Accounting Value Effects for Responsible Networking

Sileno, G.; Grosso, P.

DOI
10.1145/3472951.3473507

Publication date
2021

Document Version
Final published version

Published in
TAURIN'21

Accounting Value Effects for Responsible Networking

Giovanni Sileno
Complex Cyber Infrastructure,
University of Amsterdam
Amsterdam, The Netherlands
g.sileno@uva.nl

Paola Grosso
Multiscale Networked Systems,
University of Amsterdam
Amsterdam, The Netherlands
p.grosso@uva.nl

ABSTRACT

Networking offers an interesting standpoint to discuss aspects of responsible computing. The focus on a single high-level task (to provide secure and reliable end-to-end communications), and other characteristics such as resource/load distribution and decentralization of control, are fruitful properties to bring to the foreground fundamental dimensions which are generally overlooked in the responsible AI discourse: the pluralism of policies that may be in place, of expectations that may be taken into account, of their mechanisms of production and update, and their entrenchment in operationalization. Elaborating on a recent proposal on responsible Internet, the paper aims to provide a wider view of what is needed for networking to become a responsible type of computing, providing arguments for an architecture accounting for dynamic higher-level policies and expectation artefacts.

CCS CONCEPTS

• **Networks**; • **Social and professional topics** → **Computing / technology policy**; • **Human-centered computing**;

KEYWORDS

Responsible networking, Responsible computing, Agentive responsibility, Value effects, Policy operationalization, Normware

ACM Reference Format:
Giovanni Sileno and Paola Grosso. 2021. Accounting Value Effects for Responsible Networking. In *ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet (TAURIN '21), August 23, 2021, Virtual Event, USA.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3472951.3473507

1 INTRODUCTION

The widespread deployment of information-processing systems in all types of human activities raises reasonable concerns concerning the security, safety (when applicable), as well as the quality of services granted by these artefacts (hardware or software). These concerns are not limited to technical dimensions, but, at least implicitly, derive from the interests or preferences of users and—taking into account social acceptability and social opportunity—of any other stakeholder affected by their use.

Most of the attention on the broad topic of "responsible computing" goes nowadays to responsible AI, and in particular to machine-learning (ML) based AI. However, networks and networking, which are also a paradigmatic type of information-processing infrastructures and technology, are not exempt from these concerns. Clearly, for the consequences it may cause or facilitate, at the individual and the collective level, data sharing is not a neutral endeavour. Even at the level of infrastructural activities, selecting a certain route over another may e.g. increase the risks of data leaks, reduce effectiveness, increase the environmental impact, or violate regulations and policies in regard to jurisdiction through which the data can pass.

In fact, we believe that the focus on networking offers a novel standpoint to discuss aspects of responsible computing. Networking is concerned only by a single high-level task: to provide secure and reliable end-to-end communications. Such simplification on the diversity of possible computational tasks, and other characteristics such as the distribution of resources and computational load, and the decentralization of control, are fruitful to bring to the foreground fundamental dimensions which are generally overlooked in AI contexts. This paper will consider in particular two aspects.

Firstly, to take into account the existence of a *plurality of policies* (term used here to capture computational regulatory artefacts generally with higher abstraction than SLA agreements resulting in e.g. BGP policies or QoS configurations), spatially and temporally located, that may determine what is "right" to do under certain conditions. Similarly, the existence of the pluralism of a *plurality of expectations*: by changing point of view and observational scale any target phenomenon typically maps to different prediction/interpretative models. These two dimensions of "norms" (where the term has to be read as in normative, and as in normal, cf. the concept of *normware* [11]) are deeply entrenched in practical settings, because any operationalization of policies requires to be grounded on expectations. For instance, if users aim to reduce the carbon footprint of their activities on the network—and they specified so in their policies—the selection of the best path will require taking into account the potential environmental impact of each available path—and this will require in turn to know which networking devices are traversed, a well as their respective carbon footprint. Similar considerations apply on any other requirement users may be willing to include in their policies.

Secondly, to take into account the variety of *sources of policies* and *expectations*. Several actors intervene concurrently on different levels of regulations, either creating or modifying the policy artefacts, or by triggering events which are relevant for the policy mechanisms (e.g. emergencies typically produce a suspension of the directives in place). This raises general concerns of *sovereignty*: who

TAURIN '21, August 23, 2021, Virtual Event, USA

ACM ISBN 978-1-4503-8639-5/21/08…$15.00
https://doi.org/10.1145/3472951.3473507

is empowered to modify policies and/or expectations, and under which conditions? Moreover, contextual conditions may change the consequences of certain computational operations within and outside the infrastructure (e.g. a change from fossil to solar energy sources to power devices will entail a change of carbon footprint). This implies that a change in expectation due to novel contextual conditions may entail different operationalizations, even in absence of an explicit intervention on policies.

The distributed and decentralized nature of networks provides a natural background where to investigate these aspects. How can computation be responsible, from the perspective of the users of the network? And from the perspective of other stakeholders possibly affected by their activity? How can networking be neutral, if policy interventions constrain its operations? Elaborating on a recent proposal on the responsible Internet [6], the paper aims to investigate these questions, related categories and implicit assumptions, attempting to provide a wider view of what is needed for networking to become a form of "responsible" computing.

The paper is structured as follows. In the remainder of this section we provide a short overview of proposals on responsible AI, and a summary on the recent proposal on responsible Internet. Section 2 presents a schematic view on networking, policies and value effects. Section 3 elaborates on responsibility and related concepts, with examples of application to networking. Section 4 sketches an architecture integrating higher-level policies and expectations for continuous operationalization. Section 5 concludes the paper, discussing perspectives beyond operationalizaion.

1.1 Principles for Responsible AI

In the past few years several contributions from academia, public and private agencies have presented conceptual frameworks aiming to identify what is needed by AI to be responsible with respect to its impact on people and society. Dignum [3] presented the ART framework, consisting of three principles:

- *Accountability* — motivations (values, norms, etc.) for any system decision need to be explicit;
- *Responsibility* — the chain of (human) control (designer, manufacturer, operator, etc.) needs to be clear;
- *Transparency* — actions need to be explained in terms of algorithms and data, and those should be inspectable.

Santoni de Sio et van den Hoven [9] observed that *meaningful (human) control* can be summarized in two main requirements:

- *Tracing* — the system needs to be able to trace back the outcome of its operations to specific directives given by humans during design or operational phases;
- *Tracking* — the system needs to respond to (moral) reasons deemed relevant by directives given by humans guiding the system and to relevant facts in the environment in which the system operates.

As an example of institutional contribution, the expert group appointed by the European Commission has identified seven requirements [7]: *human oversight*, *technical robustness*, *privacy and data governance*, *transparency*, *fairness*, *well-being*, and *accountability*.

Clearly, these different accounts have overlaps, but also take different angles. The requirements of the EU expert group include non-functional, functional and extra-functional concerns, whereas the two other proposals target a more homogeneous abstraction level (i.e. functions required for responsible AI). At face value, tracing can be mapped to accountability, tracking to responsibility, although the former stresses the need for reactivity both in terms of control and of context.

1.2 Principles for Responsible Internet

Recently, Hesselman et al. [6] revisited the ART principles in the networking domain, envisioning how a responsible Internet should work. Aiming to provide guidelines for operationalization, the principles are applied in a distinct way on *data* (concerning the interaction of users with network operators) and on *infrastructure* (concerning the interaction of network operators with the various governing bodies potentially impacting the infrastructure activity). The resulting principles are:

Transparency *data*: the system is able to describe how network operators transport and process a certain data-flow; *infrastructure*: the system is able to describe properties and relationships between network operators (location, software, servers, etc.)

Accountability *data*: network operators have to explain the processing of specific data flows, e.g. their routing decisions or incidents during transmission; *infrastructure*: network operators have to explain their infrastructural design decisions.

Controllability (instead of responsibility). *data*: users are able to specify how network operators should handle their data (generally by means of path control); *infrastructure*: the infrastructure governance is able to set constraints over network operators.

Usability (necessary to facilitate interoperability). workings of the system need to be expressed in a way that enables further analysis.

By reframing responsibility in terms of controllability, the proposal goes more towards the concept of tracking. The additional principle of usability can instead be seen as a facilitator for transparency.

The authors also propose an architectural template for a possible implementation of these principles, consisting of:

- **NIP**: *network inspection plane*, enabling users to query the infrastructure for details about its internal operations in terms of network operators;
- **NCP**: *network control plane*, enabling users to specify their expectations on the data which is transmitted by network operators, based on network descriptions;
- **POL**: *policy framework*, enabling the infrastructure governance bodies to specify policies and have network operators abiding to those norms.

The responsible Internet proposal has the (strong) merit to translate higher-level concepts as those in the ART framework into a set of specific functions that need to be implemented and maintained at infrastructural level. However, this specific mapping may be not complete, or better, we do not have at this point an explicit set of criteria to evaluate in what sense the proposal would indeed make the Internet (or networking more generally) a responsible type of computing. For answering this question, we need to clarify what we mean by responsibility. Before doing that, we will frame a few concepts relevant for our discussion.

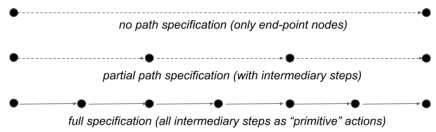

Figure 1: Different types of path specifications

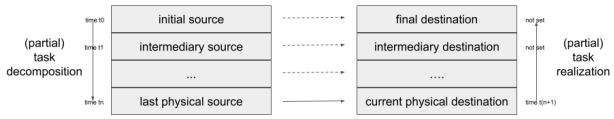

Figure 2: Routing task decomposition

2 BACKGROUND

2.1 Logistics of Data Transmission

Networking is essentially concerned with transporting data; enabling transmission between any two nodes means that the network provides at least some type of *routing* services. Most common types of routing follow a *destination-based* model. More recent types build upon some partial path specification (e.g. *segment routing* [4]), to be able to control in a more fine-grained way the traffic (see Fig. 1). In general, any transmission is defined by some *content*, and by the objective of transporting this content to a certain final *destination*. Relevant additional meta-information are the *source* or the initial point of transmission, *time* of sending and times of arrival (desired and actual, possibly not yet formed). The routing task can be thus hierarchically decomposed in sub-tasks distributed to partial paths, down to the physical level where the actual transmission occurs (see Fig. 2). The control of each level of task decomposition is typically decentralized, so that specific routing decisions can be taken at a more local scale by network operators.

2.2 Policies

In the most general meaning, policies are *artefacts used to specify and circumscribe behaviour*. They may concern interactions between individuals (norms, regulations), or the behaviour of a single individual. Under this definition, law, agreements, contracts, and non-legal normative sources can all be seen as policies. In the computational realm, we can enlist under this definition: programs (seen as sequential imperative instructions), access and usage control policies (specifying actions which are either permitted or prohibited), formal models based e.g. on deontic logics (permission, prohibition and obligation), etc. The most general framework to specify policies is plausibly defined by Hohfeld's primitive normative concepts, providing both deontic (duty, liberty), and potestative (power, disability) primitives in terms of relationships between two parties (see e.g. [12]). Using this framework, e.g. "privileges" of users in access control models, result to be both *liberties* (for the user to perform the action) and *powers* (as the user creates a duty on the system to perform the action by requesting it).

Focusing on networking, tools commonly used in inter-domain routing are Border Gateway Protocol (BGP) policies (see e.g. [2]), possibly integrating BGP community tags. These artefacts influence the routing process, by specifying which routes to advertise to external nodes (`export`) or to take into consideration from external nodes (`import`), by defining preferences across different paths. BGP policies and community tags are an example of low-level policies used for traffic control, and securing the network infrastructure.

2.3 Value effects

There are always direct and indirect demands underlying the activities performed by users of any device, tool, or designed artefact.

Computation is traditionally seen through the lens of symbol manipulation, or more recently, of data processing. However, an alternative way to see it is as "generating behaviour" (see e.g. [8]). Programs specify operations to be performed, modifying registers, memories, processing inputs, generating outputs. Part of the outputs are transmitted for further processing to other computing premises or end up to be visualized via some interface. Inputs can be provided autonomously by other entities, but can also be expected after requests sent to some other entity. Networking clearly reproduces the same template.

However, the effects of computation cannot be constrained to the digital realm. By design, at the end-point of computing flows there are typically human users, which take the results on the output interface and use it as possibly relevant information for their conduct. Less frequently, the output results in autonomously controlling some actuator, with no human in-the-loop. In both cases, some human and the environment will be eventually impacted by computational behaviour, in a positive or a negative way, which in turn may result in subsequent human behaviour. This *value effects* correlate with the drivers that promote or demote the usage of a certain device. For instance, patients may be willing to share their medical data to support the advancement of science (intended data-sharing). Even when personal data is leaked (unintended data-sharing), someone may intentionally take economic advantage at the detriment of those data-subject.

Indeed, Internet has been and is a catalyst—if not an intended instrument—for several phenomena of social, economic and cultural nature, and the resulting—generally positive—value effects on users explain its widespread utilization today. Unfortunately, the mere networking activity offers a very limited view of what is happening at the application level. What is clear is that, when nodes decide their routing policy, they may enable/disable certain routes not only to their own transmissions, but also to other nodes in the network, determining certain value effects to nodes, and eventually to users. Furthermore, the networking activity in itself has certain side effects, as for instance physical infrastructure setup and maintenance costs, and various voices related to energy consumption. At face value, networking is not a neutral endeavour: it impacts people, organizations, institutions, and the environment, positively and negatively.

3 WHAT IS RESPONSIBILITY?

3.1 Outline of core concepts

As a concept, **responsibility** is utilized in moral discourse, law, as well as software engineering (e.g. the *single-responsibility* principle: one module, one function). At the core of all these uses, responsibility is meant to help to localize in some natural or artificial social aggregate who/what caused a failure of behavioural expectations.

As a common denominator, it is an attribute which: (a) concerns some type of behaviour which is *expected*, and (b) is *ascribed* to entities (components or agents) that are deemed in charge of that behaviour within the aggregate system. Because several entities may generally concur to produce a certain outcome, rather than a binary concept, responsibility is a matter of *degree of contribution*, distributed across contextually involved components.

While responsibility is concerned primarily by actions, **accountability** is generally seen as concerned by formulating arguments explaining and justifying those actions. Instead, **liability** typically refers to potential duties (e.g. paying damages) associated to failures, or to other special contexts. Legal responsibility typically consists of liabilities attributed to legal persons.[1] In other words, acting (responsibility), explaining (accountability), and being prosecutable (liability) may be attributed to distinct agents.

Transparency is a related, yet different concept, and can have different meanings, somewhat confounded in the various proposals. Transparency of decision-making usually means having access to all the attributes determining a certain decision-making. Transparency of an environment (e.g. of a market) usually refers to the possibility of the agents to have access to information which would be usually hidden, but which would be very relevant for their tasks.

Usability is a broad notion applied in design to study the effectiveness, efficiency and satisfaction of users in interacting with an interface with respect to their tasks (see e.g. [5]).

3.2 From operational to agentive responsibility

A fundamental distinction needs to be drawn between operational uses of responsibility, as in engineering, and responsibility as an attribute of moral or legal agents. In between the technical and socio-legal uses, we refer to the notion of **agentive responsibility** [10], deemed to take into account how the agent concurs and interacts with a certain outcome, without entering the debate of its "moral" characteristics. The *agentive responsibility* of an entity results from:

- *control*: the ability to intervene upon a certain outcome (positively or negatively: producing/enabling or inhibiting it);
- *foreseeability*, the ability to anticipate relevant effects of (not) intervening upon the situation;
- *evaluation*: ability to assess those effects according to a given preferential or value structure.

The entity in focus here is an *agent*, i.e. an entity which is acting producing changes in the world, but which is also doing so in order to satisfy its preferences. Those preferences are formed privately, but may also align to collective preferences/value structures if the agent is committed to do so—for instance, those expressed via public artefacts such as policies. These three functions represents a set of criteria for any agent (natural or artificial) to be "responsible".

Agents and roles in networking. The responsible Internet proposal refers to three main roles: *users* (individuals or applications), *network operators* (e.g. ISPs), *infrastructure governance actors* (i.e. the various formal and informal governing bodies). The agents actually operating on the network are the network operators.

Today users cannot intervene on the policies of the network operators, nor those can be inspected; networking services function as black-boxes. The only possible views on operations (and more in general if SLA is satisfied) come from the behaviour observed at the end-point or from monitoring elements placed along the path. Instead, following the responsible Internet proposal, network operators should align with what requested by users and required/promoted by the infrastructure governance (for the *controllability* principle, respectively via the NCP and the POL). In terms of agentive responsibility, network operators described in [6] have control (unless there is some technical failure), can utilize the (low-level) policies expressed by users within the constraints set by the infrastructure governance to operate, but they do not have any explicit foreesability/evaluation function in place. In particular, the proposal does not consider the use of artefacts concerning expectations, so there cannot be anticipation, and consequently (user-controlled) dynamic avoidance and prevention in place. We can therefore conclude that a network abiding by the current responsible Internet proposal may fail to be agentively responsible.

3.3 Expectations and Failures

For designed entities, being "responsible" of a function means being expected to provide that function. If a failure of expectations occurs, that entity (e.g. library, module, API, service, etc.) is typically deemed faulty. For instance, if a network operator does not transmit as requested, it is at fault. If a network operator transmits where prohibited to (e.g. to untrusted nodes), it is also at fault.[2] In contrast, as autonomous entities, agents, if deemed responsible, are expected to *proactively avoid* and *prevent failures*. For doing this, three questions need then to be answered:

- how can we define that some behaviour is a failure?
- how can we foresee that a failure may occur?
- how can we state that a failure has occurred?

Possible failures in networking. To make our elaboration more concrete, let us focus on data transmission again. The primary type of failure that one can identify in this task is that the data does not arrive to the destination, or it does not arrive in time. This may for instance due to technical failure or to specific routing choices taken by one of the intermediary nodes.[3] A second type of failure makes the digital domain manifestly different from the logistics of physical goods: in networking data (and/or meta-data) can be leaked at some point of the transmission without disrupting the transmission, and so enabling nodes distinct from those intended to receive this content. To anticipate this problem, certain network domains (e.g. associated to jurisdictions) may be qualified as risky, and should therefore be avoided. Indeed, a third type of failure may concern requirements over path-dependent attributes as domains to be avoided, but also costs, energy consumption, carbon footprint to be decreased, etc.

[1] As it has been observed e.g. by Bryson [1], discussions on how to make artefacts trustworthy should not remove the fact that legal responsibility remains upon the organizations or people that produce, maintain and deploy them.

[2] This way of looking at the acting entity has similarities with the concept of *strict liability* used in certain legal contexts (e.g. product liability) in which the mental element (e.g. of the manufacturer) is not taken into consideration in judgment.

[3] Examples of reasons behind possible different local choices: (a) separating the routes for voice data on high-bandwidth from bulk content on low-bandwidth (b) responding to a DDoS attack by routing traffic through a scrubbing device.

Prototypical networking use case. A user requests to a network operator to perform a certain transmission. In terms of normative primitives, the user creates a duty upon the operator. The routing will then be performed according to e.g. BGP policies and possibly community tags provided in advance by operators of the various domains. Following those policies, network operators construct a list of possible paths by importing or not importing routes from advertisements of other operators. In terms of normative primitives, this list reifies a list of permitted transmissions. Implicitly, if the route is not in the list, it is prohibited. BGP community tags are instead used to reify preferences amongst the possible paths. Network operators are today the only entities in power to modify these operational policies, as long as they respect the transmission tasks requested by the users. In essence, the responsible Internet proposal envisions that users are entitled (empowered) to control the network operators with respect to their own transmissions, but implicitly it sets also that both users' and network operators' powers may be counterbalanced by interventions of the infrastructure governing bodies.

Using expectations. Expectations play a fundamental role in operationalizing policies. For instance, knowing the average energy consumptions of certain devices, and having a model of the type of devices used by a network operator (e.g. constructed by evidence collected via the NIP), one could compute the average energy consumption of a certain route. If minimization of energy consumption is part of the user's preferences, the resulting ranking could be stored as a community tag. Other types of expectations, e.g. about the trustworthiness of a certain operator, could be used to forbid transmission through it (i.e. to not include routes advertised by it). Unfortunately, this procedure is that it is not automated, nor now, nor in the responsible Internet proposal.

Relying on the NCP, users can provide e.g. BGP community tags expressing a numeric preference, but they can not make explicit on which bases (e.g. expectations) those numbers are constructed. This is because operational, lower-level policies as BGP policies (analogously e.g. XACML policies in access-control) do not enable users to explicitly control reasons which stay at deeper levels in the derivation process. This hard-coding is detrimental to usability, as it reduces the transparency of the control interface to the user.

4 TOWARDS EMBEDDING DYNAMIC, HIGH-LEVEL POLICIES

We will now sketch an architecture that takes into account policies specified at higher abstraction as entry point for control.

Main components. For the sake of the argument, we will assume that the target high-level policy language provides Hohfeld's primitives (duty or prohibition, liberty, power, disability), but other solutions are possible without losing generalization. Read as preferences, policy directives may be given explicit relative priority, and may be possibly conditioned to more specific contexts. Note that conditions at different levels of abstraction form in themselves a priority (just as peer-level `import` or `export` overrides a group `import` or `export` in a BGP policy). Individual directives can be aggregated in groups, which can also be prioritized. In general, policies may produce conflicting outcomes (e.g. something is both

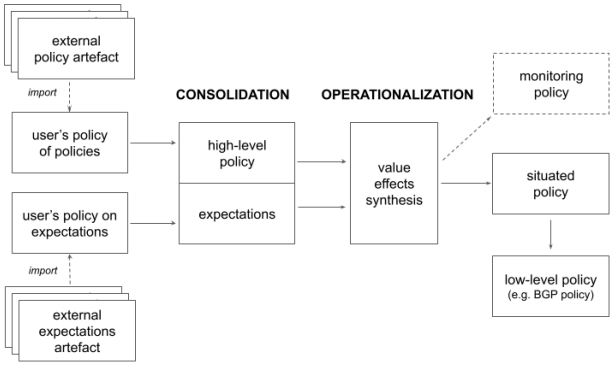

Figure 3: From higher-level polices and expectations to low-level and monitoring policies.

prohibited and obligated). However, when this occurs (during offline verification or online use) we assume users can be prompted to form a new, more specific directive that solves the conflict.

Besides internal groups of directives, users may import external groups, i.e. higher-level policies available for reuse in policy repositories. In this sense a user's policy becomes a *policy-about-policies*. For instance, users may decide to embed various policies suggested by trusted organizations, e.g. governmental or public agencies, NGOs, etc., but still override them in any conditions they deem relevant.

Similarly, expectations artefacts (in the form of models, heuristics, rules, etc.) may be made available on dedicated repositories. Users may then rely upon those sources, importing them in a *policy-about-expectations* artefact, in order to cover knowledge gaps they may have (e.g. the carbon footprint associated to devices, the number of incidents occurred on a certain routes, the associations of port numbers with type of traffic, etc.). Note that part of expectations can also be about currently holding situations; for instance, the user may rely upon a trusted third-party certifying that e.g. there is an emergency.

Operationalizaton. The combination of policy and expectations artefacts specified by the user is consolidated and then operationalized into low-level policy artefacts, as for instance a BGP policy for routing—to govern its networking activity (see Fig. 3). *Consolidation* consists in dereferencing external artefacts and resolve the various priorities into a more efficient data structure. *Operationalization* performs operations of synthesis of value effects (functionally similar to risk assessments) to decide which (group of) network operators should be forbidden transmission, or to provide a preferential ordering, *at the present conditions*. The resulting "situated" policy is then cross-compiled into a target language (e.g. BGP policy language) for actual execution. For instance, if users want to reduce the risk of data leaks, they could use statistical information about leak accidents maintained by some trusted third-party of their choice to determine an order of preference over network operators.[4] The consolidation step can also be used to set up a monitoring policy or script to check whether network operators are correctly applying

[4]Note that this could also be what operators among themselves do today. This proposal aims then to generalize and standardize this solution and give its control to users.

the required policies; this artefact is however most plausibly used at application level.

An important property of the proposed processing is that it can be triggered asynchronously to any update on the user's policy, on the user's expectations, and, more importantly, of external artefacts they import. For this reason, it offers a solution to implement a *continuous alignment* to ever changing policies and expectations.

Inertial effects. Something worth to observe is that the time of enacting a new policy and the time to see its full application may be relatively distant, because the (computational) social system on which it applies need to adapt accordingly. In routing control, for inistance, update time will be different with respect to the type of modifications, consequently to how the BGP protocol functions (another routing construction protocol may bring to different results). Modifications on `export` will take time to be effective, because all network operators need to modify their respective routing tables for the route which has been advertised, or whose advertisement has been removed. Instead, modifications on `import` and on community tags can be promptly effective.

Meta-level policies. An important advantage of this architecture is the same high-level language can be applied to specify polices about how policies can be used in the infrastructure, i.e. to *design the protocols of interaction* between users, network operators and the infrastructure governaning bodies. In the responsible Internet proposal [6], infrastructure governance plays in practice the institutional role of enforcer on the infrastructure. Activities of network operators are recorded on infrastructure premises (possibly following directives set up in POL). Users can control the network operators via NCP, and monitor their individual transactions and general performances via NIP. The infrastructure governance can put further constraints on network operators activities via POL. We can easily rephrase these relationships in Hohfeld's terms, e.g. users have the *power* over the network operators to specify their own routing rules for their transmissions; the infrastructure governance has a *power* over the network operators to impose certain constraints, etc.

Employing policies at meta-level means that the normative relationships across the three agent roles are not hard-coded and alternative solutions may be applied without requiring a full development cycle. In other words, processes associated to NCP, NIP and POL can be seen as issued after different types of collective policies, respectively dealing with protocols of user control, of inspection, and of infrastructure control. One could therefore aim to a computational infrastructure enabling a general use of policies by all parties; NCP, NIP and POL would become specific instances, plausibly associated to different low-level policy languages as end-product of operationalization (e.g. BGP for NCP). This design choice would allow us to take into account and explore further mechanisms that those three, and integrate those at run-time.

Relevant points to be explored are for instance how transmission data is collected for the NIP. A possible way would be that network operators have to provide evidence of reception and of the routing decisions for which they took responsibility. However, should they also record the rationale for that specific decision-making somewhere for accountability? In principle, if this information is maintained in their premises they may have motive to tamper with it. The infrastructure governance could for instance implement a duty over network operators to share (part of) this data with a certifying third-party, or try yet other institutional constructs.

5 PERSPECTIVES

The paper provides a reasoned selection and organization of several of the concepts related for responsible computing, contributing to the discussion over going beyond operational responsibility. By centering the focus on agentive responsibility, we have acknowledged that computation cannot be "responsible" if the computational agent has no means to evaluate the effect of its actions.

In principle, in order to have a proper analysis of the impact of a certain transmission to users or other relevant stakeholders, one needs means to evaluate the full value-chain, from the generation of a material manifestation that is captured digitally to its (context) of use. However, networks are supposed to operate blindly with respect to the content they transport; by making decisions on portions (packets) and in principle unaware of the value of the whole transactions. Part of this information is however available at the users' endpoints, as they may have expectations about what their actions will produce, and to what extent these effects are valuable. They may even decide to use resources from sources they trust in forming these expectations. Taking this information in the loop would allow to properly and continuously align the value structure and expectations of users into low-level policies to be executed at infrastructural level. The paper sketched a solution based on artefacts that may be classified as *normware* [11].

Interestingly, such proposal entails that the "dumb pipe" principle used to protect networking from being controlled by corporate or governmental agencies can be still applied. With adequate meta-rules, users can remain in control, and can opt-in/opt-out at any moment from following externally provided policies/expectations. Network operators or governing bodies may not directly intervene on users' choicess; yet, they can intervene on a second level, transparently for users, providing default policies and expectations available for imports. Thus, a network operator may in principle violate net neutrality, but this would happen after the user's mandate, thus without violating the value of the transmission (for the user) as a whole.

However, the present paper has not entered in the discussion of what is the best value structure to produce "responsible" networking, nor has taken a stance on the meta-rules specifying the power-relationships between the actors. Users, network operators, and the various governmental bodies have all legitimate interests to play a role in policy-making. An exemplary challenge in this interaction is finding the balance between *anonymity* and *accountability*. A global, definitive solution in the computational realm would be a too ambitious goal, simply because, for what we see from non-computational contexts, too many local contextual factors intervene to set which are the "right" checks and balances. Yet, we need to find a starting framework within which the various actors are able to set their policies. Rather than taking this challenge directly, the continuity between software and regulations observed here suggests rather to look at solutions developed in non-computational contexts, as in international law, or in international private law, already operative across very diverse jurisdictions.

Acknowledgments This work was partly supported by NWO (Nederlandse Organisatie voor Wetenschappelijk Onderzoek) for the DL4LD project (*Data Logistics for Logistics Data*) (628.009.001) and the HUMAINER AI project (*a HUman and Mathematical Approach to Implementing and Navigating Explainable and Responsible AI*) under contract KIVI.2019.006.

REFERENCES

[1] Joanna Bryson. 2018. AI & Global Governance: No One Should Trust AI. (2018). https://cpr.unu.edu/publications/articles/ai-global-governance-no-one-should-trust-ai.html#info

[2] Matthew Caesar and Jennifer Rexford. 2005. BGP routing policies in ISP networks. *IEEE Network* 19, 6 (2005), 5–11. https://doi.org/10.1109/MNET.2005.1541715

[3] Virginia Dignum. 2017. Responsible autonomy. *Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)* (2017), 4698–4704. https://doi.org/10.24963/ijcai.2017/655

[4] Clarence Filsfils, Nagendra Kumar Nainar, Carlos Pignataro, Juan Camilo Cardona, and Pierre Francois. 2015. The segment routing architecture. In *2015 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 1–6.

[5] Carlos Flavián, Miguel Guinalíu, and Raquel Gurrea. 2006. The role played by perceived usability, satisfaction and consumer trust on website loyalty. *Information and Management* 43, 1 (2006), 1–14. https://doi.org/10.1016/j.im.2005.01.002

[6] Cristian Hesselman, Paola Grosso, Ralph Holz, Fernando Kuipers, Janet Hui Xue, Mattijs Jonker, Joeri de Ruiter, Anna Sperotto, Roland van Rijswijk-Deij, Giovane C.M. Moura, Aiko Pras, and Cees de Laat. 2020. A Responsible Internet to Increase Trust in the Digital World. *Journal of Network and Systems Management* 28, 4 (2020), 882–922. https://doi.org/10.1007/s10922-020-09564-7

[7] AI HLEG. 2020. High-Level Expert Group on Artificial Intelligence: Ethics Guidelines for Trustworthy AI. https://ec.europa.eu/digital-single-market/en/news/ethics-guidelines-trustworthy-ai

[8] Bart Jacobs. 2016. *Introduction to coalgebra: Towards mathematics of states and observation.* Cambridge University Press. https://doi.org/10.1017/CBO9781316823187

[9] Filippo Santoni de Sio and Jeroen van den Hoven. 2018. Meaningful Human Control over Autonomous Systems: A Philosophical Account. *Frontiers in Robotics and AI* 5, February (2018). https://doi.org/10.3389/frobt.2018.00015

[10] Giovanni Sileno, Alexander Boer, Geoff Gordon, and Bernhard Rieder. 2020. Like Circles in the Water: Responsibility as a System-Level Function. *Proceedings of the 3rd XAILA workshop on eXplainable AI and Law, in conjunction with JURIX 2020* (2020).

[11] Giovanni Sileno, Alexander Boer, and Tom van Engers. 2018. The Role of Normware in Trustworthy and Explainable AI. In *Proceedings of the XAILA workshop on eXplainable AI and Law, in conjunction with JURIX 2018.*

[12] Giovanni Sileno and Matteo Pascucci. 2020. Disentangling deontic positions and abilities: a modal analysis. In *Proceedings of the 35th Edition of the Italian Conference on Computational Logic (CILC)*, Vol. 2710. 36–50.