

# Phase Transitions in Interconnection Networks with Finite Buffers

Yelena Rykalova  
Boston University  
[rykalova@bu.edu](mailto:rykalova@bu.edu)

Lev Levitin  
Boston University  
[levitin@bu.edu](mailto:levitin@bu.edu)

## Abstract

This paper presents theoretical models and simulation results for performance of a multiprocessor network modeled as a ring and as a two-dimension wraparound square lattice of nodes with local processors that generate messages with constant rate per time slot. The buffers can hold a limited number of messages. Explicit theoretical results based on first-order (independent queues) and second-order approximation of the queue distributions are obtained for small buffers sizes (1 and 2). For larger buffers the problem appears analytically intractable, and has been studied by simulations. The average queue lengths and average latency are obtained. The results show that the model of independent queues, which is valid for networks with infinite buffers, is still applicable for small generation rates, but breaks down for larger loads, which violates the Jackson theorem.

## 1. Introduction

Over the last several decades the efforts of many researchers were focused on analyzing computer communication networks as networks of queues. Most of the early important results have been obtained for networks with infinite buffers. However, this assumption in many cases does not adequately depict real networks. In practice, buffers are finite, and sometimes quite small. There is a rise of interest in the performance of the networks with small buffers in recent years. Analysis showed [1-4] that increase in buffer size does not improve the performance in wormhole routing significantly. Moreover, latest research [5-11] proved that use of smaller buffers does not decrease the link utilization for Internet routers. In interconnection networks, buffer depth of 5 packets had been shown to offer optimal performance for optical packet-switched clockwork routing [12]. Buffers of size 1, 2, and 50 messages have been considered in [13]. Larger buffers (sizes 5,10, 20,30,40, and 50) have been studied in [14]. The use of smaller buffer sizes gives some advantage in speed and provides possibilities to use SRAM or OPS (Optical Packet Switching). Therefore it is very important to study such networks, in spite of analytical difficulties this may present.

In this paper we consider several models of networks with different size of buffers, starting with size 1 (only one message can be kept in a buffer). In general, theoretical analysis of such networks is a very challenging problem. Though the independent queues assumption is not expected to give an accurate description of the performance of networks with finite buffers, it still makes sense to analyze the networks in terms of the first-order probability distributions (as if the queues were independent). In order to reflect better the correlations between nodes,

the second-order probabilities (i.e. joint probabilities of two neighboring nodes) have been obtained for certain models. Comparison with simulation results show that both the first-order probabilities, and, especially, the second-order distributions yield a reasonably good description of the system behavior until the load reaches the critical region. The simulation results show that when the load approaches the critical value, the system behavior displays typical patterns of long-range dependences between nodes, on the scale much exceeding the “interaction radius” (the distance between the source and destination). The network performance near the critical point is beyond the framework of Jackson theorem [15]; rather it is similar to the critical phenomena in systems described by statistical physics. Both second-order and first-order phase transitions have been observed in networks with finite buffers. Fluctuations with very large amplitudes, slowly changing in time, and instabilities in critical region are characteristic features of the network behavior.

All models considered below have the following common properties:

1. All nodes in the network are both the routers and hosts: every node can generate and receive messages as well as store and send further passing messages.
2. Time is discrete: all nodes simultaneously send, receive and generate messages within every time slot.
3. Each node generates at most one message per output port within every time slot with a certain probability. This probability may depend on the state of the node (the number of messages in the local output port).
4. Destination of a message is selected with equal probabilities among the nodes at exact distance  $l$  from the source.
5. Service time is deterministic (therefore, it does not have a rational Laplace transform). In the absence of queues, a message received or generated during a time slot will appear at the next node in the next time slot.
6. If possible, a message (if there is one) is sent from the output port to the neighboring node at every time slot. This depends on the number of messages in the output port (for the ring topology), or in the router (for the torus topology) in the next (neighboring) node.

In the case of limited buffers, the next state of a buffer depends not only on its previous state (as in the case of infinite buffers), but also on the previous state of its neighbor.

The difference of our work from previous studies [16-18] is in particular, the choice of model features such as discrete time and deterministic service time. While some features of our models are similar to those used by others [19], no model to-date has incorporated all of them. We believe that the combination of properties as described above, makes our model closer to a real supercomputer network, though sometimes hard for theoretical analysis. Choice of the deterministic service time implies that in our model the service time does not have a rational Laplace transform.

Section 2 deals with ring topology networks. The buffer capacities  $m = 1$  and  $m = 2$  is chosen for the theoretical model in order to obtain explicit analytical expressions for the steady-state probabilities (first-order probability distribution) of the states and the average queue length in the “mean field” theory similar to that in statistical physics. Explicit expressions for the second-order distributions have been also obtained. Experimental results based on simulation

have been obtained for the ring length of 500 routers and the distance between source and destination  $l = 5$  hops.

Section 3 provides experimental results for  $16 \times 16$  toroidal square lattice and  $l = 5$  hops.

## 2. Ring Topology

### 2.1. Buffer Limit $m=1$

In the case of buffer size  $m = 1$ , we assume that a node does not accept messages and does not generate messages when the number of messages in a buffer  $n = 1$ . If at the beginning of a time slot a node was in state 0 and then received an incoming message that was not consumed, it does not generate any new message.

#### 2.1.1. First-order probability distribution

Under the assumption of independent queues, the queue state transitions form a simple two state Markov chain and from balance equation  $p_0 p_{01} = p_1 p_{10} = (1 - p_0) p_{10}$  we obtain

$$\begin{aligned} p_0 &= \frac{1 - \lambda(l+1)}{1 - \lambda}, & \bar{n}^{(1)} &= p_1 = \frac{\lambda l}{1 - \lambda}, \\ \lambda_{crit}^{(1)} &= \frac{1}{l+1}, & \tau^{(1)} &= \frac{(1 - \lambda)l}{1 - \lambda(l+1)}. \end{aligned} \quad (1)$$

Here  $\lambda$  is the ‘‘nominal’’ probability that a message is generated by a node within one time interval,  $\bar{n}$  is the average number of messages in a buffer, and  $\tau$  is the average latency. The ‘‘actual’’ generation rate  $\lambda_{act}$  in all our models with finite buffers is different from the ‘‘nominal’’ value  $\lambda$ . Indeed, in the model with  $m=1$ , a new message can be generated in state 0 only if no passing message has arrived, and in state 1 only if the message kept in the output buffer has been sent to the next node. A calculation shows that in the first-order approximation the actual generation rate is

$$\lambda_{act}^{(1)} = p_0 \lambda \left( 1 + \frac{p_1}{l} \right) = \frac{\lambda(1 - \lambda(l+1))}{(1 - \lambda)^2}. \quad (2)$$

#### 2.1.2. Second-order probability distribution

For two neighboring nodes and their combined states the Markov chain has 4 states. After solving the system of balance equations for this Markov chain we obtain

$$\begin{aligned} p_0 &= 1 - \frac{\lambda l^2 (1 - \lambda l)}{\lambda + l(1 - 3\lambda) - \lambda^2 (l+1)(l-1)^2}, \\ \bar{n}^{(2)} &= p_1 = \frac{\lambda l^2 (1 - \lambda l)}{l - \lambda(3l-1) - \lambda^2 (l+1)(l^2-1)}. \end{aligned} \quad (3)$$

$$\lambda_{crit}^{(2)} = \frac{l^2 + 3l - 1 + \sqrt{(l^2 + 3l - 1)^2 - 4l(l^2 + 3l - 1)}}{2(l^2 + l - 1)},$$

$$\lambda_{act}^{(2)} = p_0 \lambda \left( 1 + \frac{p_1}{2l} \right), \quad (4)$$

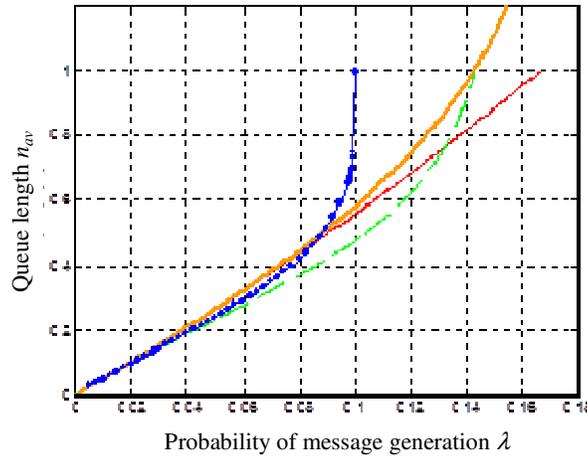
$$\tau^{(2)} = \frac{\bar{n}^{(2)}}{\lambda_{act}^{(2)}} = \frac{2p_1 l}{\lambda(1-p_1)(2l+p_1)}.$$

For comparison, the corresponding theoretical values for a system with infinite buffers are [20, 21]

$$\bar{n} = \bar{n}(\lambda) = \sum_{n=1}^{\infty} nP(n) = \frac{\lambda^2(l-1)}{1-l\lambda} + l\lambda, \quad (5)$$

$$\tau = \frac{\lambda(l-1)}{1-l\lambda} + l. \quad (6)$$

### 2.1.3. Experimental results



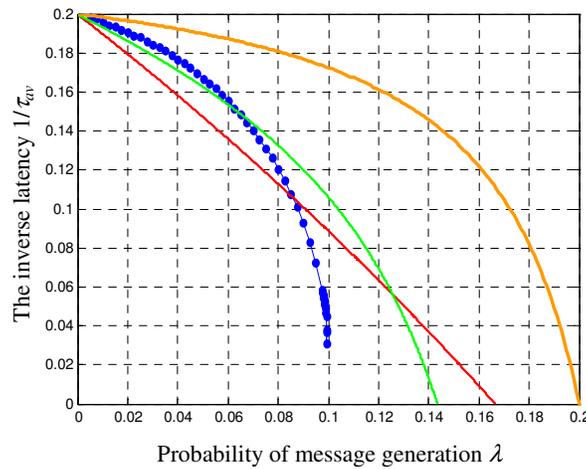
**Figure 1.** The queue length  $\bar{n}$  as a function of the network load  $\lambda$ . Red solid line: the first order approximation  $\bar{n}^{(1)}$  (1); green dashed line: the second order approximation  $\bar{n}^{(2)}$  (3); blue line with bullets: experimental data; orange dotted line: theoretical values for infinite buffers given by (5).

The values of  $\bar{n}$  from simulation together with theoretical approximations  $\bar{n}^{(1)}$  and  $\bar{n}^{(2)}$  are shown in Figure 1. One can see that second order theoretical analysis, compared to first order approximation, provides much better prediction of the network behavior for small network loads (up to  $\lambda \approx 0.7\lambda_{crit}$ ). It also can be seen that the second order analysis results in better approximation of the real critical point: second order  $\lambda_{crit}^{(2)} = 0.1435$  vs. first order  $\lambda_{crit}^{(1)} = 0.1667$ . The interesting observation is that the network performs better than theory of the independent queues predicts. Paradoxically, the first-order approximation gives an upper

bound on  $\bar{n}$  for small network loads when the probability of message generation  $\lambda < 0.8\lambda_{crit}$  while the second-order approximation provides a lower bound on  $\bar{n}$  for all  $\lambda$ .

The empirical value of the latency  $\tau$ , its analytical approximation  $\tau^{(1)}$  and  $\tau^{(2)}$ , and the latency  $\tau_{inf}$  for the network with infinite buffers are plotted in Figure 2.

A remarkable result is that the actual latency  $\tau$  is considerably smaller than  $\tau^{(1)}$  for  $\lambda < 0.6\lambda_{crit}$ , and close to, but slightly smaller than  $\tau^{(2)}$  for  $\lambda < 0.6\lambda_{crit}$ . The reason for this paradoxical behavior is that in the first-order approximation (independent queues) the probability of having a pair of neighboring nodes both in state 1, namely,  $p_{11} = p_1^2$  is substantially larger than in the reality: a pair (11) is unstable at small  $\lambda$  and most probably will become (10) pair of states at the next time interval, since a node in state 1 does not accept messages. Thus, states 0 and 1 negatively correlated for small  $\lambda$  and positively correlated for large  $\lambda$  ( $\lambda > 0.85\lambda_{crit}$ ). We are tempted to compare this phenomenon with the antiferromagnetic-ferromagnetic transition in solid-state physics.



**Figure 2.** The inverse average latency  $\tau$  as a function of the network load  $\lambda$ . Red solid line: the first order approximation  $\tau^{(1)}$  (1); green dashed line: the second order approximation  $\tau^{(2)}$  (4); blue line with bullets: experimental data; orange dotted line: theoretical values for infinite buffers given by (6).

The network behavior near the critical point displays a pattern of the second order phase transition. The average latency  $\tau$  follows a power law:

$$\tau(\lambda) = \alpha(\lambda_{crit} - \lambda)^{-\beta}, \quad (7)$$

where  $\lambda_{crit}$  is the critical network load and  $\beta$  is the critical exponent.

The data from our simulation with distance between the source and destination  $l = 5$  hops yield  $\lambda_{crit} = 0.09981$ ,  $\alpha = 4.891 \pm 0.056$ , and critical exponent  $\beta = 0.2070 \pm 0.0094$ .

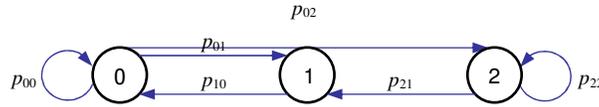
## 2.2. Buffer limit $m = 2$ , Model 1

In the case of the network with buffer capacity  $m = 2$  we consider two models:

1. A node does not accept messages and does not generate messages when  $n = 1$  or  $n = 2$ . This model is analyzed in the present section (Section 2.2).
2. A node does not accept messages when  $n = 1$  or  $n = 2$ , and does not generate messages when  $n = 2$ . This model is analyzed in Section 2.3.

### 2.2.1. First-order probability distribution

Under the assumptions described above, only a router with no messages can accept incoming message, or generate a message with probability  $\lambda$ . State diagram with transition probabilities for this model is shown in Figure 3.



**Figure 3.** The Markov chain for buffer limit  $m = 2$ , model 1.

From system of the flow balance equations

$$\begin{aligned} p_0(p_{01} + p_{02}) &= p_1 p_{10} \\ p_0 p_{02} &= p_2 p_{21} \\ p_0 + p_1 + p_2 &= 1 \end{aligned} \tag{8}$$

we have

$$\begin{aligned} p_0 &= 1 - \lambda l, & p_1 &= \lambda l - \lambda^2 (l-1) \bar{n}^{(1)} = \lambda l + \lambda^2 (l-1), & p_2 &= \lambda^2 (l-1), \\ \lambda_{act}^{(1)} &= p_0 \lambda = \lambda - \lambda^2 l, & \tau^{(1)} &= \frac{\bar{n}^{(1)}}{\lambda_{act}^{(1)}} = \frac{l + \lambda(l-1)}{1 - \lambda l}, & \lambda_{crit}^{(1)} &= \frac{1}{l}. \end{aligned} \tag{9}$$

By comparing the expression for  $\lambda_{act}^{(1)}$  and  $\tau^{(1)}$  for this model and for the model of network with buffer limit  $m = 1$ , one can see that the first-order approximation gives a smaller value of  $\lambda_{act}^{(1)}$ , and, therefore, the larger latency  $\tau^{(1)}$  for  $m = 2$  when  $\lambda < 0.4 \lambda_{crit}^{(1)}$ , but the opposite is true for larger values of  $\lambda$ . The reason for that is that  $\lambda_{crit}^{(1)}(m = 2) > \lambda_{crit}^{(1)}(m = 1)$ .

### 2.2.2. Second-order probability distribution

The second-order approximation becomes more challenging for this model comparing to the second-order approximation described in Section 2.1.2. Now the state transition diagram has nine states, which are combined states of two neighboring nodes. Note that the states are not symmetric, since the messages are always sent from left to right. The global balance equations for such Markov chain as well as probabilities of combined states and marginal probabilities  $(p_0, p_1, p_2)$  are too monstrous to present here. The equation for  $\lambda$  to be solved

to find  $\lambda^{(2)}_{crit}$  is of degree 12 (!). Solving it numerically for  $l=5$ , we obtain the smallest real nonnegative root:  $\lambda^{(2)}_{crit} = 0.13932314$ . The expression for the actual load  $\lambda^{(1)}_{act}$  has the same form as in the first-order approximation but expression for  $p_0$  is different. Correspondingly, the average queue length is  $\bar{n}^{(2)}$ :

$$\bar{n}^{(2)} = p_1 + 2p_2. \quad (10)$$

Other second-order characteristics are given below (11)

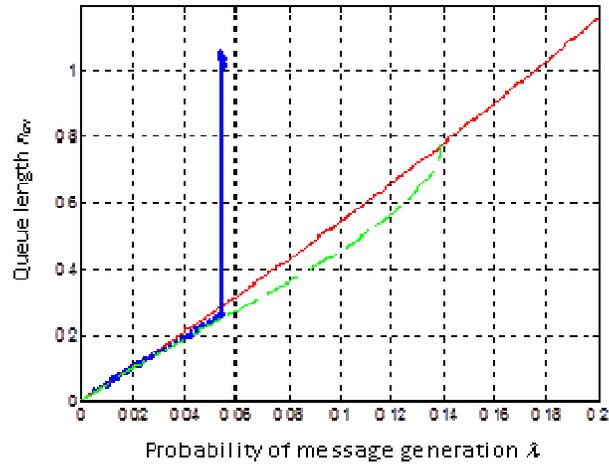
$$p_0 = \frac{(1-\lambda)^2 \pm \sqrt{(1-\lambda)^4 - 4(1-\lambda)\lambda(l-1+\lambda)}}{2(1-\lambda)},$$

$$p_1 = \frac{p_0(1-p_0)}{l\lambda} - p_0, \quad p_2 = 1 - \frac{p_0(1-p_0)}{l\lambda}, \quad (11)$$

$$\lambda^{(2)}_{act} = (p_0 + p_1)\lambda, \quad \bar{n}^{(2)} = p_0 + 2p_1, \quad \tau^{(2)} = \frac{\bar{n}^{(2)}}{\lambda^{(2)}_{act}} = \frac{p_0 + 2p_1}{(p_0 + p_1)\lambda}.$$

### 2.2.3. Numerical results

The approximations for the average queue length  $\bar{n}^{(1)}$  and  $\bar{n}^{(2)}$ , as well as experimental values are shown in Figure 4.



**Figure 4.** The queue length  $\bar{n}$  as a function of the network load  $\lambda$ . Red line: the first order approximation  $\bar{n}^{(1)}$  (9); green line: the second order approximation  $\bar{n}^{(2)}$  (10); blue line with bullets: experimental data.

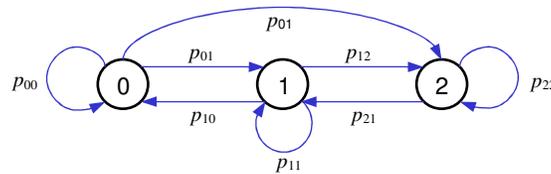
As in the previous model (Section 2.1.3), the second order approximation gives an excellent prediction for the network behavior up to the moment when network becomes saturated. Network goes to the saturation state not only much earlier than the theoretical approximations predict, but it does so abruptly, sharply departing from the second-order approximation curve. Unlike the previous model, there is no continuous transition to the

saturation. The picture we have here closely resembles a liquid-gas first-order phase transition.

Close to the critical point, the network behavior is very unstable: in one experiment, network goes to saturation; next time, even when the network load is larger then before, it does not show any sign of reaching the saturation state (Figure 3). Therefore, in this case we have the first order phase transition. We can, however, roughly estimate the critical load as  $\lambda_{crit} = 0.0546$ .

### 2.3. Buffer Limit $m = 2$ , Model 2

In this section we consider the second model for the buffer capacity  $m=2$ . Here a node can generate a message when there is one message in the buffer (state 1), but it does not accept messages when  $n = 1$  or  $n = 2$ . The Markov chain for this model is shown in Figure 5.



**Figure 5.** The Markov chain for buffer limit  $m = 2$ , model 2.

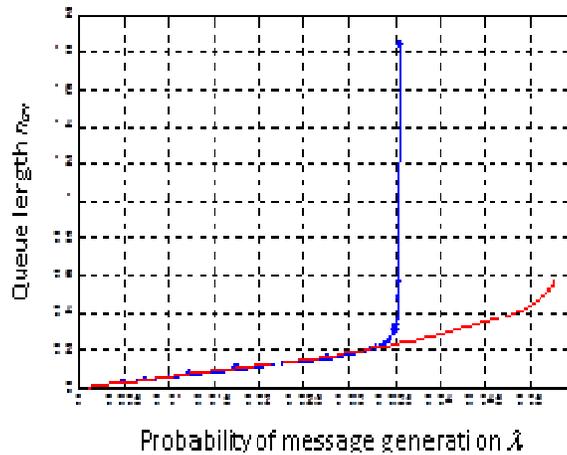
The model was described in details in [13]. The second-order probability distribution was not possible to obtain. Using the first-order analysis, we received the next expressions for actual load, average queue length and latency:

$$\begin{aligned}\lambda_{act}^{(1)} &= (p_0 + p_1)\lambda, \\ \bar{n}^{(1)} &= p_1 + 2p_2, \\ \tau^{(1)} &= \frac{\bar{n}^{(1)}}{\lambda_{act}^{(1)}} = \frac{p_1 + 2p_2}{(p_0 + p_1)\lambda}.\end{aligned}\tag{12}$$

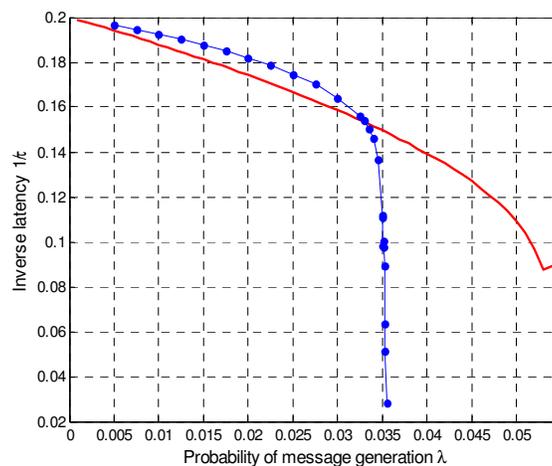
It can be seen that numerical data follows the theoretical prediction very closely until network load reaches value  $\lambda \approx 0.0345$ . Up to this point, the data from experiments show even a slightly better network performance (smaller queue length and latency) when the first-order approximation predicts. However, numerical experiments show a sharp increase in the average queue length (Figure 6) just before the network reaches saturation and long before we can expect the critical transition according to the mean field theory.

Large fluctuations in the number of messages in the network (Figure 8) and in the number of saturated buffers were observed using this model. The transition of the network to the saturation phase shows characteristics of the second-order phase transition, despite the fact that this transition is more abrupt then those described in Section 2.1. This is dramatically different from model 1 described in the previous section, which has a first-order phase transition.

By comparing model 2 for the buffer limit  $m = 2$  described here with model from Section 2.1 ( $m = 1$ ), one can see that increasing the buffer capacity does not necessarily lead to defer the saturation point. On the contrary, the former model reaches the saturation state under network load almost three times smaller than the latter. The explanation of this fact is that in model 2 a node can generate messages in both state 0 and 1, rather than in state 0 only, as in Section 2.1, which results in a larger effective load.



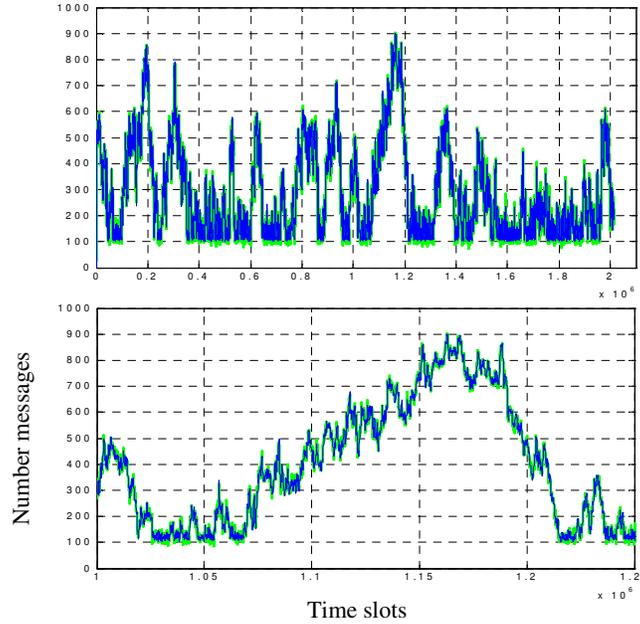
**Figure 6.** The queue length  $\bar{n}$  as a function of the network load  $\lambda$ . Red line: the first order approximation  $\bar{n}^{(1)}$  (12); blue line with bullets: experimental data.



**Figure 7.** The inverse average latency  $\tau$  as a function of the network load  $\lambda$ . Red line: the first order approximation  $\tau^{(1)}$  (12); blue line with bullets: experimental data.

The shape of the latency curve (Figure 7), the fluctuation patterns (Figure 8), and the formation of domains of different phases speak in favor of a continuous (second-order) phase-transition. The empirical data on latency agree very well with a power law (7).

Calculations for  $l=5$  yield  $\lambda_{crit} = 0.03536$ ,  $\alpha = 1.775 \pm 0.028$ , and a critical exponent  $\beta = 0.2026 \pm 0.0095$



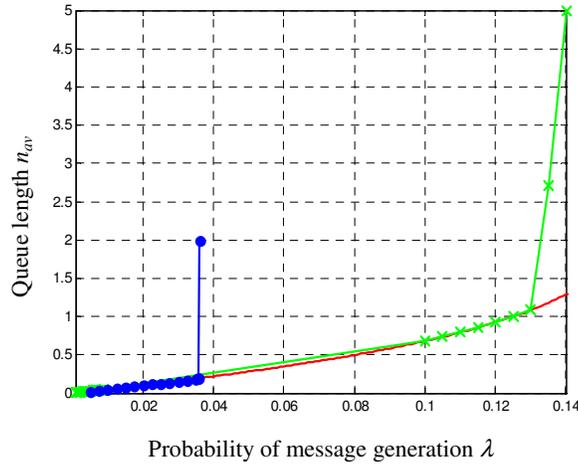
**Figure 8.** Number of messages in the network taken at every 100 clock cycles,  $\lambda = 0.03535$ ,  $l = 5$  hops. Blue line: number of messages in the network averaged over 100 time slots; green line: number of messages in the network taken every 100 time slots.

### 3. Two-Dimensional Torus Topology

We considered a  $16 \times 16$  toroidal square lattice network. We assume that messages are always routed along the shortest path to the message destination, but we do not use static routing (there is no fixed path assigned to the message). When a router receives a passing message and there is a choice between two possible buffers to keep the message, one of the buffers is chosen at random. Therefore, in the case of two-dimensional torus network topology and finite buffers, we elected to make the decision, whether to send a message to a neighboring node, based not on the queue in any single buffer at the next node along the message path, but rather on the total number of messages in the neighboring router. The limits on total number of messages held in the router have been set equal to 8 and 20. We assumed that router stops accepting incoming messages once it accumulates  $(m - 4)$  messages. Locally generated messages are generated individually and independently for all four directions, and once router accumulates  $m$  messages it stops generating messages. Results are shown in Figure 9 and 10. Red line presents the theoretical prediction for networks with infinite buffers and 2-dim torus topology [17] given by

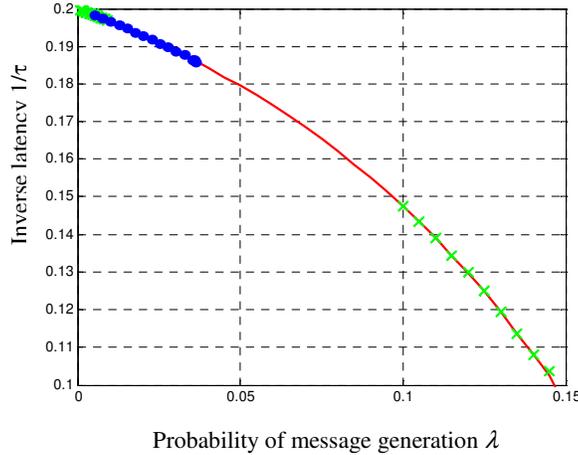
$$\bar{n} = \frac{\lambda^2 (l-1)(11+5l)}{16(1-\lambda)} + l\lambda, \quad (13)$$

$$\tau = \frac{\lambda(l-1)(11+5l)}{16(1-l\lambda)} + l. \quad (14)$$



**Figure 9.** The queue length  $\bar{n}$  as a function of the network load  $\lambda$ . Red line: theoretical values for infinite buffers given by (13); blue line with bullets: experimental data for router capacity equal 8; green line with crosses: experimental data for router capacity equal 20.

Surprisingly, until the network load comes close to the saturation value, the network behaves as if there was no buffer limit. We can see an excellent agreement between theoretical prediction and simulation results, as well as no difference in network performance for networks with different total router capacities up to their critical loads.



**Figure 10.** The latency  $\tau$  as a function of the network load  $\lambda$ . Red line: theoretical values for infinite buffers given by (14); blue line with bullets: experimental data for router capacity equal 8; green line with crosses: experimental data for router capacity equal 20.

The simulations show a sharp transition from the steady state regime to the saturated state for both router capacities. Therefore, in the case of torus topology and finite buffers, the network behaves as if it undergoes a first-order phase transition.

## 4. Conclusions

In this paper we analyzed the networks with finite buffers and proposed a few theoretical models. For small buffer size  $m = 1$  and  $m = 2$ , both first-order probability distributions (independent queue approximation) and second-order distributions (joint probabilities of pairs of neighboring nodes) were obtained. The results show that theoretical predictions given by the second-order approximation provide a better description of the network behavior compared to the first-order (mean field) approximation. But theoretical analysis, even for small value of  $m$ , proved to be very challenging, especially for the second-order approximation. Simulations data demonstrate the emergence of long-range dependences between nodes, far exceeding the “interaction radius”  $l$  when the network reaches the critical region. The most interesting observation is that the network shows both first-order and second-order phase transitions. The type of the phase transition depends not only on the buffer limit, but also on the assumptions for the message generation. In particular, with the buffer limit  $m = 2$ , we observed a second-order phase transition in model 1, while model 2 has a first-order transition. The second-order phase transitions are characterized by non-trivial values of critical exponent. An interesting feature of the first-order phase transitions is that the network behavior follows the second-order approximation (for the ring topology) and “mean field” approximation for networks with infinite buffers (for the torus topology) closely until the load reaches the critical value. When it happens, network goes to the saturation state abruptly. It is very difficult to predict when it can happen even from the experimental data: there are no significant increases in the average latency or queue length right before the critical point. Only the observation of the fluctuations of the average number of the messages in the network during the simulation can serve as an indication that the network is close to saturation. These fluctuations increase both in the amplitude and in the “wavelength” just before the critical point.

An interesting task for the future studies would be to find out how a system with finite buffers “approximate” system with infinite buffers (which has a second-order phase transition) with the increase of buffer size.

## Acknowledgement

The authors are thankful to the reviewers of the paper whose remarks have led to a number of improvements.

## References

- [1] Duato, J., Yalamanchili, S., and Ni, L. M. (2003). *Interconnection Networks: An Engineering Approach*. Morgan Kaufmann, ISBN 1558608524, 9781558608528.
- [2] Kodi, A. K., Sarathy, A., and Louri, A. (2008). Adaptive Channel Buffers in On-Chip Interconnection Networks— A Power and Performance Analysis. *IEEE Trans. Comput.* 57, 9 (Sep. 2008), 1169-1181. DOI=<http://dx.doi.org/10.1109/TC.2008.77>

- [3] Enachescu, M., Ganjali, Y., Goel, A., McKeown, N., and Roughgarden, T. (2006). Routers with Very Small Buffers. *In Proc. 25th IEEE International Conference on Computer Communications*, 1-11.
- [4] Yu Gu, Towsley, D., Hollos, C., Zhang, H. (2007). Congestion Control for Small Buffer High Speed Networks. *In Proc. 26th IEEE International Conference on Computer Communications, 6-12 May 2007*, 1037-1045.
- [5] Appenzeller, G., Keslassy, I., and McKeown, N. (2004). Sizing router buffers. *Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications (SIGCOMM '04)*. ACM, New York, NY, USA, 281-292. DOI=<http://dx.doi.org/10.1145/1015467.1015499>
- [6] Enachescu, M., Ganjali, Y., Goel, A., McKeown, N., and Roughgarden, T. (2005). Part III: routers with very small buffers. *SIGCOMM Comput. Commun. Rev.* 35, 3 (July 2005), 83-90. DOI=<http://dx.doi.org/10.1145/1070873.1070886>
- [7] Dhamdhere, A. and Dovrolis, C. (2006). Open issues in router buffer sizing. *SIGCOMM Comput. Commun. Rev.* 36, 1 (January 2006), 87-92. DOI=<http://dx.doi.org/10.1145/1111322.1111342>
- [8] Enachescu, M., Ganjali, Y., Goel, A., McKeown, N., and Roughgarden, T. (2006). Routers with Very Small Buffers. *Proceedings IEEE INFOCOM 2006. 25TH IEEE International Conference on Computer Communications*, Barcelona, Spain, 2006, 1-11. doi: 10.1109/INFOCOM.2006.240
- [9] Beheshti, N., Ganjali, Y., Ghobadi, M., McKeown, N., and Salmon, G. (2008). Experimental study of router buffer sizing. *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement (IMC '08)*. ACM, New York, NY, USA, 197-210. DOI=<http://dx.doi.org/10.1145/1452520.1452545>
- [10] Vishwanath, A., Sivaraman, V., and Rouskas, G.N. (2011). Anomalous loss performance for mixed real-time and TCP traffic in routers with very small buffers. *IEEE/ACM Trans. Netw.* 19, 4 (August 2011), 933-946. DOI=<http://dx.doi.org/10.1109/TNET.2010.2091721>
- [11] Hohlfeld, O., Pujol, E., Ciucu, F., Feldmann, A., and Barford, P. (2014). A QoS Perspective on Sizing Network Buffers. *In Proceedings of the 2014 Conference on Internet Measurement Conference (IMC '14)*. ACM, New York, NY, USA, 333-346. DOI=<http://dx.doi.org/10.1145/2663716.2663730>
- [12] Bravi, E., and Cotter, D. (2007). Optical packet-switched interconnect based on wavelength-division- multiplexed clockwork routing. *Journal of Optical Networking*, vol. 6, Issue 7, 840-853.
- [13] Rykalova, Y., Levitin, L. B., and Brower, R. (2008). Interconnection Networks with Heterogeneous Activity or Finite Buffers: Beyond Jackson's Theorem. *In Proc. of the 11th Communications and Networking Simulation Symposium, CNS08*, Ottawa, Canada, April 14 -17, 9-14.
- [14] Levitin, L.B., Rykalova, Y. (2015). Computer Networks with Finite Buffers: Beyond Jackson's Theorem. In Al-Sakib Pathan, Muhammad Monowar, and Shafiullah Khan (Eds.) *Simulation Technologies in Networking and Communications: Selecting the Best Tool for the Test*. CRC Press 2015, 31 - 68. Print ISBN: 978-1-4822-2549-5 eBook ISBN: 978-1-4822-2550-1
- [15] Jackson, P. (1963). Job shop like queueing systems. *Management Sci.*, vol. 10, no. 1, 131- 142.

- [16] Asadathorn, N., Chao, X.. (1999). A decomposition approximation for assembly-disassembly queueing networks with finite buffer and blocking. *Annals of Operations research, Volume 87, Number 1*, 1999 , pp. 247-261.
- [17] Barrenetxea, G., Beferull-Lozano, B., Vetterli, M. (2005). Efficient routing with small buffers in dense networks. *In Proc. of the 4th international symposium on Information processing in sensor networks. Los Angeles, California*, 277- 284.
- [18] Mandjes, M. and Kim, J. H. (2001). Large Deviations for Small Buffers: An Insensitivity Result. *Queueing Syst. Theory Appl.* 37, 4 (Apr. 2001), 349-362.
- [19] Chang, S. H. and Choi, D. W. (2005). Performance analysis of a finite-buffer discrete-time queue with bulk arrival, bulk service and vacations. *Comput. Oper. Res.* 32, 9 (Sep. 2005), pp. 2213-2234. DOI= <http://dx.doi.org/10.1016/j.cor.2004.01.004>
- [20] Levitin, L.B., Rykalova, Y. (2015). Analysis and Simulation of Computer Networks with Unlimited Buffers. In Al-Sakib Pathan, Muhammad Monowar, and Shafiullah Khan (Eds.) *Simulation Technologies in Networking and Communications: Selecting the Best Tool for the Test*. CRC Press 2015, 3 -30.  
Print ISBN: 978-1-4822-2549-5 eBook ISBN: 978-1-4822-2550-1
- [21] Levitin, L. B. and Rykalova, Y. (2016). Latency and Phase transitions in Interconnections Networks with Unlimited Buffers. *In Proc. of the 5th IAJC/ISAM Joint International Conference, Orlando, Fl, November 6-8, 2016*. ISBN 978-1-60643-379-9

## Biographies

LEV LEVITIN is currently a Distinguished Professor of Engineering Science with the Department of Electrical and Computer Engineering at Boston University. He has published over 190 papers, presentations, and patents. His research areas include information theory; quantum communication systems; physics of computation; quantum computing; quantum theory of measurements, mathematical linguistics; theory of complex systems; coding theory; theory of computer hardware testing, reliable computer networks, and bioinformatics. He is a Life Fellow of IEEE, a member of the International Academy of Informatics and other professional societies. Prof. Levitin may be reached at [levitin@bu.edu](mailto:levitin@bu.edu).

YELENA RYKALOVA is currently a Visiting Researcher with the Department of Electrical and Computer Engineering at Boston University Her research interests are in computer networks, in particular, in application of concepts and models of statistical physics to the analysis of network performance. She is a member of IEEE and The Society for Modeling and Simulation International (SCS). Since 2008 she has been active in organization and preparation for the Spring Simulation Multiconference (SpringSim) as Technical Committee Member, Reviewer, and Publicity and Session Chair. Dr. Rykalova may be reached at [rykalova@bu.edu](mailto:rykalova@bu.edu).