

Emergent behaviour in a model data network

T. Bruusgaard and J.H.B. Deane

Department of Electronic and Electrical Engineering,
University of Surrey,
Guildford GU2 5XH,
United Kingdom

Abstract

A simple model data network is investigated and both periodic and aperiodic behaviour are observed, the latter displaying evidence of chaos. A performance measure, the throughput, is calculated and shown to be considerably larger in the aperiodic case.

1. Introduction

We investigate a deterministic model of data transfer which has been stripped down to the essentials. It consists of a square array of connected routing cells, randomly placed hosts, and data packets that move according to given rules. We build on the work of [1] and [2], except that there is now no token in the model—packets are transferred after an appropriate delay, not when a token is at a given cell. The dynamics of the system affect its throughput, which is noticeably higher when apparently chaotic behaviour is manifested.

2. System definition

Consider a two-dimensional $N \times N$ array of routing cells. Each cell is connected to its eight nearest neighbours. Each packet in the network is assigned a source and destination cell, placed randomly but on opposite sides of the network. A packet travels repeatedly between its source and destination cells. The model is governed by the following rules: (1) A cell can only store one packet at a time. (2) A packet can be transferred only to a neighbouring cell. Ideally, it moves diagonally until in line with its destination, thereafter moving horizontally or vertically, as appropriate. Hence, the ideal path is a parallelogram. (3) While a packet is being transferred between two cells, both are occupied. This applies for a time τ_t , the transfer time. (4) After transfer, a packet waits a time τ_d , the delay time, before it can be transferred again. If all neighbouring cells are occupied, the packet remains at its current cell for a further τ_d

before next attempting transfer. (5) As soon as a packet reaches its destination it is returned to its source. Packet transfer between source and destination continues indefinitely. In the simulations presented here, $\tau_t = 0.001$ and $\tau_d = 0.01$ time units.

Each packet attempts to transfer to an adjoining cell a time τ_d after reaching its current cell. The choice of cell is made according to whether it is unoccupied, and how much in the direction of the destination it is—for details see [2].

The parameters (*e.g.* initial packet, source and destination positions) are set up randomly at the start of the simulation. To make unlikely two transfers happening at exactly the same time, a 64-bit (double precision) random number is chosen for the initial time of next transfer of each packet, giving a probability of collision of about 10^{-16} .

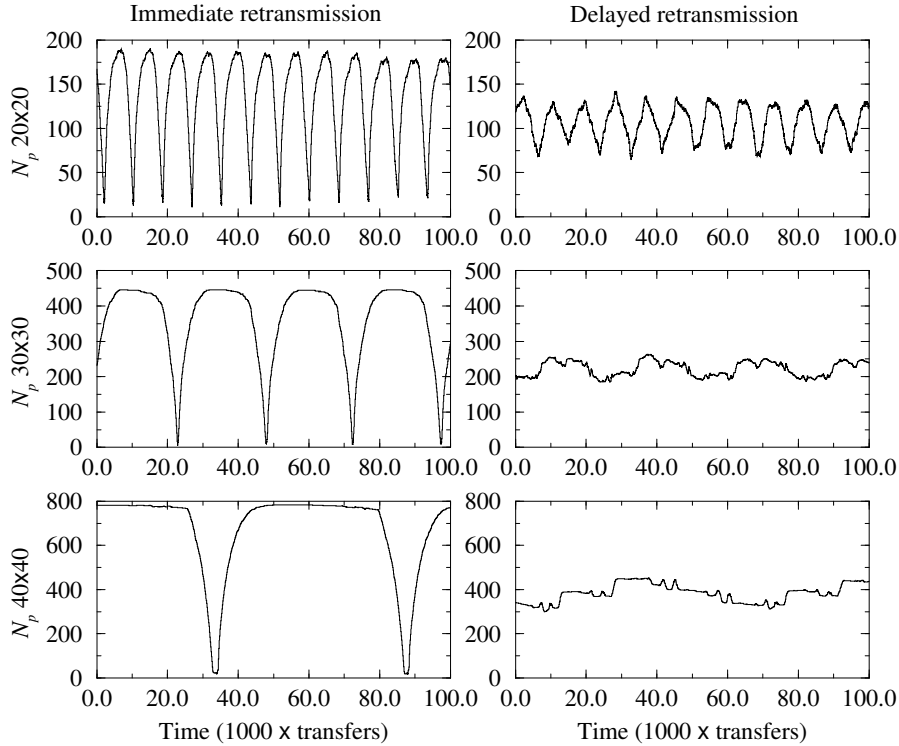


Figure 1: Comparison of the number of packets N_p as a function of time, for immediate (left) and delayed (right) retransmission, for three different sizes of network.

The number of packets in the network, N_p , varies with time. Each packet has a time-to-live, set, as it leaves the source, to one roundtrip time, the minimum time for a source-destination-source trip, which is $\tau_{rt} = 2N(\tau_d + \tau_t)$. The time-to-live decreases as the packet moves through the network. If the packet fails to complete the round trip within τ_{rt} the packet is deleted. Sources are equipped with a retransmission timer, which is set to $2\tau_{rt}$. If the packet does not return within $2\tau_{rt}$, the source retransmits the packet—provided that the routing cell at the source is not occupied. If this cell is occupied, the source will either (a) retransmit as soon as the cell becomes vacant (immediate retransmission), or (b) wait $\tau_{rt}/2$ before trying to

retransmit (delayed retransmission). The dynamics of N_p , and the throughput, using the two different retransmission rules, are now compared.

3. Results

The variation of N_p as a function of time, for immediate and delayed retransmission, is shown in figure 1. The time units used are $1000 \times$ packet transfer events. There are $N^2/2$ pairs of data sources. Note that immediate retransmission leads to nearly periodic, pulsed behaviour and delayed retransmission results in smoother, aperiodic dynamics.

We define the throughput as $(N \times \text{the number of packets reaching source and destination}) / (\text{total number of packet transfers})$. The maximum obtainable throughput, unity, would only occur if every packet followed its ideal parallelogram path. Since packets can interact, this ideal behaviour is observed only for very small N_p . The throughputs for the two retransmission rules in networks of different sizes are given below. Note that delayed outperforms immediate retransmission, substantially so in larger networks. An explanation is evident from figure 1, in which emergent behaviour for immediate retransmission consists of roughly two states: congestion, in which packets are barely able to move, alternating with underload, when packets move freely but there are very few of them. By contrast, for delayed retransmission, N_p is roughly constant with mean $\approx N^2/4$.

N	Throughput (immediate retransmission)	Throughput (delayed retransmission)
20	0.23 ± 0.01	0.24 ± 0.01
30	0.19 ± 0.01	0.24 ± 0.01
40	0.12 ± 0.01	0.24 ± 0.01

The network we have described is a nonlinear, high-dimensional, discrete dynamical system, which leads us to enquire if the aperiodic behaviour is in any sense chaotic. We make the smallest possible perturbation to the state of the system, by once moving a single packet to a neighbouring unoccupied cell. The waveforms both with and without this small change are compared in figure 2, which shows the two waveforms; their difference; and the smoothed difference, obtained by a 20,000 point running average. A least squares fitted exponential curve to the latter is also shown, giving evidence for exponential divergence of initially nearby states—a hallmark of chaos.

4. Conclusions

We have implemented a model of a data network in which packets are transferred to and from between pairs of hosts via an array of interconnected routing cells. Packets which are undelivered within a time limit are deleted. Two different retransmission rules have been

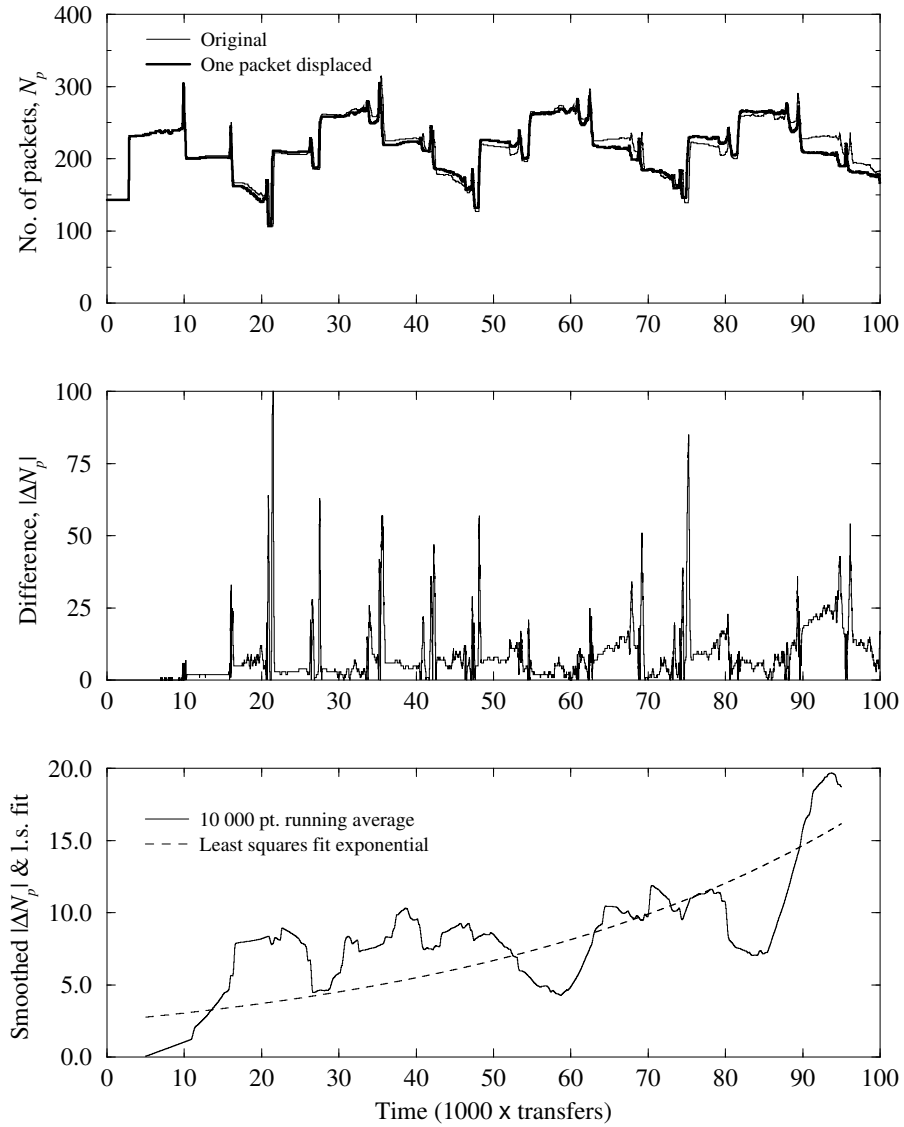


Figure 2: Evidence for exponential divergence (chaos-like behaviour) in our model. Top: N_p as a function of time for two minimally different initial conditions at time = 0. Middle: The difference between the above two waveforms. Bottom: A 20,000 point running average of the difference waveform and a least squares fitted exponential curve.

considered. Counterintuitively, the throughput is considerably higher for large networks when delayed retransmission is used. Furthermore, in this case the observed dynamics of the number of packets in the network is reminiscent of chaos: small changes in initial conditions lead to exponential divergence. For immediate retransmission this effect was not observed.

References

- [1] J.H.B. Deane, D.J. Jefferies and C. Smythe, *The dynamics of deterministic data networks*, in **Complex systems: Mechanisms of adaptation** (COMPLEX'94 Conference Proceedings), Rockhampton, Queensland, Australia, pp 345–352 (September 1994)
- [2] J.H.B. Deane, D.J. Jefferies and C. Smythe *Self-similarity in a deterministic model of data transfer*, **International Journal of Electronics**, vol. 80, no. 5, pp 677–691 (1996)