Recall MPI:

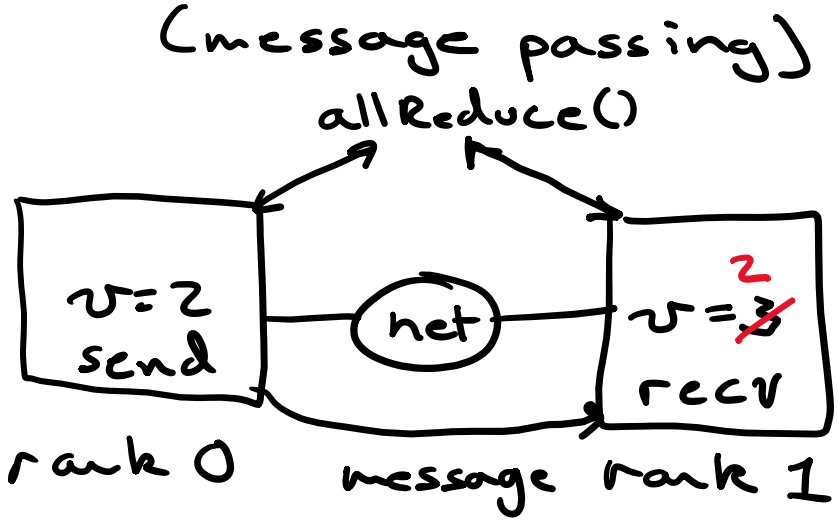

processor IO (procID)

nprocs = 2
↑
number of total ranks

Distributed Graph

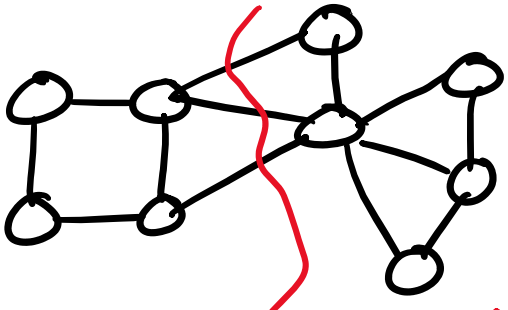

rank 0 rank 1

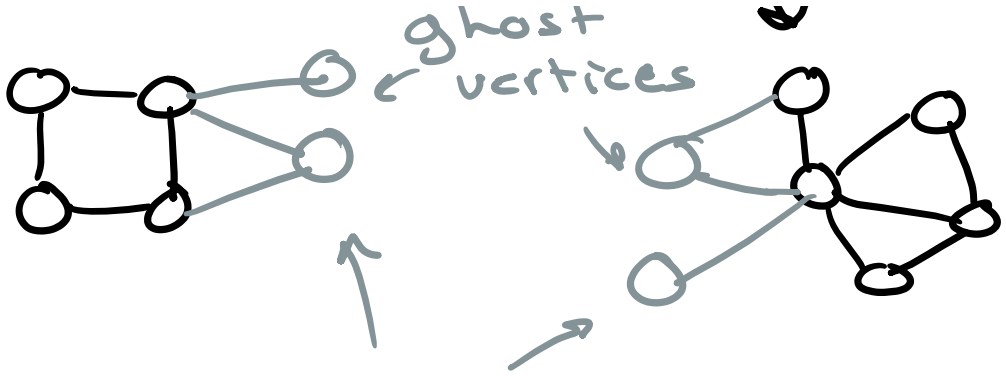

→ we compute updates on local vertices and exchange states of ghost vertices

---

BFS

```
BFS(G, root):
    For v ∈ V(G):        ← local graph owned+ghost vertices
        S[v] = undiscoverd
    Q ← root             ← only if we own root
    Q_n ← {}, level = 1, S[root] = 0
                         ← global queue not
```

level = -1, S[root] = 0

While $Q \neq \{\}$: ← global queue not empty

For all $v \in Q$:

For all $u \in N(v)$:

if S[u] = undiscovered

Q_c ← comm. queue

$Q_n \leftarrow u$ ← we need to communicate

S[u] = level

++level

$Q = Q_n$

$Q_n = \{\}$

Q = recv verts from all other ranks (using Q_c)

$Q_c = \{\}$

our communication:

For i in range(nprocs):

For j in range(nprocs):

if i == procid: ← current rank

sendbuf = [v in Q_c owned by rank j]

owned by rank j

```
If i == j:
    Q ← sendbuf
elif i == procid:
    MPI.Send(sendbuf)
elif j == procid:
    MPI.Recv(Q)
allReduce(len(Q), global_Q_length, MPI.sum)
```