FEDERATED LEARNING OVER 5G, WIFI, AND ETHERNET: MEASUREMENTS AND EVALUATION

R.J. Hayek[12], J. Chung[2], K. Comer[1], C. R. Murthy[3], R. Kettimuthu[2], I. Kadota[1]
Northwestern University[1], Argonne National Laboratory[2], Indian Institute of Science[3]

MOTIVATION

- Many advancements made recently in distributed machine learning (ML)[1, 2]
- Feasibility via wireless communications is unknown
- There is a lack of real-world implementations (only simulations and emulations)

OUR CONTRIBUTIONS [3]

- Deploying FedAvg over real networks with edge devices
- Implementing communication agnostic metrics tooling
- Measuring communication and computation metrics over the testbed over 5G, WiFi, and Ethernet
- Combining and releasing all collected data and developed software[4]

FEDERATED LEARNING

- Method of decentralized learning ensuring data privacy
- Each node has its own dataset, which it uses to train its received *local* model
- The server aggregates the models across the network to create a *global* model

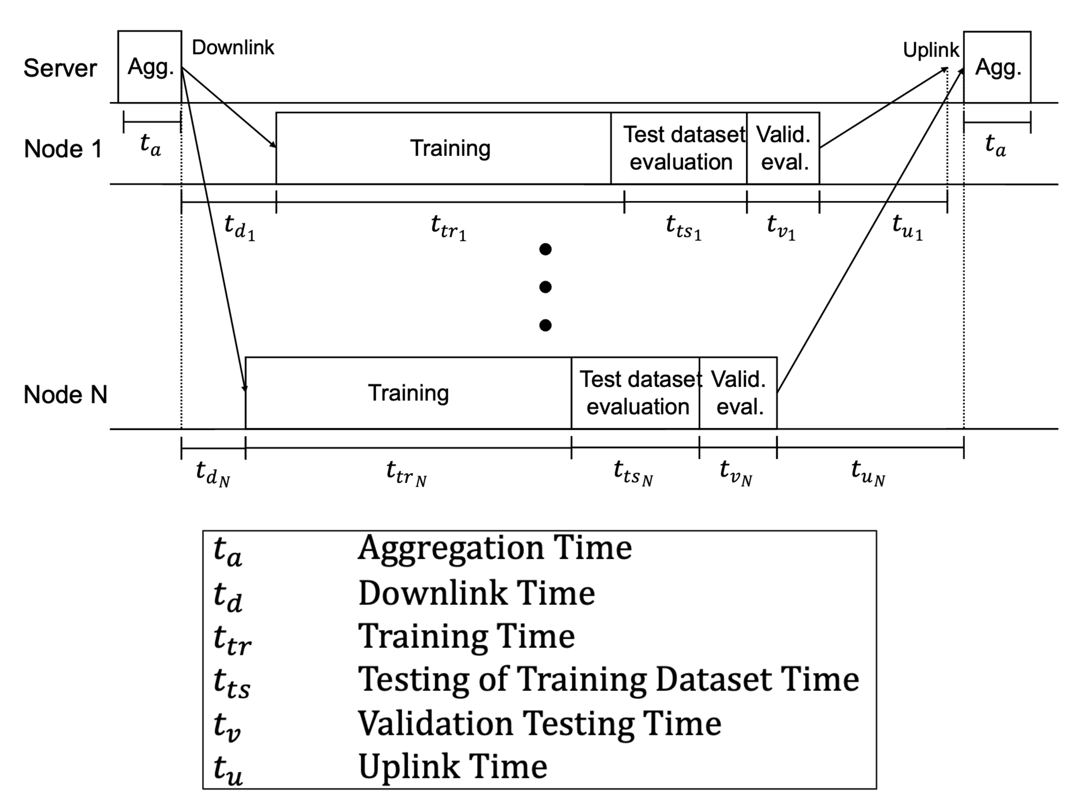

Fig. 1: One communication round

5G-NR

- Goals: high communication capacity, low latency, high reliability, and massive connectivity
- Network consists of end devices, a RAN, and the core network
- Designed as virtualized network components, allowing for software defined networks (SDN)
- Open-source solutions can be used to build a low-cost testbed using COTS devices
  - OpenAirInterface (CN + RAN)
  - Aether Onramp (Core)
  - SRS RAN (RAN)

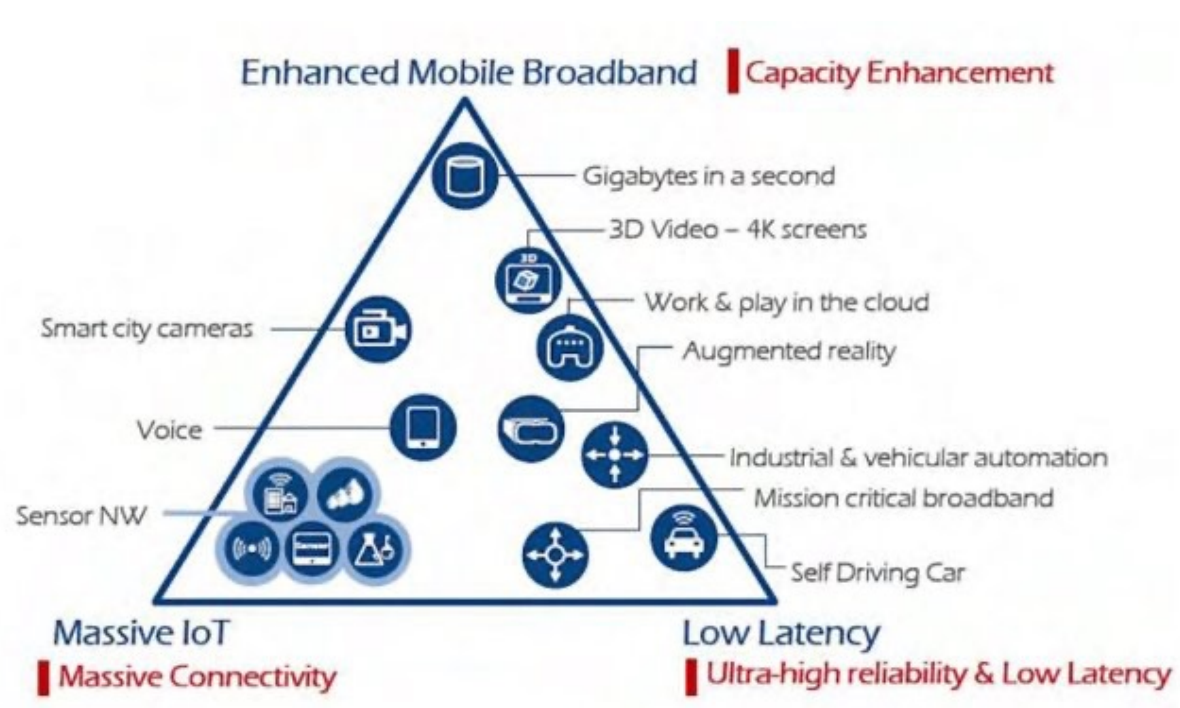

Fig. 2: Benefits of 5G-NR

IMPLEMENTATION

- Flower federated learning framework paired with the SqueezeNet[5] CNN
- Each node has access to Ethernet, WiFi, and 5G.
- 5G testbed is built using OAI CN and RAN, split between two PC's
- RAN utilizes USRP x310 as the gNB.
- UE's are Raspberry Pi 6's paired with Telit 980m modems

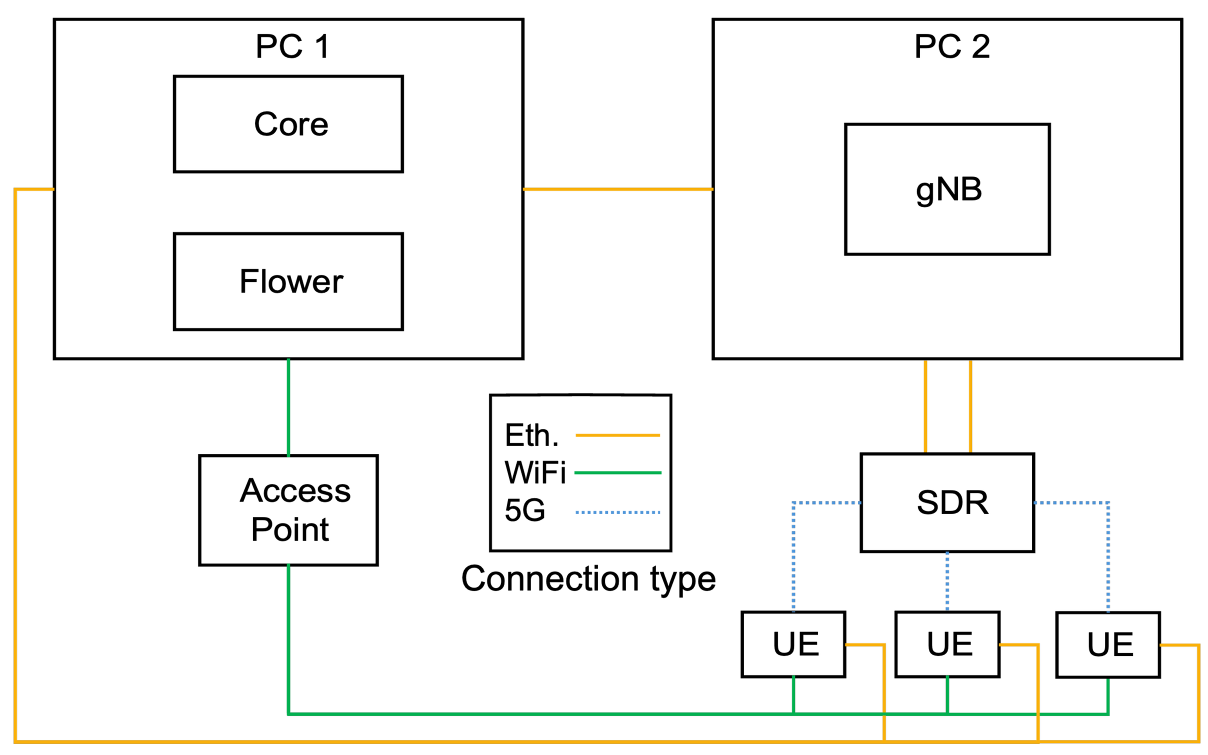

Fig. 3: Infrastructure Design

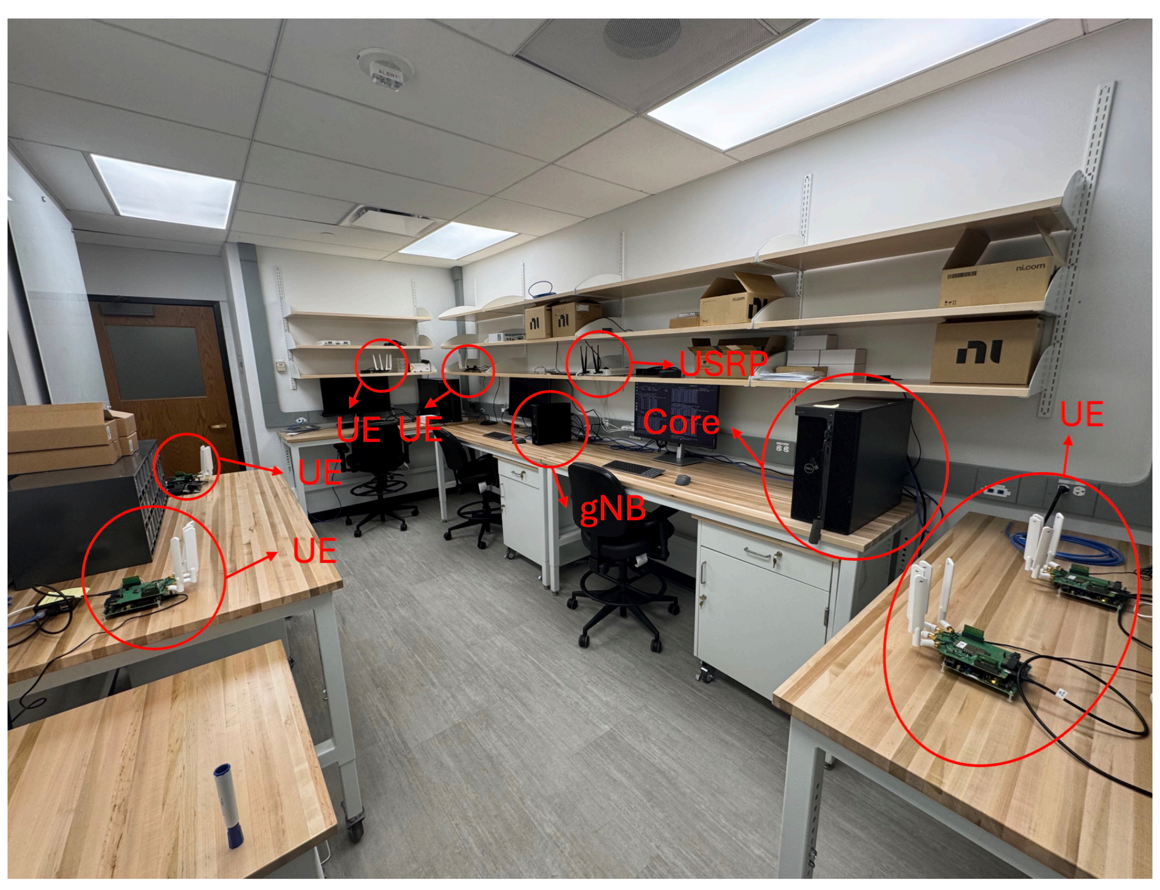

Fig. 4: Physical Testbed Infrastructure

CONVERGENCE TIME

- Defined as the overall duration of the trial required to trigger an early stopping signal.
- We **observe** Ethernet has the lowest average round time (31.46 sec.), with 108 rounds to converge.
- Compared to 5G, the rounds take 43.28 sec., 116 rounds to convergence.
- 5G shows an increase of 46% compared to Ethernet.
- Performance differences can be attributed to communication performance.

Table 1: Comparison of Ethernet, WiFi, and 5G number of rounds, round time, and convergence time

Trial	Number of Rounds			Communication Round Time			Convergence Time		
	Ethernet	WiFi	5G	Ethernet	WiFi	5G	Ethernet	WiFi	5G
1	117	127	121	31.64	32.04	44.64	3701.88	4069.08	5401.44
2	101	112	100	31.47	32.18	43.92	3178.47	3604.16	4392.00
3	121	99	106	31.41	32.18	41.97	3800.61	3185.82	4448.82
4	120	93	146	31.25	32.20	43.24	3750.00	2994.60	6313.04
5	103	89	87	31.44	32.16	43.32	3238.32	2862.24	3768.84
6	88	139	89	31.44	32.19	41.72	2766.72	4474.41	3713.08
7	79	131	114	31.43	32.28	43.80	2482.97	4228.68	4993.20
8	128	123	110	31.39	32.20	44.68	4017.92	3960.60	4914.80
9	132	122	142	31.56	32.34	43.23	4165.92	3945.48	6138.66
10	100	—	146	31.51	—	42.29	3151.00	—	6174.34
Avg.	108.9	115.0	116.1	31.46	32.20	43.28	3394.66	3657.00	4984.09

CONVERGENCE TIME CONT.

- There is no observed difference in the converged validation accuracy.
- The 5G network communication performance greatly extends the average round time.
- Attributed mainly to the communication overhead.

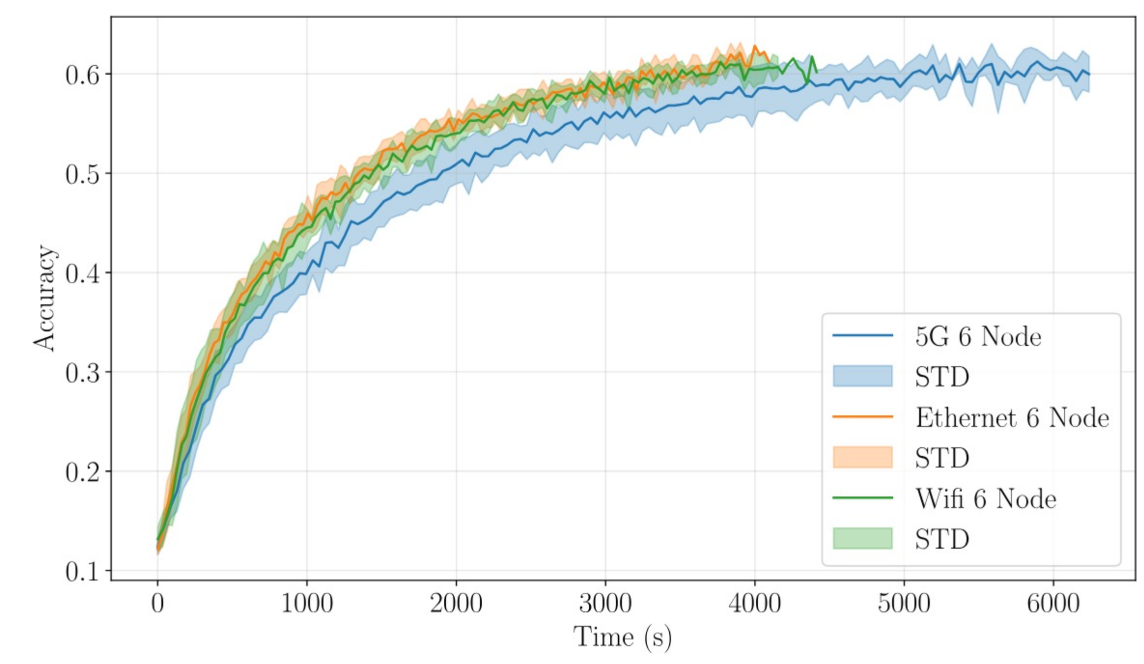

Fig. 5: Total communication round time over 10 trials comparing Ethernet, WiFi, and 5G

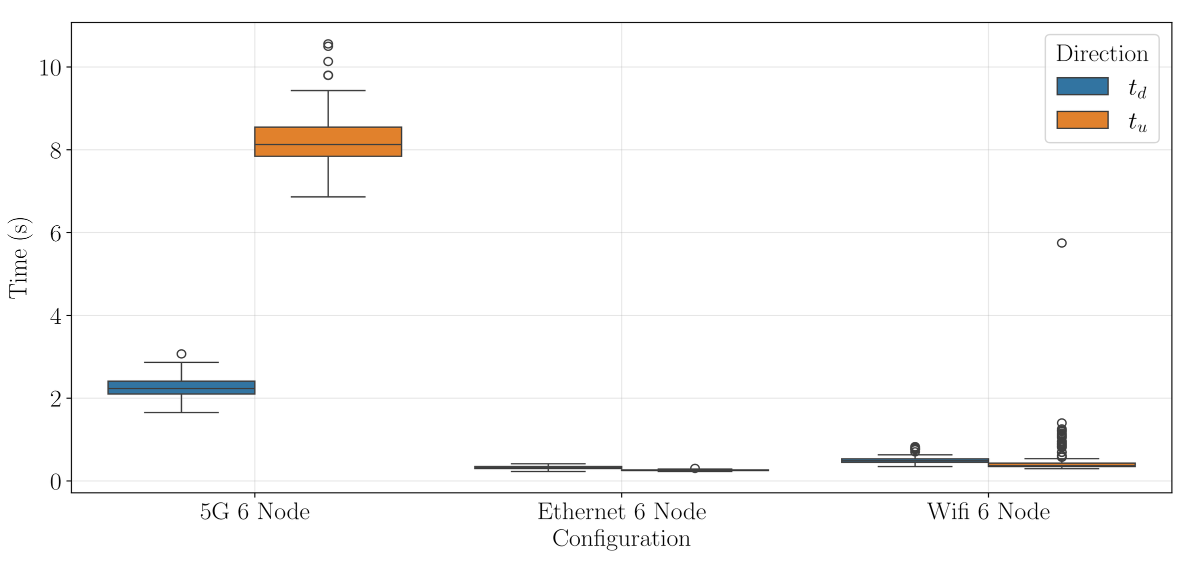

Fig. 6: Mean uplink and downlink times averaged for all nodes on each network interface

NETWORK SCALING EFFECTS

- As number of nodes increased from 3 to 6, so did uplink and downlink time (higher network use)
- However, total impact on convergence time was small – (more nodes take longer to converge)

Table 2: Communication round metrics averaged across all nodes

N	UL Time	DL Time	Round Time	Round Number	%UL Time	%DL Time
3	3.4130	2.3198	57.3190	105	5.9544	4.0472
4	4.6206	1.8415	48.2631	119	9.5738	3.8156
5	5.8487	1.7718	42.7461	128	13.6823	4.1450
6	10.3477	2.3089	43.3056	145	23.8946	5.3317

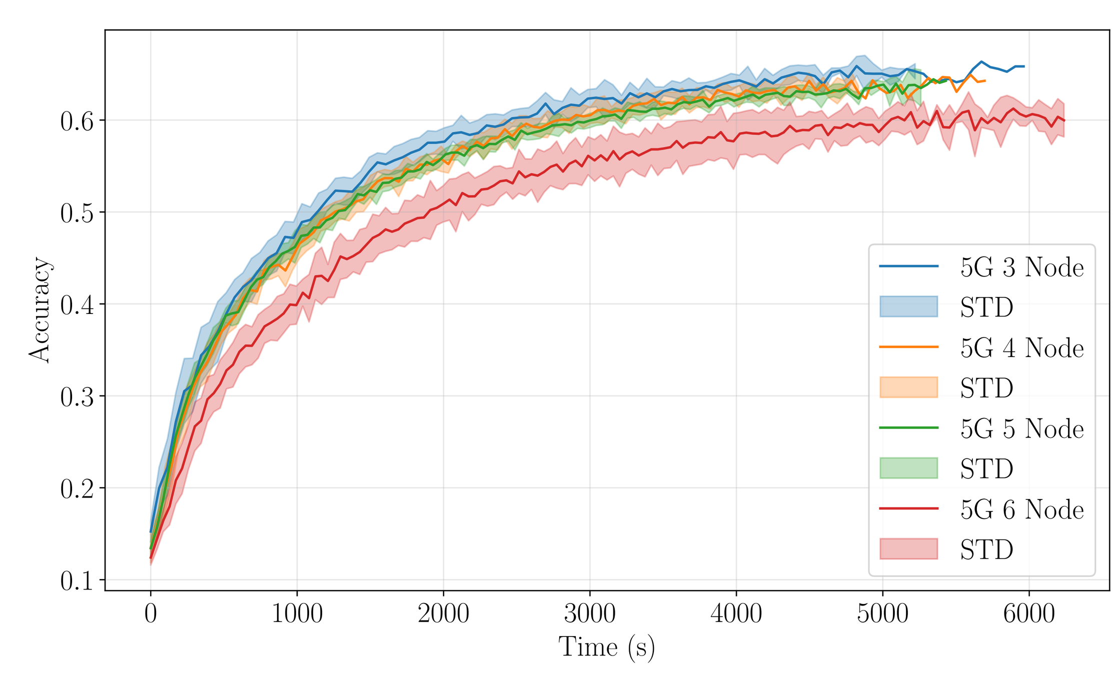

Fig. 7: Worst local validation accuracy as measured by each node on the 5G network

REFERENCES

[1] Syreen Banabilah, Moayad Aloqaily, Eitaa Alsayed, Nida Malik, and Yaser Jararweh. 2022. Federated learning review: Fundamentals, enabling technologies, and future applications. Information Processing & Management 59, 6 (Nov. 2022), 103061. https://doi.org/10.1016/j.ipm.2022.103061

[2] Wei Yang Bryan Lim, Nguyen Cong Luong, Dinh Thai Hoang, Yutao Jiao, Ying-Chang Liang, Qiang Yang, Dusit Niyato, and Chunyan Miao. 2020. Federated Learning in Mobile Edge Networks: A Comprehensive Survey. IEEE Communications Surveys & Tutorials 22, 3 (2020), 2031–2063. https://doi.org/10.1109/COMST.2020.2986024 Conference Name: IEEE Communications Surveys & Tutorials.

[3] Robert J. Hayek, Joaquin Chung, Kayla Comer, Chandra R. Murthy, Rajkumar Kettimuthu and Igor Kadota. 2025. Federated Learning over 5G, WiFi, and Ethernet: Measurements and Evaluation. arXiv:2504.04678.

[4] Robert J. Hayek, Joaquin Chung, Kayla Comer, Chandra Murthy, Rajkumar Kettimuthu, and Igor Kadota. 2025. Federated Learning for 5G Tesbed. https://github.com/Net-X-Research-Group/federated_learning_testbed.

[5] Forrest N. Iandola, Song Han, Matthew W. Moskewicz, Khalid Ashraf, William J. Dally, and Kurt Keutzer. 2016. SqueezeNet: AlexNet-level accuracy with 50x fewer parameters and <0.5MB model size. arXiv:1602.07360 (2016).

U.S. DEPARTMENT OF ENERGY
Argonne National Laboratory is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC.

Northwestern University

Argonne NATIONAL LABORATORY