Distinct Agent Activation Schemes and Their Quantitative Effect on the Distributed Averaging Model

Kenneth Comer[1], Andrew Loerch[1], and Robert Axtell[2]

[1]George Mason University, Fairfax, VA, USA
[2]George Mason University, Fairfax, VA USA; External Faculty, Santa Fe Institute, Santa Fe, NM, USA; and Oxford University, Oxford, England

Abstract. An agent-based model was constructed to extend the analytic scope of the theoretical solutions to the "Leveler" or "averaging" model, a simple finite Markov information exchange (FMIE) model. It was demonstrated that varying activation patterns leads to significant differences in the convergence rate. Only random asynchronous activation delivered convergence rates predicted by theory. Other activation patterns show different behavior, and activation based on the internal state of the agent or "endogenous" activation led to convergence rates furthest from the theoretical values.

Keywords: simulation design, software design, activation, updating schemes, replication, continuous Markov chain model, finite Markov information exchange

1 Background

Agent-based models have found wide use in the social sciences, in ecological and biological sciences where they are often called 'individual-based models' and even in the physical sciences. Constructing such models requires a number of fundamental design decisions, such as:

- How many agents?
- How will agents move, if at all?
- How the agents perceive their environment?
- Will agents be connected, e.g. on a network?

An additional key element of agent definition that must be specified is the activation scheme, sometimes referred to as 'scheduling' or 'updating'. This is simply the ordering of execution of the agents in code. For two decades, agent-based model builders have known that different activation schema produce different results [1–4] in some models, and certainly changing the activation scheme will modify the states of individual agents, but it is an open question as to whether this is a general phenomenon at the aggregate level, or one that appears rarely and is of little consequence.

adfa, p. 1, 2011.

Subsequent research has shown that activation is important to model execution and, depending on the activation scheme, to aggregate outcomes. The most consistent difference appear when activation depends on the state of the agents, a process called *endogenous* activation [5]. It is also known that activation can be important in models that seek to support actual decision, *i.e.* applied models and simulations [6, 7].

This earlier research treated the activation question either as a nuisance, hindering the accurate construction of replications of published models, or as a potential tool to replicate real-world behavior [6, 8]. It was also shown that activation may change the dynamics of a model so much that the design of the experiment needs to be reconsidered [7].

2 Co-evolution of Analytical and Computational Approaches

Varying the agent activation scheme can have another impact on modeling and system analysis: simulations with different activation schemes can allow researchers to examine system behavior that is more complex, variable, or heterogeneous than that postulated in a theoretical construct. In other words, activation can be one of the model specifications that can be varied as we harness the power of simulation and move beyond equation-based descriptions of system behavior.

For agent models of any sophistication it is usually the case that succinct mathematical models cannot be readily written down, or even if they can, they cannot be easily manipulated or solved. Computational agents help us to 'solve' such models and in this they are a substitute for analytical approaches. But once an intuitive understanding of a model is developed by making many realizations, it is often possible to then develop new mathematical abstractions for how the model behaves. It is in this sense that we can speak about the co-evolution of analytical and computational approaches.

3 An Agent Model of Distributed Exchange Processes

We focus on a simple model of interacting agents that has a convenient analog in interacting particle systems. This type of relationship has stirred considerable interest among advanced researchers who seek to combine the mathematics with complexity theory. One team, for example, claims to demonstrate that the field of ‘out-of-equilibrium’ statistical physics is uniquely appropriate for understanding complex system dynamics. Such a fusion of the fields can help to explain the ubiquitous appearance of non-stationary and non-ergodic statistical processes and inverse power-law statistical distributions [9]. Markov processes have proven useful in understanding and modeling the negotiation process [10]. More afield (and, perhaps less directly related), the broad field of Markov Chain Monte Carlo methods, and its combination with evolutionary algorithms, has also been applied to the information exchange process. While the models bear little resemblance to agent-based models, the extensive application of ideas first developed for physics and chemistry to search algorithms (among other problems) [11] show that this is a fecund combination of disciplines.

The fusion of stochastic particle physics and social systems analysis is also taking place in the opposite direction: social scientists are finding new applications that bring the rigor and mature theory to their emergent problems. For example, Cai and Ishii have started with straightforward social science questions – the formation of a consensus and the distribution of wealth – and solved the question of convergence using defined and quantized Markov chains [12]. In their conclusions, however, the authors point to a major issue in applying this extensive mathematical treatment to real world situations. In their final remark (Remark 16), Cai and Ishii note that extending their results becomes difficult if the topologies of agent interactions are less well defined. They don't mention this, but if the agent interaction topologies are inconstant in time, extension of this mathematical approach may be unachievable.

4 Theoretical Baseline: Assumptions, Derivations, Predictions

Interacting Particle Systems have been well-defined mathematically. An elegant theory based on the statistics of continuous time Markov chains provides mathematical solutions (once the system parameters are known) for convergence rates, steady-state distributions, mean arrival times (for a given state), and other outcome behaviors of interest.

Aldous, in exploring this theory, blends interacting particle systems with social systems analysis, and draws an analogy from game theory [13]. Game theory has its origins among physicists, but is now broadly applied to social science issues. Moreover, the field is characterized by a small number of simple games which have an abiding importance across a broad range of domains: Prisoner's Dilemma, Tragedy of the Commons, Battle of the Sexes, etc.

Aldous has noted that the extension of interacting particle systems is also based on the application of a small number of straightforward models. IPS are characterized by a common structure:

- A large population of agents – normally taken to represent individuals.
- A network or graph that defines the connections among these individuals. Commonly, the edges of this graph are weighted.
- A meeting model that interprets the weights of the edges as the frequency of meeting
- A meeting algorithm in which the agents exchange information, possibly changing their state in the process.

The last point has motivated Aldous to coin the name Finite Markov Information Exchanges to describe the specialized application of IPS to social science. He notes that the lack of a common name has probably limited the impact of such research, and hindered the formation of a community of practice similar to those who work in the area of game theory.

Aldous defines interaction rates as the symmetric matrix $\mathcal{N}$, which has zeroes in the diagonal (agents don't have meetings with themselves) [14]. The non-diagonal elements are defined in terms of their meeting rates, $\nu ij \geq 0$. In order to use the same techniques used to characterize Markov chains, $\mathcal{N}$ is assumed to be a stochastic matrix, with normalized rates of interaction, so that:

$$\nu_i := \sum_j \nu_{ij} = 1 \text{ for all } i. \quad (1)$$

Note that, while Aldous defines the diagonal elements equal to zero as part of the structure of his problem, this is not part of the definition of a stochastic matrix.

$\mathcal{N}$ also defines a geometric substructure for the interactions. It can take on any form, but Aldous limits his analysis to the most common form. Here we consider only his first topology, which he terms the complete graph or mean field model. In this case, every node or agent has an equal likelihood of interacting with every other. Thus, $\mathcal{N}$ is defined by:

$$\nu_{ij} = 1/(n-1), \ \ j \neq i \quad (2)$$

Aldous also considers other, more complicated topologies including small worlds or random graphs, but our analysis is limited to this straightforward case.

First consider an explicit description of how continuous time Markov chains are expressed mathematically. The objective is to define a method for stating the Markov chain transition probability matrix for a continuous-time Markov chain. Starting with an analogy that the eigenvalues of an invertible square matrix A are those values of λ_i that solve the equation: $\lambda_i A = \lambda_i v_i$ where v_i = the associated eigenvector. Now, consider a Markov transition process (and associated probability matrix) in which the system operates in continuous time (but still with a finite, countable state space). Thus, the transition matrix would not be a matrix of discrete probabilities – the probabilities of moving from one state to another in one time step. It would, rather, be a continuous function of time, $P(t)$ such that the probability the system is in state j after time t, given that it is in state i at time 0, is $p_{ij}(t)$.

In order to analyze $P(t)$, a matrix Q is defined such that $P(t) = e^{tQ}$. This notation, treating a matrix as an exponent, is a shorthand for an infinite series on the exponential of Q that is analogous to Euler's formula:

$$e^Q = \lim_{x \to \infty} \left(\sum_{k=0}^{x} \frac{Q^k}{k!} \right) \quad (3)$$

We also know the following:

$$e^{nQ} = (e^Q)^n = P^n \quad (4)$$

and

$$\frac{d}{dt} P(t) = P(t)Q \quad (5)$$

So, the Q-matrix for a complete graph pattern, in which an agent has equal probability of interacting with each of its partners, is given by:

$$\begin{bmatrix} -1 & 1/(n-1) & \cdots & 1/(n-1) \\ 1/(n-1) & -1 & 1/(n-1) & \cdots \\ \cdots & 1/(n-1) & \cdots & \cdots \\ 1/(n-1) & \cdots & 1/(n-1) & -1 \end{bmatrix}$$

The eigenvalues of this matrix are 0 and $-\left(n+1/n\right)$. It is important to note that the non-zero eigenvalues approach unity as n becomes large.

This exposition is important for follow-on analysis. The Markov chain generating matrix, Q, can be interpreted as a rate-flow matrix in a continuous time Markov chain. It is also the matrix that generates the Markov chain transition matrix, P. Note that, while P is a stochastic matrix, Q is not. [15]

5 The "Leveler Model" – Theoretical Development

Aldous used the structure of continuous-time Markov chains to complete his understanding of convergence in a social problem he deemed the "averaging process" [13] or the "Leveller" problem [14].

In Leveler, each member the population is endowed with an account of 'wealth', which normally begins as differentiated. At each meeting, the two interacting agents reset their individual wealth to the average of their two accounts. Clearly, over time, the population will converge to the point where every agent has the same wealth, especially if all $v_{ij} > 0$ if $i \neq j$. In fact, the population will converge to the average wealth in any case where all states communicate in the Markov Chain transition matrix. Additionally, with this rule all wealth in the system will remain constant, as will the mean wealth.

This leads to a theoretical result in which, given an unchanging meeting matrix, $\mathcal{N}$, the population's convergence – measured as the decay of the standard deviation of wealth to zero – is defined by the Markov chain processes. To begin, Aldous rewrites the definition of $\mathcal{N}$ such that it is a matrix of transformation rates in which the rows sum to zero. Thus, he revises the definition, defining the transition rate from i to j as $v(i,j)$. From this, he establishes the matrix as:

$$v_{ij} = v_{(i,j)}, i \neq j;\ v_{ii} = -\sum_{j\neq i} v_{ij} \tag{6}$$

This is, of course, no longer a stochastic matrix. In fact, from the theoretical development of Markov chain analysis, this is equivalent to the generating matrix, Q [15]. Aldous goes on to develop a theory of convergence rates that depend upon this new $\mathcal{N}$, which will be here denoted as Q. Aldous shows that the convergence rate (under all the previously stated conditions of stationary transition probabilities and finite, countable states), that, if the convergence is measured in terms of the standard deviation of wealth, it is bounded in its convergence to zero.

The notation used in Aldous is a bit different from that used in normal statistical treatments. In his initial conditions, Aldous assumes that the average wealth is zero. This will mean, of course, that the average wealth at all times is zero as the Leveler process does not change the mean wealth. This simplifies Aldous's mathematical notation to the more familiar statistical notation. Given a vector in which the mean value is 0, that is:

$$\frac{1}{n}\sum_{i=1}^{n} x_i = 0 \tag{7}$$

Aldous defines the "norm", which is equivalent to the standard deviation.

$$\|\mathbf{x}\|_2 := \sqrt{\frac{1}{n}\sum_i x_i^{\,2}} = \sigma_x \tag{8}$$

Thus, where σ_w is the standard deviation of the wealth at time t and σ_0 is the standard deviation of the wealth at time $t = 0$, the convergence is determined by:

$$E[\sigma_w(t)] \le \sigma_0 e^{-\lambda t/4} \tag{9}$$

where λ = the spectral gap of Q. The spectral gap is the distance between the zero eigenvalue and the next largest eigenvalue, or simply the value of the smallest non-zero eigenvalue. In this model, as noted above, as n becomes large, the non-zero eigenvalues of Q approach unity. Thus, the exponential decay rate in Equation 9 will be approximately -1/4.

It is important to consider whether this convergence rate, dependent on λ is a function of the number of agents. This depends on how Q is defined, and, thus depends on the definition of the meeting rates. Moreover, the meeting rates are determined by the definition of the unit of time. If the rates are set as above, a unit of time is defined as that amount of time such that, on average, one interaction will take place among the all the agents in a single unit time. If time were defined in such a way that each agent would initiate a meeting once per unit time, Q would be a matrix with all non-diagonal elements equal to one, and the diagonal elements equal to n – 1. The non-zero eigenvalues of such a matrix would equal –n, and the spectral gap and the convergence rate would certainly vary with the scale of the system. Thus, the definition of time units becomes a key constituent in moving from the mathematical definition of the system to its simulation.

6 Extending Analytical Results Through Computational Modeling

A commonly-used technique in operations research and systems engineering is to start with a well-developed mathematically-defined system and build a simulation. The simulation will allow the researcher to relax the assumptions of the model, through the design of the code, and examine system behavior. In the general case this allows the operations research analyst to leverage mathematical prediction and extend the range of quantitative analysis. (Simulation is also used to extend the insights gained from physical experimentation, further adding to the utility to decision-makers and the broad confidence non-academic professionals place on simulation.)

Agent-based models also have been used extensively to evaluate the diffusion of information in a population. Herrmann, et. al., have recently modeled the diffusion of urgent information (weather warnings or high-profile news events) on a network using an agent-based model [16]. Hui, et. al. simulated the diffusion of evacuation warnings to a population of agents connected via a network. A simulation approach was necessary because, as agents evacuated, the network topology would change [17]. Rosval and Sneppen explored the exchange of information in an agent-based model of a dynamic network [18]. And, Cui and Potok, using an agent-based swarm-type model of insurgency showed that information exchange among disparate, self-organized groups can be just as efficient as in a hierarchical insurgency with unified leadership and strategic planning [19].

7 Varying Activation: Computational Results

The Leveler theoretical model assumes that the meeting matrix or transition matrix $\mathcal{N}$ (or its generating matrix Q) remains unchanged during the course of the model. It has no concept of 'turn' in which a full population of agents are activated. The model evolves in 'secular' time, and all n agents activate in accordance with their own Poisson process. Most of these assumptions are made in order to make this elegant derivation of the convergence rate as a closed-form inequality possible.

Do these conditions exist in the real world? Aldous cautions researchers who extrapolate these abstract models to real-world movement of knowledge in a population. Information does not take on well-defined scalar values (such as wealth in the Leveler model), and individuals find many ways to move information beyond a simple meeting protocol [13].

To capture some of the non-abstract real-world behavior, a Leveler model was created in Python. The convergence of wealth – as measured by σ_w -- was examined using different activation schemes:

- Uniform activation creates a sequence of pairs from the population through sampling without replacement. The pairs leveled their wealth when they were activated. One turn is defined as activating the entire population (in pairs) exactly once.

(Odd-numbered populations will have one inactive agent in each turn, randomly assigned.)

- Random activation involves selecting pairs of agents from the population *with* replacement. A turn is defined as complete when a full population has been activated, or after $n/2$ pairs have been selected.
- Poisson required the determination of the activation rate, λ_i, for each individual agent. These rates were normalized at the beginning of a turn so that, on average, one population's worth of agents would be activated on each turn. Thus, the mean λ would be $1/n$. Also at the beginning of a turn, a Poisson process was populated for each agent in accordance with the individual arrival rate, λ_i. These arrival times were placed in sequence on an 'activation table'. By design, the average number of agents on each turn's activation table was one population's worth of agents. The leveling process took place by selecting the agents from the table two at a time. At the beginning of the next turn, agents' values of λ_i were recomputed. Several rules are possible to determine this λ_i. I chose to make λ_i proportional to the distance between the individual agent's wealth and the mean wealth. Those furthest from the mean wealth would activate more frequently, those closest to the mean would activate at a slower value of λ.
- Inverse Poisson activation merely reversed the above rule. Those agents closest to the mean would activate fastest while those furthest from the mean (the richest and the poorest agents) would be the least likely to 'share the wealth'.
- Natural. Poisson activation adjusted such that the rate of activation, λ, depends on the level of poverty, or the inverse of wealth. Thus, the poorest agents would activate the fastest, and the richest would be the least likely to enter the wealth-swapping process.

In order to create a distribution of wealth, each agent was endowed with a wealth 'account' equal to his index value. Thus, the first agent started with a wealth of 1 and the 1000^{th} agent began with a wealth of 1000. Thus, the average wealth was 500.5, and the initial standard deviation was 288.675. Five runs were conducted for each activation scheme.

The figure shows the results for the five activation schemes. From inspection, the different activation schemes resulted in markedly different convergence rates. As the exponent of decay was the most important for these time series, shows the average coefficient of the time variable. Note that, for this simulation, time is defined in turns. In the theoretical construct, turns do not exist and time is defined in terms of the individual agents Poisson process.

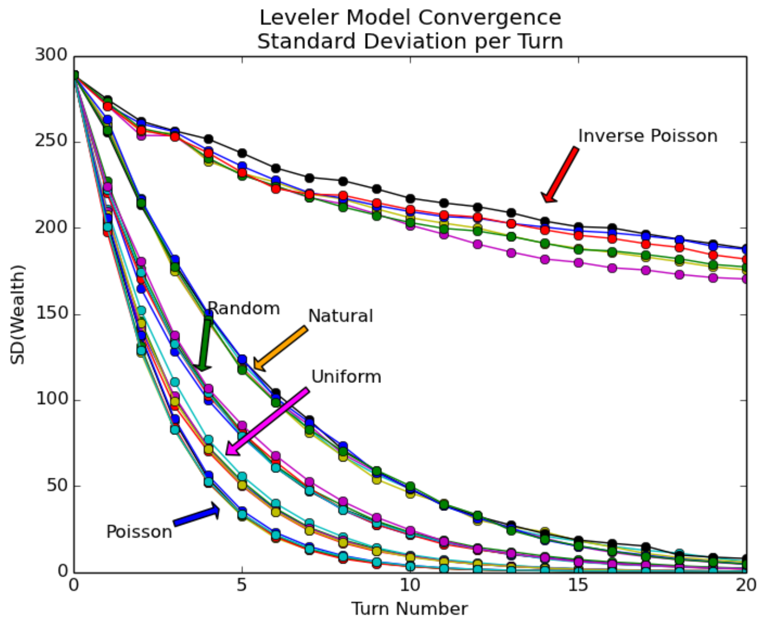

Fig. 1. Decay of standard deviation for various activation schemes in the Leveler Model; N = 1000; five runs per activation scheme

Table 1. Decay rate exponent estimate based on linear regression of log(SD)

Decay Rate			*Experimental Run*		
Activation	**1**	**2**	**3**	**4**	**5**
Random	-0.251	-0.255	-0.251	-0.254	-0.253
Uniform	-0.346	-0.352	-0.340	-0.347	-0.347
Poisson	-0.457	-0.460	-0.460	-0.454	-0.457
Inverse Poisson	-0.019	-0.020	-0.017	-0.020	-0.020
Natural	-0.090	-0.091	-0.094	-0.088	-0.089

Thus, it is clear that these average decay rates differ consistently, and that they are quite stable once activation has been set. Further, the decay rate from the random activation process closely tracks the theoretical rate of -1/4. The obvious conclusion that the runs are different can be confirmed by a Fisher Exact test of any of the five runs compared with any other of the five runs would give a *p*-value of 0.004. (This could have easily been driven smaller with more runs, but the outcome is rather obvious from **Fig. 1** and **Table 1**.)

8 Conclusions

Agent activation is a neglected aspect of agent model building and evaluation. It is important to consider multiple activation schemes because, as we have shown here, the exact method employed has quantitative impact on the model output. Without an understanding of the role of activation it will be hard to assign causes of variability in a model, as some of this will undoubtedly be due to how agents are turned on. This conclusion is valid for single-threaded models, and becomes more problematical as we move to parallel execution.

The most important result was the replication – through simulation – of the theoretical pattern of decay. The alignment of the random activation scheme decay rate with the theoretically-predicted value of -1/4 implies that random activation most closely represents the natural process described in the theoretical model. Agents interact in accordance with their own, internal "clocks", unaware of the actions of other agents. It also validates the conventional definition of a 'turn' as a population's worth of agent activations. While that definition might have seemed contrived, it does appear to conform with the system behavior predicted by the MC model.

This result is important because it allows the use of the theory-simulation analytic paradigm. The theoretical development led to the conclusion that, given the appropriate definitions of time and standard deviation, the mean convergence rate for a wealth-averaging system should be 'no greater than' -1/4. Thus, as the assumptions about homogeneous, constant activation are relaxed, the impact on convergence can be observed through simulation. Simulation, therefore, can be used to extend the analytic reach of theory in such models. And, it has been shown, changing activation does impact the outcome patterns of this simple model. It is reasonable to assume that more complex models might see similar differences and experiments should be conducted to investigate such differences.

While it is important merely to show that there are differences, it is also interesting to note that the differences are not of uniform magnitude. Clearly the inverse Poisson convergence rate is very much less (in absolute value) than the other convergence rates. Inverse Poisson activation was based on the assumption that the agents with the most extreme wealth would enter the wealth-swapping process the slowest.

These results suggest that the choice of activation pattern can become an important tool for researchers attempting to simulate real-world self-organizing systems. That is, rather than treating activation as an arbitrary and confounding choice, it can become a treatment parameter for exploring various emergence phenomena. Often agent-based models are built in an attempt to mimic real-world behavior. It would not be unusual for the model-builder to grow acquire data or insights into the real-world activation patterns of individuals. This might come from theory or there may actually be empirical data. If, in the real world this data is not stationary, then the researcher would have a tool to adjust the model structure to match behavior. This is especially true in the case of endogenous or state-based activation shown here. In fact, most intuitive expectations for real-world systems would assume that activation would be based on state: diseased individuals will interact more rarely than healthy people; wealthier people normally trade stock with greater frequency, etc. If these differences can be parameterized, the activations schemes denoted here can help create a better model.

A better connection with reality was the motivation for investigating the ‘natural’ activation process. And, its evolution about half-way between the Poisson and inverse Poisson agrees with an intuitive interpretation. The Poisson process activates the most extreme agents (the wealthiest and the poorest) the fastest. The inverse Poisson process activates them the slowest. The natural process activates half the fastest (the poorest) and half the slowest (the richest).

The clustering of the convergence rates within the activation types is somewhat of a surprise, varying much less among runs with the same activation than between runs with different activation. Clearly activation has a dominant effect on this model’s variation. Thus, if one were to choose to create a model based on this construct (Leveler has been proposed as a real-world scenario for the movement of gossip across a population), then various activation schemes should be explored and reported.

9 References

1. Huberman, B., Glance, N.: Evolutionary Games and Computer-Simulations. Proc. Natl. Acad. Sci. U. S. A. 90, 7716–7718 (1993).
2. Nowak, M., May, R.: Evolutionary Games and Spatial Chaos. Nature. 359, 826–829 (1992).
3. Axtell, R.: Effects of interaction topology and activation regime in several multi-agent systems. In: Moss, S. and Davidsson, P. (eds.) Multi-Agent-Based Simulation. pp. 33–48. Springer-Verlag Berlin, Berlin (2001).
4. Lawson, B.G., Park, S.: Asynchronous time evolution in an artificial society mode. Jasss- J. Artif. Soc. Soc. Simul. 3, U35–+ (2000).
5. Fernández-Gracia, J., Eguíluz, V.M., San Miguel, M.: Update rules and interevent time distributions: Slow ordering versus no ordering in the voter model. Phys. Rev. E. 84, 015103 (2011).
6. Comer, K.W., Loerch, A.G.: The Impact of Agent Activation on Population Behavior in an Agent-based Model of Civil Revolt. Procedia Computer Science. , Baltimore, MD (2013).
7. Comer, K.W.: Who Goes First? An Examination of the Imact of Activation on Agent Behavior in Applied Agent-based Models, (2014).
8. Page, S.E.: On Incentives and Updating in Agent Based Models. Comput. Econ. 10, 67–87 (1997).
9. West, B.J., Geneston, E.L., Grigolini, P.: Maximizing information exchange between complex networks. Phys. Rep. 468, 1–99 (2008).
10. Weingart, L.R., Prietula, M.J., Hyder, E.B., Genovese, C.R.: Knowledge and the Sequential Processes of Negotiation: A Markov Chain Analysis of Response-in-Kind. J. Exp. Soc. Psychol. 35, 366–393 (1999).
11. Laskey, K.B., Myers, J.W.: Population Markov Chain Monte Carlo. Mach. Learn. 50, 175–196 (2003).

12. Cai, K., Ishii, H.: Convergence time analysis of quantized gossip consensus on digraphs. Automatica. 48, 2344–2351 (2012).
13. Aldous, D., Lanoue, D.: Interacting Particle Systems as Stochastic Social Dynamics. Probab. Surv. 9, 90–102 (2012).
14. Aldous, D.: IPS as Stochastic Social Dynamics - PowerPoint.pdf. (2013).
15. Norris, J.R.: Markov Chains. Cambridge University Press (1998).
16. Herrmann, J., Rand, W., Schein, B., Vodopivec, N.: An Agent-Based Model of Urgent Diff usion in Social Media. Available SSRN 2297167. (2013).
17. Hui, C., Goldberg, M., Magdon-Ismail, M., Wallace, W.A.: Simulating the Diffusion of Information: An Agent-Based Modeling Approach. Int. J. Agent Technol. Syst. 2, 31–46 (2010).
18. Rosvall, M., Sneppen, K.: Modeling Dynamics of Information Networks. (2003).
19. Cui, X., Potok, T.E.: A Particle Swarm Social Model for Multi-Agent Based Insurgency Warfare Simulation. 5th ACIS International Conference on Software Engineering Research, Management Applications, 2007. SERA 2007. pp. 177–183 (2007).