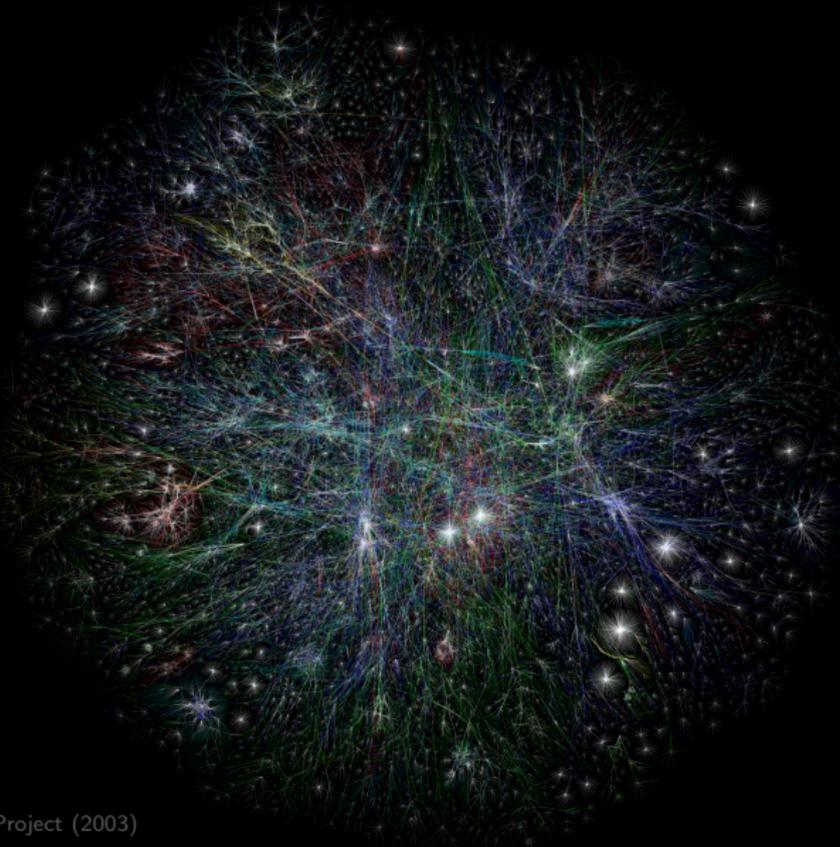


The Internet (in 2003)

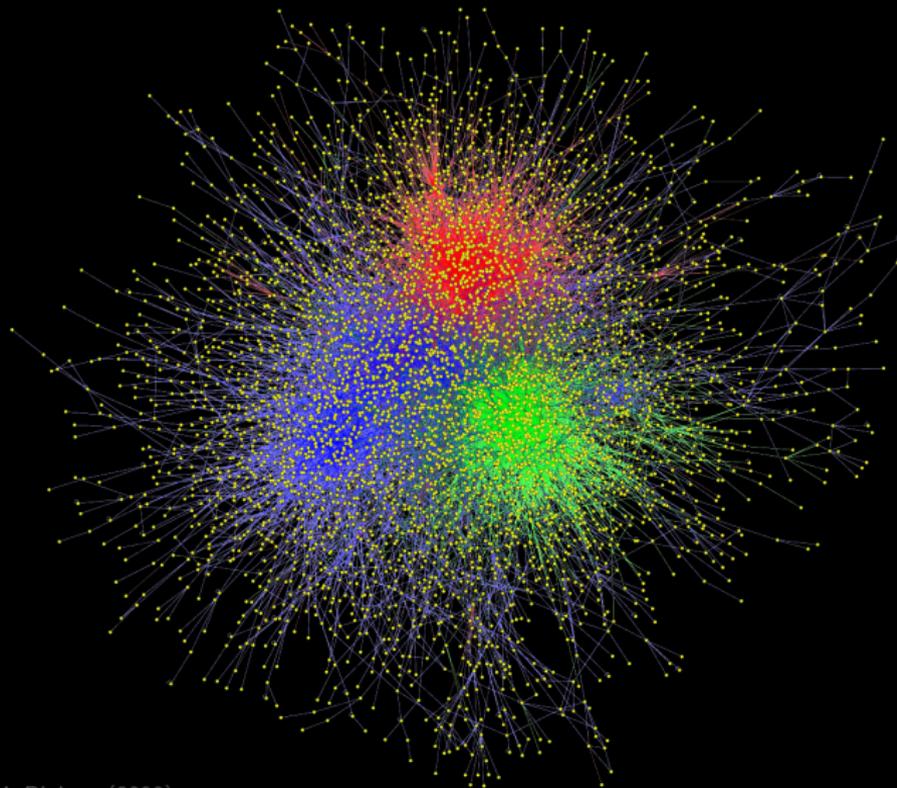


Source: The Opte Project (2003)

Facebook (in 2010)



Protein Interactions in Human Cells



Source: Ferrell, J. Biology (2009)

Graphs are increasingly massive:

- Web graphs: over 100 billion (10^{11}) edges.
- Social graphs: over 1 trillion (10^{12}) edges.
- Data sizes: 100s of GBs, if not TBs.

Graph processing systems must be **distributed**.

Must consider **time**, **memory**, and **network**.

Graphs are increasingly massive:

- Web graphs: over 100 billion (10^{11}) edges.
- Social graphs: over 1 trillion (10^{12}) edges.
- Data sizes: 100s of GBs, if not TBs.

Graph processing systems must be **distributed**.

Must consider **time**, **memory**, and **network**.

Pregel-like graph processing systems are increasingly popular.

Giraph:



facebook

GraphLab:



Pregel-like systems are **BSP**, **vertex-centric** programs.

Pregel-like systems are **BSP**, **vertex-centric** programs.

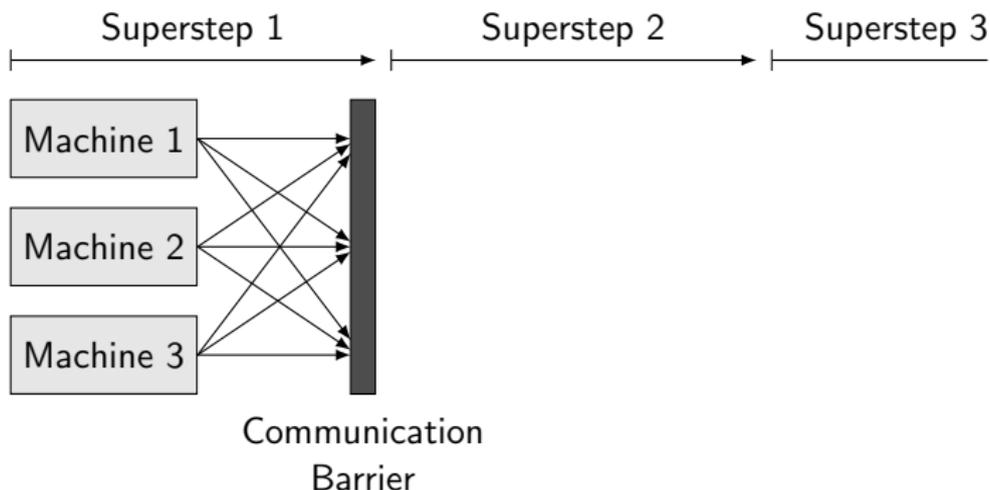
Computation



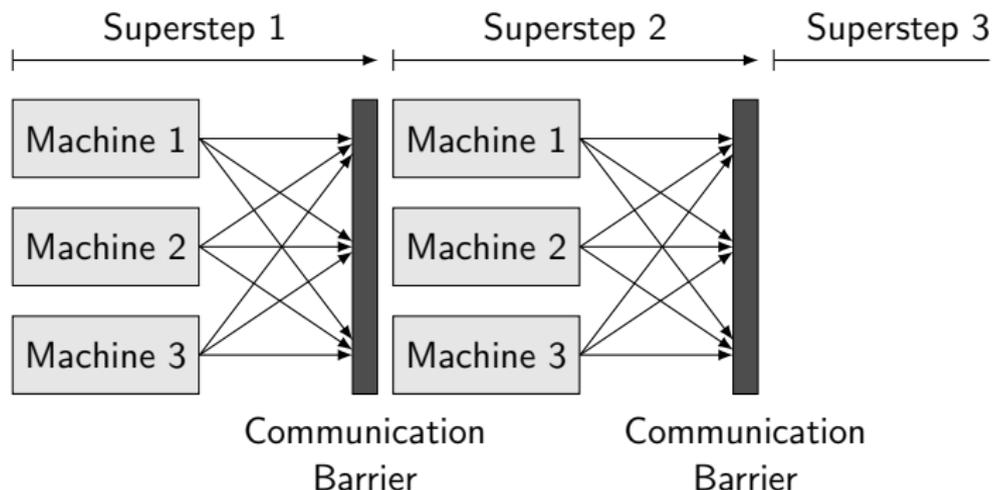
Pregel-like systems are **BSP**, **vertex-centric** programs.



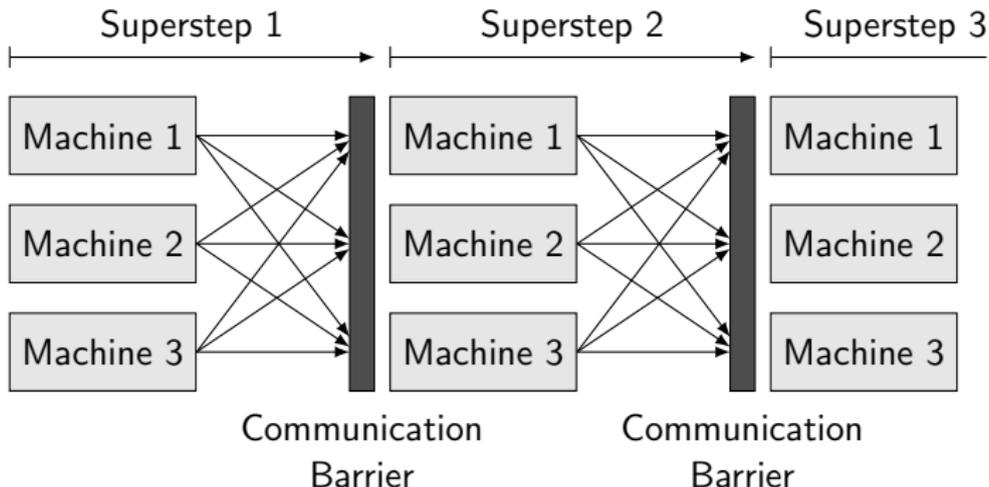
Pregel-like systems are **BSP**, **vertex-centric** programs.



Pregel-like systems are **BSP**, **vertex-centric** programs.



Pregel-like systems are **BSP**, **vertex-centric** programs.



Pregel-like systems are **BSP**, **vertex-centric** programs.

- “Think like a vertex”:

