Covering Problems like Set-Cover or Dominating Set have widespread applications in many domains of distributed computing.
In their simplest form they can be described as the following graph problem:
Consider a (possibly directed!) graph of n elements, then the goal is to select a minimal collection of elements that cover, i.e., are adjacent to, all other elements in the system.
A ususal assumption in the distributed setting is that all elements can bidirectionally communicate with all elements they potentially cover and vice versa.
In other words, the communication graph is equal to (the bidirected version of) the problem graph
The resulting algorithms are fast but usually require a lot of messages (especially in dense graphs) as they continously send messages over all available edges.
This makes it hard to apply them scenarios where communication is limited, the communication graph greatly differs from the problem graph, and/or the problem graph is changing.
In this talk we present our current state on solving Set-Cover in these settings.
In particular, the talk has three parts:

1. First, we present some general sampling techniques to estimate the number of uncovered neighbors with few messages.
2. Second, we observe a node-capacitated network where each element is fully aware of its neighbors, but can communicate with at most O(log(n)) other elements in a single round.
Here, we show that Set-Cover can be approximated in O(n^(3/4)) rounds irregardless of the problem graphs structure.
3. Last, we sketch how our ideas can be adapted to instances the problem graph is changing over time.