

Modelling web based socio-technical systems through formalising possible sequences of human experience

Rob Walton

Dept. of Engineering Science
Oxford, United Kingdom
rob.walton@eng.ox.ac.uk

David De Roure

Oxford e-Research Centre
Oxford, United Kingdom
david.deroure@oerc.ox.ac.uk

ABSTRACT

When people interact through technical infrastructure such as that of organisations or the World Wide Web, this infrastructure will change and in some cases new identifiable structures or even ecosystems may emerge. Examples of such emergent socio-technical systems on the Web include some social machines and phenomena such as echo chambers. To model these complex social systems we develop a method for formalising the moments, or occasions, of experience of an individual person by contriving a ‘chemistry’ encoding possible sequences of their external stimuli, internal experience and reactions to this internal experience. We take a process oriented approach and formalise this as a stochastic Petri net. We wire together a number of these to form a fixed social network in which experience is shared. The resulting model unfolds into many possible causal graphs of occasions of experience which we show using an interactive visualisation. We demonstrate the utility of this method by encoding models exhibiting information diffusion and what we call multiple phase diffusion, and then consensus formation, before encoding a mechanism of echo chamber formation. We then demonstrate the conflation of individuals’ positions on otherwise separate issues through emotion. The approach results in a single Petri net model which may be analysed using qualitative and quantitative techniques supporting web science research. It provides a way to describe and reason about the internal experience of individuals within multi-scale socio-technical systems.

ACM Reference format:

Rob Walton and David De Roure. 2019. Modelling web based socio-technical systems through formalising possible sequences of human experience. In *Proceedings of 11th ACM Conference on Web Science, Boston, MA, USA, June 30-July 3, 2019 (WebSci '19)*, 10 pages. <https://doi.org/10.1145/3292522.3326049>

1 INTRODUCTION

When people interact through technical infrastructure such as that of organisations or the World Wide Web, this infrastructure and the social networks that form around it may coexist and adapt together. When this occurs we may identify a socio-technical system,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WebSci '19, June 30-July 3, 2019, Boston, MA, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to Association for Computing Machinery.

ACM ISBN 978-1-4503-6202-3/19/06...\$15.00

<https://doi.org/10.1145/3292522.3326049>

examples of which, on the Web, include some social machines [21] and phenomena such as echo chambers. In some cases new identifiable structures or ecosystems may emerge [5]. Understanding the mechanisms behind these behaviours and the life-cycle of the resulting structures is important if we are to model the Web and shape its evolution.

The Web exhibits interesting and difficult to model behaviour across multiple scales. We might consider the Web itself a single piece of technical infrastructure and ask how its structure evolves alongside that of a single social network coevolving with it. At this level we see an ecosystem of social machines competing and cooperating for the attention of individuals and for access to their parts of this social network. Within some of these social machines we see the same processes forming local communities with particular interests or goals [5]. Within some of these we see similar behaviours acting to drive individuals to form short-lived communities around topics as fleeting perhaps as a passing joke or a disagreement over spelling. There is no such absolute hierarchy in reality though: topics or ideas spread through these structures in a peer to peer like way, and longer lived structures such as echo chambers form among them, some with influence from outside.

Early work to understand and learn to optimise socio-technical systems focused on organisations at the scale of businesses. The Web and the social structures on it may be considered instances of Large-Scale Socio-Technical Systems identified by Thomas Hughes’ system theory [23] and described by Igor Nikolic as “a class of systems that span technical artefacts embedded in a social network, by which a large-scale, complex socio-technical artefact emerges” [17]. At this scale, a lot of research has gone into modelling electricity systems all the way from power generation and distribution up to regulatory frameworks. Modelling these systems is complex in that more than one formalism, possibly described with different languages from different fields of study, is required to capture all relevant aspects of the system. Bringing these together requires simulation in the form of agent based modelling.

John Holland calls some of this type of modelling *exploratory modelling* which he describes as: “starting with a designated set of mechanisms, such as the various bonds between amino acids, with the objective of finding out what can happen when these mechanisms interact” [12]. Schelling showed in his study of segregation that for spatially segregated subpopulations to form there is no need for individuals to avoid others: it is sufficient that individuals have a weak preference to be surrounded by those like themselves [19]. More recently and with access to computers Starnini et al. asked which behaviours are sufficient for the metapopulations and echo chambers to emerge in populations of mobile agents [25]. They

studied the interplay between mobility, Schelling’s homophily and social influence in models of opinion dynamics.

Exploratory modelling is powerful but models with analytically derivable or empirically measurable properties have historically had more impact. For example, graph growth algorithms based on properties of generic networks, such as Watts and Strogatz’s small-world property [26] or power-law degree distributions, are in common use. Also, Girvan and Newman identified a basic property of community structure in social and biological networks [8]. Such models are powerful, as unlike exploratory models, they provide mathematically solid foundations for further work. For example, algorithms for identifying community structure in networks are often based on properties of classes of randomly growable graphs. In addition to these idealised physical or behavioural based models, statistical models such as those describing the recurrence patterns of words or topics in individual writing or collective discussion[3] are derived from empirical measurement. Holland describes exploratory modelling as a first step toward building such analytical models.

Improved models for understanding the socio-technical systems which form on the Web might help us understand and quantify the changes as they are adopted by users and society. They might help us build social machines which have desirable properties such as encouraging discourse between otherwise polarised social groups or which are not ‘owned’ centrally [10]. Better theory might help us keep such decentralised social machines answerable to society and prevent capture by external organisations. We might identify echo chambers and determine if they are natural or driven by external agenda and we might further investigate the formation mechanisms of filter bubbles [24]. We might ask in the long term how far into our lives these structures might reach, expanding as they will into Internet of Things (IoT) devices and administered increasingly efficiently by algorithms and artificial intelligence, and what the limits to their influence over our behaviour may be. This helps ensure that the new futures the Web offers are shaped with society’s most positive aspirations and not left to selection by chance or short-term thinking.

Our goal is to devise a method for formalising enough of the aspects governing a person’s internal and social behaviour on the Web to build generative models of individual and collective human experience. The method should be suitable for modelling the emergence of new structures as social networks co-evolve with technical Web infrastructure. The resulting models should be explorable through simulation but may also, in limiting cases of symmetry or low complexity, have analytically derivable properties. These properties might define for example a rigorous relationship between a model describing operation at the scale of coherent interacting communities and at the level of interactions of individuals within or across these community boundaries.

We propose a model with a formal place for capturing moments, or occasions, of human sense awareness. Given the importance of modelling the behaviours that lead to the emergence of structures on the Web this formalism is based on process rather than material or ‘solid’ systems. The model has two layers. On the top we propose the use of Alfred Whitehead’s *Philosophy of Organism* as a framework and language for modelling our large-scale socio-technical systems. This is a process philosophy reasoned from Whitehead’s

observations of the limitations of the way we perceive reality. To produce mathematically tractable models we map this layer down onto Petri nets. The mapping is incomplete which limits the expressiveness of the resulting modelling framework. However it is complete enough to capture a range of behaviours as demonstrated by the simple models developed and verified in Section 3.4. Further background on Whitehead’s *Philosophy of Organism* is given in Section 2.1 and further background on Petri net modelling is given in Section 2.2. The method is developed in Section 3.

Results from simulating the verification models of Section 3.4 are shown in Section 4 in the form of causal graphs of occasions of experience. These models demonstrate, on a fixed social network, diffusion and what we call multiple phase diffusion, percolation and then consensus formation. They also show that an accepted mechanism of echo chamber formation encourages the separation of opinion on networks with pre-existing community structure. We also demonstrated the conflation of individuals’ position on otherwise separate issues through emotion. These simple models build toward analysing the development of echo chambers. They can be considered to form an initial collection of primitives that may be combined to produce more realistic models.

2 METHOD BACKGROUND

Our method depends on two components not described above. We introduce Whitehead’s *Philosophy of Organism*[27] and Petri nets here.

2.1 Philosophy of Organism

We chose Alfred Whitehead’s process metaphysics as the top layer of our modelling framework for two reasons. First, Whitehead describes a system capable of explaining how events of sense awareness form and relate, thus providing a formal way of handling this. Second, although there are other choices that meet this goal, this moves us away from speaking of the boundary of a socio-technical system or of them as material things at all, and instead to speaking about interacting processes.

Whitehead describes the coming and going of acts of sense awareness. Each such act he considers unique. He also considers them ultimately private. For example, we may not quite know what went into building our internal experience of a mug on our desk, but when we look away and that act ends we see that it was clearly unique, and that the subjective experience constituent to it is private, not only to us as an individual, but to that act itself. Whitehead labels these moments of sense awareness *occasions*. An occasion of new experience, Whitehead argues, comes to us at some level pre-formed like a droplet. Only after we experience this may we analyse it to see what went into its formation. Whitehead describes the mechanisms behind the formation of these occasions, or droplets, of experience.

Whitehead’s *Philosophy of Organism* is a process rather than material metaphysics. We believe it provides a very useful way to understand ecosystems of socio-technical systems in which defining solid material-like boundaries between parts is not possible. In developing his system Whitehead conceived no pre-existing ‘screen’ onto which these drops of sense-awareness or occasions are projected. Instead these are continually and actively pulled as

much as pushed into a model of the present which is always based on prior experience. Whitehead says that ‘what we perceive as present is the vivid fringe of memory tinged with anticipation’. As such this model is of a short-lived slice through reality rather than an instant of zero length. Our perception of reality, and Whitehead actually argues, reality itself, being always a duration through time, is necessarily of things going on.

The only physical entities in this system of process metaphysics are occasions. These are the atomic building blocks of Whitehead’s view of both physical and perceptual reality. They form through a process Whitehead labels *conrescence*. They are the transient *creatures* of his *Philosophy of Organism*. An occasion is rooted in the branching causal networks of the upstream occasions which was its cause. In the case for mental rather than physical systems, pieces of a-priori knowledge and propositional statements will also have been folded in to these causal roots. As each occasion of awareness exists only in its capacity as an addition to a pre-existing model of the present, they often form part of a process of building some aspect of this model.

Where a procession of similar but evolving occasions form, for example, those formed while watching a leaf drop, or as a family regroups every year to recreate Christmas, Whitehead identifies a longer lived *process*. Once formed and experienced each occasion in this process dies. The leaf falling though still attracts attention and Christmas comes again. Such occasions as these then are destined to form the ancestors to the next generation of ‘themselves’. There is a definite causal relationship in which such a subsequent occasion occurs after the former. Whitehead describes across eight categories the rules governing which occasions may form and how they do so. Key to understanding the scenario presented here is that two occasions will not form from exactly the same cause.

In this way, each occasion is what Whitehead calls a *superject*, both subject and object. While living, or forming, an occasion builds a subjective view of the objective data formed by its predecessors. Once it has formed then it becomes just another object to subsequent occasions. The predominantly held material view, that one system forms a subjective view of another, or that one system measures or observes another, is no longer primary. This is a big change from the way we are used to understanding our reality.

Now the operation of a process of related occasions coming and going is actually more interesting than this. While an occasion is forming from what Whitehead calls its *actual world* (the collection of dead occasions which are its potential cause) a number of neighbours will be forming contemporaneously (that is with no way of influencing each other but with overlapping actual worlds). These neighbours will be ‘picking out’ different combinations of occasions from the actual world for inclusion. Each such combination results in not one but a set of occasions, for while forming an occasion not only ‘picks’ which upstream occasions to include, but from the elements that went into forming those. Whitehead says that each such element is *prehended*. It is to these acts of prehension themselves which a prehension reaches back toward (possible through the *medium* of a series of occasions) bringing their data forward into the occasion of which it is a part.

Whitehead’s approach results in structures of such processes of occasions, or more broadly a nest-able partial hierarchy of *societies* of occasions. Whitehead identifies a collection of causally related

occasions as a *duration*. Every duration contains other durations and every pair of durations is part a broader duration [15, chapter 13]. For example, we might pick out processes of causally related occasions associated with work and home life and section these into separate durations. Zooming ‘outwards’ we might find these contained by the duration of our entire lives and then that of life on Earth. Zooming back ‘inwards’ from here we find processes such as staying alive as a society, staying alive as an individual or maintaining good hygiene and relationships, which are all prerequisite to at least the duration containing work events. Thus we see that processes of occasions may be situated ‘next to each other’ and may depend on each other. This interdependency results in what Whitehead considers a society of occasions, or of overlapping societies of occasions. The resulting structure is suitable for modelling processes at the scale of the Web while at the same time is key in Whitehead’s understanding of individual and communal perception.

2.2 Petri net modelling

Petri nets, or place/transition nets, are a graphical tool used for the description and analysis of systems with concurrent actions. Invented by Carl Petri in 1939 for describing chemical processes [18] they were used a lot during the development of early computers. There are more general formalisms including process algebras such as Milner’s π -calculus but we feel Petri nets provide a good balance between ease of use and modelling power. Their representation in graph form provides an intuitive formalism describing the evolution of distributed systems such as the large-scale socio-technical systems on the Web we seek to develop methods to model. Beneath this they may be represented as an incidence matrix around which have developed a collection of mathematical techniques for qualitative and quantitative analysis. As shown in the remainder of this article they are also capable of encoding sufficient aspects of Whitehead’s *Philosophy of Organism* to model well characterised behaviours of the Web.

2.2.1 Petri net definition. As shown in Figure 1 a Petri net is a bipartite graph comprised of places, shown as circles, and transitions, shown as rectangles, connected with directed edges. The top half of Figure 2 show the standard elements of common Petri net definitions. The state, or *marking*, of a Petri net is represented by tokens inside the places and shown as black disks. A transition *fires* if the places feeding into it through the net’s directed arcs all have tokens. This firing is a single atomic event in which a token is removed from each input place and a token is placed in each place connected by outgoing arcs. Tokens need not be conserved. If more than one transition is enabled then the order of firing is random. Basic Petri nets have no encoding of time. If the Petri net of Figure 1 were simulated t_a would first fire causing the tokens in x_{p1} and x_{p2} to disappear as a token appears in x_a . The transition t_b would then fire causing the tokens in x_a and x_{p3} to disappear as a token appears in x_b .

A Petri net PN may be represented more formally as places P , transitions T , input and output arcs A , and an initial marking M :




Figure 1: A generative Petri net encoding the potential to experience two occasions a and b identified by blue surrounds. Each is modelled as a transition t with out place x indicating that a transition has been experienced. The input places to an occasion's transition represent its prehended occasions.




Figure 2: Common Petri net elements above shorthand used in this article

$$\begin{aligned}
 PN &= (P, T, A, M) \\
 P &= \{p_1, p_2, \dots, p_l\} \\
 T &= \{t_1, t_2, \dots, t_m\} \\
 A &\subset \{P \times T\} \cup \{T \times P\} \\
 M &= \{\mu_1, \mu_2, \dots, \mu_n\}
 \end{aligned} \tag{1}$$

where μ_i are the integer number of tokens at a place [2].

A net may be *simulated* resulting in a possible sequence of firings. There are also many analytical techniques for testing if certain markings are reachable for example, or if a net will remain live, or if the number of tokens at any place is bounded [2]. Going in the reverse direction Aalst et al. pioneered the concept of process mining, where a estimates of the net describing the interacting process that result in a given *log* of events may be inferred [1].

2.2.2 Stochastic Petri nets. Petri nets run without time, the result of a simulation being simply a sequence of changing markings or *log* of firings. Molloy extended the model of Equation (1) to include

a set of transition rates $\Lambda = \{\lambda_1, \lambda_2, \dots, \lambda_m\}$ [16]. Where λ_i are, possibly marking dependent, rates for exponentially distributed transition firing times:

$$f(\tau|\lambda) = \begin{cases} \lambda e^{-\lambda\tau} & \tau \geq 0 \\ 0 & \tau < 0 \end{cases} \tag{2}$$

In the resulting continuous-time stochastic Petri net, when a transition becomes enabled an activation time, τ , is chosen randomly and the transition will fire only if the transition remains enabled for this time. The effect is that an enabled transition will fire continuously and independently at a constant average rate while enabled—acting as a *Poisson point process* with average rate λ .

The transition firing rates of the stochastic Petri nets ultimately only affect the ordering of transition firing of the standard Petri net, but in doing so they open up new interpretations of models. Molloy showed that with some limitations a stochastic Petri net with exponentially distributed firing rates is isomorphic to a finite Markov process. These models may also be converted to a set of ordinary differential equations which may be solved numerically, or in simple or symmetric cases, analytically. For example, stable or relatively stable states may be identified using approaches from nonlinear systems analysis.

2.2.3 Coloured Petri nets. Kurt Jensen introduced a generalisation of Petri nets and showed that this can be used to describe and analyse complex systems[14]. In this extension, tokens may have different *colour* or type. This colour may, and will in this article, be associated with an instance of a repeating unit. Places and arcs are given expressions limiting which coloured tokens may enter or traverse them. Transitions may be given optional guard functions which determine if the combination of coloured tokens at their input places meet constraints additional to these.

Coloured Petri nets then can be used to model systems succinctly with repeated structure. In the field of systems biology for example Gao et al. used Coloured Petri Nets as a method for modelling multicellular organisms across multiple scales[7]. They applied this approach to modelling planar cell polarity (PCP) signalling in a *Drosophila* wing. PCP is the coordinated alignment of cell polarity across a plane of tissue, describing for example the emergence of patterns on the butterfly wing or of our fingerprints.

As an example, Figure 3 shows a coloured Petri net defined in the Snoopy software used for this work and introduced in Section 3.3. The ‘Unit’ identifier on the places is the set of colours those places may take, in this case the colours $\{1, 2, 3\}$ representing three people. The red text describes, as a multi-set, the colours of the tokens in each place: x_0 has 1 token of colour 2 and one of colour 3, while x_a has one of colour 1. The transition t_a accepts a token of any colour u from the set *Unit* and moves a token of the same colour to place x_a . Without the read arc this coloured Petri net would unfold into three separate units. The read arc presents tokens to the transition via the variable n . The *is_neighbour* function on the transition guards it from firing unless the colours of the tokens in variables u and n meet some criteria.

Looking forward to Section 3 in our work we used a function that enables the transition only if the token in n and the token in u are from neighbouring cells. Figure 4 shows this model unfolded onto the linear graph $n_1 \longleftrightarrow n_2 \longleftrightarrow n_3$. This would be done by




Figure 3: Representation of the diffusion model as displayed in the software package Snoopy. Figure 5a shows the same structure.




Figure 4: Diffusion model unfolded onto a graph
 $n_1 \longleftrightarrow n_2 \longleftrightarrow n_3$

specifying $\text{is_neighbour}(u, n) = (u=1 \ \& \ n=2) \mid (u=2 \ \& \ (n=1 \mid n=3)) \mid (u=2 \ \& \ (n=1 \mid n=3))$.

3 METHOD DEVELOPMENT AND VERIFICATION

In our goal to model large-scale socio-technical systems on the Web, particularly social machines or echo chambers, we needed to model many interacting individuals. Our approach was to develop generators of possible sequences of human experience, not unlike the n -gram language models introduced by Claude Shannon [22] but for experiences rather than words. We present this ‘idea chemistry’, based on a simple reading of Whitehead’s Philosophy of Organism, in Section 3.1. We then expand this into a ‘social idea chemistry’ presented in Section 3.2 in which the experience of a neighbour may be followed. Whitehead’s Philosophy of Organism is involved, as any attempt to fully describe reality must be, and our goal was to capture its essence.

3.1 Idea chemistry

In this method, we start building a model by identifying an alphabet O of occasions that may be experienced by an individual. Occasions representing the experience of an idea are labeled with $\{a, b, c, \dots\}$. In this paper we limit these to the experience of statements that might form part of an argumentation map. Where an occasion represents a feeling of emotion, or similar, it is taken from $\{\alpha, \beta, \gamma, \dots\}$. When two occasions represent opposing reactions to a common choice o they are labelled o^+ and o^- .

The experience of an occasion depends on the upstream occasions which are its cause. We capture this in a limited form by formalising which occasions may follow one another. As shown in Figure 1 we represent an occasion o as the combination of i) a transition t_o connected to ii) an output place x_o indicating that the occasion has been experienced and iii) a list of prehensions required for the occasion in the form of input edges to this transition. An occasion’s input places are the output places of the occasions it may be caused by. In this example the individual will first experience occasion a before then going on to experience b . This implementation does not capture the way that the formation of an occasion from a given upstream structure could result in differing subjective experience due to the inclusion or exclusion of parts of this structure. With this method we might approximate this by adding each of these possible versions of an occasion to the alphabet O .

3.2 Social idea chemistry

The mapping from Whitehead’s model onto Petri nets just proposed affords no social interaction. This mapping was expanded to attribute to each individual person a society of possible occasions which they might experience if both their past experience and now also their environment—formed exclusively by their neighbours in this article—is suitable. We encoded this ‘social idea chemistry’ as a unit U repeatable on a coloured Petri net, and coloured this with a set of colours, one per individual.

We included just one type of social occasion: the occasion of experiencing an occasion that a neighbour has experienced, or *following* a neighbour. This type of occasion depends on the underlying technical medium M capable of communicating experience—a fixed social network with n nodes throughout this work—and the underlying abilities of each individual to both share an experience and to experience what is shared. To create the final *system model* S we connected together the n coloured ‘copies’ of U by constructing an is_neighbour guard function on U ’s transitions in order to represent the social network of the medium M .

For example, and already introduced, Figure 3 shows an example of U on a network of three individuals $n_1 \longleftrightarrow n_2 \longleftrightarrow n_3$ exported from the software package Snoopy introduced in Section 3.3. It represents a single occasion of experience which may be experienced just once, such as that of first experiencing a particular rumour and possibly a reaction to this. Figure 4 shows its expansion onto a non-coloured Petri net with numerical subscripts labelling the three individuals. This model is described in more detail in Section 3.4.1 and encodes an information diffusion process.

Figure 5a shows a shorthand form of this model. The dotted arc indicates a connection from a neighbour’s place, and the open circle arrow head that it is a read arc. The one in the black circle indicates that a guard function on the transition will block firing unless one (or more) neighbours has a token in place x_a indicating that they have experienced the occasion a and are in the position to share it.

In summary the resulting mathematical objects were:

- (1) A repeatable coloured Petri net unit U encoding the possible sequences of occasions of experience of an individual. This might be ‘compiled’ from a directed graph of possible occasion orderings, but in this work was hand crafted.

- (2) A set of colours, one for each of n individuals with which to colour U .
- (3) A medium graph M representing a fixed social network with n nodes.
- (4) A complete coloured system Petri net model S obtained by calculating the `is_neighbour` guard function on U from M .

3.3 Simulation software

The Snoopy suite of Petri net modelling software¹ was used in the early development of these methods [11]. We then created a Python based package, `occmodeler`², for more quickly building models, simulating them and analysing and visualising the results. Snoopy’s scriptable component, `Spike`[4], was used for the simulation phase. Python’s `NetworkX` package³ was used for graph generation and representation [9].

3.4 Experiments

The utility of this method was explored by encoding models exhibiting information diffusion and what we call multiple phase diffusion, and then consensus formation, before then encoding a mechanism of echo chamber formation. We then demonstrated the conflation of individuals’ position on otherwise separate issues through emotion. The time step of the simulation run with `Spike` was 0.02 units, all single neighbour transitions had a firing rate λ of 1 firings per time unit, and all two neighbour transitions (introduced below) a firing rate of 2 firings per time unit. We present these models as illustrative of the potential for these methods, and as possible primitives, rather than as interesting results in their own sense.

The models capture the process nature of Whitehead’s model in that occasions depend always on upstream occasions for their cause, and once experienced then go on to form the cause of subsequent downstream occasions. None however draw out the nature of occasions to truly form part of a process—in the way that we earlier described those formed while watching a leaf drop might.

3.4.1 Diffusion. We simulated the model of Figure 5a on a graph created by joining two communities of 7 nodes with small world structure. A marking was chosen such that all nodes started with a token in place x_0 except for one which started in node x_a . We expected all individuals to follow this single individuals’ experience of a resulting in a diffusion process. The occasion a may represent an idea or emotion.

3.4.2 Multiphase diffusion. We simulated the model in Figure 5b on the same graph medium as above. This is an extension to the previous model in that the occasion b can only be experienced after the occasion a . A marking was chosen such that all individuals started with a token in place x_0 except for one which started in node x_b and three of its neighbours which started in x_a . A single 14 node small-world network was used. In this model we expect to see the experience of occasion a spreading and then the experience of occasions b spreading on top of this.




Figure 5: Diffusion models




Figure 6: Diffusion with echo model

3.4.3 Diffusion with echo. We built models demonstrating the multi-path effect that Jasny et al. identified as a mechanism occurring in social networks that contributes to echo chamber formation [13]. The authors identified a behaviour in which people are more likely to believe something (or in our case, experience it in the same way as their neighbours) if they hear it from two sources or if they hear it multiple times. We simulated the model in Figure 6 which captures both of these mechanisms. This was the same diffusion model used above but with the addition of a second transition in parallel with the first. The second transition had a guard function requiring two neighbours to have a token in x_a to fire. We set a λ of two firings per time unit for this, rather than the one firing per unit used for the single neighbour transitions elsewhere in this article.

A marking was chosen such that all nodes started with a token in place x_0 except for two which started in place x_a . We ran this simulation once with the rate of t_a^{1n} set to zero and once with it set to one. The rate of t_a^{2n} was set at two firings per time unit for both cases. In the first experiment an occasion is experienced only if two neighbours have experienced it, and in the second experiment an individual may follow either one or two neighbours, but will follow two neighbours at twice the rate.

3.4.4 Single issue consensus. We then ran simulations demonstrating consensus formation of a single issue with two possible opinions a person might choose. We simulated just the left hand issue a of Figure 7 (comprised of the competing occasions a^+ and a^-). We placed this on our two community network and started each person with a randomly chosen opinion. In this model the

¹<http://www-dssz.informatik.tu-cottbus.de/DSSZ/Software/Snoopy>

²<https://github.com/robwalton/occmodeler>

³<https://networkx.github.io>




Figure 7: Orthogonal consensus formation model

individual will experience only the occasion of changing ones mind: if they had experience a^- they may subsequently experienced a^+ . This model is the same as that introduced by Morris DeGroot in his work “Reaching a consensus” in 1974 [6].

3.4.5 Issue conflation. Above we simulated just issue a of Figure 7. We then created a model containing both issues a and b . As expected, and not shown here, an individual’s opinion on each of the two issues was uncorrelated.

We then went on to simulate the model of Figure 8 with the goal of demonstrating how our method might be used to model scenarios in which an individual’s opinion on two issues can be conflated via emotion. This model has two issues a and b like that of Figure 7, and as before, an individual may undergo the experience of believing an opinion if a neighbour has experienced it. However in this model the individual must have last experienced a compatible emotion γ . In this model to experience a^\pm or b^\pm requires γ^\pm to have been most recently experienced. To illustrate this, a possible interpretation might be: a^+ is to experience a belief in freedom of gun ownership, b^+ is to experience a belief in closed borders, and γ^+ is to experience the feeling of closed-ness or wanting to secure what one has.

In this model, the emotion γ diffuses between individuals and as just stated it must be experienced before an opinion that depends on it may be. In addition the emotion γ will also tend to follow a currently internally held belief: if an individual experiences a^+ before then experiencing the incompatible γ^- then there is tension from the place x_a^+ to cause γ^+ to be re-experienced.

For this experiment, a marking was chosen such that all individuals on one half of the graph started having most recently experienced the emotion γ^+ and the compatible opinions a^+ and b^+ and on the other half γ^- , a^- and b^- .

3.5 Result interpretation and visualisation

The result of simulating these models with Spike was, for each node on the network, a time series of token count in each place and the number of times a transition fired in the prior time interval. With this data it is easy to create a movie showing how the marking of each individual node changes across the network. However we took the approach of showing this temporal data as a causal graph such as that shown in Figure 9.

Figure 9 shows the causal graph resulting from a simulation run. The social network or medium M is shown at the bottom of the plot with initial markings indicated by colour. Time increases upwards from this plane. Following up from each vertex a new occasion is marked as a disk when the individual at that vertex has experienced it. The dashed edges indicate the local cause of




Figure 8: Issue conflation model. Transitions are not labelled and all had a λ of one firing per unit of time.

the occasion, and the solid edges the cause from a neighbour. The curvature of these edges represents nothing and is present only to make the resulting visualisation easier to interpret: these arcs leave a prehended occasion horizontally (through the ‘space’ of the network) and curve upwards (to become parallel with the time axis) as they enter a new occasion.

4 RESULTS

We present and then discuss the results from the experiments and then from the method as a whole.

4.1 Experiments

The individual experiments described in Section 3.4 are presented and discussed below. Interactive versions of the figures are available online⁴, along with results of simulating the same U models on larger networks.

4.1.1 Diffusion. Figure 9 shows a representative causal graph from the diffusion model of Section 3.4.1. As expected the result is a spreading of the experience across the whole network. We observe that, at least on these small networks, there are a countable number of possible graphs. These however expand to an infinite number of particular timings between occasions.

4.1.2 Multiphase diffusion. Figure 10 shows a representative causal graph from the multiphase diffusion model of Section 3.4.2. Here there is a progression of each individual experiencing first a and then b .

4.1.3 Diffusion with echo. Figure 11 shows a representative causal graph from the first run of the diffusion with echo model of Section 3.4.3. The b message can be seen to spread requiring always the prehension of two neighbours having most recently experienced b . In this run the rate of firing of the transition which follows a single neighbour was set to zero. We see, as we would

⁴<https://robwalton.github.io/posts/2019/websci19/>




Figure 9: Causal graph for a run of the *diffusion* model




Figure 10: Causal graph for a run of the *multiphase diffusion* model

expect, that the causal path cannot hop from the initial community to its neighbour.

Figure 12 shows the same model extended to also include single neighbour sharing. Two instances of single neighbour sharing are indicated with arrows.

4.1.4 Single issue consensus formation. Figure 13 shows a representative causal graph from the single issue consensus model of Section 3.4.4. This model is the same as that introduced by Morris DeGroot in his work “Reaching a consensus” in 1974 [6]. DeGroot showed that such a model will always converge on a connected graph and modelled the dynamics of opinion across the graph. In about half of the runs the two communities had opposing views. The extreme separation of the two communities reduces the chance of convergence to a single view if the initial random marking results in disagreement. However, all a^+ and all a^- are the only completely stable states and given long enough the system will find itself in one of these. It is interesting that there are two long term stable




Figure 11: Causal graph for a run of the *diffusion with echo* model with only two neighbour following




Figure 12: Causal graph for a run of the *diffusion with echo* model with both one and two neighbour following. Two instances of single neighbour sharing are indicated with arrows.

states and two short term stable states. We might liken these to attractors from non-linear dynamics. We present the result from this particular simulation run because the a^- opinion made an exciting but brief foray into a^+ territory.

4.1.5 Issue conflation. Fig 14 shows a representative causal graph from the conflation through emotion model of Section 4.1.5. Figure 15 shows, as a function of time, the number of nodes with positive rather than negative opinions for this run. In this run, where both communities converged to the same opinions, the positive opinions were both chosen along with the positive emotion. In all cases the communities tended to either all positively labelled or all negatively labelled choices.




Figure 13: Causal graph for a run of the *single issue consensus* model showing two communities fighting out differing opinions




Figure 14: Causal graph for the *conflation through emotion* model

4.2 Method

We developed a two step method of modelling behaviours found in some socio-technical system scenarios on the Web. The mapping from a scenario onto Whitehead’s concepts of occasions, with individuals represented as societies of occasions formalised by the possible paths through these occasions was natural. The subsequent mapping from this to Petri nets and from simulated results to causal graphs worked well. The resulting models’ basis in Whitehead’s metaphysics provides room to attach richly nuanced interpretation—although this may not be evident from looking at the shape of the




Figure 15: Place count for graph for the *conflation through emotion* model. Traces are shown only for positive opinions (negative counts and positive counts sum to the network’s size).

resulting Petri nets themselves or of the unfolded causal graphs of experience. We did not validate the results of experiment against other modelling methods, although the models were simple and the results were as expected. These models should be considered primitives which might be combined together to form more complex sequences of experience.

In some cases, when generating a causal graph from the Spike software’s simulation results, multiple firings of the same transition occurred in the same time bin. In these cases there was not enough information to generate a causal graph with all certainty and a heuristic approach was used.

5 DISCUSSION

In this paper we developed and tested a method for the exploratory modelling of emergent socio-technical systems on the Web. We proposed a mapping for describing a fixed social network of interacting people as a partially overlapping collection of Whitehead’s societies of occasions. We created a single model for each person which formalised every sequence of occasions they might conceivably experience by contriving a ‘chemistry’ relating their external stimuli, internal experience and reactions to this internal experience. Interaction between people was limited to one person experiencing what a neighbour had experienced, but only if their own history was compatible. We proposed a mapping from this representation onto a single stochastic coloured Petri net. This was done by creating a unit describing a single ‘chemistry’ applicable to all, and then encoding the shape of a fixed network into the model in a way that it can be potentially unfolded into a single non-coloured stochastic Petri net in the shape of that network. When simulated the resulting models unfolded into one of many possible causal graphs of occasions of experience which we showed using an interactive visualisation.

We verified the method by developing open source software built on top of the Petri net simulation software Spike. With this we encoded exploratory models exhibiting information diffusion and what we called multiple phase diffusion, and then consensus formation, before then encoding a mechanism of echo chamber formation. We then demonstrated the conflation of individuals’ positions on otherwise separate issues through emotion. The results were as expected.

A limitation to the work is that the network individuals interact on is fixed. There is a mechanism for an edge between two people to break in that they may reach markings which prevent future interaction. However there is no way for new edges to form. Although the obvious solution of providing a fully connected network provides a good place to reason from it would result in prohibitively large simulations. Also, at first glance the requirement that every individual be identical appears limiting. However the single ‘chemistry’ describing the possible sequences of experience might include places with tokens placed during the initial marking which determine permanent characteristics such as character. Although there is no limit to how big models might be in theory the software in its current form will not scale past running on a single machine.

Andrew Selbst et al. identify a number of traps discovered while developing approaches to designing socio-technical systems that behave fairly—for example when performing an automated resume screen or making legal judgements [20]. Among more detailed recommendations they note that system designers must draw abstraction boundaries around social actors rather than just technical ones. The method presented here offers one way of doing this.

6 CONCLUSION

In this paper we developed and tested a method for the exploratory modelling of emergent large-scale socio-technical systems on the Web. Based around modelling possible sequences of human experience with Petri nets the method provides a formalism not seen in other work. It was shown that models exhibiting simple diffusion or consensus formation were simple to build. Highlighting the method’s intended use for exploratory modelling we showed that an accepted mechanism for echo chamber formation resulted in relatively stable polarised views on a network with pre-existing community structure. We were also able to demonstrate a plausible model showing how opinion on otherwise unrelated issues may be coupled through an emotion common to both.

This method may be useful to those interested in understanding the mechanisms behind the emergence of complex ecosystems of social machines, echo chambers and other structures on the web. The use of Petri nets at the base of the method blurs the lines between ad-hoc agent based modelling and more analytical modelling. Properties of small or suitably symmetric models might be derived analytically through a mapping onto sets of ordinary differential equations or Markov models. Some properties might help for example build multi-scale models comprised of a perspective formalising operation at the scale of coherent interacting communities and other perspectives formalising interactions of individuals within or across these community boundaries. More generally, this method and the open source occmodeller software which implements it provides a good place from which to explore process based approaches to modelling social systems in general.

REFERENCES

- [1] Wil Van Der Aalst, Ton Weijters, and Laura Maruster. 2004. Workflow mining: Discovering process models from event logs. *IEEE Trans. on Knowl. and Data Eng.* (2004).
- [2] Tilak Agerwala. 1979. Putting Petri nets to work. *Computer* 12, 12 (1979), 85–94.
- [3] Eduardo G. Altmann, Janet B. Pierrehumbert, and Adilson E. Motter. 2009. Beyond Word Frequency: Bursts, Lulls, and Scaling in the Temporal Distributions of Words. *PLOS ONE* 4, 11 (Nov. 2009), e7678. <https://doi.org/10.1371/journal.pone.0007678>
- [4] J Chodak and M Heiner. 2018. Spike - a command line tool for continuous, stochastic & hybrid simulation of (coloured) Petri nets. In *Proc. 21th German Workshop on Algorithms and Tools for Petri Nets (AWPN 2018)*. University of Augsburg, 1–6. <https://opus.bibliothek.uni-augsburg.de/opus4/frontdoor/deliver/index/docId/41861/file/awpn18-lorenz-metzger-OPUS.pdf#page=9>
- [5] David De Roure, Clare Hooper, Megan Meredith-Lobay, Kevin Page, Ségolène Tarte, Don Cruickshank, and Catherine De Roure. 2013. Observing Social Machines Part 1: What to Observe?. In *Proceedings of the 22Nd International Conference on World Wide Web (WWW '13 Companion)*. ACM, New York, NY, USA, 901–904. <https://doi.org/10.1145/2487788.2488077>
- [6] Morris H. DeGroot. 1974. Reaching a Consensus. *J. Amer. Statist. Assoc.* 69, 345 (March 1974), 118. <https://doi.org/10.2307/2285509>
- [7] Qian Gao, Fei Liu, David Gilbert, Monika Heiner, and David Tree. 2011. A multiscale approach to modelling planar cell polarity in Drosophila wing using hierarchically coloured Petri nets. In *Proceedings of the 9th International Conference on Computational Methods in Systems Biology - CMSB '11*. ACM Press, Paris, France, 209. <https://doi.org/10.1145/2037509.2037538>
- [8] M. Girvan and M. E. J. Newman. 2002. Community structure in social and biological networks. *Proceedings of the National Academy of Sciences* 99, 12 (June 2002), 7821–7826. <https://doi.org/10.1073/pnas.122653799>
- [9] Aric A. Hagberg, Daniel A. Schult, and Pieter J. Swart. 2008. Exploring Network Structure, Dynamics, and Function using NetworkX. In *Proceedings of the 7th Python in Science Conference*, Gaël Varoquaux, Travis Vaught, and Jarrod Millman (Eds.), 11–15. event-place: Pasadena, CA USA.
- [10] Mark Hartswood and Luc Moreau. 2016. A Social Charter for Smart Platforms. (2016). https://eprints.soton.ac.uk/410307/1/SmartSocietySocialCharterforSmartPlatforms_final.pdf
- [11] Monika Heiner, Mostafa Herajy, Fei Liu, Christian Rohr, and Martin Schwarick. 2012. Snoopy – A Unifying Petri Net Tool. In *Application and Theory of Petri Nets (Lecture Notes in Computer Science)*, Serge Haddad and Lucia Pomello (Eds.). Springer Berlin Heidelberg, 398–407.
- [12] John H. Holland. 2012. *Signals and Boundaries : Building Blocks for Complex Adaptive Systems*. MIT Press, Cambridge, Mass., UNITED STATES.
- [13] Lorien Jasny, Joseph Waggle, and Dana R. Fisher. 2015. An empirical examination of echo chambers in US climate policy networks. *Nature Climate Change* 5, 8 (Aug. 2015), 782–786. <https://doi.org/10.1038/nclimate2666>
- [14] Kurt Jensen. 1981. Coloured petri nets and the invariant-method. *Theoretical Computer Science* 14, 3 (Jan. 1981), 317–336. [https://doi.org/10.1016/0304-3975\(81\)90049-9](https://doi.org/10.1016/0304-3975(81)90049-9)
- [15] G. T. Kneebone. 1963. The application of mathematics to the natural world. In *Mathematical logic and the foundations of mathematics : an introductory survey*.
- [16] Michael K. Molloy. 1982. Performance analysis using stochastic Petri nets. *IEEE Transactions on computers* 9 (1982), 913–917.
- [17] Nikolic, I. 2009. *Co-Evolutionary Method For Modelling Large Scale Socio-Technical Systems Evolution*. Ph.D. Dissertation. NGInfra / TU Delft. <http://resolver.tudelft.nl/uuid:b6855afa-e8ab-442d-ac5a-f645f7639c73>
- [18] C. Adam Petri and W. Reisig. 2008. Petri net. *Scholarpedia* 3, 4 (2008), 6477. <https://doi.org/10.4249/scholarpedia.6477> revision #91647.
- [19] Tucmas C. Schelling. 1969. Models of Segregation. *American Economic Review* 59, 2 (May 1969), 488.
- [20] Andrew D. Selbst, Danah Boyd, Sorelle Friedler, Suresh Venkatasubramanian, and Janet Vertesi. 2018. *Fairness and Abstraction in Sociotechnical Systems*. SSRN Scholarly Paper ID 3265913. Social Science Research Network, Rochester, NY. <https://papers.ssrn.com/abstract=3265913>
- [21] Nigel Shadbolt, Kieron O’Hara, David De Roure, and Dame Wendy Hall. 2019. *The Theory and Practice of Social Machines*. Springer International Publishing. <https://www.springer.com/us/book/9783030108885>
- [22] C. E. Shannon. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27, 3 (July 1948), 379–423. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- [23] John Kenly Smith. 1988. The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology. Edited by Wiebe E. Bijker, Thomas P. Hughes, and Trevor Pinch. Cambridge, Mass.: MIT Press, 1987. 405 pp. Illustrations, charts, notes, bibliography, and index. \$35.00. *Business History Review* 62, 2 (1988), 341–342. <https://doi.org/10.2307/3116018>
- [24] Larissa Spinelli and Mark Crovella. 2017. Closed-Loop Opinion Formation. In *the 2017 ACM*. ACM Press, New York, New York, USA, 73–82.
- [25] Michele Starnini, Mattia Frasca, and Andrea Baronchelli. 2016. Emergence of metapopulations and echo chambers in mobile agents. *Scientific Reports* 6, 1 (Oct. 2016). <https://doi.org/10.1038/srep31834> arXiv: 1603.04789.
- [26] Duncan J Watts and Steven H Strogatz. 1998. Collective dynamics of ‘small-world’ networks. *nature* 393, 6684 (1998), 440. bibtext*[publisher=Nature Publishing Group].
- [27] Alfred North Whitehead and Donald W Sherburne. 1957. *Process and reality*. Macmillan New York, NY.