The Environmental Cost of Our Conferences: The CO2 Emissions due to Travel at PODC and DISC

Laurent Feuilloley and Tijn de Vos

1 Introduction

The symptoms of the ongoing climate crisis are becoming more and more prevalent, with floods, droughts, and other extreme weather events taking place more and more often [12]. It is clear that action has to be taken now. In the Paris Agreement, world leaders have committed to limiting global warming to 1.5°C [16]. This means emissions must peak before 2025, should decline by 45% by 2030, and reach net zero by 2050. In the years since the Paris Agreement (2015), computer science conferences have been taking little to no action on this front, although the issue has been put on the table several years ago [17]. A commonly-voiced opinion is that our conference-based system is bad for the climate, but is difficult to change. Discussions on this topic, *e.g.* at business meetings, often mention that we do not have enough data, and that some alternatives could exist, but would have some drawbacks, and would maybe not be effective.

In this column, we give a detailed account of the work done by the authors to gather data, estimate the current emissions, and predict the effect of different changes in the conference system. Since we belong to the community that studies the theory of distributed computing, sometimes shortened in PODC/DISC community (for its two largest conferences), our study is about this community.

Our interest was triggered by discussions about the impact of the conference location on the average emissions of the participants. This is why we have a focus on geographical data, which might not be very relevant for another research community. Nevertheless, we believe that the comparison of the different alternatives should give a good insight for any conference where usually a fair share of the participants take transcontinental flights to attend.

Last but not least, we acknowledge the following disclaimer: we are computer scientists, not carbon footprint experts. Thus, the precise numbers we provided may well be questioned, and potentially corrected. But we have used several expert sources, made the computation as precise as possible, and estimated the sensitivity of the model. We therefore believe that, broadly-speaking, the trends we provide in our study are reliable. Also, all our computations are reproducible, and our (anonymized) dataset is public.[1]

In the rest of this introductory section, we introduce the quantities we are estimating, state the results of our work and our own conclusions on the topic. In the rest of the column, we will detail the methodology we followed to obtain these results.

What we estimate and what we compare against. We start from the reasonable premise that almost the entire climate impact of a typical international TCS conference comes from participant travel. Hence,

[1]The spreadsheet is available here: `https://perso.liris.cnrs.fr/lfeuilloley/autre/carbon-emissions-disc-podc.xlsx`. Any comments or questions are welcome.

we will focus on this aspect[2]. To measure this impact, we estimate the average CO2 equivalent emissions per DISC/PODC participant for the return trip to the conference location. The CO2 equivalent emission, denoted CO2e, gathers the impact of all greenhouse gases in one number (see Section 2.5 for more details). All numbers in this report will be *CO2e tons*.

Before we give our results, let us list some CO2e emissions per year, to get points of reference.

- Annual emissions per capita (world): 5 t CO2e.
- Annual emissions per capita (Europe): 8 t CO2e.
- Annual emissions per capita (USA): 15 t CO2e.
- Annual emissions for one car (12,000 km; middle class model): 2 t CO2e.
- Annual emissions for running an average fridge: 0.3 t CO2e.
- Annual emissions budget for one person for Paris Agreement: **1.5 t CO2e**.

The last number is a rough indication. One will find several such estimates while browsing the Internet. This estimate of 1.5 t CO2e/person is obtained as follows. The Paris Agreement sets the objective of being carbon neutral in 2050, that is, 0 t CO2e/person. To stay below a temperature increase of 1.5°C, the Intergovernmental Panel on Climate Change (IPCC) reported in 2018 that we have a CO2e budget of roughly 420 billion tons [2]. Divided out over 30 years and 8.8 billion people (taking population growth into account) gives 1.5 t CO2e/person for the years until 2050. See also *e.g.* [5].

In this report, we will often compare the emissions of an average participant going to a conference to this figure of 1.5 CO2e tons. Of course, one has to take into account that conference participants have other sources of emissions (their daily life, vacations, and other conference trips), but we find it useful to provide a concrete number to compare to.

Results on the impact of the conference location. We estimated the average emissions for various locations of PODC or DISC. Figure 1 reports the estimates of our model. The locations chosen are the ones from our dataset (which allows validating the model), and a few cities that are either on the table for future PODC/DISC locations, or are classic locations.

We see that different locations correspond to very different average emissions per participant: it almost triples between the lowest ones and the highest ones. Setting the objective to below 3 tons (which is already twice the threshold for Paris Agreement) already prevents locations that are not East-Coast US/Canada or Europe. Of course picking conference locations is a complicated topic, with many criteria, but given these numbers it appears that the carbon emissions should be one of the main criteria, and that even choosing the best places for this criteria alone is far from enough.

Impact of two relaxations of the mandatory attendance. A key feature of conferences is mandatory attendance: at least one of the authors of each paper has to be on-site, and this is one of the key points in terms of climate impact (see [3] for a discussion of this aspect in the Bulletin of the EATCS). We measured how two relaxations of this rule impact CO2e emissions. In the theory of distributed computing, most of the participants come either from North America or from Europe (including Israel), and we use this observation to exemplify the two changes. Of course, this has to be adapted for any specific research community. Also, the text is written as if each paper had only one author, and again, this has to be adapted when there are several authors, especially if they come from different continents.

[2]We have made more precise computations for the case of DISC 2023, but simple back-of-the-envelope calculations for other conference venues also lead to the same conclusion.

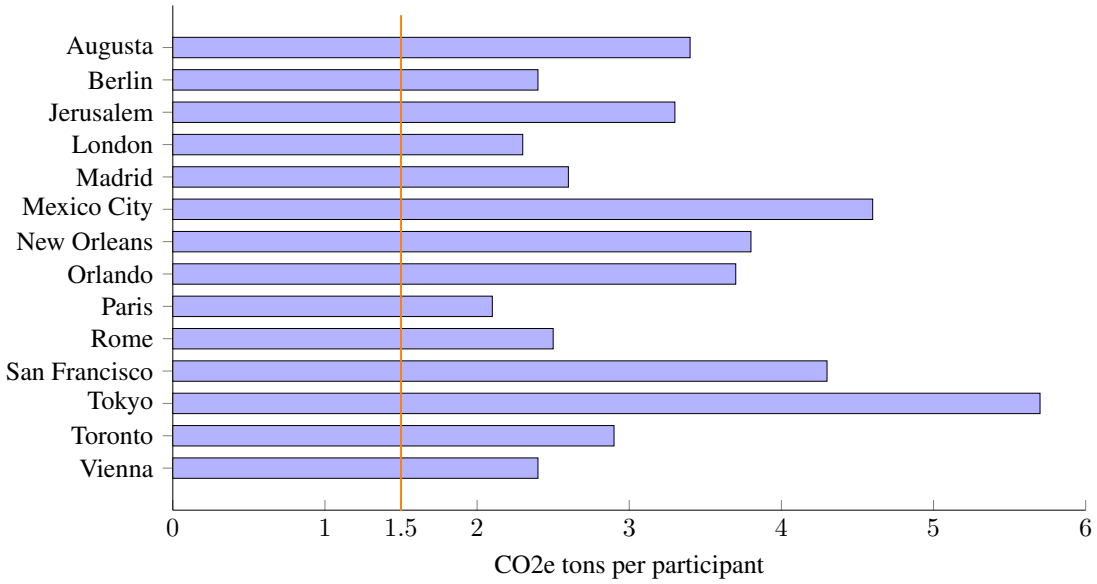

Figure 1: Average CO2e emissions per participant by conference location (alphabetically sorted).

- Limited exemption to mandatory attendance: a few authors (*e.g.* 10%) are allowed not to attend the conference, because the flight would be very long, *e.g.* researchers East Asia, or South America. (This exemption could come with some requirement, such as giving the talk online, finding someone on-site to present the paper, having a prerecorded video, etc.)
- No mandatory transatlantic flights: no author from Europe or North America is forced to take a transatlantic flight, but attendance is mandatory when the conference is on the same continent. This can be organized in different ways, either having the conference taking place in two places at the same time (the 'two hub' option), or allowing the presentation to be delayed to the next edition, or to be given at another conference (*e.g.* PODC instead of DISC), having videos, online talks, etc.

Quick overview of other alternatives. We list below some other alternatives that are aware of. We encourage the community to consider and propose other alternatives.

- Carbon offsets: this consists in paying companies to 'compensate for' carbon emissions, by investing in sectors that remove CO2 from the atmosphere (*e.g.* planting trees). Companies often like to claim that something is carbon-neutral, since they have purchased the equivalent number of carbon offsets. This phenomenon also appears in the realm of conferences [19]. Unfortunately, climate advocates seem to agree that the current system of voluntary carbon offsets is not working, and could potentially even be harmful. See, for example, this Guardian article [9].
- Reducing the number of conference editions: that is, having a conference every x years, instead of every year, or even stopping some conferences.
- Co-locating conferences: having several conferences at the same time, or one after the other, in the same place. In the same spirit of making the most of a flight, one can encourage research visits close to the conference place. (See the Extended Stay Support Scheme implemented by ICALP 2022, DISC

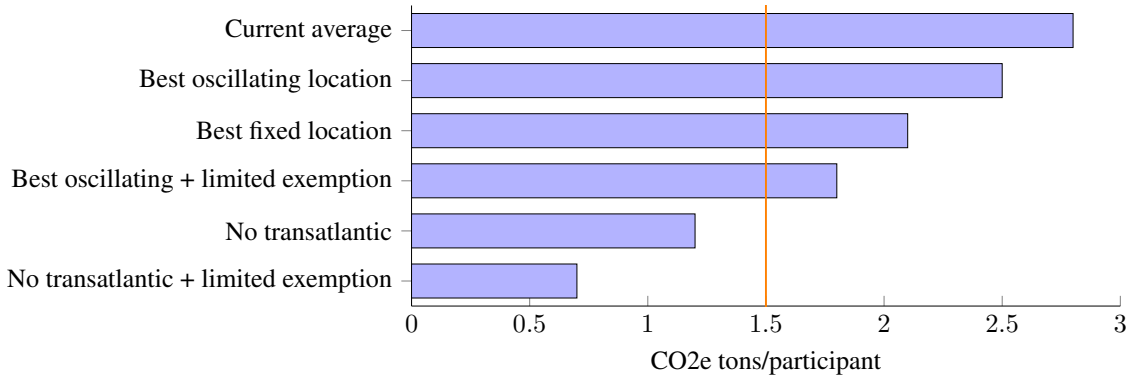

Figure 2: Predicted average CO2e emission per participant for different scenarios. The first bar corresponds to the current average (over 9 editions of PODC/DISC). *Oscillating* corresponds to a scenario where the conference takes place alternatively in North America and Europe. *Limited exemption* and *No transatlantic* refer to the two system changes listed.

2023, STACS, etc.) The effect of such actions is harder to measure, but by itself will not be enough, since yearly long-distance flights are simply too high in emissions.

- Having a hybrid set-up, with no constraint of attendance at all. In a hybrid set-up, interaction between on-site and online participants is facilitated (as opposed to only streaming the conference). We note that this option has many benefits, but requires more work from the local organizers. If for example the 50% furthest participants choose to attend online and we have best oscillating locations, the emissions are 0.35 CO2e tons/participants.
- Having a fully online set-up. Although there would still be some emissions, these are negligible compared to the other numbers considered.
- Switching to a journal-based system.

Opinion. On the one hand, there are benefits to the current conference-based system instead of a journal-based, as listed *e.g.* in [7]: it encourages good talks, it prevents excessive reviewing time, and it creates a rotation in the pool of reviewers. It also allows to better maintain a community, where people know and trust each other. In addition, it is often hazardous to drastically change a system very quickly: it could harm the community, and in particular some specific members, and it could also be that a fraction of the community just do not agree and start new conferences, spoiling the efforts of the others.

On the other hand, we cannot continue with the current system, and the changes have to be substantial and fast. For example, co-locations and extended stay support are not effective by themselves, when only one transcontinental conference trip a year is already too much. It is a real challenge for our community, and we do not claim that we know the perfect way to face it. In our opinion, the very first step is to set emission targets for the years to come, and adapt the conferences to meet these goals. Having some sort of task force dedicated to the topic seems reasonable, given the major changes required, and the multiple options to be explored. The task force can monitor that these changes do not harm our community in the long run, for example it can develop strategies to make sure that we do not split into Europe and North America parts that no longer interact.

As a matter of personal opinion, the authors advocate for the delayed presentation, allowing for online presentations, and streaming the conference. In particular, anyone could request for an online presentation and on a first-come-first-serve basis the first 10% of accepted papers can be granted. This would result in a significant decrease in carbon emissions ('no transatlantic + limited exemption' in Figure 1), while, we argue, minimally impacting the conference experience. Having the conference in two places simultaneously, or having a fully hybrid[3] conference (with say 50% online) are also good options, but require much more work from the organizers.

Since changes are hard to make, we want to highlight that they will have good impact on other aspects of our community. Many people cannot benefit from the current conference system, and suffer from it, for example: people who cannot get visas to the countries where conferences happen, people with limited funding, and people whose duties do not allow traveling far away (*e.g.* taking care of relatives). Having a larger portion of the event available without traveling will make a great difference for them.

Finally, in our opinion, scientists should take a leading role in battling the climate crisis, both in research and as role models. As these numbers show, attending multiple conferences a year puts the carbon emissions of a researcher already at worse-than-average, while this is not even including private travel. Action is necessary on every level (individual, group, department, university, etc.), in this column we focus on what conferences can (and should) do.

Recommendations for Steering Committees. Based on our findings and the discussion above, we have the following concrete recommendations for conference Steering Committees.

1. Formulate concrete goals; set milestones with concrete emission targets. For example, we can follow the Paris Agreement and set -50% for 2030. Having additional goals for the intermediate and following years would be helpful.

2. Experiment with the *next* edition: do not wait too long. For example, we can try limited exemptions for one year to see how it impacts the conference.

3. Take a data-driven approach, by appointing an environmental chair for each edition, whose tasks include gathering participant data, computing the estimated environmental impact of that year's edition, and exploring secondary goals, *e.g.* reducing waste, extended stay schemes, or vegetarian food.

4. Install a task force to oversee long term progress.

2 Methodology

This part has two purposes: explaining the methods to support our findings and providing some guidelines for people who want to do similar calculations (*e.g.* for another conference). Later in the document, we will justify our choices by a basic sensitivity analysis.

2.1 Data collection

Choice of scope. Our raw data are the lists of (affiliations of the) participants of past editions of DISC and PODC. Since it is a common belief that DISC and PODC have basically the same community, we looked for the data for the two conferences, and considered them as one.

[3] As opposed to streaming talks, in a fully hybrid conference interaction between online participants and with the onsite participants is facilitated.

The years for which we gathered participants lists are 2016, 2017, 2018, 2019, 2022 and 2023. Years before 2016 were not considered, since they might not be representative of the community today. We chose to exclude the data from 2020 and 2021, since the conference were fully or partially online these years, due to COVID. We expect that the participant lists of 2020-21 will contain people who would normally not attend these conferences, and are not part of the 'core community'. Of course, making the conferences available to these people as well is a separated discussion, beyond the scope of this report.

We note that PODC 2023 was part of FCRC, which most probably influences the set of participants. However, it is only one of our four PODC datasets, and PODC is planning to take part in FCRC every four years, so we thought it would be fair to include it.

We did not include DISC 2023 (the conference happened while writing this document). The registration list was significantly more European (incl. Israel) than in other years. However, due to current developments in the Middle East, many Israelis were not able to attend in the end. Since this covered about 10% of conference registrations, we decided to leave the data out as it does not seem to be representative for a typical DISC.

Gathering the data. We contacted the general/local chairs for all of these years for both PODC and DISC. We could get all the lists for DISC (except 2023), but only 2018, 2019, 2022 and 2023 for PODC. The chairs of the other years replied, but the data has been lost, or deleted for privacy reasons.

These are the locations of these conferences.

- DISC 2016, Paris.
- DISC 2017, Vienna.
- DISC 2018, New Orleans.
- DISC 2019, Budapest.
- DISC 2022, Augusta.
- PODC 2018, London.
- PODC 2019, Toronto.
- PODC 2022, Salerno.
- PODC 2023, Orlando.

2.2 Data preprocessing

Basic anonymization. A first step was to remove the participants names from the files, since we only needed the affiliations, and affiliations are better than names in terms of privacy.

Affiliation standardization. A second step was to standardize the affiliations. Many places have different names, and it makes later steps more complicated. For example, *CNRS, IRIF* and *Université Paris Cité* refer to the same place. For each standardized affiliation, we looked for the country and city.

Mitigating the bias of local participants. We noted that for some conference editions there were many participants from the host city, and for some a suspicious zero. We assume this corresponds to a local chair including the local organizers in registration or not. To make the impact of these differences smaller, we replaced the number of participants from the host city by the average over all years (including the year itself)

Clustering. For both computation and privacy purposes, we decided to group cities by areas. For each area, we chose a city, that would give its name to the area, and also be the city from which the distance estimates would be made. There are several constraints on this clustering:

- It should be fine-grained enough such that it does not induce too much error in the final carbon emissions estimates.
- It should be coarse enough to avoid privacy issues (recovering too easily the names of the participants from the areas). This allows to make the processed spreadsheet public.
- Ideally, the number of clusters is not too large, in order to save on distance computations (which requires a specific manual processing, see below).

In order to satisfy Item 1, the rule of thumb was that traveling inside the cluster would be much cheaper (in CO2) than between clusters, which typically means that people would not fly within a cluster. For a few participants, it was difficult to satisfy both Item 1 and 2. E.g. participants from Qatar or Iceland. We decided to remove them from the dataset. This corresponds to less than 1% of the participants.

In the end, our clusters correspond to the following 30 cities.

- Atlanta
- Bangalore
- Berlin
- Boston
- Calgary
- Chicago
- Dallas
- Denver
- Frankfurt
- Jerusalem
- London
- Madrid
- New York
- Nicosia
- Paris
- Phoenix
- Pittsburgh
- Rio de Janeiro
- Rome
- San Francisco
- Seoul
- Shanghai
- Singapore
- Stockholm
- Sydney
- Tokyo
- Toronto
- Vienna
- Zurich

2.3 Typical participants list

To evaluate emissions for future locations, we need a virtual participants list. An issue here is that depending on the locations, this list has to be different, because the participant distribution is quite different, see Table 1.

Participants origins	**Participants in average**	**Participants if in Europe**	**Participants if in North America**
Europe + Israel	54	60	45
Canada + USA	34	26	46
Middle and South America	1	1	1
Asia and Oceania	11	13	8

Table 1: Percentages of participants per origins, for conference location on various continents. We round the percentage for a clear overview, hence they no longer add up to exactly 100%.

We used three populations: one for North America locations, one for Europe locations, and one for other locations. For the two first ones, we used the sums of the participants for the past conferences in the relevant area. For the third, since we miss data, we just use the sums of all participants from our data. Note that since our end results are emissions per participant, it is not necessary to scale down these populations.

2.4 Choice of cities for conference locations

Our goal is to evaluate how different carbon emissions are for various conference locations. We chose to evaluate the locations of the conferences of our data set, both for validating some of our simplifications (see later) and because conferences tend to return to close places (because this is where there is higher density of researchers from the community).

We also evaluated the locations for next years, either already decided or with current bids. Finally, we added Tokyo, San Francisco, and Jerusalem, since they are in areas that are popular for conferences, and do not appear yet in our list of locations.

The final list of 14 target cities is the following.

- Paris
- Vienna
- New Orleans
- Augusta
- London
- Toronto
- Rome (Salerno)
- Orlando
- Mexico City
- Madrid
- Tokyo
- San Francisco
- Berlin
- Jerusalem

2.5 Evaluation of CO2e emissions of traveling between two cities

What we count in this document are *equivalent CO2 masses* (CO2e), and do not correspond to actual CO2 emissions. The reason is that plane travel, in addition to its raw effects of releasing CO2, also contributes to global warming by releasing other substances including nitrogen oxides and water vapor, and releasing them at high altitudes [13]. For some details, see for example: 'Environmental effects of aviation' on wikipedia [18]. For this reason, the environmental footprint of plane trips is usually expressed as an equivalent mass of conventional CO2 emissions.

We computed the estimated CO2e emissions of traveling from any of the 30 source cities to any of the 14 target cities, and back. This required some choices, that we believe do not impact the results in an important way. Below is the method we use.

- If there exists a train connection that would take less than 8 hours, we consider that the travel will be by train, otherwise by plane. (We made exceptions for Vienna-Berlin, because there is a convenient, direct 8h20 train, and Pittsburgh-Toronto that is 5 hours by car).

- For train the CO2e emissions were approximated by: (1) evaluating the distance, taking Google maps distance by car (back and forth) and (2) multiplying by the average CO2e emissions by kilometer for trains, for which we used 0.03 kg. In reality high speed train are mostly in the 0.006-0.012 regime, local train can reach up to 0.06. For details on this see for example [1, 6].

- For planes, we use atmostfair's calculator [4] roundtrip estimates (for an average airline). This includes the indirect CO2 effects not just as a multiplicative factor 2 (an approximation recommended by Jungbluth and Meili [10] as a first order estimate), but makes this factor dependent of the flight altitude (as recommended as a next order accuracy therein). Additionally, atmosfair's calculator uses a database which analyzed the aircraft types, their fuel consumption and passenger loads typically flown on specific routes. We assume economy class for every participant. Note that airlines or ticket retailers such as Google flights often presents lower numbers, not taking all the indirect effect into account.

2.6 Further reading on conference emission calculations.

- Bousema et al. [8] calculated the emissions of the 2019 edition of the annual conference of the American Society of Tropical Medicine and Hygiene.
- Stroud et al. [15] calculated the emissions for four editions of the conference of the International Biogeography Society.
- Spinellis and Louridas [14] calculated the total emissions of scientific conferences, using the Scopus digital library.
- Leochico et al. [11] survey the climate impact of academic conferences and general alternatives.

3 Detailed Results and Sensitivity Analysis

3.1 Results on conference locations

Let us restate our main results, the average CO2e emissions per participant for the 14 target cities. We estimate our margin of error to be below 10%, see the next section.

Conference location	CO2e per participant (tons)
Augusta	3.4
Berlin	2.4
Jerusalem	3.3
London	2.3
Madrid	2.6
Mexico City	4.6
New Orleans	3.8
Orlando	3.7
Paris	2.1
Rome	2.5
San Francisco	4.3
Tokyo	5.7
Toronto	2.9
Vienna	2.4

The way we interpret the fact that Europe location are in average better than others is that: (1) the community has a large portion of people in Europe (see Table 1), (2) travel within Europe can more easily be done by train, and (3) it is geographically more central (just as a reminder, the Pacific Ocean is much bigger than the Atlantic, *e.g.* Tokyo-San Francisco is > 8000 km, while Paris-Toronto is 6000 km).

3.2 Results on alternatives

In this section, we discuss possible measures and their effects on the carbon emissions.

Limited exemption. A significant portion of the CO2e emissions come from few participants that are further away. A fully hybrid conference, with say 50% online participants, will have a much much lower environmental impact (see below). However, several senior community members have expressed that they oppose such fully hybrid conferences, because of organization overhead and degraded conference atmosphere. Instead, in this section, we explore what happens with a light hybrid set-up when participants who live far away from the conference location are allowed to attend/present online (just like people who could not get visas, for example). As an example, we let the furthest 10% of participants attend online[4]. This would have minimal impact on the feeling of the conference, while having significant impact on the CO2 emissions. Below, we average the CO2e emissions again per participant (including the online participants). Alternatively, this 10% can also publish in the proceedings without presentation. Of course the CO2e emissions for both policies are the same.

Of course we do not argue from *excluding* the furthest 10%; online presentation should always be voluntary. This can for example be done in a first-come-first-serve manner, where everyone who has to take a long-distance flight can apply for an online presentation and the first 10% of accepted papers gets approval.

Conference location	**CO2e per participant (tons) with 10% online participants**	**change**	**relative change**
Augusta	2.6	0.8	-26%
Berlin	1.7	0.7	-27%
Jerusalem	2.5	0.8	-25%
London	1.6	0.7	-30%
Madrid	1.9	0.7	-30%
Mexico City	3.6	1.0	-22%
New Orleans	2.9	0.9	-24%
Orlando	2.7	1.0	-26%
Paris	1.5	0.6	-30%
Rome	1.8	0.7	-28%
San Francisco	3.4	0.9	-21%
Tokyo	4.9	0.8	-14%
Toronto	2.1	0.8	-27%
Vienna	1.7	0.7	-28%

[4]Of course, by letting the average 10% attend online, instead of the 10% furthest, the cost would go down by 10%.

Hybrid with 50% online participants. We computed the CO2e emissions for hosting a conference in Paris or Toronto with the furthest 50% of the participants attending online.

- Paris: 0.2 t CO2e/participant, i.e., -91%.
- Toronto: 0.5 t CO2e/participant, i.e., -84%.

No transatlantic flights. The idea is to make sure no participants need to take long flights. Since most of our participants are located in North America and Europe, we particularly try to avoid transatlantic flights. This can be realized in two ways: (1) allow authors to present their paper in the next x years after publication, or (2) organize the conference simultaneously in two locations, with an online connection between the two. As an example, we present here the case of hosting in Paris and Toronto (either simultaneously or in consecutive years). This would decrease the carbon emissions to

- 1.2 t CO2e/participant.

Another option is to combine this with limited exemption. As before, we assume 10% online participation from the furthest participants. In our data set that means roughly that people from East Asia, Oceania, and South America attend online, while people from Europe and North America do not have to fly across the Atlantic. In this case, the carbon emissions decrease further to

- 0.7 t CO2e/participant.

3.3 Sensitivity analysis

Typical participants lists versus real lists. For the conferences of which we received the participant data, we also calculated the CO2e emissions with the real participants list (after clustering) of that year. The difference between those outcomes and the typical participant lists is small: about 0-10%.

Removed participants. Since we removed very few participants ($\leq$1 %) this has a negligible effect.

Clustering. We aimed at creating clusters where, within a cluster, people take the train or other local low-emission transport. These emissions would have negligible impact when traveling to another cluster. This was relatively straightforward in Europe, since the density of participants is high there. Outside of Europe, we let the radius of the cluster grow until it contains a non-trivial amount of people. E.g., the participants from Los Angeles and San Francisco are clustered in San Francisco, although this trip would, according to our assumptions, always be taken by plane. We argue that this is not a big issue, since it concerns a way smaller number of participants. For comparison: we have an average of 15 participants per year clustered in Paris, and 4 in San Francisco.

In other words, our clustering is such that

$$\frac{\text{(additional local travel)} \cdot \text{(number of local participants)}}{\text{total emissions}} = \text{small}.$$

Ignoring 'the last mile'. In our calculations, we often ignored the last part of the journey, near the conference location. This means for example that we treat both Salerno and L'Aquila as 'Rome'. Rome is the closest airport to L'Aquila, and one of the closest airports to Salerno. The additional train/bus journey from there is negligible compared to any flight. Similarly, we treated Budapest as Vienna. These cities are a two-hour train ride apart, which again is negligible compared to any flights.

Train-plane cut-off. We assumed that participants would take the plane if a train ride would be more than 8 hours. This number is an educated guess based on conversations with community members. To see if it is in the right ballpark, we surveyed the DISC '23 participants. We received 36 answers, with a median of 10 hours and a average of 11 hours. Since the number of answers in very low, and perhaps biased by the pool of surveyed people, we are hesitant to draw any significant conclusions from this. We do want to conclude that 8 hours seems to be the right ballpark. We estimate that changing the number 8 slightly up or down by 2 hours will have an impact of less than 5%.

Acknowledgments

We would like to thank all the conference chairs for providing us with the data, and our colleagues who gave feedback on earlier versions of this report, and in general people who have been supporting this initiative.

References

[1] International Energy Agency. Well-to-wheel (wake/wing) ghg intensity of motorised passenger transport modes. `https://www.iea.org/data-and-statistics/charts/well-to-wheel-wake-wing-ghg-intensity-of-motorised-passenger-transport-modes-2`.

[2] M Allen, OP Dube, W Solecki, F Aragón-Durand, W Cramer, S Humphreys, M Kainuma, et al. Special report: Global warming of 1.5 c. *Intergovernmental Panel on Climate Change (IPCC)*, 2018.

[3] Antoine Amirilli. Should conferences still require mandatory attendance? *Bull. EATCS*, 139:30–34, 2023.

[4] Atmosfair. Flight emissions calculator. `https://www.atmosfair.de/en/offset/`.

[5] Atmosfair. Paris aligned annual carbon budget. `https://www.atmosfair.de/en/green_travel/annual_climate_budget/`.

[6] Deutsche Bahn. Facts and figures on environmental protection in 2016. `https://www.deutschebahn.com/resource/blob/6925828/f3c7c9bc0d0766230e64d0464f299b56/report_environmental-data.pdf`, 2017.

[7] Boaz Barak. Computer science should stay young. *Commun. ACM*, 59(6):37–38, 2016.

[8] Teun Bousema, Prashanth Selvaraj, Abdoulaye A Djimde, Derya Yakar, Brittany Hagedorn, Abigail Pratt, Didier Barret, Kate Whitfield, and Justin M Cohen. Reducing the carbon footprint of academic conferences: the example of the american society of tropical medicine and hygiene. *The American Journal of Tropical Medicine and Hygiene*, 103(5):1758, 2020.

[9] Patrick Greenfield. Carbon credit speculators could lose billions as offsets deemed 'worthless'. *The Guardian*, 2023.

[10] Niels Jungbluth and Christoph Meili. Recommendations for calculation of the global warming potential of aviation including the radiative forcing index. *The International Journal of Life Cycle Assessment*, 24:404–411, 2019.

[11] Carl Froilan D Leochico, Melina Longoni Di Giusto, and Ramiro Mitre. Impact of scientific conferences on climate change and how to make them eco-friendly and inclusive: A scoping review. *The Journal of Climate Change and Health*, 4:100042, 2021.

[12] Isabella O'Malley. Scientists confirm global floods and droughts worsened by climate change. https://www.pbs.org/newshour/science/scientists-confirm-global-floods-and-droughts-worsened-by-climate-change.

[13] Joyce E Penner, David Lister, David J Griggs, David J Dokken, and Mack McFarland. *Aviation and the global atmosphere: a special report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 1999.

[14] Diomidis Spinellis and Panos Louridas. The carbon footprint of conference papers. *PloS one*, 8(6):e66508, 2013.

[15] James T Stroud and Kenneth J Feeley. Responsible academia: optimizing conference locations to minimize greenhouse gas emissions. *Ecography*, 38(4):402–404, 2015.

[16] UNFCCC. The paris agreement. https://unfccc.int/process-and-meetings/the-paris-agreement.

[17] Moshe Y. Vardi. Publish *and* perish. *Commun. ACM*, 63(1):7, 2020.

[18] Wikipedia. Environmental effects of aviation. https://en.wikipedia.org/wiki/Environmental_effects_of_aviation.

[19] Greg Wood. Next steps towards a net zero ietf. https://www.ietf.org/blog/towards-a-net-zero-ietf-next-steps/.