Trusted Data: A New Framework for Identity and Data Sharing, MIT Press 2019

MIT Open Algorithms

THOMAS HARDJONO, MIT Connection Science
ALEX PENTLAND, MIT Connection Science

1 OPEN ALGORITHMS: A NEW PARADIGM FOR SHARING INSIGHTS

The Open Algorithms (OPAL) paradigm seeks to address the increasing need for individuals and organizations to share data in a privacy-preserving manner [1]. Data is crucial to the proper functioning of communities, businesses and government. Previous research has indicated that data increases in value when it is combined. Better insight is obtained when different types of data from different areas or domains are combined [2]. These insights allows communities to begin addressing the difficult social challenges of today, including better urban design, containing the spread of diseases, detecting factors that impact the economy, and other challenges of the data-driven society [3, 4].

Today there are a number of open challenges with regards to the information sharing ecosystem:

- *Data is siloed*: Today data is siloed within organizational boundaries, and the sharing of raw data with parties outside the organization remains unattainable, either due to regulatory constraints or due to business risk exposures.
- *Privacy is inadequately addressed*: The 2011 World Economic Forum (WEF) report [5] on personal data as a new asset class finds that the current ecosystems that access and use personal data is fragmented and inefficient.
  For many participants, the risks and liabilities exceed the economic returns and personal privacy concerns are inadequately addressed. Current technologies and laws fall short of providing the legal and technical infrastructure needed to support a well-functioning digital economy. The rapid rate of technological change and commercialization in using personal data is undermining end-user confidence and trust.
- *Regulatory and compliance*: The introduction of the EU General Data Protection Regulations (GDPR) [6] will impact global organizations that rely on the Internet for trans-border flow of raw data. This includes cloud-based processing sites that are spread across the globe.

Thus, we are facing an interesting dilemma with regards to data-driven decision making for individuals, organizations and communities. On one hand, individuals, organizations and communities need "access to data" in order to perform computations as part of decision-making. The promise is that better insights can be obtained by combining data from different domains in interesting and innovative ways. On the other hand, however, there is considerable risk to privacy when "data is shared" across entities. And the WEF report [5] clearly points to inadequate care given today to personal data – with evidence abound with regards to theft or misuse of personal data reported in the media [7–9]. It is with this backdrop that open algorithms model is put forward as an alternative paradigm in which to view and treat data.

In this chapter we discuss the OPAL principles and put forward the architecture developed at MIT to implement these principles. We discuss the issue of authorization and consent in the context of "consent to execute a vetted algorithm" and contrast this to the prevailing interpretation of consent as being "consent to copy and move data". We also briefly put forward a basic model for multiple data providers to collaborate in a trust network founded on the principles of open algorithms.

Authors' addresses: Thomas Hardjono, MIT Connection Science, hardjono@mit.edu; Alex Pentland, MIT Connection Science, pentland@mit.edu.

2 OPEN ALGORITHMS PRINCIPLES

The concept of *Open Algorithms* (OPAL) evolved from several research projects over the past decade within the Human Dynamics Group at the MIT Media Lab, particularly the thesis work of Yves-Alexandre de Montjoye (now at Imperial College) and Guy Zyskind (now CEO and founder of Enigma.co). From various research results it was increasingly becoming apparent that an individual's privacy could be affected through the correlation of just small amounts of data [10, 11].

One noteworthy seed project was *OpenPDS* that sought to develop further the concept of personal data stores (PDS) [12–14], by incorporating the idea of analytics on personal data and the notion of "safe answers" as being the norm for responses generated by a personal data store.

However, beyond the world of personal data stores there remains the pressing challenges around how large data stores are to secure their data, safeguard privacy and comply to regulations (e.g. GDPR [6]) – while at the same time enable productive collaborative data sharing. The larger the data repository, the more attractive it would become to hackers and attackers. As such, it became evident that the current mindset of performing data analytics at a centralized location needed to be replaced with a new paradigm for thinking about data sharing in a distributed manner.

The following are the fundamental principles of open algorithms and the treatment of data:

- *Move the algorithm to the data*: Instead of pulling data from various repositories into a centralized location for processing, it is the algorithm that should be sent to the data repositories for processing there. The goal here is to *share insights* instead of sharing raw data.
- *Data must never leave its repository*: Data must never be exported from (or copied from) its repository. This is consistent with the previous principle and enforces that principle.
  Exceptions to this rule are when the user requests a download of their data, and when there is a legally valid court order to obtain a copy of the data.
- *Vetted algorithms*: Algorithms should be studied, reviewed and vetted by experts.
  The goal here is to provide all entities in the ecosystem with a better understanding and assessment of the quality of algorithms from the perspective of bias, unfairness and other possible unintended/unforseen side effects.
- *Default to safe answers*: The default behavior of data repositories when returning responses should be that of protecting privacy as the primary goal. This applies to individual as well as organizational data privacy.
  For aggregate computations data repositories should place additional filters on responses to detect and resolve potential privacy leakages. For subject-specific computations (i.e. about a specific individual or organization) data repositories should obtain explicit and informed consent from the subject. We believe this principle is consistent with Article 7 of the GDPR [6].

There are a number of corollary principles from the above principles that enhance the protection of data and therefore enhance privacy:

- *Data always in encrypted state*: Data must remain encrypted during computation and in storage.
  The notion here is that in order to protect data repositories from attacks and theft (e.g. theft by insider), data should never be decrypted. This implies that algorithms sent to a data repository must be executed by the repository on its encrypted data without first decrypting it into plaintext. We believe that in the future this principle will be crucial and unavoidable from the perspective of infrastructure cybersecurity.
  There are a number of emerging technologies – such as homomorphic encryption [15] and secure multi-party computation [16–18] – that may provide the future foundations to address this principle.

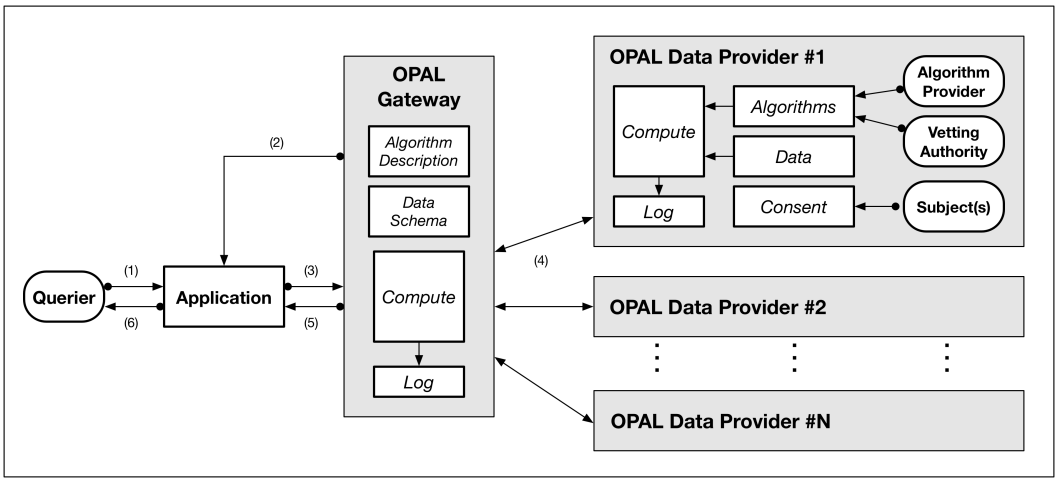

Fig. 1. Overview of the Open Algorithms (OPAL) Ecosystem

- *Decentralized Data Architectures*: Data repositories should adopt decentralized and distributed data architectures for infrastructure security and resiliency.
  Cryptographic techniques such as *secret sharing* [19] can be applied to data, which yields multiple encrypted "shards" of the data. These shards can in-turn be distributed physically across a network of repositories belonging to the same data provider [11]. This approach increases the resiliency of the data provider infrastructure because an attacker would need to compromise a minimal number of repositories (N out of M nodes) in the data provider's network before the attacker can obtain access to the data item. This approach increases the attack surface, and makes the task of attacking considerably harder. Combining this approach with secure multi-party computation (e.g. see MIT Enigma [20]) provides a possible future direction for the infrastructure resiliency of data providers.

The open algorithm principles also applies to individual *personal data stores* (PDS) [12–14, 21], independent of whether the PDS is operated by the individual or by a third party service provider (e.g. hosted model). The basic idea is that in order to include the individual citizen in the open algorithms ecosystem, they must have sufficient interest, empowerment and incentive to be a participant [22]. The ecosystem must therefore respect personal data stores as legitimate OPAL data repository end-points. New models for computations across highly distributed personal data repositories need to be developed following the open algorithms principles.

Furthermore, new service provisioning architecture must be envisaged that allows individuals to automatically relocate their OPAL personal data stores from one operator to another. One such proposal [23] uses smart-contract technologies together with legally binding terms to ensure that the service provider does not retain copies of the data embedded within the portable PDS. For communities of individuals and for organizations, the notion of *decentralized autonomous organizations* [24] can be further developed to become the basis for *community data stores* (CDS). A community data store combines individual data belonging to multiple community members, and operates under a well-defined governance model and legal trust framework [25–27]. The financial Credit Union model in the U.S. may provide a suitable legal model for community data stores [28], one in which the community as a legal entity has *information fiduciary* obligations to its individual members [29, 30].

3 THE MIT OPAL DESIGN

In the following sections we describe the OPAL development project at MIT that implements the open algorithms paradigm. The high-level design of MIT OPAL is summarized in Figure 1, and consists of the following high-level entities, services and functions:

- *OPAL Data Service*: The OPAL Data Service is the service that allows a caller (Querier) to request algorithm(s) to be executed on data located at the Data Provider.
  The data service makes available a description of the algorithms and the schema of the data-sets available at the backend data providers. Additionally, in some deployment situations the data service may perform the task of collating and merging responses from various data providers.
  Note that in the single data provider scenario, the data service may be owned and operated by the data provider, and may even be collapsed into the data provider's infrastructure. In the case of a consortium of data providers, they may collectively own and operate the data service (see Section 5).
- *Data Provider*: The Data Provider represents the data repository which holds the relevant data and algorithms.
  The source of the data and algorithms can be the data provider itself, or they may have been obtained by the data provider from external sources. The data provider may publish (directly or through to the data service) the *data schemas* of its available data and *algorithm description* of its algorithms .
- *Algorithm Provider*: The Algorithm Provider is the entity that supplies algorithms specifically for data held at the data provider.
  In some cases the algorithm provider may not a separate entity from the data provider (i.e. the data provider authors its own algorithms). In other situations, algorithm provider could be an outsourced entity whose task is to create custom algorithms for the data provider. The algorithm provider is called out as a separate function in Figure 1 because there are some circumstances in which the data provider may not wish to create or own algorithms due to liabilities that may be incurred.
- *Querier*: The Querier is the entity wishing to obtain information or insights by having (requesting) a specific algorithm be computed over data held by the data provider entity. The querier is assumed to remunerate the data service operator and the data provider in some manner (i.e. fees).
- *Application*: The Application is the tool or system used by the Querier to interact with the OPAL Data Service.
  The application may be owned and operated by a single data provider, by a consortium of data providers, or by an independent 3rd party service provider. This later case of a 3rd party operated application is the basis for many current web-applications [31], where the web-applications is often referred to as the "client".
- *Data Subject*: The Data Subject is the legal individual or organization whose data is present within the larger data collection held by the Data Provider. The data subject is the entity providing consent for the algorithm execution over the data-set which may contain their data.
- *Vetting Authority*: The Vetting Authority is the entity that provides assurance regarding the quality of a given algorithm intended for a specific data. The vetting authority provides expert review based on a well defined *fairness criteria* that is relevant to the data provider, data subject(s) and the querier.

The general interaction flow among the entities is shown in Figure 1. The Querier (individual or organization) seeking information employs the Application in Step 1 to select one or more algorithms and their intended data (Step 2). The Querier uses the Application to convey these selections to the Data Service in Step 3. Payment may accompany this request from the Querier. The Data Service interacts with the relevant Data Providers in Step 4 in order to complete the request. The Data Service returns the response to the Application and Querier in Step 5.

As mentioned previously, an algorithm intended for a given data must be vetted by experts in order to obtain some measure of fairness of the algorithm as used for a given data. Although the topic of fairness is outside the scope of the current work, it is worthwhile to mention specific challenges that are relevant. The term "fair" is used loosely to denote a set of broad categories of issues, and not any specific technical approaches. As AI and machine learning approaches get more adoption in the real world, their impact will affect different parts of society in different ways.

Although outside the scope of this work, associated with algorithmic fairness are the issues of *transparency* and *accountability*. In OPAL transparency refers to the history of transactions, namely the precise tracking of *which algorithm* was executed on *which data*, by *which entity* at *which moment in time* for *what purposes*. This tracking information should be visible to data subjects (individuals and organizations) whose data is involved. Authentic information regarding tracking should be the basis for accountability.

Today the OPAL architecture is being deployed at national scale in Senegal and Colombia, by the DataPop Alliance, Imperial College, the authors here at MIT, and the French telcom company Orange. These deployments are supported by the French AFD, Orange, the governments of Colombia and Sengal, and telcos Sonatel and Telefonica.

4 AUTHORIZATION AND CONSENT MANAGEMENT

Within the MIT OPAL design we used the notion of *consent for execution* in contrast to the usual ambiguous concept of "consent to access". Consent for execution means permission to execute a given vetted algorithm over data for a duration of time for a stated purpose, without moving the data from its repository. We believe is approach is substantially different from the usual "consent for access" which is most commonly interpreted as permission to read (copy) data – something which violates the open algorithms principles. In order for a subject (individual or organization) to have a meaningful understanding of the implications of "consent", sufficient, clear and unambigous *notice* must be provided to the subject. This notice must never be modified post-event without the subject re-consent.

In addressing the issue of consent, the MIT OPAL design takes into account the following three (3) models for consent:

- *Subject consent for data participation*: Here the Subject is giving permission to the Data Provider to include the subject's data within the broader data-set for algorithm execution. This question of inclusion is separate from (but may be dependent on) the question of (i) whether the algorithm computes aggregates only, and (ii) whether the subject is allowing themselves to be re-identified as a result.
- *Subject consent for execution*: Here the Subject is giving permission to the Data Provider to run a specific algorithm on a data-set, within which the subject's data resides. A key aspect is the subject's own understanding of what the algorithm computes, and whether the algorithm has been vetted by entities accepted (trusted) by the subject.
- *Subject consent delegation*: Here the Subject is giving permission to another entity to make decisions regarding consent for data participation and consent for execution. This aspect is relevant in cases when the subject does not have the capacity to decide on a per-execution

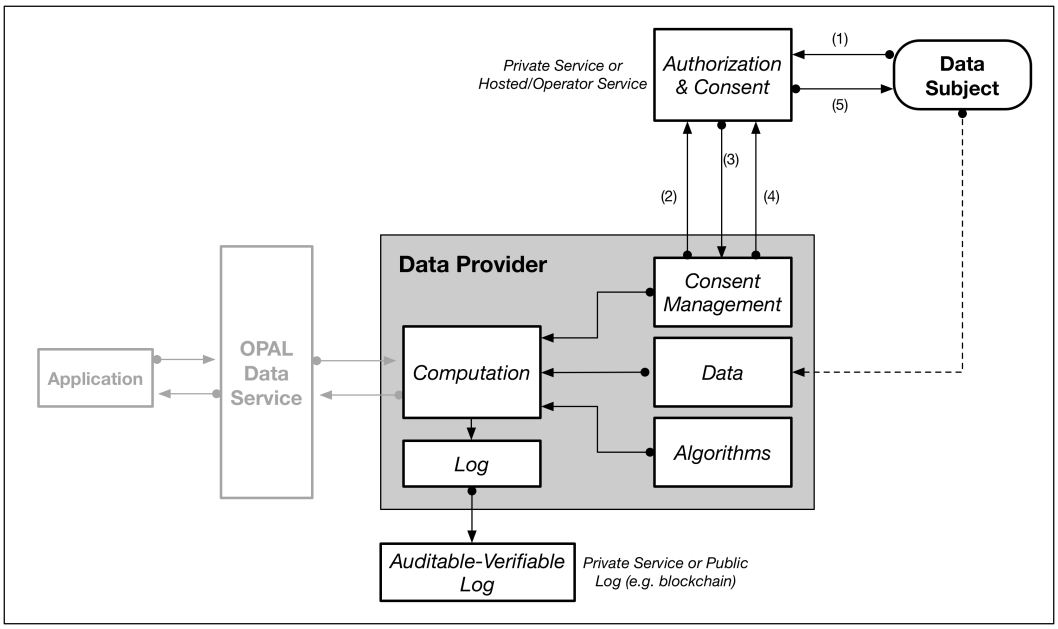

Fig. 2. Authorization and Consent in MIT OPAL

basis (e.g. medical situation) or the case when the subject is deceased (e.g. personal data bequeathed to a trust).

The technical construct used in the MIT OPAL design is the classic *access ticket* (access token) that was first popularized in the 1980s by the MIT Kerberos authentication system [32–34], and which is now the basis for the majority of token-based access control systems in industry. In the context of web applications the same notion of tickets or tokens is used in the XACML standard [35] as well as the OAuth2.0 authorization framework [31]. For consent management, the MIT OPAL design follows closely the authorization model of the *User Managed Access* (UMA) standard [36, 37]. UMA extends the OAuth2.0 framework by introducing new functions and services which contributes in the following way:

- *Recognition of service operators as 3rd party legal entities*: An important contribution of UMA is the recognition that in the real world deployments of services there are entities which provide services to the user and the data/resource owner but which may be "opaque" to them in that they have not provided consent for these service operators to gain access to data and information handled by or routed through the operator. This problem becomes acute when the data or resource is of high value. For example, UMA recognizes that the OAuth2.0 *client* entity is typically a 3rd party service that is owned and operated by a legal entity with whom the resource owner (subject or data owner) does not have any relationship.
- *Service end-points as legally binding points of access*: Another important contribution of UMA is the recognition that a given access service-endpoint (e.g. REST API) and the transaction flow provides a natural point to bind transacting entities (and their service-operators) into a legal agreement. More specifically, in a given transaction handshake (e.g. between Alice and Bob) consisting of several transaction flows (i.e. atomically inseparable set of messages), there is a point where continuing to the next transaction flow implies to both parties the acceptance

of legally binding outcomes. Thus, the technical implementation of access service-endpoint gives rise to legally binding obligations.

More specifically, in MIT OPAL we further distinguish the following types of tokens in the MIT OPAL design:

- *Execution Consent Token*: This token represents the permission granted to run a specific algorithm over a specific data for a specific purpose. This token must be digitally signed by the grantor of the permission.
  In the case of personal data, the grantor is the individual data subject and the grantee is the data provider. In the case of organizational data owned by a legal organization, the grantor and grantee maybe persons (e.g. employees, entities) inside the organization. In such cases the token may be part of an internal broader privileges-management system (e.g. Microsoft PAC structure [38]).
- *Delegation token*: This token represents permission by a data subject to another party to perform decision-making (on behalf of the subject) with regards to the execution a specific algorithm on a specific data. The delegation token must be digitally signed by the subject and ideally should indicate a time duration of validity.
  The notion of delegation tokens or tickets has been explored extensively in different systems (e.g. proxiable and forwardable tickets in Kerberos [33, 38, 39] and more recently in the XACML standard [35]).

The authorization and consent flow is summarized in Figure 2. In Step 1 the subject registers the existence, availability and location of data. This registration is performed at an Authorization and Consent Service which maybe a service or system operated by the subject, operated by the data provider, or a hosted service operated by a third party. Here it is important to note that no data is exported to Authorization and Consent Service. A data provider must obtain an execution-consent token from the subject to include the subject's data within a given algorithm execution. This request is shown as Step 2 in Figure 2, with the issuance of the execution-consent token in Step 3. In turn the data provider must issue an execution consent-receipt in Step 4 and Step 5.

With regards to receipts, one promising construct that has been standardized is the *Consent Receipt* [40] structure. This receipt is signed by the data provider, and in effect it gives the subject a record of what the subject has consented to. Related to consent, Figure 2 also shows a audit-log that is external to the data provider entity. The purpose of this log to to allow the subject and other relevant entities to obtain insight as to which algorithms were executed on which data at which point in time (i.e. transparency and accountability).

5 ENABLING DATA PROVIDER COMMUNITIES IN INDUSTRY

Data increases in value when it is combined across different domains or verticals, yielding insights that were previously unattainable [2, 3]. Data is crucial for the running of communities, business, society and government. The key challenge is how different entities in the community can share insights meaningfully, without compromising the privacy and individuals and organizations.

We believe that the open algorithms paradigm points to a new direction in which the sharing of insights among organizations and institutions can be achieved at scale. We refer to groups of entities sharing insights as *data provider communities* or *data communities* (or "circle of trust" for short). The idea is that a group of data providers across relevant data-domains agree to create a "consortium" which operate following clear rules, with privacy as major requirement (see Figure 3). The idea is not new, and has been used in smaller scale in fixed-attribute sharing among a federated group of Identity Providers [41, 42].

In the context of open algorithms, there are a number important aspects to the notion of data communities:

- *Adherence to open algorithms principles*: Members of a data community must agree to adhere to the principles of open algorithms, emphasizing the need for preserving the privacy of individuals and organizations whose data are held by the members.
- *Collaborative creation of new algorithms*: New insights can only be obtained by creating new algorithms which analyze data from different domains. Members of a data community should invest resources towards this end in order for them to realize the benefits of collaboration.
- *Mutual agreement to execute algorithms*: Depending on the composition of membership, the members of a data community should agree to honor the request from other members to execute group-shared algorithms against data in their respective repositories.
- *Governance by systems rules*: Members of a data community must operate by adhering to *system rules* that define the various dimensions of operations based on the open algorithms principles.
- *Common auditability*: Part of the system rules of the data community must address the function and implementation of common audit of actions and transactions occurring among the members. Such auditability will be needed to avoid potential conflicts from occuring when the membership consists of competing entities (e.g. competing businesses).
- *Enforceable legal obligations and liabilities*: The governance of the data community must be based on legal contracts that clearly define the obligations of each member, as well as liabilities in cases of non-adherence to the systems rules.

In order for a group of data providers to collaborate as a OPAL-based data community, there must be sufficient and clear business advantages and gains from doing so. That is, there must be a business case for data providers to collaborate. Some examples of these benefits are as follows:

- *Creation new data-related services*: which allow participants in a data community to offer unique and previously-unavailable insights and information derived using shared algorithms on their respective (private) data-sets.
- *Broadening of market adoption*: of a participant's existing services by enhancing it through combining with algorithms and data-sets from other participants in a data community.
- *Standardization of technical or functional operations*: to allow for improved reusability and more efficient certification, thus lowering cost burdens (for participants and their customers).

More formally, we define *system rules* for OPAL-based data communities as follows:

> *The set of rules, methods, procedures and routines, technology, standards, policies, and processes, applicable to a group of participating entities, governing the collection, verification, storage, exchange of algorithms (for specific data-sets) which provide information and insights about an individual, a community, or organization under their consent for the purpose of facilitating risk-based decisions.*

There are two broad purposes of system rules in the context of OPAL. The first is *achieving functionality* which allows the shared system to operate, while the second being *establishing trustworthiness* of the system as a whole among the members of the community of data providers.

The first purpose of systems rules – namely to achieve functionality – refers to the technical aspects of the operations of a data community:

- *Proper operations*: The system rules provides some form of governance to ensure the system function properly for its intended purpose (i.e. it works)

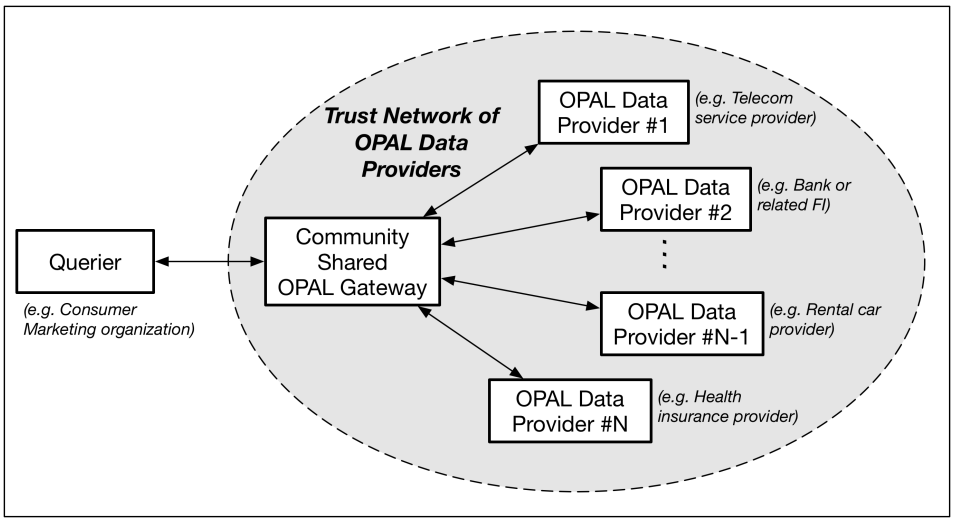

Fig. 3. Data Provider Communities based on Open Algorithms

- *Compliance*: To ensure that the system and its participants operate in accordance with the requirements of any applicable law. In the OPAL case, this includes the various privacy regulations in the jurisdiction within which the community of data providers operate.

The second purpose of systems rules – namely to establis trustworthiness – refers to the ability for parties participating in a data community to obtain a measurable degree of certainty as a function of risk management.

- *Risk Management*: The system rules allows entities to address and manages risks inherent in participating in the OPAL-based data community system.
- *Legal Certainty and Predictability*: The system rules addresses the legal rights, responsibilities, and liabilities of the participants, and thus eliminates the uncertainty of the application of existing law not written for OPAL-based data community systems.
- *Transparency*: The availability of system rules makes the terms of the specifications, rules, and agreements comprising the system rules to be accessible to all participants

6 THE PERSONAL INFOMEDIARY

As mentioned previously the MIT OPAL Project grew out of several projects that sought to address the question of privacy in this Big Data world. One such project was *OpenPDS* [14] that sought to develop further the concept of personal data stores (PDS) [12, 13]. In OpenPDS the idea was for individuals to install an agent software on the mobile devices that would retain copies of all data sent and received by the user on that device. Seeing that the mobile devices typically have limited memory/storage, the data would be periodically downloaded to some personal data store (e.g. home storage server, storage in the cloud, etc). The data in the personal data store could then be made accessible to external queriers, where local analytics would always return "safe answers". This is illustrated in Figure 4.

Currently, an extension to the basic OpenPDS concept is being developed based on the OPAL concept, and which is referred to as the *OPAL Personal Infomediary* (or simply "Infomediary" for short). We define the personal infomediary as an active software agent that acts as a two-way proxy between the user (owner of the infomediary) and a given social media platform with which the

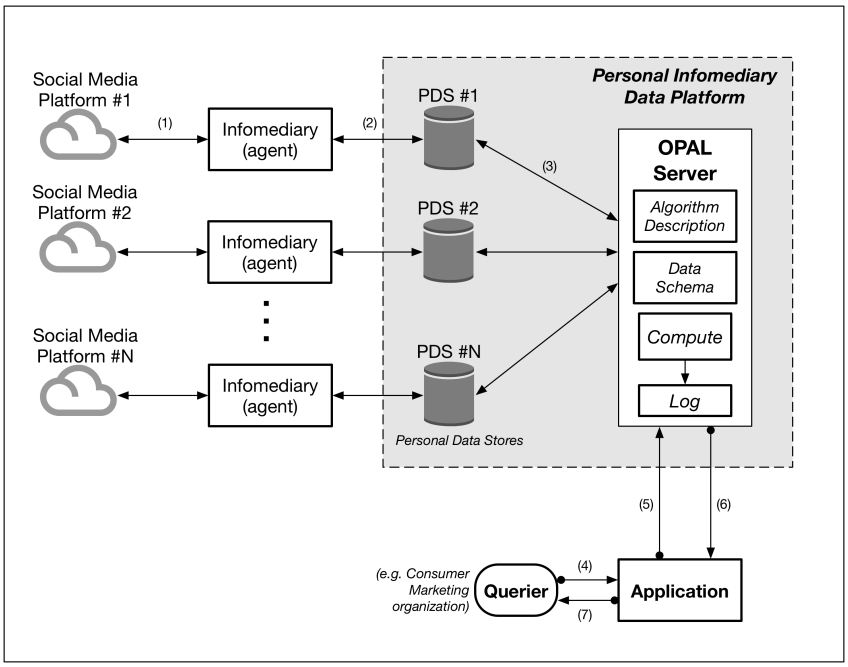

Fig. 4. The Personal Infomediary based on OPAL

user interacts on a daily basis. For simplicity, we assume there will be a unique infomediary for each social media platform or other interactive data sources.

In being an active proxy to the social media platform (or other interactive data sources) the the goal of the MIT personal infomediary is as follows:

- *Intelligent personal mediation*: Provide an intelligent agent that mediate interaction between the user at the social media platform. The agent can be enhanced to incorporate various advanced AI and Machine Learning techniques to enhance the user experience.
- *Personal data collection*: Collect copies of all traffic between the user and the social media platforms. The user is free to deploy one infomediary for multiple social media platforms (1-to-Many), or one infomediary for each social media platform. Data from the infomediaries are connected into the user's personal data store.
- *Personal analytics*: Provide a source of user-owned data such that the user can perform analytics (e.g. using software tools) in order to: (a) gain insight about the user's behavior for a given social media platform; (b) gain insight about the user's behavior across multiple social media platforms; and (c) gain insight about the behavior of the social media platform and detect possible unfair targeting feeds from the social media platform.
- *Personal revenue*: Provide a source of revenue for the individual user by providing OPAL-based access to the user's data collected in the various personal data stores.

Cooperatives such as credit unions or other trusted third parties may act as *information fiduciaries* [29], offering to host infomediary services for their community. By aggregating user data for the benefit of the users, such a fiduciary service can act as a "data union", analogous to a labor union, balancing the power of large corporations.

7 CONCLUSIONS

The goal of this chapter has been to present and discuss the *open algorithms* (OPAL) paradigm, discuss its key principles and provide an example of an implementation in the form of the MIT OPAL Project.

Society is currently facing an interesting dilemma with regards to data-driven decision making. On one hand, individuals, organizations and communities need access to data in order to perform computations as part of decision-making. On the other hand, however, there is considerable risk to privacy when data is shared (copied) across entities. The open algorithm paradigms seeks to address this by changing the way we view data processing and privacy-preserving computations.

The key principles – move the algorithm to the data, data must never leave its repository, and using vetted algorithms – are very much in-line with the goals of the GDPR regulations, notably in placing emphasis on individual data privacy. We have discussed the MIT OPAL design as an example of an implementation of the open algorithms principles. The design recognizes a number of entities in the OPAL ecosystem, with specific roles that together must meet the goals of the OPAL principles. From the user consent perspective, the OPAL paradigm simplifies consent by requesting *consent for algorithm execution* over the user's data. This is considerable better than the current norm in industry where "consent to access" is interpreted as permission to copy/move raw data.

Data increases in value when it is combined across different domains or verticals, yielding insights that were previously unattainable. To this end we believe that a trust network of data providers – operating based on a common legal framework – provides a promising avenue for data providers in industry to collaborate while maintaining data privacy. Part of this legal framework are the system rules that define the legal obligations and liabilities of each entity in the trust network.

Finally, we venture into the future by presenting the MIT *personal infomediary* project that builds on OPAL. The personal infomediary is an active software agent that acts as a two-way proxy between the user (owner of the infomediary) and a given social media platform with which the user interacts on a daily basis. The infomediary collects copies of all traffic between the user and the social media platforms, and provides a source of user-owned data such that the user can perform analytics in order to gain insight about the interactions of the user with the various social media platforms. When a trusted third party offers to host user infomediaries as a fiduciary service, the result can be a "data union" or "data cooperative" that can balance the power of large corporations.

REFERENCES

[1] A. Pentland, D. Shrier, T. Hardjono, and I. Wladawsky-Berger, "Towards an Internet of Trusted Data: Input to the Whitehouse Commission on Enhancing National Cybersecurity," in *Trust::Data - A New Framework for Identity and Data Sharing*, T. Hardjono, A. Pentland, and D. Shrier, Eds. Visionary Future, 2016, pp. 21–49.

[2] V. K. Singh, B. Bozkaya, and A. Pentland, "Money Walks: Implicit Mobility Behavior and Financial Well-Being," *PLOS ONE*, vol. 10, no. 8, pp. 1–17, 08 2015. [Online]. Available: https://doi.org/10.1371/journal.pone.0136628

[3] A. Pentland, *Social Physics: How Social Networks Can Make Us Smarter*. Penguin Books, 2015.

[4] A. Pentland, T. Reid, and T. Heibeck, "Big Data and Health - Revolutionizing Medicine and Public Health: Report of the Big Data and Health Working Group 2013," World Innovation Summit for Health, Qatar Foundation., Tech. Rep., December 2013, http://www.wish-qatar.org/app/media/382.

[5] World Economic Forum, "Personal Data: The Emergence of a New Asset Class," 2011, http://www.weforum.org/reports/personal-data-emergence-new-asset-class.

[6] European Commission, "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation)," *Official Journal of the European Union*, vol. L119, pp. 1–88, 2016.

[7] R. Abelson and M. Goldstein, "Millions of Anthem customers targeted in cyberattack," *New York Times*, February 2015. [Online]. Available: https://www.nytimes.com/2015/02/05/business/hackers-breached-data-of-millions-insurer-says.html

[8] T. S. Bernard, T. Hsu, N. Perlroth, and R. Lieber, "Equifax says cyberattack may have affected 143 million in the U.S." *New York Times*, September 2017. [Online]. Available: https://www.nytimes.com/2017/09/07/business/equifax-cyberattack.html

[9] K. Gramville, "Facebook and Cambridge Analytica: What you need to know as fallout widens," *New York Times*, March 2018. [Online]. Available: https://www.nytimes.com/2018/03/19/technology/facebook-cambridge-analytica-explained.html

[10] Y. A. de Montjoye, J. Quoidbach, F. Robic, and A. Pentland, "Predicting personality using novel mobile phone-based metrics," in *Social Computing, Behavioral-Cultural Modeling and Prediction (LNCS Vol. 7812)*. Springer, 2013, pp. 48–55.

[11] A. Pentland, "Saving Big Data from Itself," *Scientific American*, pp. 65–68, August 2014.

[12] T. Hardjono and J. Seberry, "Strongboxes for Electronic Commerce," in *Proceedings of the Second USENIX Workshop on Electronic Commerce*. Berkeley, CA, USA: USENIX Association, 1996.

[13] ——, "Secure Access to Electronic Strongboxes in Electronic Commerce," in *Proceedings of 2nd International Small Systems Security Conference (IFIP WG 11.2)*, J. Eloff and R. von Solms, Eds. Copenhagen, Denmark: IFIP, May 1997, pp. 1–13.

[14] Y. A. de Montjoye, E. Shmueli, S. Wang, and A. Pentland, "openPDS: Protecting the Privacy of Metadata through SafeAnswers," *PLoS ONE 9(7)*, pp. 13–18, July 2014, https://doi.org/10.1371/journal.pone.0098790.

[15] C. Gentry, "Fully Homomorphic Encryption using Ideal Lattices," in *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC'09)*. New York, NY, USA: ACM, 2009, pp. 169–178.

[16] A. C. Yao, "Protocols for Secure Computations," in *Proceedings of the 23rd Annual Symposium on Foundations of Computer Science*, ser. SFCS '82. Washington, DC, USA: IEEE Computer Society, 1982, pp. 160–164.

[17] A. C.-C. Yao, "How to generate and exchange secrets," in *Proceedings of the 27th Annual Symposium on Foundations of Computer Science*, ser. SFCS '86. Washington, DC, USA: IEEE Computer Society, 1986, pp. 162–167.

[18] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game," in *Proceedings of the Nineteenth Annual ACM Symposium on Theory of Computing*, ser. STOC '87. New York, NY, USA: ACM, 1987, pp. 218–229.

[19] A. Shamir, "How to Share a Secret," *Communications of the ACM*, vol. 22, no. 11, pp. 612–613, Nov. 1979.

[20] G. Zyskind, O. Nathan, and A. Pentland, "Decentralizing privacy: Using blockchain to protect personal data," in *Proceedings of the 2015 IEEE Security and Privacy Workshops*. IEEE, 2015.

[21] A. Mitchell, I. Henderson, and D. Searls, "Reinventing direct marketing âĂŤ with VRM inside," *Journal of Direct, Data and Digital Marketing Practice*, vol. 10, no. 1, pp. 3–15, July 2008.

[22] D. Searls, *The Intention Economy*. Harvard Business Review Press, 2012.

[23] T. Hardjono, "Decentralized Service Architecture for OAuth2.0," Internet Engineering Task Force, draft-hardjono-oauth-decentralized-00, February 2017, https://tools.ietf.org/html/draft-hardjono-oauth-decentralized-00.

[24] J. H. Clippinger, "The Next Great Internet Disruption: Authority and Governance," in *From Bitcoin to Burning Man and Beyond*, D. Bollier and J. Clippinger, Eds. ID3, 2013.

[25] T. Hardjono, P. Deegan, and J. Clippinger, "The ID3 Open Mustard Seed Platform," in *From Bitcoin to Burning Man and Beyond*, D. Bollier and J. Clippinger, Eds. ID3, 2013.

[26] ——, "On the Design of Trustworthy Compute Frameworks for Self-Organizing Digital Institutions," in *Proceedings of the 6th International Conference on Social Computing and Social Media and 16t International Conference on Human-Computer Interaction (LNCS 8531)*, G. Meiselwitz, Ed. Springer-Verlag, 2014, pp. 342–353.

[27] ——, "Social Use Cases for the ID3 Open Mustard Seed Platform," *IEEE Technology and Society Magazine*, vol. 33, no. 3, pp. 48–54, 2014.

[28] D. Greenwood, A. Stopczynski, B. Sweatt, T. Hardjono, and A. Pentland, "The New Deal on Data: A Framework for Institutional Controls," in *Privacy, Big Data and the Public Good: Frameworks for Engagement*, J. Lane, V. Stodden, S. Bender, and H. Nissenbaum, Eds. Cambridge University Press, 2014.

[29] J. M. Balkin, "Information Fiduciaries and the First Amendment," *UC Davis Law Review*, vol. 49, no. 4, pp. 1183–1234, April 2016.

[30] C. F. Kerry, "Why protecting privacy is a losing game today – and how to change the game," Brookings Institution, Report - Center for Technology Innovation, July 2018, https://www.brookings.edu/research/why-protecting-privacy-is-a-losing-game-today-and-how-to-change-the-game.

[31] D. Hardt, "The OAuth 2.0 Authorization Framework," RFC 6749 (Proposed Standard), Internet Engineering Task Force, Oct. 2012. [Online]. Available: http://www.ietf.org/rfc/rfc6749.txt

[32] J. G. Steiner, C. Neuman, and J. I. Schiller, "Kerberos: An authentication service for open network systems," in *Proceedings of USENIX*, March 1988.
[33] J. Kohl and C. Neuman, "The Kerberos Network Authentication Service (V5)," RFC 1510 (Historic), Internet Engineering Task Force, Sep. 1993, obsoleted by RFCs 4120, 6649. [Online]. Available: http://www.ietf.org/rfc/rfc1510.txt
[34] T. Hardjono, "Oauth 2.0 support for the kerberos v5 authentication protocol," Internet Engineering Task Force, draft-hardjono-oauth-kerberos-01, December 2010, https://tools.ietf.org/html/draft-hardjono-oauth-kerberos-01.
[35] OASIS, "eXtensible Access Control Markup Language (XACML) Version 3.0," January 2013, available on http://docs.oasis-open.org/xacml/3.0/ xacml-3.0-core-spec-os-en.html.
[36] T. Hardjono, E. Maler, M. Machulak, and D. Catalano, "User-Managed Access (UMA) Profile of OAuth2.0 – Specification Version 1.0," Kantara Initiative, Kantara Published Specification, April 2015, https://docs.kantarainitiative.org/uma/rec-uma-core.html.
[37] E. Maler, M. Machulak, and J. Richer, "User-Managed Access (UMA) 2.0," Kantara Initiative, Kantara Published Specification, January 2017, https://docs.kantarainitiative.org/uma/ed/uma-core-2.0-10.html.
[38] Microsoft Corporation, "Microsoft Privilege Attribute Certificate Data Structure," Microsoft Corporation, MS-PAC Specification v20140502, May 2014.
[39] ——, "Microsoft Kerberos Protocol Extensions," Microsoft Corporation, MS-KILE Specification v20140502, May 2014.
[40] M. Lizar and D. Turner, "Consent Receipt Specification Version 1.0," March 2017, https://kantarainitiative.org/confluence/display/infosharing/Home.
[41] T. Hardjono and A. Pentland, "On Privacy-Preserving Identity within Future Blockchain Systems," in *W3C Workshop on Distributed Ledgers on the Web*. Cambridge, MA, USA: W3C, June 2016, https://www.w3.org/2016/04/blockchain-workshop.
[42] ——, "Open Algorithms for Identity Federation," in *Proc IEEE Future of Information and Communication Conference*, Singapore, April 2018, https://arxiv.org/pdf/1705.10880.pdf.