Motivation

- Adapt the award policy to the growing number of submissions
- Provide guidelines for consistency in the way the awards are decided.

Outcome

Once discussed, the contents of this document will be added to the appropriate documents (e.g. Program Policy, and Handbook).

Current Practice

The current practice is the following:

1. The Program co-Chairs (PC) select a short list of papers (~10) based on the evaluation they received during the review process (overall evaluation, whether they received a rejection from one reviewer, the recommendation of AC/SAC, the “award” flag in the (meta-)review form, etc.
2. The PCs select an Award Committee (AwC) that has no conflict of interest with the short-listed papers.
3. An oral session of the short-listed papers is organised, and the AwC attends this session.
4. The AwC discusses taking all the information provided by the AC and gathered during the session, and outputs a Best Paper Award, and a Best Student Paper Award.

Issues

- Two awards for 1200+ presented papers seems unreasonable.
- The relationship between the (general/program) chairs and the AwC is not clarified.

Proposal

Aim: Give a reasonable number of “Outstanding Paper Awards” to a percentage of the accepted papers, and build a process to identify several “Best Paper Awards”. The rationale is for every 100 papers to give 5 Outstanding Paper Awards and one Best Paper Award.

Process:

1. The PCs select the top 5% (ceiling it) of the accepted papers. This selection is based on quantitative criteria (such as the ones used in current practices) and is done “per topic” (see the table for an exemple).
2. The PCs select two AwC co-chairs that have no conflict of interest (as defined in the ACM texts) with the PCs and with the General co-Chairs. The AwC and the PCs select/invite a few (~5) senior members of the community to take part in the AwC.
3. The criteria to select the Best Paper Awards must be scientific, and can include any information in the review process, the paper itself, and even the presentation (the contents of the oral presentation and of the poster). The best paper awards can only be given to the top selection of Outstanding Paper Awards (top 5%, see above).

4. The number of Best Paper Awards is also computed "per topic" using the following rules to ensure a minimum and maximum of competition.
    a. The number of Best Paper Awards per topic is tentatively computed as ⅕ of the number of Outstanding Paper Awards.
    b. Topics for which this computation is "0" will be grouped and considered as a single meta-topic. The number of Best Paper Awards for this meta-topic is ⅕ of the number of Outstanding Paper Awards (with a minimum of 1).

The following two tables provide two examples to compute the # of Outstanding and Best Paper Awards per topic. We see that the overall probability of Outstanding Paper Award is around 5%, and of Best Paper Award, around 1%.

Example 1: Four topics (A, B, C, D). The number of Outstanding Paper Awards being smaller than 5 for topics C, and D, they are grouped into one meta-topic.

Topic	# Accepted	#Outstanding (~5%)	Probability Award	#Best Papers (~1%)	Probability Award
A	327	17	5.20%	3	0.92%
B	128	7	5.47%	1	0.78%
C	15	1	6.67%	1	3.33%
D	13	1	7.69%		3.85%
Total	483	26	5.38%	5	1.04%

Example 2: Statistics from ACM MM 2022. Topics A, B, C would have had enough Outstanding Paper Awards to get a Best Paper Award on their own. The other topics would have been grouped and a fifth of the total number of Outstanding Paper Awards (15/5=3) would have become Best Paper Awards.

Topic	# Accepted	#Outstanding (~5%)	Probability Award	#Best Papers (~1%)	Probability Award
A	155	8	5.16%	1	0.65%
B	149	8	5.37%	1	0.67%
C	103	6	5.83%	1	0.97%
D	75	4	5.33%	3	1.07%
E	45	3	6.67%		1.33%
F	21	2	9.52%		1.90%
G	20	1	5.00%		1.00%
H	18	1	5.56%		1.11%
I	13	1	7.69%		1.54%
J	9	1	11.11%		2.22%
K	6	1	16.67%		3.33%
L	5	1	20.00%		4.00%
Total	619	37	5.98%	6	0.97%