Towards Taming the Tussles in Tomorrow's Internet

ABSTRACT
The Internet has long moved from being a research curiosity to being embedded well into the social fabric called society. Any evolutionary Internet extensions as well as any complete overhaul in some form of clean slate approach cannot ignore that the wider socio-economic conflicts of society will decidedly determine the outcome. As a research community, we need the ammunition to identify, formalize and eventually tame the various tussles, the conflicts between the involved parties, which will impose themselves on our attempts to design successful future Internet(s). This paper outlines a methodology for doing so. We utilize system dynamics modeling to analytically capture socio-economic causalities that represent the various tussles that might unfold. We also illustrate an example of its usage, paving the way for building an analytical approach to taming the tussles in cyberspace.

Categories and Subject Descriptors
C.2.1 [**Computer-Communication Networks**]: Network Architecture and Design – *distributed networks.*

General Terms
Design, Economics.

Keywords
Tussles, system dynamics, system design process.

1. INTRODUCTION
The Internet has evolved from a purely technical artifact, in which all creators shared a common goal of interconnecting computers in the world, towards a central element of our social fabric. The milieu that makes up the Internet of today includes powerful players, ranging from equipment manufacturers over ISPs to application developers but also including various governmental organizations chartered with enforcing laws and protecting various interest groups through regulatory intervention, as well as organizations representing various interest groups.[1] Given the increase in players and their often conflicting interests in the Internet and its evolution, contentions among these players come as no surprise. These contentions are more often defined on the basis of socio-economic constraints than technological superiority of solutions.

There are many examples of these contentions throughout the history of the Internet, some of them having influenced design choices of Internet deployments. For instance, the desire to freely exchange information is constrained by corporations and governments alike by the introduction of firewall and deep packet inspection technology. The battle of interests in the privacy area is carried out among end users, interest groups and governments, each imposing its interests in the deployment of final solutions, often leading to an arms race between technology choices.

The impact of these conflicts of interests, these *tussles in cyberspace*, on the various design stages has been recognized by Clark et al. [1] in their seminal work that aims for a *design for tussle*. They recognize that mechanisms for resolving contentions in the design and redesign phase have been employed effectively in various engineering disciplines, including the networking community. Standardization approaches, requirements engineering and all associated methods aim to identify the various conflicts. They all result in an engineered solution that embodies the solutions to these conflicts in the form of a technical specification for the artifact being deployed. But it is the Internet's distinctive nature, as pointed out in [1], that it embodies continued tussles at system *runtime*, such as peering decisions. And the authors in [1] conclude that it is the lack of a final outcome of these (runtime) tussles that forces us to "think about design differently" [1] in order to find an answer to the question of *how can a system be built in the presence of runtime tussles*?

Let us think a moment longer about this question, in the advent of the Internet being part of our social fabric rather than a mere technical artifact. The wider socio-economic constraints are crucial for the (temporary) outcomes of these tussles that result in the evolving design of the Internet. These constraints are difficult to be optimized towards a single objective. Instead, the design process becomes one of weighing constraints against each other, reflecting compromises along the way and reaching several possible outcomes that the design in question will need to accommodate. It is the process of *satisficing*[2] constraints more than satisfying all of them that interests us here.

It is this recognition of a wider socio-economic angle of any large-scale system design that in many case creates a resistance in our community to see this as our problem rather than a problem of economics, sociology or the like. Instead, as also pointed out in [3], we have assumed a design autonomy for ourselves and have ignored that "the *market* determinates of *both* the underlying mechanisms as well as the strategic behaviors" [3], i.e., it influences the design more than most technological decisions. Hence, any large-scale system design is futile without breaking out of this design autonomy that we as a community have assumed for too long. As Clark et al. put it [1]

> *We, as technical designers, should not try to deny the reality of the tussle, but instead recognize our power to shape it.*

[1] See [1] and [2] for a more detailed discussion on the nature of these players and their impact on the Internet.

[2] A hybrid between the words *satisfying* and *suffice*, see http://en.wikipedia.org/wiki/Satisficing

This sentence outlines a mandate for the wider research community to develop methodologies that help understand, shape and possibly tame the tussles that impose themselves on the design of the current and any future Internet[3] - and it places us technical designers in the middle of the required efforts to create such means.

This paper addresses this mandate by outlining a design methodology that captures causalities in a socio-economic environment; causalities that impose themselves on a set of architectural solutions. This methodology paves the way for developing an analytical model that informs about feasible (as well as infeasible) strategies and the possible outcomes they might trigger. It is important, however, to not misunderstand our work as a resurrection of dialectical materialism, striving for a crystal ball that predicts the future. We merely aim at aiding the many decision-making processes within our society. We do so by accompanying today's design codification, in the form of specification and programming code, with another form of code that captures the socio-economic environment. With that, we enable a constant evolution of the understanding of the design's socio-economic viability as we learn about the unfolding changes; we enable a form of 'socio-economic code maintenance'. Hence, this paper not only recognizes the power to shape the tussle, it provides a tool to do so.

The remainder of this paper is organized as follows. Section 2 provides an overview of some of the related work in this field. Section 3 then introduces the field of systems dynamics, a technique that allows for capturing system level causalities and that is therefore key to our approach. We outline our design methodology in Section 4 along two possible thrusts of design and exemplify the methodology in Section 5 with the problem of designing large-scale discovery solutions in the Internet. We provide an outlook for going forward in the debate on tussle in Section 6 before concluding in Section 7.

2. RELATED WORK

Central to understanding complex systems is the notion of *emergence* [5], describing how complex systems arise out of simple transactions. The common characteristics of *emergent behavior* are (i) features not previously observed in systems, (ii) providing a coherence or correlation over a period of time, (iii) constituting a macro level of behavior, (iv) being the product of a dynamic process, and (v) it being perceived. Within natural sciences, physical and biological systems are the classical examples for complex systems that are studied based on emergent behavior.

The notion of emergent behavior has also been applied in social science as well as engineering. One such discipline at the crossroad of engineering and sociology is that of *HCI modeling* [6][7] in which socio-technical concerns are taken into account from the perspective of multiple stakeholders. The Open System Task Analysis (OSTA) method [7], for instance, captures the various stakeholder concerns around a primary task. It studies the organizational integration as well as the organizational change that is possibly required for adopting any solution. At its heart, the objective of this modeling technique is very close to our desire, namely a holistic understanding of the system's emergent behavior. However, much of the HCI modeling is directed at the qualitative understanding without giving the designer the ability to quantitatively argue about various emergent behaviors.

Economic theory has pushed forward the quantitative understanding of the behavior of actors in a complex system. *Multi-agent techniques* [8] as well as *game theory* [9][10] are approaches developed by the economics community. These capture the strategies of players in an economic setting and the impact of their behavior on the overall outcome. Such methodologies allow for a deeper analytical insight into the impact that various design decisions might have. But expanding the games to a larger number of individually acting players proves to be challenging in terms of modeling, not the least due to the player-oriented formulation of the problem[4]. Olson's work on institutional economics [11] provides the macro-economic insight that allows for better understanding the impact social groups (or entire nations) have on social processes, such as regulation or general policy-making.

Requirements engineering [12] is a sub-discipline of systems engineering that gains more and more importance. Its goal is to capture and verify the fulfillment of a set of requirements for a system's design. Similar to our goals is the aspect of stakeholder engagement in order to capture the requirements. The focus, however, is neither to gain an understanding of the evolving behavior of the system nor to uncover the impact of conflicting requirements not only on the system, but even more on the socio-economic environment in which these conflicts unfold. With that, one can relate our work to requirements engineering as providing a better understanding of the impact of given requirements rather than determining the requirements per se. Within requirements engineering, *Ockham's razor* [13] is often used as a methodological principle to ensure simplicity in a design. As argued in [13], however, its applicability to system design is limited due to its aspect of prejudice in the definition of *simplicity*.

The field of *systems dynamics* (SD) [14] aims at capturing the evolving behavior of a system, focusing on states of the system rather than strategies of players within. The latter, however, are captured as causalities that influence system

[3] The reader may recognize our indifferent position regarding the applicability of such methodology to both *clean-slate* and *evolutionary* research [4]. We believe that both approaches require a socio-economic approach to design, as argued here, with a scientific foundation for its solution.

[4] Game theory formulates the problem as a game between selfish actors in a system, transforming the system-level view of design into a player-specific view of selfish outcome optimization.

state. Used for a variety of problems (see [14] for examples), systems dynamics provides a comprehensive understanding of the various categories of influences that are exerted on the system. It furthermore allows for quantitatively evaluating the impact of these forces through its underlying analytical basis. It is this aspect of capturing system-level causalities with its possible analytical grounding that drives the development of our methodology.

3. CAPTURING CAUSALITIES

In our goal to recognize the tussles between various actors, it is crucial to identify and describe the causalities that underlie these tussles. This enables formulizing (and quantifying the likelihood of) possible outcomes under the influence of various parameters for the causalities.

Techniques like abstract dynamical systems, feedback control and *systems dynamics (SD) modeling* [14] have been developed for this exact purpose. We focus on the latter throughout this paper. SD modeling captures the various dynamics that occur within a systems-level problem, graphically describes the causalities driving these dynamics and provides an understanding as to what causalities drive the possible outcomes of the system. These causal structures are ultimately turned into an analytical model in the form of time-varying differential equations. It is therefore a seemingly perfect match to our objectives. In order to better understand this match, let us now outline the main characteristics of this method.[5]

At the heart of SD modeling is the notion of *feedback processes* that describe the individual dynamics within the system. It is the interactions between the various feedback processes that influence the overall system behavior over time. Within a defined *problem*, the *stock* defines the state of the system that best describes the underlying problem. The rate of change with respect to this state is defined as the *flow*, usually existing as an inflow and outflow. The causalities that influence these key variables of the system are captured as negative or positive feedback loops, annotated with '+' or '-' signs to denote the type of feedback. Not only the causalities towards the main variables stock and flow, but also the influences towards *auxiliary variables* in the system, are captured. Such notion of feedback loops lends itself to graphical depiction and therefore represents the qualitative part of the SD method.

An important step in creating the causal structure itself is that of developing the so-called *reference mode* of the system – that is the expected time-dependent behavior in an idealized form [14]. Typical forms of expected behavior are linear or exponential growth (or decay) as well as hyperbolic curves [14] or combinations of them. Outlining the reference mode is important since it captures the understanding of the modeler as to how the system is expected to behave, based on the currently captured feedback. A drastic divergence from this expectation is an important input into refining the model. It is an apparent signal that the seemingly understood system behavior is not reflected in the actual, simulated, one. Reasons for this could be incomplete or incorrect parameterization but also causalities that have not been captured appropriately. In Section 5, we will see examples of such reference modes. Another important function of outlining reference modes is to support the development of *scenarios* defining the various parameter sets that are used in the simulations.

As an example, take the system in Figure 1. Here, the problem is that of some population development over time and the various influences that it faces. The stock is defined as 'population', while 'birth' and 'death' represent the in- and outflows of our system. We only depict simple feedback here, i.e., from the overall population directly into the in- and outflows. The natural rate of death and birth is treated as an *exogenous*, i.e., system-external, parameter. The feedback loops are annotated with a 'R' sign to illustrate the reinforcing nature of this feedback (as opposed to a 'B' sign for a balancing force). This example shows the qualitative nature of SD modeling in that it provides a simple graphical visualization of the dynamics in our problem. Extensions of the model could include more elaborate feedback, such as from the population stock towards parameters that capture overcrowding, which in turn might influence birth and death rates.

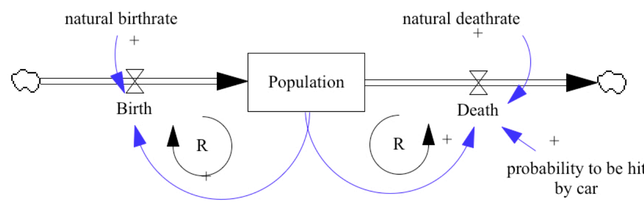

Figure 1: Systems Dynamics Model

On the analytical, i.e., formal, level, each variable in the SD model, be it auxiliary or stock/flows, is defined through a function that describes the nature of the causality. For instance, a natural birthrate could be described based on national statistics. With that, the qualitative nature of the model moves towards an analytical one, which can be used to simulate the system's behavior, depending on a set of parameters for the various functions that are being tested.

Putting all these pieces together results in a parameterized analytical model for a given problem, in which the varying outcomes depend on scenario-based input. But how does this bring us further forward in our quest to tame the tussles that are possible within architectural designs? To answer this question, we embed the SD modeling technique into a design methodology throughout the following section.

4. DESIGN METHDOLOGY

The applicability of SD modeling to system-level problems motivates the integration of this method into a methodology for evaluating system designs. We first present the methodology itself, including how we move from a design time towards an evaluation of tussles and their outcome at

[5] The reader is referred to [14] for a more elaborate introduction into systems dynamics and its various applications.

runtime of the designed system. Last but not least, we present the tools developed for this approach.

4.1 Methodology

In our thinking, the *viability of a design* constitutes two different foci within our methodology: *market focus* and *design focus*. For the former, we aim at understanding the various market outcomes enabled or inhibited by a particular design under a set of possible development scenarios. For the design focus, where the strategies are in the forefront, we aim at understanding what would make particular design choices successful or fail in the presence of various socio-economic influences.

With that, the focus of our methodology becomes that of weeding out bad design choices in the presence of certain socio-economic scenarios that the designer (or his customers) deems as being important. These 'bad' choices relate to designs that enable market outcomes that are seen as undesirable (or prohibit those being desired) or whose design choices require design strategies that are unlikely to be successful. It is this aspect of selection that positions our methodology as a *tool for architecting* on a scientific basis, rather than leaving this important step to the deployment of the selected solution.

Figure 2 presents the steps of our methodology within these two foci. We first identify important *design characteristics* within the design space that is being considered. While such design characteristics may emerge from architectural discussions, characterizing their time-dependent behavior is where SD modeling will be of help. For instance, design characteristics may be reflected as the number of players for a particular function, the degree of collaboration between particular players, or anything else that represents a particular characteristic of interest within the design space[6]. This first step is connected to SD modeling by formulating the design characteristics as system dynamics problems, i.e., representing them as *stocks* and *flows* in to-be-developed causal loop diagrams.

For each stock and flow model, a *reference* mode is developed. This allows for outlining potential *socio-economic outcomes* that are enabled (or prohibited) by particular design choices. In order to understand the likelihood of these outcomes, the various influences are captured along multiple socio-economic dimensions, ranging from user behavior over business strategies to regulation. Capturing these influences is usually done through desk research or interviews with various parties, such as regulators, incumbents, investors, end users, etc. The uncovered causalities are then modeled as *causal loop* diagrams. This results in a SD model for each considered problem in which each auxiliary variable is defined through an analytical expression of the causality that it represents. Existing work can be consulted for capturing this analytical relationship, describing, for instance, viral campaigns [14] or the uptake of investments in new technologies [15].

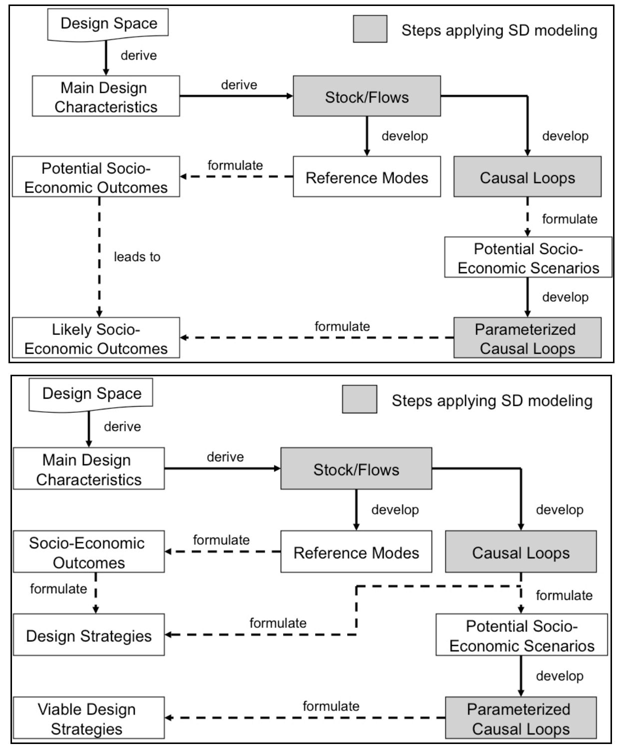

Figure 2: Methodology – Market View (top) & Design View (bottom)

Within the design view, the captured causalities help us identify *design strategies* that are potentially successful or not within a range of considered scenarios. From here, we move on to formulating relevant *socio-economic scenarios* under which the design choices are to be evaluated. These scenarios help us parameterize the auxiliary variables, i.e., the scenarios provide the parameter sets for running simulations of the developed SD models. Running the simulations, i.e., solving the equations that underlie our individual SD models, leads us to the *set of likely socio-economic outcomes* (or the set of *viable design strategies* in our design view) as a subset of the possible outcomes, as described through the reference mode for each model, under this given set of scenarios.

In the market view, the results of the simulations lead to making statements on the various markets being created while in the design view, this will enable statements on design strategies that are likely to be successful (e.g., regulatory influence to make a certain design happen).

It is our ambition to evaluate the viability of system designs not only at design time, but even after deployment. The methodology in Figure 2, together with the resulting SD models, provides a codification of the wider socio-economic causalities as the basis for such continuous evaluation. In other words, when moving from design to

[6] One must understand that the selection of the design characteristics is an important but also very subjective step that depends on the particular angle of the evaluation. It is therefore crucial to record the reasoning behind the choices being made, something our design tools in Section 4.2 specifically aim at.

runtime, we see the role of the methodology as that of *capturing evolving tussles*, including the ones that were not identified at design time. This leads to a repeatedly refined model, which in turn could result in changes to the original design (which is in turn incorporated into a refined model and so on). This evolution aims to capture the emergence and role of new actors that are often impossible to identify at the time of the original design.

4.2 No Methodology Without Tools

It seems obvious that there cannot be a successful methodology without proper tools. This is even more important in our case due to the ambition to provide some form of codification of the findings within the methodology that can constantly evolve through a refined understanding of the design space. Hence, we place a strong focus in our work on developing design tools that allow for capturing the various steps in our methodology.

It is well understood that the development of SD models is an ambitious and difficult task. It requires that an SD model developer understands the problem space from many angles, usually through a series of interviews with various stakeholders and experts, each of which increases the pool of knowledge for the particular problem at hand. Throughout these interviews, the relevant causalities and influences need to be uncovered. The results of these interviews are usually recorded, either audio-visually or through note taking (see [14][15]). There are no recognized methods or best practices for how and what should be recorded in what form. Much is left to the experience and skill of the modeler. Conducting these interviews and capturing the essential information is crucial since setting up interviews with various (often busy) experts is difficult and time-consuming. Clarifications are often difficult after the interview has ended. Hence, it is imperative to provide a tool for the modeler that supports the recording of findings, interviews, anecdotes and desk research in a coherent and guided way. Such recordings also help in the actual engagement with the various experts, i.e., a graphical visualization possibly utilizes the already captured information in future dialogues.

As a result of these considerations, we develop a toolkit that captures the findings of the various steps in our methodology. In order to aid engagement with the various parties, we utilize *mind mapping* techniques, complemented with the ability to add comprehensive notes, including multimedia annotations (e.g., recordings of interviews). The toolkit is implemented with the open source software XMind [16]. As part of our efforts, we plan to release the toolkit under an open source license.

Figure 3 outlines the various steps of the toolkit, addressing different questions regarding the particular use case, the actors and components, the points of control and many others. Behind each leaf in the mind map lies a separate mind map sheet (indicated by the small "**c**" symbols in Figure 3), with more detailed instructions and specific representations of the step implementation – for space reasons we omit these specific sheets here. Each mind map can be extended according to the instructions.

The steps in the toolkit start with specifying the particular focus of the problem, formulated as specific questions that directly relate to the design in question. The identification of the use case aids the designer in focusing on a particular part of the architecture. The use case also captures the assumptions being made, often for simplification of the problem space[7]. The *Sketch & Scope* step captures the various actors, components, and services that are required to implement the desired design. This step helps identifying the *functional control points* of the socio-economic environment, i.e., the ones directly implementing the technical design. A *control point* here is defined as a point in the environment where control of some sort can be applied, e.g., through centralizing a particular component or enforcing a particular regulatory requirement. This step is extended by the *Deconstruct* step, which lists all control points within the socio-economic environment, extending the functional ones from the previous step.

Until now, the toolkit has provided the designer with an understanding of points where socio-economic triggers could influence the workings of the system. The next step captures the influences on these control points. These *triggers*, like control points themselves, are divided into socio-economic categories that range from user behavior over regulation and business strategy to technology.

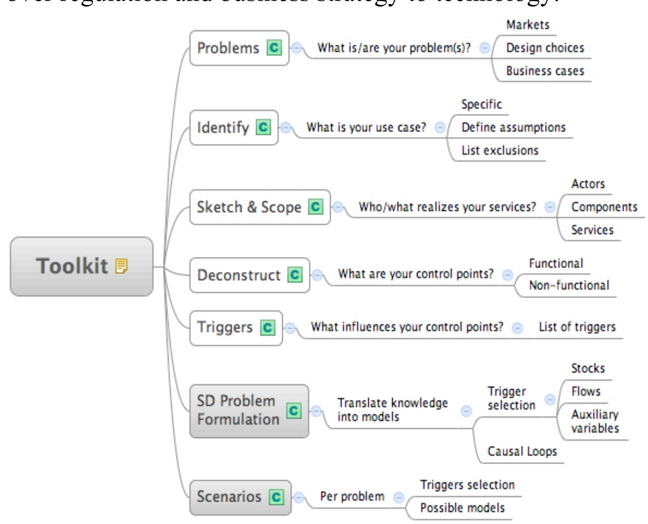

Figure 3: Steps of Our Methodology Toolkit

With the help of the toolkit and its various mind maps, the case developer gathers the necessary information to translate the knowledge into SD models. For that, a return to the Problem step helps selecting the appropriate triggers to formulate an initial set of stock and flow models that represents the formulated design problems (following our methodology in Figure 2). This step also identifies the auxiliary variables for an initial set of causal loop diagrams.

[7] It is this set of assumptions that needs to be carefully recorded for a later revision of the model.

The development of the causal loop diagrams is done separately, based on the information captured via the toolkit. For our purposes, we use the Vensim tool [17] with an academic license, resulting in an interaction between the SD modeling and the recordings captured by the toolkit. Vensim allows for a graphical representation of stock and flow models with causal loops, each of which has an underlying set of equations for each set of auxiliary variables. The tool further allows for manipulation of exogenous factors based on text file input or in real-time. Reference modes can also be constructed in Vensim although we use Powerpoint for our purposes since it is easier to annotate the various phases of the reference mode.

Our design tool work provides the system designer with a set of tools that not only enable our methodology at design time, but also allow for repeatedly evolving the understanding of the system's viability over time. This is achieved through recording the relevant knowledge, which allows for refinements in secondary modeling steps, for instance, after the initial desk research has been performed or parts of the socio-economic environment have evolved beyond the initial setting.

5. EXAMPLE: DISCOVERY

We now walk through an example for applying our methodology with a market focus. For this, we choose the problem of global-scale discovery of, for instance, services or content. We provide a strawman architecture that, as we argue, is representative of a variety of efforts, followed by models and results derived from our methodology.

5.1 Strawman Architecture

Large-scale discovery (or rendezvous) in communication solutions comes with many faces and for many purposes. The main structure, however, is similar throughout most if not all of these solutions. Firstly, discovery is usually implemented in a tiered manner. In other words, a request is sent to a well-known local entity for resolution (tier 3). If the request cannot be resolved, it is forwarded to a local federation (tier 2), which usually represents some form of administrative boundary, such as a corporate environment, an administrative network, or similar. If the request still cannot be resolved, it is sent to other local federations via some form of interconnection structure. The entity performing this interconnection represents tier 1 in the process of discovery. Secondly, tier 3 and 2 entities might choose several next tier entities in the process of resolution. This includes forwarding to different parents for different requests (e.g., based on some local policy). Although such choice does not always exist, it might be important to consider for certain solutions. Requests that cannot be resolved within the interconnection structure are forwarded to other interconnection providers for resolution, i.e., there exists an assumption of interconnecting at tier 1 so as to finally be able to resolve any request (we discuss the impact of this assumption on the socio-economic outcomes later in this paper). This choice can be implicitly implemented through, for instance, address space management or business arrangements, resulting in publishing relevant information in a variety of repositories.

These observations lead us to the strawman architecture shown in Figure 4. The described 3-tiered structure is represented as the local rendezvous point (RP) being tier 3 with its local rendezvous network (RENE) being tier 2, which in turn connects to at least one interconnection overlay (IO) provider as tier 1 in order to deliver the request to the grey RP on the right side of Figure 4.

We assert that this architecture is generic enough to encompass existing solutions but also accommodates future solutions in this space. Examples for existing one are event overlay systems like Siena [18], CORBA [19], the DNS or UDDI for web service discovery. But also recent efforts in new internetworking architectures, such as DONA [20], CCN [21] or PSIRP [22] utilize some form of global rendezvous mechanism[8]. Hence, we can see with these examples and through the study of others that our strawman architecture in Figure 4 adequately describes the structure of most solutions in this space, including the potential choices of parallel tier structures for interconnection. The central role of discovery in all of these efforts makes it an ideal target for showcasing our methodology.

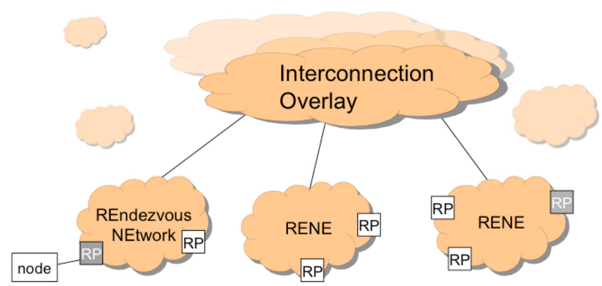

Figure 4: Discovery Example Strawman Architecture

5.2 Main Design Characteristics

Looking more closely at the architecture in Figure 4, we can recognize two major characteristics that will influence the design of most solutions, namely the *existence of distinct players* as well as the *degree of interconnection* between these players[9]. For instance, a higher number of IO providers favors designs with manageable (or low) cost for providing the overlay, while solutions with higher costs for overlay provisioning might still be viable in scenarios with a low number of IO providers. On the other hand, the number of rendezvous networks could provide guidance on required scalability and load balancing for the technical solutions. In addition, the degree of interconnection

[8] Although CCN does not explicitly define a rendezvous service, its inter-domain routing of interest requests represents a similar lookup mechanism as outlined in Figure 4.

[9] Other possible characteristics could be the required degree of deployment and the collaboration between systems having deployed the system and the ones that have not. However, we leave these aspects out of our considerations by assuming a full non-discriminating deployment of our architecture.

between these players determines the fragmentation of regions and therefore markets. This fragmentation could possibly be reflected in the design, for instance, through utilizing hierarchical DHTs or similar. In the following, we focus on these design characteristics, which we can now formulate as *problems* within our methodology as follows:

1. How many IO providers will there be?
2. How many rendezvous networks will there be?
3. What is the incentive to interconnect (either within a single or between several interconnection overlays)?

As outlined in Figure 2, we now translate these problems into a set of stock and flow models from which we derive our possible as well as likely socio-economic outcomes. In other words, the number of interconnection overlay providers as well as the number of rendezvous network providers, together with the incentive to interconnect (normalized between 0 and 1), represents the stocks in our system dynamics models while the flows are represented by the changes in these stocks.

The following paragraphs present the system dynamics models for each of these stocks together with the reference modes[10]. The latter allow for formulating potential socio-economic outcomes while the former capture the major influences within the dynamic environment in which the system will develop over time.

5.3 Model for Interconnection Providers

As outlined in our strawman architecture of Figure 4, an interconnection provider forwards discovery requests, originating in one rendezvous network (and which could not be resolved in that local rendezvous network), to another rendezvous network that is able to resolve the request. More than one interconnection provider can exist in the overall deployment while the number of providers is not limited, as long as technology does not limit the number for performance reasons.

With that in mind, Figure 5 outlines the possible behavior of our SD model for the interconnection provider. Note that the timeline as well as the total number on the y-axis are only indicative and by no means final – such annotation can be found in the actual scenario-based evaluation cases. After the initial deployment, we expect a phase of market interest that is reflected by a growing number of providers. At some tipping point, a phase of commercialization occurs, leading to a competitive market with entrants and exits stabilizing. A phase of consolidation leads to either (a) a single dominant player (lower, dotted curve) (b) a stable but limited number of market players (dashed curve) or (c) commoditization of the interconnection provider function over time (e.g., due to technological advances) as shown with the lower, solid curve. The latter commoditization can also occur without commercialization phase, reflected in the uppermost solid curve, rapidly devaluing the initial market interest.

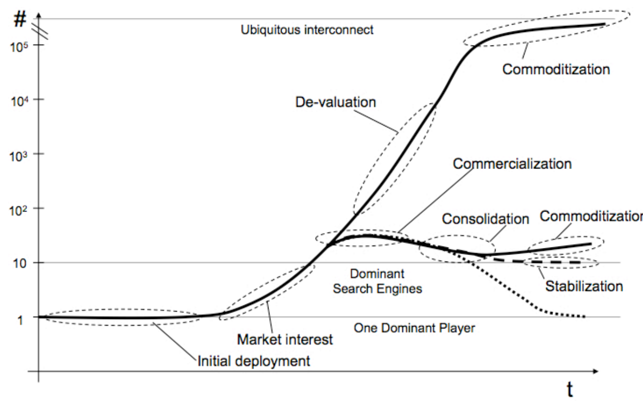

Figure 5: Reference Mode for IO Providers

The outcomes in Figure 5 can be identified as various types of markets that are eventually enabled (or prohibited). Monopoly and oligarchy markets are defined by the lower outcomes in Figure 5, while commoditization of the market is captured by the upper outcomes. Fragmented markets are captured through outcomes for the interconnection incentive in Section 5.5.

Let us move on to the underlying system dynamics model utilized in our evaluation, shown in Figure 6. The model is divided into two distinct parts.

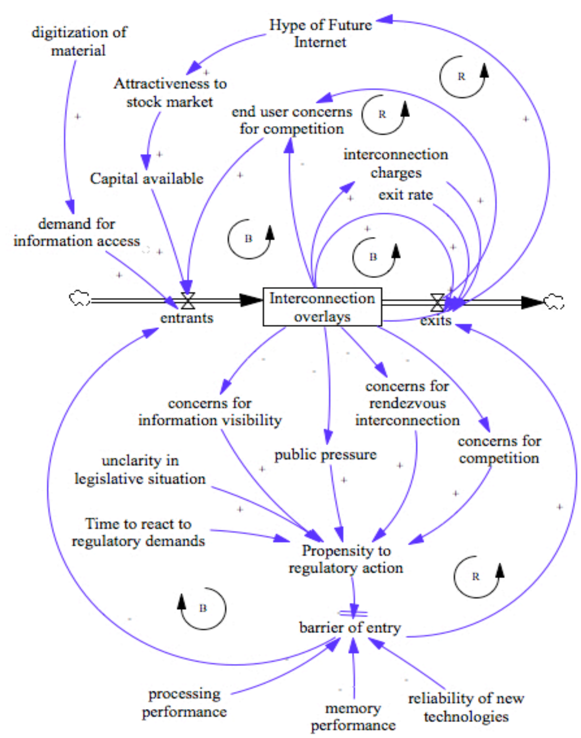

Figure 6: SD Model for Interconnection Providers

[10] The knowledge captured with our toolkit, the resulting SD models as well as the detailed equations are available at http://www.megaupload.com/?d=2ROOGD9U (anonymously uploaded for blind review purposes). The knowledge was captured through a series of interviews with representatives of (two major European) ISPs, a content provider, two major equipment vendors as well as a large European retailer.

The lower part models the various factors that influence the barrier of entry for players in the (IO provider) market. As exogenous factors, memory and processing performance as well as reliability of technology utilized for implementation are influencing the barrier to entry with a variable weight. Furthermore, a set of four major concerns influences the propensity to change the barrier of entry through regulatory action. These are information visibility (e.g., through insufficient interconnection), public pressure, (sufficient) rendezvous interconnection and competition.

All these concerns are modeled as being linearly dependent on the number of IO providers up to a given level of IO providers (after which the concern remains constant at a low level). The propensity to regulatory action is influenced by an exogenous factor that represents the lack of clarity in the legislative situation (e.g., through introducing new concerns, change in procedures). A delayed action is modeled through a smoothened delay, determined by the exogenous factor that determines the time to react to regulatory demands. The resulting barrier to entry multiplicatively influences the entrant and exit rate.

The upper part of the model captures several causal loops. The outer loop represents the availability of investment, influenced by the hype for Future Internet technologies. Furthermore, the entrants are influenced by end user concerns for competition (which in itself depends on the number of IO providers and rate of exit from the market[11]).

As an exogenous factor, the demand for information access influences the entrants in a weighted manner (while such demand is driven by the digitization of material, an exogenous factor in our model). Last but not least, we model a causal relation of interconnection charges and exits from the market, while an exogenous factor determining the general exit rate (e.g., due to capital burnout or other factors) is factored into the exit flow in a weighted manner. We also represent the desire to not interconnect as an exogenous factor into our model. This factor is driven by our model for the interconnection incentive, representing a fully interconnected market through a value of 1.

With this model, we capture regulatory, user-centric, and market causalities as well as certain exogenous drivers that influence the number of IO providers.

5.4 Model for RENE Providers

Local rendezvous networks are responsible for resolving rendezvous requests either sent directly to them within their own network or forwarded by the IO provider from another RENE. Hence, they form regions within the global search space. Such regions may be determined, for instance, by geography or organizational boundaries. Figure 7 shows the reference mode for the number of RENE providers.

The shapes of the various curves are similar to that of the number of IO providers; the difference lies in the annotation of the phases of the curves. Similar to the IO providers, there is an initial deployment phase for rendezvous networks, followed by an uptake in numbers that represents an initial regionalization of rendezvous networks. This could be, for instance, driven by various ISPs providing regional resolution services (being interconnected via an IO provider). The dotted curve represents the demise of regions towards a single (global) region of rendezvous networks. This case questions the need for interconnection through any provider since such RENE would be the only resolution network on tier 2. The middle curves represent the formation of stable regions (note that the actual number here is only indicative and will depend on actual parameterized simulation runs). The dashed curve converges to a stable number over time while the upper, solid curve increases after a time of stability through commoditization of the function. Finally, the uppermost solid curve represents the commoditization (and therefore devaluation) of the RENE provider function over time. This outcome represents an extreme regionalization of the search space although it is important to understand that this does not need to happen uniformly (since certain regions can be fairly large and also competitive).

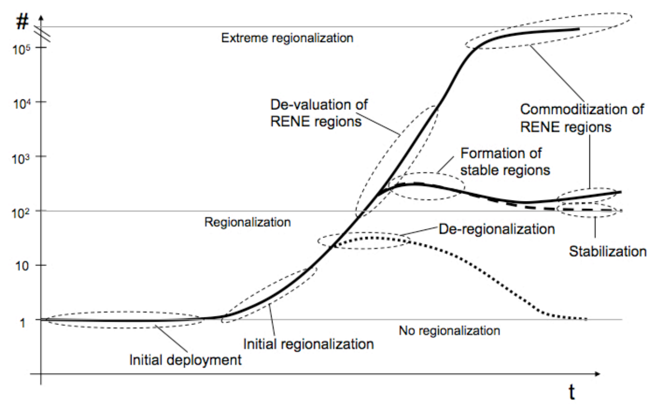

Figure 7: Reference Mode for RENE Providers

The causal loops for the number of RENE providers are shown in Figure 8. The model is divided similarly to that of IO providers. Hence, the lower part of the model represents the regulatory and technological drivers that determine the barrier to entry for new entrants into the market. Here, however, only public pressure and competition concerns are considered on regulatory level, with delay and unclarity similar to those influences in Figure 6.

The upper part of the model includes a capital market loop similar to Figure 6, i.e., the number of providers driving hype in the market, which in turn drives the capital available for new entrants. As exogenous drivers, digitization plays a role to that in Figure 6, while we add the perceived usefulness of this function as an exogenous factor. This is due to the user-facing character of this function compared to the IO provider function. Last but not least, we also include the end user concern for competition

[11] It is important to note that such dependence is not directly linear but instead indirectly defined through a utility function that captures perceived quality through the choice of providers. This utility function is linearized in our equations for the model.

in the model. The model utilizes similar weighted equations to those of the model for IO providers.

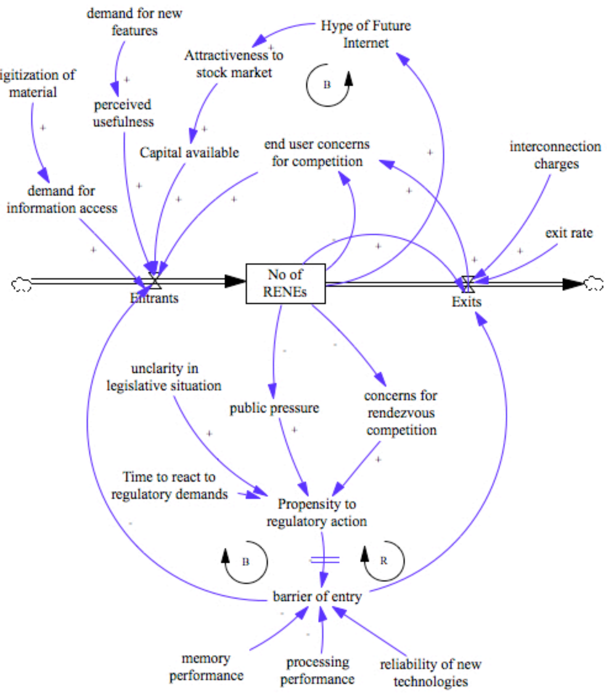

Figure 8: SD Model for RENE Providers

5.5 Model for Interconnection Incentive

Our third main characteristic is the incentive to interconnect individual IO providers with other providers, in the presence of more than one IO provider. The higher such an incentive is, the fuller reachability within the search space is achieved. This is shown in Figure 9 as the upper curve of our reference mode for this stock, with an initial deployment and growing deployment phase similar to the other models. The two dotted curves represent cases in which the adoption either fails entirely (incentives converge towards zero) or almost entirely (convergence is greater but still close to zero). The latter case might occur within an adoption as an intra-domain solution only, such as within largely closed enterprise service buses.

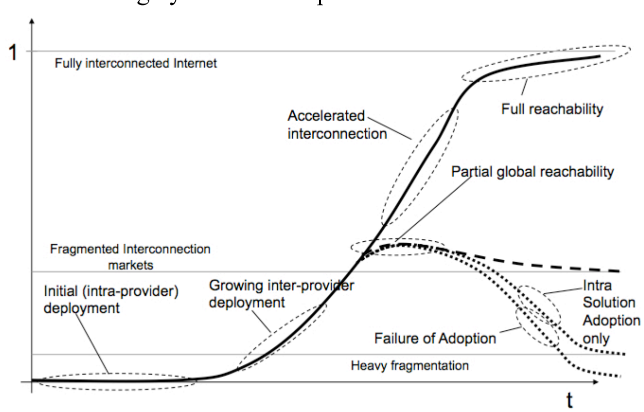

Figure 9: Reference Mode for Interconnection Incentive

In the case of convergence to zero, these buses do not interconnect at all (fully closed) while the second case represents some although very limited interconnection between such systems, for instance, in selected industries like finance. The dashed curve represents that of a fragmented interconnection market. Such an outcome could represent strong interconnections within otherwise largely isolated vertical industries, such as transport, retail. Hence, such fragmentation could occur at a fairly high incentive value, representing strong fragmentation within otherwise highly interconnected markets. This effectively represents isolation of these markets. One reason for such isolation could be regulation. For instance, an isolated search space could be established for emergency services (with virtually no interconnection to the 'rest' of the Internet) for national security reasons. The finance sector is another example of possibly isolated markets (in terms of interconnection).

Figure 10 shows the causal loop diagram for the incentive characteristics. Again, we present the regulatory component in the lower part of the model.

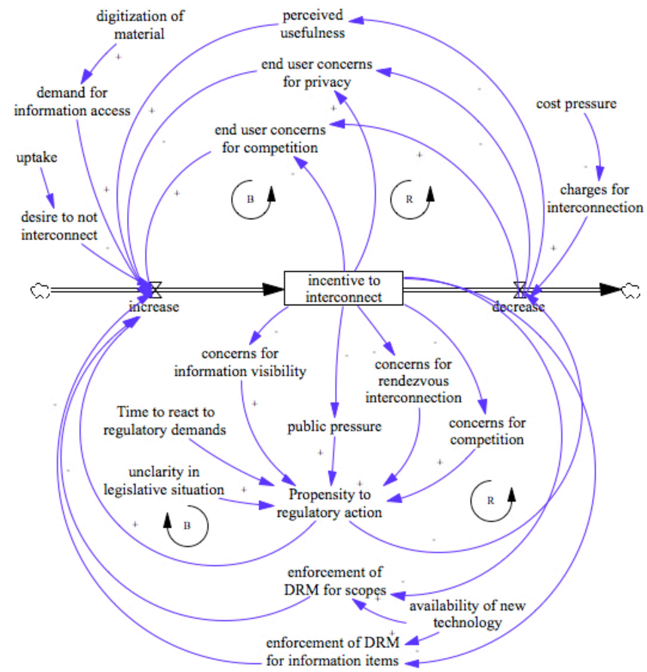

Figure 10: SD Model for Interconnection Incentive

Here, however, the influence of technology is omitted since we do not see a causal relation of technology development and incentives for interconnection due to the economic nature of this decision (to interconnect). However, we included the influence of digital rights enforcement into the model since the interconnection of search spaces determines the dissemination of information (which DRM enforcement intends to control or limit). We differentiate DRM enforcement for information items (which are individual pieces of information) and scopes (which are sets of information, like a collection of movies) in order to accommodate technical solutions that provision for this difference, such as the PSIRP efforts [22].

The upper part of the model largely focuses on end user influences. Perceived usefulness of the rendezvous function to be interconnected, i.e., information being searchable beyond a local rendezvous network, is one driver, while concerns for competition and privacy are captured as end user concerns. As a market driver, cost pressures are taken as exogenous drivers. Digitization of information, as in our other models, is taken as another exogenous factor, i.e., making this information available in an interconnected network. Finally, an exogenous factor for the desire to not interconnect represents cases in which isolation of rendezvous networks and IO providers are defined through, for instance, regulatory means.

5.6 Scenarios and Likely Outcomes

Given the possible socio-economic outcomes for our main design characteristics, presented in Figures 5, 7, and 9, we now highlight scenarios that demonstrate variations in the size and nature of markets (in the case of IO providers), the regionalization of markets (in the case of RENE networks) and the extent of collaboration in terms of interconnection. We present the main drivers of each scenario from a stakeholder perspective before outlining the main parameterization of the models resulting from these drivers. The evaluation results are then tied back to the reference modes, allowing us to reflect on the likely socio-economic outcomes for each scenario. It is important to remember, however, that our evaluation focuses on the system-level behavior under a range of parameter changes. Hence, it is the shape and final outcome of the curves that is important rather than the exact numerical result of each simulation.

For each of our scenarios, we assume the following basic configuration. Each model runs over a lifetime of 20 years. Since we assume a non-discriminating deployment of the solution, an initial delay similar to the current Internet (i.e., the various phases of transfer from a research network over a pre-commercial towards a fully commercial environment) is left out of our consideration. With that in mind, our simulation lifetime represents the lifetime of a commercial market under full deployment capability. For the capital investment available, we assume $20 million per month. This represents a venture capital market driven model of entrants. We omit here the consideration of dividing the entrants into incumbents and new venture-driven entrants. This is left for future work, coupling this consideration with that of incumbent versus new entrant strategies. As the upper limit of entrants in the IO provider market, we assume 20 per month, a number taken from the competitive VPN (virtual private networks) market. We assume a ten-times higher number for RENE providers, given the locality of investment. The reliability of technology factor is assumed to be linearly increasing from 0.1 towards 0.8 over the lifetime of the simulation. This represents the assumed venture-driven entrance while simulating increased maturity of technology over time. For the exit rate of players from the market due to capital burnout issues, we assume 10%.

In the following, we outline each scenario and its motivation. We then describe which model plays a role in each scenario and describe the parameters that each scenario influences in these models. Finally, we present the socio-economic outcomes, i.e., the simulation runs, for each scenario under the varying parameters.

5.6.1 Scenario 1: Anti-Monopoly Movement

This scenario assumes an increasing movement away from various monopolies. This might result, for instance, in increasing number of grass root movements of various forms, as for instance seen in the wireless access space with a number of community schemes [23][24][25].

The parties driving this trend are end users (through public campaigns), legislators (through increased public pressure and the need to address international monopolies), regional powers (not accepting monopolies imposed by other regional powers) and corporations not being successful in establishing themselves as monopolies. The IO provider market is a clear target for such a scenario since possible monopoly or even oligarchy market constellations are not unlikely to occur (see Figure 5 for the possible socio-economic outcomes of this market). In addition, the RENE provider market is also evaluated against this scenario.

Figure 11 shows the outcome of our simulations. In a) and b), we vary the weight of end user concerns into the investment decisions for new entrants from 10% to 80% while changing the regulatory concern for competition from 25% to 75%. We can see that in all runs, an increasing end user concern weight clearly influences the number of IO providers in their final outcome. For the highest concern weight of 80%, we can even see a phase of consolidation towards a slightly lower number. This is not surprising since the end user concerns influence the investment decisions (although not linearly). However, we can still see a substantial investment into the market despite a significant increase of the weight from 10% to 80%, for instance. Hence, a public anti-monopoly movement does have an impact on the market size in terms of players although it does not change its overall outcome, namely that of a significant and stable number of market players. We can see in Figure 11 c) that there is virtually no impact from public pressure on the overall market size (shown for an end user concern weight of 10%).

What is striking is that the number of players decreases with an increasing end user concern for competition. This is probably in contrast to the expectation of seeing an increase of players (and therefore competition) instead. The reason for this behavior is the equation used for the causality that describes the end user concern for competition. As can be seen in Figure 6, both the number of IO providers and the exit rate of players play a role in the concern. In our equation for the end user concern, we weighted the number of exits more strongly than the impact that lower numbers of IO providers would have, since we saw the rate of exits as being stronger in influence (due to its immediate impact on public opinion compared to a longer term trend of change in overall numbers). This explains the behavior observed in Figure 11.

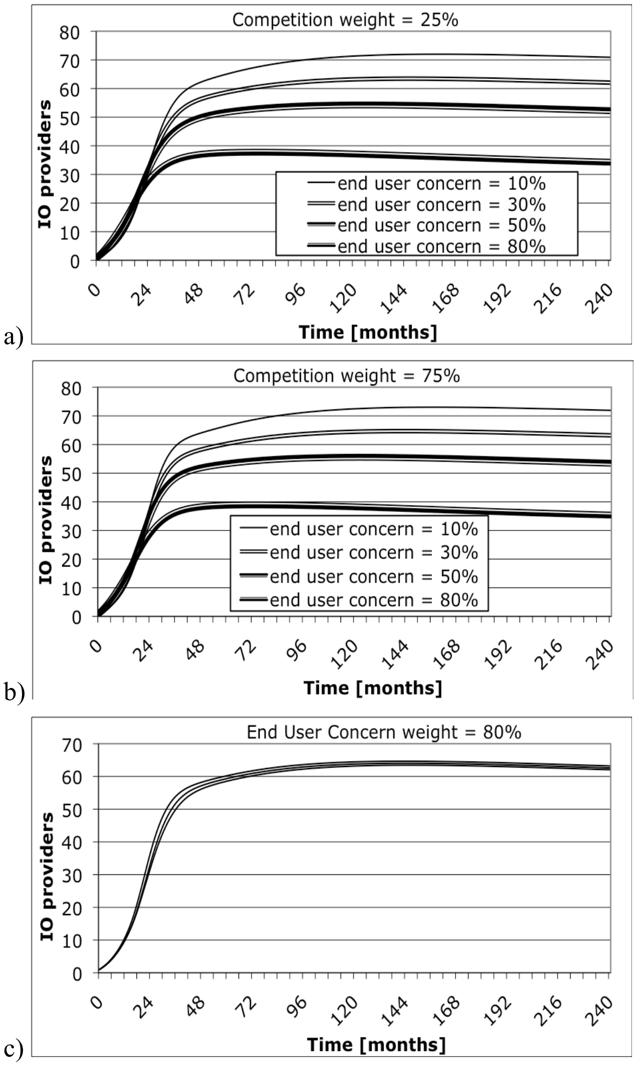

Figure 11: IO provider simulations for Scenario 1

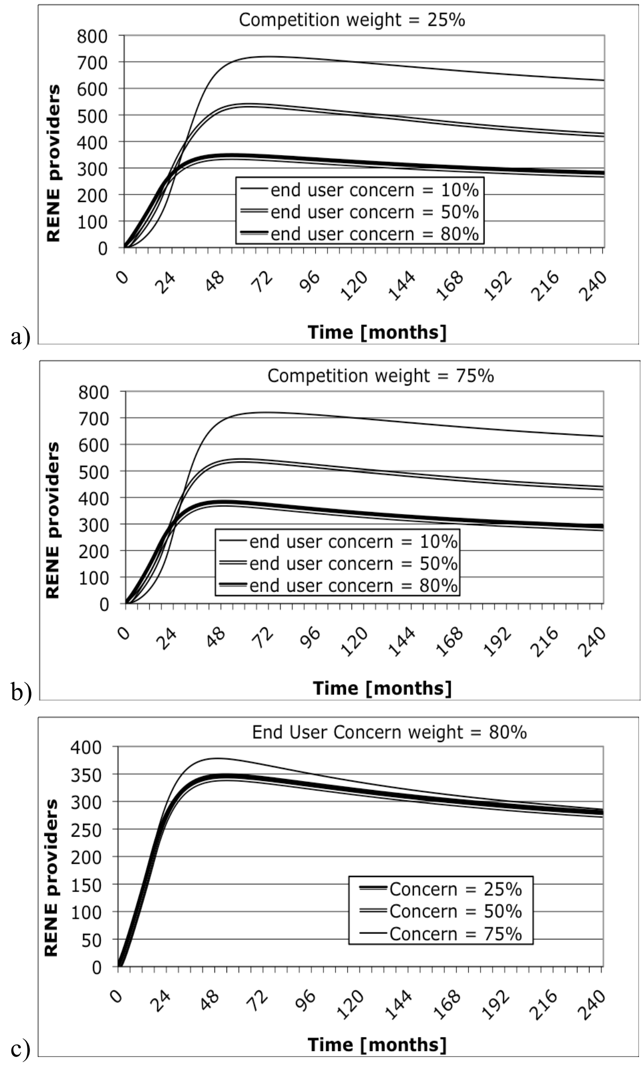

Figure 12: RENE provider simulations for Scenario 1

Figure 12 shows the outcome of our simulations for the number of RENE providers. Similar to the IO provider case, a) and b) show the variation of the weighted end user concerns into the investment decisions for new entrants from 10% to 80% while changing the regulatory concern for competition from 25% to 75%. Similar to the IO provider case, an increasing end user concern weight clearly influences the number of providers.

For the highest concern weight of 80%, we can see in Figure 12 c) that there is a small dependency on the increase in regulatory concern. Overall, the dependencies in behavior are similar to the IO provider market, which can be explained by the similar equations used for the RENE provider model, including the decreasing overall number of providers with increasing end user concern.

Another interesting case is the dependency on differently weighting public pressure on regulation. Due to the similar input into the propensity to regulatory action, the system behavior is similar to that of changing the concern for competition. Hence, we omit the output of the simulations for this case.

5.6.2 Scenario 2: Regional Power Struggles

Regional power struggles can already be observed in the current Internet, such as in name and address space management [26][27]. In addition, the setting of standards for key technologies is often a sign of regional power struggles. This scenario investigates the potential impact of such power struggles on the structure of the rendezvous market. Stakeholders driving this scenario are end users (through perceived superiority of regional values), legislators (through setting policies for strengthening local structures in disadvantage of global ones), corporations (attempting to benefit from such struggles) as well as stock markets (speculating on the outcomes of such struggles).

We model the regional power struggles as changes in various parameters within our interconnection incentive model. More specifically, we demonstrate the impact of regulatory concerns by varying the balance that regulatory actions take in the incentives inflow. In addition, we vary the public pressure concern under these changing influences of regulatory actions. Furthermore, we study the influence that the enforcement of DRM (e.g., due to regional variations in digital rights) has on the overall interconnection incentive.

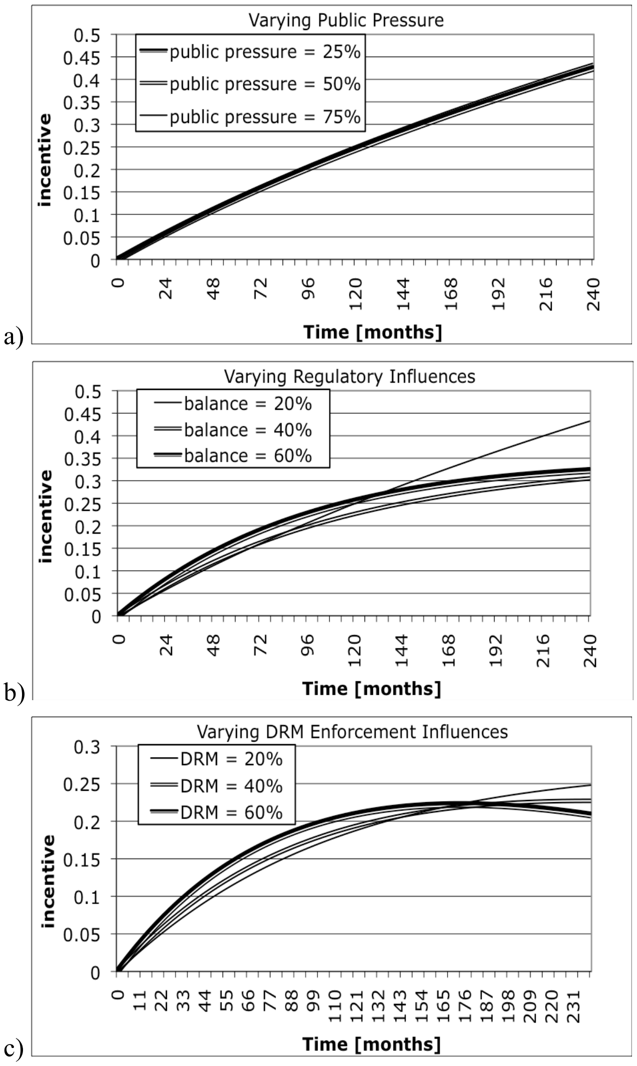

Figure 13: Incentive simulations for Scenario 2

Figure 13 shows the results for the interconnection incentive. In a), we vary the public pressure influence from 25% over 50 to 75%. We can see that there is hardly any difference in incentive uptake for the given change in public pressure. In Figure 13 b), the balance of regulatory influences is changed from 20% over 40 to 60%, i.e., the influence on the increase in interconnection incentive is changed. We can see that higher regulatory influences change the overall outcome of the curve, i.e., it leads to lower interconnection incentives due to the higher regulatory intervention. However, the difference between 40 and 60% is relatively small with the highest balance curve even leading to a slightly higher incentive.

Figure 13 c) shows the influence of DRM enforcement (we only model the enforcement of scope-based DRM since item-based DRM has a similar dependence, given the equations used). In our model, DRM enforcement directly influences the incentives through a weighted factor in the inflow. While the general number in all three cases is similar, the shape of the curve subtly changes from an increase in the 20% case to a downward sloping curve in the 60% case. Running the simulation with a longer lifetime reveals that the 60% case converges to a fragmented market while the 20% case increases towards saturation. Hence, we can observe a strong dependence of our model on DRM enforcement. This is not surprising since DRM is often intended for fragmenting markets.

5.7 Feeding Back Into Design

One of the main goals for applying our methodology is to feed knowledge back into the overall solution design. Although there is no particular design that we have based our work on, there are design considerations that can be derived from our initial results.

Design for flexible interconnection: Our current models assume an interconnection similar to that of the current Internet, i.e., everybody is interconnected with everybody else on agreed charging terms. However, current trends in the Internet, most notably in the search space, treat the exchange of information only as a means of charging for another transaction, namely that of providing advertisement. Such a model can easily be imagined for interconnection in our case. Other charging models could focus on providing the 'right' information in a timely manner rather than 'just' information, for instance, by filling caches in peered networks beforehand rather than on-demand. Such interconnection would significantly change the way interconnection would be charged for.

We assert that the current focus on an Internet-like interconnection charging model is the main reason for not observing commoditization and monopoly outcomes. Technologically however, it is not unlikely that providing rendezvous overlays could become commoditized. We believe that the currently assumed interconnection keeps the market stable in terms of players entering and exiting since it provides a barrier of entry that prevents uptake towards larger numbers of rendezvous overlays (and rendezvous networks). Proxy solutions, i.e., IO providers freely interconnecting on one end and participating in the charged market on the other, could stimulate the market towards different uptakes than those observed in our results. Different charging models, in particular those including caching, would require consideration in the overall design by, for instance, localizing rendezvous functionality for particular information. The solution in [28], for instance, provides such an ability to implement several, possibly more localized, rendezvous points for a sub-set of the information space.

Design for choice: In our models, various causalities on the regulatory and user side emphasize the need for choice in interconnection. This is driven by the fact that rendezvous interconnects information that inherently carries value for end users. This is somewhat different from interconnection in the Internet today, which focuses on resource pooling and therefore cost minimization – differentiation on service or content level is hardly provided. In our discovery example, one could imagine end users publishing certain information in one IO provider, while other information is published via another IO provider[12]. Any solution should consider mechanisms for such choice.

Design for isolation: One expression of choice is a desired isolation of information spaces, each of which is interconnected by its own IO provider. Enforcement of digital rights influences the incentive to widely interconnect within an isolated island of policy enforcement. Such regional power struggles already exist today and are likely to exist in the future. Any design must accommodate these influences on markets.

Design for flexible deployment: The need for evolvability of solutions has long been recognized. Hence, any design should consider various deployment scenarios. One such scenario could assume vertical industries, such as the content industry, retail, health or others, being the drivers for initial adoption. Only after a period of isolation, widespread interconnection might occur. In contrast to the full adoption model, i.e., every player adopting the technology, this is a viable alternative for bringing about the design in a real-world setting. Considering various deployment options should be accompanied by a proper understanding of their market impacts through, for instance, utilizing the methodology of this paper[13].

Decouple business models: Another aspect is that of de-coupling interconnection models on the bit and information level. Coupling occurs, for instance, by routing discovery requests on the same upgraph connections that are being established through bit-level interconnection (such as in [28]). This creates a strong alignment of the business models underlying both interconnections, an alignment that is not necessarily upheld in reality (see our discussion on flexible interconnection models above).

There is a wider architectural context in which our recommendations need to be seen. As Clark et al. outline in [1], the modularity of functions in an architecture is crucial in order for solutions to be flexible within a wide range of socio-economic scenarios. We believe that our methodology has the potential to further the understanding as to what good or bad modularity is, while providing some indicative evidence for such judgment. Part of this understanding is the appropriateness of processes on institutional level, such as standardization and regulation, as well as their potentially required changes (see [15] for an example discussing changes in regulation). We argue that our methodology may provide additional insight by understanding the influences of these processes.

[12] The reader is referred to [29], which discusses the lack of price differentiation in the Internet with an outlook towards an architecture that could provide such differentiation.

[13] Our models can easily be parameterized to result in an 'edge' deployment scenario by artificially choking the number of IO providers, while opening the market to full interconnection after a period of time. We omit this scenario for space reasons.

6. GOING FORWARD

While our example provided a first glimpse into the potential of our methodology, it is clear, however, that we are only at the beginning of a much bigger challenge here.

System design is a truly multi-disciplinary challenge. This has often been said. But it is the societal importance of the Internet today that emphasizes the truth of this statement in any future Internet design endeavor. We recognize that many communities have developed their own methodologies for solution designs, such as [30][31][32], often based on roots similar to the one presented here. What is needed now is to engage in a dialogue with these communities, exploring these various methods and cross-fertilizing them with a growing understanding of how Internet-scale systems could be designed rigorously. The outcome of this dialogue must be a scientific foundation of our understanding of the Internet – possibly leading to the emergence of an *Internet Science* discipline in its own right, with the technical networking community taking a leading role in its development.

This new science needs an educational basis. For such scientific foundation to diffuse, a growing curriculum needs to be developed that educates the future generation in the (growing) scientific understanding and rigor of designing solutions. This educational basis has to be seeded with the junior faculty that is currently defining its research agenda – influencing this new generation is crucial.

This new science needs its codification. We position our work as more rigorously codifying the socio-economic viability of a design, similar to the well-used methods of codifying design through programming code. This will require a growing repository of common methods, an understanding of their underlying formats, and the various visualization methods being used by different communities. Undoubtedly, this will need certain forms of agreement similar to that of programming guidelines or even standardization of common formats. Without such agreement on codification, little progress will be made in improving the knowledge of a design's viability over time.

This new science needs its successes. Our example can only highlight the potential benefits that a rigorous analysis of a design might bring. More is clearly needed. Applying methodologies like ours to a growing set of use cases is the only way of convincing even the grimmest opponents. It provides the anecdotal basis from which one can argue about applicability and success of these methodologies. A set of such use cases is that of addressing key architectural questions, such as the one raised in [33]. A *control point analysis*, as argued for in [33] and provided by our work, can bring about the desired progress in such architectural debates. Combining this need for successes with the educational challenge is possibly the shortest route to success, having a growing set of students examine various designs through a growing number of methodologies. This could provide the required leap to prove the claimed rigor that is introduced by methodologies like ours.

7. CONCLUSIONS

As technical designers, we not only create technical but eventually social artifacts that are immersed in our society. Understanding the social fabric and the conflicts between members of this fabric is crucial when trying to determine the potential success of one's design. As a community, we need to recognize the power that designing solutions inherently brings, and we need to find ways to formalize that power and articulate its potential outcomes.

This paper has provided means for both the recognition of and the means for formalizing the tussles and their possible outcomes. We have provided a methodology for design as well as an insight into tools developed for gathering evidence within such methodology. We also exemplified the application of this methodology with a use case that is relevant in many ongoing architectural efforts.

But possibly the most important outcome is the recognition that we need to move from stove-piped design solutions and methodologies towards a concerted multi-disciplinary effort. The outcome of this effort could potentially formulate the basis of a new Internet Science discipline; a discipline that provides the required rigor to argue about the viability of large-scale system designs. We as technical designer must take a role in its creation and provide our input into its growing pool of methods.

8. REFERENCES

[1] D. Clark, J. Wroclawski, K. R. Sollins, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," in IEEE/ACM Transactions on Networking, Vol. 13, No. 3, 2005

[2] M. S. Blumenthal, D.D. Clark, "Rethinking the design of the Internet: The end-to-end arguments versus the brave new world", ACM Transactions Internet Technologies, vol. 1, no. 1, Aug 2001

[3] P. Faratin, "Economics of Overlay Networks: An Industrial Organization Perspective on Network Economics," NetEcon '07: Workshop on the Economics of Networks, Systems and Computation, June 2007

[4] J. Rexford, C. Dovrolis, "Point/Counterpoint Future Internet Architecture: Clean Slate Versus Evolutionary Research", Comm. of the ACM Vol. 53, No. 09, Sept. 2010

[5] I. Licata, A. Sakaji (eds), "Physics of Emergence and Organization", World Scientific & Imperial College Press, 2008

[6] K. Eason, "Information Technology and Organizational Change", Taylor &. Francis, London, 1988

[7] H. R. Hertson, D. Hix, "Toward empirically derived methodologies and tools for HCI development", International Journal of Man-Machine Studies no. 31, 1989

[8] V. K. Murthy, E. V. Krishnamurthy, "Multiset of Agents in a Network for Simulation of Complex Systems", in "Recent advances in Nonlinear Dynamics and synchronization, (NDS-1) -Theory and applications", Springer, 2009

[9] J. W. Weibull, "Evolutionary Game Theory", MIT Press, 1997

[10] L. F. Fernandez, H. S. Bierman, "Game theory with economic applications", Addison-Wesley, 1998

[11] M. Olson, "The Rise and Decline of Nations: Economic Growth, Stagflation, and Social Rigidities", Yale University Press, 1982

[12] K. E. Wiegers, "Software Requirements (2nd ed.)", Microsoft Press, Redmond, WA, 2003

[13] D. Gernert, "Ockham's Razor and Its Improper Use", Journal of Scientific Exploration, Vol. 21, No. 1, pp. 135–140, 2007

[14] J. D. Sterman, "Business Dynamics: Systems Thinking and Modeling for a Complex World", McGraw-Hill Higher Education, Boston, 2000

[15] C. Vaishnav, A. Khakifirooz, "Does Technology Disruption Always Mean Industry Disruption?", International Conference of System Dynamics, 2008

[16] XMind software at http://www.xmind.net, 2010

[17] Vensim software at http://www.vensim.com, 2010

[18] D. Rosenblum, "A Tour of Siena, an Interoperability Infra-structure for Internet-scale Distributed Architectures," Ground System Architectures Workshop, 2001

[19] OMG, "The Common Object Request Broker Architecture", at http://www.omg.org, 2009

[20] T. Koponen et al., "A Data-Oriented Network Architecture", Proceedings of ACM SIGCOMM, 2007

[21] V. Jacobson, D. K. Smetters, J. D. Thornton, M. Plass, N. Briggs, R. L. Braynard, "Networking Named Content", Proceedings of ACM CoNext, 2009

[22] D. Trossen, M. Sarela, K. Sollins, "Arguments for an Information-Centric Internetworking Architecture", ACM Computer Communication Review, April 2010

[23] P. Mahonen (ed), "EIFFEL whitepaper", EIFFEL group,2006

[24] Ian Brown (ed), "Towards a Future Internet – State of the Art report", Study for the European Commission, 2009

[25] Rand Corporation, "Trends in connectivity technologies and their socioeconomic impacts", Technical Report 776, 2009

[26] J.Q. Anderson, "The Impact of the Internet on Institutions in the Future", Pew Research Center, 2010

[27] J. M. Rao, "Internet Governance: The Struggle over the Political Economy of Cyberspace", On The Internet, 1999

[28] J. Rajahalme, M. Särelä, K. Visala, J. Riihijarvi, "Inter-Domain Rendezvous Scalability in Content Networking Architectures", TR09-003 at http://www.psirp.org, 2009

[29] D. Trossen, G. Biczok, "Not Paying the Truck Driver: Differentiated Pricing for the Future Internet", Proceedings of ReArch workshop at the ACM CoNext conference, 2010

[30] W. Bijker, T. Hughes, T. Pinch (eds), "The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology, MIT Press, 1987

[31] R. Volti, "Society and Technological Change", 5th ed., St. Martin's Press, 2005

[32] D. Tapscott, "Grown Up Digital: How the Net Generation is Changing Your World", McGraw Hill, 2009

[33] I. Brown, D. Clark, D. Trossen, "Should Specific Values Be Embedded In The Internet Architecture?", Proceedings of ReArch workshop at the ACM CoNext conference, 2010