RAIN

A Bio-Inspired Communication and Data Storage Infrastructure

Monti, Matteo; Rasmussen, Steen

Published in:
Artificial Life

DOI:
10.1162/ARTL_a_00247

Publication date:
2017

Document version
Final published version

Citation for pulished version (APA):
Monti, M., & Rasmussen, S. (2017). RAIN: A Bio-Inspired Communication and Data Storage Infrastructure. *Artificial Life*, *23*(4), 552-557. https://doi.org/10.1162/ARTL_a_00247

RAIN: A Bio-Inspired Communication and Data Storage Infrastructure

Matteo Monti[*,**]
University of Southern Denmark
École Polytechnique Fédéral de Lausanne
University of Bologna

Steen Rasmussen[*,†]
University of Southern Denmark
Santa Fe Institute

Abstract We summarize the results and perspectives from a companion article, where we presented and evaluated an alternative architecture for data storage in distributed networks. We name the bio-inspired architecture RAIN, and it offers file storage service that, in contrast with current centralized cloud storage, has privacy by design, is open source, is more secure, is scalable, is more sustainable, has community ownership, is inexpensive, and is potentially faster, more efficient, and more reliable. We propose that a RAIN-style architecture could form the backbone of the Internet of Things that likely will integrate multiple current and future infrastructures ranging from online services and cryptocurrency to parts of government administration.

Keywords: Distributed storage, privacy by design, Internet of Things, community ownership, sustainability

1 Background

Recently our physical technologies (e.g., the converging bio-, info-, nano-, and cognotechnologies) have started to advance beyond our social technologies (e.g., governance, laws, educational systems, and social norms). This rapidly growing gap generates challenges and opportunities within most areas of modern society [9], including privacy and security in cyberspace as well as environmental issues.

* Contact author.
** Center for Fundamental Living Technology, University of Southern Denmark; École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland; Complex Systems Group, University of Bologna, Bologna, Italy. E-mail: matteo.monti@msoftprogramming.com
† Center for Fundamental Living Technology, University of Southern Denmark; Santa Fe Institute, Santa Fe, NM 87501, USA. E-mail: steen@sdu.dk

The Internet was originally designed with robustness in mind, as a means to guarantee communications in times of war. Instead of focusing on the protection of central points of failures, its protocols allowed redundancy, self-repair, and self-organization: While single nodes can fail, and new nodes can be connected, the overall functionality of the network is guaranteed by a resilience rooted in ecology.

Despite the ecosystematic nature of the infrastructure of Internet services, they are becoming progressively more centralized, with fewer and fewer organizations in charge of managing information on a planetary scale, thus creating monopolies and raising significant issues of privacy, security, and democracy.

The Internet data storage services provided today violate privacy, are expensive, and come at a high environmental cost. Today more than 3% of the world's power consumption is attributed to data centers, with a CO_2 footprint surpassing that of global air traffic and a rapidly growing power consumption rate [1]. The high entrance cost to the data storage market creates monopolies, in that only the largest companies are capable of offering scalable, cost-efficient services (e.g., [11]).

2 Basic Design Concepts

This article summarizes the results and perspectives from [7], outlining how current Internet of Things (IoT) technology could enable further decentralization and a more bio-inspired, distributed paradigm not only for information delivery, but also for storage and processing. We offer preliminary results on the development of RAIN,[1] an alternative and potentially superior software backbone for storage of data in distributed networks.

Our network architecture offers a distributed file storage service that is faster, is more efficient and reliable, is more secure, offers privacy by design as well as community ownership, and is open source, scalable, and more sustainable and less expensive than the current, centralized paradigm.

Owned by the community of its users (e.g., citizens, businesses, and organizations), this network service will be lower cost, democratic, and designed to guarantee the privacy of the data it stores. Embedded in citizen-owned computing devices (e.g., inexpensive Raspberry Pis with flash drives), it is now possible to have cheap, energy-efficient, always online computing nodes in our homes and businesses. The RAIN network design leverages on the collective storage power of these devices: Every node will store parts of other nodes' data to guarantee redundancy and reliability, and an elegant cryptographic designed architecture will prevent unwanted access to the stored data.

Such a bio-inspired architecture offers redundancy, distributed control, error correction, self-repair, and obvious potential for autonomous adaptation (learning) in later versions, with no central point of failure or trusted third parties. Each node operates via local interactions with a limited set of other nodes that it does not need to trust a priori.

Similarly to blockchains eliminating many banks as middlemen for standard financial transactions (see, e.g., [8]), RAIN could disrupt current cloud storage facilities and eliminate the need for centralized data centers overseeing many market segments, offering a solution to growing concerns about personal privacy and democracy, stemming from increasingly pervasive and unnecessary surveillance by private and public organizations.

Our preliminary results (see [7]) include: a feasibility study, where we quantitatively estimate the reliability of a decentralized storage network in comparison with a data-center-based architecture; provide an overview of the main security challenges to developing this infrastructure; identify how

1 RAIN is a metaphor for what comes after the clouds.

new security mechanisms can be designed to guarantee data security—even against government-grade attackers; and offer an outlook for the potential applications of this network to a broader spectrum of services than cloud-based storage can provide.

In particular, we estimate network size requirements to port to a distributed paradigm: a content delivery network for public Web content (with a more in-depth study of the resources needed to host Wikipedia); an end-to-end encrypted, peer-to-peer messaging platform; a social network without centralized control, free from targeted advertising and surveillance; and a distributed search engine (we argue on the one hand the high performance of distributed Web crawling, and on the other the limited querying capabilities of a high-latency, peer-to-peer distributed database).

Finally, we discuss how high-uptime, low-power nodes enable the development of a highly efficient cryptocurrency based on authenticated hash tables (see, e.g., [6]) instead of blockchains, with logarithmic space, time, and communication complexity, and no need for proof-of-work-based mining for initial currency distribution. Thus such a cryptocurrency should be significantly more memory and energy efficient than blockchains.

3 How is RAIN Different?

RAIN lies at the intersection of the well-explored field of decentralized and distributed systems security and that of low-cost, pervasive networked computation. A paradigm shift from software instances running on personal computers to permanently online, but still unreliable, dedicated low-energy nodes may seem minor, but allows us to ground our architecture design in far more stringent reliability assumptions. Until a few years ago, only expensive, dedicated servers could guarantee such reliabilities.

Peer-to-peer file distribution, for example, is a well-known technology that today aids the distribution of open-source operating systems and creative commons media. The challenge of translating this download-only paradigm to one where data can be reliably *uploaded* to a network of nodes has so far been undertaken only by storage-trading projects (like Storj; see [12]) that make stronger reliability assumptions than those offered by personal computers.

As we have seen, globally used, blockchain-based cryptocurrencies and distributed ledgers exist today, but limited uptime assumptions force their architectures to a paradigm where consensus needs to be verifiable asynchronously. This often leads to CPU-intensive security procedures and limited overall transaction throughput. Our preliminary results show that using proofs of space (see, e.g., [3]) on semi-reliable devices, we can guarantee security at a significantly smaller hardware and energy cost.

Finally, as is often seen in biological systems, subsystems integration and multipurpose interaction play a significant role in RAIN. This is in contrast, for example, to the bitcoin mining process. It has to run dedicated hardware whose sole purpose is to solve costly and otherwise useless computational challenges. RAIN, a community-owned, distributed storage network, could make use of its spare storage space to collectively guarantee its own security, while offering a variety of useful services to the community of its users.

4 RAIN Architecture Highlights

Optimal erasure codes (e.g., [10]) exist based on polynomial oversampling and interpolation that allow us to organize an S-byte-long string of data in $K = rN$ (with $r > 1$) blocks of size S/N, so that S can be recovered by any N of those blocks. The design of our network (which leverages only local, scalable interactions between the nodes and requires no mediation of a central decision-making authority) organizes embedded computers, persistently connected to home-grade Internet connections in *villages* of size K. Within the same village, each node trades its storage space with the others, offering to store redundancy blocks for the other nodes in exchange for space to store its own in a peer-to-peer fashion.

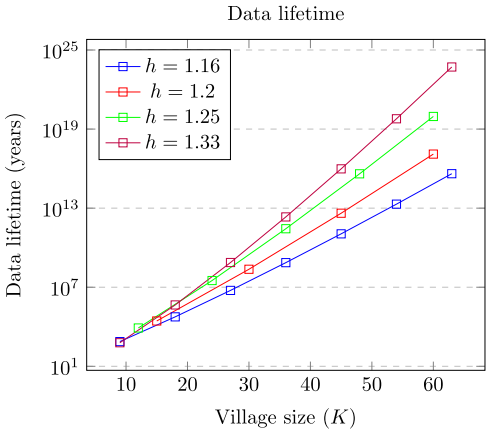

Figure 1. Data lifetime (i.e., expected time before a piece of data stored by RAIN becomes permanently unavailable due to random hardware failures; the higher the better) for different values of the recovery ratio h (i.e., the fraction of chunks the network affords to lose before triggering a recovery procedure for a piece of data; the lower the better, as recovery procedures are network-intensive), as a function of the village size K (i.e., the number of physical nodes that simultaneously contribute to the redundancy of each piece of data; the lower the better, as distributed bookkeeping increases in complexity with the number of nodes involved). Here we have $r = 1.5$ (i.e., the ratio between space occupied and actual file size; the lower the better) and $Z = 100$ GB (i.e., the amount of space each node contributes to each redundancy pool; the larger the better, as it reduces the metadata overhead).

A village-wide distributed ledger is kept between the villagers to keep each file under real-time control. Nodes securely monitor each other's data availability (which can be done with logarithmic time and communication complexity using Merkle tree hashes; see [5]) to readily detect failures. When a node experiences an unrecoverable failure (e.g., hardware failure or permanent disconnection),

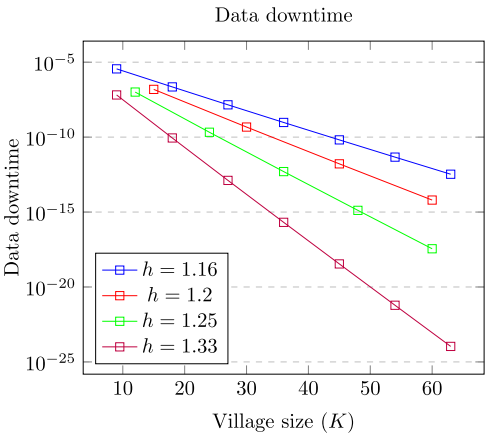

Figure 2. Data downtime (i.e., expected fraction of time that a piece of data stored by RAIN is unavailable due to temporary network malfunctioning; the lower the better) for different values of the recovery ratio h (i.e., the fraction of chunks the network affords to lose before triggering a recovery procedure for a piece of data; the lower the better, as recovery procedures are network-intensive), as a function of the village size K (i.e., the number of physical nodes that simultaneously contribute to the redundancy of each piece of data; the lower the better, as distributed bookkeeping increases in complexity with the number of nodes involved). Here we have $r = 1.5$ (i.e., the ratio between space occupied and actual file size; the lower the better) and $Z = 100$ GB (i.e., the amount of space each node contributes to each redundancy pool; the larger the better, as it reduces the metadata overhead).

the village signals new nodes to join it. When the availability of data reaches a threshold value $T = hN$ (with $1 < h < r$), a self-repairing, distributed recovery procedure is triggered and new redundancy blocks are generated.

Using experimental data for hard disk drive (HDD) and solid state drive (SSD) failure rates (see [2] and [4]) and gathering experimental data on home-grade Internet connection uptime and speed (see [7]), from the above model we could determine the expected lifetime L^* and the expected downtime (i.e., the fraction of time something is unreachable due to temporary malfunctioning of its connection) d^* of a file in our network.

Figures 1 and 2 show the expected data lifetime and downtime for a file stored by a village, determined by our analytical model, as a function of its size K and its recovery ratio $h = T/N$. Here each node is contributing with $Z = 100$ GB of storage space. Note how, without having to affect the storage ratio r (which determines how efficiently data is stored on the network), we can make the data lifetime arbitrarily large, and the data downtime arbitrarily small, just by changing the size of the village.

5 Discussion

Our proposed bottom-up, low-energy, bio-inspired technology offers a more cooperative, civic-centered ownership structure to preserve critical aspects of online privacy as well as freedom from the steering power of today's invasive marketing, behavior manipulation, and high-financed data attackers. We have demonstrated, for example, the feasibility of our proposed architecture, based on Solomon-Reed redundancy, in which 36 nodes provide an expected data lifetime of the same order of magnitude as the age of the Earth [7].

Additionally, RAIN could support the development of communitarian services, including telecommunication, content delivery, cryptocurrency, and distributed administration (nation-state and regional governmental), which currently are services managed in a centralized manner through trusted third parties [7]. Implementation of a RAIN-style architecture could thus distribute the power from global centralized trusted third parties to local citizens and businesses, while at the same time presumably reducing the significant energy requirement and resulting CO_2 burden of centralized data storage.

Acknowledgment

We are grateful for constructive suggestions from Alex Penn and Piper Stover, and we thank Lucinda Voldsgaard for proofreading the manuscript. Partial financial support was provided by the European Commission-sponsored SYNENERGENE project.

References

1. Bawden, T. (2016). Global warming: Data centres to consume three times as much energy in next decade, experts warn. *The Independent*, 23 January.
2. Beach, B. (2013). How long do disk drives last? Backblaze, Inc. (www.backblaze.com).
3. Dziembowski, S., Faust, S., Kolmogorov, V., & Pietrzak, K. (2015). Proofs of space. In *Advances in cryptology—CRYPTO 2015* (pp. 585–605).
4. Gasior, G. (2015). The SSD endurance experiment: They're all dead. *The Tech Report* (techreport.com).
5. Merkle, R. C. (1988). A digital signature based on a conventional encryption function. In *Advances in cryptology—CRYPTO 1987* (p. 369).
6. Miller, A., Hicks, M., Katz, J., & Shi, E. (2014). Authenticated data structures, generically. *SIGPLAN Notices*, *49*(1), 411–423.
7. Monti, M., Rasmussen, S., Moschettini, M., & Posani, L. (2017). *An alternative information plan* (Working paper). Santa Fe Institute.

8. Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. https://bitcoin.org/bitcoin.pdf.
9. Rasmussen, S. (2016). The BINC manifesto: Technology drive societal changes, science, policy & stakeholder engagement. In *Proceedings of ALife XV: The Fifteenth International Conference on the Simulation and Synthesis of Living Systems, Artificial Life* (pp. 53–54). Cambridge, MA: MIT Press.
10. Reed, I. S., & Solomon, G. (1960). Polynomial codes over certain finite fields. *Journal of the Society for Industrial and Applied Mathematics*, *8*, 300–304.
11. Scott, M. (2017). What U.S. tech giants face in Europe in 2017. *The New York Times*, 1 January.
12. Wilkinson, S., Boshevski, T., Brandof, J., Prestwich, J., Hall, G., Gerbes, P., Hutchins, P., & Pollard, C. (2016). Storj—a peer-to-peer cloud storage network. https://storj.io/storj.pdf.