

Decentralized Slicing in Mobile Low-Power Wireless Networks

Piotr Jaszowski, Pawel Sienkowski, Konrad Iwanicki
University of Warsaw

pj306249@students.mimuw.edu.pl
ps319383@students.mimuw.edu.pl
iwanicki@mimuw.edu.pl



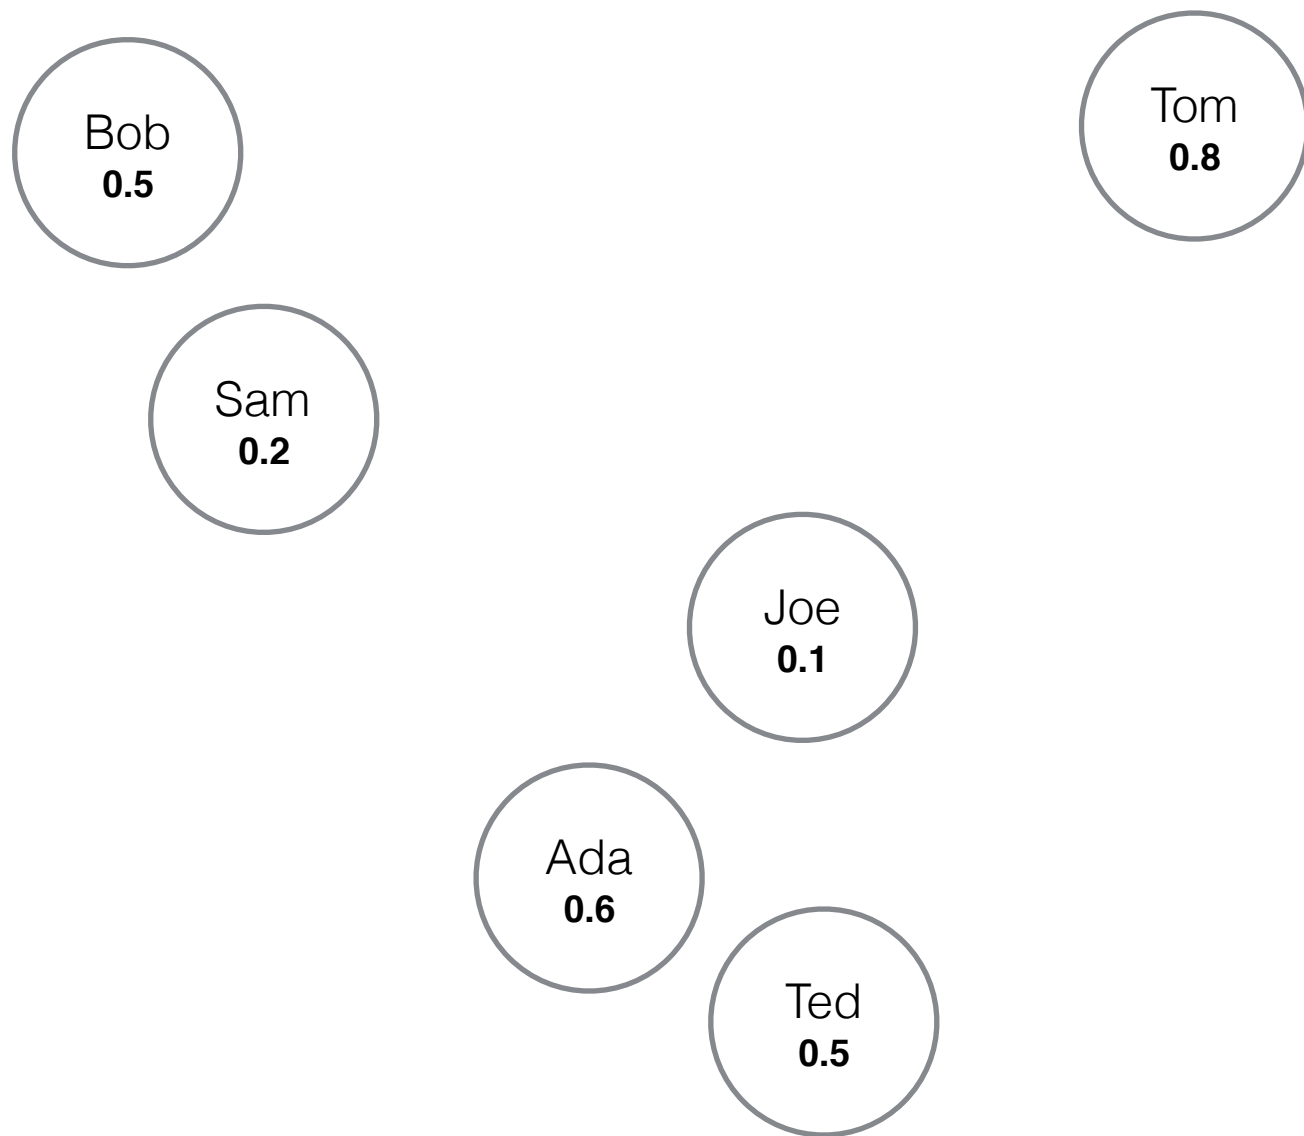
Decentralized Slicing Problem



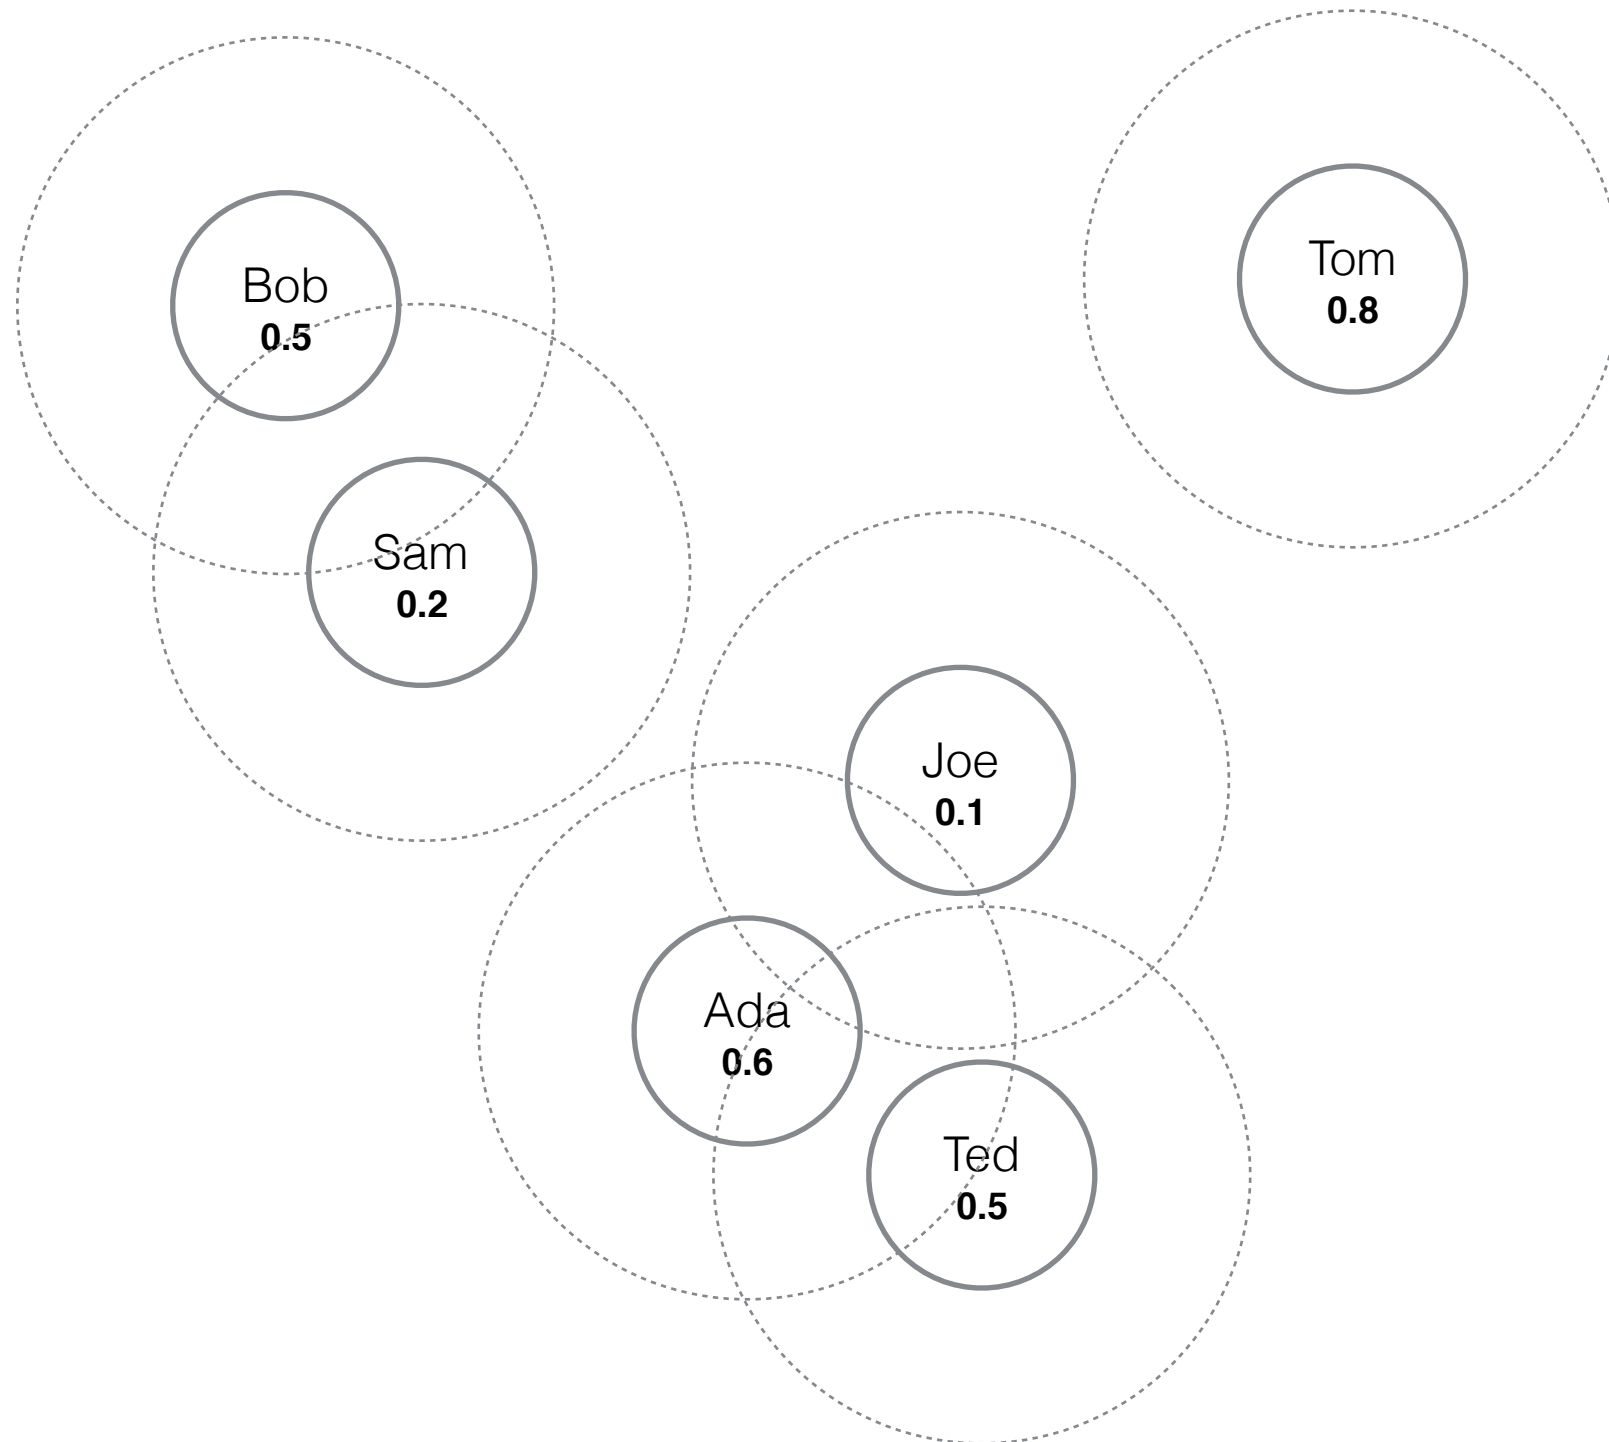
Decentralized Slicing Problem



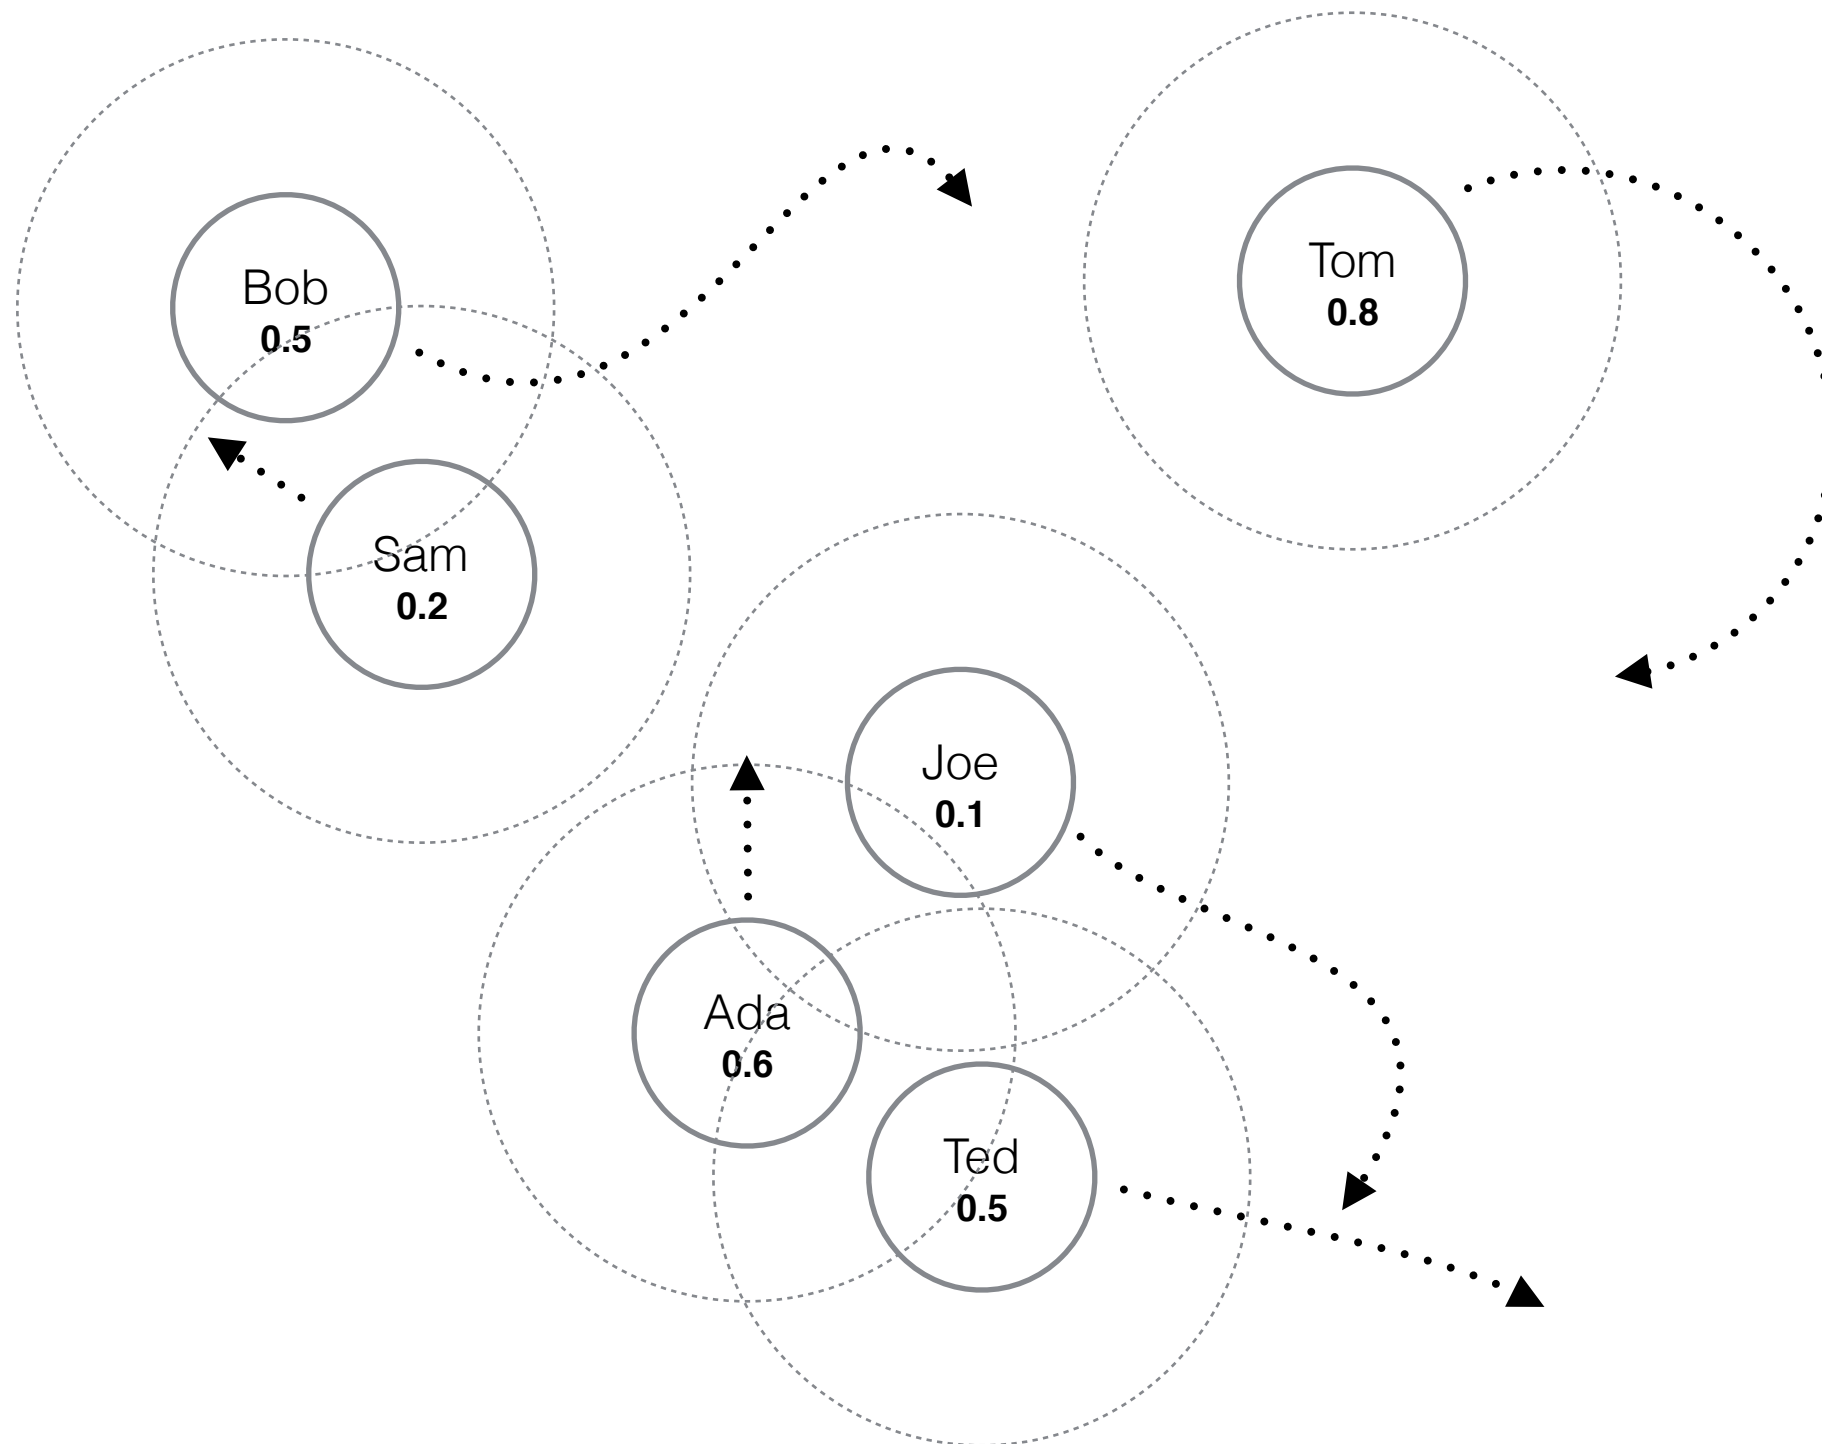
Decentralized Slicing Problem



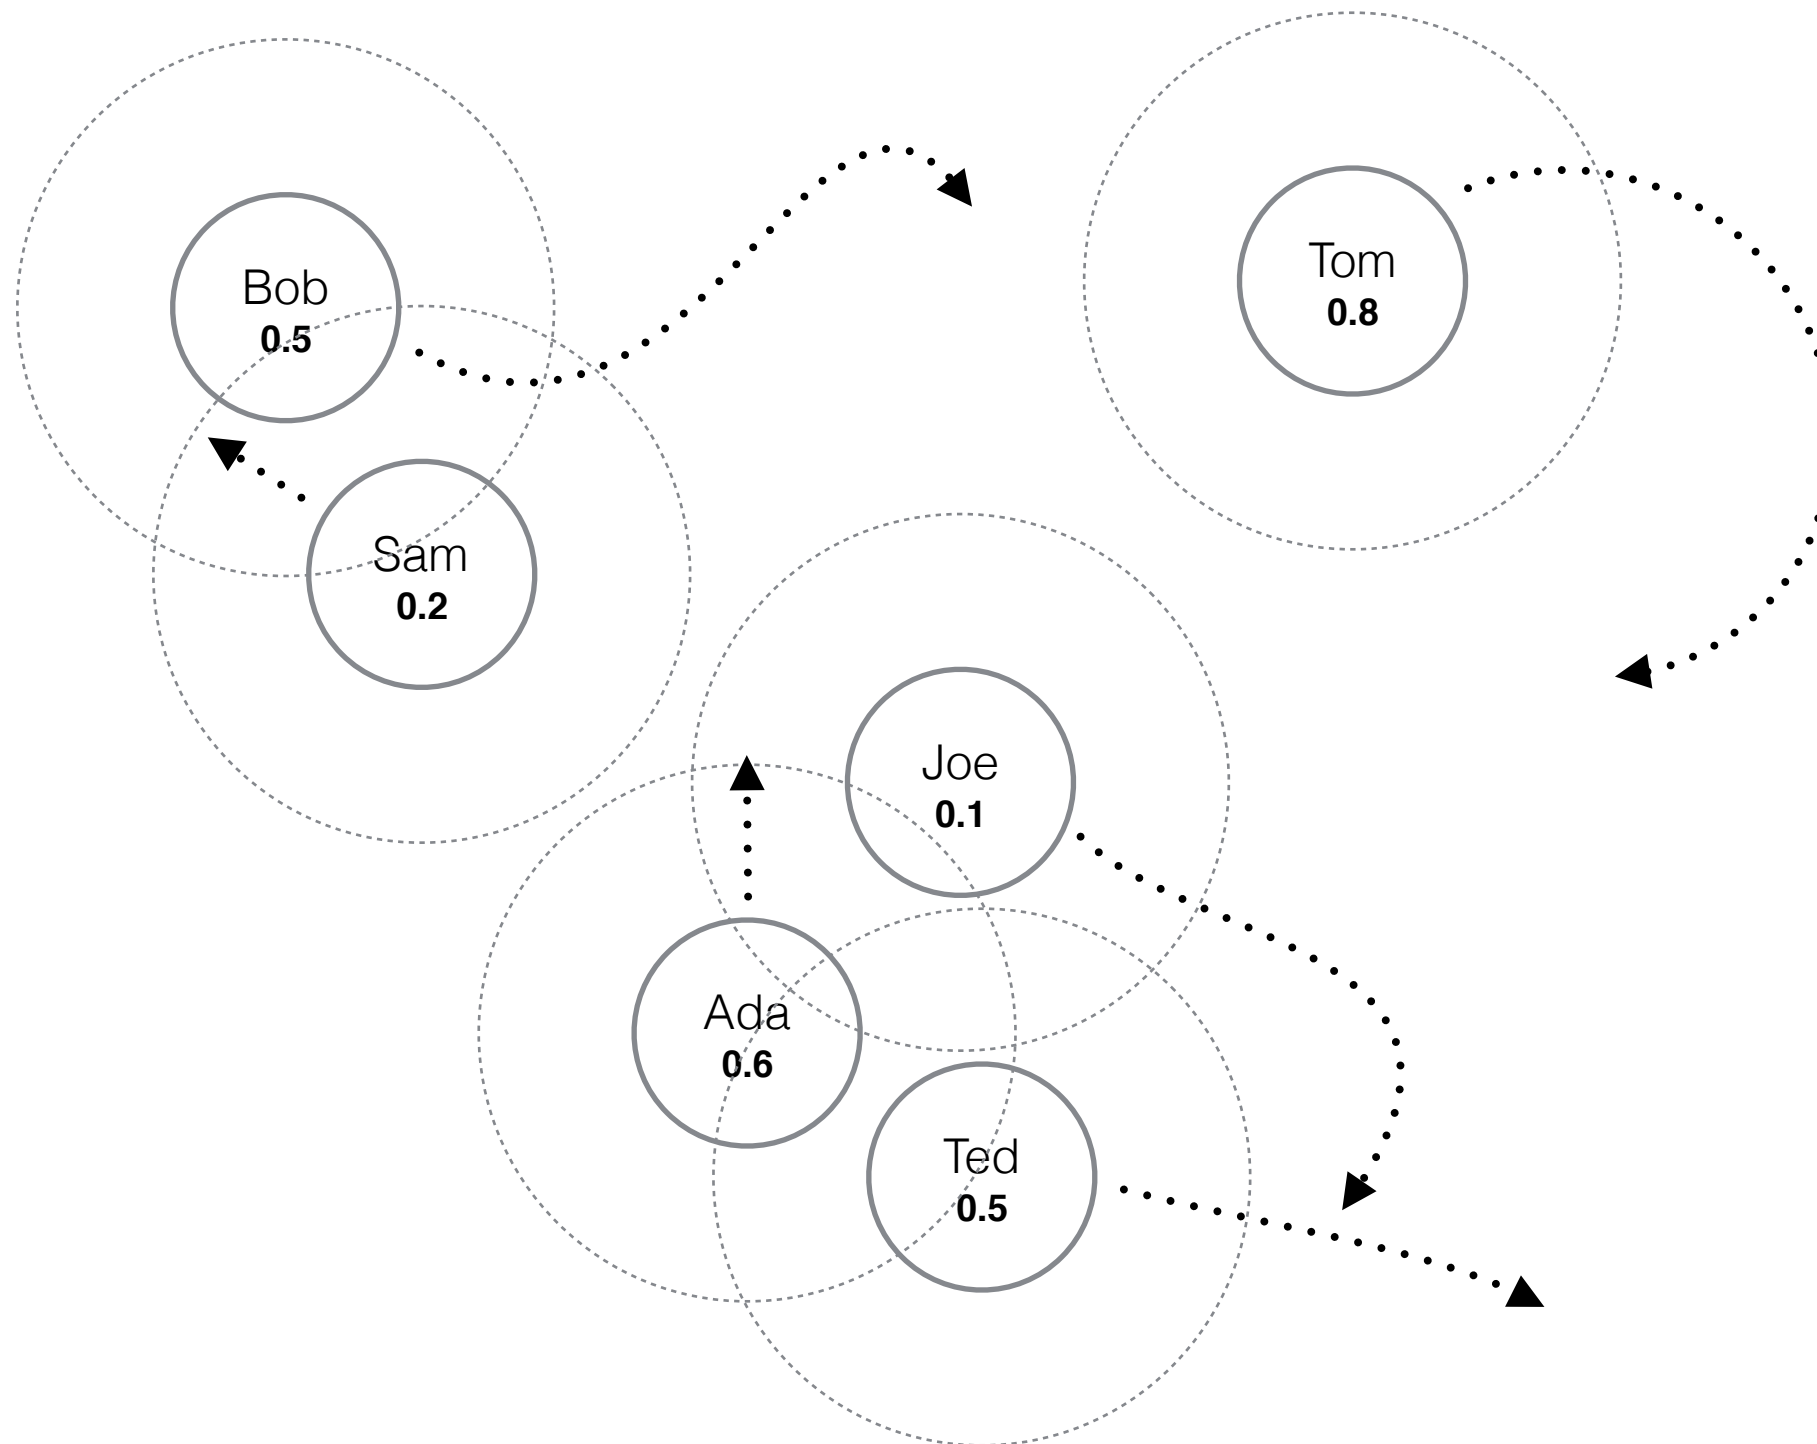
Decentralized Slicing Problem



Decentralized Slicing Problem

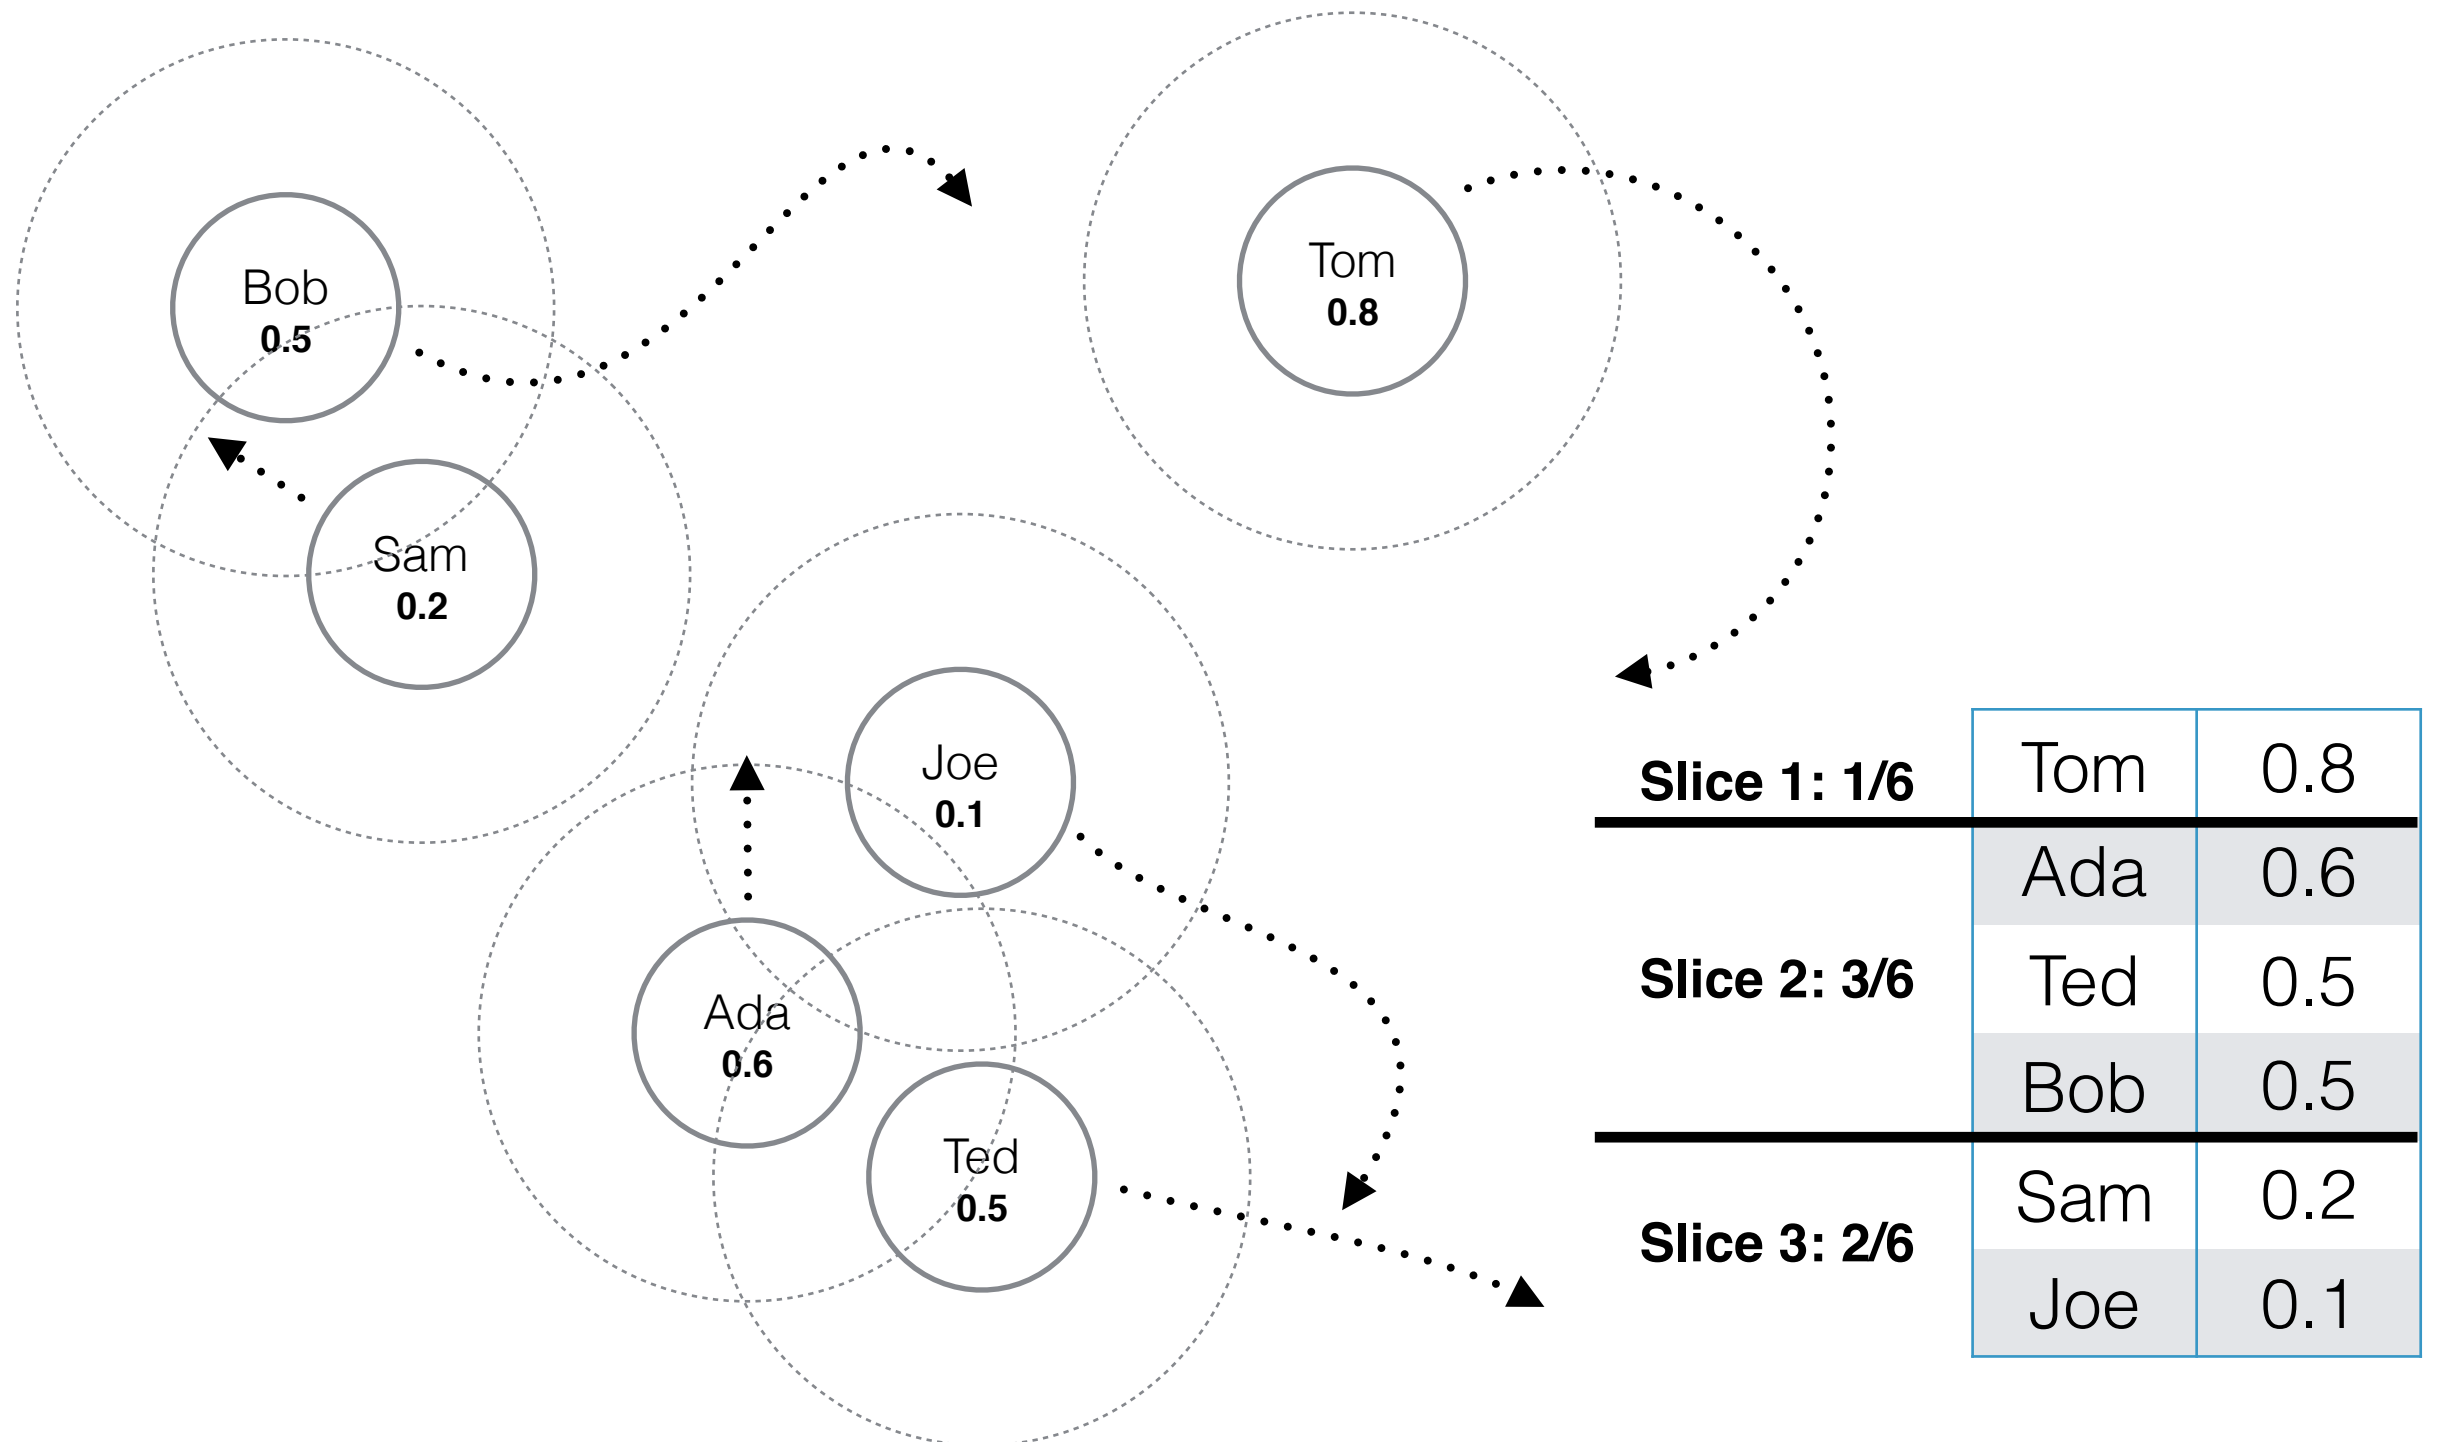


Decentralized Slicing Problem



Tom	0.8
Ada	0.6
Ted	0.5
Bob	0.5
Sam	0.2
Joe	0.1

Decentralized Slicing Problem



Slice Disorder Measure

Definition

$$SDM = \sum_{node} |actual_slice(node) - slice_estimate(node)|$$

Slice Disorder Measure

Definition

$$SDM = \sum_{node} |actual_slice(node) - slice_estimate(node)|$$

Example

	id	value	estimate
Slice 1	Tom	0.8	Slice 2
Slice 2	Ada	0.6	Slice 1
	Ted	0.5	Slice 2
	Bob	0.5	Slice 3
Slice 3	Sam	0.2	Slice 3
	Joe	0.1	Slice 3

$$SDM = |1 - 2| + |2 - 1| + |2 - 2| + |2 - 3| + |3 - 3| + |3 - 3| = 1 + 1 + 1 = 3$$

Tom Ada Ted Bob Sam Joe

Applications

- gamification mechanisms
- self-division of a robobee swarm
- finding potential cluster heads in an area hierarchy over sensors



Related work

- To this date, a few algorithms have been proposed:

JK, mod-JK, Dynamic ranking by sampling, Sliver, Q-digest

- All of the solutions proposed so far either:
 - relay on a global connectivity of a network (point-to-point communication)
 - assume that nodes are static (so an overlay network can be created)
 - were not designed for resource-constrained devices (memory or bandwidth)

Our algorithms: *BSort*



Our algorithms: *BSort*

Sam!

Bob

local: 0.5
random: 0.2

Sam

local: 0.2
random: 0.4

Ada

local: 0.2
random: 0.3

Joe

local: 0.1
random: 0.1

Our algorithms: *BSort*

Hey Sam!
my local > your local,
but
my random < your random.

Let's swap!

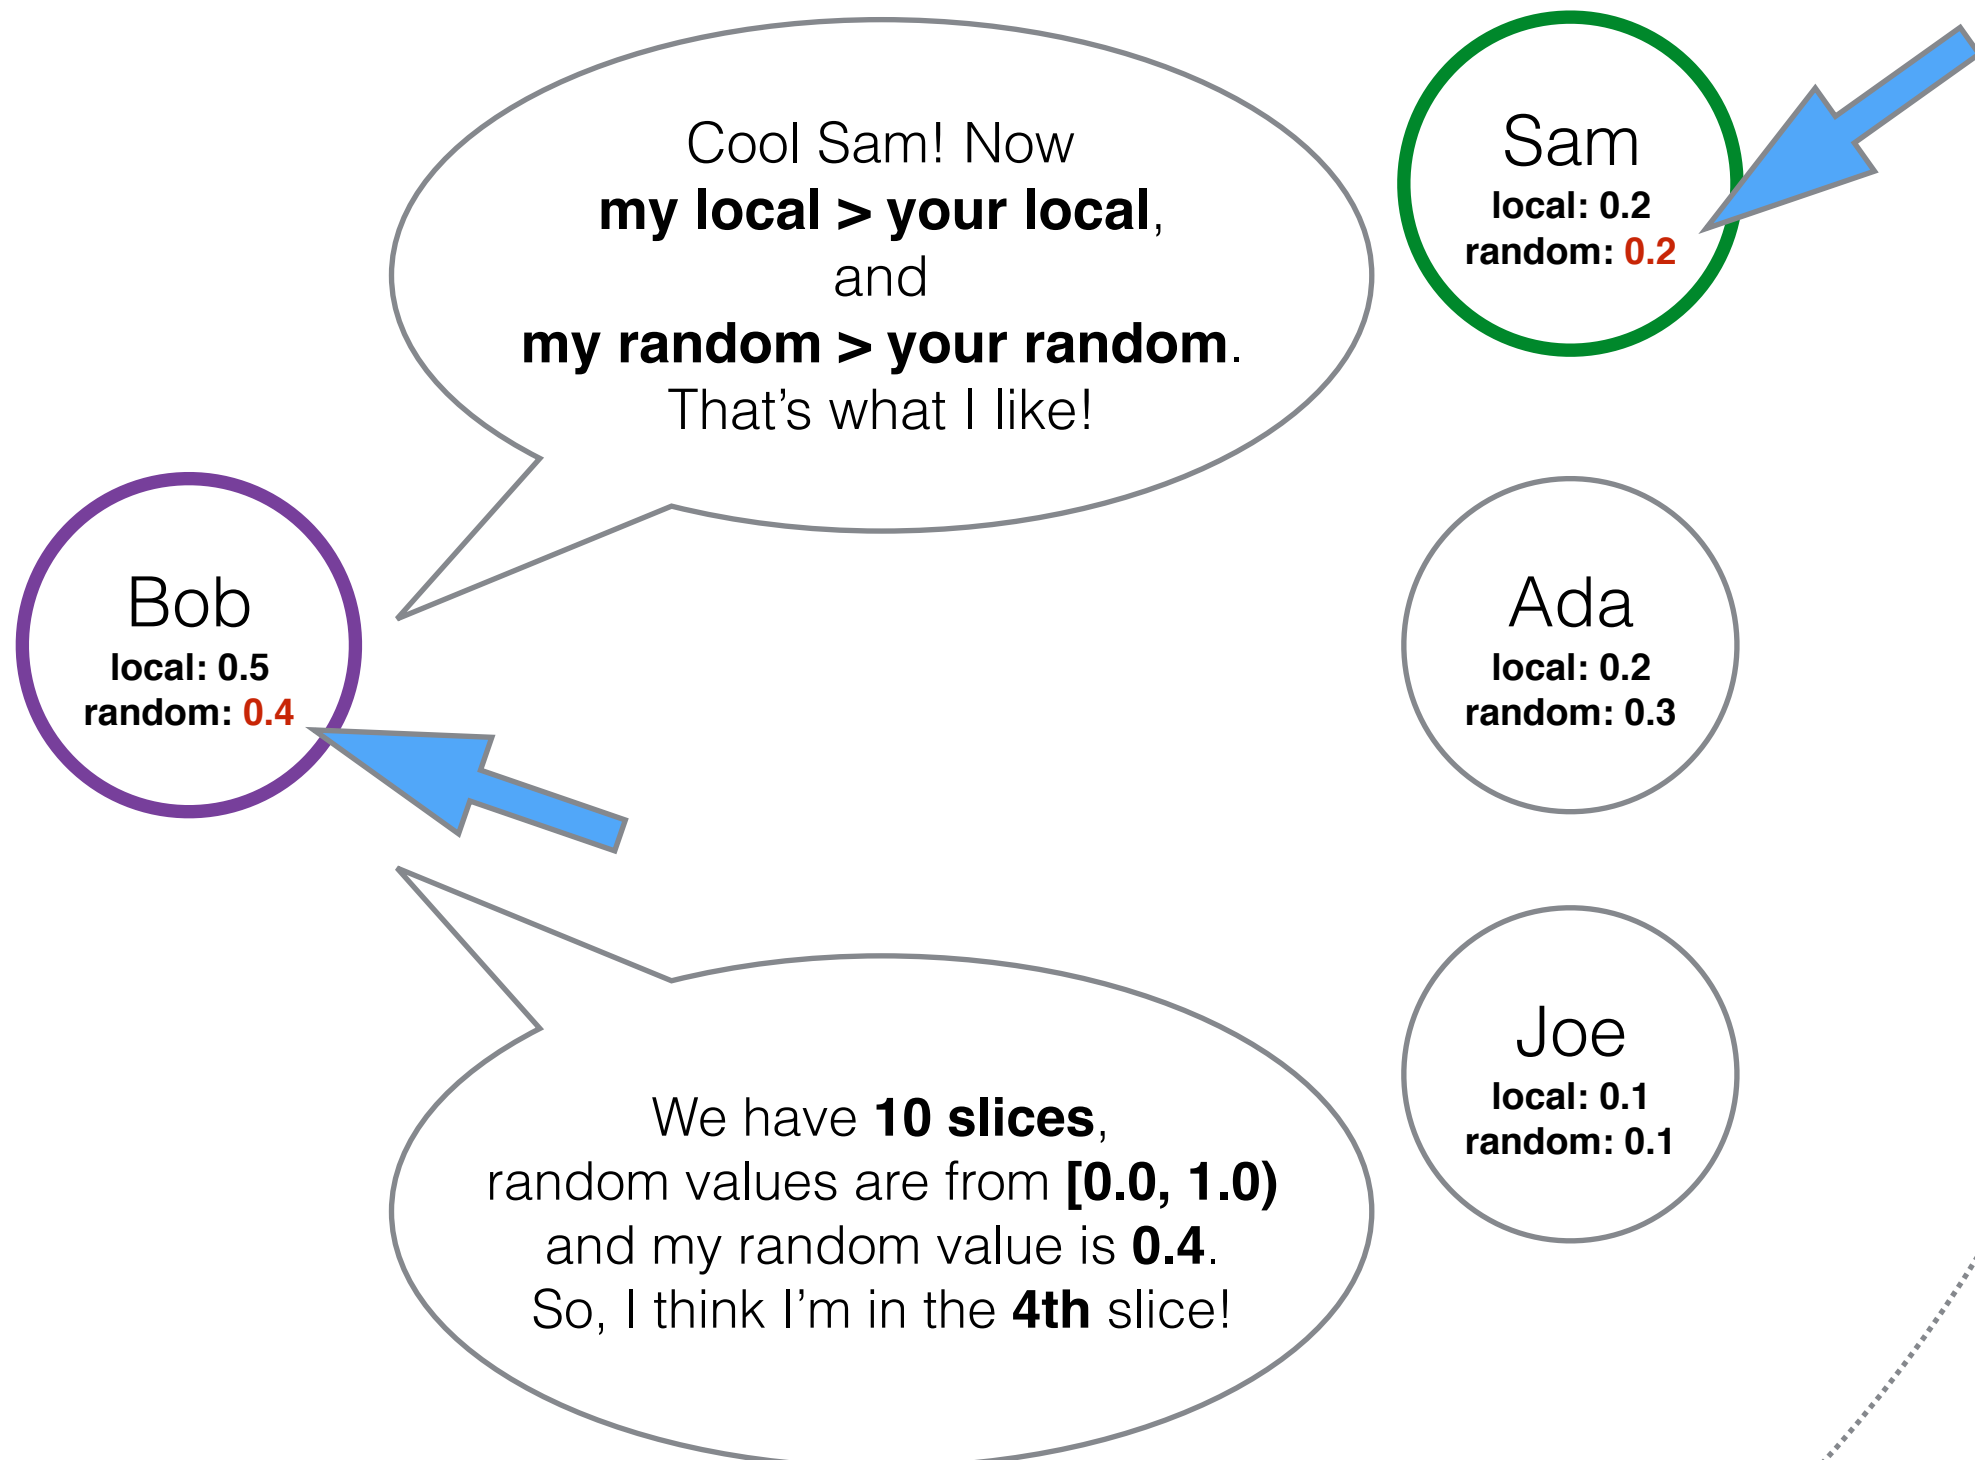
Bob
local: 0.5
random: 0.2

Sam
local: 0.2
random: 0.4


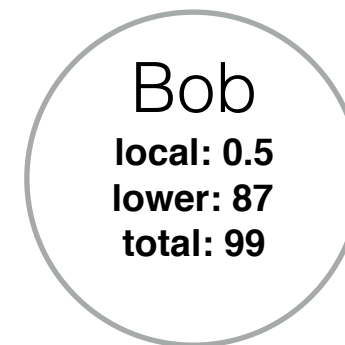
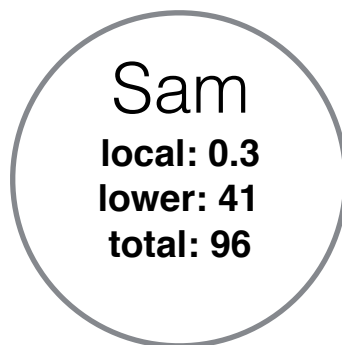
Ada
local: 0.2
random: 0.3

Joe
local: 0.1
random: 0.1

Our algorithms: *BSort*



Our algorithms: *ICount*



Our algorithms: *ICount*

Hey all, my local value is **0.3!**

Sam

local: 0.3
lower: 41
total: 96

Bob

local: 0.5
lower: 87
total: 99

Ada

local: 0.2
lower: 10
total: 62

Our algorithms: *ICount*

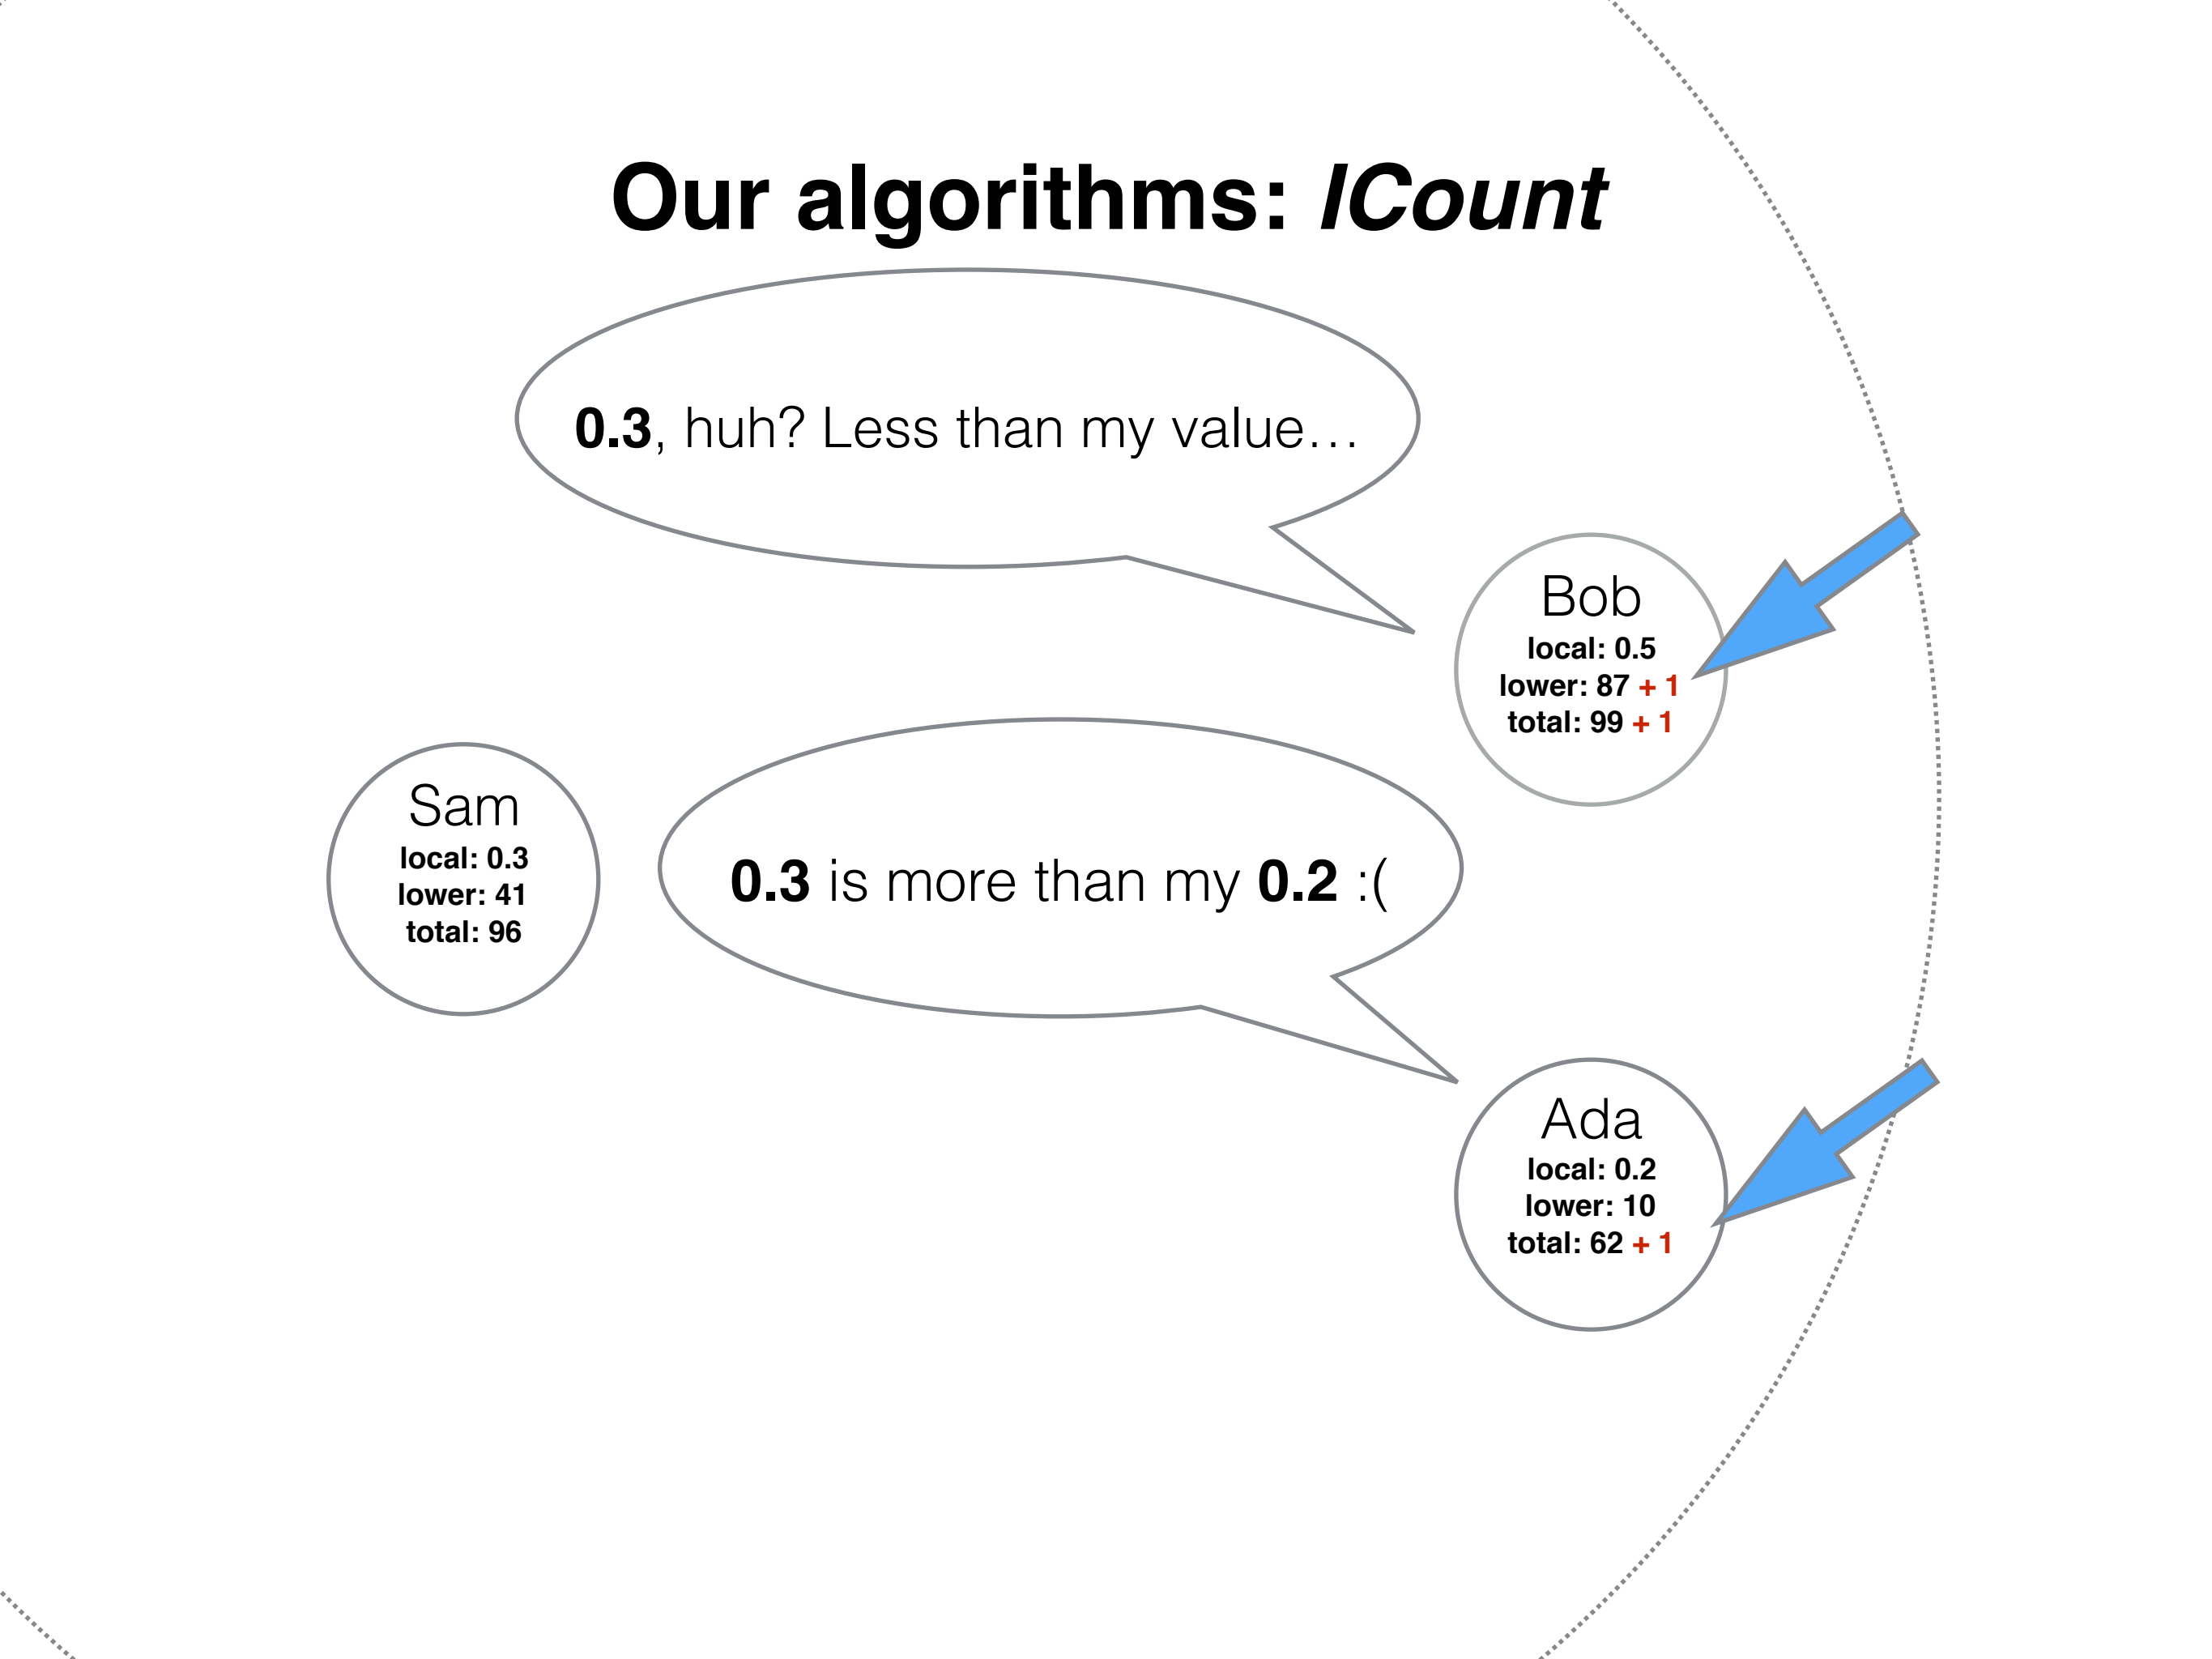
0.3, huh? Less than my value...

Bob
local: 0.5
lower: 87 + 1
total: 99 + 1


Sam
local: 0.3
lower: 41
total: 96

0.3 is more than my **0.2** :(

Ada
local: 0.2
lower: 10
total: 62 + 1



Our algorithms: *ICount*



Digression: SharedState

- A scheme for data distribution in a wireless network
- Idea:
 - Each node maintains a set of values
 - Initially nodes store only own values in their sets
 - Each node periodically broadcasts a subset of its set
 - Recipients merge local and received sets
 - Random entries are discarded to meet size limits

Our algorithms: *LCount*

Sam

local: 0.3

lower: 12

total: 54

{(Sam, 0.3), (Ada, 0.6), (Joe, 0.2)}

Bob

local: 0.5

lower: 30

total: 64

{(Bob, 0.3), (Ted, 0.6), (Joe, 0.2)}

Our algorithms: *LCount*

Hey all,
here is a sample of our population:
{(Sam, 0.3), (Ada, 0.6)}

Sam

local: 0.3
lower: 12
total: 54

{(Sam, 0.3), (Ada, 0.6), (Joe, 0.2)}

Bob

local: 0.5
lower: 30
total: 64

{(Bob, 0.3), (Ted, 0.6), (Joe, 0.2)}

Our algorithms: *LCount*

Sam

local: 0.3
lower: 12
total: 54


{(Sam, 0.3), (Ada, 0.6), (Joe, 0.2)}

Sam's value is lower,
Ada's is greater.

Bob

local: 0.5
lower: 30 + 1
total: 64 + 2

{(Bob, 0.3), (Ted, 0.6), (Joe, 0.2),
(Sam, 0.3), (Ada, 0.6)}



Our algorithms: *LCount*

Sam

local: 0.3
lower: 12
total: 54

{(Sam, 0.3), (Ada, 0.6), (Joe, 0.2)}

I need to discard
some entries from my local
set to satisfy the limit.

Bob

local: 0.5
lower: 31
total: 66

{(Bob, 0.3), ~~(Ted, 0.6)~~, (Joe, 0.2),
~~(Sam, 0.3)~~, (Ada, 0.6)}



Our algorithms: *LCount*

Sam

local: 0.3
lower: 12
total: 54

{(Sam, 0.3), (Ada, 0.6), (Joe, 0.2)}

We have **10** slices,
53% of nodes I met
had greater values than
mine, so I think I am in
the **6th** slice.

Bob

local: 0.5
lower: 31
total: 66

{(Bob, 0.3), (Joe, 0.2), (Ada, 0.6)}

Digression: Counting Sketch

- Probabilistic data structure
- Aims cardinality estimation problem
- Uses sublinear-space
- Operations:
 - *add(element)* - idempotent
 - *merge(sketch)* - idempotent, associative and commutative
 - *count()* - retrieves cardinality approximation

Our algorithms: *SCount*

Sam

local: 0.3
sketches: lower, greater

{(Sam, 0.3, lower, greater),
(Ada, 0.6, lower, greater)}

Hey all, here is some info:
{(Sam, 0.3, lower, greater)}

Bob

local: 0.5
sketches: lower, greater

{(Bob, 0.5, lower, greater),
(Joe, 0.2, lower, greater)}

Our algorithms: *SCount*

$\text{local}_{\text{Sam}} < \text{local}_{\text{Bob}}$, so $\text{lower}_{\text{Bob}} := \text{merge}(\text{lower}_{\text{Bob}}, \text{lower}_{\text{Sam}})$
and

$\text{local}_{\text{Sam}} > \text{local}_{\text{Joe}}$, so $\text{greater}_{\text{Joe}} := \text{merge}(\text{greater}_{\text{Joe}}, \text{greater}_{\text{Sam}})$

Sam

local: 0.3

sketches: lower, greater

{(Sam, 0.3, lower, greater),
(Ada, 0.6, lower, greater)}

Updating sketches...

Bob

local: 0.5

sketches: lower, greater

{(Bob, 0.5, **lower**, greater),
(Joe, 0.2, lower, **greater**)}

Our algorithms: *SCount*

Sam

local: 0.3

sketches: lower, greater

{(Sam, 0.3, lower, greater),
(Ada, 0.6, lower, greater)}

Adding received entries...

Bob

local: 0.5

sketches: lower, greater

{(Bob, 0.5, lower, greater),
(Joe, 0.2, lower, greater),
(Sam, 0.3, lower, greater)}}



Our algorithms: *SCount*

Sam

local: 0.3

sketches: lower, greater

{(Sam, 0.3, lower, greater),
(Ada, 0.6, lower, greater)}


Discarding entries to meet the limit.

Bob

local: 0.5

sketches: lower, greater

{(Bob, 0.5, lower, greater),
~~(Joe, 0.2, lower, greater)~~,
(Sam, 0.3, lower, greater)}



Our algorithms: *SCount*

$$\text{count}(\text{greater}_{\text{Bob}}) / (\text{count}(\text{lower}_{\text{Bob}}) + \text{count}(\text{greater}_{\text{Bob}})) = 12 / (12 + 68) = \mathbf{15\%}$$

Sam

local: 0.3

sketches: lower, greater

{(Sam, 0.3, lower, greater),
(Ada, 0.6, lower, greater)}

We have **10** slices,
15% of nodes I met had greater values.
I think I'm in the **2nd** slice.

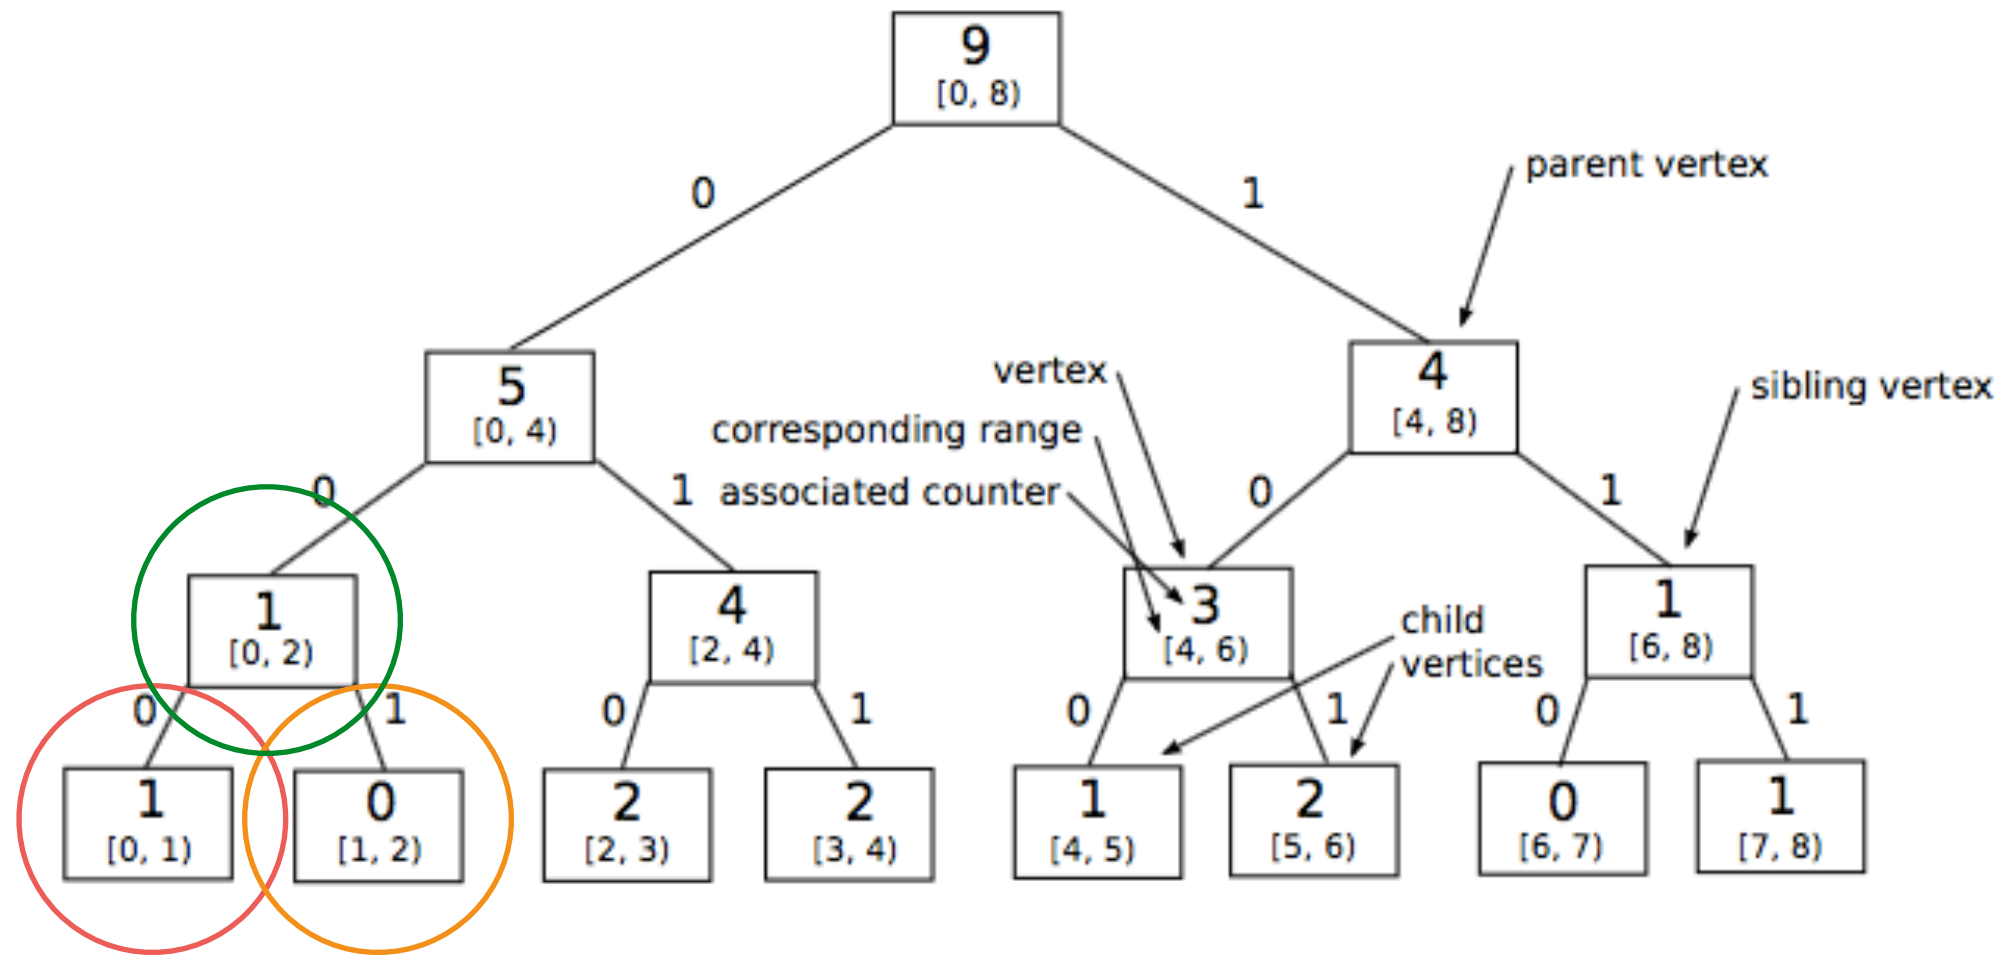
Bob

local: 0.5

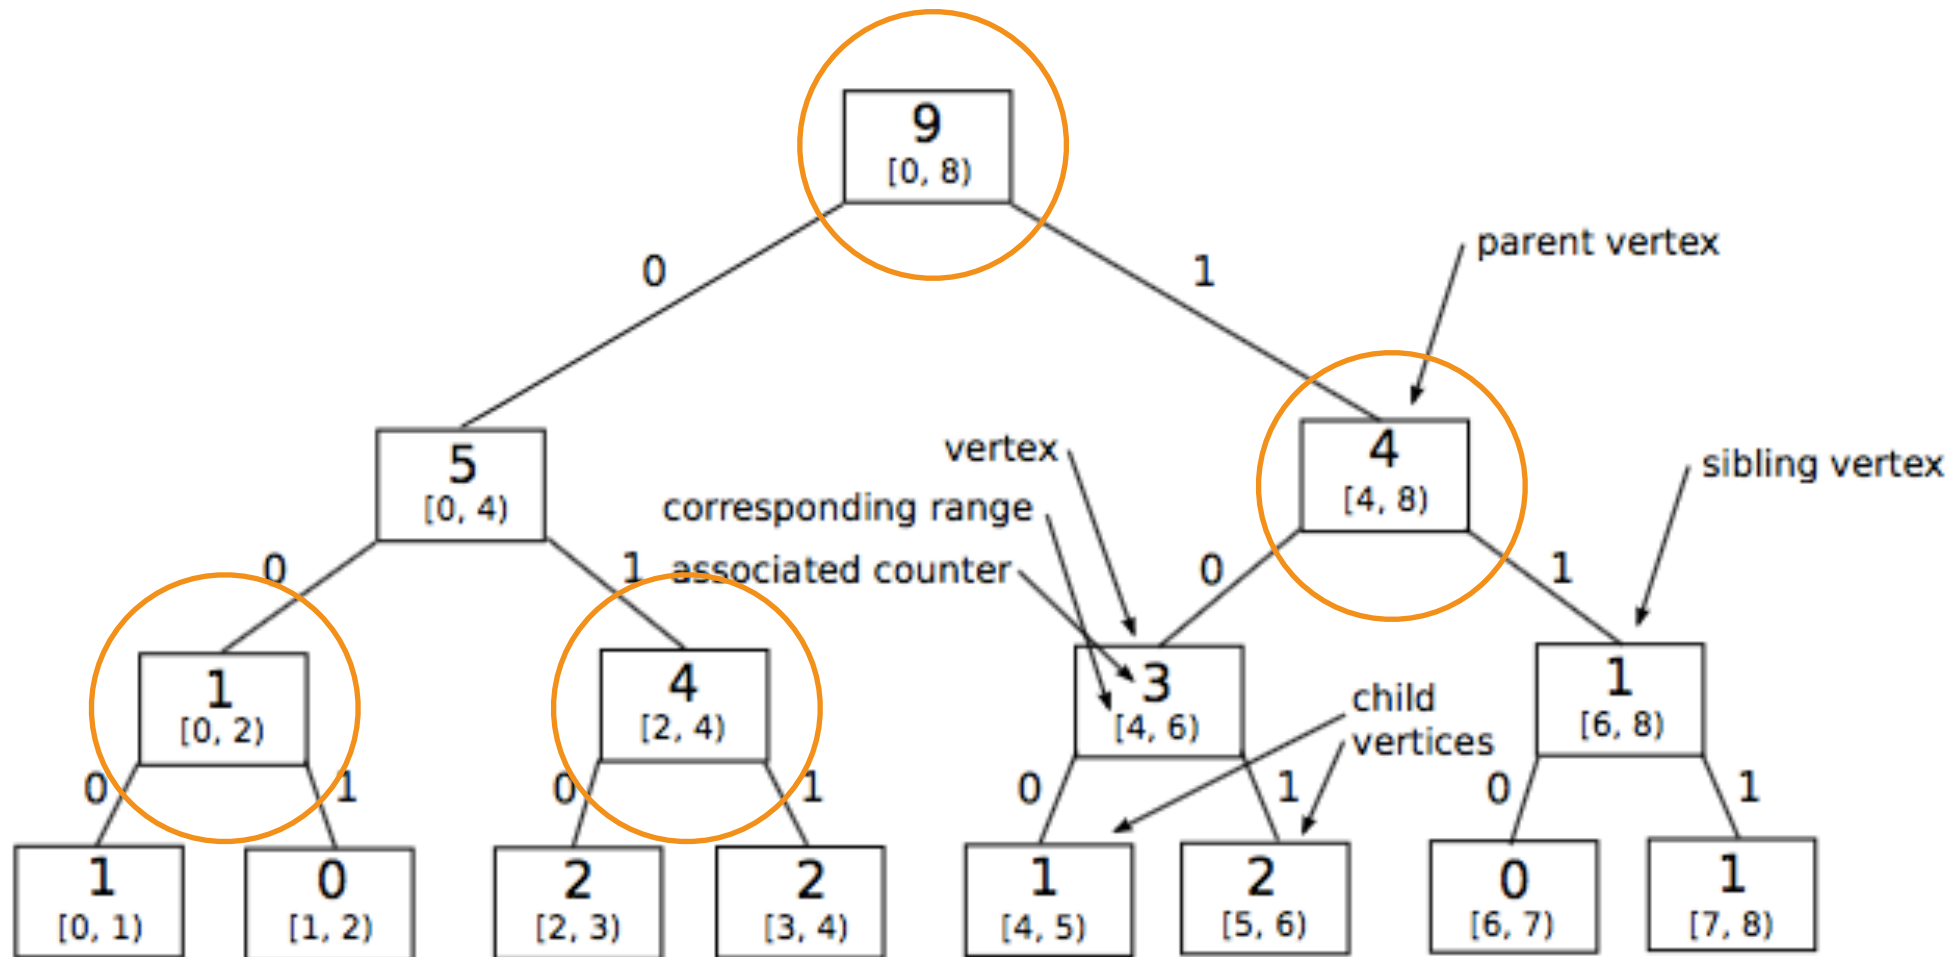
sketches: lower, greater

{(Bob, 0.5, lower, greater),
(Sam, 0.3, lower, greater)}

Our algorithms: *DTree*



Our algorithms: *DTree*




Simulations

- 1024 nodes
- effective radio range: 50 meters, 100 bytes per message
- square-shaped area, side length: 1024 meters
- 100 slices
- 3 mobility patterns:
 - static grid
 - Reference Point Group Mobility
 - Random Waypoint Mobility


Performance comparison

No mobility (grid)




Performance comparison

Reference Point Group Mobility

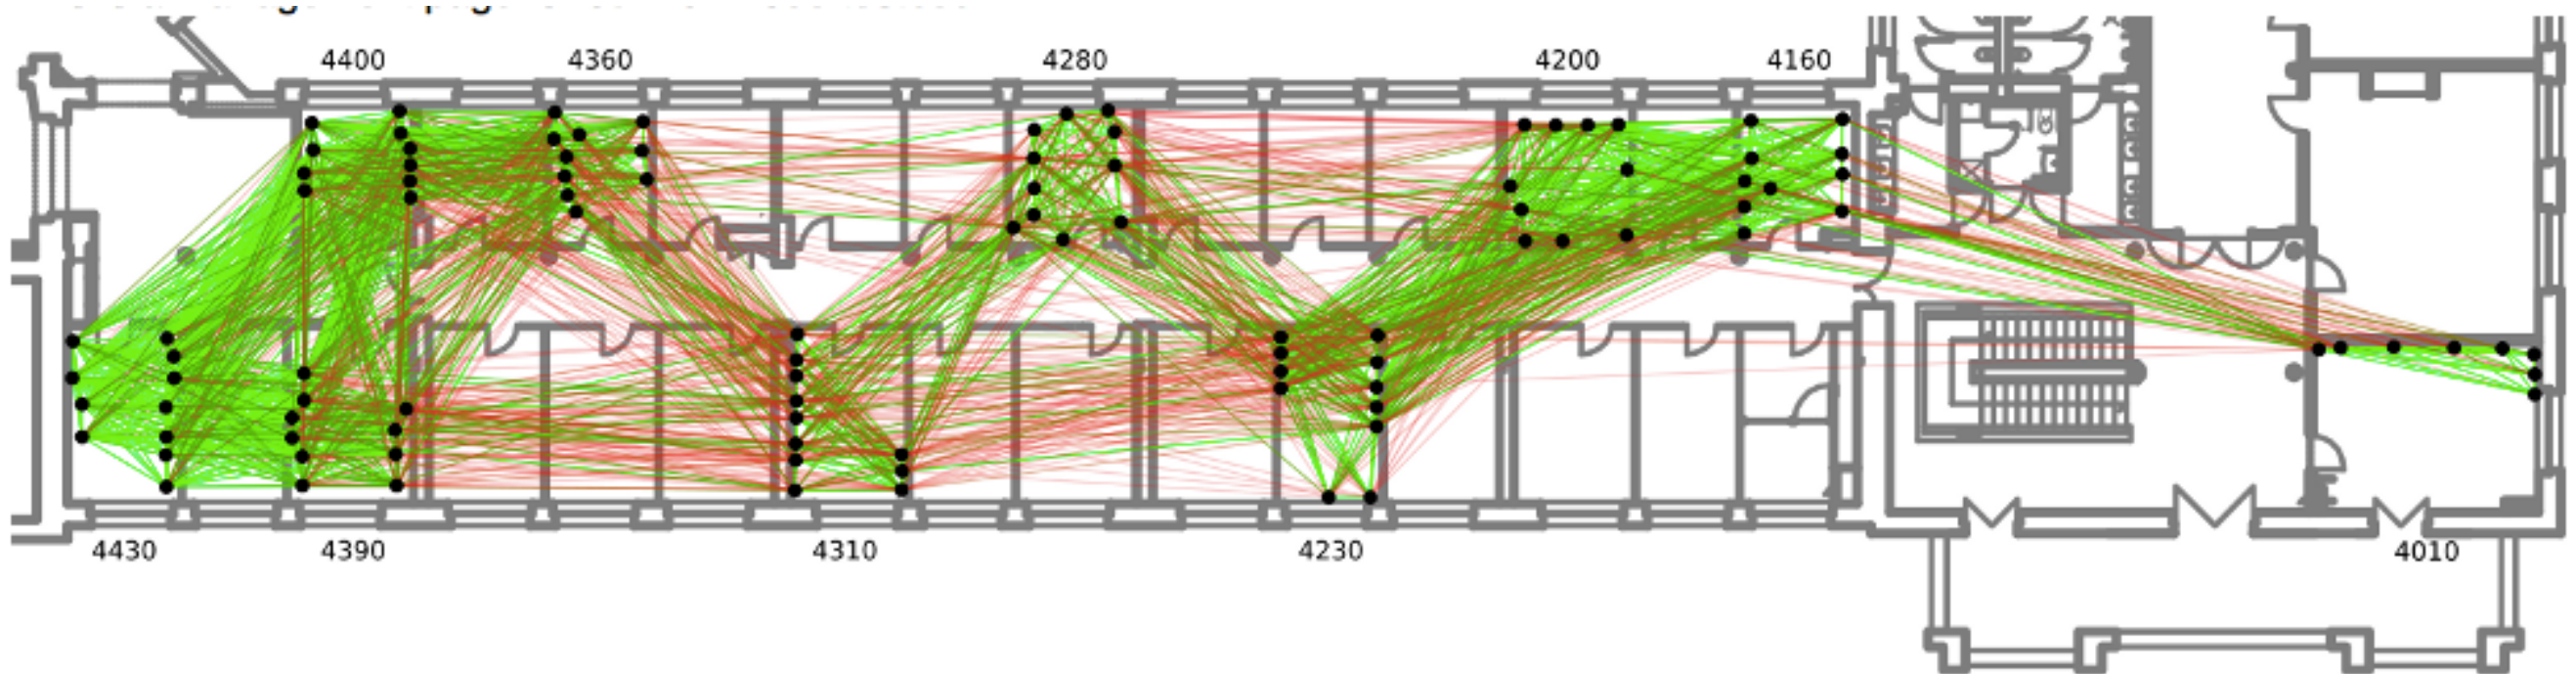


Performance comparison

Random Waypoint Mobility

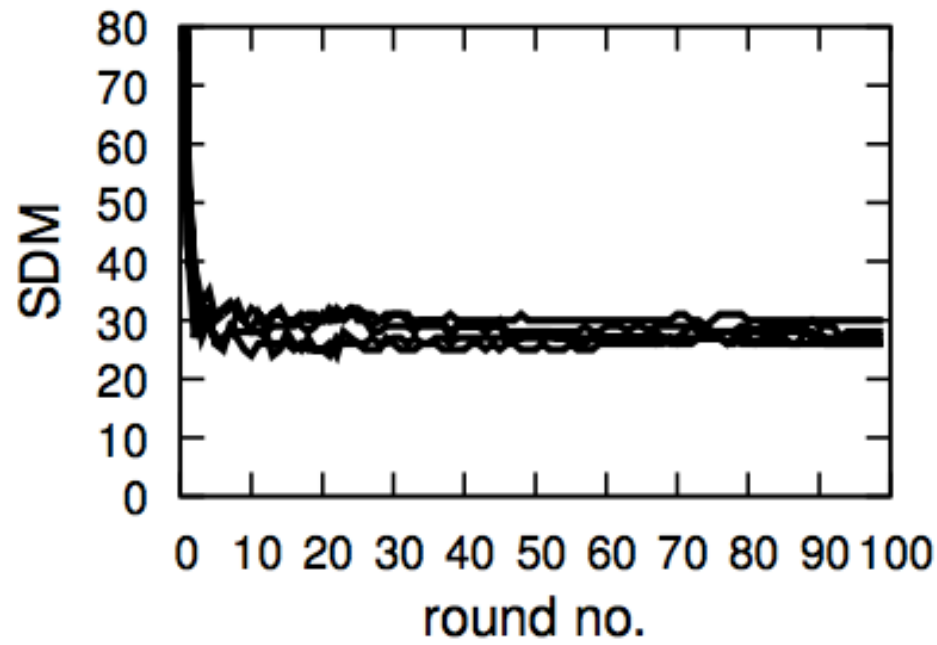


Testbed experiments




G-Node	
CPU	8 MHz
RAM	8 kb
ROM	116 kb
Radio	500 kb/s

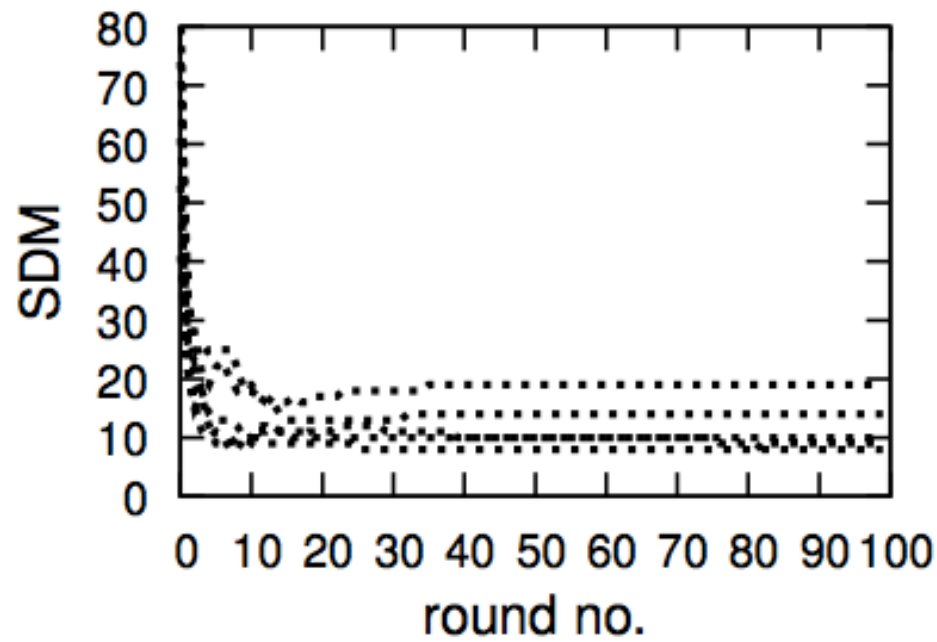
Testbed experiments




(a) *ICount*



(b) *LCount*



(c) *SCount*



(d) *DTree*

Conclusions

- There were some solutions to the Decentralised Slicing Problem
- None of them worked in highly dynamic, low-powered networks
- 2 algorithms have been adapted, 3 new has been designed
- Our new algorithms yield promising results
- Unfortunately, there is no best algorithm - performance depends on configuration, network size and mobility patterns

Thank You

Questions?

Supported by the (Polish) National Science Centre (NCN) within the SONATA programme under grant no. DEC-2012/05/D/ST6/03582. K. Iwanicki was additionally supported by a scholarship from the (Polish) Ministry of Science and Higher Education for outstanding young scientists.



NATIONAL SCIENCE CENTRE



Ministry of Science
and Higher Education
Republic of Poland

Conclusions

Algorithm	Advantages	Drawbacks
<i>BSort</i>	Uses a fixed amount of memory and bandwidth. Guarantees that slice estimates finally converge.	Does not perform well in static networks. Converges slowly. May be difficult to implement.
<i>ICount</i>	Uses a fixed amount of memory and bandwidth. Is easy to implement.	Does not perform well in static and mobile clustered networks. Converges slowly. Has its slice estimates fluctuate even if <i>SDM</i> is low.
<i>LCount</i>	Converges fast initially. Has a relatively low final <i>SDM</i> . Is easy to implement. Benefits from an increased network throughput.	Converges slowly in later stages. Has its slice estimates fluctuate even if <i>SDM</i> is low.
<i>SCount</i>	Converges fast. Has its final <i>SDM</i> independent of the device density. Benefits from an increased network throughput. Is relatively easy to implement. Guarantees that slice estimates finally converge.	Has its final <i>SDM</i> limited by the accuracy of counting sketches.
<i>DTree</i>	Converges relatively fast. Has its final <i>SDM</i> independent of the device density. Benefits from an increased network throughput. Guarantees that slice estimates finally converge.	Has its final <i>SDM</i> limited by the accuracy of counting sketches. Is difficult to implement.