

Simulating Large-Scale Networks with Analytical Models

Matthias Scheidegger (corresponding author), Florian Baumgartner, Torsten Braun

Institute of Computer Science and Applied Mathematics
University of Bern, Neubrückestrasse 10, CH-3012 Bern
Email: (mscheid|baumgart|braun)@iam.unibe.ch
Phone: +41-31-6318692 Fax: +41-31-6313261

Abstract

Discrete event simulation of computer networks has significant scalability issues, which makes simulating large scale networks problematic. We propose a high-level abstraction modeling network domains, inter-domain links and traffic with highly scalable analytical models, which is much more efficient but slightly less accurate than node-by-node models. Thus, simulation scenarios containing several ISP networks become feasible. We also propose a way to combine this modeling approach with traditional packet-based simulators and present some preliminary evaluation results of the concept.

1 Introduction

In traditional packet-based simulators networks are modeled in terms of nodes and links with individual

capacities and delay characteristics. When simulating whole Internet domains this approach quickly becomes problematic, due to the huge amount of events to be processed. Many approaches to this scalability problem have been proposed, each with slightly different application ranges. Parallel simulation ([1], [2]) is probably the most prominent one, but there are also approaches such as fluid flow simulation ([3], [4], [5]), time stepped hybrid simulation [6] and packet trains [7], amongst others.

Scalability in network simulation is generally achieved by reducing the level of detail of the simulation scenario or of the simulation algorithm. Carefully chosen, such abstractions of the simulated network can significantly reduce the complexity of large scale simulations. In this paper we propose a model that aims for far more efficient simulations than traditional approaches while still giving a good approximation of real

network behavior.

This model is based on the assumption that, over certain time spans, networks like the Internet can be divided into areas where congestion is negligible, interconnected by bottleneck links. We treat congestion free areas as black boxes, which we call *domain models*. Modelling congestion free areas has the advantage that we can neglect packet losses and excessive queuing in large parts of the network and restrict the model to quasi-static delay behavior. Apart from its scalability advantage this approach may be useful to model network areas of which we do not know the exact topology. Domain models are based on empirical cumulative distribution functions (ECDFs) to simulate the delays of packets crossing the domain. The ECDF is chosen depending on the ingress and egress nodes on which the packet enters and leaves the domain, respectively.

The bottleneck links between two domains of a simulation scenario are represented by *inter-domain link models*. Here, packet loss and the effects of queuing on delay are simulated. The basic parameters of an inter-domain link model are similar to those of a link in a packet-based simulator. Nonetheless, inter-domain link models are not event-driven but rely on parameters like offered load and link capacity. Fig. 1 shows this model-




Figure 1: The basic modeling view

ing view.

One useful application of this abstraction may be to partition backbone networks into autonomous systems (\Rightarrow domain models) and their inter-connecting links (\Rightarrow inter-domain link models). This partitioning is reasonable since autonomous systems are usually controlled. For example, internal routes can be changed to distribute traffic load, and ingress routers may police flows to prevent congestion inside the AS. Moreover, the bandwidths inside an AS network are usually bigger than the bandwidths of inter-domain links.

Further components of this model system are the *application traffic models* concerned with traffic load. They serve as scalable models for large aggregates of application traffic like VoIP, Video, HTTP, etc. They take the form of a function that yields the load gen-

erated by the traffic aggregate given a (monotonously rising) point in time.

By combining domain, inter-domain link and application traffic models we create a *multi-domain model*. Multi-domain models can be viewed as an equivalent to simulation scenarios in packet-based simulations. A simulator could be written based solely on these models. For several reasons it is desirable to combine these models with packet-based simulation, however. The behavior of an individual flow is easier to describe as a packet-based model, and many of today's protocols and applications have been implemented for packet-based simulators. Furthermore, a combination of fine grained packet-based simulation and coarse grained analytical models could be very useful in scenarios like a multi-site virtual private network, for example.

The remainder of the paper is organized as follows: Section 2 describes how domain, inter-domain link and application traffic models are combined to multi-domain models. Sections 3 and 4 go into further detail on domain and inter-domain link models, respectively. Section 5 discusses the combination of analytical models with packet-based simulation, and in Section 6 we present some preliminary evaluation of the concept. Section 7 concludes the paper.

2 Multi-Domain Models

The purpose of a multi-domain model is to organize and control domain models, inter-domain link models and application traffic models to form a single analytical model. Thus, the basis of a multi-domain model is a set of such models and their parameters, e.g. delay characteristics for domain models and the link capacity for inter-domain link models. In order to combine these model sets to a multi-domain model additional information is required. The topology of a multi-domain model is a directed graph, where the domain and application traffic models are the vertices and the inter-domain link models—which are always simplex—are the edges. Accordingly, standard ways to represent graph topologies can be used, e.g. vertex and edge tables. Routing information is required to map the application traffic models' generated load to the correct inter-domain links. Each route is stored as a sequence of inter-domain links, which again can be implemented using tables. In combination with the topology this is sufficient to resolve all models along a routing path.

2.1 Multi-Domain Load

While the inter-domain link models simulate the effects of network load on a single link, it is the task of multi-

domain models to simulate the distribution of network load among the its inter-domain links. This is the basis for packet loss and delay behavior in the modeled network. Note that unlike packet-based simulators where events trigger an update of the system, we do not have to update unless we want to inspect the system's state.

Given we want to inspect the system at simulation time t . Let P_s be the routing path (a sequence of inter-domain links) of the traffic originating at source s (an application traffic model), and let $s(t)$ be the load generated by the traffic aggregate at time t . The processed load of an inter-domain link L depending on the offered load λ is written as $L(\lambda)$. Here, inter-domain links take the role of a function with $0 \leq L(\lambda) \leq L$. If link L_i directly follows link L_j on a path we call L_j a predecessor of L_i .

Now, we calculate the along a path $P_s = \{L_1, L_2, \dots\}$ using the sequence

$$s(t), \quad L_1(s(t)), \quad L_2(L_1(s(t))), \quad \dots$$

until a link on the path has more than one predecessor, or until the path ends. Then, we start over with the next path, and so forth, until all paths either have ended or have reached a link with more than one predecessor. Now we can return to the first path. The offered load λ_{L_i} on the link in question is given by the sum of the

processed loads of all predecessors. If the last calculated element of the path's load sequence was λ we can now calculate the next element with

$$\lambda' = \frac{L_i(\lambda_{L_i})}{\lambda_{L_i}} \lambda.$$

We continue with this procedure until all paths have been followed to their end and the offered and processed loads of all inter-domain links are known.

The above algorithm may be optimized in several ways. First, when updating the system we only have to pursue changes in the offered load as far as the make a difference for the whole system. For example, if a traffic model overloads the first link on its path on one update, any additional load in the next update will influence only this first link. The processed load of this link stays the same. Furthermore, changes in the offered load may be marginal, in which case we can ignore this change at the cost of reduced accuracy. However, in order not to accumulate errors we then have to force updates in regular intervals.

2.2 Multi-Domain Loss and Delay

Based on the load distribution calculations above, the delay distributions and packet loss ratios of a multi-domain model's paths can be found. The packet loss

ratio along the path $P = L_1, \dots, L_n$ is given by

$$1 - \prod_{i=1}^n \frac{\lambda_{L_i} - L_i(\lambda_{L_i})}{\lambda_{L_i}}$$

where λ_{L_i} is again the offered load on link L_i .

Delays along a path are similarly modeled. The time it takes for a packet to traverse a domain or an inter-domain link can be described as a random variable. Let δ_L be the random variable of the delay caused by inter-domain link L , and let $\delta_{L,K}$ be the random variable of delay in the domain between the inter-domain links L and K ($\delta_{L,K}$ is only defined if L is a predecessor of K). Then the delay distribution on the routing path $P = L_1, \dots, L_n$ is given by

$$\delta_P = \sum_{i=1}^n \delta_{L_i} + \sum_{i=1}^{n-1} \delta_{L_i, L_{i+1}}$$

In a simulation we need to generate random values accordingly. This can be easily done by generating random values for each of the random variables and summing them up. The fact that the delay distributions of domain models do not change can be used to make this procedure much more efficient, however. Since these delay distributions are discrete their distribution functions can easily be convoluted into a single one, which reduces the task of simulating the domain delays to the generation of a single random value. The convolution can be performed efficiently by using the fast

fourier transform algorithm. Convoluting the link delay distributions isn't efficient in normal scenarios as they change rather rapidly according the load distribution. Having a random variable of a path's delay further allows to easily calculate moments like the mean delay or the path's jitter, which would be $\text{Var}(\delta_P)$ if interpreted as delay variation.

3 Domain Models

Domain models represent network “clouds” in a simulation scenario, e.g. ISP networks or BGP autonomous systems, or parts of them. The partitioning of a topology into domains and inter-domain links can be freely configured but it must be chosen so as to satisfy the basic assumptions—packet loss occurs only in inter-domain links—as well as possible. Network “clouds” of nodes under a common management (e.g. an ISP) are good candidates. Moreover, experience suggests that the further away from the network edge a domain is located, the better the basic assumption holds. This is due to the traffic smoothing effect observed in backbone networks.

The chosen abstraction allows that domain models only simulate the delay behavior of a network cloud and

do not react to changes of network load. Domain models are black boxes; their interior structure is not explicitly modeled. The highest level of detail in a domain model is the distinction of paths through the domain. A model of a domain with n edge nodes can thus contain $n(n-1)$ delay models, one for each ingress-egress node combination. Simpler cases with only one common delay model for all paths are useful if there isn't enough information available about the network area. Using simple models can also significantly reduce the memory consumption of a simulation.

During preliminary evaluation we found that empirical cumulative distribution functions (ECDFs) are well suited to model the delay behavior of network domains. They can be easily built from a series of delay measurements taken from a real network. In the optimal case, one-way delays should be used, but as this requires clock synchronisation of the measurement endpoints we can also approximate them by taking roundtrip times divided by two. This requires a nearly symmetrical path, however.

An ECDF is built from a sample by storing the observations in a table. Random values following the same distribution are then generated by randomly selecting table entries using a uniform distribution. Given a suf-




Figure 2: Generating random values using an interpolated ECDF

ficient sample size, this approach yields very good results if the basic assumption is not violated. The size of large tables can be reduced by using linear interpolation. The procedure can be seen in Fig. 2. We start by generating a random value x , uniformly distributed on $[0, 1]$, which designates a position in the sorted observation table (seen as a step function in the Figure). The two nearest observations are then interpolated to get a simulated delay value $f(x)$.

It is important to note that ECDF models, while giving good reproductions of observed first and second order moments in measurements, ignore any non-stationarity of the sample.

4 Inter-Domain Link Models

Inter-domain link models cover the dynamic parts of network behavior, like the effect of queuing and overload on delay and packet loss. Since they represent a single physical link between the interfaces of two nodes it is an obvious approach to model them as an analytical queuing system. Queuing theory is the traditional approach to this. Although creating a queuing model for a given system is often non-trivial, the results are both accurate and efficient.

We chose the relatively simple M/M/1/K queue as a first approximation, that is, a queue with markovian (a.k.a. poisson) arrival and service processes, a single “processing station” (the physical link) and system capacity K . Recent work [8] suggests that the arrival process would be better modeled as a batch markovian process M[k]. The arrival and services rates λ and μ depend on the offered load on the link and the links capacity, respectively. K —the queue length plus one—can be set to a typical value.

In order to model the behavior of the inter-domain link we have to find the probability p_i of the system to be in state i , where state K means the queue is full, and state 0 means the system is empty and does not send. The M/M/1/K queue is a birth and death process




Figure 3: Birth and Death Process

as shown in Figure 3. For a birth and death process of this kind the probabilities p_i are given by

$$p_i = \begin{cases} \frac{1-\lambda/\mu}{1-(\lambda/\mu)^{K+1}}, & i = 0 \\ (\lambda/\mu)^i p_0, & i > 0 \end{cases} \quad (1)$$

if $\lambda \neq \mu$, and

$$p_0 = p_1 = \dots = p_K = \frac{1}{K+1} \quad (2)$$

if $\lambda = \mu$. As states above, p_K is the probability of the system being full. Therefore, p_K is also the loss rate of the link. The functional representation of the inter-domain link used in Section 2.1 can thus be written as $L(\lambda) = (1 - p_K)\lambda$, with p_K calculated according to formulas above.

From the probabilities p_i we can further construct a discrete density function of the link’s delay distribution. The number of bytes that are in the system when another byte arrives is proportional to the time this byte has to wait before it is sent to the link. If δ_{pr} is the propagation delay on the link the discrete delay distribution looks like this

$$\begin{pmatrix} p_0 & p_1 & \dots & p_K \\ \delta_{pr} & \delta_{pr} + \frac{1}{\mu} & \dots & \delta_{pr} + \frac{K}{\mu} \end{pmatrix} \quad (3)$$

5 Hybrid Simulation

Creating a hybrid simulation of packet-based and analytical models enables makes it possible to combine large scale, coarse-grained topologies with fine-grained models for points of special interest in the scenario. Especially simulations of multi-site corporate VPNs may benefit from this: The intranet components can be modeled using the packet-based approach, while the inter-site connections over the public internet benefit from the efficiency boost of analytical simulation.

We propose to enhance traditional packet-based simulators by enabling their nodes to contain analytical multi-domain models. In this way a simulator node can stand for and behave like a whole network cloud (typically a multi-domain model). Figure 4 shows an example for this. When a simulated packet reaches an enhanced node, it triggers an inspection of the underlying multi-domain model to determine how much the packet should be delayed and whether it should be forwarded at all. Both decisions are based on the cumulative forwarding probability and delay distribution calculations described in Section 2.

This approach necessitates a new load generator in the multi-domain models: the bandwidth estimator (BE). It converts packet reception events to a bandwidth




Figure 4: Enhanced node in a packet-based simulator

estimate for every routing path between an ingress and an egress node of the multi-domain model. A good way to estimate bandwidths from packet events is to use a sliding time window algorithm. The number of bytes received in the time window Δt is added up and divided by Δt .

While packets generated in the event-driven simulator influence the analytical models inside enhanced nodes, load generated by the application traffic models of a multi-domain model do not create additional packets outside of the enhanced node. Our approach only allows packets to go through enhanced nodes, not to be created by them. The reason for that is the higher level of abstraction used in multi-domain models.

6 Evaluation

We implemented the concept of hybrid simulation in the `ns2` [10] simulator by extending the simulator with a mechanism that makes it possible to overload the behavior of the simulator nodes with arbitrary loadable modules. The analytical models presented in this paper were implemented in such a module. All parameters of the models can be configured using XML files. During preliminary evaluation we tested the behavior of the implemented inter-domain link and domain models.




Figure 5: Comparison of `ns2` and analytical link: transfer rates

We compared the inter-domain link model a standard `ns2` link. For this we used a scenario with three consecutive links, of which the middle one was the 2 Mbps bottleneck and studied the behavior of this link under five kinds of traffic load: 1 Mbps, 2 Mbps and 4 Mbps CBR traffic, FTP traffic (5 sources), and a mix of FTP traffic (3 sources) and 1 Mbps CBR traffic. Fig. 5 shows a comparison of the transfer rates achieved with the

`ns2` link and the analytical link model. While in the CBR and FTP cases the performance is good, the mix of CBR and FTP (i.e. TCP) seems to be more problematic: With both link types the transfer rate decreases but not by equal amounts. We believe this is due to the stochastic nature of dropping in the analytical model, similar to the behavior of random early detection (RED) queues, which are known to enable higher transfer rates with TCP than traditional drop-tail queues do.

As a preliminary evaluation of the domain model, the delay characteristics between the network of the University of Bern and the ETH Zürich have been measured. In a first step the delay between two hosts in the networks was measured. Both networks are connected by the Swiss scientific network SWITCH [9], and the distance between the measurement hosts was nine hops. Based on the measurements an empirical distribution was computed and used to configure the domain model. For the simulation the simple `ns2` network in Figure 6 with three nodes was set up. While the two outer nodes act as source and sink, the central node has the domain model attached.

Figure 7 shows a comparison between the measured delays and the delays in the simulation. Both graphs show almost exactly the same delay behavior for the




Figure 6: ns2 setup simulate the delay of a single ISP

measurement and the simulation.




Figure 7: Delay histograms from measurements (upper graph) and simulation (lower graph)

7 Conclusion

In this paper we have presented a scalable approach to simulating large scale inter-domain networks. This scalability is achieved by partitioning the simulation scenario into congested bottleneck links and the congestion-free areas in between, and by creating analytical models for both (inter-domain link models and

domain models, respectively). These models are configured by measuring the characteristics of a live network and can then predict delay and dropping behavior of this network. We have further presented a concept to combine these high-level models with traditional packet-based simulators, which we implemented in the ns2 simulator. Some preliminary evaluation was also done for the basic models, comparing an inter-domain link model with a link model of the ns2 simulator, and comparing the measured delay between two real network nodes to the simulated delay of a correspondingly configured domain model in ns2 .

Acknowledgements

This work was funded by the Swiss Federal Office for Education and Science (BBW, contract number 01.0551) and was done in the context of the InterMON project [11], which is part of the European Union’s 5th Framework Programme.

References

- [1] K. M. Chandy and J. Misra, “Asynchronous distributed simulation via a sequence of parallel com-

- putations,” *Communications of the ACM*, vol. 11, no. 24, pp. 198–205, April 1981.
- [2] M. H. Ammar, G. F. Riley, and R. M. Fujimoto, “A generic framework for parallelization of network simulations,” in *MASCOTS’99*, College Park, MD, October 1999.
- [3] A. Yan and W. B. Gong, “Fluid simulation for high speed networks with flow-based routing,” *IEEE Transactions on Information Theory*, pp. 1588–1599, 1999.
- [4] B. Liu, Y. Guo, J. Kurose, D. Towsley, and W. Gong, “Fluid simulation of large scale networks: Issues and tradeoffs,” in *PDPTA’99*, Las Vegas, NV, June 1999, pp. 2136–2142.
- [5] B. Liu, D. R. Figueirido, Y. Guo, J. Kurose, and D. Towsley, “A study of networks simulation efficiency: Fluid simulation vs. packet-level simulation,” in *Proceedings of IEEE Infocom*, April 2001.
- [6] Y. Guo, W. Gong, and D. Towsley, “Time-stepped hybrid simulation (TSHS) for large scale networks,” in *Proceedings of IEEE Infocom*, March 2000.
- [7] J. S. Ahn and P. B. Danzig, “Packet network simulation: speedup and accuracy versus timing granularity,” *IEEE/ACM Transactions on Networking*, vol. 4, no. 5, pp. 743–757, October 1996.
- [8] A. Klemm, C. Lindemann, and M. Lohmann, “Modeling ip traffic using the batch markovian arrival process,” *Performance Evaluation*, no. 52, pp. 149–173, 2003.
- [9] “Switch – the swiss education & and research network,” <http://www.switch.ch>, February 2003.
- [10] “The network simulator - ns-2, Information Science Institute, University of Southern California, <http://www.isi.edu/nsnam/ns>,” June 2002.
- [11] “InterMON, European Union project IST-2001-34123,” <http://www.ist-intermon.org>.