

# Implementation of PI<sup>2</sup> Queuing Discipline for Classic TCP Traffic in ns-3

Rohit P. Tahiliani, Hitesh Tewari  
School of Computer Science & Statistics  
Trinity College Dublin  
Dublin 2, Ireland  
tahiliar@tcd.ie,htewari@cs.tcd.ie

## ABSTRACT

This paper presents the implementation and validation of PI<sup>2</sup> Active Queue Management (AQM) algorithm in ns-3. PI<sup>2</sup> provides an alternate design and implementation to Proportional Integral controller Enhanced (PIE) algorithm without affecting the performance benefits it provides in tackling the problem of *bufferbloat*. Bufferbloat is a situation arising due to the presence of large unmanaged buffers in the network. It results in increased latency and therefore, degrades the performance of delay-sensitive traffic. PIE algorithm tries to minimize the queuing delay by auto-tuning its control parameters. However, with PI<sup>2</sup>, this auto-tuning can be replaced by just squaring the packet drop probability. In this paper, we implement a model for PI<sup>2</sup> in ns-3 and verify its correctness by comparing the results obtained from it to those obtained from the PIE model in ns-3. The results indicate that PI<sup>2</sup> offers a simple design and achieves similar responsiveness and stability, and in some cases, better than PIE.

## CCS CONCEPTS

•**Networks** → *Network algorithms; Network performance evaluation*; •**Network algorithms** → *Control path algorithms*; •**Network performance evaluation** → *Network simulations*; •**Computing methodologies** → *Modeling and simulation*; •**Modeling and simulation** → *Simulation evaluation*;

## KEYWORDS

Queuing Disciplines, Bufferbloat, Proportional Integral Controller Enhanced (PIE), PI<sup>2</sup>

### ACM Reference format:

Rohit P. Tahiliani, Hitesh Tewari. 2017. Implementation of PI<sup>2</sup> Queuing Discipline for Classic TCP Traffic in ns-3. In *Proceedings of Workshop on ns-3, Porto, Portugal, June 2017 (WNS-3'17)*, 8 pages.  
DOI: 10.475/123.4

---

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

WNS-3'17, Porto, Portugal

© 2017 Copyright held by the owner/author(s). 978-1-4503-4216-26...\$15.00  
DOI: 10.475/123.4

## 1 INTRODUCTION

The proliferation of delay-sensitive applications on the Internet has given rise to new challenges for queue management. On the other hand, reduced memory costs and the need to accommodate large bursts have encouraged the vendors to increase the router buffer sizes. Although this solves the issue of packet loss and improves TCP throughput, it leads to increased queuing latency. Management of large buffers is indispensable because the unmanaged buffers result in a number of problems such as bufferbloat [1], lock-out [2] and global synchronization [3].

AQM algorithms are being re-investigated with a focus on controlling the queuing latency. Algorithms such as Controlled Delay (CoDel) [4] and Proportional Integral controller Enhanced (PIE) [5] have been designed to minimize queue delay and retain high link utilization. Recently, a new AQM algorithm called PI<sup>2</sup> [6] has been proposed which offers same responsiveness and stability as PIE, but has a simpler design and implementation.

Our contributions in this paper are twofold. Firstly, we propose a new model for PI<sup>2</sup> algorithm in ns-3 [7] and provide details about its design and implementation. Our proposed model is based on the Linux code of the authors of PI<sup>2</sup>.<sup>1</sup> Secondly, we validate the implementation of our PI<sup>2</sup> model in ns-3 by comparing its results to those obtained from PIE model in ns-3, because both are expected to deliver similar performance.

The rest of the paper is organized as follows: Section 2 provides a brief background on PIE, PI<sup>2</sup> and the differences between both. Section 3 details the design and implementation of PI<sup>2</sup> model. Section 4 presents the validation of our PI<sup>2</sup> model in ns-3. Section 5 summarizes and concludes the paper.

## 2 BACKGROUND

### 2.1 PIE

PIE is a recommended AQM algorithm for DOCSIS cable modems [8] and is also being documented as an IETF specification [9]. It uses the Proportional Integral (PI) [10] controller to keep the queuing delay to a specified target value by updating the drop probability at regular intervals. Unlike PI, PIE controls the queuing latency instead of the queue length. Moreover, it auto-tunes the internal parameters based on the

---

<sup>1</sup>[https://github.com/olgabo/dualpi2/blob/master/sch\\_pi2/sch\\_pi2.c](https://github.com/olgabo/dualpi2/blob/master/sch_pi2/sch_pi2.c)

level of congestion. Following are the major components of PIE.

**Random Dropping:** On packet arrival, PIE enqueues or drops the packet based on the drop probability,  $p$ .  $p$  is compared with a random value  $u$  obtained from `UniformRandomVariable` class in ns-3. The packet is enqueued if  $p < u$ , else it is dropped.

**Drop Probability Calculation:** Drop probability provides a measure of the congestion level. PIE calculates the drop probability at every *tupdate* interval and auto-tunes the internal parameters based on level of congestion. The drop probability is calculated as [11]:

$$\Delta p = \alpha * (qdelay - target) + \beta * (qdelay - qdelay\_old)$$

$$p = p + \Delta p$$

where:

- *qdelay*: queuing delay during the current sample.
- *qdelay.old*: queuing delay during the previous sample.
- *target*: desired queuing delay.
- $\alpha$  and  $\beta$ : weights in the drop probability calculation.

**Queuing delay estimate:** PIE uses Little's law [12] for calculating the current queuing latency.

**Burst Tolerance:** PIE allows the short term packet bursts to pass through for a specified interval. This interval duration is determined by user configurable parameter.

## 2.2 PI<sup>2</sup>

Like PIE, PI<sup>2</sup> uses PI controller to keep the queuing delay within a specified target value. However, unlike PIE, it removes the scaling block and makes the drop decision by applying the squared drop probability. Furthermore, it extends PIE to support both Classic (e.g., Reno) and Scalable (e.g., Data Center TCP [13]) congestion controls. In this paper we limit our discussion to implementing PI<sup>2</sup> for Classic TCP traffic in ns-3 because the differentiation between Classic TCP traffic and Scalable TCP traffic is achieved by using Explicit Congestion Notification (ECN) [14] which is not yet completely supported in the main line of ns-3. The components discussed in Section 2.1 apply even to PI<sup>2</sup> with minor changes. These differences are listed in the following subsection.

## 2.3 Differences between PI<sup>2</sup> and PIE

**Drop decision:** PIE drops the packets by comparing the drop probability,  $p$  with the uniform random variable,  $u$ . On the other hand, PI<sup>2</sup> drops the packets by comparing  $p^2$  with  $u$ . Squaring the drop probability helps PI<sup>2</sup> offer a simple design and eliminate the corrective heuristics of PIE without the risking responsiveness and stability [6].

**Burst allowance:** PIE allows the short term packet bursts to pass through. However, PI<sup>2</sup> disables the burst allowance as to avoid an impact on the Data Center TCP fairness [6].

**Other heuristics:** PI<sup>2</sup> chooses to remove a few more heuristics which are a part of Linux Implementation of PIE.

Details and justifications on removing these heuristics have been provided in Section 5 of [6].

## 2.4 Related Work

The implementation and validation of PIE algorithm in ns-3 has been presented in [15]. To verify the correctness of implementation, results obtained from PIE model of ns-3 have been compared to those obtained from PIE model of ns-2 [16]. This model of PIE has been merged in the main development tree of ns-3 since ns-3.26 release.

Another work which is relevant to our paper is the implementation and evaluation of Controlled Delay (CoDel) [17] queuing discipline in ns-3. This was completed as a part of the Google Summer of Code in 2014. CoDel model in ns-3 has been validated by providing unit test cases that compare its results to those obtained from Linux model of CoDel.

## 3 PI<sup>2</sup> MODEL IN NS-3

This section provides insights into the implementation of PI<sup>2</sup> algorithm in ns-3. PI<sup>2</sup> algorithm has been implemented in a new class called `PiSquareQueueDisc` which is inherited from `QueueDisc`. `QueueDisc` is an abstract base class provided by the traffic control layer and has been subclassed to implement queuing disciplines such as Random Early Detection (RED) [3], PIE and CoDel. The following `virtual` methods provided in `QueueDisc` should be implemented in the respective classes of every queuing discipline:


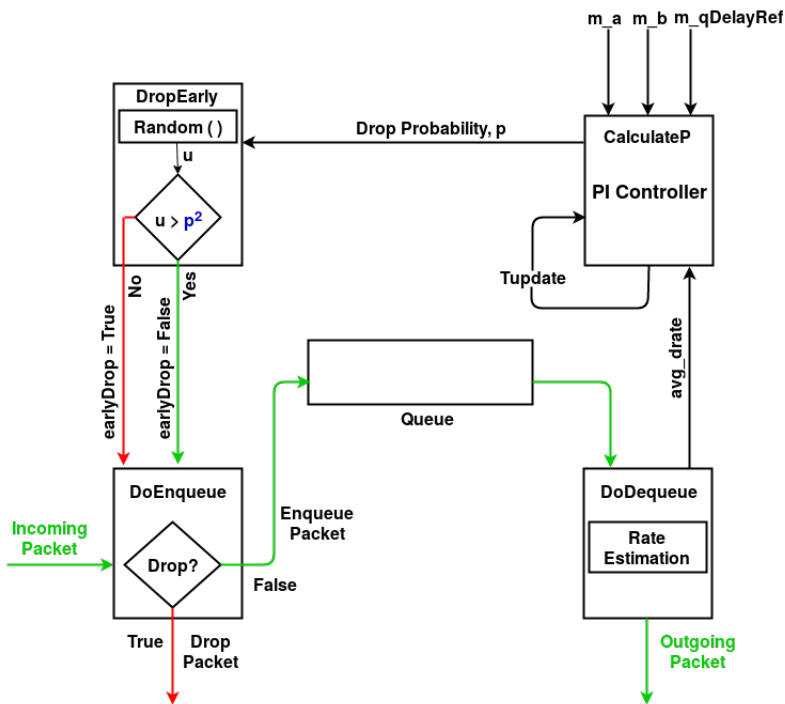
- `bool DoEnqueue (Ptr<QueueDiscItem> item)`: enqueues or drops the incoming packet.
- `Ptr<QueueDiscItem> DoDequeue (void)`: dequeues the packet.
- `Ptr<const QueueDiscItem> DoPeek (void)`  
`const`: peeks into the first item of the queue.
- `bool CheckConfig (void) const`: this method is implemented to check whether the configuration of the queue disc is correct or not.
- `void InitializeParams (void)`: initializes the parameters of the queue disc.

Figure 1 shows the relation between the parent class `QueueDisc` and the derived class `PiSquareQueueDisc`. In addition to the methods mentioned above, `PiSquareQueueDisc` implements the following two methods: `CalculateP` and `DropEarly`. These are specific to the PI<sup>2</sup> algorithm. Figure 2 depicts the interactions among the core components of PI<sup>2</sup>.

On packet arrival, `DoEnqueue` is invoked which thereafter invokes `DropEarly` to check if the incoming packet should be dropped or enqueued. `CalculateP` calculates the drop probability at regular intervals (*tupdate*). `DoDequeue` is invoked on the packet departure.

### 3.1 Dropping Packets Randomly

This functionality is implemented in `DoEnqueue` method in `PiSquareQueueDisc`. Like PIE, PI<sup>2</sup> drops the packets randomly based on the drop probability,  $p$  obtained from `CalculateP`. PI<sup>2</sup> applies the squared drop probability. The squaring is implemented by multiplying  $p$  by itself.

Figure 1: Class Diagram for PI<sup>2</sup> model in ns-3.Figure 2: Interactions among components of PI<sup>2</sup> in ns-3.

DropEarly therefore, makes the drop decision based on the comparison between the squared drop probability and a random value  $u$  obtained from `UniformRandomVariable` class in ns-3. On packet arrival, `DoEnqueue` invokes `DropEarly`.

The packet is enqueued if `DropEarly` returns false, otherwise dropped.

### 3.2 Drop Probability Calculation

This functionality is implemented in `CalculateP` method in `PiSquareQueueDisc` class. PI<sup>2</sup> periodically calculates the drop probability based on the average dequeue rate ( $m\_avgDqRate$ ) and updates the old queuing delay ( $m\_qDelayOld$ ). Table 1 provides a list of parameters used in the calculation of drop probability. Variables used in PI<sup>2</sup> Linux implementation are mapped onto corresponding variables used in ns-3 model.

Table 1: PI<sup>2</sup> parameters to calculate  $p$ .

PI <sup>2</sup> parameter	ns-3 variable
$tupdate$	<code>m_tUpdate</code>
$qdelay$	<code>m_qDelay</code>
$qdelay\_old$	<code>m_qDelayOld</code>
$target$	<code>m_qDelayRef</code>
$alpha$	<code>m_a</code>
$beta$	<code>m_b</code>
$avg\_dq\_rate$	<code>m_avgDqRate</code>

### 3.3 Estimation of Average Departure Rate

This functionality is implemented in `DoDequeue` method in `PiSquareQueueDisc` class. On the packet departure, `DoDequeue` calculates the average departure rate ( $m\_avgDqRate$ ) if the queue is in the measurement cycle. Table 2 provides a list of parameters required to calculate  $m\_avgDqRate$ . Variables used in PI<sup>2</sup> Linux implementation are mapped onto corresponding variables used in ns-3 model.

Table 2: PI<sup>2</sup> parameters to estimate  $avg\_drate$ .

PI <sup>2</sup> parameter	ns-3 variable
$qlen$	<code>m_packets / m_bytesInQueue</code>
$QUEUE\_THRESHOLD$	<code>m_dqThreshold</code>
$dq\_count$	<code>m_dqCount</code>
$dq\_tstamp$	<code>m_dqStart</code>
$dtime$	<code>tmp</code>
$\epsilon$	fixed to 0.5

All the variables are set internally and updated by PI<sup>2</sup>. The only configurable parameter provided by the user is  $m\_qDelayRef$ .

## 4 MODEL VALIDATION

We have designed a test suite with unit tests for verifying the implementation of PI<sup>2</sup> model in ns-3, which is a mandatory step in the process of merging new models into `ns-3-dev`. Our implementation of PI<sup>2</sup> model along with test suite is currently under review.<sup>2</sup>

<sup>2</sup><https://codereview.appspot.com/314290043/>

To further verify the correctness of our implementation, we compare the results obtained from our model of  $PI^2$  to those obtained from the PIE model in ns-3. The simulation scenarios considered for comparison are: (i) varying the traffic and (ii) comparing the CDF of queue delay. These scenarios are in line with the ones used by the authors of  $PI^2$  [6]. However, due to the unavailability of CUBIC [18] and ECN models in ns-3, we have used TCP NewReno [19] without ECN for the evaluation. Our aim is to ensure that our implementation exhibits the key characteristics of the  $PI^2$  algorithm. The performance parameters used for comparison are throughput and queue delay. Table 3 presents the details of simulation setup. Results are presented in the sections that follow.

**Table 3: Simulation setup.**

Parameter	Value
Topology	Dumbbell
Bottleneck RTT	100ms
Bottleneck buffer size	200KB
Bottleneck bandwidth	10Mbps
Bottleneck queue	$PI^2$
Non-bottleneck RTT	10ms
Non-bottleneck bandwidth	10Mbps
Non-bottleneck queue	DropTail
Mean packet size	1000B
TCP	NewReno
<i>target</i>	20ms
<i>tupdate</i>	30ms
<i>alpha</i>	0.125Hz
<i>beta</i>	1.25Hz
<i>dq_threshold</i>	10KB
Application start time	0s
Application stop time	100s
Simulation stop time	100s

### Scenario 1: Light TCP Traffic

In this scenario, a dumbbell topology is used to simulate 5 TCP flows that start at the same time and pass through the same bottleneck link. Other simulation parameters are set as shown in Table 3. Figure 3 shows the variations in queuing delay over time. We can observe the initial peak in the instantaneous queuing delay for both  $PI^2$  and PIE results. This is attributed to the burst traffic generated due to all 5 TCP sources starting at the same time. Moreover, it can be observed that  $PI^2$  to some extent provides better control on the queuing delay. The initial peak in PIE goes to 160+ ms. However,  $PI^2$  keeps it to a maximum of 160ms. Both  $PI^2$  and PIE bring down the queuing delay quickly and maintain it around the reference delay for the rest of the simulation. We can infer that both  $PI^2$  and PIE produce similar results and control the queuing delay to a desired target value.

Figure 4 shows the instantaneous throughput. Initially the throughput degrades for both the results due to packets

being dropped by  $PI^2$  and PIE in an effort to control the queuing delay and maintain it around the desired target delay. Moreover, in contrast to PIE,  $PI^2$  improves the throughput rapidly. This is attributed to  $PI^2$  controlling the queuing delay as mentioned before.

### Scenario 2: Heavy TCP Traffic

This scenario is same as Scenario 1, but configures 50 TCP flows instead of 5 TCP flows. Figure 5 shows the variations in queuing delay over time. Similar to the previous scenario, we can observe that  $PI^2$ , like PIE, quickly brings down the queuing delay and keeps it around the desired target value despite heavy TCP traffic. The results obtained from  $PI^2$  and PIE models are similar. However, under heavy traffic, PIE keeps the queuing delay to 160ms whereas the queuing delay with  $PI^2$  goes little beyond 160ms.

Figure 6 shows the instantaneous throughput. Apart from the initial drop in throughput due to the burst, we can observe that  $PI^2$  controls the queuing delaying without penalizing the link throughput, like PIE.

### Scenario 3: Mix TCP and UDP Traffic

This simulation scenario is to determine whether  $PI^2$  can function normally with unresponsive UDP traffic. We use dumbbell topology and simulate 5 TCP and 2 UDP flows passing through the same bottleneck link. All TCP and UDP flows begin transmission at the same time. UDP sources transmit at a rate of 10 Mbps. Other simulation parameters are same as mentioned in Table 3.

Even in the presence of unresponsive UDP flows, we observe that the results obtained for  $PI^2$  and PIE are similar. Figure 7 shows that  $PI^2$  and PIE control the queuing delay successfully. Moreover, in Figure 8 we can observe that the bottleneck bandwidth is completely utilized with both the algorithms.

### Scenario 4: CDF of queue delay

In this scenario, we compare the CDF of queuing delay obtained for  $PI^2$  and PIE. We conduct two experiments using different traffic loads as done in [6]. First, we use 20 TCP flows with target delay of 5ms and 20ms. Next, we use a mix traffic consisting of 5 TCP and 2 UDP flows with target delay of 5ms and 20ms. Rest of the simulation parameters are same as listed in Table 3. Figure 9 and 10 show the CDF plots comparing the queuing delay of  $PI^2$  and PIE. In line with the observations made by the authors of  $PI^2$ , we observe that  $PI^2$  performs no worse and infact, offers slight improvement over PIE in some cases (see Figure 9 (b)).

## 5 CONCLUSIONS AND FUTURE WORK

In this paper, we describe the implementation of  $PI^2$  algorithm in ns-3 for Classic TCP flows. We present the design of our model and the interactions among different components




Figure 3: Queue delay with light TCP traffic.




Figure 4: Link throughput with light TCP traffic.




Figure 5: Queue delay with heavy TCP traffic.

of  $PI^2$ . Furthermore, we evaluate the effectiveness of our implementation by comparing the results obtained from it to those obtained from the PIE model of ns-3. We note that  $PI^2$

with its simple design can deliver similar performance as PIE. Our implementation of  $PI^2$  has been submitted for review.<sup>3</sup>

<sup>3</sup><https://codereview.appspot.com/314290043/>




Figure 6: Link throughput with heavy TCP traffic.




Figure 7: Queue delay with mix TCP and UDP traffic.





Figure 8: Link throughput with mix TCP and UDP traffic.


and the GitHub repository<sup>4</sup> provides details on reproducing the simulation results. On the availability of ECN in main

distribution of ns-3, we plan to extend PI<sup>2</sup> to work alongside ECN. Moreover, PI<sup>2</sup> model in ns-3 can be further extended


<sup>4</sup><https://github.com/sneaker-rohit/PI2-ns-3>




(a) 20 TCP Flows and target delay = 5ms



(b) 20 TCP Flows and target delay = 20ms

**Figure 9:** CDF of queuing delay with 5 TCP and 2 UDP flows.

(a) 5 TCP + 2 UDP Flows and target delay = 5ms



(b) 5 TCP + 2 UDP Flows and target delay = 20ms

**Figure 10:** CDF of queuing delay with 5 TCP and 2 UDP flows.

to work for scalable congestion control algorithms like Data Center TCP after they are available in the main distribution.

## REFERENCES

- [1] J. Gettys and K. Nichols. Bufferbloat: Dark buffers in the internet. *Communications of the ACM*, 55(1):57–65, 2012.
- [2] M. Hassan and R. Jain. *High Performance TCP/IP Networking*, volume 29. Prentice Hall, 2003.
- [3] S. Floyd and V. Jacobson. Random Early Detection Gateways for Congestion Avoidance. *Networking, IEEE/ACM Transactions on*, 1:397–413, 1993.
- [4] K. Nichols and V. Jacobson. Controlling Queue Delay. *Communications of the ACM*, 55(7):42–50, 2012.
- [5] R. Pan, P. Natarajan, C. Piglion, M. S. Prabhu, V. Subramanian, F. Baker, and B. VerSteeg. PIE: A Lightweight Control Scheme to Address the Bufferbloat Problem. In *High Performance Switching and Routing (HPSR), 2013 IEEE 14th International Conference on*, pages 148–155. IEEE, 2013.
- [6] Koen De Schepper, Olga Bondarenko, Jyh Tsang, and Bob Briscoe. PI2: A Linearized AQM for both Classic and Scalable TCP. In *Proceedings of the 12th International on Conference on emerging Networking EXperiments and Technologies*, pages 105–119. ACM, 2016.
- [7] Network Simulator 3. <https://www.nsnam.org>, 2011.
- [8] CableLabs. Data-Over-Cable Service Interface Specifications DOCSIS<sup>®</sup> 3.1; MAC and upper layer protocols interface specification. specification CM-SP-MULPIv3.1-I01-131029, CableLabs. Oct. 2013.
- [9] F. Baker G. White B. Ver Steeg M. Prabhu C. Piglion R. Pan, P. Natarajan and V. Subramanian. PIE: A Lightweight Control Scheme To Address the Bufferbloat Problem. internet draft draft-ietf-aqm-pie-10, internet engineering task force. Sept. 2016.
- [10] C. V. Hollot, V. Misra, D. Towsley, and W. Gong. On designing improved controllers for AQM routers supporting TCP flows. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1726–1734. IEEE, 2001.
- [11] Koen De Schepper, Olga Bondarenko, Jyh Tsang, and Bob Briscoe. PI2 AQM for Classic and Scalable Congestion Control. Sept. 2016.
- [12] J. D. C. Little and S. C. Graves. Little’s law. In *Building intuition*, pages 81–100. Springer, 2008.
- [13] Mohammad Alizadeh, Albert Greenberg, David A Maltz, Jitendra Padhye, Parveen Patel, Balaji Prabhakar, Sudipta Sengupta, and Murari Sridharan. Data Center TCP (DCTCP). In *ACM SIGCOMM computer communication review*, volume 40, pages 63–74. ACM, 2010.
- [14] K. Ramakrishnan, S. Floyd, D. Black, et al. The addition of Explicit Congestion Notification (ECN) to IP, 2001.
- [15] Smriti Murali, Mohit P Tahiliani, et al. Implementation and Evaluation of Proportional Integral controller Enhanced (PIE) algorithm in ns-3. In *Proceedings of the Workshop on ns-3*, pages 9–16. ACM, 2016.
- [16] Network Simulator 2. <http://www.isi.edu/nsnam/ns>, 1995.
- [17] Truc Anh N Nguyen. Understanding bufferbloat through simulations in ns-3. 2014.

- [18] Sangtae Ha, Injong Rhee, and Lisong Xu. CUBIC: a new TCP-Friendly High-Speed TCP Variant. *ACM SIGOPS Operating Systems Review*, 42(5):64–74, 2008.
- [19] Tom Henderson, Sally Floyd, Andrei Gurtov, and Yoshifumi Nishida. The NewReno Modification to TCP's Fast Recovery Algorithm. Technical report, 2012.