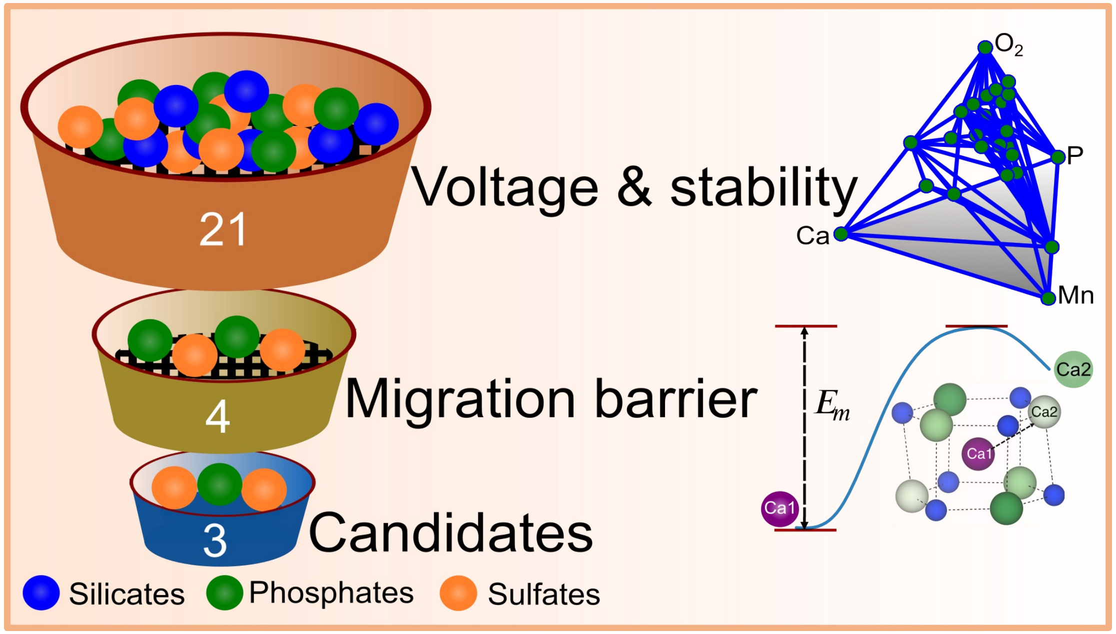

“NaSICONs”
$Ca_xM_2(ZO_4)_3$

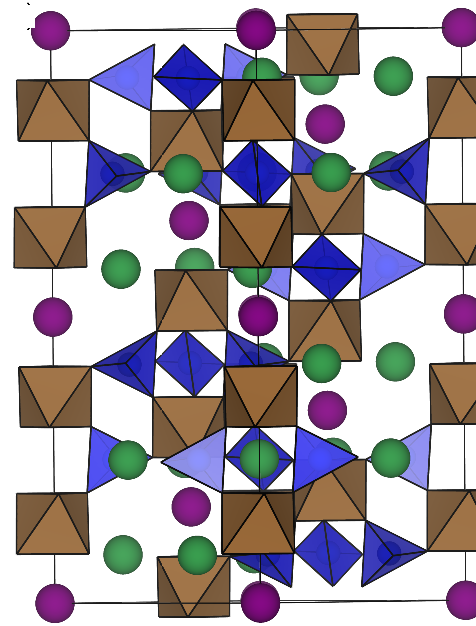

High-throughput DFT calculations: 3 candidates

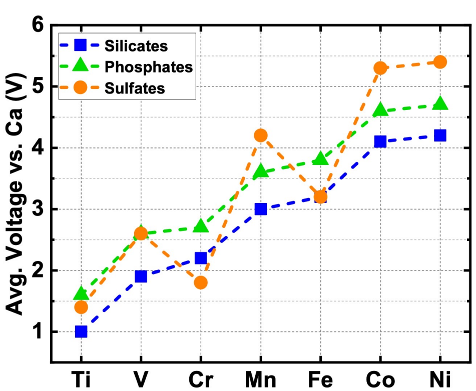

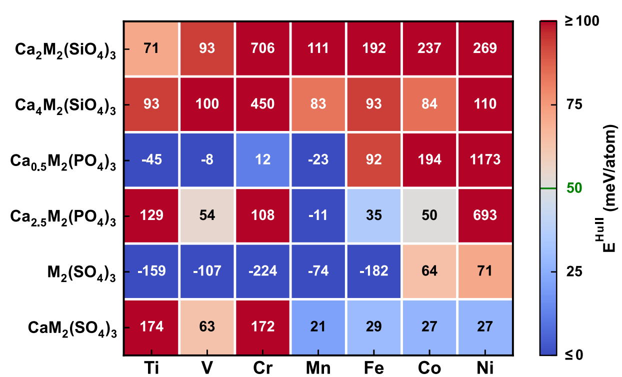

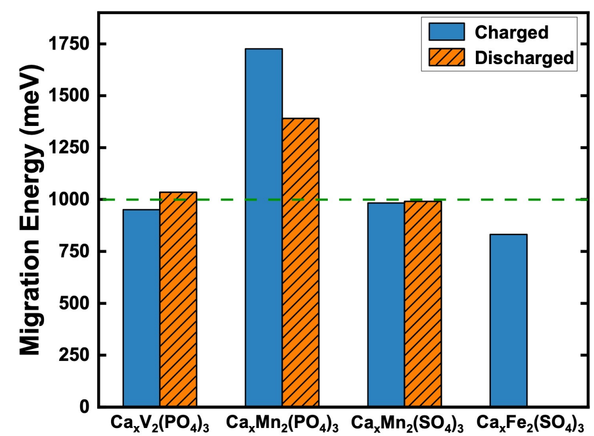

$Ca_xV_2(PO_4)_3$, $Ca_xMn_2(SO_4)_3$, and $Ca_xFe_2(SO_4)_3$

D.B. Tekliye, G.Sai Gautam, et al., **Chem. Mater. 2022**, *34*, 10133-10143

Pnictides as possible photovoltaics

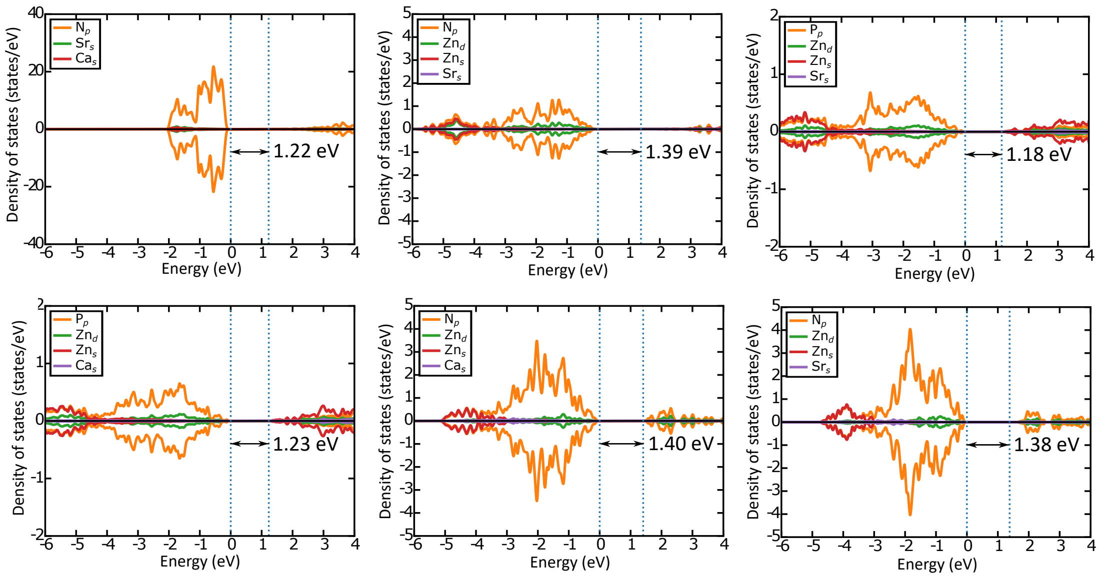

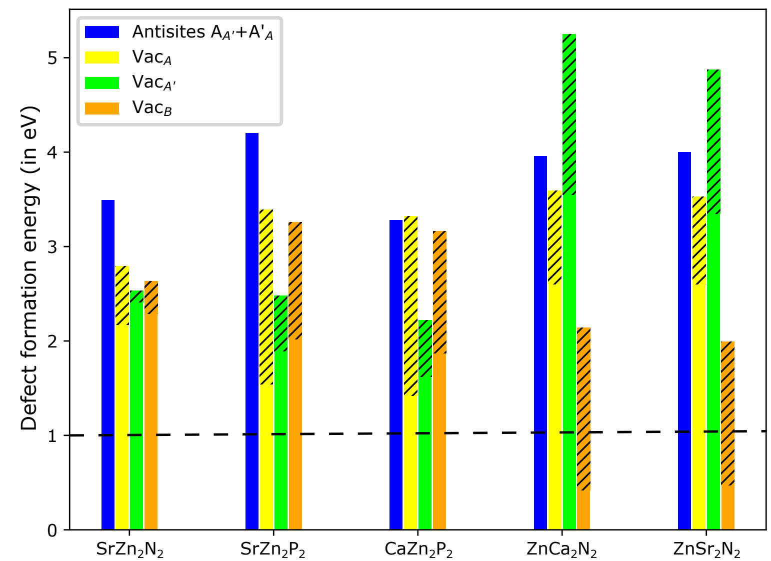

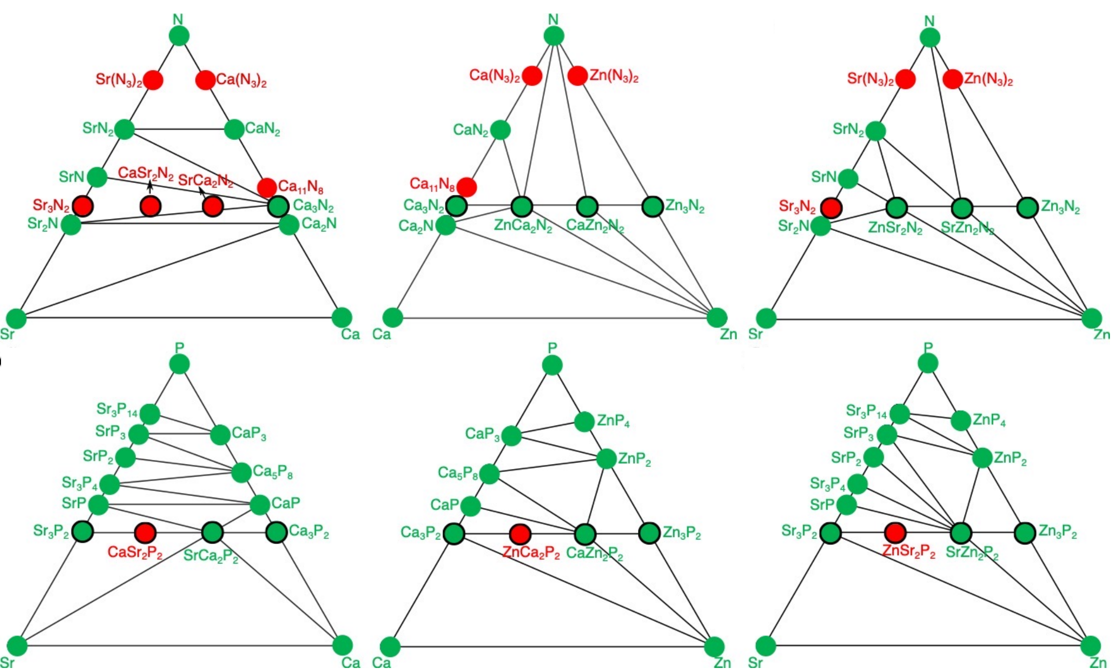

Band gap estimates
+ 0 K thermodynamic stability screening
+ resistance to point defects
= candidate beyond-Si photovoltaics

$SrZn_2N_2$, $SrZn_2P_2$, and $CaZn_2P_2$: predicted candidates

J. Kumar, and G.Sai Gautam, **under review**

Using ML for predicting water-splitters

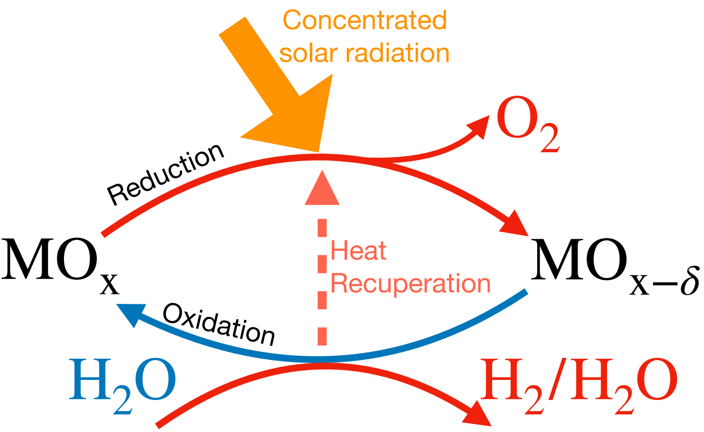

Need oxides with "optimal" oxygen vacancy formation energies

Build a "simple" ML model with "physically intuitive" descriptors

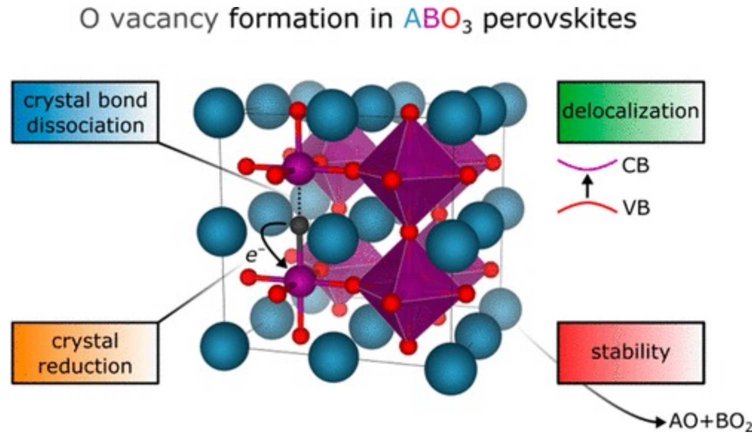

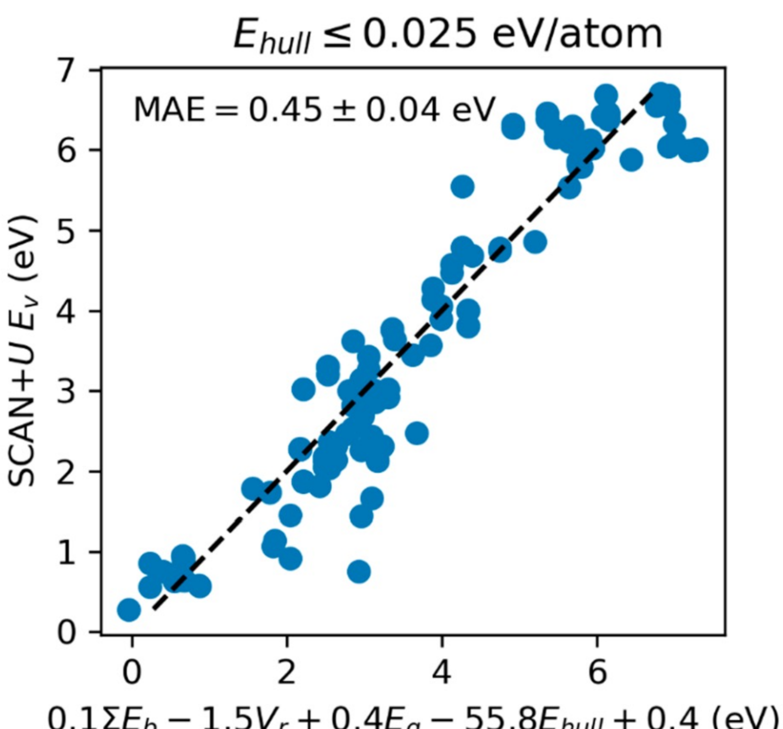

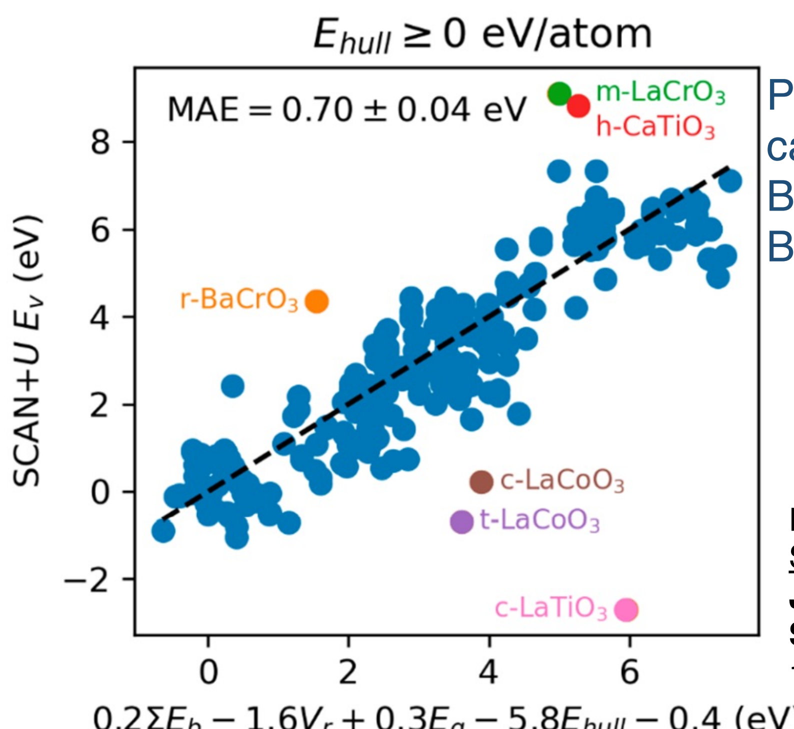

Promising new candidates:
$BiFeO_3$
$BiCoO_3$

R.B. Wexler, G. Sai Gautam et al., **J. Am. Chem. Soc. 2021**, *143*, 13212-13227

Understand (possible) anionic redox in bisulfate Li-cathodes

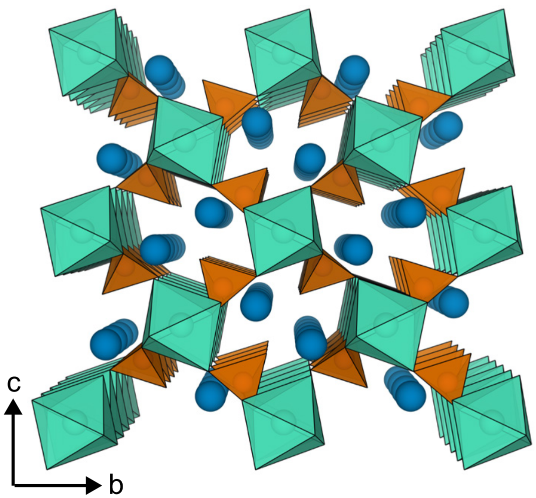

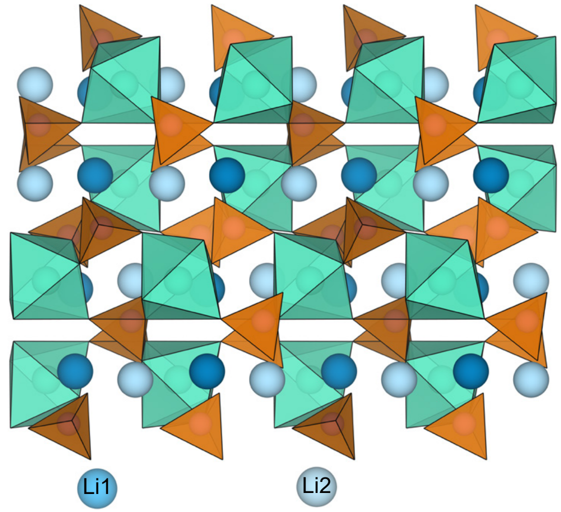

$Li_2M(SO_4)_2$, M = Mn, Fe, Co, Ni

Two polymorphs: orthorhombic and monoclinic

Both Li can be theoretically removed

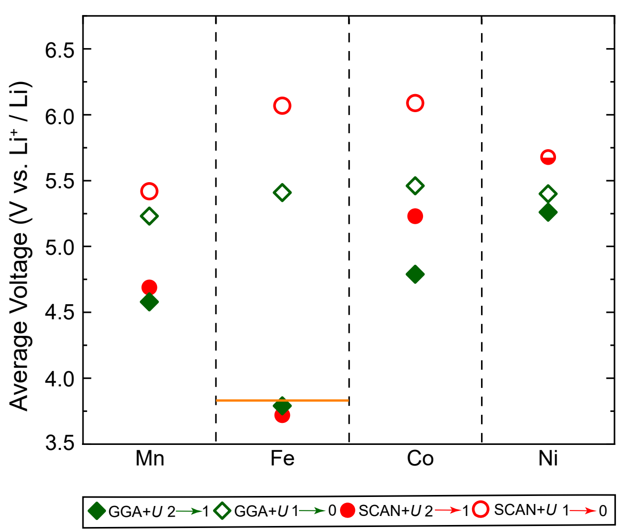

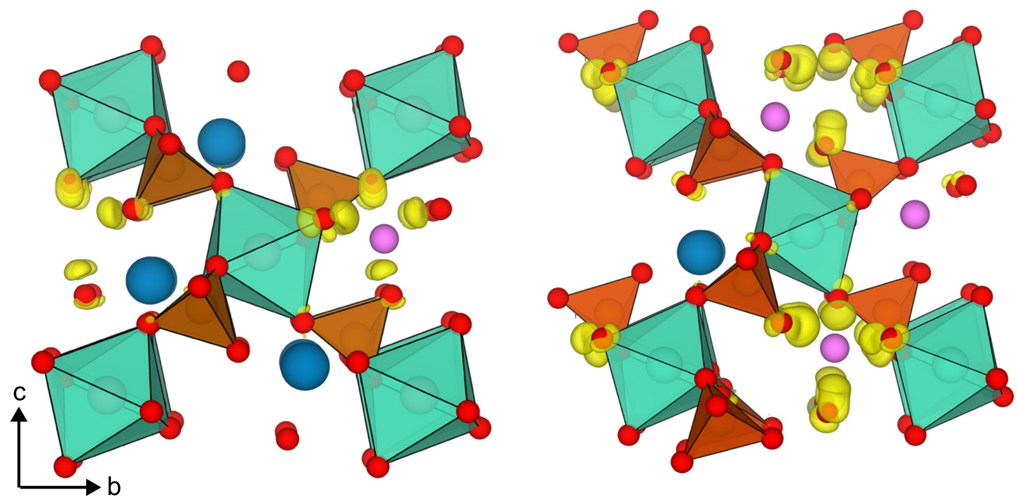

Robust evidence for anionic redox in a polyanionic framework

P.K. Jha, S. Singh, M. Shrivastava, P. Barpanda and G. Sai Gautam, **Phys. Chem. Chem. Phys. 2022**, Advance Article

Use (ML) molecular dynamics to understand interfacial transport bottlenecks

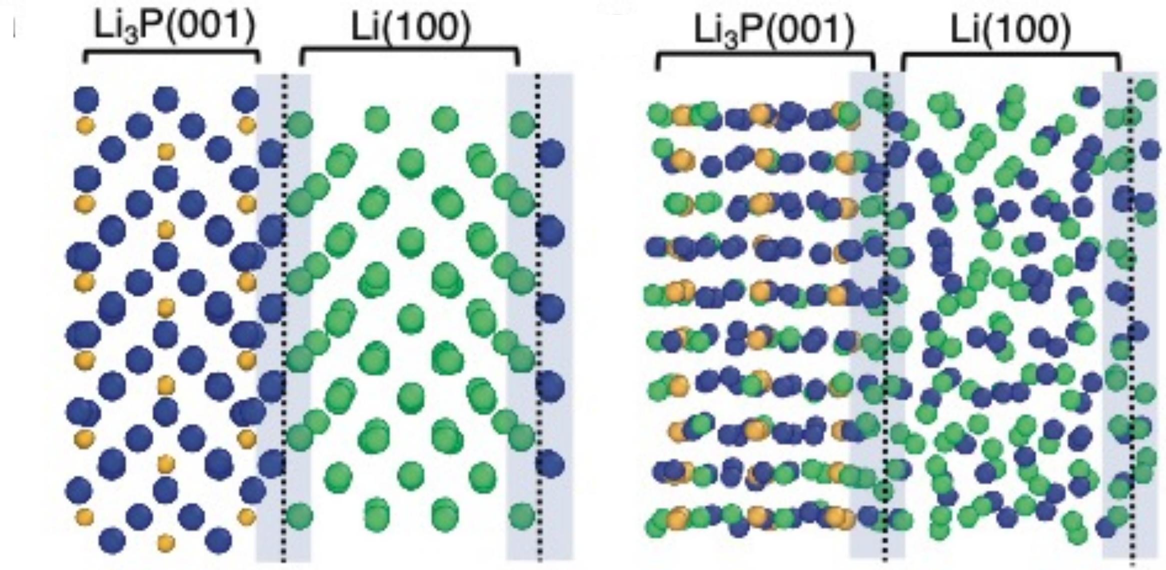

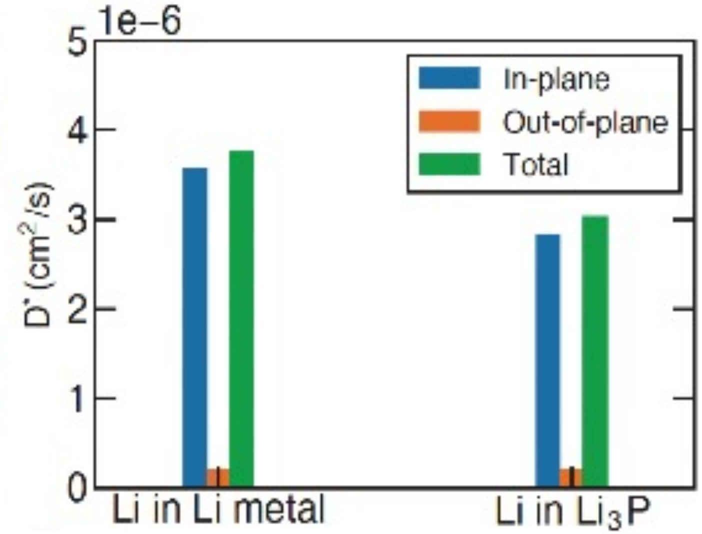

~8000 atoms; 10 ns

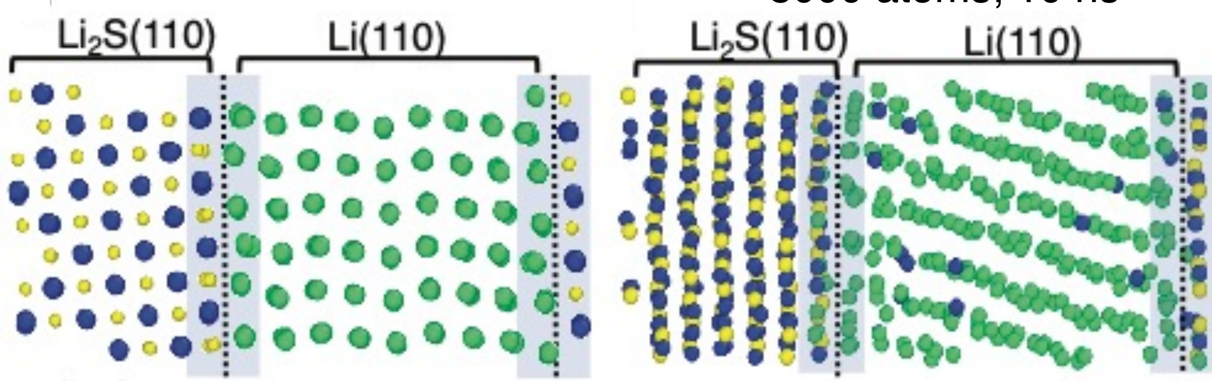

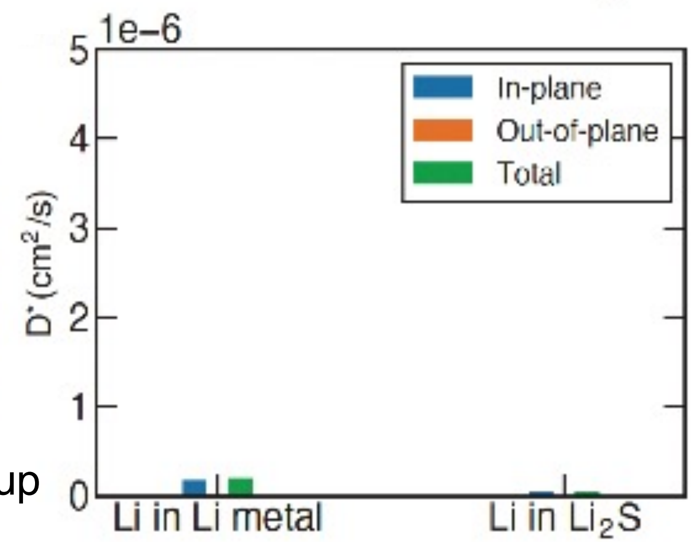

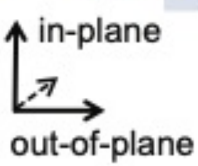

Model Li interface with possible argyrodite decomposition products to explain impedance build-up

Li_3P: conducive to Li-transport across interface

J. Wang, A.A. Panchal, G.Sai Gautam, and P. Canepa, **J. Mater. Chem. A 2022**, *10*, 19732-19742

Which ”functional” predicts migration barriers well?

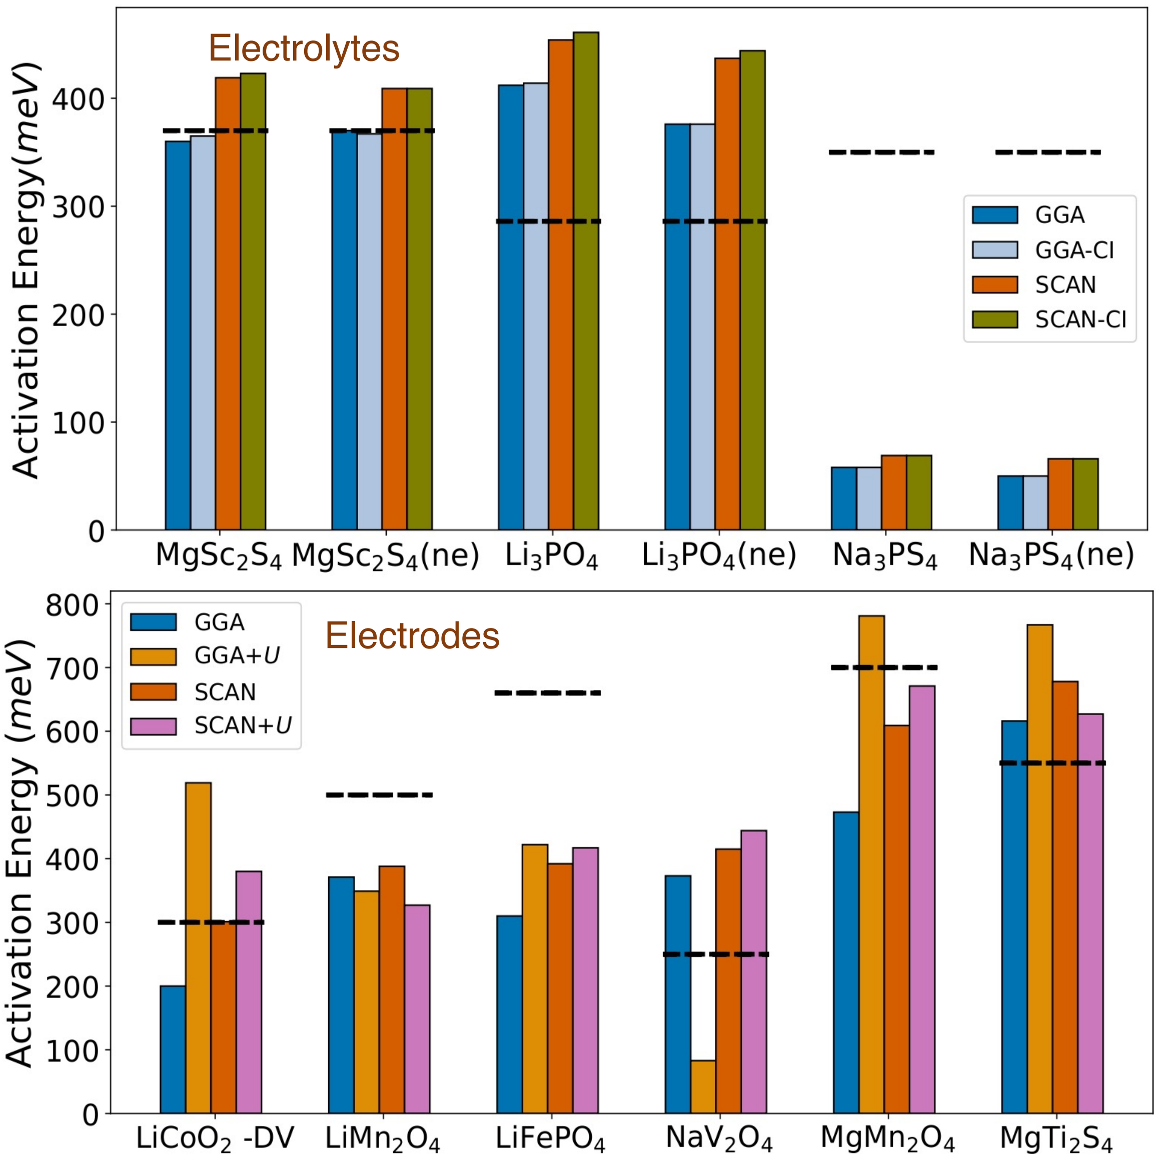

Migration barriers: crucial for power performance

Which exchange-correlation functional is best suited for migration barrier predictions in battery materials?

Strongly constrained and appropriately normed (SCAN) more accurate on average

- Describes right electronic structure
- Computationally expensive and difficult to converge
- Generalized gradient approximation (GGA): not bad either

R. Devi, B. Singh, P. Canepa, and G.Sai Gautam, **npj Comput. Mater. 2022**, *8*, 160

Usage statistics and cost of computations

Where do we run our computations?

- Resources we have dedicated to our group
  - Computing cluster (512 cores, up since May 1, 2021)
    - Currently has 768 cores
  - Workstation (48 cores, up since May 1, 2021)
  - 9 high-performance desktops (16 cores each, up since ~March 2021)

- Resources we have that are shared
  - Param Pravega (maintained by SERC)
  - Roddam Narasimha Cluster (maintained by SERC)

- Other resources
  - Amazon web services (AWS) cloud high-performance computing (Jan-Apr 2021)
  - Collaborators

Our DFT calculations are parallelized, memory-heavy, run on CPUs
Machine learning jobs are usually serial, light, run on CPUs (but portable to GPUs)

How much computations did we run?

- Usage on SERC: ~4.5 Million CPU hours
  - Period: Aug 2021-July 2022
  - Regular charges (2 accounts): **INR 3,40,000**
    - Slab system
    - Typical queue time: ~3 weeks for a "small" (< 512 core) calculation
  - SERC has a high-priority queue
    - Typical wait time: 2-3 days for a small job
    - Significantly more expensive: INR 1.5/(CPU hour)
    - If our usage was run entirely on high-priority, our charges would have been **INR 67,50,000**

- Usage on computing cluster: ~4.1 Million CPU hours
  - Assuming ~1 month downtime due to power cuts
  - Dedicated to group, no explicit payments

For any fully computational group, ~10 Million CPU hours is a good ball-park usage per year

Budgets for computations: Scenario 1

- Capital cost for a 1024-core cluster
  - INR 1,40,00,000 – INR 1,60,00,000
  - Variation due to technical specifications
  - Typical life of a cluster: 5 years
    - Can go to 10 years with good maintenance
    - Warranty given for 3 years
  - Annualized capital cost: ~INR 32,00,000 (upper limit)
    - Not taking into account inflation/interest

Total capital+operating budget for a standalone cluster per year (assuming 5 year life cycle): ~INR 39,00,000

- Air conditioning (highly approximate estimate)
  - 5-ton split air-conditioning unit may be needed
    - Capital cost: ~INR 5,00,000
    - Annualized cost (for 5 years): INR 1,00,000

At Indian Institute of Science, capital+operating costs (assuming 10 year life cycle): ~INR 16,50,000!

- Electricity charges
  - Cluster: ~5500 W on average; ~130 units/day; 4000 units/month
  - Charges: ~INR 4000/month (BESCOM); ~INR 48,000/year
  - Air-conditioning unit's power consumption: ~2200 units/month (~12 hour operation per day)
  - Charges: ~INR 2300/month (BESCOM); ~INR 28,000/year

- Technical assistant for managing cluster: ~INR 5,00,000/year

Budgets for computations: Scenario 2

- Do all calculations at Param Pravega
  - Slab model annual rate for 10 Million CPU hours: **INR 4,70,000**
  - May not be possible to run 10 Million CPU hours in regular queues in a year!
  - Wait time for ~2400 CPU hour job is ~3 weeks
    - Assume running 2400 CPU hours per day
    - For a full year: usage is 876,000 CPU hours only!

- Do all calculations at Param Pravega in high-priority queue
  - Typical wait time is 2-3 days, so can feasibly run 10 Million CPU hours over a year
    - Cost for usage: **INR 1,50,00,000**!
    - Practically impossible

Practically: have to use both local computing clusters and Param Pravega to reach 10 Million CPU hours

Cloud computing solutions can compete for high-performance computing if operating costs are in the INR 17,00,000 – INR 40,00,000 range for 10 Million CPU hours

Basics and costs of cloud computing

What is cloud computing?

Generic term that refers to any "service" that is hosted over the internet

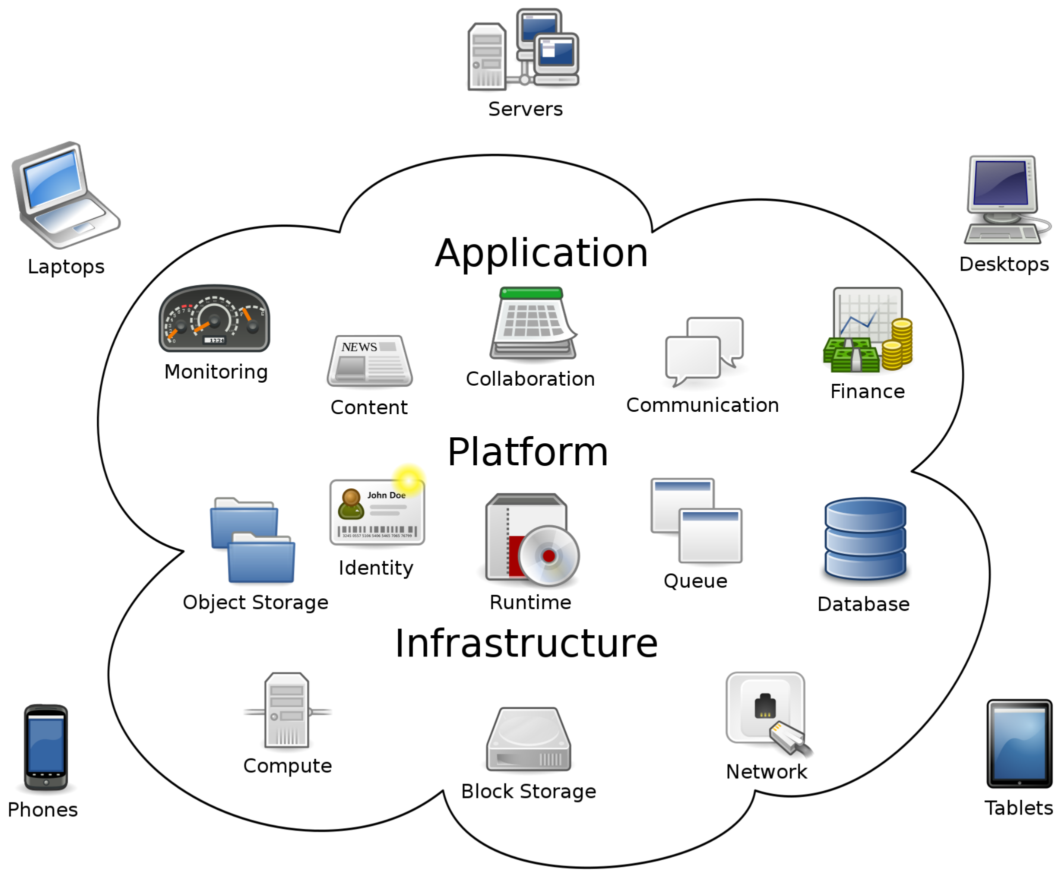

3-types of services common:

- Infrastructure
  - Virtual machines, servers, storage, networking
  - Bundled as a "container", can compile custom software
  - AWS, Azure, Google Compute
- Platform
  - For hosting applications, websites, software development, machine learning
  - Limited access to machines that host platform, can scale
  - Google App Engine, Heroku
- Software
  - Everyday use for consumers
  - Applications such as Dropbox, Gmail, Teams, Zoom, iCloud, etc.
  - Salesforce, SAP, Netsuite

Dominant players: Amazon, Microsoft, Google
Other players: Salesforce, SAP, IBM, Oracle, etc.

Cloud can be private/public/hybrid

- Regular Google is public
- Google "Workspace" is private

Image: https://en.wikipedia.org/wiki/Cloud_computing

Where do we use cloud in our group?

Outlook

OneDrive

GitHub

https://sai-mat-group.github.io

Overleaf

How do you setup a calculation on cloud?

Create a virtual private cloud (VPC) or an equivalent virtual machine: collection of servers
Usually done via a web/command-line interface
Some services operate via APIs

Compile custom software

Run calculations

Pay-per-use

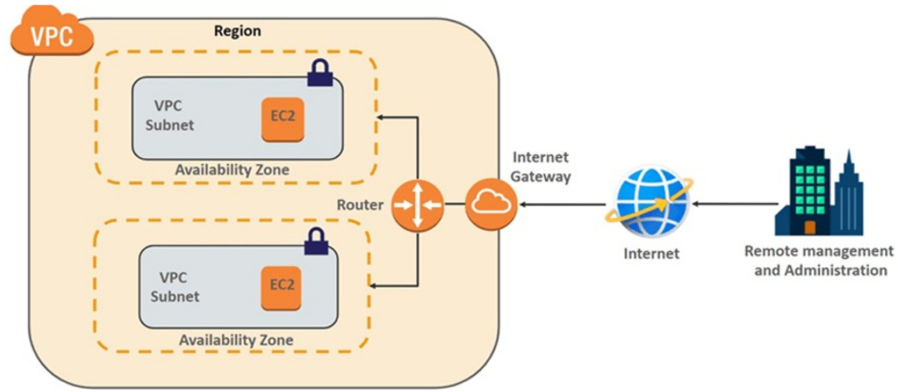

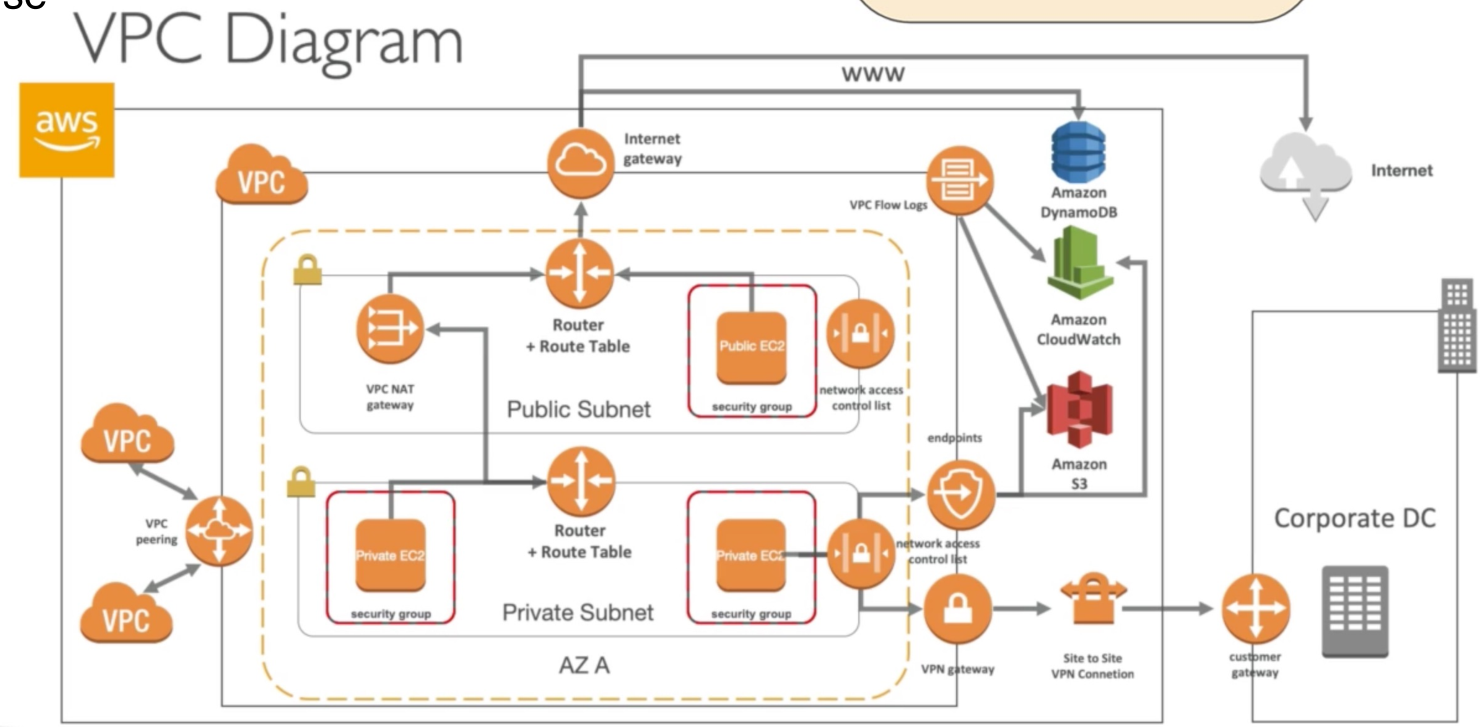

Budget for cloud high-performance computing?

- Varies depending on type of virtual machine requested (number of cores, RAM, etc.)
  - We were charged between $0.12-0.15 per CPU hour (~INR 9.6-12 per CPU hour); 36-core Intel Xeon CPUs
  - So for a usage of 10 Million CPU hours, this will be ~INR 12,00,00,000 (upper limit)
    - This is 3.5-4 times the cost of a typical computing cluster+operating costs
    - Highly impractical

Thus, the current charges levied for high-performance computing on cloud is highly impractical for an academic group or industry to invest in cloud for all of their computing needs

Can be useful to run a "few" "large" jobs, beyond the scale available in-house

https://aws.amazon.com/vpc/pricing/

Pros and cons of cloud computing

Pros of cloud computing: economies of scale

Equipment and resources:

- No requirement of dedicated infrastructure and management personnel
- Can be scaled up or down as needed – reduce redundancy
- Migration flexibility – especially with changing to updating hardware and software needs

Data generation and management:

- Ease of data access – anywhere on the internet (upload/download)
- Facilitate collaboration – internally in an organization or externally (can dictate which data is to be shared)
- Data resilience – storage in multiple physical locations, reduced risk of data loss

Budget:

- No capital or personnel costs – only operating costs
- Budget flexibility – can be scaled up or down
- Platform non-permanence – can always change cloud platform for better deals

Cons of cloud computing: lack of control

Equipment and resources:

- Requires internet 24x7 – else no access to data or resources
- Vendor "lock in" – limited flexibility/control on how resources can be used
- Decrease in performance with increase in size (both computations and data syncing/backup) – internet speeds can be slower than in-house cabling speeds

Data generation and management:

- Data migration difficulties – especially when changing vendors or moving data to local servers
- Non-availability of data – data is never stored locally, so power/internet/vendor outages to be managed
- Security – How does vendor handle data? Who owns data? Vendor can access data fully!

Budget and Security:

- Operating costs scale significantly with scale of computing/storage
- Security – Data breaches, API hacks, Authentication Issues, Law and order compliance issues
- Privacy, confidentiality and encryption – can your data be truly encrypted on cloud?

So, should I use cloud computing?

- It depends…

- Cloud services are already being used one-way or another

- Cloud computing is likely useful for
  - Product (software) development
  - Large-scale testing of product before deployment
  - Big data analytics (use built-in machine learning tools)
  - Collection of data from diverse resources (e.g., sensors)
  - Data archiving (long-term storage)

- Cloud for high-performance computing is not worth the cost at this stage
  - Definitely in academic settings, not feasible!
  - Could improve with time/economies-of-scale

Summary and conclusions

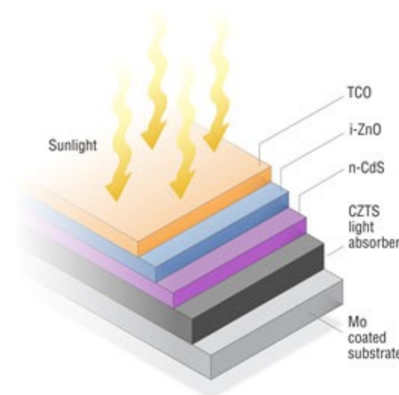

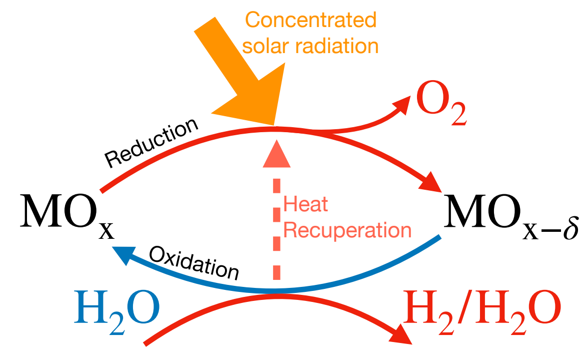

Density functional theory calculations, augmented by machine learning: quite useful for energy applications

Cloud computing: partially relevant for academia and industry for once-in-a-while usage, given cost constraints

Thanks for your attention! Questions?

saigautamg@iisc.ac.in
https://sai-mat-group.github.io