

Thermal and repulsive traffic flow

I. Campos and A. Tarancón*

Departamento de Física Teórica, Universidad de Zaragoza, 50009 Zaragoza, Spain

F. Clérot†

France Télécom, Centre National d'Etudes des Télécommunications, 2 avenue Pierre Marzin, 22307 Lannion Cedex, France

L. A. Fernández‡

Departamento de Física Teórica, Universidad Complutense de Madrid, 28040 Madrid, Spain

(Received 1 June 1995)

We study a particle flow model that may be used to gain insight into various real traffic problems. The model is implemented using a discrete lattice, in which particles move towards their destination, fluctuating about the minimal distance path. A repulsive interaction between particles is introduced so as to avoid the appearance of a traffic jam. We have studied the parameter space finding regions of fluid traffic, and saturated ones, both regions being separated by abrupt changes. The improvement of the system performance is also explored by introducing a nonconstant potential acting on the particles. Finally, we deal with the behavior of the system when temporary failures in transmission occur.

PACS number(s): 05.40.+j

I. INTRODUCTION

The formulation of adequate models for describing message passing through a computer network [1], traffic flow in a city [2], adsorption of molecules in a crystal [3], fast ionic conductor [4], or the mechanism of high temperature superconductors [5] has been getting increasing attention over the last few years. All these problems share the fact that they can be studied from the point of view of the theory of random walks. However, the high number of degrees of freedom makes it impractical to extract the relevant information from analytical calculations. The alternative then is to use numerical simulations. One can think of formulating models as close to reality as possible including as many degrees of freedom as one is able to handle. The resulting approaches are too complicated to study, not only in what concerns the description of relevant phenomena, but also because they lack predictive power. Therefore the approach we follow is to start with a simpler formulation, less realistic, but retaining the main features of the physical system. Such models are described by fewer parameters, being in this way easier to handle, and making a global study of the parameter space feasible, without the loss of physical intuition.

In this paper, our hope is to obtain relevant results on complex systems by studying relatively simple mod-

els as done in statistical mechanics (SM) where systems as complex as a real ferromagnetic lattice are found to share many properties with models as simple as the Ising model. In SM the relation between both systems is understood by means of such techniques as the renormalization group. Unfortunately we are far from proving that kind of relation for traffic models but nevertheless the study of simple models is still necessary. When we deal with problems involving traffic flows, it is necessary to face problems related to congestion processes. These systems undergo a transition from a situation of fluid traffic to another one characterized by a traffic jam. This change can be related to a phase transition in SM, and hence the same tools developed for statistical systems could be used. First of all though, it is necessary to search for the variables governing this transition.

The transmission in the lattice is dependent on the maximum number of particles that a site can contain simultaneously, the interaction between particles, and the geometric constraints such as the coordination number (number of nearest neighbors).

In our model particles are continuously injected into the whole lattice and a final destination is randomly assigned to each particle. The movement of particles is implemented by allowing particles to move from their present position to their final destination according to some rules which will be given below. In real traffic, particles are often stopped by obstacles in their path. These obstacles are simulated here by limiting the number of particles that a site can contain simultaneously.

At this point, we return to the discussion of the problem from the point of view of SM. Let us focus on the situation where the density is high. It is clear that, from an individual point of view, a particle should try to avoid

*Electronic address: isabel, tarancon, @sol.unizar.es

†Electronic address: fabrice.clerot@lannion.cnet.fr

‡Electronic address: laf@lattice.fis.ucm.es

zones of traffic congestion. This can be done by allowing the trajectory of a particle to fluctuate around its minimal distance trajectory (as determined in an empty lattice). We shall show that this policy indeed brings an improvement on a global level: under some conditions, the system will be globally less congested and the throughput will be improved, despite the temporary misrouting of the particles. This idea is not only of benefit in situations where traffic jams arise, but also in the general case: allowing fluctuations about the minimal distance path improves the characteristics of the traffic.

In terms of SM, the above discussion means that by heating the system ($T \neq 0$) we will obtain better performances [6].

With regard to the particle interactions, up to now we have introduced a contact interaction only: fully occupied sites represent obstacles to the propagation of the particles. The consequences of the contact interaction can be seen as an infinite potential acting on saturated sites. An appropriate longer-range interaction term can be implemented by assigning an *electrical* charge to every particle. The repulsion will move the particles away from regions of high charge. Hereafter we shall only consider interactions among nearest neighbors. Such repulsive forces are usually considered in the case of directed polymers [7].

We have introduced a term to control the size of the thermal fluctuations, and another one to simulate a repulsive force. The parameters controlling these terms are the temperature and the charge, respectively. We will see how, by tuning both parameters, it is possible to improve the performance of the transmission.

II. THE MODEL

We consider a two-dimensional lattice with coordination number 4, and periodic boundary conditions. The particles exist in the lattice sites, labeled by $n \equiv (n_0, n_1)$, and each particle can move from its site to one of its four nearest neighbors at every time step. The maximum number of particles that a site can contain will be denoted by B , and will be kept fixed for the entire simulation at a value $B = 5$. We denote the occupation number of the site n by $\sigma(n)$. In this notation, the particles are prevented from moving to sites with $\sigma(n) = B$. In addition we consider an input queue at each site for the particles waiting to be put on the lattice.

By analogy with SM systems, we work with the inverse of the temperature, β . We denote the charge of the particles by κ and the probability of particle injection to the lattice by p .

The dynamic of the system is as follows.

- (1) A lattice site is chosen at random.
- (2) A particle is added, with fixed probability p , to the queue of the site, waiting to be introduced to the lattice. (The size of the queue is as big as necessary, so that it contains all particles waiting to be fed at that lattice site.)
- (3) If the queue of the site is not empty, and $\sigma(n) < B$, a new particle is introduced in the lattice, and an end

point assigned to it at random.

(4) Each particle at the considered site tends to move towards one of its four neighbors. For a given particle located at the position n we must assign a probability of it jumping to each of its neighboring locations. This probability is given by

$$P(\pm\mu) = N \exp[\pm\beta \text{sgn}(n^f_\mu - n_\mu) - \kappa\sigma(n_\mu)], \quad (1)$$

where n_μ here signifies the μ coordinate of site n , n^f is the end point of the particle being considered, and N is the normalization constant. To choose between possible destinations we use a *heat bath* [8] algorithm.

The factor multiplying β is a potential term. It implies a constant force acting on the particle driving it to its end point. The κ term produces a repulsion between particles sitting in nearer neighbor sites. Obviously there is a wide range of potentials that could be considered, in order to produce more effective forces, and partial improvements will be expected.

(5) Movement is allowed if the chosen site has $\sigma(n) < B$, otherwise the particle remains at its original site until a new movement is attempted in the next iteration.

(6) A particle is removed from the lattice when it reaches its end point.

The smaller p is, the weaker is the effect of the interaction between the particles. For fixed (β, κ) the fluid flow of particles on the lattice only takes place for those values of p being less than a certain threshold which is (β, κ) dependent. Above this threshold, the density becomes too high and the transmission process is prevented. We say the system is saturated. The saturation mechanism begins with the appearance of saturated domains, which grow in size, making the movement of the particles more and more difficult.

Our purpose is to quantify this threshold density, as well as to describe the flow properties along the parameter space (β, κ, p) .

From the point of view of telecommunication networks, the lattice sites can be thought of as nodes of such networks, and *particles* here can be interpreted as *messages* and/or *packets*. Within this approach, B can be identified with the available buffer space at the considered node. In this way, our model can be considered as a first step towards an *abstract* modeling of packet-message telecommunication networks. In such networks, packets are routed by each node according to routing tables dynamically maintained by the network so as to minimize some *cost function* along the trajectory of the packets [9]. In the present model, this cost function is simply the traveled distance, and, as the network state does not change during simulations, routing tables need not be dynamically calculated. A natural extension of this model would be to allow variable link lengths between the nodes and to compute routing tables according to the path lengths. The role of the temperature parameter is to quantify how firmly the network will try to stick to its policy of minimizing the cost along a trajectory: as will be shown below, implementing the routing according to the cost minimization procedure in a rigid way (*zero*

temperature), although this is the best *naive* choice from an individual user point of view, can be detrimental to the global behavior of the network and thus to its collective utility (the total throughput offered by the network).

Congestion is an unavoidable phenomenon in uncontrolled packet networks, since packets introduced at one node have no guarantee of finding the necessary resources (available buffer space) in the transit nodes. If uncontrolled, congestion results in packet losses which, in the case of data transmission, are detected by upper layer protocols (the transport control protocol, TCP, for instance, in the case of the Internet) and are corrected by retransmission request. Hence these packet losses will trigger the introduction of more packets into the network, thus further amplifying the congestion. Two approaches are possible to avoid this vicious circle: reducing the source emission rate when a congestion is detected (this is the TCP approach) or trying to control the congestion inside the network by preventing packets from accessing overloaded areas (this approach does not rely on any source behavior: in the extreme congestion case the source is simply denied any access to the network). This last option, congestion control inside the network, is presently a very active area of research in the networking community [10]. To implement such a control, the overload information should be transmitted by a node to its neighbors at least. This is precisely what is modeled by the repulsive charge term in the present model: access to a loaded node is discouraged, and access to an overloaded node is simply forbidden. It should be noted that in the model as it stands, a node is instantaneously aware of the load state of its neighbors. A more complex model should take into account the latency introduced by the node-to-node transit time, leading to a description of the congestion control by some kind of retarded potential. It will be particularly important to take this latency aspect into account for modeling high speed networks [11], but this is left for future work.

The system we are considering here is also related to a lattice gas with fixed number of particles in a constant electric field \vec{E} [4]. The electric current is related to the traffic speed. Also in this model a transition appears, between a disordered phase and another one where particles move collectively in the \vec{E} direction. In our case the electric field is no longer constant, it depends dynamically on the particle environment: the surrounding electrical charge and the end point associated with the particle. As is usual for traffic models, the number of particles is not constant; as a consequence, the transition mechanism is different.

III. NUMERICAL SIMULATION

We have performed Monte Carlo (MC) simulations in order to study the parameter space. Here we present results obtained using an $L \times L$ lattice with $L = 32$ and periodic boundary conditions. The computations have been carried out on workstations.

The starting configuration is obtained by generating a particle with probability p at each site of the lattice,

therefore the total number of sites occupied initially is about $p \times L^2$. The random number generator is based on one described in [12]. The time step is identified with a MC iteration, that is the update of L^2 lattice sites in the way described in Sec. II.

The temporal evolution of the system exhibits a transient regime, characterized by the instability in the observables (see next section). After this, the system falls into an extremely long-living metastable state, where the flow properties do not change significantly with the temporal evolution. We say the system has reached the asymptotic regime.

The time the system spends in the transient regime depends on the parameter space point (β, κ, p) . In this regime, in the nonsaturated region, the particle density is initially low and grows with the particle injection.

This transit time also grows near the parameter space points where an abrupt change in the properties is observed (e.g., near the threshold density), reaching in this case up to 2×10^4 .

For each value of the parameters, we have performed typically 8×10^4 MC iterations. For β and $\kappa \in (0, 0.4)$ the transient regime takes around 5000 iterations, while above 0.5 for both parameters, this time falls to 400–600 iterations. We have also performed the simulations starting from different configurations, allowing the system to evolve for up to 2×10^5 iterations. We have computed the errors by calculating the dispersion between the results obtained from starting from these different configurations.

IV. OBSERVABLES

A correct description of the system is obtained from the measurement of the relevant observables. From their temporal evolution we are able to evaluate when the asymptotic regime is reached, or even whether this regime will or will not be saturated. From the averaged value of these magnitudes we obtain a quantitative description of the flow process.

For a given configuration, we define the occupation M as

$$M = (1/V) \sum_{i=1}^V \sigma(i), \quad (2)$$

where $V = L \times L$.

The statistical average over the configurations (labeled by j) is

$$\langle M \rangle = \lim_{N \rightarrow \infty} (1/N) \sum_{j=1}^N M_j, \quad (3)$$

N being the number of averaged configurations.

During the time interval $t_j - t_i \equiv \Delta t$, n_e particles will reach their end point. We define the bandwidth (B_W) as the number of particles that arrive at the end point per time unit.

$$B_W(\Delta t) = \frac{n_e}{\Delta t}. \quad (4)$$

The statistical average for B_W is obtained from its mean value over a number of time intervals N_T :

$$\langle B_W \rangle = \lim_{N_T \rightarrow \infty} (1/N_T) \sum_{j=1}^{N_T} B_W(j). \quad (5)$$

We can obtain a good measure of the system performance from the mean time taken by the particles to reach their end point, T_M . We compute T_M in the time interval Δt by adding the individual time spent by each particle (delay time), divided by n_f :

$$T_M(\Delta t) = \frac{\sum_{i=1}^{n_f} T_i}{n_f}. \quad (6)$$

In the same way we define the statistical average over N_T intervals as

$$\langle T_M \rangle = \lim_{N_T \rightarrow \infty} (1/N_T) \sum_{j=1}^{N_T} T_M(j). \quad (7)$$

The occupation frequency of a certain occupation number, $\sigma(n)$, is defined as the number of times that the occupation $\sigma(n)$ appears at any lattice site.

V. PHASE DIAGRAM

We examine the parameter space (β, κ, p) searching for regions where sharp changes in the temporal evolution arise. At each (β, κ) value there is a p value denoted by p_c such that for $p < p_c$ the asymptotic regime presents a stationary flow, and for $p > p_c$ the asymptotic regime is saturated and the flow is no longer possible. We plot the temporal evolution of M below p_c (Fig. 1) and above p_c (Fig. 2) for some values of the parameters.




FIG. 1. Temporal evolution of M below the critical injection p_c .




FIG. 2. Temporal evolution of M above p_c .

This change is similar to a phase transition. The temporal evolution leads the system to one or another phase depending on the parameter space point. Once in the asymptotic regime, the nonsaturated phase exhibits a dynamic equilibrium: the number of injected particles equals the number of those arriving at their end position. This is reflected by $\langle B_W \rangle = pL^2$. In this phase, $\langle T_M \rangle$ is constant, as well as $\langle M \rangle$ which is always less than B .

The parameter space is divided into two regions by the surface defined by (β, κ, p_c) . Figure 3 shows two sections, for fixed β and κ , respectively.

Above the surface, after the transient regime $\langle M \rangle = B$ and $\langle B_W \rangle = 0$, $\langle T_M \rangle$ diverges.

It is possible to give a simple interpretation of this congestion phenomenon; when the load increases so does the probability that two particles residing in two neighboring saturated sites will want to *exchange* their sites, but of course they cannot do so since one of the particles would




FIG. 3. Sections through the phase space cube. On the left side is plotted p_c versus β , on the right side, p_c versus κ .

have to move first and then have nowhere to go as no resource will be available at the desired site. This mutual blocking then tends to propagate to other neighboring sites and once above a given load a complete congestion will develop at the scale of the network.

Such mutual blocking is well known in most *no loss* routing schemes such as *wormhole* routing and is called *deadlock* [13]. Deadlock avoidance is a very active research area, especially for massively parallel system interconnections [14].

β dependence

The thermal fluctuations move particles away from their minimal distance path. The larger β is, the less important these fluctuations are.

As shown in Fig. 3 (left side), there is an interval of β values where the reduction of fluctuations has a positive influence on the throughput: p_c rises, as well as $\langle B_W \rangle$, while $\langle T_M \rangle$ decreases, as shown in Fig. 4.

When going to larger β values, the situation does not persist. The absence of thermal fluctuations damages the throughput because particles are not able to go around obstacles. As a consequence p_c and $\langle B_W \rangle$ decrease. $\langle T_M \rangle$ does not appreciably change from its minimal value, corresponding to the infinite directionality one.

Figure 2 shows how the thermal fluctuations influence the saturation mechanism. If they are important, M grows slowly until it reaches B . If they are not significant a sharp jump appears in the temporal evolution of M , between its value in the transient regime and B . We conclude from this that the saturated domains grow faster in the absence of fluctuations, as would be expected from the earlier discussions.

In particular, it is intuitive that a mutual deadlock will last longer if thermal fluctuations are disallowed (particles in mutual deadlock will repeatedly attempt to exchange their sites) hence creating a larger local congestion area which can eventually evolve towards a global congestion. This was clearly observed in [6] where only one buffer was available per site and this behavior is also




FIG. 4. $\langle T_M \rangle$ dependence on β for $p \approx p_c$.




FIG. 5. (β, κ_{opt}) line. The errors are of the size of the κ step measured ($\Delta\kappa = 0.5$).

exhibited in this model, hence illustrating the well-known fact [15] that deadlock is a consequence of our *no loss* routing scheme (movement is granted if and only if the available resource is available at the low end) and not a consequence of insufficient resources.

κ dependence

The dependence of p_c on κ presents two different regions (see Fig. 3).

In the first one, p_c rises with increasing κ . For these κ values, the inclusion of the repulsion term helps the system to avoid congested regions. We denote by κ_{opt} the value at which p_c reaches its maximal value.

In the second region, $\kappa > \kappa_{opt}$, p_c decreases with κ . Above κ_{opt} on, the repulsion is too strong and the particles move far away from their minimal paths. As a result, the collapse appears for smaller injection densities.

In Fig. 5 we give the values of κ_{opt} for some β values. The more a particle is restricted to its minimal path, the stronger the repulsion needed to avoid obstacles in the lattice.

Therefore there is a maximal value for the particle injection supported by each β value, which is reached for κ in a neighborhood of κ_{opt} .

VI. GENERAL BEHAVIOR AND OPTIMIZATION

Let us focus on the behavior of the other significant observables for the traffic flow, when β and κ have been tuned to obtain a maximal p_c .

Behavior of B_W

As we have already pointed out, B_W in the asymptotic regime is proportional to p by a factor L^2 , as may be seen in Fig. 6. Hence, at each (β, κ) value the maximum for




FIG. 6. $\langle B_W \rangle$ (upper part) and $\langle T_M \rangle$ (lower part) versus p .

$\langle B_W \rangle$ is reached when $p = p_c$.

Above p_c , $\langle B_W \rangle$ remains constant during the transient regime, and falls sharply to zero after it. B_W is a good measure to determine when the system has reached the asymptotic regime.

Behavior of T_M

The greater p the stronger the effect of the interaction over the particles: they move further and further away from their minimal path with increasing $\langle T_M \rangle$ (Fig. 6). In Fig. 7 we observe how $\langle T_M \rangle$ slightly increases with κ until κ_{opt} is reached. From κ_{opt} on, $\langle T_M \rangle$ increases with a high slope.

We conclude that better performances are obtained for the maximal injection supported and for the bandwidth on the line (β, κ_{opt}) . The repulsion term damages T_M : the greater κ , the larger the time the particles take to reach their destination. However, it is only for $\kappa > \kappa_{opt}$ that $\langle T_M \rangle$ increases in a dramatic fashion.

Figure 8 shows the particle distribution in the simulation for $\beta = 0.2$ and $\beta = 4.0$. The plot of $\sigma(n)\text{freq}[\sigma(n)]$




FIG. 7. $\langle T_M \rangle$ dependence on κ .




FIG. 8. Distribution of the occupation number $\sigma(n)$ times the frequency of this state versus $\sigma(n)$ after 5×10^5 MC iterations. Transient regime contributions have been discarded.

exhibits a maximum around $\langle M \rangle$, and also reveals a wider distribution of particles for small β values.

VII. IMPROVING THE THROUGHPUT: FORCE DEPENDING ON THE DISTANCE

Although our implementation of the physical system has been inspired by the search for the simplest model accomplishing the desired features, partial improvements are to be expected by taking into account details not included up to this point. As an example, we develop a possible improvement: quenching thermal fluctuations when they are no longer useful, that is, at the neighborhood of the end points.




FIG. 9. β dependence on p_c for the distance dependent force (dashed line). Dotted line represents the values obtained for the constant force.




FIG. 10. $\langle T_M \rangle$ versus β for fixed $\kappa = 1$ in $p \approx p_c$ for the distance dependent force (dashed line) and constant force (dotted line).

In the primary implementation, we have used a constant force and therefore the particles support thermal fluctuations of the same strength no matter what the remaining distance to the end point is. We can expect a globally less congested lattice if particles which are only a few sites away from their destination are prevented from fluctuating. We will see how the inclusion of this feature does not spoil the good properties of the throughput, and also that the general behavior of the system is not altered.

A possible implementation to take this fact into account is obtained by introducing a distance dependence

in the probability distribution, such that contributions of the fluctuations decrease with decreasing distance to the end point.

A straightforward way of doing this is to include a dependence on the relative distance to the end point: the probability distribution now reads

$$P(\pm\mu) = N \exp[\pm\beta \text{sgn}(n_{\mu}^f - n_{\mu}) - \kappa\sigma(n_{\mu})] \left(\frac{r_n}{r_{n+\mu}} \right), \quad (8)$$

where r_n is defined by

$$r_{(n_0, n_1)} = \sqrt{(n_{f_0}^f - n_0)^2 + (n_{f_1}^f - n_1)^2}. \quad (9)$$

In Fig. 9 the evolution of p_c with β is shown. A global throughput improvement is reflected by higher p_c values. This effect is more remarkable when the size of the thermal fluctuations is important (small β values), corroborating our first intuition on the effect of fluctuations in the steps preceding the end point.

In Fig. 10 we plot the β dependence of T_M . We observe a global decrease in this time for all β values.

The improvements concerning the delay time are not restricted to a smaller $\langle T_M \rangle$. Figure 11 shows the delay distribution for all particles, compared with the delay distribution obtained with the constant potential. We see at $\beta = 0.2$ how the dispersion of the distribution strongly decreases when using a distance dependent force. So, the particles arrive in more similar times, increasing the uniformity and the reliability of the traffic flow.

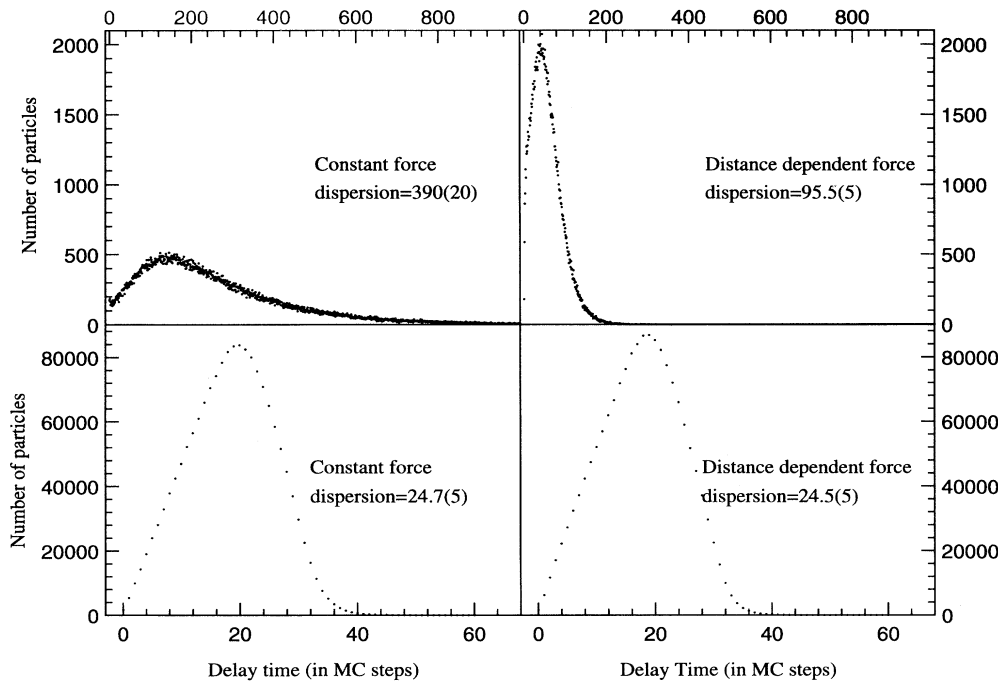


FIG. 11. Delay distribution for the whole particles after 4×10^4 MC iterations. Graphics on the left (constant force) must be compared with those on the right side (distance dependent force). The distributions in the top are calculated in ($\beta = 0.2, \kappa = 1, p = 0.08$), in the bottom they correspond to ($\beta = 4, \kappa = 1, p = 0.008$).

VIII. FAULT TOLERANCE

As we have pointed out, the lattice sites may simulate the control nodes of an information flow system. Needless to say, the sites here are idealizations of nodes in real system because the possibility of communication failures is not allowed. Real nodes are subject to external factors, such as technical constraints or outside influences, that often damage the communication ability between some nodes. These nodes are then temporarily out of order, and cannot communicate or receive information from other nodes. It is therefore of interest to have an estimate of how robust a system is when temporary problems in the transmission occur.

To implement the occurrence of communication problems, information exchange is prevented at n randomly chosen sites during an interval of $\Delta t \equiv 10$ MC iterations in the simulation. After this time interval these n sites are again allowed to communicate and another n sites are broken at random. (In this section the probability distribution (1) is used.)

In general, the measured values for observables will be shifted by an amount depending on n and in the parameter region. For each p value, we define the shift in the observable O as its value relative to the one obtained in the ideal system:

$$\Delta_r O(p) = \frac{O(p, n=0) - O(p, n)}{O(p, n=0)}. \quad (10)$$

We have studied the influence on the relevant param-

TABLE I. Shifted values of relevant observables for some values of the number of failures n .

Shift	$n = 10$	$n = 100$	$n = 200$
$\beta = 4.0, \kappa = 1.0$			
$\Delta_r p_c$	0.09	0.40	0.58
$\Delta_r \langle T_M \rangle (\approx p_c)$	-0.05	-0.38	-0.89
$\Delta_r \langle M \rangle (\approx p_c)$	0.00	0.37	0.55
$\beta = 0.2, \kappa = 1.0$			
$\Delta_r p_c$	0.04	0.24	0.41
$\Delta_r \langle T_M \rangle (\approx p_c)$	-0.00	-0.08	-0.37
$\Delta_r \langle M \rangle (\approx p_c)$	0.00	0.43	0.50

eters of communication failures for $n = 10, 100$, and 200 with $\beta = 0.2$ and $\beta = 4.0$.

Figure 12 shows the comparative evolution of $\langle M \rangle$ for $n = 10$ and $n = 100$.

The first consequence of a bad transmission is that the system must support globally higher occupation numbers. We measure smaller $\Delta_r p_c$ values for lower β , because, as has already been demonstrated, the lattice supports higher occupation numbers when the fluctuations are important.

$\Delta_r \langle M \rangle$ is almost zero for both β values with $n = 10$ (around 1% of nodes out of order), while with $n = 100$ and 200 $\Delta_r \langle M \rangle$ is always large.

In Table I we give the results obtained. As would be expected $\langle T_M \rangle$ increases for all n values even though it is almost zero for $n = 10$.




FIG. 12. $\langle M \rangle$ versus p for values of the fault percentage. Dotted lines represent $\langle M \rangle$ values obtained for an ideal lattice. In the top diagram the parameter space point is ($\beta = 4, \kappa = 1$), and in the bottom frames the plotted point is ($\beta = 0.2, \kappa = 1$).

IX. SUMMARY AND OUTLOOK

We have studied a model useful for describing the relevant processes occurring in traffic flow systems with immediate applications to network message passing and traffic problems in general.

The introduction of the parameters β and κ as controllers of the system behavior allows us to go a step further from purely descriptive models, because we are able to give prescriptions to improve the performance of the flow process.

Deeper studies are also possible. Concretely an accurate study of the scaling with L of the relevant magnitudes as well as a detailed description of how the saturation time behaves could lead us to the definition of quantities analogous to critical exponents. Also non-equilibrium states could be studied in order to monitor

the parameters controlling the saturation process.

The introduction of the κ parameter has implied that the particles are able to avoid congested regions. The study has been limited to short-ranged interactions, because the particles at a site are only affected by the occupation of their nearest neighbors. By informing the particles about the occupation of wider surrounding regions improvements in the throughput are expected.

ACKNOWLEDGMENTS

This work has been partially supported by CICYT Contracts No. AEN93-0604-C03, No. AEN94-0218, and No. AEN93-0776. One of us (I.C.) is grateful for a CERN grant, through the TSF program, and computer facilities, where part of this work was performed.

-
- [1] A. G. Greenberg and B. Hajek, *IEEE Trans. Commun.* **40**, 1070 (1992).
 - [2] J. A. Cuesta, F. C. Martínez, J. M. Molera, and A. Sánchez, *Phys. Rev. E* **48**, R4175 (1993).
 - [3] When-Shyan Sheu and Katja Lindenberg, *Phys. Lett. A* **147**, 437 (1991).
 - [4] S. Katz, J. L. Lebowitz, and H. Spohn, *Phys. Rev. B* **28**, 1655 (1983).
 - [5] O. F. Schilling, *Phys. Rev. B* **44**, 2784 (1991).
 - [6] I. Campos and A. Tarancón, *Phys. Rev. E* **50**, 91 (1994).
 - [7] S. Mukherji and S. M. Bhattacharjee, *Phys. Rev. E* **48**, 3427 (1993).
 - [8] A. Sokal, *Quantum Fields on the Computer* (World Scientific, Singapore, 1992).
 - [9] D. Bertsekas and R. Gallager, *Data Networks* (Prentice-Hall, Englewood Cliffs, NJ, 1987).
 - [10] F. Bonomi and K. W. Fendick, *IEEE Network* **9**, 25 (1995); H. T. Kung and R. Morris, *ibid.* **9**, 40 (1995); K. K. Ramakrishnan and P. Newman, *ibid.* **9**, 49 (1995).
 - [11] L. Kleinrock, *IEEE Commun. Mag.* **30**, 36 (1992).
 - [12] G. Parisi and F. Rapuano, *Phys. Lett.* **157B**, 301 (1985).
 - [13] L. M. Li and P. K. McKinley, *Computer* **26** (2), 62 (1993).
 - [14] D. J. Pritchard and D. A. Nicole, *IEEE Trans. Parallel Distributed Syst.* **4**, 111 (1993).
 - [15] A. Tanenbaum, *Modern Operating Systems* (Prentice-Hall, Englewood Cliffs, NJ, 1992).