



Graph Neural Networks

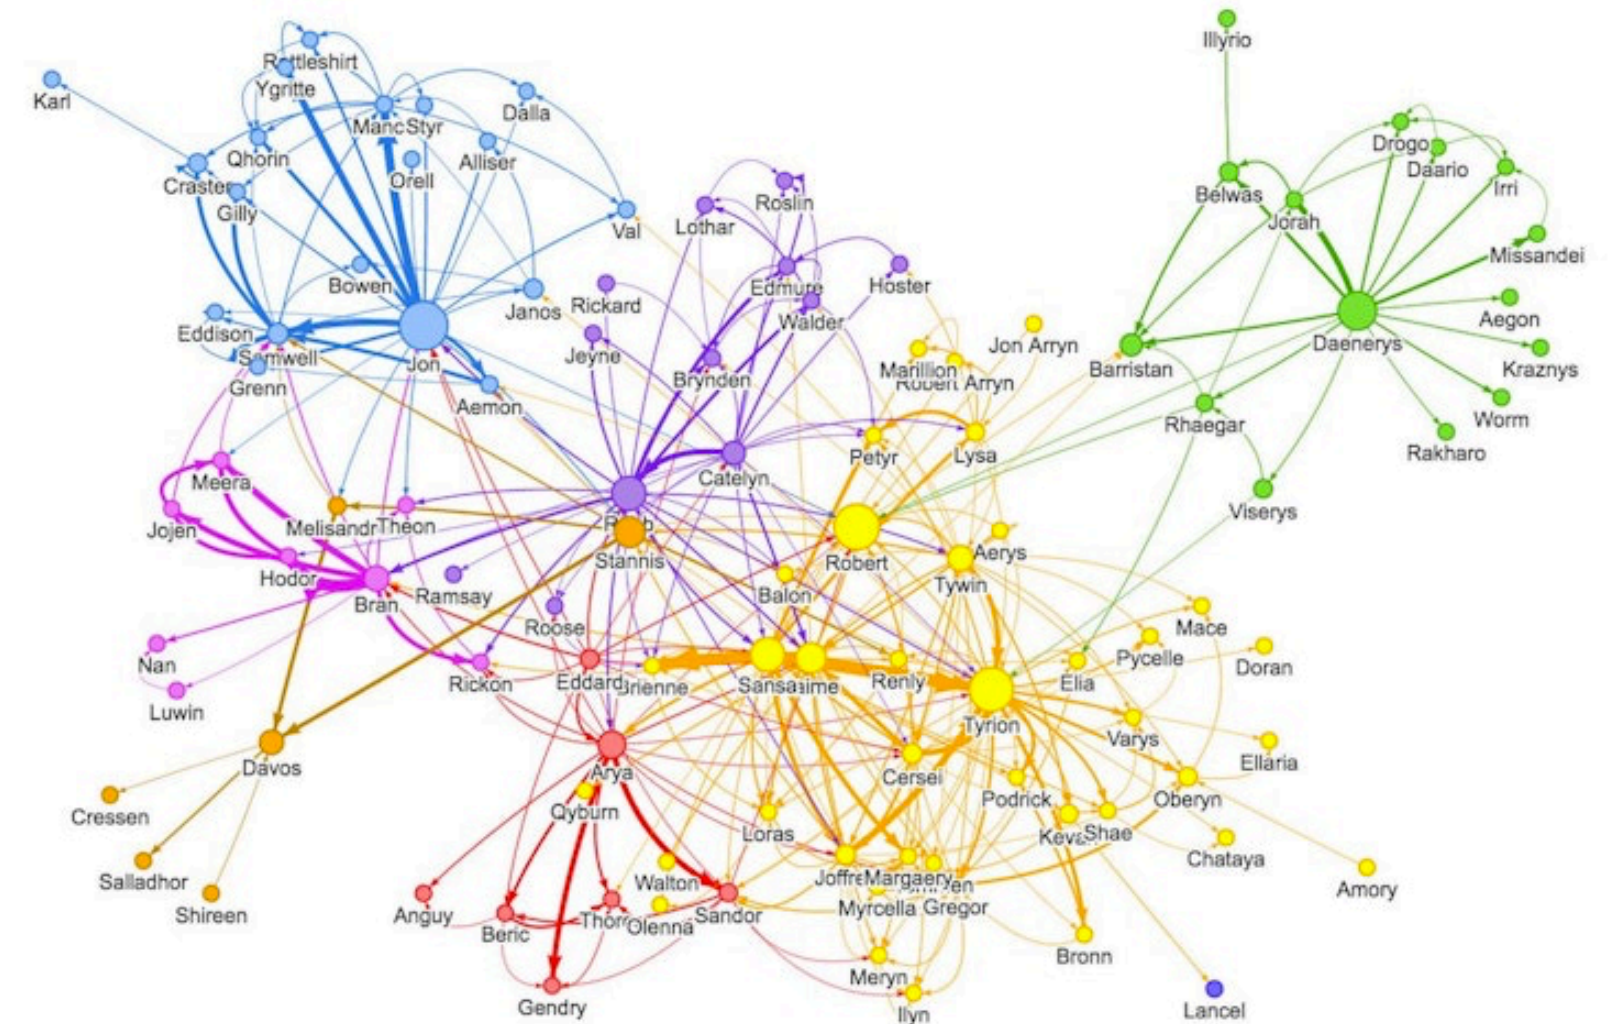



Abhinav Bhatele, Department of Computer Science



UNIVERSITY OF
MARYLAND

Real-world graphs

- Graphs are everywhere: financial transaction networks, protein-protein interactions, social networks
- Learn complex properties of and relationships within graphs
- Use cases: fraud detection, bioinformatics, recommendation systems



<https://doi.org/10.1057/s41599-022-01075-x01>

<https://www.cylynx.io/blog/a-comparison-of-javascript-graph-network-visualisation-libraries/>

Tasks on graphs

- Node-level: predicting the class of individual nodes
- Edge-level: determine if an edge exist or predict the type of edge
- Graph-level: even categorize entire graphs

Graph Neural Networks



Thomas Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. ICLR 2017

Graph Neural Networks

- Type of neural network to learn from graph datasets
- Graph Convolutional Networks (GCNs) have become widely popular in recent years for studying properties of graphs

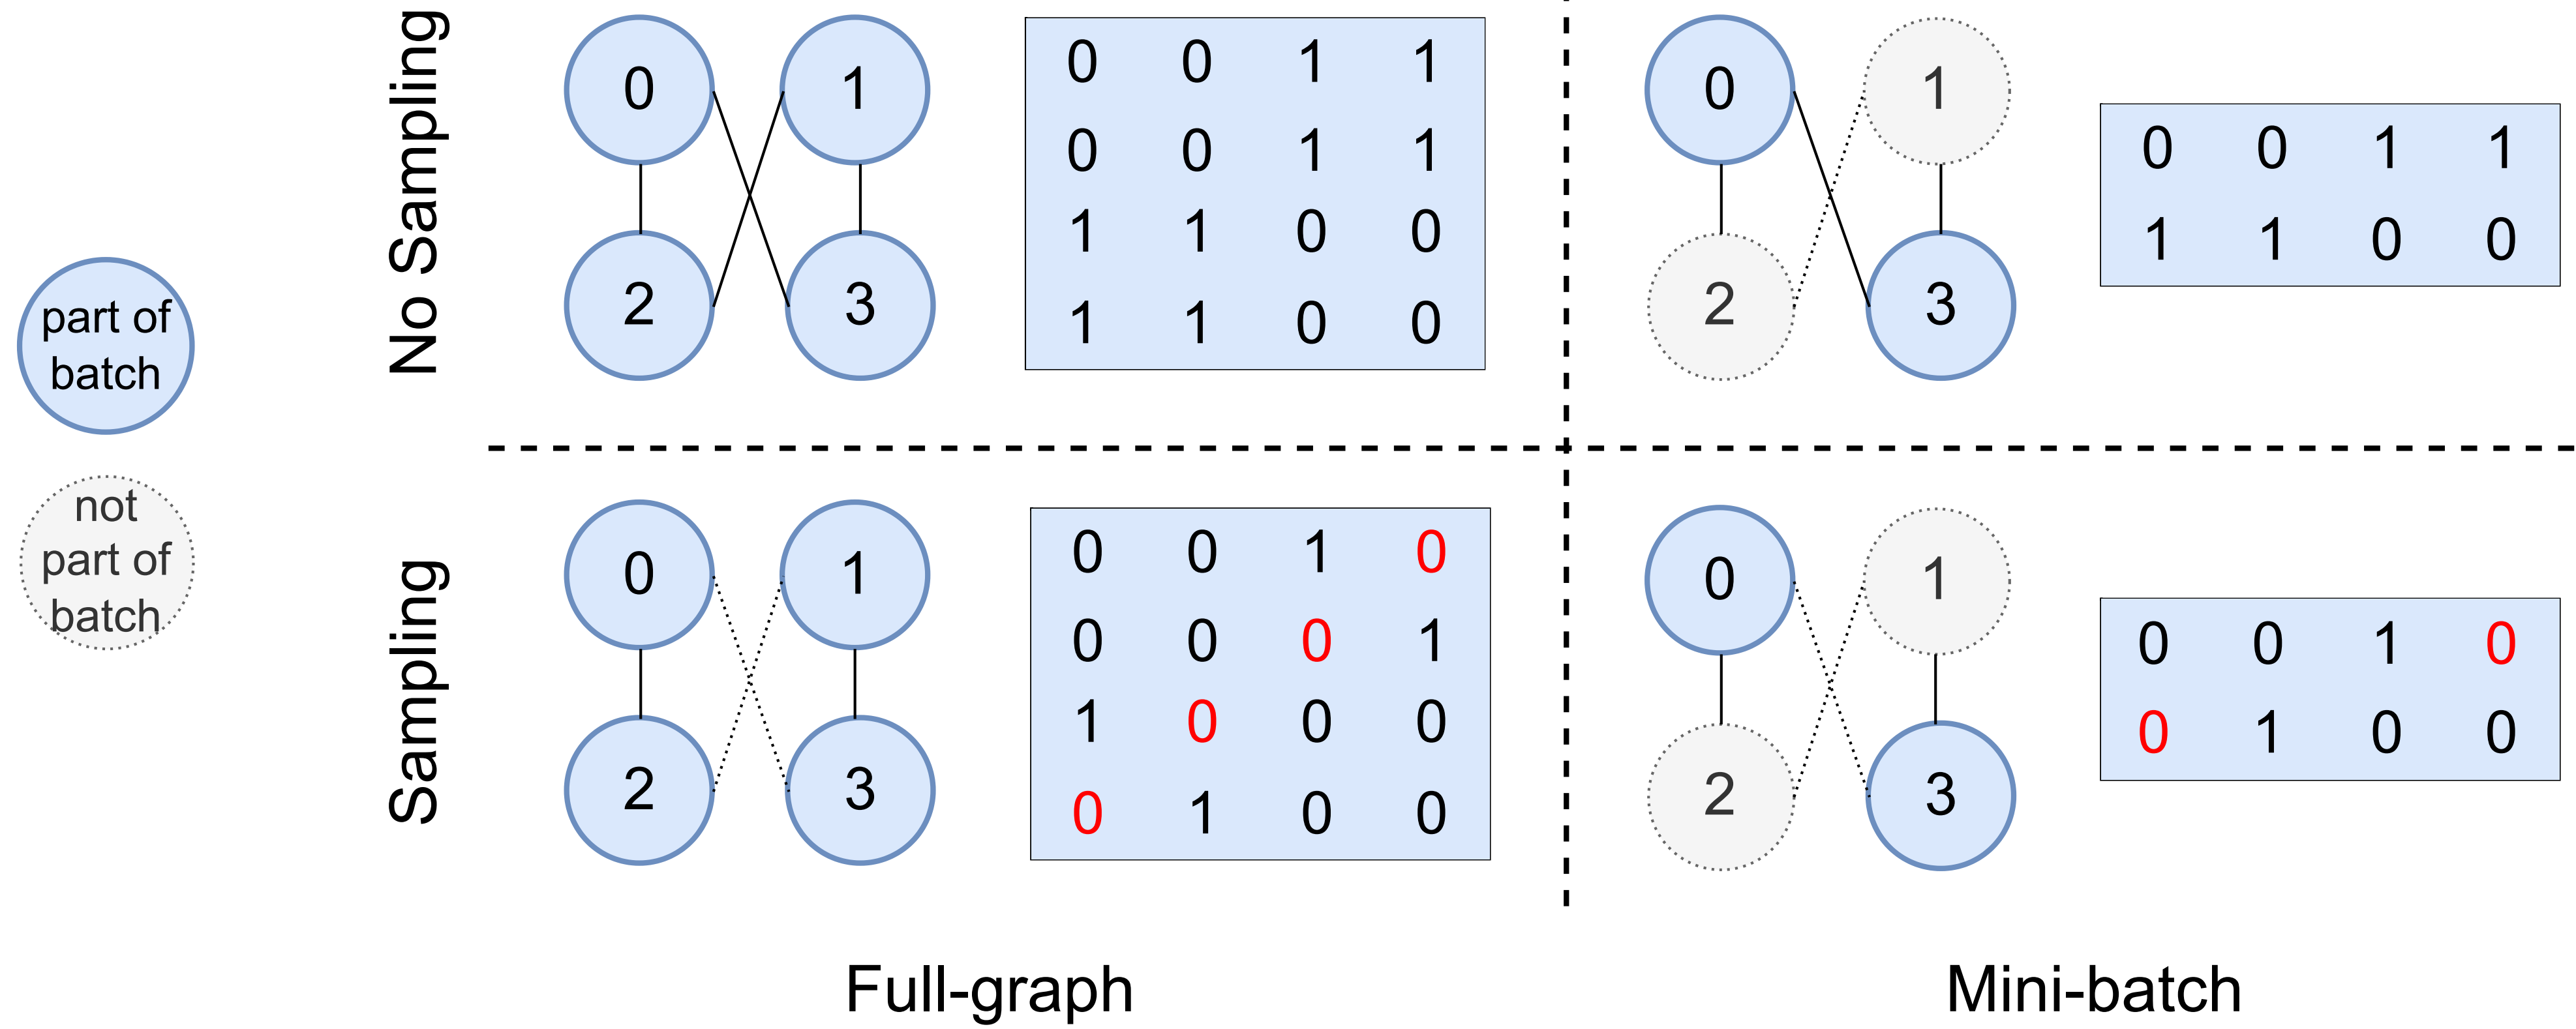
Thomas Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. ICLR 2017

Graph Neural Networks

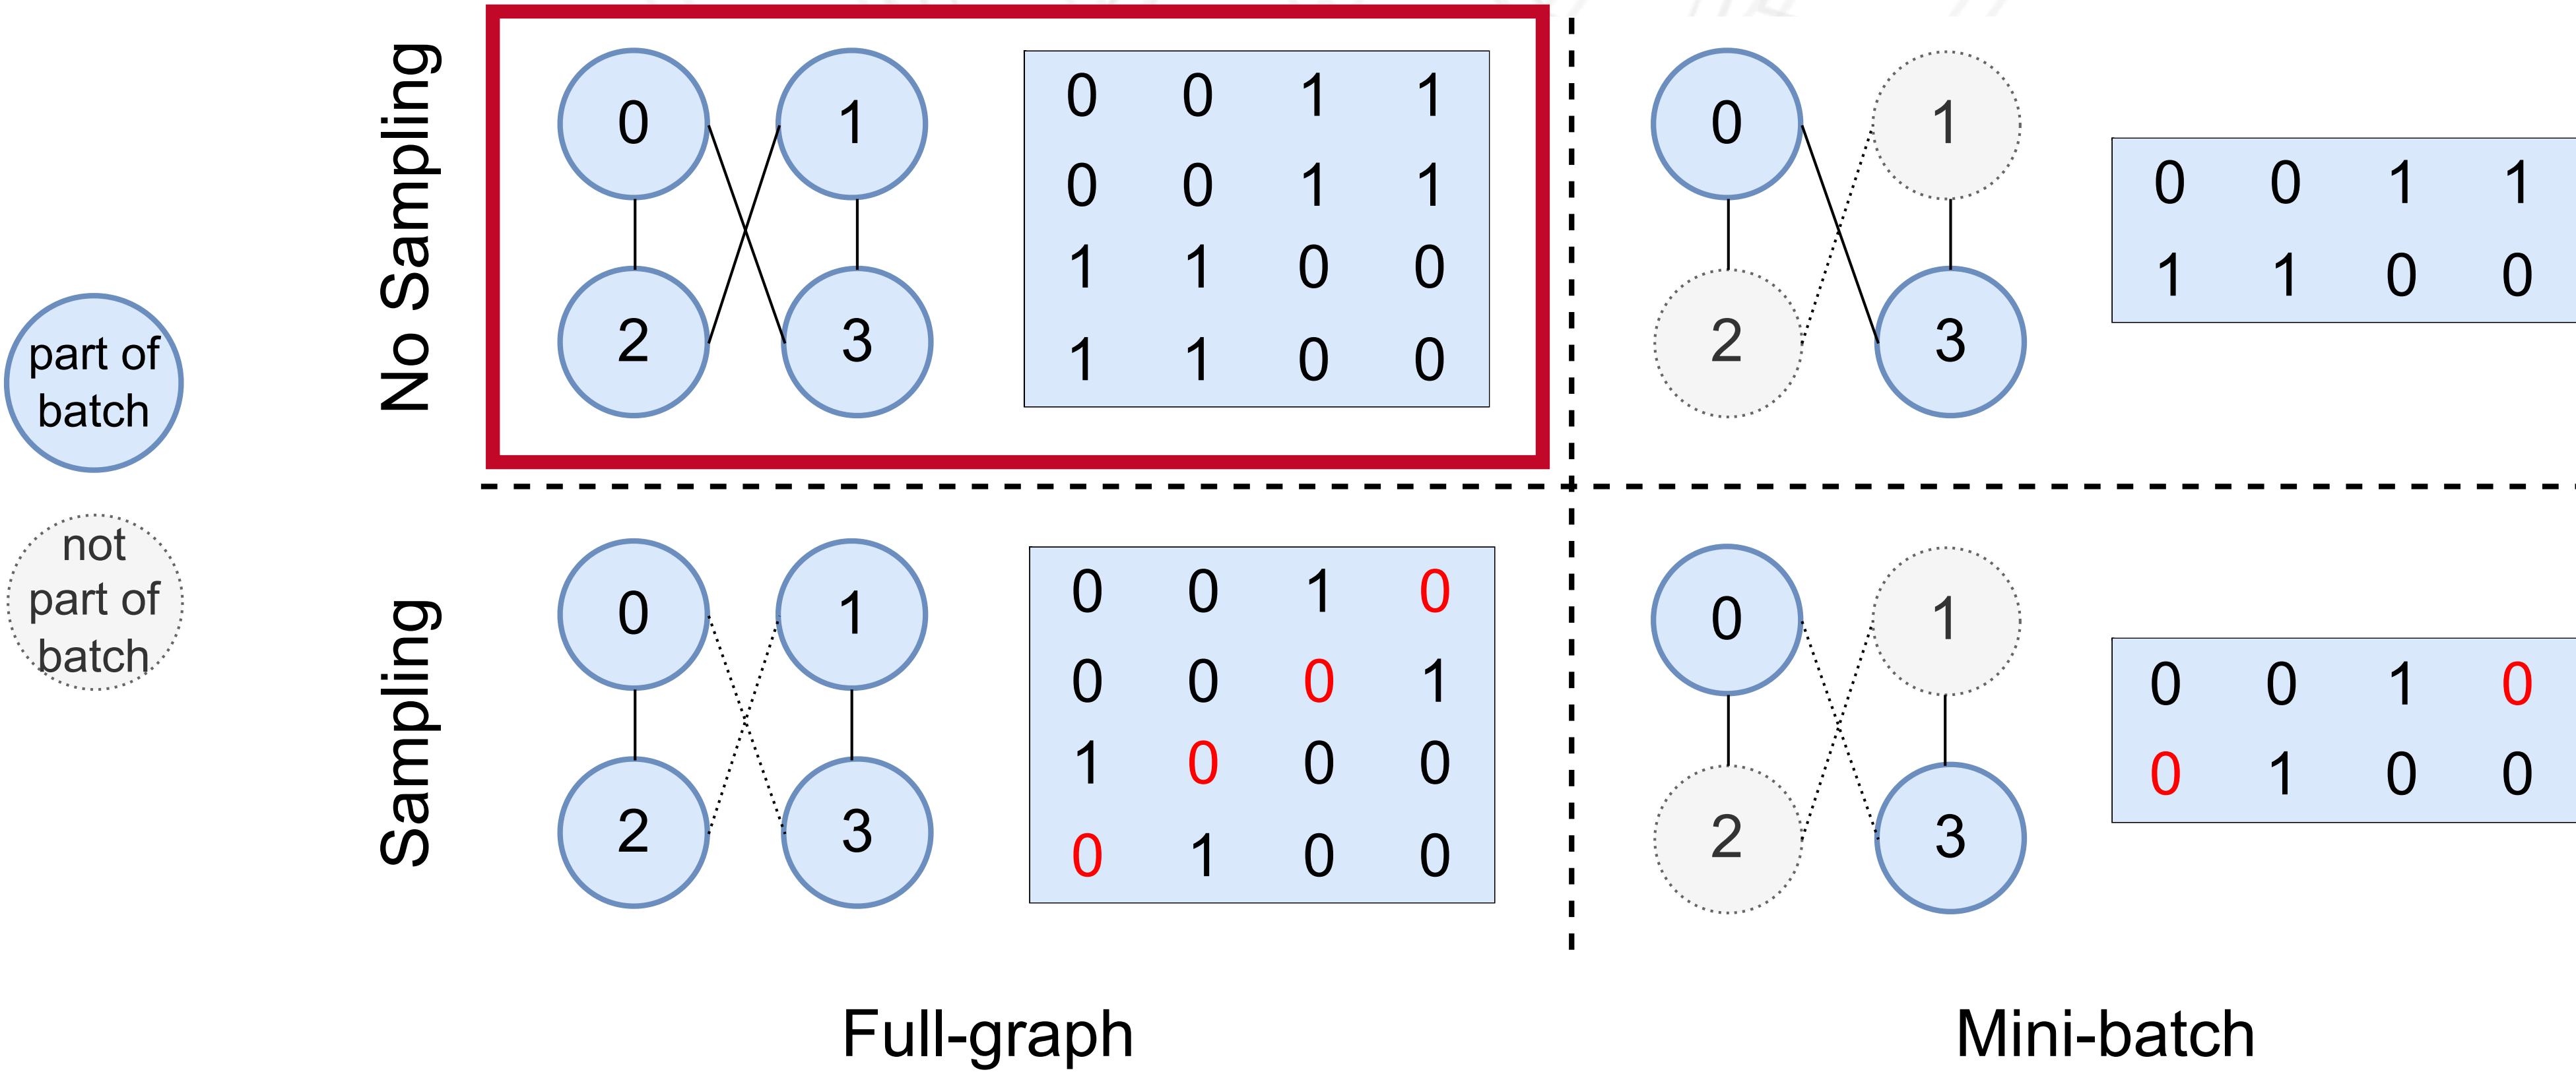
- Type of neural network to learn from graph datasets
- Graph Convolutional Networks (GCNs) have become widely popular in recent years for studying properties of graphs
- Systems challenges with GNNs:
 - Graphs are irregular as opposed to images or text
 - Input datasets representing extremely large graphs do not fit within memory

Thomas Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. ICLR 2017


Full-graph versus mini-batching



Full-graph versus mini-batching



Full-graph versus mini-batching



Challenges with Parallel GNN Training

- Very large graphs require effective parallelization over multiple GPUs
- Irregularity in graph structure leads to highly imbalanced and highly sparse adjacency matrices
 - Leads to load imbalance when distributing work
- Sparse kernels can be slow on GPUs
- Significant communication for synchronizing activations and gradients: $N \times D$



UNIVERSITY OF
MARYLAND

Abhinav Bhatele

5218 Brendan Iribe Center (IRB) / College Park, MD 20742

phone: 301.405.4507 / e-mail: bhatele@cs.umd.edu