

An Evaluation Framework For Active Queue Management Schemes

Arkaitz Bitorika, Mathieu Robin, Meriel Huggard
Department of Computer Science
Trinity College Dublin, Ireland
Email: {bitorika, robinm, huggardm}@cs.tcd.ie

Abstract

Over the last decade numerous Active Queue Management (AQM) schemes have been proposed in the literature. Many of these studies have been directed towards improving congestion control in best-effort networks. However, there has been a notable lack of standardised performance evaluation of AQM schemes. A rigorous study of the influence of parameterisation on specific schemes and the establishment of common comparison criteria is essential for objective evaluation of the different approaches.

A framework for the detailed evaluation of AQM schemes is described in this paper. This provides a deceptively simple user interface whilst maximally exploiting relevant features of the NS2 simulator. Traffic models and network topologies are carefully chosen to characterise the target simulation environment. The credibility of the results obtained is enhanced by vigilant treatment of the simulation data. The impact of AQM schemes on global network performance is assessed using five carefully selected metrics. Thus, a comprehensive evaluation of AQM schemes may be achieved using the proposed framework.

1. Introduction

More than fifty AQM schemes have appeared in the literature since the original *Random Early Detection* (RED) proposal by Floyd and Jacobson in 1993 [11]. (These schemes are “active” because they employ a dynamic packet dropping procedure based on a feedback mechanism.) However, AQM is still not widely utilised in the Internet and the majority of routers are based on basic Drop-Tail queues. The Internet Engineering Task Force (IETF) recommended the deployment of RED for the prevention of congestion collapse in 1998 [4].

It has been suggested [12] that the slow pace of AQM deployment is due to a lack of detailed objective evaluation of the various schemes. The lack of consistent evaluation criteria for use in such studies has contributed to this problem.

Recently there have been some studies comparing a relatively small selection of AQM schemes (e.g. [13,24]). However there is little consistency between the various studies. This may be explained by the size of the parameter space to be considered: intrinsic parameters, those relating to evaluation metrics and those arising from the network scenario being considered must all be taken into account. Furthermore, apparently different results may be obtained from functionally identical, high-level parameterisations if differing underlying models or simulation techniques are used.

An integrated framework for benchmarking AQM schemes is proposed. This provides sophisticated capabilities for the evaluation of AQM schemes using a variety of simulation scenarios. The framework provides an intuitive, powerful interface to the network simulator NS2 [5]. To obtain statistically correct results the framework calculates the required number of simulations, together with the required time for each simulation, dynamically. It provides features found in advanced simulation tools, such as AKAROA [7] (e.g. multiple replications in parallel, output data analysis), while using NS2 as its simulation engine.

This paper is organised as follows: Section 2 presents the metrics chosen to evaluate the AQM schemes, while Section 3 outlines the network model (topology and traffic) used to simulate these schemes. In Section 4 the implementation of the framework is outlined, and ways of improving the framework credibility are discussed in Section 5. Section 6 presents sample results and the paper concludes with a discussion on possible refinements to the proposed framework.

2. Metrics

At the core of the framework are a small number of carefully chosen performance metrics. These:

- are *representative of the end-user experience* of network performance. For example, end-to-end jitter is considered rather than queue length stability. If router-based metrics are relevant to the understanding of a

given scheme's behaviour, such comparison should be based on end-user experience.

- *address a wide range of performance issues* with a small number of meaningful metrics; thus reducing the computational overhead and facilitating the comparison of schemes.
- *reflect the performance of the entire network* i.e. they do not depend on measurements at a single point of the network and so are independent of the network scenario being considered.

Five metrics representing network utilisation, delay, jitter, drop rate, and fairness have been chosen (see [3] for a more detailed discussion):

Utilisation metric: The percentage of total network capacity utilised during a simulation run. For a given traffic mix and topology, the network capacity is defined as the maximum total flow in the network. This provides an upper bound on the total goodput achievable. To compute the final utilisation metric, the sum of the aggregate goodput of all flows is divided by the computed network capacity.

Delay metric: The average end-to-end delay experienced by the packets. The average over all the packets in an individual flow is first computed and from these the average over all flows is calculated.

Jitter metric: This is based on the variance of the overall delay of packets. First the coefficient of variation of the delay for each flow is computed and these values are then averaged over all flows [3].

Drop metric: This represents the percentage of packets dropped. It includes packets dropped due to the AQM policy (early drops) as well as those dropped due to buffer overflow. We do not differentiate between these as they have the same impact on the end-user experience.

Fairness metric: The "fair share" for each flow is computed statically by reference to the capacity of the network using the fairness metric proposed by Jain et al. [14]. If each flow gets its fair share of the network capacity, the metric value is 1; this value decreases if resources are shared unevenly between the flows (see [14] for details).

3. Simulation Scenarios For AQM Evaluation

One critical factor in the design of an evaluation framework is the choice of simulation environment. A network scenario may be defined by two elements: a static network topology and the traffic mix that flows amongst its nodes. The topology and the traffic mix offer a wide parameter space to be modelled. The model chosen to address a specific research issue – AQM performance in this case – should incorporate as many of the essential characteristics of the Internet as possible.

3.1. Network Topology




Figure 1. Dumbbell Topology

The effects of different queueing algorithms on performance in a single-gateway scenario can be evaluated using a Dumbbell topology, see Figure 1. This configuration is well understood and widely implemented, and hence has been chosen for preliminary testing of the evaluation framework.

Three parameters are used to define the Dumbbell topology: the bandwidth of the bottleneck link, its propagation delay and the range of per-flow round-trip times (RTTs). The latter is one of the parameters frequently overlooked in common simulation scenarios. The different RTTs between TCP flows impacts on the share of a router queue that a flow obtains [9]. Measurement studies have shown that RTTs found on the Internet can vary widely, with most RTTs in a range between 15 and 500 ms [1].

3.2. Traffic Mix

Creating a realistic model of Internet traffic necessitates the use of a large parameter set. These parameters are obtained from measurement studies [1, 6] of Internet traffic.

A typical traffic mix used for the evaluation of router-queue management algorithms is that of long-lived FTP bulk-transfers. This mix is not representative of current Internet traffic: this is dominated by a high proportion of short-lived Web flows, or *dragonflies*, and a low number of long-lived transfers, or *tortoises* (which generate a large percentage of the traffic on the network) [6].

The effect of traffic that does not use end-to-end congestion control on network performance and throughput fairness should not be neglected [10]. The evaluation framework provides for the inclusion of such UDP traffic in any simulation scenario: the traffic mix may contain Web-like traffic with long-lived FTP and UDP flows.

4. Framework Implementation

The architecture of the framework is made up of two key components: a high-level interface for specifying simula-

tion scenarios and experiments, and an engine for running several simulations in parallel.

4.1. High-level Interface

NS2 is the *de-facto* standard network simulator for Internet research. However, NS2 limits its higher-level functionality to a scripting interface. One of the requirements of the proposed evaluation framework is that it should be able to address research issues with a minimum of configuration or scripting complexity. In doing so it should address questions like:

- What is the effect of increasing non-responsive UDP traffic on network performance and flow throughput fairness?
- How well do Active Queue Management schemes perform under different Web traffic loads?
- How sensitive to parameter settings is the performance of a given AQM scheme?

Configuring and running the simulations that will address these issues using NS2 is a time-consuming process. It involves writing OTcl scripts for establishing the simulation scenarios in NS2, running numerous simulations and processing the output from NS2 to obtain results. Moreover, NS2 does not provide any functionality to perform the output data analysis necessary to obtain valid simulation results.

Poorly chosen network scenarios may further limit the usefulness of the results obtained and may even invalidate the results if the parameters chosen are unrealistic.

Our framework uses NS2 as its “simulation engine” but provides a high level interface to it. The system has been initially targeted for AQM evaluation but can equally be used to address other network performance evaluation questions. It incorporates a Python API [22] offering two main functional groupings:

- An *Experiment* comprises a set of simulations that have a common set of fixed parameters and one variable parameter. The fixed parameters may be either user-specified or predefined default values. Experiments provide functionality for running the associated simulations and they calculate network performance metrics using the output data. For each simulation model a set of replications is run with independent streams of random numbers. From metrics values obtained in those replications estimates of mean values and their associated relative error level is reported.
- The *Report* provides the main interface to the framework. It comprises a simplified way of specifying a set

of experiments and obtaining results: The user specifies one or more parameter spaces to explore and the framework configures and runs the required NS2 simulations.

The following sample script code illustrates how the framework may be used to obtain a report showing how a UDP traffic flow with increasing load affects network performance:

```
report = Report()
report.addPlotParam('scenario.cbrBitrate',
                    range(0, 10.1))
report.addFixedParam('stopTime', 20)
report.generate()
```

This five-line script will create and run eleven NS2 simulations for each supported AQM scheme. For each scheme the rate of offered UDP traffic will vary between 0 and 10 MB/s. Once all the simulations have finished, and their output parsed, a report is generated summarising the network performance.

The framework supports numerous parameters, making it a powerful, high-level interface for carrying out simulation-based network performance experiments.

4.2. Parallel Simulation Engine

The evaluation framework may require numerous, lengthy simulations to be performed for each set of results generated. Hence, using a single workstation to generate reports is extremely time consuming. Sequential network simulators have been extended to run on parallel computers [23]. Our approach differs in that we obtain economies of scale by parallelising each experiment at the granularity of one simulation. Substantial speed advantages may be obtained from distributing the load amongst a set of networked workstations. Individual simulations still run using a standard sequential NS2 simulator thus providing good scalability without the need for modification of the NS2 software.

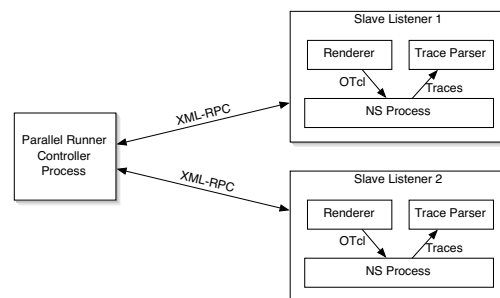


Figure 2. Simulation engine architecture

Figure 2 illustrates the architecture of the framework when using the parallel simulation engine with two “slave”

hosts. The “slaves” execute the simulations and parse the output traces. The controller process communicates with the slaves using XML-RPC, a remote procedure call protocol based on HTTP and XML. The controller sends scenario parameters with the requests and the slaves reply with a summary of the results of the simulation. Once all the simulations have finished and their results have been collected, the framework calculates experiment-level network performance metrics.

5. Improving the framework credibility

While some have advocated better simulation models for Internet research [12], there seems to be little concern regarding the analysis and credibility of simulation results. It has been argued [20] that a credibility crisis exists due to the widespread lack of correct analysis of simulation output data. In this section the statistical techniques and procedures implemented in the framework are described. These may be used to obtain meaningful results from simulation-based network performance studies.

5.1. Analysis of a single run

The framework should obtain steady-state values for the five metrics described in section 2, however for one single run of the simulation, two significant problems are encountered:

- The estimation of the steady-state value is biased by use of an atypical initial state (i.e. empty queues and links).
- Once the initial, transient, phase is over, it is necessary to determine how long the simulation run should be in order to obtain a good approximation to the steady-state mean.

Possible methods for overcoming these problems are now discussed.

5.1.1. Initial Transient

The network model simulated is empty and idle at the beginning of the simulation: all queues are empty and no traffic enters the network. These initial conditions usually introduce a bias as they are unrepresentative of the desired steady state. Different methods have been proposed to remove this warm-up period: Pawlikowski [19] presents 11 different methods which can be used to estimate the initial transient, while recent studies [17,21] present graphical, heuristic or statistical methods.

To compute the length of the initial transient, a simple heuristic developed by White [15] is used. This method,

called the Marginal Confidence Rule (MCR), selects the truncation point that minimises the width of the confidence interval about the truncate sample mean. If x_1, x_2, \dots, x_N is the set of points measured during the simulated time, the optimal truncation point $i \in [1, N]$ may be given by:

$$i = \arg \min_{0 \leq i < N} \left[\frac{\sigma_i^2(X)}{N-i} \right], \quad (1)$$

where $\sigma_i(X)$ is the sample standard deviation of the sequence x_{i+1}, \dots, x_N .




Figure 3. MCR truncation of the initial bias

Figure 3 shows how the convergence to steady-state value is ensured by the use of MCR to remove the initial, biased values. The sample average calculated without initial bias converges after 10 seconds, while after 40 seconds of simulated time, the effect of the initial biasing on the sample average is still apparent. As may be seen in figure 4, the bias introduced by the initial transient is different for each metric. The truncation point used is taken as the maximum of the five truncation points.

5.1.2. Simulation Run-length

For the simulation methodology adopted, Law and Kelton [16] recommend a simulation time *much larger* than the length of the initial transient, to ensure that sufficient data has been collected.

The truncation point is computed every k points, thus avoiding the computation of the MCR at every point. The simulation is halted when the truncation point i is such as $N > p \times i$. While the steady-state is reached at point i , we need a data set which is sufficiently large so as to capture the steady-state behaviour of the system. Thus we need a data set “*much larger*” than the initial transient [16]. p has been set empirically to 5 in the current version of the framework.




Figure 4. Convergence of the five metrics, with initial transient (the sample averages are normalised to their value at $t = 40s$).

5.2. Multiple Replications

Many network simulation studies construct a model, select performance metrics and conclude with a single simulation run to obtain results. The use of random values from particular distributions as simulator inputs means that the results obtained from a single simulation run are just instantiations of random processes [16].




Figure 5. Evolution of the Fairness metric over simulated time for different replications

Figure 5 illustrates how the point values obtained for an individual performance metric can vary considerably between statistically independent simulation runs. Hence, it is essential to use multiple, independent replications of the

same simulation model and proper statistical analysis of the resultant output data. This lowers the probability of obtaining erroneous results to an acceptable level and improves the credibility of the results obtained.

When using multiple independent replications, the number of replications n can be established beforehand, using the *fixed-sample-size procedure* [16] method. A point value $\bar{X}(n)$ can be used to estimate the mean of a metric and hence a $100(1 - \alpha)\%$ confidence interval for the mean can be found:

$$\delta(n, \alpha) = \bar{X}(n) \pm t_{n-1, 1-\alpha/2} \sqrt{\frac{S^2(n)}{n}}, \quad (2)$$

where $\bar{X}(n)$ is the estimated mean, $S^2(n)$ is the sample variance and t is from the Student's t distribution. The main disadvantage of the fixed-sample-size procedure is that it does not allow a predefined confidence interval half-length to be calculated. If a specific error or precision is desired for the estimated mean $\bar{X}(n)$ then the number of replications n must be calculated dynamically.

The framework implements a sequential procedure to dynamically decide the number of replications required. A desired relative error γ , and a confidence level α , are specified initially. The procedure starts with an initial number of replications $n_0 = 10$, sets $n = n_0$ and computes the estimated mean $\bar{X}(n)$ and the confidence interval half-length $\delta(n, \alpha)$ given by (2). If $\delta(n, \alpha)/|\bar{X}(n)| \leq \gamma$, the relative error is below our desired threshold, so the current $\bar{X}(n)$ is used as the point estimate for the mean metric value μ and no more replications are added. The current approximate $100(1 - \alpha)\%$ confidence interval is given by:

$$I(\alpha, \gamma) = [\bar{X}(n) - \delta(n, \alpha), \bar{X}(n) + \delta(n, \alpha)]. \quad (3)$$

The number of replications n is incremented by 1 until a value of n is reached for which $\delta(n, \alpha)$ is below the predefined threshold.

Figure 6 shows how the metric values evolve over the simulation run length once the bias introduced by the initial warm-up period has been eliminated. The estimated mean for the fairness metric with its associated confidence interval is also plotted on this graph. When comparing this graph with figure 5 it can be seen how the deletion of the initial transient improves the stability of the fairness values.

The sequential procedure for determining the number of replications, when used in conjunction with the technique for eliminating initial transient bias, as described above, is known as the *replication/deletion approach* [16].

6. Study of the fairness of RED

The main purpose of the framework is to study the differences between AQM schemes. A full evaluation of these algorithms is outside the scope of this paper, rather we wish to

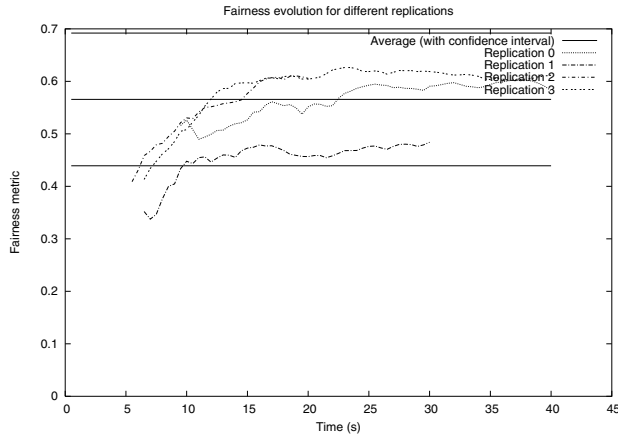


Figure 6. Fairness metric value over time without initial transient bias and estimated average with confidence interval.

demonstrate the use of the framework and its performance evaluation metrics. The RED algorithm [11] and its more recent variant, RED-PD [18] (RED with preferential dropping), are studied in this experiment. Drop-tail queue is also considered as it is the basic scheme implemented in current Internet routers. The algorithms considered are:

- Drop-Tail: standard drop-tail queue. Packets are dropped only when buffer overflow occurs.
- RED: the “current version” of RED, implemented in NS2 by its authors. RED computes an exponential moving weighted average (EMWA) of the queue length, and drops incoming packets randomly: the drop rate is an increasing function of the EMWA value (see [11] for details). The version of RED used here includes the gentle version, which has no discontinuities in the drop probability function, and improved parameter settings.
- RED-PD: RED with preferential dropping [18], adds another layer to RED, detecting and penalising non-responsive flows.

The experiment is based on a simple dumbbell scenario. The parameters used are:

- Bottleneck bandwidth: 10Mb/s
- Bottleneck delay: 10ms
- RTT range: $20\text{ms} - 460\text{ms}$

This experiment examines the effect of the proportion of unresponsive traffic (UDP in this case) on performance. We focus here on the drop and fairness metrics.




Figure 7. Fairness and Drop Metrics

Figure 7 show that RED-PD achieved its design objectives: RED with preferential dropping provides for better fairness by the more aggressive dropping of incoming UDP packets. While RED is more aggressive than regular Drop-Tail, it exhibits comparable fairness.

The high level framework interface allows for the easy setup of such experiments: the graphs produced in this section are taken from the report generated by a ten-line python script.

7. Future Work

The framework may be extended in several ways. The caching of results would yield significant performance enhancements: even the simplest simulation can take several minutes; while detailed experiments, involving higher bandwidth flows, can take several hours to run. Caching previous results would enable the framework to be used for much more complex network simulation scenarios, thus increasing its utility as a research tool.

More complex network topologies should also be supported. These may include the “reverse-dumbbell” topol-

ogy with multiple congested gateways [2] and realistic Internet-like topologies such as power-laws [8]. It is planned to extend the framework so that it can interface with other network simulators such as GloMoSim or OP-NET. This will allow the framework to seamlessly validate results across multiple simulators.

8. Conclusion

A framework for objective comparison and evaluation of AQM schemes is described. It incorporates five carefully chosen metrics characterising overall network performance. The highly configurable tool allows for easy exploration of the parameter space and provides consistent assessment of different AQM schemes making it an invaluable resource for AQM performance evaluation.

Acknowledgements

This work has been supported by Enterprise Ireland Basic Research Grant, SC/2002/293.

References

- [1] M. Allman. A web server's view of the transport layer. *ACM Computer Communication Review*, 30(5):10–20, Oct. 2000.
- [2] K. Anagnostakis, M. Greenwald, and R. Ryger. On the sensitivity of network simulation to topology. In *Proc. MAS-COTS'02.*, pages 117–126. IEEE, Oct. 2002.
- [3] A. Bitorika, M. Robin, and M. Huggard. A framework for evaluating active queue management schemes. Technical Report TCD-CS-2003-34, Trinity College Dublin, Department of Computer Science, July 2003.
- [4] B. Braden et al. Recommendations on queue management and congestion avoidance in the Internet. RFC 2309, Apr. 1998.
- [5] L. Breslau, D. Estrin, K. Fall, S. Floyd, J. Heidemann, A. Helmy, P. Huang, S. McCanne, K. Varadhan, Y. Xu, and H. Yu. Advances in network simulation. *IEEE Computer*, 33(5):59–67, May 2000.
- [6] N. Brownlee and K. Claffy. Understanding Internet traffic streams: Dragonflies and tortoises. *IEEE Communications Magazine*, 40(10):110–117, Oct. 2002.
- [7] G. Ewing, K. Pawlikowski, and D. McNickle. Akaroa2: Exploiting network computing by distributing stochastic simulation. In *Proc. ESM'99*, pages 175–181. International Society for Computer Simulation, June 1999.
- [8] M. Faloutsos, P. Faloutsos, and C. Faloutsos. On power-law relationships of the internet topology. In *SIGCOMM'99*, volume 29 of 4, pages 251–262. ACM, Sept. 1999.
- [9] S. Floyd. Connections with multiple congested gateways in packet-switched networks - part 1: One-way traffic. *ACM Computer Communication Review*, 21(5):30–47, Oct. 1991.
- [10] S. Floyd and K. Fall. Promoting the use of end-to-end congestion control in the Internet. *IEEE/ACM Transactions on Networking*, 7(4):458–472, Aug. 1999.
- [11] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, Aug. 1993.
- [12] S. Floyd and E. Kohler. Internet research needs better models. *ACM Computer Communication Review*, 33(1):29–34, Jan. 2003.
- [13] G. Iannaccone, C. Brandauer, T. Ziegler, C. Diot, S. Fdida, and M. May. Comparison of tail drop and active queue management performance for bulk-data and web-like internet traffic. In *Proc. ISCC*, pages 122–129. IEEE, July 2001.
- [14] R. Jain, D. Chiu, and W. Hawe. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. Technical Report DEC-TR-301, Digital Equipment Corporation, Sept. 1984.
- [15] J. K. Preston White. An effective truncation heuristic for bias reduction in simulation output. *Simulation*, 69(6):323–334, Dec. 1997.
- [16] A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw Hill, third edition, 2000.
- [17] J. R. Linton and C. M. Harmonosky. A comparison of selective initialization bias elimination methods. In *Winter Simulation Conference*, volume 2, pages 1951–1957, Dec. 2002.
- [18] R. Mahajan and S. Floyd. Controlling high bandwidth flows at the congested router. Technical Report TR-01-001, AT&T Center for Internet Research at ICSI (ACIRI), Apr. 2001.
- [19] K. Pawlikowski. Steady-state simulation of queueing processes: A survey of problems and solutions. *ACM Computing Surveys*, 22(2):123–169, June 1990.
- [20] K. Pawlikowski, H.-D. Jeong, and J.-S. Lee. On credibility of simulation studies of telecommunication networks. *IEEE Communications Magazine*, 40(1):132–139, Jan. 2002.
- [21] S. Robinson. A statistical process control approach for estimating the warm-up period. In *Winter Simulation Conference*, volume 1, pages 439–446, Dec. 2002.
- [22] G. van Rossum. Python reference manual. Technical Report CS-R9525, CWI, Netherlands, 1995.
- [23] H. Wu, R. M. Fujimoto, and G. Riley. Experiences parallelizing a commercial network simulator. In *Winter Simulation Conference*, pages 1353–1360, Dec. 2001.
- [24] C. Zhu, O. Yang, J. Aweya, M. Ouellette, and D. Montuno. A comparison of active queue management algorithms using the OPNET modeler. *IEEE Communications Magazine*, 40(6):158–167, June 2002.