

Report of the Task force on Applied and Experimental Papers at SoCG April 10, 2025

Computational geometry has been long recognized as a field providing theoretically sound solutions to relevant practical problems. As we know, it has strong connections to fields such as computer graphics, computer-aided design, robotics, machine learning, and many others. However, despite these strong connections, the core results published at SoCG have mainly focused on mathematical results, and for many years it has been difficult for researchers both within and outside of our community to feel supported in work on implementation and application results. It is important for the future of our community as a whole that we conduct a serious effort to not only focus on proving theorems, but also to demonstrate progress in facilitating actual computation: Only then can we ever hope to be more than just a specialized subcategory of a top theory conference such as SODA.

Over the last decades, efforts from members of the SoCG community have also resulted in tremendous progress in understanding how geometric algorithms should be implemented, along with experience in successfully applying such implementations. Practical work is celebrated, as evidenced by the SoCG 2023 Test of Time award for the CGAL library, but many have trouble getting that work published in our flagship conference. Although experts in our community are capable of producing high-quality papers in this area, many of those experts believe that the current structure of SoCG does not always reward or support submissions of this type. In particular, despite separate paper types for Implementations and Applications, such work is sometimes judged primarily by theoretical criteria and thus ultimately rejected. (We defer a discussion on what constitutes an applied/experimental paper to the appendix, as we realize there has been considerable discussion even on the definitions of such papers.)

This situation has several negative long-term consequences for the well-being of SoCG. First, the scientific impact and perceived relevance of SoCG is diminished – as evidenced by the recent rejected proposal for the A+ status of SoCG. Strong implementation and application papers open an entirely new avenue for citations and impact in other fields, and even the more theoretical researchers could benefit from this increased impact. Moreover, the difficulty of getting recognition for this practical aspect of our work discourages many students from pursuing practical results, so we are losing strong intellectual talent with a preference for practical computation, who move on to other fields. While this may lead to some personal connections that can be beneficial at times, all members of the task force agree that SoCG should aim to be the undisputed home of geometric work, with researchers from other communities following developments published here. Finally, a lack of regard for practical aspects also diminishes the potential broader societal impact of what we do and teach (including the value of our deep theoretical insights), as the majority of our students will eventually leave academia and may feel a disconnect from the “real world” challenges they are facing.

While we can disagree about the severity of some of these impacts, it seems clear that it is time for the community to revisit the topic of applied and implementation-based contributions at

SoCG and how to best support this critical work. With a long preceding history, the current discussion began in SoCG 2024 at the discussion forum, followed by an overwhelming vote in support of such a discourse. Thus, this task force was convened in the fall of 2024, and has assembled a slate of suggestions for possible revisions to SoCG.

The committee reviewed models that have been adapted in other communities. For example, SODA has a separate conference (Alenex) for engineering and experiments. After discussions, the committee did not support such an option for SoCG, as we feel the risk with separate conferences, such as this, is that implementation and application papers may look second class. Also, the PC might feel a pressure to accept papers even if they are not very strong, in order to make sure that they have enough talks to make a program.

Proposal: SoCG track B

Following these objectives, we propose to strengthen SoCG as a whole by splitting the existing four paper types into two tracks. The second track will focus on implementations and applications, ensuring that such work will be reviewed by experts in such topics. No changes to the mathematical foundations and complexity papers will be made.

To achieve the desirable significance of this new track, we favor the tried and tested format of ESA, which offers a Track A (aimed at algorithm theory in ESA) and a track B (aimed at algorithm engineering in ESA). Each track has its own program committee and separate chairs, but the PC chairs work together to jointly create a program for ESA, with no explicit distinction in the proceedings and in the program. Scope and objectives of this SoCG track B will be adjusted to fit the needs of our community, including the submission of data and software, and the addition of 3 to 4 “software testers” for evaluating the viability of computational contributions. However, SoCG will keep its high quality standards, and the SoCG B PC will not have any pressure to accept submissions. Weak submissions will be rejected at SoCG B, just as weak theory papers are rejected at SoCG (A) currently, but the addition of the extra track will allow for the applied papers to be reviewed by a group with expertise in this domain and desiderata to focus on these aspects of the submissions.

An (unofficial) rough estimate is that in recent years the size of ESA B has been about one sixth of ESA in the number of papers, with a PC for Track B of roughly one third of the size of Track A. Following this estimate, the number of papers in SoCG would increase from about 75 to about 90 in total. Accommodating this growth could be achieved without any loss for the SoCG track A by simply adding a third parallel session to the program, without diminishing the space for theoretical work at SoCG. Of course, these numbers may increase with time, and adjustments can be made as needed.

Ensuring the strength of SoCG B could be achieved by incorporating a spectrum of guidelines, including but not limited to

- the use of real-world data

- consulting reviewers from other fields to ensure that the approach is relevant
- making code available
- checking the correctness of code as part of the review process
- requiring comprehensive benchmarking (for implementation work focusing on efficiency)

We propose that the details of these guidelines should be specified by the first PC, with input from the community, and evaluated and adjusted each year as needed. We expect a period of flexibility and adaptation being necessary, as there will be a period of adjusting details based on the success and volume of submissions.

Task force members: Erin Chambers, Danny Halperin, Sándor Fekete, Jeff Phillips, Monique Teillaud, and Hubert Wagner

Appendix:

Discussion: What is an applied or experimental paper?

Paper types were added in SoCG 2016, and despite several years of usage and many discussions, it seems clear that there is confusion surrounding the “application” and “implementation” paper types. We did discuss various aspects to clarify the criteria, which we enclose below. We emphasize that this is just an initial framework, which we hope will guide discussion on how to best support these papers moving forward.

A strong implementation paper:

- presents novel insights for an algorithm that are necessary to make it work in practice, explains the difficulties/obstacles to implement an algorithm, how they can be overcome, the choices made,
- describes the way relevant input data can be found or generated (if there are no existing benchmark datasets, creating one is a big plus)
- shows experimental results and analyses them
- compares with previous implementations (if any) ideally on a wide range of inputs
- provides an actual practical tool, which is publicly shared and successful on a wide range of relevant data types

Such a paper might not target any specific application. It can be a companion of a theory paper published earlier at SoCG or elsewhere. The motivations for the work are diverse: present a first implementation to solve a problem, compare algorithms in practice, check a theorem in practice, explore conjectures on the complexity of a problem or an algorithm, propose heuristic for the solution of a problem for which no exact solution is known, provide significant improved empirical runtime over established methods, etc.

A strong application paper:

- presents the modeling of an applied problem into a CG/CT problem,
- justifies why the modeling is faithful to the actual problem,
- shows and analyses results on data coming from the application domain
- explains the geometry and topology theoretical tools used to solve the problem
- compares with previous work on the same problem. If there are no established methods for a given problem, the paper must make a compelling case that the proposed solution is relevant.
- potentially ‘imports’ a new nontrivial geometric problem from another field

A successful applications paper can also be a companion of a paper published in an applied conference, in which primarily results matter, but which includes neither the modeling step nor the theory – or significantly simplified versions of them. Such a work could also be a proof-of-concept paper (that could have a follow up in a more applied venue) so that the modelling/theory/algorithms get the SoCG-stamp-of-approval.

Regardless of the adoption of our proposed changes, we hope these definitions can be discussed and refined, with more careful evaluation criteria and discussion of the type, as well as care taken in choosing reviewers who are skilled in each area. That said, we propose the best model for evaluating such work is a separate PC, who can determine review criteria and continue to evaluate the process.