

Two Decades of IETF Affiliations: Evolution and Impact

Yangjun Zhang
Queen Mary University of London
London, UK
yangjun.zhang@qmul.ac.uk

Hugo Ramirez-Centeno
Queen Mary University of London
London, UK
h.e.ramirez-centeno@qmul.ac.uk

Stephen McQuistin
University of St Andrews
St Andrews, UK
sm@smcquistin.uk

Colin Perkins
University of Glasgow
Glasgow, UK
csp@csperkins.org

Vanja Karan
University of Vienna
Vienna, Austria
vanja.karan@univie.ac.at

Gareth Tyson*
Hong Kong University of Science
and Technology (GZ)
Guangzhou, China
gtyson@ust.hk

Ignacio Castro
Queen Mary University of London
London, UK
i.castro@qmul.ac.uk

Abstract

Standards developed by the Internet Engineering Task Force (IETF) are critical to the operation of the Internet. Understanding which organisations employ or support participants in the IETF is therefore critical to understanding the development of the Internet. In this paper, we present a longitudinal analysis of affiliation changes in the period 2001–2023, covering 73,764 individuals affiliated with 6,940 organisations. We show that organisational diversity first grew, before remaining broadly constant or even declining (e.g., meeting attendance and RFC authorship). Affiliations have gradually shifted away from North America-based organisations, with an increase in European and Asian organisations. We also observe that a change in affiliation has a short-lived positive impact on participant output and engagement, and increases the chance that an RFC is published, but slows its publication.

CCS Concepts

• **General and reference** → **Computing standards, RFCs and guidelines**; • **Networks** → **Public Internet**.

Keywords

Internet standards, IETF

*Also affiliated with Queen Mary University of London.



This work is licensed under a Creative Commons Attribution 4.0 International License.

ANRW '25, Madrid, Spain

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2009-3/25/07

<https://doi.org/10.1145/3744200.3744757>

ACM Reference Format:

Yangjun Zhang, Stephen McQuistin, Vanja Karan, Hugo Ramirez-Centeno, Colin Perkins, Gareth Tyson*, and Ignacio Castro. 2025. Two Decades of IETF Affiliations: Evolution and Impact. In *Applied Networking Research Workshop (ANRW '25)*, July 22, 2025, Madrid, Spain. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3744200.3744757>

1 Introduction

Technical standards are essential for the operation of the Internet. The IETF plays a pivotal role in the development and maintenance of these standards, bringing together people with a range of affiliations in an open forum where they can engage in discussion, contribute documents, and eventually reach consensus to publish standards documents [5].

Over the past two decades, the set of IETF participants, the organisations they are affiliated with, and the dynamics between them, have evolved. These changes have affected various aspects of the IETF process. For instance, publication of RFCs now takes longer, and the set of active organisations has changed [11, 16].

Understanding this is challenging. Participants in IETF engage as individuals, not as representatives of organisations, and affiliation data is incomplete and based on self-declarations. As a result, there is limited visibility into the role and relevance of organisational affiliation in IETF.

We examine IETF participation over time and address the following research questions. **RQ1:** What organisations are IETF participants affiliated with, and how does this evolve over time? **RQ2:** How often do participants change their affiliation, and how does it affect their behaviour within the IETF? **RQ3:** How do affiliation transitions impact the likelihood of RFC publication, and the time taken to do so?

To answer these questions, we combine data from multiple sources, including meeting registration records, document submissions, and mailing lists, to infer participant affiliations over the 2001–2023 period. Our dataset includes 73,764 unique participants, with 6,940 unique affiliations. We release the code and instructions.¹

To our knowledge, ours is the first large-scale study to systematically analyse the affiliation of IETF participants. Our findings reveal that organisational diversity within the IETF first increased, then became broadly constant and even decreased for meeting attendance and RFC authorship. We also find that overall participation has decreased. The total number of individuals participating in mailing lists and plenary meetings has been decreasing, and the individuals authoring drafts and RFCs initially increased, but has decreased in recent years. Moreover, we find that IETF activity has declined, with less email discussion and a smaller number of RFCs being published, although there are signs that the number of drafts is again increasing. We show that participant affiliation changes declined during periods of global crisis, including the 2008 financial crisis and the COVID-19 pandemic around 2020. Changes in affiliation are associated with a short-term increase in individual activity levels, but this effect tends to diminish over time. Finally, we find that affiliation transitions are associated with both a higher probability of RFC publication and a longer time span from initial submission to RFC publication.

2 Data and Methodology

The IETF develops standards through a consensus process that starts with the submission of an Internet-draft. Following discussion, the draft may be adopted by an IETF working group for development and eventual publication in the RFC series. There are several sub-streams of RFCs, for documents from the IETF, IRTF, and IAB, plus independent submissions, editorial, and legacy RFCs. Our analysis focuses on the IETF stream, which is the source of standards-related documents, but we use from all streams to help infer affiliations.

We collect IETF data from 2001–01–01 to 2023–12–31 (see Table 1) as follows: (i) **Documents**. We collect 30.4k drafts, comprising 29.5k from the IETF stream and 0.9k from non-IETF stream drafts, without considering the draft version, along with relevant metadata, including document names, version numbers, author affiliation and email address, and corresponding RFC for completed work, from the IETF Datatracker. (ii) **Mailing Lists**. We collect 2.3M emails from 1,139 mailing lists, including metadata (e.g., message ID, email address, date) and email content (e.g., subject, email content) from the IETF mail archive. (iii) **Meetings**. We extract 90.2k

Group	Total	Document	Mailing	Meeting
Record	2,638,157	351,681	2,196,208	90,268
With affil.	1,810,659	342,495	1,386,169	81,995
With affil. (%)	68.6%	97.4%	63.1%	90.8%
Individual	73,764	13,605	49,614	27,241
With affil.	30,096	12,443	12,675	20,328
With affil. (%)	40.8%	91.5%	25.5%	74.6%
Trans.	3,658	2,845	2,385	2,849
Trans. (%)	12.2%	22.9%	18.8%	14.0%
Trans.(-MA)	3,540	2,775	2,323	2,752
Trans.(-MA) (%)	11.8%	22.3%	18.3%	13.5%

Table 1: Dataset statistics.

meeting registration records relating to the 69 plenary meetings that took place during the period via the Datatracker.² For the purpose of analysis, we exclusively focus on registrations to the plenary meetings. To extract affiliations, we also leverage data from Hackathons and the Applied Networking Research Workshop (ANRW) –as some attendants only register to these.

Then, we exclude non-personal IETF official email accounts (e.g., ietf-announce@ietf.org) and 26k spam messages (<1%) using a Naive Bayes classifier (96.0% accuracy) trained on a combined dataset from SpamAssassin, EnronSpam [17], and LingSpam [2]. Finally, we integrate all raw data and construct a temporal dataset with the affiliations of the IETF participants. Table 1 summarises the collected data.

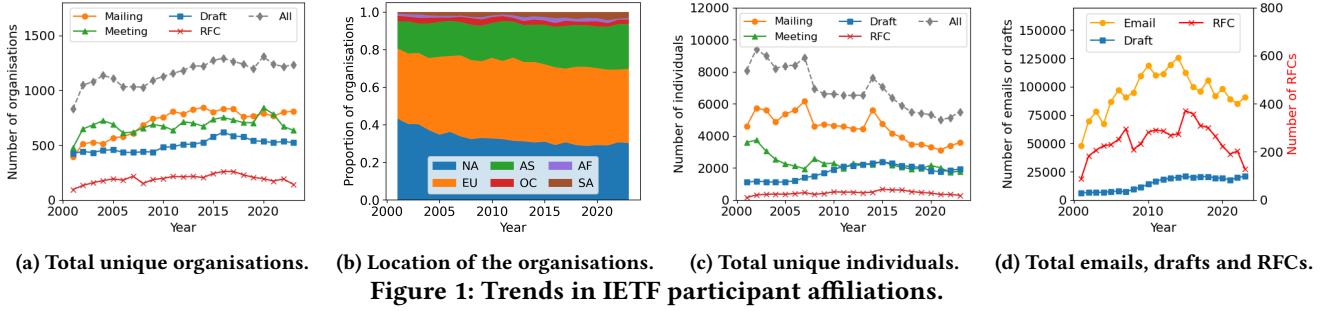
Identifying unique participants. We generate a global person ID, following McQuistin et al. [16], based on email address, name, and Datatracker profile. We consider the three main types of IETF activity including document authorship, email interaction, and meeting attendance. In total, we obtain 73,764 unique person IDs. For each person ID, we infer the participant’s affiliation and construct a time series of affiliation records, *i.e.*, a “record” refers to every observation of the affiliation of a participant. Overall, we obtain 2.6M records (*i.e.*, an average of 35.8 records per individual).

Inferring affiliations. Affiliation records are based on self-declared data from explicit (*i.e.*, draft authorship, meeting registration) and implicit statements (*i.e.*, email addresses).

We take the following steps to infer affiliations: (i) we extract affiliations as declared in the registration of each IETF plenary *meeting*. (ii) we extract the author-declared affiliations from Internet-drafts and RFCs, using the Datatracker or from the draft text if the affiliation is missing in the Datatracker. (iii) we *consolidate* the different observed variants of the same organisation into a canonical form (e.g., we combine “Cisco Systems” and “Cisco LTD” into “Cisco”)

¹Available at: https://sodestream.github.io/data-docs/ietf_docs.html and <https://github.com/sodestream>.

²<https://datatracker.ietf.org> and <https://mailarchive.ietf.org>



through a semi-automated process combining Levenshtein string similarity and expert validation. (iv) we infer participant affiliations from email domain suffixes in their mailing list posts. We first map email suffixes to affiliations based on the records for which we have the name of the organisation and the email (e.g., we have a participant that has declared to be employed by Cisco and that uses a name@cisco email). Then we associate individuals for which we have email but no information on their affiliation. In doing this, we consider suffixes for known affiliations and ignore shared hosting providers (e.g., gmail).

Constructing affiliation timelines. At the end of this process, we obtain a timeline of the observed affiliation records for IETF participants. These timelines naturally have gaps (e.g., a participant was not active for a period of time), which we fill in via interpolation, i.e., if an individual has the same affiliation at two records within a three-month window, we impute the same affiliation for the intermediate period.

We identify 6,940 unique affiliations, and 426 individuals with multiple simultaneous affiliations, e.g., “Cisco” and “Keio University.” Since it is not possible to know which/whether there is a main affiliation, we record these as unordered sets. **M&A.** We identify affiliation changes due to Mergers and Acquisitions (M&A) using data from Bureau van Dijk (BvD) Zephyr, widely used by researchers in finance [1, 22]. We find 54 M&A-induced affiliation changes involving 205 people.

Region and affiliation type. We also label affiliations as academic or non-academic. We define “academic affiliations” as those whose names include “University,” “Institute,” “Center,” “Association,” or “College,” or whose email domains are identified as educational (<https://github.com/JetBrains/swot>). We also associate organisations to regions (North America, Asia, Europe, South America, Oceania, Africa) depending of where it is headquartered according to Wikipedia.

Topic modelling. We compute topics for both email discussions (subject and email) and documents (title and abstract).

We employ a BERT-based topic model using the pre-trained all-mpnet-base-v2 embedding model to capture sentence-level semantics. For dimensionality reduction we employ UMAP followed by HDBSCAN clustering. To reduce topic

granularity, we apply K-means clustering (with $k = 2000$ selected via the elbow method) on the HDBSCAN outputs. We obtain a final coherence score of 0.967.

Ethics. We use only public data, following the IETF privacy and data usage policies, with no access to private data. We consulted with IETF leadership to ensure compliance. Data is securely stored and used solely for research.

3 Affiliations of IETF Participants (RQ1)

We now examine the evolution of the IETF from the perspective of the affiliation of its participants. Here we complement earlier work that studied similar questions from the perspective of individual participants [4, 11, 16].

Organisational involvement. We illustrate the involvement of the different organisations in Figure 1a, which shows the number of unique participant affiliations. The number of participating organisations has increased over time, however, there are relevant differences across the different modes of participation: (i) the diversity of affiliations for participants in the IETF meetings increased in early years, and then stabilised; (ii) the number of affiliations in the email discussions increased, then stabilised; (iii) the number of different affiliations in the authorship of drafts and RFCs first grew to then decrease. Meeting participation was on an upward trajectory that peaked around 2020. The increase is substantial and likely a result of the remote participation approach introduced as a result of COVID-19. Later years, where hybrid participation became common, have seen a decrease in the number of participants.

Regional representation has also varied. Figure 1b shows the headquarter location of the organisations. As observed in prior work [16], Asia and Europe have gained presence with North America declining in prominence. The IETF increased the number of meetings in Europe, and added the Asia-Pacific region to its regular rotation, starting in the early 2000s.

Interestingly, while we observe a relatively stable number of unique affiliations, the total number of participants is clearly decreasing. Figure 1c presents the total number of individuals (i.e., regardless of their affiliation) engaged in the




Figure 2: Affiliation activity and concentration analysis (based on distinct draft).

different modes of participation (*i.e.*, for drafts and RFCs we refer to authors). This trend towards a decreasing number of participants has been reversed in recent years everywhere except in terms of the number of RFC authors. As the time required to produce new RFCs is increasing [16], it is possible that there is a lagged effect and the trend will reverse.

Finally, we examine the total number of emails, drafts, and RFCs in Figure 1d. This shows a clear decrease in the number of emails and RFCs published. COVID-19 might be partially responsible for the recent sharp decrease in the number of RFCs published, since there was a decrease in new work started during the initial phase of pandemic. The downward trend, however, started in the mid-2010s. The trend might partially reverse in the future as the number of drafts seems to be increasing and some will likely become RFCs.

We observe growing participation of smaller organisations relative to larger ones. We categorize the participating organisations with regards to the size of their participation in the IETF as small, medium, or large based on the 33rd and 66th percentiles of the number of unique participants per organisation. We observe that the share of small organisations increased from 61.5% in 2001 to 71.3% in 2023. In contrast, medium-sized organisations declined from 30.2% to 23.8%, and large organisations decreased from 8.2% to 4.9%.

Activity per organisation. We now consider activity of participants when categorised by their affiliation. Figure 2a shows the median number of documents and emails per organisation. Since drafts are versioned, we exclusively consider the last available version of a given draft. Note that this inflates the last year of the series, and we exclude it from trend analysis. The number of documents per affiliation has slightly increased while the number of emails per organisation grew and then decreased, remaining stable in the last few years. Figure 2b shows the mean number of authors, organisations, and regions per document. We observe a clear increase from 2005, which has recently stabilized. Inter-regional collaboration is highest between North America and Europe, followed by North America-Asia and Europe-Asia. Other continents show less collaboration.

To examine the level of concentration in the different modes of participation, we calculate the Gini coefficient over time. A Gini coefficient close to one, indicates that participation is dominated by a small number of affiliations, whereas the opposite is true when it is close to zero. We obtain values of 0.886 for mailing list, 0.813 for individual drafts, 0.774 for meeting, and 0.767 for WG adopted drafts. These values indicate that organisations are most concentrated in mailing lists and least concentrated in WG-adopted drafts. While the concentration levels for mailing lists remain relatively stable over time (Figure 2c), the concentration in meetings initially increases and then declines. The concentration in drafts remains stable overall, with a recent trend toward decreasing concentration. The results suggests that although a small number of affiliations lead draft submissions and mailing list discussions, the final WG adopted drafts increasingly incorporate contributions from a wider array of affiliations. A possible reason is that dominant organisations may propose many drafts that are ultimately not adopted by WGs.

4 Changes in Participant Affiliation (RQ2)

Individual affiliation transition trend. To assess how often individuals change their affiliation, that we describe as an “affiliation transition”, Figure 3a presents the annual percentage of participants who change affiliation (*i.e.*, transition rate). The analysis includes all individuals, including those where an affiliation change is due to M&A.

We observe clear drops in the rate of affiliation changes coinciding with known economic downturns (*e.g.*, the dot-com recession in 2002, the great recession of 2008-2009, and the one of the COVID-19 pandemic in 2020). This is a common behaviour where labour markets become unstable in periods of crisis [6, 14, 15]. In 2006 and 2016, M&A were responsible for a substantial number of affiliation changes, with 81 individuals affected by acquisitions in 2006, notably the merger of Alcatel and Lucent to form Alcatel-Lucent, and 48 individuals affected in 2016 primarily due to Nokia’s acquisition of Alcatel-Lucent.

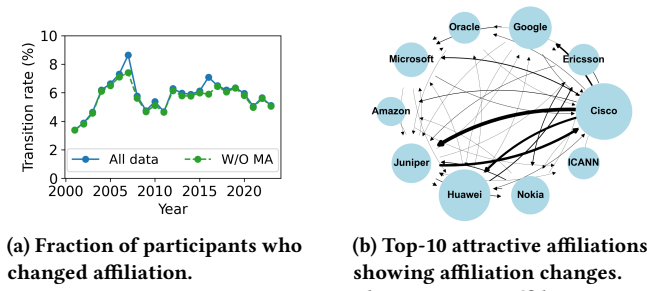


Figure 3: Transition rates and attractive affiliations.

To better understand affiliation changes, we categorize the active length of an affiliation and size of the organisation to which a participant is affiliated. We divide active length into three groups using 8-year bins, individuals with 1–8 active years are classified as short, 9–16 years as medium, and 17–23 years as long. We classify organisation size into small, medium, and large based on the 33rd and 66th quantiles, as shown in the previous section. We then perform T-tests on the PageRank scores of affiliations across groups, *i.e.*, location, size, active length. The PageRank metric captures how often organisations attract people from other organisations, particularly highly connected ones, and thus serves as a proxy for organisational attractiveness. Participants often move toward organisations in North America ($T=3.32$, $p<0.001$ vs. Europe; $T=2.08$, $p=0.04$ vs. Asia; and $T=2.99$, $P=0.003$ vs. other continents), that have had a long active length in the IETF, respectively ($T=6.05$, $p<0.001$ vs. medium, $T=7.44$, $p<0.001$ vs. short) and large number of IETF participants ($T=3.67$, $p=0.004$ vs. medium and $T=4.46$, $p<0.001$ vs. small). Figure 3b shows the ten most attractive affiliations based on PageRank. Arrows represent affiliation transitions between them, with edge widths proportional to the number of transitions.

Honeymoon-hangover Effect. The honeymoon-hangover effect is a phenomenon where individuals’ job satisfaction, engagement, or other behavioural metrics increase when they first enter their new job, and decrease later on [30].

To explore this, we measure participation on a given year by counting the number of documents authored and emails sent in that year for individuals with affiliation transitions. For each individual and each affiliation change, we then compute metrics over four time intervals: one year prior to the change $[t-1, t]$, six months following the change $[t, t+0.5]$, as well as one year $[t, t+1]$ and two years $[t, t+2]$ after the change. Figure 4 shows the honeymoon-hangover effect in the IETF. The dots are the average of the total number of documents and emails. Blue shows the honeymoon effect (*i.e.*, the increase in activity immediately after the affiliation change), and gray shows the hangover effect (*i.e.*, the decline in activity after the affiliation change). We observe that IETF

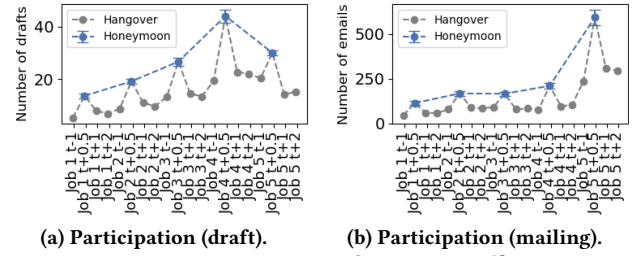


Figure 4: Honeymoon-hangover effect.

participants indeed experience a substantial rise in their engagement soon after changing that affiliation but that the increase is short-lived. The honeymoon effects increase for the first five jobs except for document participation. For document participation, after four job transitions, the honeymoon effect after the current affiliation transition does not guarantee an increase compared to the last affiliation transition. The reason might be that the individuals are more accustomed to changing jobs or reflect a decay in activity as they are closer to retiring. At the IETF level, recognizing this pattern can inform efforts to support continued participation through affiliation transitions.

5 Influence of Affiliation Transitions (RQ3)

To further understand the impact of affiliation transitions, we examine whether they affect the likelihood of RFC publication. Given that RFC publication has become increasingly time-consuming [16], we also study if affiliation transitions affect the time from initial submission to publication.

RFC publication prediction. To explore how affiliation transitions influence RFC publication, we train a logistic regression model where the dependent variable is the RFC publication, indicating whether a version of the draft is published as an RFC or not.

We first train a baseline model using the features described in §2 and excluding affiliation transition. We then train a model with the feature indicating the number of affiliation transitions for the authors. We use a train-test split of 8:2. As shown in Table 2, the model that includes affiliation transitions outperforms the baseline across all evaluation metrics. These consistent improvements suggest that affiliation transition contributes meaningful predictive power to the likelihood of RFC publication.

The coefficient of affiliation transition is positive and statistically significant ($\beta = 0.14$, $p = 0.007$), suggesting affiliation transition correlates with RFC publication. This suggests that affiliation mobility may serve as a facilitator of standardization success. This may be explained by increased visibility, renewed motivation, and improved access to resources. We also find that the coefficient for transitions from small or medium to large organisations is statistically significant ($p <$

	Precision	Recall	F1	Accuracy
w/o trans.	0.768	0.670	0.696	0.819
w/ trans.	0.772	0.682	0.708	0.823

Table 2: RFC publication predicting performance.




Figure 5: Content change after affiliation transition.

0.001), suggesting that such affiliation changes, which typically reflect increased resource availability, are associated with a higher likelihood of RFC publication.

Effect on the time to publish. To examine whether affiliation transitions influence the efficiency of the standardization process, we build a linear regression model with the number of days from the submission of the initial draft (version '00') to the latest draft version for drafts that were eventually published as RFCs as the dependent variable. The main explanatory variable is the number of affiliation transitions experienced by individuals prior to the last draft. We include several control variables to account for potential confounders (§2). The coefficient for affiliation transitions is positive and statistically significant ($\beta = 0.25$, $p < 0.001$), indicating that drafts with a greater number of author affiliation transitions during the drafting process take longer to reach RFC publication. A shift in technical focus is a possible explanation for the longer drafting time. To prove the assumption, we build a Difference-in-Difference (DiD) model [9]:

$$Y = \beta_0 + \beta_1 \text{Treatment} + \beta_2 \text{Post} + \beta_3 (\text{Treatment} \times \text{Post}) + \epsilon,$$

where $\text{Treatment} = 1$ for records with affiliation transitions, $\text{Treatment} = 0$ for records without transitions; and $\text{Post} = 1$ for records in the 0.5-, 1-, or 2-year window after the transition, $\text{Post} = 0$ for the matching window before transition. The coefficient $\hat{\beta}_2$ is negative and significant, implying that authors who remain in the same affiliation exhibit topic change over time. The interaction term $\hat{\beta}_3$ is negative and significant, indicating that individuals who change affiliations diverge even further from their previous output than do non-transition group. The results suggest that affiliation transition induces an additional decline in topic similarity, confirming a causal effect of affiliations transition on topic divergence, with time trend controlled.

Interestingly, affiliation transitions are associated with both a longer publication timeline and a higher likelihood of eventual publication. This suggests that while transitions may introduce delays, possibly due to technical realignment, as Table 5 suggests, they may also provide access to greater support, new collaborators, or strategic resources that increase the chances of standardization success.

6 Related Work

Longitudinal affiliation datasets usually rely on *e.g.*, social media, online professional network (OPN) websites, or third-party data collection [21]. A post-PhD career dataset with 10k individuals first analysed affiliation transition in CS [23]; Microsoft Academic Graph (MAG), with affiliation transitions for 10k scientists [28], is also widely used. Most works investigate employees' opinion about the organisation they are affiliated with [3, 10, 13]. Our dataset contains the output (documents) and interactions (emails) of IETF participants, and we study this instead of their opinions. Using similar metrics (*e.g.*, attractiveness, influence) as in earlier work [7, 8, 12, 19, 20, 23–27], we find similar results Zhou et al. [29]: a new job is related to positive outcomes and this effect decreases over time. Niedermayer et al. [18] analyse IETF documents and identify which affiliations are more active between 2009 to 2015. McQuistin et al. [16] investigate author affiliations, identifying key characteristics and observing that IETF continued to attract new affiliations from 2001 to 2020. Khare et al. [11] find that individuals affiliated with influential organisations exhibit a higher rate of draft adoption. The focus of these studies, is however restricted to document authorship and looks at general IETF dynamics.

7 Conclusion

We studied the evolution of participation in the IETF, by considering the organisations its participants are affiliated with, integrating affiliation data across IETF documents, emails, meetings, and M&A events. To our knowledge, this is the first systematic study of affiliation trend and affiliation change in the IETF. We show how the IETF has become increasingly diverse in terms of the number of organisations involved, but at the same time, its activity has declined, with fewer meeting participants, emails, RFCs, and individuals authoring them. We also show how many IETF participants change affiliations during their participation, and discuss how this impacts their engagement and their output.

Acknowledgements

This work is supported in part by UK EPSRC (EP/S033564/1, EP/S036075/1), DNSRF, UKRI (EP/W032473/1), Guangzhou Science and Technology Bureau (2024A03J0684).

References

- [1] Yan Alperovych, Douglas Cumming, Veronika Czellar, and Alexander Groh. 2021. M&A rumors about unlisted firms. *Journal of Financial Economics* 142, 3 (2021), 1324–1339.
- [2] Ion Androutsopoulos, John Koutsias, Konstantinos V Chandrinos, George Paliouras, and Constantine D Spyropoulos. 2000. An evaluation of naive bayesian anti-spam filtering. In *Proceedings of the Workshop on Machine Learning in the New Information Age, 11th European Conference on Machine Learning*. ACM, Barcelona, 9–17.
- [3] Rajiv Bajpai, Devamanyu Hazarika, Kunal Singh, Sruthi Gorantla, Erik Cambria, and Roger Zimmermann. 2018. Aspect-sentiment embeddings for company profiling and employee opinion mining. In *International Conference on Computational Linguistics and Intelligent Text Processing*. Springer, 142–160.
- [4] Matthew Russell Barnes, Mladen Karan, Stephen McQuistin, Colin Perkins, Gareth Tyson, Matthew Purver, Ignacio Castro, and Richard G Clegg. 2024. Temporal network analysis of email communication patterns in a long standing hierarchy. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 18. 126–138.
- [5] Scott O. Bradner. 1996. The Internet Standards Process – Revision 3. RFC 2026. doi:10.17487/RFC2026
- [6] Olivier Coibion, Yuriy Gorodnichenko, and Michael Weber. 2020. *Labor markets during the COVID-19 crisis: A preliminary view*. Technical Report. National Bureau of economic research.
- [7] Le Dai, Yu Yin, Chuan Qin, Tong Xu, Xiangnan He, Enhong Chen, and Hui Xiong. 2020. Enterprise cooperation and competition analysis with a sign-oriented preference network. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 774–782.
- [8] Sana Elouaer. 2006. A social network analysis of interlocking directorates in French firms. *Observatoire e Politico-Economique des Structures du Capiralisme* (2006).
- [9] David Gamba, Yulin Yu, Yuan Yuan, Grant Schoenebeck, and Daniel M Romero. 2024. Exit Ripple Effects: Understanding the Disruption of Socialization Networks Following Employee Departures. In *Proceedings of the ACM Web Conference 2024*. 211–222.
- [10] Victoria Ikoru, Maria Sharmina, Khaleel Malik, and Riza Batista-Navarro. 2018. Analyzing sentiments expressed on Twitter by UK energy company consumers. In *2018 Fifth international conference on social networks analysis, management and security (SNAMS)*. IEEE, 95–98.
- [11] Prashant Khare, Mladen Karan, Stephen McQuistin, Colin Perkins, Gareth Tyson, Matthew Purver, Patrick Healey, and Ignacio Castro. 2022. The web we weave: Untangling the social graph of the IETF. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 16. 500–511.
- [12] Lun Li, Theodoros Lappas, and Rong Liu. 2024. Measuring employer attractiveness in diverse talent markets. *Decision Support Systems* 177 (2024), 114079.
- [13] Hao Lin, Hengshu Zhu, Junjie Wu, Yuan Zuo, Chen Zhu, and Hui Xiong. 2020. Enhancing employer brand evaluation with collaborative topic regression models. *ACM Transactions on Information Systems (TOIS)* 38, 4 (2020), 1–33.
- [14] Alex Jiahong Lu, Anna Hihool, Joey Chiao-Yin Hsiao, and Tawanna R Dillahunt. 2022. Emotional labor in everyday resilience: Class-based experiences of navigating unemployment amid the COVID-19 pandemic in the US. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–27.
- [15] Yannis Markovits, Diana Boer, and Rolf van Dick. 2014. Economic crisis and the employee: The effects of economic crisis on employee job satisfaction, commitment, and self-regulation. *European Management Journal* 32, 3 (2014), 413–422.
- [16] Stephen McQuistin, Mladen Karan, Prashant Khare, Colin Perkins, Gareth Tyson, Matthew Purver, Patrick Healey, Waleed Iqbal, Junaid Qadir, and Ignacio Castro. 2021. Characterising the IETF through the lens of RFC deployment. In *Proceedings of the 21st ACM Internet Measurement Conference*. 137–149.
- [17] Vangelis Metsis, Ion Androutsopoulos, and Georgios Paliouras. 2006. Spam filtering with naive bayes-which naive bayes?. In *CEAS*, Vol. 17. Mountain View, CA, 28–69.
- [18] Heiko Niedermayer, Daniel Raumer, Nikolai Schweltnus, Edwin Cordeiro, and Georg Carle. 2016. An analysis of IETF activities using mailing lists and social media. In *Internet Science: Third International Conference, INSCI 2016, Florence, Italy, September 12–14, 2016, Proceedings 3*. Springer, 218–230.
- [19] Richard J Oentaryo, Xavier Jayaraj Siddarth Ashok, Ee-Peng Lim, and Philips Kokoh Prasetyo. 2017. On analyzing job hop behavior and talent flow networks. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*. IEEE, 207–214.
- [20] Richard J Oentaryo, Ee-Peng Lim, Xavier Jayaraj Siddarth Ashok, Philips Kokoh Prasetyo, Koon Han Ong, and Zi Quan Lau. 2018. Talent flow analytics in online professional network. *Data Science and Engineering* 3 (2018), 199–220.
- [21] Chuan Qin, Le Zhang, Yihang Cheng, Rui Zha, Dazhong Shen, Qi Zhang, Xi Chen, Ying Sun, Chen Zhu, Hengshu Zhu, et al. 2023. A comprehensive survey of artificial intelligence techniques for talent analytics. *arXiv preprint arXiv:2307.03195* (2023).
- [22] Lauren Reiter. 2013. Zephyr. *Journal of Business & Finance Librarianship* 18, 3 (2013), 259–263.
- [23] Tara Safavi, Maryam Davoodi, and Danai Koutra. 2018. Career transitions and trajectories: A case study in computing. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 675–684.
- [24] Huang Xu, Zhiwen Yu, Jingyuan Yang, Hui Xiong, and Hengshu Zhu. 2016. Talent circle detection in job transition networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 655–664.
- [25] Le Zhang, Tong Xu, Hengshu Zhu, Chuan Qin, Qingxin Meng, Hui Xiong, and Enhong Chen. 2020. Large-scale talent flow embedding for company competitive analysis. In *Proceedings of The Web Conference 2020*. 2354–2364.
- [26] Yuxia Zhang, Klaas-Jan Stol, Hui Liu, and Minghui Zhou. 2022. Corporate dominance in open source ecosystems: a case study of OpenStack. In *Proceedings of the 30th ACM Joint European Software Engineering Conference and Symposium on the Foundations of Software Engineering*. 1048–1060.
- [27] Yuxia Zhang, Minghui Zhou, Klaas-Jan Stol, Jianyu Wu, and Zhi Jin. 2020. How do companies collaborate in open source ecosystems? an empirical study of openstack. In *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering*. 1196–1208.
- [28] Zhenyue Zhao, Yi Bu, and Jiang Li. 2021. Characterizing scientists leaving science before their time: Evidence from mathematics. *Information Processing & Management* 58, 5 (2021), 102661.
- [29] Ying Zhou, Min Zou, Chia-Huei Wu, Sharon K Parker, and Mark Griffin. 2023. A study of new labor market entrants’ job satisfaction trajectories during a series of consecutive job changes. *Journal of applied psychology* (2023).
- [30] Ying Zhou, Min Zou, Chia-Huei Wu, Sharon K Parker, and Mark Griffin. 2024. A study of new labor market entrants’ job satisfaction trajectories during a series of consecutive job changes. *Journal of applied psychology* 109, 2 (2024), 293.