A Proposal for Reorganizing the IACR Publications Landscape

1 Introduction

The IACR has seen a lot of growth since its foundation in 1982: the community roughly doubles in size every 15 years. Until 2005, this growth has been possible by regularly creating new venues (or integrating venues into the IACR) after the first CRYPTO conference in 1981: EUROCRYPT in 1984, the JoC in 1988, ASIACRYPT/AUSCRYPT in 1990, FSE in 1993, PKC in 1998, CHES in 1999, TCC in 2004. In the last 10 years, growth has been possible by having each conference accept more papers, with various adjustment to the schedule (shorter talks, more conference days, parallel tracks, ...).

Today, we have a fragmented landscape with eight different IACR publication venues. This is complex for authors having to choose where to submit their papers, and the conference model requires them to travel to present their results. The system with many independent venues also creates a large workload for reviewers serving in the program committees (in particular, resubmitted papers are treated as new submissions). Reviewing for conferences has other drawbacks: a bias towards some areas of cryptography, not enough space to publish full proofs, not enough time to review long papers with proofs.

We believe it is time to have a discussion about the IACR publication model, to rethink and rationalize it. Recent developments such as the move to open access journal/conference hybrids (ToSC and TCHES), and the switch to virtual or physical/virtual hybrid conferences during the covid pandemic open the way for better publication models. We must take into account the growth of the community, but we should avoid the creation of new outlets that make the system more complex without solving the current issues.

Vision. In this note, we propose to reorganize IACR publications as journal/conference hybrids with two well-identified tiers:

- A single flagship journal accepting contributions from all areas of cryptology, whose papers are invited to be presented at CRYPTO, EUROCRYPT and ASIACRYPT;
- A series of *IACR Transactions*, with a corresponding area conference, managed by each community (ToSC, TCHES, TTCC, TPKC, TRWC);

The move to journal hybrids was discussed within the IACR board in 2013[1], inspired by PVLBD; it lead to a straw-man proposition by Nigel Smart[2]. Our proposal is modelled on the *American Physical Society* system with a flagship journal (*Physical Review Letters*), and several area journals (*Physical Review A, B, C, D, E*). It clarifies the relation between the different venues, and provides the benefits of journal publication to all IACR members, with higher-quality reviews, better indexing, and open access publishing.

A journal also offers flexibility to decouple paper publication from conference presentation. We could make presentation of accepted papers optional to offer the same publication opportunities to authors who cannot or do not want to travel, and to leave some room for future growth of the community without over-packing the conferences.

[1] https://iacr.org/news/item/2684
[2] https://www.iacr.org/news/files/2013-08-05PublicationReform.pdf

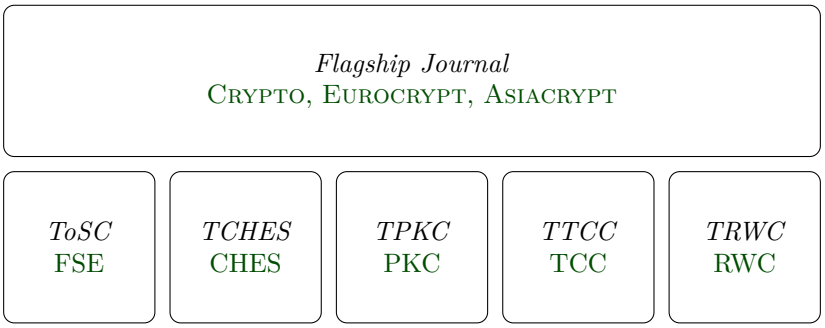

Figure 1: Proposed publication landscape

Outline. We first review the current IACR venues, and point out several shortcomings felt by various members of the community (Section 2). Then we introduce a series of proposals to improve the current system, and to reorganize it around journal/conference hybrids.

We start with short-term tweaks (Section 3), that can solve some issues of the current venues, such as guidelines to aim for a better representation of all areas of the community, a move to diamond open access, or a long-term commitment to physical/virtual hybrid conferences for researchers who do not travel.

Then, we propose measures to implement a journal/conference hybrid model (Section 4). This requires deeper changes to the system, but is necessary to simplify the publication landscape, improve the reviewing system, and decouple papers publication from conference presentation.

This note does not aim to provide a fully specified solution, but rather to initiate a discussion around possible improvements. Each proposed measure has a clearly defined goal, and a minimal impact outside this goal. We suggest having a discussion inside the community, with polling to identify the main issues, and which changes would be acceptable for most members. Then we can implement measures to fix the issues that are felt as important. Since the measures are modular, they can be implemented gradually.

2 Current Situation and Perceived Problems

The IACR runs ten different venues:

- **The Journal of Cryptology** (JoC)
- **3 flagship conferences**: Eurocrypt, Crypto, Asiacrypt
- **4 area conferences**: CHES, FSE, PKC, TCC
  - Two of them are journal/conference hybrids: CHES papers are published in *IACR Transactions on Cryptographic Hardware and Embedded Systems* (TCHES) while FSE papers are published in *IACR Transactions on Symmetric Cryptology* (ToSC)
  - Two of them publish proceedings with Springer: PKC and TCC
- **A symposium**: RWC (without proceedings)
- **The ePrint** archive for preprints

In addition, many smaller events are organized "in cooperation with IACR" and are strongly related to our community.

Positioning. It is generally accepted that the flagship conferences publish the best papers in cryptography in general, while the area conferences publish strong papers in their respective areas. Papers that are rejected from the area conferences are typically published by less prestigious events held in cooperation with IACR (such as SAC, PQCrypto, SCN, CANS, Indocrypt, Latincrypt, Africacrypt, ...), or outside IACR either in conferences (such as CT-RSA, CARDIS, ProvSec, ACNS, ...), or in journals (such as DCC, IJACT, ...). There is no strict ordering of venues (the best papers in area conferences are often above the bar of acceptance for flagship conferences), but there is a clear difference between papers submitted at Crypto and papers submitted at Latincrypt.

The Journal of Cryptology publishes long versions of some conference papers (in particular, the best three papers in each conference are invited to submit to the JoC), as well as original submissions (with quality requirements similar to the flagship conferences).

Some excellent cryptography papers are also published in flagship venues outside the IACR, such as IEEE S&P, USENIX Security, or ACM CCS (for more practical results), and FOCS, STOC, or IEEE TIT (for more theoretical contributions).

Publication. Most of the IACR conferences publish papers as proceedings with Springer[3]. Springer sells access to the papers, and pays a fee to the IACR. Authors are encouraged to upload preprints in the ePrint archive (Green Open Access). IACR members also have access to the Springer version for free. Moreover, the IACR publishes the author version on the IACR archive after two years, and the Springer version becomes available for free one year later.

On the other hand, the *Transactions* are published by Ruhr University Bochum and the publisher version is accessible freely from Day one under a CC-BY licence (Diamond Open Access).

RWC is a symposium and does not publish papers (talks that present a new research finding are typically published in one of the flagship conferences, or outside the IACR). The ePrint archive is not considered as a formal publication.

We now review some problems of the current system, and desirable goals for potential improvements.

2.1 Limitations of Conference Reviewing

The conference model implies some limitations of the reviewing system:

- Papers can either be accepted as they are (potentially with shepherding to fix small issues), or rejected and later submitted to a different conference.

  This creates a heavy review load, because many papers are submitted several times to the flagship conferences, in the hope of eventually being accepted (*e.g.* rejected from Crypto, re-submitted to the following Asiacrypt and rejected, re-submitted to the following Eurocrypt and finally accepted). Since each conference has its own program committee, this is treated as three independent submissions, and each one has to be evaluated by three reviewers. This is a waste of reviewers time and author time. The introduction of "sticky reviews" in 2014 was an attempt to mitigate this, but it doesn't seem to be used a lot.

- Since there is a limited number of slots, papers are competing against each others, and this tends to make reviews more adversarial than in a journal (where papers compete against a perceived quality threshold). A good paper can also be rejected just because there are more good submissions than usual.

[3] https://iacr.org/docs/pub_2013-16.html

- Due to the strict deadlines of conferences, papers are limited in lengths, and technical elements (*e.g.* proofs) are sometimes omitted from the reviewed and published paper. Having both a short and long paper with the same scientific result makes the scientific record more complex, and is not ideal for citations. In some cases, the proof is never published.

2.2 Travel Requirements

The conference model requires authors to attend the conference in order to have a publication. This is problematic because some authors cannot or do not want to travel, for a number of different reasons (limited budget, visa issues, environmental impact, family obligations, teaching duties, ...). In particular, conference travel is a significant fraction of the environmental impact of computer science research.

2.3 Impact of Conference Publications

Conference publications are sometimes considered as inferior to journal publication by evaluation and hiring committees. This is especially true in mathematics, and this is a problem as cryptology research is somehow in between computer science and mathematics. For instance, citation indexing databases such as the Web of Science and Scopus are primarily focused on journal publications. Scopus includes LNCS proceedings, but considers the whole series as a single entity.

2.4 Issues with the *Journal of Cryptology*

Being a journal, the JoC does not suffer the issues mentioned above. However, it is not an attractive venue for publication because of its long review time. The situation has improved in the last years (after significant efforts from the editors and switching to an online submission system), but median time between submission and acceptance was still of 16 months for papers published in 2021.

As a result, the JoC publishes few original research results. Efforts have also been made to improve this situation with special issues (called "topical collection"), but with limited results. In 2021 the JoC published 42 papers, including 5 in a special issue on TLS 1.3 and 3 in a special issue on CAESAR. Two thirds of those 42 papers are long versions of papers previously published in a conference (inside or outside IACR).

2.5 Growth of the Community

Currently, the flagship conferences publish about 80–100 papers per year each, while the area conferences publish around 40–60 papers per year each. The Journal of Cryptology publishes less than 50 papers per year. Looking at the past 20 years, the number of papers published by the IACR follows an exponential growth, doubling roughly every 15 years (Figure 3 page 12). The number of submitted papers to IACR conference also doubles every 15 years, and other key indicators grow at a similar rate (Figure 2 page 12).

This can be an issue if the community grows faster than the number of papers published, leading to more selective conferences and more difficulties for authors to publish their results. Looking at the acceptance rate of IACR events (Figure 4 page 13), this was a problem between 2004 and 2012, with the acceptance rate of flagship conferences dropping below 20%. However, the number of papers accepted has increased after this period, and the acceptance rate in flagship conferences reaches an average of 24% since the switch to parallel tracks (2015–2021). We need to anticipate future growth, but there does not seem to be a high pressure currently[4].

[4]During the covid pandemic, some virtual conferences (Crypto 2021, Asiacrypt 2021, Crypto 2022)

2.6 Respecting All Areas of Cryptography

Some areas of cryptography are more represented than others in the flagship conferences (in general, there seems to be a bias towards theoretical papers). In particular, CHES is the largest area conference (with more than 400 attendees regularly), but very few papers from this area are accepted in the flagship conferences. RWC was also created because of a perceived difficulty to present applied cryptography papers in the flagship conferences.

More generally, members of many areas also express a feeling that their area is under-represented in the flagship conferences (*e.g.* symmetric cryptography, code-based cryptography, quantum cryptography, ...). Some areas also fall out of the scope of the four area conferences, and do not have a fallback venue within the IACR landscape.

2.7 Open Access Publication

There is a growing trend to require scientific publications funded by public money to be freely available to the public and to other researchers. For instance, publications resulting from an ERC grant funded under a Horizon Europe ERC Work Programme must be available under a CC-BY licence or equivalent at the time of publication (author version). In the current systems, publishers like Springer and Elsevier extract large profits from the research community but provide relatively few services.

Currently, IACR proceedings published by Springer are not freely available at the time of publication. Only the author version is available freely (typically on ePrint); this is referred to as Green Open Access. A move to Gold or Diamond Open Access would make the publisher version available freely from day one (this is the case for ToSC and TCHES).

2.8 Minimize Impact on Existing IACR Venues

While there are clearly identified problems with the current venues, the system is nonetheless quite successful. Therefore, any proposed change should carefully evaluate the impact of the change on the larger IACR publication landscape and weight the advantages and the risks. In particular, specific measures should be taken to limit the impact of a given change on venues that are not affected by the particular problem it is supposed to address. More generally, we should minimize the impact on the larger cryptology research ecosystem, including events that are run "in cooperation with IACR" and non-IACR events.

We now explore several solutions to solve the identified problem. We first propose simple tweaks that can implemented within the current system, then we propose deeper changes that could more significantly improve the current situation.

3 Short-term tweaks

Some problems can be solved with minor changes that only address one specific problem with limited side effects. If the IACR community agrees that the underlying problems must be solved, we suggest adopting the following measures.

Proposal 1: Limiting travel

Keep the physical/virtual hybrid model for IACR conferences in the long run:

a) Continue to stream conferences online for remote attendees.

have taken advantage of being organized online to noticeably increase the number of accepted papers. This did not seem to lower the quality of accepted papers, and could be a sign that there is some pressure.

b) Allow remote presentations in IACR conferences.

The experience of virtual and hybrid conferences shows that it is possible to have a conference with some remote attendees and/or presenters. If there is a consensus that we want to make conferences more accessible to members of the community that don't travel, we could keep this hybrid model (either only for attendees, or both for attendees and speakers).
Listening to a remote talk is a degraded experience for physical attendees, but most authors prefer to present their papers in person. We can use the recent hybrid conferences as a benchmark: EUROCRYPT 2022 had 62 papers presented in person, and 23 presented online, even though intercontinental travel was still restricted at the time.

Proposal 2: Reviewing guidelines

Create IACR policies to improve the reviewing in all IACR venues.
We suggest guidelines to encourage program committees to:

a) respect all areas in flagship conferences

b) focus reviews on positive aspects of papers rather than negative aspects

c) expand the scope of each area conference to avoid gaps in between

We could try to have a better representation of the many areas of cryptography in the flagship conferences by asking program chairs of the flagship conferences to take proactive measure to ensure that papers from different communities are treated fairly. Some program chairs already make significant efforts to enforce a good diversity of papers, but this is not always the case. The program chair report[a] from EUROCRYPT 2021 is a good example of what can be done.
We could also ask the committees of the conferences to enlarge the scope of the respective venues. In particular, we could ask that papers that are between the topics of two area conferences (*e.g.* FSE and CHES) are considered to be in-scope of both conferences rather than being in-scope for neither.

[a] https://iacr.org/docs/EC21_report.pdf

Proposal 3: Improve the *Journal of cryptology* review time

Reduce the review time of the Journal of Cryptology with a larger editorial board and strict reviewing deadlines.

The JoC still has a long publication delay, resulting in few original papers being submitted. If we want to offer good options for journal publication within the IACR landscape, the most natural way forward is to improve this situation.
This is not an easy problem to solve, but maybe it can be improved with a larger editorial board and more strict deadlines for reviewers. A possibility would be to move to a conference-like schedule, with a submission dead-line, and a pre-defined period for reviewing.

Proposal 4: Open access

Move all IACR publications to Diamond Open Access, with the underlying costs taken in charge by the IACR.

The move to open access publication is conceptually simple. We just have to pay fees to Springer[a] or to move to a different publisher. If there is a consensus that IACR members want Diamond Open Access, we should consider such a move.
Diamond Open Access with a commercial publisher might be too expensive for the IACR, but another option is to move to self-publication or to an academic publisher. This is doable because the authors already do most of the publishing work (by providing "camera ready" versions). This work that was already done for ToSC and TCHES (published by Ruhr University Bochum), and the work done by the "New Journal" committee could be used as a basis to select a new publisher. As a rough estimate, if each paper requires one hour of work for finalization, this would represent about 500 hours of work per year and could be handled by a (part-time) IACR employee. Moving to an academic publisher or self-publication also gives us more control of the format of published papers. In particular we can better match what is reviewed and what is published.
There are some drawbacks of changing publisher and leaving Springer, in particular in terms of indexing. We discuss them in Section 4.1 and we estimate that they are limited (in particular, we don't think that the recognition of the IACR conferences is linked to the publisher).

[a] 30€ per page if the full proceeding is open access: https://www.springer.com/gp/computer-science/lncs/open-access-publishing-in-computer-proceedings

Proposal 5: Review forwarding

When a paper is resubmitted to an IACR venue after rejection from an IACR venue, forward the reviews to the new program committee.

Since 2014, the IACR encourages authors to use "sticky reviews", *i.e.* to include reviews from the previous conference when a paper is resubmitted. However, this does not seem widely used by authors.
Having a way to forward directly the reviews (and discussion) to the new program committee would reduce the review load caused by resubmissions, and improve the consistency of decisions. In order to avoid bias, the reviews should be forwarded during the discussion phase, after each reviewer has independently evaluated the paper.
To limit the number of reviews, we could start with only two reviews for resubmitted papers. Based on the two new reviews and the forwarded reviews, we would decide to either ask for a third review, or to directly reject the paper.

Proposal 6: Scaling conferences to follow growth of the community

Increase the number of papers published in the flagship conferences to roughly 100.

Conferences organized during the covid pandemic have accepted more papers than previously (up to 103 for Crypto 2021), using a format with shorter talks. This did not seem to strongly reduce the quality of accepted papers.

There is a limit to how much we can scale the current venues, but there are some possibilities to modify the organization of the physical events to accommodate more papers. For instance, we could use the "short talks" format that was used for virtual conferences and some hybrid conferences (*e.g.* EUROCRYPT 2021, TCC 2021, FSE 2021), where accepted papers have a 10 minutes slot at the conference, and provide a video of a longer presentation. The short talk format could be used for all papers, or only for papers presented remotely in physical/virtual hybrid conferences.
Reducing the length of conference talks has drawbacks but it is worth discussing further and could be implemented if the community considers that it is worth it.

4 Moving IACR Conferences to Journal Hybrids

Some problems are inherent to the conference model, in particular the issues related to reviewing. In order to address them, we propose to move all IACR conferences to the conference/journal hybrid model, following the move of ToSC and TCHES in the last years. Many computer science conferences had made a similar move, starting with PVLDB in 2008[5] and PoPETS in 2013[6]. Moving a conference to a journal hybrid has important benefits:

- Several deadlines per year. This allows authors to choose where to submit independently of calendar issues.
- More interactions between authors and reviewers (minor and major revisions). Authors and reviewers can work together to improve the quality of a paper.
- Papers being judged against a perceived quality level rather than competing against each other.
- Possibility of longer review time for longer papers.
- Better indexing and recognition of journal publications.

Proposal 7: Generalize the *Transactions* model

Encourage the TCC and PKC steering committees to move to the *Transactions* model.
Encourage the RWC steering committee to create a *Transactions on Real World Crypto*.

FSE switched to the *Transactions* model in 2016 with ToSC, and CHES switched in 2018 with the creation of TCHES. Many other conferences in computer security have made a similar move, and communities generally consider it as a success: USENIX Security, ACM CCS, IEEE S&P, NDSS, ...
Since the choice of moving to a hybrid model for FSE was discussed at length within the IACR board, maybe it's now time to discuss the choice of TCC and PKC to keep the conference model while everybody is moving to the hybrid model.

The transaction model lets each sub-community control its publication outlet. In particular, they can decide how many papers to accept, and how to articulate publication in the journal and presentation at the conference. Presentation can be mandatory (as it is

[5] https://www.vldb.org/pvldb/
[6] https://petsymposium.org/popets/

generally now), optional (as it was for the ToSC special issue on the NIST lightweight standardisation process), or even restricted to a subset of published papers invited by the program committee (this would be a natural choice for RWC to keep the event as a symposium with most of the talks not linked to a publication).

Proposal 8: Create a single journal for the flagship conferences

Create a flagship journal, replacing the proceedings of Crypto, Eurocrypt and Asiacrypt.
Authors would submit to a single journal, and a rule would decide which papers are presented in each conference; we suggest several options:

a) Authors of accepted papers choose where they want to present;

b) Papers are presented at the first conference following acceptation;

c) Papers are presented at the conference corresponding to the region of authors;

d) Local conference committees invite papers from the journal for presentation.

Merging the three flagship conferences is the only way to avoid the review load caused by iterative resubmissions: many papers are submitted repeatedly to the flagship conferences until they are accepted or authors decide to try a less prestigious venue. With a single journal, resubmissions are turned into major revisions, which reduces the review load, and improves the feedback between authors and reviewers, typically leading to higher quality papers.
Papers will not be rejected because of a lack of slots, but because they don't have a sufficient quality level. Therefore, re-submission of rejected papers will be marginal. Papers that show good potential but are not yet suitable for publication (because of editorial or technical issues) will be asked to do a major revision and the revision will be reviewed by the same reviewers as the initial submission.

Concrete details of the flagship journal should be discussed further if the IACR community wants to go in this direction, but we make a tentative concrete proposition as a starting point in Appendix B page 14.

The journal/conference hybrid model offers more flexibility to decouple paper publication from conference presentation.

Proposal 9: Optional presentation

Make conference presentation optional for papers published in journal hybrids.

Making presentation optional can solve two problems: it permits authors to publish without travelling (2.2), and gives some room for growth by having more papers published than presented (2.5).
Making presentation optional will reduce attendance to conferences, but probably in a rather limited way; the impact will likely be smaller than the impact of physical/virtual hybrid conference (see Proposal 1). For instance, at Eurocrypt 2022 73% of the papers have been presented in person.

4.1 Discussion

Moving to a journal/conference hybrid model provides the benefits of journal publication for all papers published by the IACR, and solves the majority of the issues discussed in Section 2. If there is a sub-community that does not feel represented by the existing area conferences, it is possible to create additional *Transactions* for the corresponding field (we must also make sure that the flagship conferences include all areas as discussed in Proposal 2).

Having a separate journal for each area and a common journal for the flagship conferences creates journals with a clear ranking inherited from the existing conferences, giving authors a venue of known quality and known ranking.

Each journal would aim to have the same level of selectivity as the conferences it is replacing. However, they wouldn't have a fixed number of slots, but only a perceived quality threshold for accepting papers. Each paper is evaluated against this bar, rather than being compared to other submissions.

Conference size. With a journal/conference hybrid, it is hard to predict in advance how many papers will be presented at the conference (a traditional conference typically fixes the number of slots in advance). The general chair has to deal with this uncertainty and adapt the conference if needed. However, our experience with ToSC and TCHES shows that it is manageable.

For the flagship journal, there is further uncertainty about where each accepted paper will be presented. If this is an issue, a strict rule can be implemented to limit author's choice.

Change of publisher. Moving our publications out of Springer has some negative effects.

First we will lose the current indexing. This could reduce the status of the conferences, but we estimate the risk to be low because the current indexing of the LNCS series is low, and the recognition of IACR conferences is most likely attached to the conference name rather than to the publisher. After a few years, the indexing will probably be better than it is currently, as was the case with ToSC and TCHES. Indeed, ToSC[7] and TCHES[8] are now indexed by Scopus, with a citescore of 4.8 and 9.0 respectively (for comparison, the JoC[9] has a citescore of 4.1, and LNCS proceedings have a citescore of 2.1[10]). ToSC is also indexed in the Clarivate Web of Science[11], but does not yet have an impact factor. The CORE conference rankings[12] have not been affected by the move to a journal/conference hybrid: FSE is still ranked B, and CHES was promoted from C to A in 2018. Google Scholar includes both journals and conference proceedings in the same ranking[13]. However, when ToSC and TCHES moved to a journal hybrid, Google Scholar did not consider the conference proceedings and the new journal to be the same entity. Therefore, during a transition period the journals had no good papers indexed and were ranked poorly (Google Scholar uses papers in the last five years for ranking).

The move could also impact the ranking of the JoC, because citations from conferences proceedings will not be included during the transition period when the hybrid journals are not indexed.

[7] https://www.scopus.com/sourceid/21100898676
[8] https://www.scopus.com/sourceid/21101064700
[9] https://www.scopus.com/sourceid/21101064700
[10] https://www.scopus.com/sourceid/25674
[11] https://mjl.clarivate.com:/search-results?issn=2519-173X
[12] http://portal.core.edu.au/conf-ranks
[13] https://scholar.google.com/citations?view_op=top_venues&hl=en&vq=eng_computersecuritycryptography

Second, if we don't renew the IACR contract with Springer, the Springer version of papers publisher before 2013 will no longer be freely readable (but the IACR version will still be available on the IACR archive).

The *Journal of Cryptology*. The new flagship journal will fill out most of the current roles of the JoC (assuming we also implement Proposal 1 or Proposal 9). Authors that want a journal publication for better indexing or because they don't travel will be able to use the flagship journal, and journal versions of conference papers will no longer be required. Therefore, the JoC should probably be retired, or be re-purposed as the flagship journal.

Authorship

This document was written by Gaëtan Leurent.

Supporters

This proposal is supported by (in alphabetical order):

- Christof Beierle, Ruhr University Bochum, Germany
- Olivier Blazy, Ecole Polytechnique, France
- Anne Canteaut, Inria, France
- Itai Dinur, Ben-Gurion University, Israel
- Tetsu Iwata, Nagoya University, Japan
- Bart Mennink, Radboud University, The Netherlands
- Léo Perrin, Inria, France
- Thomas Peyrin, NTU, Singapore
- Yu Sasaki, NTT, Japan
- Jean-Pierre Tillich, Inria, France
- *⟨Add your name here⟩*

Acknowledgement

I would like to thank Christof Beierle, Anne Canteaut, Orr Dunkelman, and Maria Naya-Placensia for comments and suggestions to improve the proposal.

A Statistics

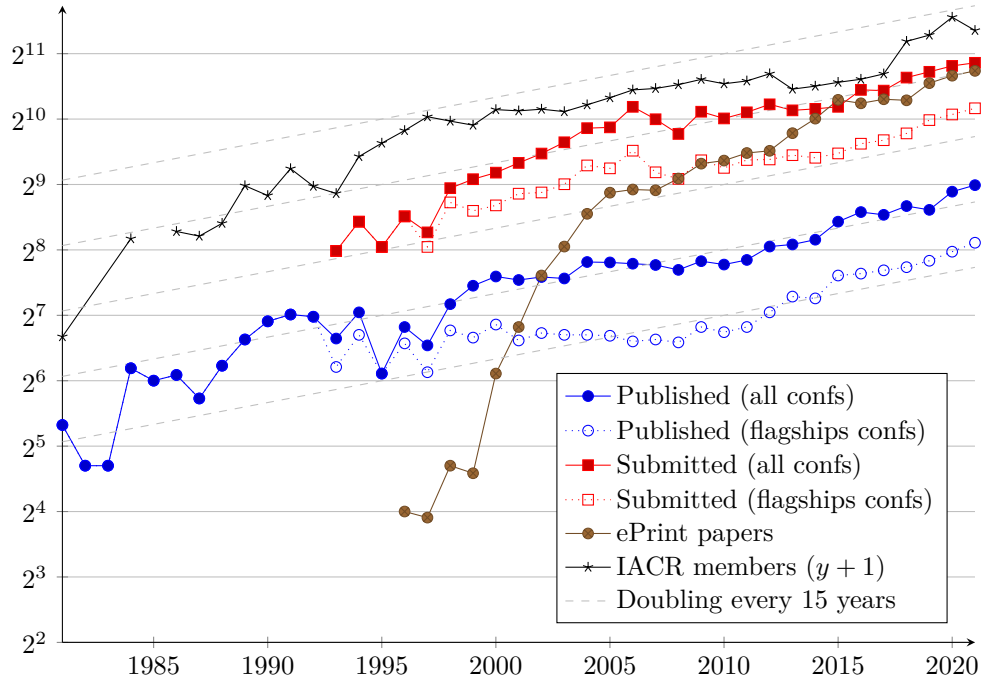

Figure 2: Growth metrics

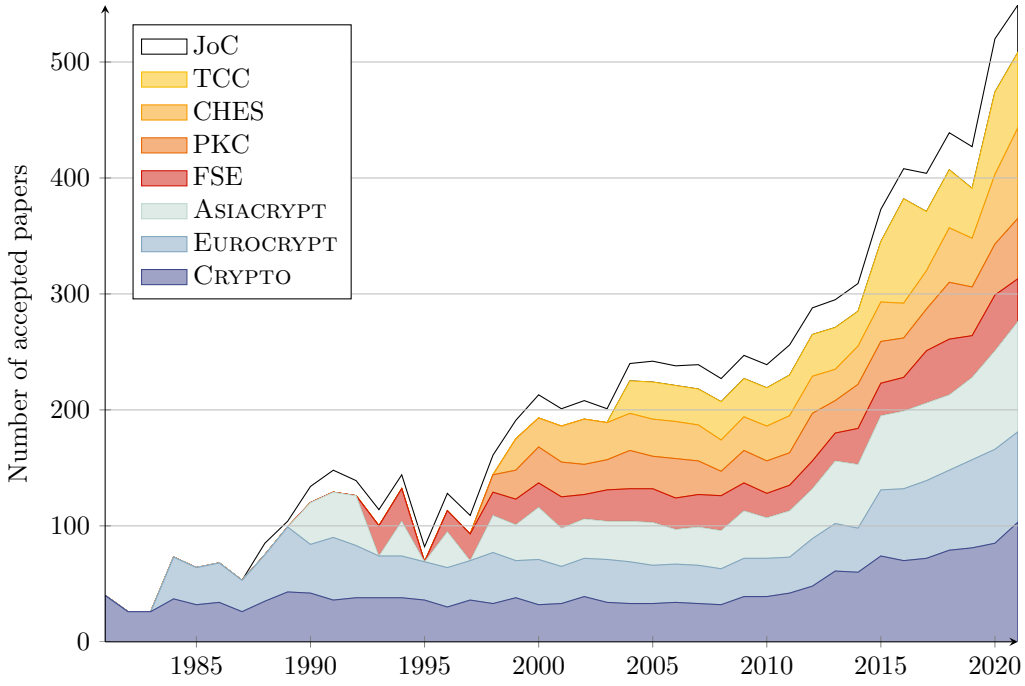

Figure 3: Accepted papers per year

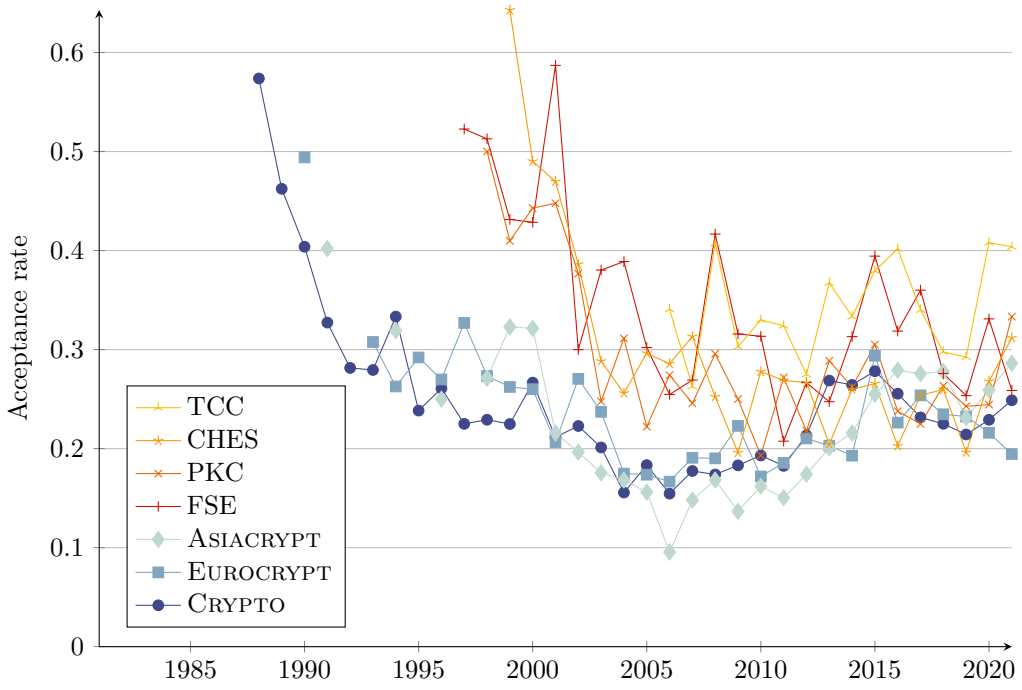

Figure 4: Acceptance rate

B A concrete proposal

A tentative proposal for the flagship journal, modelled on ToCS and TCHES

Editorial board

- Roughly 200 people, 3 editors-in-chief and 5 area editors.
- EB members and area editors renewed every year.
- Rolling EiC serving for one year, nominated every 4 months.
- EiC nominated by the IACR BoD, with a rotation by continent.
- Area editors nominated by steering committees of area conferences.
- EB members can submit at most 4 papers per year.
- EiC and area chairs can submit at most 1 paper (during one year term).

Submissions and reviews

- 3 deadlines per year, following the current calendar.
- 3 reviews per paper, authors rebuttals, and discussion between reviewers.
- Editors do most of the reviews themselves, without subreviewers.
- Decisions: accept/minor revision/major revision/reject.
  (major revisions resubmitted to one of the 2 following deadlines,
  rejected papers cannot be resubmitted in the 2 following deadlines)
- Short papers reviewed in one cycle, long papers in several cycles.
- Threshold for acceptance based on current level of flagship conferences.

Publication

- Diamond Open Access, published by Ruhr University Bochum or by IACR.
- No publication fee for authors.

Conference presentation

- Authors choose one of three conferences following acceptance to present, or don't present the paper.

Discussion. In 2021, we had 1147 submissions to the flagship conferences. Assuming that one quartet of the submissions are resubmissions, there should be about 900 submissions per year in the flagship journal. With an editorial board of 200 members (the size of the combined PC of the flagship conferences today), this amount to roughly 5 reviews per editorial board member per dead-line. With a bar for acceptance similar to the current level of the flagship conferences, we would have around 300 accepted papers per year.

The workload for the chairs will be quite consequent, with 3 dead-lines per year with 300 papers each. If this is too high, an alternative could be to have a different set of chairs for each issue (i.e. for each submission deadline).

IACR could hire an editor to finalize the camera-ready versions sent by the authors and to manage the publication and reviewing system. In order to (partially) offset the publication cost, we could require that at least one author of each accepted paper registers as an IACR member.

We suggest letting authors decide the venue where they present their paper (if any).

We anticipate that authors will have different preferences,[14] which would roughly average out. Other time, conferences with too many papers would become less desirable because of the shorter time allotted to present a paper.

The general chair would be in charge of organizing the conference, and adapting it to the number of papers to be presented (as done today with ToSC and TCHES). If necessary, conferences could resort to the short talk model used by some online and hybrid conferences (*e.g.* Eurocrypt 2021).

C Recent Changes in IACR Publications

Over the years, the number of IACR venues as increased to reach the current situation, and the number of papers accepted per venue has also increased. In the last years, some venues have also changed the way they operate, we discuss some of the most recent changes.

Real World Crypto. The RWC symposium was launched in 2012 and became an IACR event in 2018.

Sticky reviews. Since 2014, the IACR encourages authors to use "sticky reviews", *i.e.* to include reviews from the previous conference when a paper is resubmitted after being rejected.

Parallel tracks. Since 2015 flagship conferences have parallel tracks, allowing more papers to be presented.

ToSC and TCHES. FSE moved to the journal/conference hybrid model in 2016 with the creation of ToSC, and CHES moved to hybrid model in 2018 with the creation of TCHES.

Virtual and hybrid conferences. The covid pandemic also had a strong influence on IACR conferences: they were all organized online during a two-year period, and they still run as physical/virtual hybrid events at the time of writing. Virtual conferences (and some hybrid conferences) have used a different format with short live talks (5 minutes plus questions), and longer videos available online.

D Comparison with The "New Journal" Proposal

Two proposals have already been in discussion with the IACR members and the board in the last years:

- In 2013, Nigel Smart put out a proposal to turn some or all of the IACR conferences into a single journal[2]. This proposal didn't go through at the time.
- In early 2020, a proposal for a new journal that would come in addition to the current venues was submitted to the board and a "New Journal" committee started working on the proposal[15]. The proposal was first presented to the community in October 2021[16].

We agree with the issues identified in both proposals, and we agree that moving to a more journal-oriented model will solve many of them. However, we believe that the "New

[14] Some authors will choose the first conference after the notification, some will choose the conference closer to their location, some will choose based on some perceived notion of prestige, ...

[15] https://www.iacr.org/docs/minutes/virtual-5_2020bod.pdf

[16] https://iacr.org/docs/minutes/ec2021mem-slides.pdf

Journal" proposal is not up to the challenge and will not solve the issues in a satisfactory way.

Instead, we propose to rethink the publication strategy of the IACR, and to revisit the plan proposed in 2013 taking into account what has happened since. In particular, the creation of ToSC and TCHES have shown that moving a conference to the hybrid model works well, and has confirmed that journals have clear benefits.

D.1 The "New Journal" Proposal

The "New Journal" proposal consists in creating a new journal in addition to the current venues. The stated goals of the proposal are:

- Reducing the review load
- Allowing scaling to accommodate the growth of the community
- Respecting all areas of cryptography
- Minimizing the impact of the new journal on existing IACR venues

As discussed in Section 2 and 3, these are the hardest issues to solve with the existing system, so creating a new journal might be a way to address them.

The "New Journal" proposal is not completely fleshed out, so we evaluate it to the best of our understanding. In particular, at the time of writing, it is still unclear how the new journal would be positioned compared to existing IACR venues. We believe the journal would have a very different role depending on how selective it is: a journal with a high threshold for acceptance is likely to get a better ranking and to attract higher quality papers. In the following, we consider two options: a new journal with a clear positioning below to area conferences, and a new journal that aims to publish papers with a quality level similar to the area conferences or higher.

D.1.1 New Journal Below the Area Conferences

A new journal positioned below the area conferences would give authors an outlet to published results that are correct but maybe not considered interesting enough by the program committees of conferences (for instance, somewhat incremental results). It would mostly take papers that are currently published outside the IACR in conferences like SAC, CT-RSA, Indocrypt, Latincrypt, ...

However, a journal that accepts papers that would be rejected from the area conferences is unlikely to be judged as a high quality venue, and most authors will probably prefer to submit to other venues first. Therefore, it will only solve a small number of the problems identified:

Scaling. If the community grows significantly, the new journal can absorb the growth, leading to more selective conference, and papers of lower quality going to the journal.

Limiting travel. If authors publish papers in the new journal instead of a conference, they will travel less. Authors that do not want to travel are offered an alternative to conferences, but they will be unhappy if the only IACR option is a low-ranking one.

Respecting all areas. If the main acceptance criterion is that the paper must be correct, it will naturally accept papers from areas that are currently under-represented in IACR venues. However, authors will probably prefer a higher ranking venue and complain if IACR conferences do not respect their area.

Minimizing impact. If the new journal is clearly identified as below the area conferences, it will have a minimal impact on the existing IACR venues. At the same time, it will likely have a strong effect on lower-tier conferences (we cannot simultaneously reduce overall travel and maintain those conferences).

Review load. Most authors will still submit their papers iteratively to the flagship conferences and/or to the area conferences before submitting to the new journal.

Other risks. There is a risk to damage the reputation of IACR by publishing lower quality papers. There is also a risk that outsiders specifically associate IACR journals with a lower perceived quality, and that this impacts ToSC and TCHES negatively.

D.1.2 New Journal Comparable to or Above Area Conferences

A journal with a high bar for acceptance is likely to get a better ranking and to attract higher quality papers. This leads to a higher impact on the existing systems, both for positive and negative impact.

Scaling. If the bar is high, the journal will publish a smaller number of papers per year, and will maybe not absorb the full growth of the community.

Limiting travel. If authors publish papers in the new journal instead of a conference, they will travel less. Authors that do not want to travel are offered a high-ranking alternative to conferences.

Respecting all areas. If the journal is selective, it will face the same difficulties to fairly evaluate papers from all areas as the flagship conferences. Having a perceived quality bar rather than having papers compete each other could make the process more fair, but the program committee of the new venue will likely have the same biases as the current committees. Any measure that can be taken to improve the balance in a new venue can and should also be taken for existing venues. (There is no reason to believe that it will be easier to respect all areas by creating a new venue).

Minimizing impact. If the journal accepts papers of the same quality as existing venues, it will directly compete with them. There is a risk that the new journal absorbs a significant number of papers from some area conferences, and forces them to move to a different format.

Review load. If the journal absorbs a significant number of submissions from the existing venues it will potentially reduce the overall review load. However, this is only possible by having a significant impact on the existing venues, and will not benefit areas that already use a journal/conference hybrid.

Other risks. The new journal would create a complicated publication landscape where we have area conferences with proceeding, area conferences with hybrid journals, and a new journal, all with similar papers and goals.

Evaluation. Both options (with a high bar or a low bar) seem to have significant drawbacks, and will hardly meet the stated goals (reduce the review load, respect all areas, minimize impact).

The New Journal makes the publication landscape more complex for authors (one more option to choose from), and for the IACR (one more venue to run with its own board). There is also a significant risk to existing venues if the journal is successful: it would strongly affect area conferences if the bar for acceptance is similar, or smaller events (if particular, many events run "in cooperation with IACR") if the bar is lower. More generally, the problems that have been identified are problems of the existing venues,

therefore, solving them will require modifying their organization rather than adding a new venue to complement them.

Having said that, a new journal that is clearly positioned below the existing venues would fill a niche that is not served by existing IACR venues: publishing results that are correct, but have trouble convincing reviewers of their interest. Such a journal would compete with events run in cooperation with the IACR and its impact should be discussed with the community.