

# A Survey on Artifacts from CoNEXT, ICN, IMC, and SIGCOMM Conferences in 2017

Matthias Flittner<sup>†</sup>, Mohamed Naoufal Mahfoudi<sup>‡</sup>, Damien Saucez<sup>‡</sup>, Matthias Wählisch<sup>◇</sup>, Luigi Iannone<sup>♡</sup>, Vaibhav Bajpai<sup>♣</sup>, Alex Afanasyev<sup>♣</sup>

<sup>†</sup>Karlsruhe Institute of Technology, <sup>‡</sup>Université Côte d’Azur, Inria, <sup>◇</sup>Freie Universität Berlin, <sup>♡</sup>Telecom Paristech, <sup>♣</sup>Technical University of Munich, <sup>♣</sup>Florida International University

<sup>†</sup>matthias.flittner@kit.edu, <sup>‡</sup>{first.last}@inria.fr, <sup>◇</sup>m.waehlich@fu-berlin.de,

<sup>♡</sup>luigi.iannone@telecom-paristech.fr, <sup>♣</sup>bajpai@in.tum.de, <sup>♣</sup>aa@cs.fiu.edu

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 It has NOT been peer reviewed.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

Reproducibility of artifacts is a cornerstone of most scientific publications. To improve the current state and strengthen ongoing community efforts towards reproducibility by design, we conducted a survey among the papers published at leading ACM computer networking conferences in 2017: CoNEXT, ICN, IMC, and SIGCOMM.

The objective of this paper is to assess the current state of artifact availability and reproducibility based on a survey. We hope that it will serve as a starting point for further discussions to encourage researchers to ease the reproduction of scientific work published within the SIGCOMM community. Furthermore, we hope this work will inspire program chairs of future conferences to emphasize reproducibility within the ACM SIGCOMM community as well as will strengthen awareness of researchers.

## CCS Concepts

•General and reference → *General conference proceedings; General literature;*

## Keywords

Reproducibility, Artifacts, Survey

## 1. INTRODUCTION

Research on computer networks studies human-made systems. Compared to high energy physics, for example, it is relatively inexpensive to produce research artifacts such as measurements data or software. According to the ACM definition of artifacts, an artifact is “*a digital object that was either created by the authors to be used as part of the study or generated by the experiment itself*” [53]. In this survey, we report on a follow-up of the ACM SIGCOMM 2017 Reproducibility Workshop [54] and show a brief overview of the nature of the artifacts that were produced in 2017 in four leading conferences of the ACM SIGCOMM, namely SIGCOMM, CoNEXT, IMC, and ICN. To that aim, we asked authors of papers in these conferences to fill an online survey to describe their artifacts.

The remainder of this report is structured as follows: In Section 2, we describe the survey and recruiting of participants in more detail. In Section 3, we analyze the submitted artifacts. We present a comprehensive discussion and conclusion, in Sections 4 and 5 respectively.

## 2. ARTIFACT SURVEY

In order to assess the current state of artifact availability and current practices in the networking community, we conducted an informal survey. We collected data from networking researchers that published scientific work in 2017. We focused on authors of conferences which are sponsored by ACM and ACM SIGCOMM, including venues with broad topics (SIGCOMM, CoNEXT), as well as more domain-specific venues (IMC and ICN). CoNEXT is a conference on novel and emerging networking technologies; ICN is a conference on Information-Centric Networking; IMC is a conference on Internet measurement and analysis, and SIGCOMM is a major generic conference in the field of communications and computer networks.

*Questionnaire.* Our questionnaire was implemented using Google Forms. We grouped the questions in three parts. The first three questions asked for paper title, conference name, and author email. Then, in order to allow the authors to explain their artifacts in their own words, we asked for a brief but precise and complete description of the provided tools and data. The participants were also required to provide an URL to access the artifacts in question. All responders provided a link to their artifacts and in general a description of the artifacts more precise than what could be found in their paper. Among all responses we noticed that 3 papers were not listing any link to their artifacts in the paper itself. The questionnaire finished with a scaled question about the easiness to reproduce the paper. This self-assessment allowed the researchers judge their publications on a scale from 1 (“easy – an undergraduate can do that”) to 10 (“hard – only I can do that”).

*Recruiting Participants.* To solicit participation in the survey, we asked the chairs of the technical program committees to contact all authors of accepted papers. Invitations were sent out after the conferences. It is worth noting that some authors did not receive the email, in case they disabled notifications in the conference submission system. However, at least one author per paper was reached, and we removed duplicate submissions. Furthermore, we should note that we increased the number of survey participants of ACM ICN by sending a reminder.

After we collected the provided data, we analyzed the artifact descriptions and the actual artifacts in more detail.

Conference	Potential Responses [# papers]	Actual Responses [# papers]	Response Rate [%]
CoNEXT	40	8	20
ICN	19	12	63
IMC	42	17	40
SIGCOMM	36	12	33
<b>Total</b>	137	49	35.8

**Table 1: Summary of artifact survey, compared to the overall number of published papers per conference.**

Importantly, those authors who replied provide at least the same amount of information in their published paper.

*Participants.* Out of the 137 potential respondents, 49 researchers (35.8%) participated in the survey. Surprisingly, the response rate was quite diverse among the conferences. Most of the ICN authors (63%) were responsive, followed by IMC (40%) and SIGCOMM (33%). Only 20% of the CoNEXT authors participated in the survey. We summarize these observations in Table 1. We provide a detailed analysis of the survey replies in the next section.

The form and all the data we collected are available online [55].

### 3. ARTIFACTS AND REPRODUCIBILITY BY RESEARCH TOPIC

#### 3.1 Grouping of Artifacts

We categorized the 49 responses into the following topics:

**Architectural** are papers aiming at providing a new network algorithm, protocol, or architecture.

**Measurements** are papers that focus on measuring an already installed system.

**Miscellaneous (misc.)** papers that do not fit directly in the other topics, typically optical networks or security.

We identified 22 architectural papers [1–3, 5, 6, 10, 12, 13, 15, 16, 24, 39–41, 43–50]. The scientific work described in 19 papers is categorized as measurements [4, 7, 8, 17, 19, 25–38]. Finally, 8 papers are classified as miscellaneous [9, 11, 14, 18, 20, 22, 23, 42].

We then introduced an orthogonal classification of the artifacts themselves:

**Tools** groups all artifacts that are significantly based on software which was developed or used to conduct the research presented in the paper.

**Hardware** groups artifacts which depend on specialized hardware.

**Simulation** groups artifacts which are obtained by numerical evaluation, simulation, or emulation.

**Dataset** groups artifacts which are based on an external dataset.

	Arch.	Measurements	Misc.
<b>Tools</b>			
New	8	11	6
NDN [63]	3	–	–
CCN-Lite [59]	2	–	–
Linux/RIOT [58]	6	–	–
Other	8	4	2
<b>Hardware</b>			
New	–	–	1
Smartphones	3	–	1
Specific	3	1	1
<b>Simulation</b>			
New	1	–	–
Matlab [62]	–	–	2
ndnSim [57]	2	–	–
Other	4	–	1
<b>Dataset</b>			
New	–	12	3
CAIDA [61]	2	4	–
Other	4	4	1
<b>Testbed</b>			
Private	2	5	2
IoT-Lab [56]	2	–	–
RIPE [60]	–	4	–
Other	1	6	1
<b>Average rank</b>	4.2	3.5	2.0

**Table 2: Summary of artifact nature. Please note: Some artifacts are counted in multiple rows (if applicable). But only once per column.**

**Testbed** groups artifacts which are based on a testbed or a specific infrastructure.

For each type of artifact, we identified three options. Either the artifact is *new* (i.e., researchers had to build the artifact on their own) or the artifact is built upon existing material. When a previously existing material has been used by at least two papers, we highlight the artifact by naming it explicitly in the table. If existing material is used by only one paper, we summarize those artifacts by *Other*. It is worth noting that we count an artifact in the table only if it is made available (by some sort) to the community. The only exception are private testbeds, which have been used by users but that cannot be shared.

#### 3.2 Analysis & Observations

In the following, we briefly summarize our observations. Table 2 shows the number of papers for each research topic and the applied methodology. Significant differences are visible among the fields. For architectural papers, researchers tend to use existing tools or modify the operating system directly. On the other hand, researchers in the measurement domain mostly created their own tools (typically automation scripts). A trade-off is followed by the ICN community, which extends libraries and well-established tools or create their own new tools from scratch.

Researchers who publish at ICN tend to use more spe-

cific hardware, compared to other conferences. This is not surprising for two reasons. First, the ICN community usually evaluates their work in experiments. Second, a major topic at ACM ICN 2017 was IoT, which involves in experiments special hardware such as constrained devices or smartphones. This might restraint the possibility for anyone to reproduce the work as some hardware must be purchased (or borrowed) first. One paper introduced the design and implementation of its own hardware platform but provided a dataset with all measurements made on the system, to allow third parties to evaluate the work even though they do not have access to the same hardware.

Without surprise simulations are not used in measurement papers. We notice that the ICN community seems to have a dedicated toolbox of simulators and system implementations, which is much less pronounced for the other communities. This is not surprising, as the ICN community worked on real-world implementations right from the beginning, and the number of default ready-to-use software is low.

When it comes to datasets, we observe that all communities rely on well-known datasets. This shows the importance to release data and to make them publicly available, so that other researchers can use them. However, it is also important to give more explicit incentives. This approach is successfully followed by the measurement community. IMC gives a *community contribution award* that “recognizes a paper with an outstanding contribution to the community in the form of a novel dataset, source code distribution, open platform, or other noteworthy service to the community [21].”

Networking community relies on existing infrastructures to perform research, either in testbeds or on measurement platforms. For example, Internet measurement studies leverage multiple vantage points to improve visibility on the measured data and strengthen the conclusiveness of the analysis. For the ICN community, we also notice a general usage of testbeds and additional infrastructures. This emphasizes the need to provide high quality testbeds, not only to allow external users to conduct their own experiments but also to allow comparison of solutions by using the exact same infrastructure in multiple studies. In our survey, architectural publications often focused on system aspects and thus only needed commodity hardware deployed in small settings, instead of large testbeds. It is important to note that many papers still use their own infrastructure or testbed. In particular in many measurement papers, data is based on real-world infrastructure but this infrastructure is private and not publicly accessible.

Finally, even though self-assessment of reproducibility is highly subjective and has the potential of being biased, two researchers admitted that their papers were hardly reproducible. One study needs a specific testbed; the other study consistently crawls websites. Nevertheless, researchers in the measurement and simulation domains are much more confident in the ability to reproduce their work, compared to other researchers. The least confident researchers are those who worked with complex platforms or testbeds.

## 4. DISCUSSION

*Caveat.* Drawing final conclusions and recommendations based on a limited dataset such as ours is always a sensitive exercise. Nevertheless, in the following we discuss some

recommendations that we believe are reasonable, based on what we learned while doing our brief artifact meta-analysis.

*Storage of Data.* All of the papers for which we received an answer provide information to help readers to reproduce the results of those papers. In general, papers provide links to webpages that contain some of the tools or data that have been used. Interestingly, less than 20% of researchers store artifacts on their personal or project website, instead they use popular public code platforms such as GitHub. We already noticed that artifacts links were broken for four papers, and thus argue that well-maintained platforms such as the ACM Digital Library [52] should be preferred to guarantee the durability data access, together with a snapshot of the status of the artifacts at publication time.

*Completeness of Tools.* Only the minority of papers (i.e., three publications) provide scripts to produce figures or compute numerical data, which is presented in the papers. In the majority of the papers, artifacts do not cover 100% of the results in the paper. That is mostly due to the fact that papers often require specific hardware or testbeds, which challenges self-contained artifacts. Remarkably, three papers provide the virtual image of the environment they used to produce the results of the paper.

*Research Cultures.* Considering research by field, we can clearly identify that for work on network architectures researchers prefer to modify real systems (e.g., the Linux kernel) while measurement work massively rely on datasets.

In IoT, wireless, and optical networking, researchers often have to rely on specific hardware and testbeds but they make sure to specify clearly the type of environment they used, which is much less clearly documented for researchers working on network architectures. Publications from the measurement community massively use well known public datasets and public measurement platforms and tend to make their collected data available to everyone. However, there is also a larger set of measurement papers that use confidential data, which prevents publication of (at least) raw data. When it comes to datasets, we have to make the distinction between raw and aggregated data. In case of raw data, the dataset contains the information as produced directly by the authors of the dataset (e.g., delay measurements) while aggregated data are the result of some processing meaning that some information is lost between the actual information and the one that is in the dataset. To allow the best usage of data one would recommend to always provide at least the raw data. Nevertheless, we noticed that a small fraction of papers makes only aggregated data or partial datasets publicly available, instead of their raw data. In particular, in large measurement projects, data is often provided by companies. These data are usually confidential because of business reasons or cannot be shared easily because of size.

As seen by categories, we could say that when it comes to developing new network architectures, researchers emphasize on implementing real systems and make sure their code is available, putting slightly aside the actual data used in their evaluation or the precise description of their evaluation environment. On the contrary, measurement papers insist on the data they used more than on the tools themselves. Finally, when it comes to system relying on specific

hardware, researchers generally provide fair description of the hardware and software they used.

Finally, a large number of papers in our survey used already existing public datasets and testbeds or infrastructures. This observation emphasizes the importance of making high quality datasets and testbeds available to the community. More generally, we see that most papers are built on top of artifacts made by other researchers in other papers, which shows the importance for researchers to provide their artifacts but also to keep them available for long period of time as we noticed that some researchers use tools and testbeds that were produced almost a decade ago.

## 5. CONCLUSION

Artifacts are indeed reused by others. Providing a detailed description which explains the data set and the usage of tools pays off. We thus encourage the SIGCOMM community to work further on releasing their artifacts and to highlight examples of good reproducibility and artifact availability. Forming an Artifact Evaluation Committee can be a first step in this direction. The artifacts produced by our survey are available online [55]. In addition, we are in the process of adding the discovered artifacts to *FindResearch.org* [51] as encouraged by Christian Collberg during his keynote at the ACM SIGCOMM 2017 Reproducibility Workshop. This editorial may serve as a starting point to build an Artifact Evaluation Committee for the SIGCOMM community.

## Acknowledgments

First and foremost, we thank the researchers that participated in this survey, as well as the TPC chairs who supported our efforts to contact the authors of the conferences addressed in this publication.

The authors gratefully acknowledge Olivier Bonaventure, who suggested to conduct this survey.

This work was partly funded by several organizations: the French National Research Agency (ANR) through the program “Investments for the Future” (ANR-11-LABX-0031-01) and the project REFLEXION (ANR-14-CE28-0019); NewNet@Paris, Cisco’s Chair “Networks for the Future” at Telecom ParisTech (<http://newnet.telecom-paristech.fr>); in the framework of the CELTIC EUREKA project SENDATE-PLANETS (project ID C2015/3-1); and the German Ministry of Education and Research (BMBF) in the projects 16KIS0460K, 16KIS0528K, 16KIS0550K, and 03VP02041.

Any opinions, findings, or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of project partners. Authors alone are responsible for the content of this paper.

## References

### ACM SIGCOMM 2017

- [1] BASAT, R. B., EINZIGER, G., FRIEDMAN, R., LUIZELLI, M. C., AND WAISBARD, E. Constant time updates in hierarchical heavy hitters. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [2] CHO, I., JANG, K., AND HAN, D. Credit-scheduled delay-bounded congestion control for datacenters. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).

- [3] HOLTERBACH, T., VISSICCHIO, S., DAINOTTI, A., AND VANBEVER, L. Swift: Predictive fast reroute. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [4] IORDANOU, C., SORIENTE, C., SIRIVIANOS, M., AND LAOUTARIS, N. Who is fiddling with prices? building and deploying a watchdog service for e-commerce. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [5] KASSING, S., VALADARSKY, A., SHAHAF, G., SCHAPIRA, M., AND SINGLA, A. Beyond fat-trees without antennae, mirrors, and disco-balls. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [6] LANGLEY, A., RIDDOCH, A., WILK, A., VICENTE, A., KRASIC, C., ZHANG, D., YANG, F., KOURANOV, F., SWETT, I., IYENGAR, J., BAILEY, J., DORFMAN, J., ROSKIND, J., KULIK, J., WESTIN, P., TENNETI, R., SHADE, R., HAMILTON, R., VASILIEV, V., CHANG, W.-T., AND SHI, Z. The quic transport protocol: Design and internet-scale deployment. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [7] LUCKIE, M., AND BEVERLY, R. The impact of router outages on the as-level internet. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [8] MAO, H., NETRAVALI, R., AND ALIZADEH, M. Neural adaptive video streaming with pensieve. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [9] MELLETTE, W., MCGUINNESS, R., ROY, A., FORENICH, A., PAPAN, G., SNOEREN, A., AND PORTER, G. Rotornet: A scalable, low-complexity, optical datacenter network. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [10] SAMBASIVAN, R., TRAM-LAM, D., AKELLA, A., AND P. STEENKISTE. Bootstrapping evolvability for inter-domain routing with d-bgp. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [11] ZAOSTROVNYKH, A., PIRELLI, S., PEDROSA, L., ARGYRAKI, K., AND CANDEA, G. A formally verified nat. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).
- [12] ZAVE, P., R.A.FERREIRA, ZOU, X., MORIMOTO, M., AND REXFORD, J. Dynamic service chaining with dysco. In *ACM SIGCOMM 2017* (Los Angeles, USA, August 2017).

### ACM CoNEXT 2017

- [13] CONINCK, Q. D., AND BONAVENTURE, O. Multipath quic: Design and evaluation. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [14] GILAD, Y., SAGGA, O., AND GOLDBERG, S. Maxlength considered harmful to the rpki. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).

- [15] LEONG, W., ZIXIAO, W., AND LEONG, B. Tcp congestion control beyond bandwidth-delay product for mobile cellular networks. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [16] LIM, Y.-S., NAHUM, E., TOWSLEY, D., AND GIBBENS, R. Ecf: An mptcp path scheduler to manage heterogeneous paths. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [17] RAZAGHPANAH, A., NIAKI, A., VALLINA-RODRIGUEZ, N., SUNDARESAN, S., AMANN, J., AND GILL, P. Studying tls usage in android apps. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [18] STEINMETZER, D., WEGEMER, D., SCHULZ, M., WIDMER, J., AND HOLLICK, M. Compressive millimeter-wave sector selection in off-the-shelf ieee 802.11ad devices. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [19] ZHANG, C., OUYANG, X., AND PATRAS, P. Zipnet-gan: Inferring fine-grained mobile traffic patterns via a generative adversarial neural network. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [20] ZHANG, R., AND PRENEEL, B. On the necessity of a prescribed block validity consensus: Analyzing bitcoin unlimited mining protocol. In *The 13th International Conference on emerging Networking EXperiments and Technologies – CoNEXT’17* (Seoul/Icheon, South Korea, December 2017).
- [21] ACM IMC 2017 – Call for Papers. <https://conferences.sigcomm.org/imc/2017/cfp/>.
- [22] AMANN, J., GASSER, O., SCHEITL, Q., BRENT, L., CARLE, G., AND HOLZ, R. Mission accomplished? https security after diginotar. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [23] ANDERSON, D., BEVAN, P., LANG, K., LIBERTY, E., RHODES, L., AND THALER, J. A high-performance algorithm for identifying frequent items in data streams. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [24] CHO, K. . Recursive lattice search: Hierarchical heavy hitters revisited. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [25] DE VRIES, W. B., DE O. SCHMIDT, R., HARDAKER, W., HEIDEMANN, J., DE BOER, P.-T., AND PRAS, A. Broad and load-aware anycast mapping using verfloeter. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [26] DEBLASIO, J., SAVAGE, S., VOELKER, G., AND SNODEN, A. Tripwire: Inferring internet site compromise. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [27] FONTUGNE, R., ABEN, E., AND PELSSER, C. Pinpointing delay and forwarding anomalies using large-scale traceroute measurements. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [28] GOODCHILD, B., CHIU, Y.-C., LU, H., HANSEN, R., CALDER, M., CHOFFNES, D., LLOYD, W., LUCKIE, M., AND KATZ-BASSETT, E. The record route option is an option! In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [29] KOTRONIS, V., NOMIKOS, G., MANASSAKIS, L., MAVROMMATIS, D., AND DIMITROPOULOS, X. Shortcuts through colocation facilities. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [30] LI, F., RAZAGHPANAH, A., KAKHKI, A., NIAKI, A., CHOFFNES, D., GILL, P., AND MISLOVE, A. lib-erate, (n): A library for exposing (traffic-classification) rules and avoiding them efficiently. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [31] MI, X., QIAN, F., ZHANG, Y., AND WANG, X. An empirical characterization of ifttt: Ecosystem, usage, and performance. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [32] MÜLLER, M., MOURA, G., DE O. SCHMIDT, R., AND HEIDEMANN, J. Recursives in the wild: Engineering authoritative dns servers. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [33] RÜTH, J., BORMANN, C., AND HOHLFELD, O. Large-scale scanning of tcp’s initial window. In *Proceedings of the 2017 Internet Measurement Conference* (New York, NY, USA, 2017), IMC ’17, ACM, pp. 304–310.
- [34] SNYDER, P., DOERFLER, P., KANICH, C., AND MCCOY, D. Fifteen minutes of unwanted fame: Detecting and characterizing doxing. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [35] SOMMERS, J., DURAIRAJAN, R., AND BARFORD, P. Automatic metadata generation for active measurement. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [36] VANAUBEL, Y., MÉRINDOL, P., PANSIOT, J.-J., AND DONNET, B. Through the wormhole: Tracking invisible mpls tunnels. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).

- [37] WANG, Z., CAO, Y., QIAN, Z., SONG, C., AND KRISHNAMURTHY, S. Your state is not mine: A closer look at evading stateful internet censorship. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- [38] ZHANG, Q., LIU, V., ZENG, H., AND KRISHNAMURTHY, A. High-resolution measurement of data center microbursts. In *ACM Internet Measurement Conference 2017 – IMC’17* (London, UK, November 2017).
- ## ACM ICN 2017
- [39] APONTE, O., AND MENDES, P. Now@ - content sharing application over ndn. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [40] ASCIGIL, O., SOURLAS, V., PSARAS, I., AND PAVLOU, G. A native content discovery for the information-centric networks. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [41] AUGÉ, H., CAROFIGLIO, G., ENGUEHARD, M., MUSCARIELLO, L., AND SARDARA, M. Virtualized icn (vicn): Towards a unified network virtualization framework for icn experimentation. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [42] BENEDETTO, J. D., WISSINGH, B., AND FU, X. Icn personalized global-scale testbed using gts. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [43] DYNEROWICZ, S., AND MENDES, P. Named-data networking in opportunistic networks. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [44] GÜNDOGAN, C., KIETZMANN, P., SCHMIDT, T., LENDERS, M., PETERSEN, H., WÄHLISCH, M., FREY, M., AND SHZU-JURASCHEK, F. Information-centric networking for the industrial iot. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [45] HAHM, O., BACCELLI, E., SCHMIDT, T., WÄHLISCH, M., ADJIH, C., AND MASSOULIÉ, L. Low-power internet of things with ndn and cooperative caching. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [46] KIETZMANN, P., GÜNDOGAN, C., SCHMIDT, T., HAHM, O., AND WÄHLISCH, M. The need for a name to mac address mapping in ndn: Towards quantifying the resource gain. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [47] KROL, M., AND PSARAS, I. Nfaas: Named function as a service. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [48] LIU, M., SONG, T., YANG, Y., AND ZHANG, B. A unified data structure of name lookup for ndn data plane. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [49] MOLL, P., JANDA, J., AND HELLWAGNER, H. Adaptive forwarding of persistent interests in named data networking. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- [50] ZURANIEWSKI, P., VAN ADRICHEM, N., RAVESTELJN, D., IJNTEMA, W., PAPADOPOULOS, C., AND FAN, C. Facilitating icn deployment with an extended openflow protocol. In *The 4th ACM Conference on Information-Centric Networking – ICN’17* (Berlin, Germany, September 2017).
- ## Others
- [51] A Catalog of Research Artifacts for Computer Science. <http://www.findresearch.org>.
- [52] ACM Digital Library. <https://dl.acm.org/dl.cfm>.
- [53] Artifact Review and Badging. <https://www.acm.org/publications/policies/artifact-review-badging>. Accessed: 2017-09-11.
- [54] Proceedings of ACM SIGCOMM Reproducibility Workshop. <https://dl.acm.org/citation.cfm?id=3097766>.
- [55] SIGCOMM 2017 Artifacts Survey data. <https://team.inria.fr/diana/reproducibility/>.
- [56] ADJIH, C., BACCELLI, E., FLEURY, E., HARTER, G., MITTON, N., NOEL, T., PISSARD-GIBOLLET, R., SAINT-MARCEL, F., SCHREINER, G., VANDAELE, J., ET AL. Fit iot-lab: A large scale open experimental iot testbed. In *Internet of Things (WF-IoT), 2015 IEEE 2nd World Forum on* (2015), IEEE, pp. 459–464.
- [57] AFANASYEV, A., MOISEENKO, I., ZHANG, L., ET AL. ndnsim: Ndn simulator for ns-3. *University of California, Los Angeles, Tech. Rep 4* (2012).
- [58] BACCELLI, E., HAHM, O., GUNES, M., WAHLISCH, M., AND SCHMIDT, T. C. Riot os: Towards an os for the internet of things. In *Computer Communications Workshops (INFOCOM WKSHPS), 2013 IEEE Conference on* (2013), IEEE, pp. 79–80.
- [59] SCHERB, C., SIFALAKIS, M., AND TSCHUDIN, C. Ccn-lite, 2013.
- [60] STAFF, R. Ripe atlas: A global internet measurement network. *Internet Protocol Journal* 18, 3 (2015).
- [61] TEAM, C. Center for Applied Internet Data Analysis, 2018.
- [62] TEAM, M. MATLAB - MathWorks, 2018.
- [63] ZHANG, L., ESTRIN, D., BURKE, J., JACOBSON, V., THORNTON, J. D., SMETTERS, D. K., ZHANG, B., TSUDIK, G., MASSEY, D., PAPADOPOULOS, C., ET AL. Named data networking (ndn) project. *Relatório Técnico NDN-0001, Xerox Palo Alto Research Center-PARC* (2010).