

# Evaluating the artifacts of SIGCOMM papers

Damien Saucez  
Université Côte d’Azur, Inria, France  
damien.saucez@inria.fr

Luigi Iannone  
Telecom ParisTech, France  
luigi.iannone@telecom-paristech.fr

Olivier Bonaventure  
UCLouvain, Belgium  
Olivier.Bonaventure@uclouvain.be

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.  
The authors take full responsibility for this article’s technical content. Comments can be posted through CCR Online.

## ABSTRACT

A growing fraction of the papers published by CCR and at SIGCOMM-sponsored conferences include artifacts such as software or datasets. Besides CCR, these artifacts were rarely evaluated. During the last months of 2018, we organised two different Artifacts Evaluation Committees to which authors could submit the artifacts of their papers for evaluation. The first one evaluated the papers accepted by Conext’18 shortly after the TPC decision. It assigned ACM reproducibility badges to 12 different papers. The second one evaluated papers accepted by CCR and any SIGCOMM-sponsored conference. 28 papers received ACM reproducibility badges. We report on the results of a short survey among artifacts authors and reviewers and provide some suggestions for future artifacts evaluations.

## CCS CONCEPTS

• **General and reference** → **Evaluation**;

## KEYWORDS

Artifacts, Reproducibility

## 1 INTRODUCTION

Latest years have witnessed a steadily growing number of the papers, accepted by Computer Communication Review and the SIGCOMM-sponsored conferences, including artifacts such as simulation models, measurement datasets, software implementations, etc. These artifacts are an essential part of many of these papers, and artifacts’ availability encourages other researchers to build upon and reproduce and extend previous results.

The ACM has proposed guidelines for assessing the quality of artifacts in publications<sup>1</sup>

These two evaluations focused on assessing if artifacts were **available**, **functional**, or **reusable**; which definitions are given by the ACM as follows.

- **Artifacts Available**: author-created artifacts relevant to this paper have been placed on a publicly accessible archival repository.
- **Artifacts Evaluated - Functional**: the artifacts associated with the research are found to be documented, consistent, complete, exercisable, and include appropriate evidence of verification and validation.
- **Artifacts Evaluated - Reusable**: the artifacts associated with the paper are of a quality that significantly exceeds minimal functionality.

<sup>1</sup><https://www.acm.org/publications/policies/artifact-review-badging>



Figure 1: Artifacts badges used for SIGCOMM evaluation.

The ACM proposes two additional definitions for results validation, **Results Replicated** and **Results Reproduced**. In an ideal world the evaluation committee should also have validated results. However, validating results is time consuming and the committees were not having enough resource to accomplish this mission. Meanwhile, it is worth to mention that most of the time when artifacts were evaluated as functional or reusable in our two evaluations, the paper results were also replicated. However, as we didn’t define strict guidelines for results validation we could not conclude on the actual validity of results. Hence the choice of focusing on the artifacts only.

The ACM associates a badging system to these definitions. These badges can be used to visually indicate the conclusions of the artifacts evaluation committee. Badges used in our evaluations are presented in Fig. 1.

As the objective was to promote reproducibility and open science, the evaluation process was incremental with interactions with the authors to improve the quality of artifacts when possible. For that reason, the artifacts study was optional and authors had to expressly apply in order to have their artifact evaluated. Therefore, the absence of badge on a 2018 SIGCOMM-sponsored venue paper doesn’t indicate a lack of reproducibility of a paper.

## 2 CONEXT’18 ARTIFACTS EVALUATION RESULTS

The evaluation of CoNEXT’18 papers’ artifacts was carried out shortly after the acceptance notification. Out of 14 accepted papers proposing an artifact, 12 have been awarded a badge. Seven of them received the *Artifacts Available* badge.

- DenseVLC: A Cell-Free Massive MIMO System with Distributed LEDs [6]

- Dynam-IX: a Dynamic Interconnection eXchange [20]
- Intent-Driven Composition of Resource-Management SDN Applications [12]
- Memento: Making Sliding Windows Efficient for Heavy Hitters [5]
- P4Visor: lightweight virtualization and composition primitives for building and testing modular programs<sup>2</sup> [38]
- P-Rex: Fast Verification of MPLS Networks with Multiple Link Failures [15]
- REINFORCE: Achieving Efficient Failure Resiliency for Network Function Virtualization based Services [16]

Two papers received both the *Artifacts available* and the *Artifacts Evaluated - Functional* badges.

- Boosting fine-grained activity sensing by embracing wireless multipath effects [26]
- Robustly Disjoint Paths with Segment Routing [1]

The CoNEXT artifacts reviewers agreed to assign the *Artifacts available*, *Artifacts Evaluated - Functional* and *Artifacts Evaluated - Reusable* badges to two papers.

- Verification of P4 Programs in Feasible Time using Assertions [25]
- Leveraging eBPF for programmable network functions with IPv6 Segment Routing [36]

Finally, the *Artifacts Evaluated - Reusable* badge was assigned to the following paper that describes optimisations that are included in the Linux kernel. It has been decided to provide only the *Artifacts Evaluated - Reusable* badge to this paper as some artifacts used in the paper were not “author-created” but were coming from the community.

- The eXpress Data Path: Fast Programmable Packet Processing in the Operating System Kernel [13]

### 3 SIGCOMM-SPONSORED ARTIFACTS EVALUATION RESULTS

This evaluation was organised during the winter 2018-2019. It was open to all accepted 6+ pages papers of SIGCOMM sponsored 2018 conferences and journal. The Artifact Evaluation Committee received 32 papers for evaluation in 2018, out of which the following 28 have been awarded with badges. The following nine articles received the *Artifacts Available* badge.

- A First Look at Certification Authority Authorization (CAA) [31]
- A Formally Verified NAT Stack [28]
- Inferring Persistent Interdomain Congestion [10]
- Network-Wide Routing-Oblivious Heavy Hitters [4]
- On the Origins of Memes by Means of Fringe Web Communities [37]
- Scanning the Internet for Liveness [2]
- Studying TLS Usage in Android Apps [29]
- The Rise of Certificate Transparency and Its Implications on the Internet Ecosystem [32]
- Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering [30]

<sup>2</sup>This paper has been re-evaluated later, see Sec. 3

Seven papers were tagged with both the *Artifacts Available* and the *Artifacts Evaluated - Functional* badges.

- Accelerating Network Measurement in Software [39]
- Data-driven Resource Flexing for Network Functions Virtualization [8]
- Homa: A Receiver-Driven Low-Latency Transport Protocol Using Network Priorities [22]
- Mobility Support in Cellular Networks: A Measurement Study on Its Configurations and Implications [9]
- On Collaborative Predictive Blacklisting [21]
- Restructuring Endpoint Congestion Control [23]
- YATES: Rapid Prototyping for Traffic Engineering Systems [17]

Finally, the reviewers agreed to assign the three badges to twelve papers.

- A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists [33]
- Automated Synthesis of Adversarial Workloads for Network Functions [27]
- Cuckoo++ Hash Tables: High-Performance Hash Tables for Networking Applications [34]
- HIPE – An Energy-Status-Data Set from Industrial Production [7]
- How much demand side flexibility do we need? Analyzing where to exploit flexibility in industrial processes [3]
- Multilevel MDA-Lite Paris Traceroute [35]
- On low-latency-capable topologies, and their impact on the design of intra-domain routing [11]
- P4Visor: lightweight virtualization and composition primitives for building and testing modular programs [38]
- SketchLearn: Relieving User Burdens in Approximate Measurement with Automated Statistical Inference [14]
- Understanding PCIe performance for end host networking [24]
- Understanding Tor Usage with Privacy-Preserving Measurement [19]
- Want to Reduce Energy Consumption, Whom should we call? [18]

As one can see the P4Visor [38] paper has been re-evaluated. Indeed, authors have significantly reworked their artifacts and requested to be re-evaluated given their efforts. As our ultimate objective is not to evaluate artifacts but to promote reproducibility in general, we accepted and even encouraged this request.

### 4 AUTHORS AND REVIEWERS FEEDBACK

After the publication of the badges, we sent a short survey to the authors that applied to the SIGCOMM-sponsored artifacts evaluation and to the reviewers who analysed them.

We received 26 responses from the authors of paper artifacts. These artifacts were mainly software (23 out of 26) followed by datasets (13 out of 26) and scripts (10 out of 26). Sixty percent of the authors agreed that the reviewers who analysed their artifacts were competent. The remaining authors disagree or strongly disagreed on the competence of the reviewers. This indicates that finding the right reviewers to evaluate artifacts is not simple and indeed we sometimes had to send many emails to try to find a candidate

reviewer. An Artifact Evaluation Committee associated to a specific conference might not have this problem. More than three quarters of the authors strongly agreed or agreed that the comments that they received have helped them to improve the quality of their artifacts. A majority of 88% of the authors would recommend their colleagues to also send their artifacts for evaluation.

We received 16 responses from the reviewers of the paper artifacts. Among them, 56.3% agreed that they learned useful information while reviewing artifacts and 25% of them strongly agreed with this statement. A total of 80% of the reviewers think that the reviews and the discussions with the authors have improved the quality of the artifacts. They all agree or strongly agree to recommend their colleagues to also participate in the evaluation of paper artifacts.

Finally, we asked the same question to both authors and reviewers: *How should the evaluation of artifacts be organised in the future?* Less than 20% of the authors and the reviewers considered that organising one artifacts evaluation per year was a good approach. Among the authors, more than half of them were in favor of evaluating the artifacts between paper acceptance and the conference. 20% of them were in favor of evaluating the artifacts after the conference. 43% of the reviewers were also in favor of this organisation. From an editorial perspective, badging papers after their publication is much more complicated since the ACM Digital Library needs to be updated manually.

Some survey replies also provided interesting feedback which could help the organisers of future Artifacts Evaluation Committees. A first point is that it would be useful to provide a public version of the review form before the conference so that authors know what reviewers will assess. A second point is that many artifacts reviewers “discussed” with the authors of the artifacts to solve practical problems and hotcrp was not perfect in handling such frequent discussions. A third point was that there is not yet a consensus within the community on what reviewers should expect from paper artifacts. For software, some reviewers had difficulties in having the right modules and libraries to recompile/use the software provided by the artifacts. Some authors released docker containers and virtual machine images to simplify the installation of their artifacts. This is probably a good idea for some types of artifacts, but another burden on authors who release their artifacts. The same applies for datasets. Some of them were released as raw data while others included scripts or software libraries to easily extract data from them.

## 5 CONCLUSION

We provided a brief summary of the results of the evaluation of artifacts for Conext’18 and other SIGCOMM venues last year.

Authors and reviewers consider that evaluating the artifacts is useful. SIGCOMM should probably encourage its sponsored conferences to organise an evaluation committee for the papers’ artifacts.

From a workflow viewpoint, it appears that evaluating the artifacts between paper acceptance and paper presentation seems to be the best compromise. However, the artifacts evaluation should not start too early as authors might need to update their artifacts based on the comments received from the regular TPC. The camera-ready

version deadline could be a good time to start the evaluation of the artifacts so that it can be finished before the conference.

Overall, our experience in evaluating artifacts is that this is useful, interesting but sometimes time-consuming for the reviewers. We encourage the conference organisers to create Artifacts Evaluation Committees to evaluate papers accepted by the conference TPC.

## ACKNOWLEDGMENTS

We would like to thank all the artifacts reviewers: Ronaldo Alves Ferreira, Omid Ardakanian, Mathieu Bouet, Anna Brunstrom, Marc Bruyere, Lianjie Cao, Kenjiro Cho, Walid Dabbous, Quentin De Coninck, Thibault Delmas, Giuseppe Di Lena, Benoit Donnet, Ramakrishnan Durairajan, Paul Emmerich, Marcel Enguehard, Christian Esteve Rothenberg, Oliver Gasser, Yossi Gilad, Thomas Given-Wilson, Mohammed Hawari, Ralph Holz, Michio Honda, Costas Iordanou, Mathieu Jadin, Praveen Kumar, Sebastian Lehnhoff, Simon Leinen, Keunwoo Lim, Leonardo Linguaglossa, Jed Liu, Nicole Ludwig, Andra Lutu, Francois Michel, Cristel Pelsser, Haroon Rashid, Quirin Scheitle, Rute Sofia, Joel Sommers, Holger Trittenbach, Adrien Wion, Arseniy Zaostrovnykh, Noa Zilberman, Thomas Zinner. We would also like to thank Theophilus Benson and Laurent Vanbever for supporting the artifacts evaluation committee for CoNEXT 2018.

## REFERENCES

- [1] AUBRY, F., VISSICCHIO, S., BONAVENTURE, O., AND DEVILLE, Y. Robustly disjoint paths with segment routing. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT ’18, ACM, pp. 204–216.
- [2] BANO, S., RICHTER, P., JAVED, M., SUNDARESAN, S., DURUMERIC, Z., MURDOCH, S. J., MORTIER, R., AND PAXSON, V. Scanning the internet for liveness. *SIGCOMM Comput. Commun. Rev.* 48, 2 (May 2018), 2–9.
- [3] BARTH, L., HAGENMEYER, V., LUDWIG, N., AND WAGNER, D. How much demand side flexibility do we need?: Analyzing where to exploit flexibility in industrial processes. In *Proceedings of the Ninth International Conference on Future Energy Systems* (New York, NY, USA, 2018), e-Energy ’18, ACM, pp. 43–62.
- [4] BASAT, R. B., EINZIGER, G., FEIBISH, S. L., MORANEY, J., AND RAZ, D. Network-wide routing-oblivious heavy hitters. In *Proceedings of the 2018 Symposium on Architectures for Networking and Communications Systems* (New York, NY, USA, 2018), ANCS ’18, ACM, pp. 66–73.
- [5] BASAT, R. B., EINZIGER, G., KESLASSY, I., ORDA, A., VARGAFTIK, S., AND WAISBARD, E. Memento: Making sliding windows efficient for heavy hitters. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT ’18, ACM, pp. 254–266.
- [6] BEYSENS, J., GALISTEO, A., WANG, Q., JUARA, D., GIUSTINIANO, D., AND POLLIN, S. Densevcl: A cell-free massive mimo system with distributed leds. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT ’18, ACM, pp. 320–332.
- [7] BISCHOF, S., TRITTENBACH, H., VOLLMER, M., WERLE, D., BLANK, T., AND BÖHM, K. Hipe: An energy-status-data set from industrial production. In *Proceedings of the Ninth International Conference on Future Energy Systems* (New York, NY, USA, 2018), e-Energy ’18, ACM, pp. 599–603.
- [8] CAO, L., FAHMY, S., SHARMA, P., AND ZHE, S. Data-driven resource flexing for network functions visualization. In *Proceedings of the 2018 Symposium on Architectures for Networking and Communications Systems* (New York, NY, USA, 2018), ANCS ’18, ACM, pp. 111–124.
- [9] DENG, H., PENG, C., FIDA, A., MENG, J., AND HU, Y. C. Mobility support in cellular networks: A measurement study on its configurations and implications. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC ’18, ACM, pp. 147–160.
- [10] DHAMDHARE, A., CLARK, D. D., GAMERO-GARRIDO, A., LUCKIE, M., MOK, R. K. P., AKIWATE, G., GOGIA, K., BAJPAI, V., SNOEREN, A. C., AND CLAFFY, K. Inferring persistent interdomain congestion. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM ’18, ACM, pp. 1–15.
- [11] GVOZDIEV, N., VISSICCHIO, S., KARP, B., AND HANDLEY, M. On low-latency-capable topologies, and their impact on the design of intra-domain routing. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM ’18, ACM, pp. 88–102.

- [12] HEORHIADI, V., CHANDRASEKARAN, S., REITER, M. K., AND SEKAR, V. Intent-driven composition of resource-management sdn applications. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 86–97.
- [13] HØILAND-JØRGENSEN, T., BROUER, J. D., BORKMANN, D., FASTABEND, J., HERBERT, T., AHERN, D., AND MILLER, D. The express data path: Fast programmable packet processing in the operating system kernel. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 54–66.
- [14] HUANG, Q., LEE, P. P. C., AND BAO, Y. Sketchlearn: Relieving user burdens in approximate measurement with automated statistical inference. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM '18, ACM, pp. 576–590.
- [15] JENSEN, J. S., KRØGH, T. B., MADSEN, J. S., SCHMID, S., SRBA, J., AND THORGERSEN, M. T. P-rex: Fast verification of mpls networks with multiple link failures. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 217–227.
- [16] KULKARNI, S. G., LIU, G., RAMAKRISHNAN, K., ARUMAITHURAI, M., WOOD, T., AND FU, X. Reinforce: Achieving efficient failure resiliency for network function virtualization based services. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 41–53.
- [17] KUMAR, P., YU, C., YUAN, Y., FOSTER, N., KLEINBERG, R., AND SOULÉ, R. Yates: Rapid prototyping for traffic engineering systems. In *Proceedings of the Symposium on SDN Research* (New York, NY, USA, 2018), SOSR '18, ACM, pp. 11:1–11:7.
- [18] MAMMEN, P. M., KUMAR, H., RAMAMRITHAM, K., AND RASHID, H. Want to reduce energy consumption, whom should we call? In *Proceedings of the Ninth International Conference on Future Energy Systems* (New York, NY, USA, 2018), e-Energy '18, ACM, pp. 12–20.
- [19] MANI, A., WILSON-BROWN, T., JANSEN, R., JOHNSON, A., AND SHERR, M. Understanding tor usage with privacy-preserving measurement. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC '18, ACM, pp. 175–187.
- [20] MARCOS, P., CHIESA, M., MÜLLER, L., KATHIRAVELU, P., DIETZEL, C., CANINI, M., AND BARCELLOS, M. Dynam-ix: A dynamic interconnection exchange. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 228–240.
- [21] MELIS, L., PYRGELIS, A., AND DE CRISTOFARO, E. On collaborative predictive blacklisting. *SIGCOMM Comput. Commun. Rev.* 48, 5 (Jan. 2019), 9–20.
- [22] MONTAZERI, B., LI, Y., ALIZADEH, M., AND OUSTERHOUT, J. Homa: A receiver-driven low-latency transport protocol using network priorities. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM '18, ACM, pp. 221–235.
- [23] NARAYAN, A., CANGIALOSI, F., RAGHAVAN, D., GOYAL, P., NARAYANA, S., MITTAL, R., ALIZADEH, M., AND BALAKRISHNAN, H. Restructuring endpoint congestion control. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM '18, ACM, pp. 30–43.
- [24] NEUGEBAUER, R., ANTICHI, G., ZAZO, J. F., AUDZEVICH, Y., LÓPEZ-BUENO, S., AND MOORE, A. W. Understanding pcie performance for end host networking. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM '18, ACM, pp. 327–341.
- [25] NEVES, M., FREIRE, L., SCHAEFFER-FILHO, A., AND BARCELLOS, M. Verification of p4 programs in feasible time using assertions. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 73–85.
- [26] NIU, K., ZHANG, F., XIONG, J., LI, X., YI, E., AND ZHANG, D. Boosting fine-grained activity sensing by embracing wireless multipath effects. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 139–151.
- [27] PEDROSA, L., IYER, R., ZAOSTROVNYKH, A., FIETZ, J., AND ARGYRAKI, K. Automated synthesis of adversarial workloads for network functions. In *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication* (New York, NY, USA, 2018), SIGCOMM '18, ACM, pp. 372–385.
- [28] PIRELLI, S., ZAOSTROVNYKH, A., AND CANDEA, G. A formally verified nat stack. *SIGCOMM Comput. Commun. Rev.* 48, 5 (Jan. 2019), 77–83.
- [29] RAZAGHPANAH, A., NIAKI, A. A., VALLINA-RODRIGUEZ, N., SUNDARESAN, S., AMANN, J., AND GILL, P. Studying tls usage in android apps. In *Proceedings of the Applied Networking Research Workshop* (New York, NY, USA, 2018), ANRW '18, ACM, pp. 5–5.
- [30] REUTER, A., BUSH, R., CUNHA, I., KATZ-BASSETT, E., SCHMIDT, T. C., AND WÄHLISCH, M. Towards a rigorous methodology for measuring adoption of rpki route validation and filtering. *SIGCOMM Comput. Commun. Rev.* 48, 1 (Apr. 2018), 19–27.
- [31] SCHEITTE, Q., CHUNG, T., HILLER, J., GASSER, O., NAAB, J., VAN RIJSWIJK-DEIJ, R., HOHLFELD, O., HOLZ, R., CHOFFNES, D., MISLOVE, A., AND CARLE, G. A first look at certification authority authorization (caa). *SIGCOMM Comput. Commun. Rev.* 48, 2 (May 2018), 10–23.
- [32] SCHEITTE, Q., GASSER, O., NOLTE, T., AMANN, J., BRENT, L., CARLE, G., HOLZ, R., SCHMIDT, T. C., AND WÄHLISCH, M. The rise of certificate transparency and its implications on the internet ecosystem. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC '18, ACM, pp. 343–349.
- [33] SCHEITTE, Q., HOHLFELD, O., GAMBA, J., JELTEN, J., ZIMMERMANN, T., STROWES, S. D., AND VALLINA-RODRIGUEZ, N. A long way to the top: Significance, structure, and stability of internet top lists. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC '18, ACM, pp. 478–493.
- [34] SCOURNEC, N. L. Cuckoo++ hash tables: High-performance hash tables for networking applications. In *Proceedings of the 2018 Symposium on Architectures for Networking and Communications Systems* (New York, NY, USA, 2018), ANCS '18, ACM, pp. 41–54.
- [35] VERMEULEN, K., STROWES, S. D., FOURMAUX, O., AND FRIEDMAN, T. Multilevel mda-lite paris traceroute. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC '18, ACM, pp. 29–42.
- [36] XHONNEUX, M., DUCHENE, F., AND BONAVENTURE, O. Leveraging ebpf for programmable network functions with ipv6 segment routing. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 67–72.
- [37] ZANNETTOU, S., CAULFIELD, T., BLACKBURN, J., DE CRISTOFARO, E., SIRIVIANOS, M., STRINGHINI, G., AND SUAREZ-TANGLI, G. On the origins of memes by means of fringe web communities. In *Proceedings of the Internet Measurement Conference 2018* (New York, NY, USA, 2018), IMC '18, ACM, pp. 188–202.
- [38] ZHENG, P., BENSON, T., AND HU, C. P4visor: Lightweight virtualization and composition primitives for building and testing modular programs. In *Proceedings of the 14th International Conference on Emerging Networking Experiments and Technologies* (New York, NY, USA, 2018), CoNEXT '18, ACM, pp. 98–111.
- [39] ZHOU, Y., ALIPOURFARD, O., YU, M., AND YANG, T. Accelerating network measurement in software. *SIGCOMM Comput. Commun. Rev.* 48, 3 (Sept. 2018), 2–12.