

An integrated model of traffic, geography and economy in the Internet

Petter Holme
University of New Mexico
Royal Institute of Technology
pholme@kth.se

Josh Karlin
University of New Mexico
karlinjf@cs.unm.edu

Stephanie Forrest
University of New Mexico
Santa Fe Institute
forrest@cs.unm.edu

ABSTRACT

Modeling Internet growth is important both for understanding the current network and to predict and improve its future. To date, Internet models have typically attempted to explain a subset of the following characteristics: network structure, traffic flow, geography, and economy. In this paper we present a discrete, agent-based model, that integrates all of them. We show that the model generates networks with topologies, dynamics, and more speculatively spatial distributions that are similar to the Internet.

Keywords

Internet — agent-based model — complex networks — Autonomous System — networks

1. INTRODUCTION

As one of the most complex human artifacts, the Internet is a challenging system to model. Dynamic processes of different time scales operate simultaneously—from slow processes, like the development of new hardware to the transport of data, which often occurs at the speed of light.

These phenomena are to some extent interdependent. Traffic provides income to the service providers, which is then invested in infrastructure, which can lead to changes in traffic patterns. We describe a mechanistic, agent-based model (ABM) [4] to study how these phenomena interact to produce the macroscopic features of the Autonomous System (AS) level Internet. Instead of simply reproducing a macroscopic pattern using statistical fitting or phenomenological models, mechanistic models specify a set of primitive components (known as agents) and interaction rules that mimic the architecture of the real system. The models are judged on their ability to generate realistic macroscopic behaviors from these primitive components. The goal is to provide a parsimonious explanation of how a system works by hypothesizing a small set of simple but relevant mechanisms. In this spirit our model attempts to reproduce large-scale features of the Autonomous System level of the Internet by modeling localized and well-understood network interactions.

The ASes of the Internet lend themselves naturally to ABM modeling. Each AS is an economic agent, comprised of a discrete network that can have spatial extent. Over time, ASes create new links to other ASes, upgrade their carrying capacity, and compete for customer traffic. The agents in the model described here, behave similarly, although we have simplified as much as possible. The model is designed to be general enough to simulate any spatially

extended communication network built by subnetworks of economically driven agents.

In previous work, Chang *et al.* showed that incorporating economics and geography into the Highly-Optimized Tolerance (HOT) [6] model increases the model's accuracy [7]. A related ABM model of the AS graph produces degree distributions similar to empirical observations [8]. Bar *et al.* proposed a similar model [2] that incorporates another aspect of the real Internet—that the agents are spatially extended objects. Our model is similar in scope to this earlier work but differs in the details, most importantly by adding explicit economics in the form of cost. Other differences include accounting for population density, simplifying the treatment of traffic flow, and not assuming a HOT framework. The previous work in this area, like much research on network models, focuses almost exclusively on degree distributions of the graphs. In this paper, we compare our results to Internet data using several topological measures [19], including degree distributions, as well as geography and traffic dynamics.

The remainder of the paper is organized as follows. First, we describe and motivate the model. Then, we characterize the time evolution, network topology, correlation between network structure and traffic flow, packet routing statistics, and geographical aspects of the networks produced by the model. Where possible, we compare the properties of these synthetic networks to observed data from the Internet.

2. AS SIMULATION MODEL (ASIM)

We begin with the fundamental unit responsible for network growth, an agent with economic interests [15]. These agents manage traffic over a geographically extended network (which we refer to as a *sub-network* to distinguish it from the network of ASes) and profit from the traffic that flows through their network.

We compare the agents to the ASes that comprise the Internet. This is not an exact mapping—some of the Internet Service Providers (ISPs) have many AS numbers (e.g., AT&T), while other ASes are shared by several organizations. We make the common simplifying assumption that once an agent is introduced, it does not merge with another agent or go bankrupt [23, 25, 8]. This is partially justified by the fact that the Internet, from its inception, has grown monotonically, and we seek to capture this dynamic in our model. Most models of the AS graph enforce strict growth [23] as well and are, as ours, justified by their *a posteriori* ability to reproduce measured features.

We assume a network user population distributed over




Figure 1: Illustration of the network growth algorithm. (a) shows the locations of four agents on the geographic grid. These are assumed to be connected by a physical network administrated by the agent, but is not explicit in the model. (b) is an example graph resulting from (a). That two agents are present in the same pixel is a necessary, but not sufficient condition for a link to form between the agents. (c) illustrates the area that each hypothetical agent can afford to expand to (the shaded region).

a two-dimensional area. Traffic is simulated by a packet-exchange model, where a packet’s source and destination are generated with a probability that is a function of the population profile. The model is initialized with one agent comprised of a network (a sub-network in our terminology) that spans one grid location (referred to as a *pixel* of the landscape). As time progresses, the agent may extend its subnetwork to other pixels, so that the sub-networks reach a larger fraction of the population. This creates more traffic, which generates profit, which is then reinvested into further network expansion. Through positive feedback, the network grows until it covers the entire population. In this section we describe the assumptions and most of the details of the model; the source code is publicly available from www.csc.kth.se/~pholme/asim/.

An agent i is associated with a set of locations Λ_i (representing sources or end-points of traffic, and peering points), a capacity K_i (limiting the rate of packets that can pass through the agent), a packet-queue Q_i , and a set of neighbor agents Γ_i . A necessary, but not sufficient, condition for two agents to be connected is that their locations overlap in at least one pixel. The locations exist on an $L_x \times L_y$ square grid. A pixel of the grid is characterized by its population $p(x, y)$ and the set of agents with a presence there $\mathcal{A}(x, y)$. The total number of agents in the simulation is denoted by n , and the number of links between agents by m . These quantities, except L_x and L_y , depend on the simulation time. The outer loop of the model then iterates over the following steps:

1. *Network growth.* The number of agents is increased. Existing agents expand geographically, and their capacities are adjusted.
2. *Network traffic.* Packets are created, propagated toward their targets, and delivered. This process is repeated N_{traffic} times before the next network-growth step.

We measure simulation time τ as the number of times Step 1 is executed (the time unit between packet movements is $1/N_{\text{traffic}}$). In the remainder of this section we describe the growth and traffic steps in greater detail.

2.1 Network growth

The income of an agent, during a time step, is proportional to the traffic propagated by the agent during the period. This is a simplification. For example, income could depend both on the amount of traffic and the prices for forwarding the packets set by business agreements. Assume an agent i has a budget B_i that it invests so that it can increase its traffic, and thus its profit. Since there is a possibility of congestion in the model, agent i tries first to remove bottlenecks by increasing its capacity K_i (the number of packets that the agent can transit during one time step). When the capacity is sufficient, the agent spends the rest of its budget on increasing its traffic by expanding geographically. There are three prices associated with network growth. The capacity price C_{capacity} is the price of increasing K_i one unit. For simplicity we let C_{capacity} be independent of the size of the agent’s subnetwork. The wire price C_{wire} is the price per pixel between a new location and the agent’s closest existing location. Finally, C_{connect} is the cost of connecting two agents with locations at the same pixel.

The average degree (number of neighbors of an AS) in the AS graph has been relatively constant over time [23, 11] (increasing about 5% from 2001 to 2007).¹ We take this as a constraint in the model and let the desired average degree k_D be a control parameter. We also assume that each agent tries to spend all of its budget, but not more than that, whenever it is updated.

The network growth step iterates over the following steps:

1. *Increase of the number of agents.* As long as the network is too dense (i.e. if $2m > k_D n$), new agents are added. New agents are situated in the pixel (x, y) that has the highest available population $p(x, y)/(A(x, y) + 1)$ where $A(x, y)$ is the cardinality of $\mathcal{A}(x, y)$ and $A(x, y) \geq 1$. The budget and capacity of the new agents are initialized to B_{init} and K_{init} respectively.
If the network is small, $n < k_D + 1$, it is not dense enough for new agents to be added in step 1. Thus, we do not apply this condition when n is less than a threshold n_0 and call the time when $n = n_0$ is reached t_0 .
2. *Capacity increase.* Each agent synchronously increases its subnetwork’s capacity based upon traffic from the last time step (but not more than the agent can afford). Agent i invests the minimum of $(B_i, C_{\text{capacity}} \Delta T_i, 0, 0)$ to increase capacity (ΔT_i is the change in traffic propagated by i since the last update).
3. *Link addition.* While $2m \leq nk_D$ (which usually means $k_D - 1$ times), choose two agents randomly that are not already connected and share a common pixel. If the budgets of both agents are larger than C_{connect} , then connect them.
4. *Spatial extension.* Let the agents with remaining budget extend their networks. Iterate through all agents i and add a location at the pixel, not in Λ_i , that has

¹This calculation is based on data from Oregon Route Views, www.routeviews.org. Although more edges of the AS graph can be identified by combining multiple data sources, the Route Views data set has been compiled in a consistent way over the years, so we believe that the relative degree increase is reliable.




Figure 2: Illustration of traffic simulation. (a) A packet is created with source pixel s and target pixel t with probability proportional to the product of populations at s and t . One of the agents at the target pixel is randomly chosen as the target agent. The propagation of the packet is shown in the graph. Each agent i is associated with a queue Q_i and a capacity K_i . When a packet reaches an agent, it is appended to Q_i . K_i packets in the queue are relayed to neighboring agents and i 's budget is credited one unit. The arrows in (b) symbolize the packet's route from source to destination agent. The package is routed to a neighboring agent j with probability $\exp((d(i,t) - d(j,t))/\lambda)$ (where t is the packet's target, $d(\cdot, \cdot)$ gives the graph distance, and λ is a parameter).

the highest available population $p(x,y)/(L(x,y) + 1)$, and is not further than $(B_i - C_{\text{connect}})/C_{\text{wire}}$ from a location in Λ_i (i.e., not further from i than i can afford). (See Figure 1(b)). Alternatively, the algorithm could select the point with the lowest cost per unit of population. However, such an algorithm is computationally prohibitive for studying networks of the Internet's scale.

Each agent's budget is updated immediately after each modification.

2.2 Network traffic

We model traffic with a discrete, packet-exchange model [18, 12]. The packets are generated with specific source and target pixels, but the routing takes place on the network of agents. We neglect intradomain routing among the agent's locations, assuming that the time it takes for a packet to pass through an agent is independent of the specific locations it visits. The dynamics are defined as follows:

1. *Packet generation.* We assume that most traffic originates from direct communication between individuals and does not depend on the distance between them. For each pair of points $[(x,y), (x',y')]$ on the grid, we create a packet with source (x,y) and destination (x',y') with probability $P_{\text{pkg}} p(x,y) p(x',y')$, where P_{pkg} is a parameter that controls the rate at which new packets are created. Then, an agent is selected at random from those at the source pixel to become the

source node. The destination agent is randomly chosen from the agents at the destination pixel. Finally, one unit of credit is added to the sender's budget.

2. *Packet propagation.* Each agent i propagates the first K_i packets from its queue (of length l_i) each time step and receives one unit credit for each propagated packet. A packet can travel only one hop (inter-AS transmission) per time step. A packet at agent i is propagated to a neighbor j with probability $\exp(\lambda(d(i,t) - d(j,t)))$ (where t is the recipient AS, $d(\cdot, \cdot)$ is the graph distance, and λ is a parameter controlling the deviation from shortest-path routing [26] observed in Ref. [16]).
3. *Packet delivery.* For all agents, delete all packets that have reached their target.

The assumption in step 1 that the probability of two agents communicating is independent of their spatial separation agrees with the (somewhat debated) "death of distance" in the Internet age [5]. We also tested communication rates that decay with the square of the distance, as observed in conventional trade firms [20], with qualitatively similar results. Our traffic propagation model is simplified from reality, and it more closely resembles peer-to-peer traffic than user-to-service traffic. We also assume that temporal fluctuations in packet generation are negligible and ignore peak levels of congestion. Because the economy of the agents grows as function of accumulated traffic through their sub-networks, average traffic load is a reasonable approximation. Given the level of abstraction in our model, we believe these traffic propagation assumptions are reasonable.

Business agreements between ASes are an important factor in the Border Gateway Protocol (BGP) [24] (the Internet's largest scale routing protocol). Next hops are often selected by cost, rather than path length. We do not explicitly include inter-AS contractual agreements. However, our probabilistic propagation method 2 has a similar effect on average path length, the excess distance of real paths traveled compared to the shortest graph distance, as that observed for real Internet traffic [16].

3. NUMERICAL SIMULATIONS

3.1 Parameter values

Before presenting the simulation results, we describe the experimental design, and choice of parameters. First, we specify a population profile $p(x,y)$. We primarily model population distributions, but we also study specific geographic populations (e.g. U.S.A. census data). To simplify the generation of population distributions, we neglect spatial correlations and simply model the frequency of population densities. This frequency has two important features: it is skewed (pixels with low population densities are more frequent than highly populated pixels) and fat-tailed (there are pixels with a population density many orders of magnitude larger than the average). One probability distribution with such features is the power-law distribution $\text{Prob } p \sim p^{-\chi}$. To reduce the fluctuations between different realizations of $\{p(x,y)\}$, and prevent unrealistically high populations within a pixel, we sample the power-law distribution in the bounded interval $[1, (L_x L_y)^{1/(1-\chi)}]$ [10] with




Figure 3: Time evolution of an example run. In panel (a) the number of agents and the number of inter-agent links as a function of simulation time. In (b) the fraction of the landscape with network coverage, and the fraction of the population reached by the network, is plotted against time. Panel (c) shows the average travel time $\langle \tau_p \rangle$ for packets and the average distance (number of inter-agent hops) in the network $\langle d \rangle$, as functions of the number of agents.

$\chi = 3$. Our results do not depend strongly on the distribution of $p(x, y)$. We obtain qualitatively similar results with normally distributed p values and real population-density maps (data not shown).

In multiparameter, agent-based models, such as ASIM, a systematic investigation of the full parameter space is infeasible. Parameters are, if possible, obtained from real systems. We set the desired degree $k_D = 5.52$, the same average degree reported in Ref. [19]. Unless otherwise stated, the desired size of the network is $n_D = 16,000$, which is the same order of magnitude as the current AS graph. Other parameters are balanced to keep runtime low (less than one day) while still engaging all aspects of the algorithm. This means, for example, that between every network update, a significant number of packets are routed through even the smallest agents, and enough packages to cause congestion pass through larger agents. Unless otherwise stated, we use the parameter set given in Table 1. Many of the results we show are taken from a single run, but we confirmed that the results are representative by comparing them to 20 other runs.

3.2 Network Growth

We begin by studying the growth of the network over time. Fig. 3(a) plots the number of agents and links as a function of simulation time for one representative run. At $\tau = \tau_0 \sim 4 \times 10^5$ the graph is sparser than k_D . Initially, the agents spend their budget on new links and increasing capacity. At $\tau \sim 1.5 \times 10^6$, the budget of the wealthier agents is sufficient to invest in wires to new locations (see Fig. 3(b)). This creates new traffic, which causes positive feedback accelerating the traffic flow, coverage, budget, and also more congestion. At $\tau \sim 1.9 \times 10^6$, $n(\tau)$ and $m(\tau)$ change from exponential to sub-exponential growth. As we see below, this is also the time when a significant level of congestion appears in the network. At about the same time, the

the network has expanded to serve entire population. With the current model, the network would continue to grow indefinitely with decreasing returns for the agents. A plausible extension would be to introduce maintenance costs that are proportional to network size, in which case the network would reach a steady state where the budgets of the agents are balanced and no further investments can be made. For $\tau \gtrsim 1.9 \times 10^6$ the increase of $n(\tau)$ is slower than exponential. This is explained by the increasing level of congestion in the system. In Fig. 3(c) we plot the average time $\langle \tau_p \rangle$ for a packet to travel from source to destination. $\langle \tau_p \rangle$ is bounded from below by the average distance (number of links in the shortest path, averaged over pairs of nodes) $\langle d \rangle$. The two curves diverge, i.e. a significant level of congestion appears, around $N = 1000$. The growth of $n(\tau)$ and $m(\tau)$ slows down at the same point. We conclude that the growth slowing is arises from a congestion-driven negative feedback. The most striking feature of network growth over time is the transition from a small network, almost constant in size, to a rapidly increasing system (around $\tau \sim 1.8 \times 10^6$). This effect is typical for technologies emerging from the interactions of a large number of agents—they need a *critical mass* of users to reach a significant fraction of the total population. One can argue that the Internet reached this critical mass in the early 1980’s when it started to span the globe. Another important point in the Internet’s history was the advent of the World Wide Web (WWW) in the early 1990’s, and with it commercial applications and access to the general public. Our model does not include applications, such as the WWW, that undeniably affect network growth. Such effects could be included by adopting a different traffic model, but for this paper we aim at simplicity and generality. In the Internet, ASes growth has been slower than the exponential increase of agents predicted by the model (bgp.potaroo.net/cidr/; read January 7, 2008). This discrepancy arises in part because we do not assume that maintenance costs are proportional to income. If maintenance costs grew super-linearly, then negative feedback could dampen growth. Other external factors, such as the centralized method for allocating and assigning AS numbers (Internet Assigned Numbers Authority, www.iana.org), might also influence the actual rate of growth experienced by the Internet.

3.3 Degree distribution

One of the most conspicuous features of AS-graphs is their skewed degree distribution (first observed in Ref. [14]), compatible with a power-law functional form [9]. In Fig. 4(a) we compare the cumulative degree distribution of our model with that of the Internet’s. We use the model network from the example run described earlier (taking data from the simulation when $N = 16,000$), and the “AS06” network of Ref. [19] (an AS-graph constructed from www.routeviews.org and www.ripe.net, with $N = 22,688$). The match between the model and the real networks is striking. Preliminary studies suggest that the slope of the curve is largely insensitive to changes in parameter values. The complexity of ASIM raises the question of what causes this emergent degree distribution. By comparing ASIM to two simpler models, we provide evidence that this is a combined effect of geographic and economic factors. The two models are: the Barabási-Albert (BA) model [3] (a general network model that explains power-law degree distribution as a “rich-gets-richer” phenomena), and the Fabrikant, Koutsoupias, and

Parameter	Interpretation	Value
$L_x = L_y$	Number of pixels in the x (and y) direction	50
N_{traffic}	Number of packets sent per simulation step	1×10^4
P_{pkg}	Constant to determine packet source and dest.	0.001
n_0	Agent growth threshold	35
K_{init}	Initial capacity of an agent	5
C_{wire}	Price per pixel for new wire	500
B_{init}	Initial budget for a new agent	3×10^5
λ	Parameter in exponential distribution	75

Table 1: Default parameters values for simulation experiments.




Figure 4: The degree distribution (cumulative mass function) of a real AS-graph (AS06) together with degree distribution of a network generated with the model (a), the BA (b) and the FKP models (c). Panel (d) is a density plot that illustrates the correlation between traffic and degree in our model runs.

Papadimitriou (FKP) [13] (explaining how power-law degree distributions can appear from trade-offs in spatial optimization).

The BA model is a growth model in which one node (and m links to attach it with the rest of the network) is added every time step. *Preferential attachment* is used to determine the endpoints of the new links—the probability of attaching to a node of degree k is proportional to k .

The FKP model is also a simple growth-model. Each time step, one node, and a link attached to it, is added to the graph. A new node i is assigned random coordinates in the unit square and attached to the old node j that minimizes $d_0(j) + \alpha|\mathbf{r}_i - \mathbf{r}_j|$. Here $d_0(j)$ is the graph distance (the number of edges in the shortest path) between j and the node added first, $|\mathbf{r}_i - \mathbf{r}_j|$ is the Euclidean distance between (the coordinates of) i and j , and α is a parameter setting the cost-balance between making new physical connections or using the existing network).

In Figs. 4(b) and (c) we plot the cumulative mass function of degree for one BA and one FKP network. The model parameter values were chosen to give networks as close as possible to the real AS-graph ($m = 5$ for the BA model, $\alpha = 4$ for the FKP model, and $N = 22,688$ for both). The slope of the BA model is steeper than the real network, and the curve for the FKP-model is flatter than the real data. To compare the goodness-of-fit, since the curves have a similar range in $\log p_k$, we measure the ratio θ of the area between the curves and the area (in the $\log p_k, \log k$ -space) spanned by the extreme values of $\log k$ and $\log p_k$. We find $\theta = 0.95\%$ for our model, 4.0% for the BA model, and 11% for the FKP model. Although both the BA and FKP models have been extended to yield better data fits [29, 1], the original forms of the models illustrate two important components of Internet growth, namely the rich-gets-richer effect driving the growth of the BA model and the spatial trade-off effect of the FKP model. A combination of these effects may explain why our model's degree distribution, and the curve of the real network, lies between those of the original BA and FKP models. In ASIM, the degrees of nodes do not directly affect the creation of new links. However, preferential attachment occurs indirectly via positive feedback—nodes with large degree acquire more traffic, and thus more budget which they can reinvest in more connections, thus increasing their degree. The effect of preferential attachment in the model is shown in Fig. 4(d), which is a plot of the probability density of a node's traffic load given its degree. Because an agent's income is correlated with the traffic that it propagates, and a larger budget will increase the possibility of creating new links, there is positive feedback between the degree and the rate of degree increase, i.e. a form of preferential attachment.




Figure 5: Radial statistics for real and model networks. Panels (a)–(c) show the radial densities of nodes for the real AS-graph and our algorithm (a), the BA (b) and FKP (c) model. Panels (d)–(f) show the average degree vs. average distance $\bar{d}$ for our algorithm, the BA, and the FKP model respectively. The data of panels (b), (c), (e), and (f) are plotted in Ref. [19] as well.

Note that the correlation in Fig. 4(d) is not linear (the slope is different from the solid line). It is known that nonlinear preferential attachment does not give a power-law degree distribution [21] (which we seem to have), so preferential attachment is not the only factor affecting the network's growth. Also, if we had linear preferential attachment, the slope of $P(k)$ would be the same as the BA model.

3.4 Radial structure

Structurally, the AS graph is hierarchically ordered [28]—engineers and network operators speak of the first, second and third tier. For the model networks, we measure a node's position in the hierarchy by its network centrality [19]. In Fig. 5 we diagram the average fraction of nodes and the average degree as functions of the average distance $\bar{d}$ to other nodes in the network ($\bar{d}$ is the inverse of a centrality measure, known as *closeness centrality*, so more central nodes are to the left in the diagrams). By this method we can get a *radial* picture of the AS graph structure from the center to the periphery. In Fig. 5(a)–(c) we plot the fraction of vertices at different $\bar{d}$ -values. We note that the curves of the model match the real AS-graph more closely than the BA and FKP models do. For example, they have peaks roughly corresponding to the tiers of the Internet, similar to the observed AS-graph. The shift to the left of the model curve in Fig. 5(a) can to some extent be explained by its smaller size because larger networks have larger average distances, leading to a curve displaced to the right. In brief, the BA model lacks the complex periphery of the real AS-graph. This can be seen by noting that its density is more balanced compared to the left-skewed curve of the observed network. The average degree as a function of $\bar{d}$ is less right-skewed in the BA model compared to the empirical network. As with the degree distribution, the FKP model deviates from the real network in the opposite direction compared to the BA model. Specifically, the high degree nodes of the FKP model




Figure 6: Traffic patterns of the model. (a) displays the number of extra steps d_+ in packet navigation in the real Internet compared to our model. Panel (b) shows the probability density of agents having betweenness C_B and traffic density ρ . The data is collected from twenty independent runs.

are extremely concentrated in the center of the network.

3.5 Traffic flow and congestion patterns

Section 3.2 investigated network topology and its growth. In this section we study traffic flow and how network topology affects it. In the Internet, packets do not necessarily travel the shortest distances between source and destination. Most importantly, business agreements between agents arrange agents into a hierarchy [15]. The business contracts put constraints on how packets are routed. For example, in the hierarchy, a packet normally cannot first be routed downwards (to customers), then upwards (to providers), even if that is a shorter path (this is known as the *valley free rule*). Gao and Wang [16] investigated the extra distance d_+ packets need to travel as a result of constraints such as these. They found a decaying probability distribution of d_+ , meaning that most of the traffic actually travels via shortest paths. ASIM does not have explicit business agreements that force hierarchical routing first into the core of the network and then out again. However, in most graphs a vast majority of shortest paths pass through a restricted core of the graph [17], and our traffic model routes most traffic via short (if not the shortest) paths. The d_+ distribution of our model (shown in Fig. 6(a)) matches the observation of Gao and Wang [16] ($\theta = 8.1\%$).

We proceed to investigate the relationship between graph centrality and traffic density. This can reveal how congestion and fluctuations affect routing [18]. If all agents have sufficient capacity for packets to always route along shortest paths, then traffic density along a link l will be proportional

to its *betweenness centrality*

$$C_B(l) = \sum_{i,j} \sigma_l(i,j) / \sum_{i,j} \sigma(i,j) \quad (1)$$

where $\sigma_l(i,j)$ is the number of shortest paths between nodes i and j passing through the link l , and $\sigma(i,j)$ is the total number of shortest paths between i and j . If an AS is congested, the traffic through its links will be lower than anticipated by the betweenness of the edge. Thus, congestion patterns can be illustrated by studying betweenness and traffic load. Fig. 6(b) is a density plot of the actual traffic density as a function of betweenness of the links of the model network. For more central nodes (higher betweenness), there is a strong correlation between betweenness and traffic density—the vertices with $C_B \approx 4 \times 10^5$ spans half a decade of ρ . For the more peripheral nodes the correlation is less clear (vertices with $C_B \approx 5 \times 10^4$ can have ρ -values of almost three orders of magnitude). Indeed, there seems to be a separation of agents into two classes, one comprised of agents with the capacity to keep traffic flowing and another with inadequate capacity. For links of low betweenness the traffic/betweenness correlation is weak. To summarize, congestion does affect the system, and it is most pronounced for nodes carrying little or intermediate traffic levels.

3.6 Geographic structure

We briefly discuss the spatial network structure—another feature that emerges from the model. As an example, we ran the simulation on the population density profile of the United States. In Fig. 7(a)–(d) we show the growth of the largest agent for a run with $n_D = 20$, $L_x = 513$ and $L_y = 323$. Lines are drawn between each node (pixel) and the agent’s nearest node at the time of the node’s addition. In this representation the length of the lines are proportional to the wire cost. Fig. 7(e) and (f) plot the locations of Tier 1 exchange points of two major Internet providers Sprint and AT&T (adapted from Ref. [27]). There are some similarities between these real networks and the model network of Fig. 7(d)—all networks span the whole continent and have locations concentrated in urban areas. In Ref. [22] the authors observe a super linear scaling relationship between the density of servers and the population density with an exponent between 1.2 and 1.7. Our model is consistent with this observation (with an exponent in the lower range of this observation). Studying spatial aspects of the model more carefully is an area of future research.

4. DISCUSSION

We have presented a mechanistic model of communication networks that, like the AS-level Internet, is built of spatially extended subnetworks that have an interest in increasing the traffic running through them. Our model networks grow slowly until they reach a critical mass where an approximately exponential growth begins; they match the degree distribution of real networks and the radial statistics closely. The degree distributions of both the model and the real world lie between the distributions of the pure BA and FKP models. Because ASIM incorporates aspects of both the BA and FKP models we hypothesize that this macro-feature arises from the combination of preferential attachment (of the BA model) and geographically constrained optimization (of the FKP model). ASIM recreates important

traffic characteristics observed in real Internet traffic. And, when we run the model on the US population density map many features of the backbone of existing large agents are recreated.

The different aspects of the model (traffic, geography, and economy) all affect the output. In this paper we do not scrutinize the model’s parameter dependence, although preliminary studies suggest that the speed of growth (quantified by e.g. the time to reach the critical density) is strongly dependent on both the wire and attachment prices, the population density profile (a more clumped population distribution produces faster growth), and their desire to communicate. On the other hand, network topology is rather insensitive to the population distribution, and also not very dependent on how sources and destinations are generated (e.g., introducing a distance dependence does not matter much). The actual layout of the network, however, does depend on the population profile.

Many interesting extensions of ASIM are possible. For example, we could include business agreements between the different agents (similar to Ref. [25, 8]), or change the traffic patterns from person-to-person communication to a situation with more traffic originating from central servers. We could also model intra-AS routing. Many of today’s ASes employ “hot-potato” routing and transfer packets to the next AS as quickly as possible, to reduce cost. Alternative intra-AS routing strategies, such as routing the packet as close to the destination as possible, could be tested within the model’s framework.

Acknowledgements

The authors would like to thank Allen Downey for his helpful comments. PH acknowledges financial support from the Swedish Foundation for Strategic Research. SF acknowledges the support of the National Science Foundation (grants CCF 0621900 and CCR-0331580), the Air Force Office of Scientific Research MURI grant FA9550-07-1-0532, and the Santa Fe Institute.

5. REFERENCES

- [1] J. I. Alvarez-Hamelin and N. Schabanel. An Internet graph model based on trade-off optimization. *Eur. Phys. J. B*, 38:231–237, 2004.
- [2] S. Bar, M. Gonena, and A. Wool. A geographic directed preferential Internet topology model. *Computer Networks*, 51:4174–4188, 2007.
- [3] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [4] E. Bonabeau. Agent-based modeling: Methods and techniques for simulating human systems. *Proc Natl Acad Sci*, 99:7280–7287, 2002.
- [5] F. Cairncross. *The death of distance*. Harvard Business School Press, Boston, MA, 1997.
- [6] J. M. Carlson and J. Doyle. Highly optimized tolerance: a mechanism for power laws in designed systems. *Phys. Rev. E*, 60:1412–1427, August 1999.
- [7] H. Chang, S. Jamin, and W. Willinger. Internet connectivity at the AS-level: an optimization-driven modeling approach. In *MoMeTools ’03: Proceedings of the ACM SIGCOMM workshop on Models, methods and tools for reproducible network research*, pages 33–46, New York, NY, USA, 2003. ACM.




Figure 7: The spatial expansion of a single agent with the US population density as model input. The simulation parameters are the same as the rest of the paper, except $n_D = 20$, $L_x = 513$ and $L_y = 323$. Panels (e) and (f) represent the points of presence of AT&T and Sprint within the United States. This data was adapted from Ref. [27].

- [8] H. Chang, S. Jamin, and W. Willinger. To peer or not to peer: Modeling the evolution of the Internet's AS-level topology. In *Proc. IEEE INFOCOM*, 2006.
- [9] A. Clauset, C. R. Shalizi, and M. E. J. Newman. Power-law distributions in empirical data. e-print arXiv:0706.1062, 2007.
- [10] R. Cohen, K. Erez, D. ben Avraham, and S. Havlin. Resilience of the Internet to random breakdowns. *Phys. Rev. Lett.*, 85:4626–4628, 2000.
- [11] I. Daubechies, K. Drakakis, and T. Khovanova. A detailed study of the attachment strategies of new autonomous systems in the AS connectivity graph. *Internet Mathematics*, 2:185–246, 2006.
- [12] P. Echenique, J. Gómez-Gardēnes, and Y. Moreno. Dynamics of jamming transitions in complex networks. *Europhys. Lett.*, 71:325–331, 2005.
- [13] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou. Heuristically optimized trade-offs: A new paradigm for power laws in the Internet. In *Proceedings of the 29th International Conference on Automata, Languages, and Programming*, volume 2380 of *Lecture notes in Computer science*, pages 110–122, Heidelberg, 2002. Springer.
- [14] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the Internet topology. *Comput. Commun. Rev.*, 29:251–262, 1999.
- [15] L. Gao. On inferring autonomous system relationships in the Internet. *IEEE / ACM Transactions on Networking*, 9:733–745, 2001.
- [16] L. Gao and F. Wang. The extent of AS path inflation by routing policies. In *Proceedings of GLOBECOM '02*, volume 3, pages 2180–2184, 2002.
- [17] K.-I. Goh, E. Oh, H. Jeong, B. Kahng, and D. Kim. Classification of scale-free networks. *Proc. Natl. Acad. Sci. USA*, 99:12583–12588, 2002.
- [18] P. Holme. Congestion and centrality in traffic flow on complex networks. *Advances in Complex Systems*, 6:163–176, 2003.
- [19] P. Holme, J. Karlin, and S. Forrest. Radial structure of the Internet. *Proc. R. Soc. A*, 463:1231–1246, 2007.
- [20] W. Isard. *Location and space economy*. MIT Press, Cambridge MA, 1956.
- [21] P. L. Krapivsky, S. Redner, and F. Leyvraz. Connectivity of growing random networks. *Phys. Rev. Lett.*, 85:4629 – 4632, 2000.
- [22] A. Lakhina, J. W. Byers, M. Crovella, and I. Matta. On the geographic location of Internet resources. Technical Report BUCS-TR-2002-015, Boston University, 2002.
- [23] R. Pastor-Santorrás and A. Vespignani. *Evolution and structure of the Internet: a statistical physics approach*. Cambridge University Press, Cambridge, 2004.
- [24] Y. Rekhter and T. Li. A Border Gateway Protocol 4 (BGP-4). RFC 1771 (Draft Standard), Mar. 1995. Obsoleted by RFC 4271.
- [25] S. Shakkottai, T. Vest, D. Krioukov, and K. C. Claffy. Economic evolution of the Internet AS-level ecosystem. e-print arxiv:cs.NI/0608058, 2006.
- [26] V. Sood and P. Grassberger. Localization transition of

- biased random walks on random networks. *Phys. Rev. Lett.*, 99:098701, 2007.
- [27] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with Rocketfuel. *IEEE / ACM Transactions of Networking*, 12:2–16, 2004.
- [28] L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz. Characterizing the Internet hierarchy from multiple vantage points. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 618–627, 2002.
- [29] S.-H. Yook, H. Jeong, and A.-L. Barabási. Modeling the Internet’s large-scale topology. *Proc. Natl. Acad. Sci. USA*, 99:13382–13386, 2002.