

An Agent-Based Model for Simultaneous Phone and SMS Traffic over Time

Kenneth Joseph^{1,2}, Wei Wei^{1,2}, and Kathleen M. Carley^{1,2}

¹ iLab, Heinz College, Carnegie Mellon University, Pittsburgh, PA, USA

² Institute for Software Research, Carnegie Mellon University, Pittsburgh, PA, USA
{kjoseph, weiwei, kathleen.carley}@cs.cmu.edu

Abstract. The present work describes a utility-based, multi-agent, dynamic network model of phone call and SMS traffic in a population. The simulation is novel in its ability to generate interactions from both an asymmetric and a symmetric media simultaneously. Within the model, we develop and test a simple extension to the theory of media multiplexity, a well-known theory of how humans use the communication media available to them with different alters (friends). Model output qualitatively matches patterns in real data at the network-level and with respect to how humans use SMS and voice calls with different alters and thus shows general support for our theoretical claim.

Keywords: agent-based modeling, dynamic networks, interpersonal communication theory.

1 Introduction

Effects on how people use the repertoire of communication media (e.g. voice calling, emails and SMS) available to them to interact with others exist within a complex, nonlinear system across the individual, instantaneous, media, dyadic and network levels of analysis (see, e.g., [1]). The dynamic, multi-level nature of these effects has led to a variety of both competing and complementary theoretical and empirical claims of how humans use different communication media.

Precisely because of its ability to control effects at different levels of analysis, agent-based modeling presents a useful tool to develop and test theories of media use. In the present work, we describe an agent and utility-based dynamic network simulation model to analyze effects on media usage at three levels of analysis - the node, the relation, and the affordances of different media¹.

To show the practicality of such a model, we use it to test a new theoretical proposition that extends Haythornthwaite's argument for media multiplexity [2], a theory of human media usage. Haythornthwaite argues that strong social ties will interact on a variety of media, while weak ties will tend to stay on more "established" media within their social context- for example, email is an established

¹ Code for the model is available upon request to kjoseph@cs.cmu.edu

media within many organizations². As a corollary to this argument, Haythornthwaite shows that an increase in tie strength promotes increased interaction on all media.

While an important contribution to the field, Haythornthwaite’s work does not give clear guidance for analysis when no obvious social context or established media exist. This is often the case with datasets of call detail records (CDRs), which are typically anonymized sets of data from a large and diverse population. In the present work, we develop, implement and test a theory to extend media multiplexity to such situations. We hypothesize that where no obvious social context or established media exist, people will increasingly rely on their preferred medium as tie strength decreases- correspondingly, they will be more open to using a variety of media when interacting with stronger ties. We title this theoretical claim the *weak-tie exaggerated choice theory* (WT-ECT).

We implement our theory in a model where agents develop a preference for one of two media, voice calling or SMS, based on what they have gained emotionally from communicating on these media in previous interactions. While each media has specific affordances in the model that present an obviously rational choice, agents are *boundedly rational* [3], meaning they make the most rational decision based on their view of the world, which may or may not be correct. Bounded rationality, combined with our implementation of WT-ECT, suggests a means by which people might prefer interacting on certain media even when it is not entirely rational to do so, and will rely on their (possibly irrational) preferences more heavily with weaker ties.

Our theory thus falls in line with recent thinking regarding the notion of the “domestication” of communication technology [4] and presents a possible means of connecting this theory with the concepts presented by Haythornthwaite. Additionally, it provides another perspective from which to interpret recent literature suggesting a strong effect of tie strength on the usage of SMS, in particular its use as a tool for communicating only with stronger ties [5,6].

Our main contributions in the present work are thus three-fold. First, we extend a previous model to simulate asymmetric interactions within an evolving social network. Second, our extensions allow us to simultaneously model interactions on two different media with distinct properties. Third, we propose, test, and find evidence for an extension of a well-known theory in the interpersonal communication literature, with ties to other relevant research on mobile communication.

2 Simulation Model

The model we present is grounded in the work of Du et al. [7], who develop the Pay-and-Call (PaC) Model, an agent and utility-based model of network creation over time. In the original PaC model, each agent wants to maximize their supply of “emotional capital”, which they can do via interaction with others. Each agent

² When we refer to a social tie, we mean any connection between two people. The strength of a tie can be modeled in various ways - for example, by the number of interactions between them.

is defined by a uniformly distributed *friendliness* value that gives their aptness in social situations and a uniformly distributed *lifetime* value representing the likelihood that they are replaced on each turn of the model, as described below. For each interaction, agents receive a payoff, defined in Equation 1.

$$payoff = \sqrt{Fr_i * Fr_j} * \frac{1 - \alpha^{intLen+1}}{1 - \alpha} - CPM * intLen - initCost \quad (1)$$

In Equation 1, Fr_i stands for the friendliness of agent i and $intLen$ is the length of the interaction (in minutes). The value α is a model parameter that represents the ability of a medium to convey the true benefit of social interaction to the agents. As time progresses, the effect of α increases, causing a limit in the amount of capital gained by agents on an interaction. Our model differs from the one presented in [7] by a factor of α - this is done in order to ensure that interactions lasting only one minute are differentiated by unique α s, as explained below. The benefit of an interaction is also mitigated by a linear cost defined by an intercept $initCost$, the cost of initializing an interaction, and CPM , the cost per minute of the interaction. In addition to the CPM mitigating the benefit from an interaction, it is also subtracted from the sender's capital each minute the interaction continues. An interaction ends when the interaction sender has no more capital to continue, or the payoff decreases from one minute to the next.

In our model, we differentiate the α values of voice calls and SMS- that is, we differentiate between their ability to relay emotional benefits as an interaction progresses. Early research in media choice suggested that the more "socially present" a medium is, the better it is able to convey emotion between sender and receiver [8]. Though these claims have since been questioned, it has also been shown that Americans interact more using voice calls than SMS when talking to core ties [9]. Thus, we would expect that agents in our model should be limited in the amount they can benefit emotionally from an SMS as compared to a voice call. From Equation 1, we see that by lowering the α value for SMS, we can model this theoretical concept- the precise values of α for each media are set after calibration, as discussed in the following section.

On each turn of the original PaC model, each agent, in succession, is tested to see whether or not they should be replaced based on their *lifetime*. If not replaced, an agent interacts with his alters in order of the payoff the agent received from each alter the last time the pair interacted. An agent will continue interacting with alters until his capital is exhausted or until the alter he will next interact with has a remembered payoff lower than the agent's expected payoff from talking to a new tie. This expected payoff is calculated as the average of the first payoff an agent received from his alters. At this point, the agent, i , will then interact with one new agent, j , that had the maximum payoff of all of i 's alter's friends (i.e. friends of friends).

One vital difference between various communication media that cannot be modeled by the generator in [7] is their level of "synchronicity" [1] - phone calls are synchronous in that both agents must be present for a phone call to occur, while other media, such as email, SMS and IM, are asynchronous in that sends and replies need not occur concurrently. As shown in Algorithm 1, we modify

Algorithm 1. Time-scale Based PaC Model

```

1 foreach Minute, m do
2   foreach Agent, ai in A (in random order) do
3     ai.doTurn()
4     if (m modulo MinutesInDay) == 0 then
5       foreach agent, aj, in A do
6         if RandomZeroToOne() < agent.lifetime then
7           replace aj with new agent

```

Algorithm 2. *Agent.doTurn*

```

1 if currInteraction != NULL then
2   Update currInteraction by one minute according to Equation 1 (only if sender)
3   if currInteraction is finished then
4     If asynchronous, only the sender obtains the payoff and then places the interaction
     on the receiver's queue. If synchronous, the recipient and sender both obtain the
     payoff
5     currInteraction = NULL
6   return
7 if capital < needed then
8   capital+ = storedCapital
9   storedCapital = 0
10  return
11 if Interaction Queue is Empty then
12  currInteraction = InteractionFactory.getNewInteraction()
13  Start currInteraction
14  return
15 nextInt = InteractionQueue.pop()
16 storedCapital+ = nextInt.getReceiverPayoff()
17 if nextInt.conversation is not over then
18  currInteraction = Interaction.replyTo(nextInt)
19  Start currInteraction

```

the original model to allow each agent one chance to act in each minute of the simulation. In making this straightforward modification, agents are now able to carry interactions across multiple turns, and may wait an indiscriminate amount of time before replying to an asynchronous interaction.

Algorithm 2 shows the process for each simulation turn for each agent. If the agent is currently on an interaction, the interaction is simply “updated” by one minute (the payoff equation is recomputed at the next minute). When an agent is on an asynchronous interaction, he is said to be constructing the message. This message construction, like in real life, occurs without the message receiver being made aware that it is occurring. In contrast, synchronous messages require two-way interaction - both the sender and the receiver must actively be on the interaction for it to commence. Consequently, when an agent begins a synchronous interaction, it is either accepted by the receiver if he is not currently on another interaction, or ignored otherwise. If the call is accepted, the receiver sets his current interaction to that call, and can have no other interactions until the call is over.

Ignored calls and completed asynchronous interactions are placed on the receiver’s *interaction queue*. This interaction queue, akin to a “social to-do list”, represents all interactions that the agent might respond to. The queue is sorted by a uniform probability associated with each interaction (as is done in [10]).

If an agent is not currently on an interaction, he will immediately attempt to begin a new one. In order to do so, however, he must have enough capital- if that is not the case, the agent will refresh his capital with the supply he has obtained from previous interactions. This delay between obtaining capital and being able to use it is done in order to ensure consistency with the original PaC model. Thus, agents act in “cycles”, where a cycle is defined as the period between when an agent refills his capital from payoffs earned from previous interactions and when he does so again. In order to keep the distributional properties of the original PaC model with respect to the network, an agent can only interact with each alter once per cycle.

If the agent is not on a current interaction and has enough capital to begin a new one, he can do so in one of two ways. If the agent’s interaction queue is not empty, he will obtain the first interaction off the queue, collect the payoff from the message, and determine whether or not to reply. The decision of whether or not to reply is based on the work of Wu et al. [11], who find that SMS conversation durations can be approximated by a power-law with respect to their temporal distributions. We therefore model the likelihood of agents responding to an interaction off of their queue as a power-law (with exponent 1) based on the number of times the interaction has been “bounced” back and forth between agents, using this as a proxy for conversation length. Thus, the likelihood of Agent b replying to an initial interaction from Agent a is 1. When a sends a reply message to b , b responds with probability .5, and so forth. If an agent obtains an interaction from his queue and chooses to reply, he will begin a reply using the same media on which the interaction was initially sent. Note that two agents may (and often do) end one conversation and begin a new one later in the simulation.

If the agent’s interaction queue is empty, he will begin a new interaction. The agent will first select a new alter, using the mechanism described in the model from [7] (as described above). Once a partner has been selected, the agent will then determine whether or not to begin a phone call or text message with that alter. This decision is based on Equations 2 and 3, which model this decision as a function of both node (agent) and edge-level preferences. The agent, i , first calculates his average payoff from phone calls and SMS and uses these values to determine $p_i(Call)$, his base probability of making a phone call, as shown in Equation 2.

$$p_i(Call) = \frac{avgPayoffCall_i}{avgPayoffCall_i + avgPayoffSMS_i} \quad (2)$$

According to the WT-ECT theory we propose, agent media preferences are exacerbated as the tie strength between two interacting agents decreases. We model tie strength here as the inverse of the “rank” of an alter in the agent’s remembered payoffs. That is, the higher an alter’s *tieRank*, the lower the strength of the tie. Each alter, j , of agent i has a dynamic *tieRank* $_{ij}$ which indicates j ’s place in i ’s list of remembered payoffs. For example, if i had alters j with last payoff .3 and k with last payoff .2, *tieRank* $_{ij}$ would be 1, and *tieRank* $_{ik}$ would be 2. The rank of a new tie is computed based on the agent’s expected payoff from a new tie - thus, for

example, if i 's expected payoff from a new tie was .25, then a new alter, n would have a $tieRank_{in}$ of 2 (and $tieRank_{ik}$ would actually be 3).

$$p_{i,j}(Call) = \begin{cases} p_i(Call) \sqrt{tieRank_{ij}} & p_i(Call) < .5 \\ \frac{1}{\sqrt{tieRank_{ij}}} & \text{otherwise} \end{cases} \quad (3)$$

Given this definition of tie strength, Equation 3 models the effect of tie strength on media preference to obtain a final likelihood of i calling (as opposed to text messaging) j , $p_{i,j}(Call)$. The equation, in accordance with our WT-ECT, increases an agent's likelihood to use their preferred media as tie strength decreases (or correspondingly, as $tieRank$ increases). If i 's preference for (equivalently, likelihood to make a) phone call is less than .5, his preference for SMS moves towards 1 as $tieRank$ increases. Similarly, the likelihood of i making a phone call moves towards 1 if $p_i(Call) \geq .5$. For brevity, edge cases are omitted, however, they behave in the model as one would expect.

3 Results

In order to calibrate and test our model's relevance to real-world data, we utilize a dataset of approximately 110 million phone and SMS interactions from approximately 430,000 people spanning three months in early 2008 in an Asian nation. In the present work, we consider only moderately heavy users, which we define as those users having between 5 and 200 alters and having sent at least 30 text messages and 30 phone calls. Though this means we cannot extrapolate our findings to the entire population of study, we find that it is difficult to understand usage patterns for those people with less than 5 alters and 30 interactions per media, and that those having greater than 200 ties were relatively unlikely to be a single human. After pruning, we are left with approximately 65,000 users, which we split evenly into a training set for calibration and a testing set for evaluation.

Calibration was completed with a chief focus on α_{sms} and α_{phone} , as they are, in the present work, the only theoretically relevant parameters. However, we note that, with one exception, the model was reasonably robust to changes in the other parameters. The exception is sensitivity to moderately large changes in the CPM - though the reasons why are clear from Equation 1, modifications in future work are necessary to lower sensitivity to this parameter. During the calibration process, we also experimented with a variety of functional forms for Equation 3- we found that the square root function, our initial hypothesis, actually gave the best fit to our training data. Calibration resulted in parameter settings of $\alpha_{sms} = .6$ and $\alpha_{phone} = .8$. The CPM and $initCost$ parameters were kept at the values used by Du et al. [7]. We use a model with these settings and simulate interactions between 100,000 agents over 30 simulation "days" when contrasting results with the held-out (test) data.

Comparisons to the held-out data were made to understand the extent to which our simulation could generate a realistic social network and a realistic distribution of media preferences and the model's ability to capture evidence of the WT-ECT existent in the testing data. We here consider only one measure of




Fig. 1. a) The CCDF, in log-log scale, of the degree distributions of the real and simulated data; b) Number of phone calls (y-axis) versus SMS messages (x-axis) per person (agent). Each dot is a single node, and the red is a 2D density estimator; c) A random sample of all edges from 7,000 nodes in each dataset, where the y-axis is proportion of a node’s total text messages sent that alter, and the x-axis is the proportion of a node’s total phone calls made to that alter. In a), b) and c), the gray bar at the top indicates the dataset.

our model’s ability to generate realistic social networks across both phone and SMS, and do so by plotting the degree distribution of the aggregate networks from the real and simulated data. As we see in Figure 1a, degree distributions are qualitatively quite similar. However, the graph belies one important difference that may exist between model output and the real data- due to our parsing mechanisms, it is unclear whether or not the simulated data correctly captures the head of the distribution or whether it generates too many low-degree nodes.

Regardless of the head of the distribution, our model captures the heavy-tailed nature of degree distributions well-known to exist in aggregate social networks from CDRs (e.g. [7]). We now consider aspects of our model (and the real-world data) that have received less attention in the literature. First, we consider how well the model represents media preferences at the node level. In order to do so, we consider the distribution of the number of phone calls agents (in the model) and people (in the real data) make in comparison to the number of text messages they send. As we see in Figure 1b, which plots a point for each node (agent or person) as well as a density estimator for the entire dataset, the vast majority of nodes preferred voice calling to SMS. As we know, a range of effects at various levels of analysis may have caused this in the real data, and though the point spread is quite similar, these two plots clearly show that our model does not capture significant effects. However, it is interesting to note that even when modeling node-level media choice as the resultant of only the technological affordances of the media and bounded rationality, the model suggests a qualitatively similar conclusion.

Our final comparison between the simulated and real data considers how well the two datasets conform to Haythornthwaite’s theory of media multiplexity, and to the WT-ECT extension. The first claim of media multiplexity is that strong ties will tend to interact more on all media than weak ties - to this end, Figure 1c shows




Fig. 2. Seven sample users from the real (top row) and simulated (bottom row) data, and their interactions via phone and SMS with alters of different tie ranks

that in both the real data and the simulated data, there exists an approximately linear correlation between the (log) proportion of a node’s total SMS messages and the (log) proportion of their phone calls that go to a specific alter. Though the plot only show a small sub-sample of the data to avoid over-plotting, we find that the log proportion of a node’s total calls to an alter explain 34.6% and 57.4% of the variance in the log proportion of text messages sent to that same alter in a simple linear model in the real and simulated data, respectively.

The fact that this simple linear model has such a clear predictive power lends significant support to Haythornthwaite’s theory, and as is clear, our model definitively captures this aspect of interaction across multiple media. It is important to note, however, that the linear model only predicts 57.4% of the variance in our simulated data. Thus, our model contains a reasonable and desirable amount of stochasticity in agent interactions with different partners on different media. This variability can be seen in Figure 2, which shows sample users from the real (top) and simulated (bottom) datasets and the number of interactions they had via phone and SMS with alters of different “tie ranks” (discussed below).

The second claim of media multiplexity is that weaker ties will tend to stick to more established media. WT-ECT, however, suggests that with weaker ties, people will more heavily rely on the media they have a stronger preference for. If this were to be the case, then we would expect that with weaker ties, preferences would be more obvious- that is, we would see people use one media more heavily than they would under normal circumstances.

We use a three step process to obtain a set of statistics that can be used to understand the extent to which WT-ECT is supported. We first compute, for each edge $\langle i, j \rangle$ in the network, the log odds of i calling (as opposed to texting) j as $\log\left(\frac{NumCalls_{i,j}}{NumSMS_{i,j}}\right)$ and subtract i ’s log odds of making a phone call to any of his alters. Thus, a positive value for an edge means that i used voice calling more heavily with j than he would on average, and a negative value the opposite. We then bin edges by their strength. To differentiate between ties of different strengths, we bin using the notion of tie “ranks”, as above. Because there is no




Fig. 3. The DLO statistic (y-axis) as tie rank increases (tie strength decreases) for the real and simulated data for agents (people) preferring both SMS and phone calling. All 95% confidence intervals shown are determined using the bootstrap method with 1,000 iterations.

notion of payoff in the real data, we compute rank here based on the number of interactions sent to each alter. Thus, the alter an agent communicated with most (on both SMS and phone combined) would have a tie rank of one, and so on (edges with the same strength are placed into the same bin).

Finally, we compute the mean of this value (for convenience, the DLO, or difference of log odds, statistic) for all edges in each tie bin. As is clear, if the WT-ECT theory were to be supported, we would expect that as tie rank increases, the log-odds of an agent using their “preferred” media increases as well. Figure 3 plots our DLO statistic for the first 35 tie bins for the real and simulated data for agents preferring SMS (the triangles) and those preferring voice calling (the circles). We consider only the first 35 tie bins both because model estimates were highly varied outside this range and because we feel it is difficult to estimate differences in preference for the real data outside of this general range (the number of communications was, on average, around two).

The figure shows that our simulated data obtains pattern (and at some tie ranks, statistical) validity- this lends credence to our model of human communication, regardless of the underlying phenomenon. More importantly, though, the plot largely shows qualitative support for our WT-ECT- as tie rank increases (tie strength decreases), the likelihood of an agent using their preferred media increases to some extent before finally appearing to level out, as our square root function in Equation 3 predicts. In further analysis, we find that this “leveling out” extends to larger tie ranks, though as mentioned, such estimates are somewhat unreliable.

The one exception to support for the WT-ECT appears to be with agents who prefer phone calls and their ties to alters ranking from 10 to 20- in this range, the DLO statistic actually regresses back towards the node-level preference. Interestingly, there is also evidence of this regression in the simulated data, in which the WT-ECT is an explicit functional piece of the model. This observation can be accounted for by the fact that SMS is asymmetric, and thus will, by its nature, tend to have more interactions per conversation than phone calls. However, it also goes lengths to indicate the complexities and intricacies

of human communication across media and, to this end, suggests the power of agent-based modeling in theory development.

4 Conclusion

The current work has several limitations- we provide mostly qualitative evidence of fits with real data and only a cursory model calibration. In addition, we model tie strength as a dynamic element- though it is true that the strength of human social ties with those around us are constantly fluctuating, our model may allow for too much fluctuation of tie strength in the described time span. Limitations aside, however, the model we develop shows promise, as it is the first we are aware of to attempt to model communication on different media simultaneously over time in an evolving network. Our results show that the WT-ETC deserves further attention, particularly because a derivative of the theory suggests that depending on the perceptions (preferences) of a person, (s)he may irrationally choose an affectively weak communication media (like SMS) when interacting with newer and/or weaker ties, leading to difficulties in forming strong social ties for emotional support.

References

1. Baym, N.K.: Personal Connections in the Digital Age. In: Polity (July 2010)
2. Haythornthwaite, C.: Strong, weak, and latent ties and the impact of new media. *The Information Society* 18(5), 385–401 (2002)
3. Simon, H.A.: Bounded rationality. *The New Palgrave: Utility and Probability*, 15–18 (1987)
4. Silverstone, R., Haddon, L.: Design and the domestication of ICTs: technical change and everyday life. In: *Communicating by Design: The Politics of Information and Communication Technologies.*, pp. 44–74. Oxford University Press, Oxford (1996)
5. Ling, R., Bertel, T.F., Sundsy, P.R.: The socio-demographics of texting: An analysis of traffic data. *New Media & Society* 14(2), 281–298 (2012)
6. Xu, X.K., Wang, J.B., Wu, Y., Small, M.: Pairwise interaction pattern in the weighted communication network (March 2012), arXiv:1203.1105
7. Du, N., Faloutsos, C., Wang, B., Akoglu, L.: Large human communication networks: Patterns and a utility-driven generator. In: *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2009*, pp. 269–278. ACM, New York (2009)
8. Short, J., Williams, E., Christie, B.: *The Social Psychology of Telecommunications.* John Wiley and Sons Ltd. (September 1976)
9. Hampton, K.N., Sessions, L.F., Her, E.J., Raine, L.: *Social Isolation and New Technology.* Pew Internet and American Life Project (November 2008) pewinternet.org/Reports/2009/18Social-Isolation-and-New-Technology.aspx (accessed January 10, 2011)
10. Barabasi, A.L.: The origin of bursts and heavy tails in human dynamics. *Nature* 435(7039), 207–211 (2005)
11. Wu, Y., Zhou, C., Xiao, J., Kurths, J., Schellnhuber, H.J.: Evidence for a bimodal distribution in human communication. *Proceedings of the National Academy of Sciences* (October 2010)