left. Meanwhile, it is now possible to completely separate jams and lamellar flow. A system with density $\rho < \rho(q_{max.det}) = 0.5$ can be prepared to behave completely deterministically (Fig. 7.4). Then, a single disturbance initiates, via a chain reaction, a jam, and depending if one has super-critical, sub-critical, or exactly critical density, the emergent jam is, in the average, growing, shrinking, or "undecided".

Since $v_{max} = 1$ means that the jam is dense, this completely corresponds to a random walk description (Chapter 5), i.e. the number of cars in the jam, $n(t)$, behaves as the trajectory of a random walk. If the inflow is higher than the outflow (middle part of Fig. 7.4), then cars enter the jam with a higher rate than they leave, the average jam grows linearly in time, and the corresponding random walk is biased. If the inflow is lower than the outflow, the jam quickly dissolves. If the inflow exactly matches the outflow (which happens generically, because of the outflow behavior; see, e.g., in the lower part of Fig. 7.4), then the random walk is unbiased and shows its typical fluctuations and exponents, as explained earlier.[1]

This instability definitely goes beyond a Burgers equation, since the Burgers equation can never be instable. One has to include the effect of momentum into the partial differential equations.

It is, however, from plots such as Fig. 7.3, not clear why the cruise control limit behavior, i.e. singular jams and the phase separation, should play an important role. This only becomes clear at higher maximum velocities. For that reason, it is worthwhile to spend some more time describing the $v_{max} = 1$ cruise control case, because the following remarks remain valid for arbitrary v_{max}.

It is useful to compare the CA cruise control case results with results from partial differential equation (PDE) models for traffic flow [80, 87, 143, 156]. This is easiest when starting from homogeneous initial conditions, as is done in all figures of this chapter.

- For low densities, i.e. $\rho < \rho(q_{max})$, a disturbance decays, although some survive longer than others. The *average* behavior corresponds to the decay of the wave amplitude of the PDE.
- For high densities, i.e. $\rho > \rho(q_{max})$ but $\rho < \rho(q_{max.det})$, a homogeneous initial state is a meta-stable state. Here, a disturbance *in the average* grows linearly in time (growth of the wave amplitude). However, an *individual* disturbance can still die out.

The picture for the *average* CA wave is this similar to the linear stability analysis of the PDEs for traffic flow, where disturbances become instable when the density is higher than a critical value. There is, however, a difference in the growth rate: PDE waves grow exponentially, CA waves grow linearly in time.

Thus, the flow restriction is not caused by an intrinsic density limit for lamellar flow, but by the fact that there is a density above which the *waves* no longer decay. The rate of inflow which a wave can barely support without growing (in the average) is just the rate of release on the other, the outflow, side. And this rate of release is given by the acceleration characteristics. This is why, in Chapter 6, only the change in the acceleration had a remarkable influence on throughput.

In consequence, a higher flow than the outflow from a jam can only be sustained for a limited time. Measures for stabilizing traffic flow would thus only be useful if the disturbance could be suppressed during the whole rush hour. Otherwise, a disturbance would eventually happen, it would grow because the inflow would be supercritical, and the outflow from the jam would reset the traffic system to the same lower flow as before the introduction of the new technology. From traffic science measurements, this phenomenon is known as hysteresis [169].

[1]Actually, for Fig. 7.4 this is not precisely true. There is a correlation via the periodic boundary conditions.

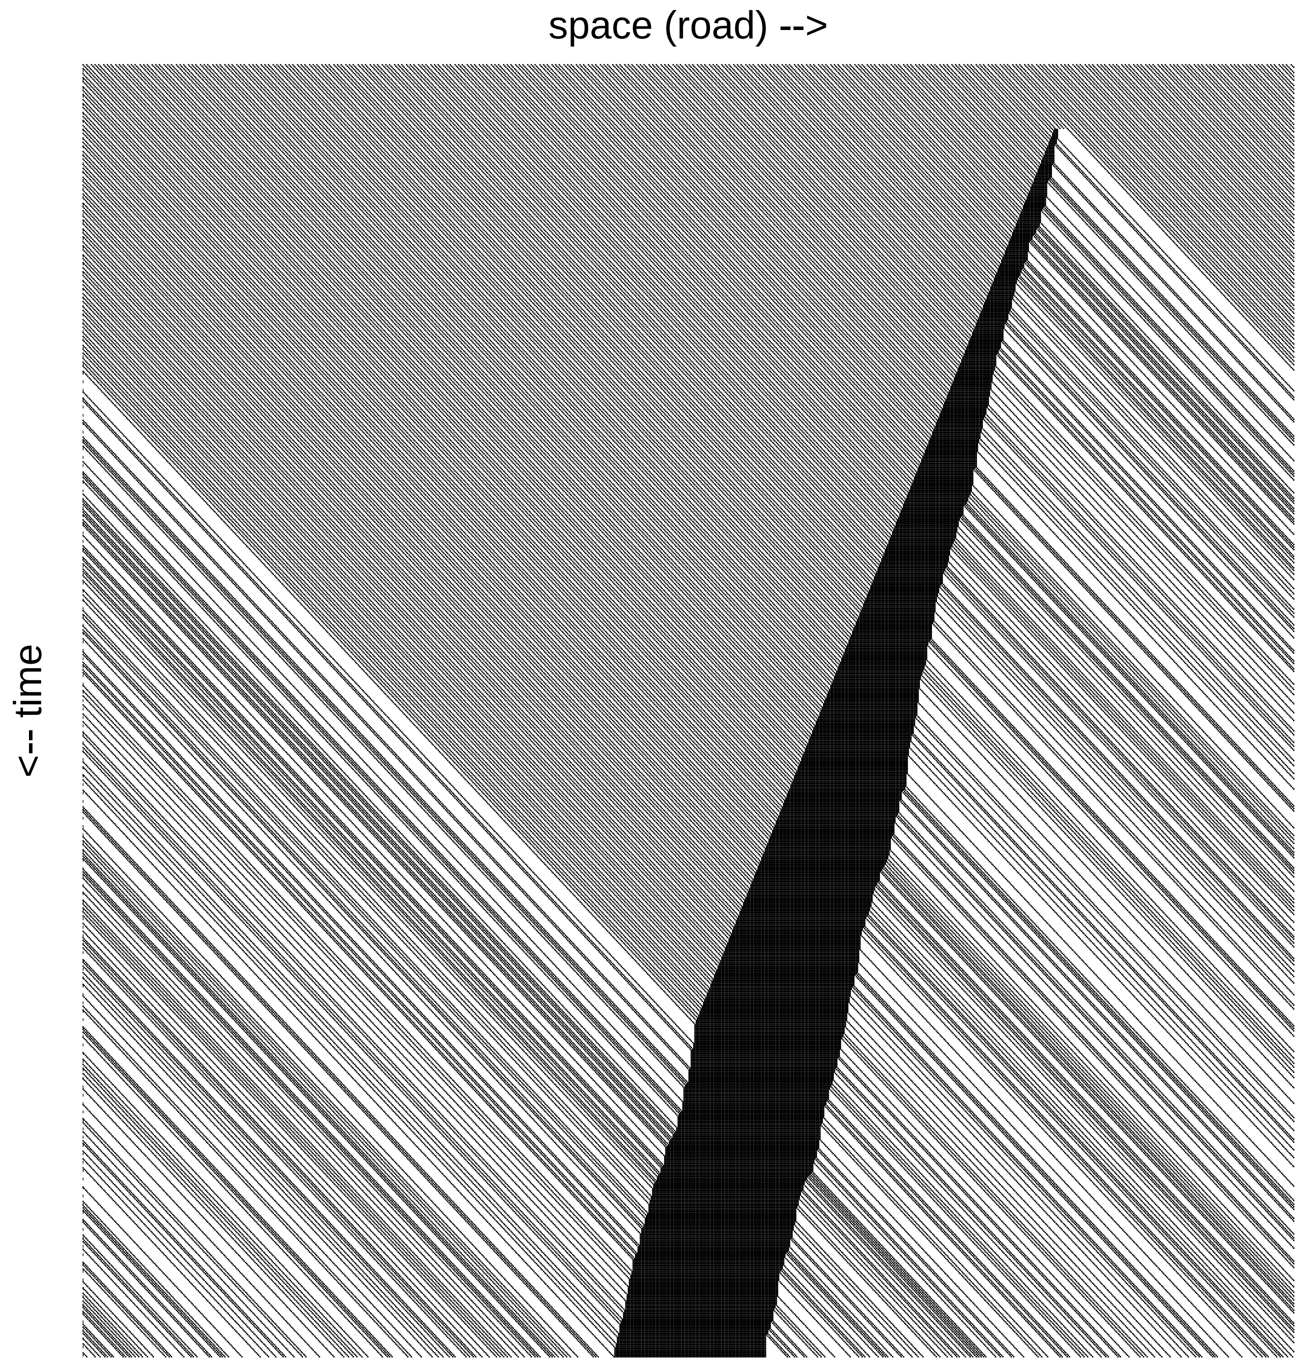

Figure 7.4: Space-time plot for parallel update (cruise control limit) and $v_{max} = 1$. The first 50 time steps show deterministic supercritical flow; after the disturbance, the jam clearly grows; when the outflow reaches the jam again via the periodic boundary conditions, both jam fronts describe a random walk (but see footnote in text).—In order to be clearer, the figure uses a lower acceleration probability than normal.

In the cruise control limit, the flow-density-relation is non-smooth at q_{max}. This just reflects the phase separation:

- Below $\rho(q_{max})$, in the steady state there are only cars with maximum velocity v_{max}, leading to a linear flow-density relation $q_< = \rho \cdot v_{max}$. The velocity of the kinematic waves $c_< = q' = v_{max}$, i.e., disturbances move with the traffic.
- Above $\rho(q_{max})$, traffic separates into laminar regions operating at $(\rho(q_{max}), q_{max})$ and jammed regions operating at $(\rho = 1, q = 0)$. The fundamental diagram is just the linear interpolation between both: $q_> = q_{max}\,(1-\rho)/(1-\rho(q_{max}))$. The velocity of the kinematic waves

$$c_> = \frac{\Delta q}{\Delta \rho} = -\frac{q_{max}}{1-\rho(q_{max})} \; .$$

No other wave velocity is possible in the steady state. The effect is similar (although smeared out) for higher velocities, which explains why the waves in the model all move backwards with similar velocities.

Takayasu and Takayasu [167] have shown that the bi-stability can also be obtained by a deterministic model—although in this case one replaces the random walk picture by its deterministic average.

7.4 The full CA model for traffic flow

For higher maximum velocities, e.g. $v_{max} = 5$, there are three new observations: The cruise control limit dynamics now visibly plays a role, i.e., singular jams are an important feature of the dynamics; jams can branch; and analytic work is more difficult.

- **Cruise control dynamics.** The cruise control regime now visibly plays a role, i.e. simulated traffic flow separates into laminar and jammed regions, as in reality. However, for this, there is no *dynamical* difference to the $v_{max} = 1$ case, it is only that the natural parameters are closer to the cruise control case.

  Together with the knowledge that the $v_{max} = 1$ case looks like random sequential update, and the knowledge of the cut-off in the lifetime measurements (see Chapter 6), one comes to the conjecture that a correct description of traffic contains a cross-over from cruise-control dynamics to Burgers/KPZ dynamics. An associated correlation length ξ describes up to which length the cruise control description is valid; there is also a similar correlation time. Judging from Treiterers space-time-plots (see Fig. 3.1), for traffic this length scale has to be larger than 10 km.

  It is this effect which is responsible for the special shape of the fundamental diagram. If ξ were infinity (cruise control limit), the fundamental diagram would be non-smooth at $\rho(q_{max})$. However, a $\xi < \infty$ smoothes out the peak region.

- **Jams can branch.** Completely new for $v_{max} > 1$ is the effect that jams now can branch. This behavior enters *naturally* into the model when one attempts to obtain a more realistic fundamental diagram by using a maximum speed v_{max} larger than one. The branching is due to the memory effect of $v_{max} > 1$, where cars leaving a jam need some time to accelerate to v_{max}. During this acceleration phase, they follow different rules than when driving at maximum speed, which may lead to a sub-branch of the jam (Fig. 7.5).

  I am not aware of any traffic measurements which directly support this picture, but it seems intuitively reasonable—after an accident site has been cleared, one usually drives through many start-stop-waves with free traffic in between—, and it is actually very difficult to suppress it in the simulations as soon as one allows fluctuations in the acceleration.

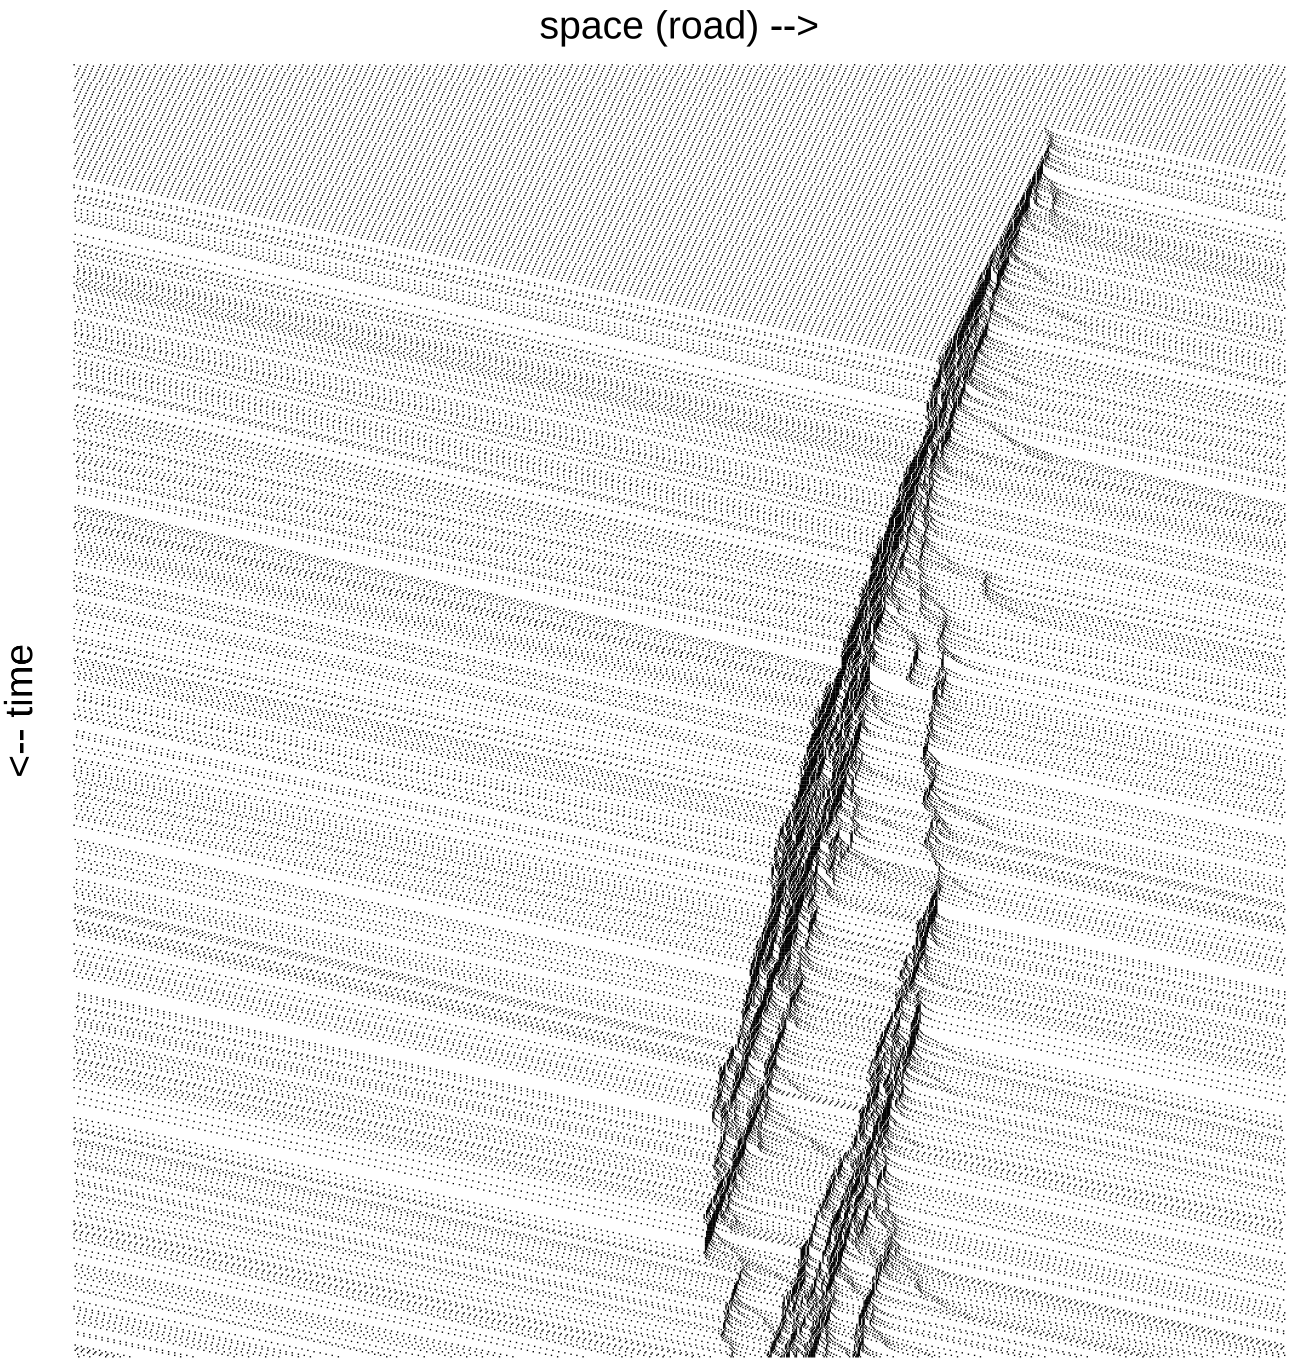

Figure 7.5: Space-time plot for parallel update (cruise control limit), $v_{max} = 5$, $\rho = 0.09$, i.e. slightly above critical. The flow is started in a deterministic, supercritical configuration, but from a single disturbance separates into a jam and a region of exactly critical density.—This is phenomenologically the same plot as Fig. 7.4 except that $v_{max} = 5$.

Chapter 9 will contain an example how branching will make the density estimation for congestion pricing much harder than one could expect.

- **Analytic results.** The fact that the internal jam dynamics becomes more complicated is reflected by the fact that the analytic approximation becomes more complicated [147, 150]. The

n-sites approximations only converge towards the correct solution for $n \to \infty$, but do not seem to become exact.

Even with its natural parameters, the model behaves very similar to the fluid-dynamical results of the Daimler-Benz research group [80, 88, 143, 156]: Starting from ordered initial conditions (Fig. 7.6), dynamic instabilities develop when $\rho >\sim \rho_c$, which transform into kinematic waves, which survive for a long time. See Ref. [96] for a thorough treatment of the differences and connections between dynamic and kinematic waves. These models use a smooth flow-density relation; the correct behavior comes out due to the nonlinear effects in the equations.

It is interesting to note that both the fluid-dynamical and the particle hopping approach go in the same way beyond the Lighthill-Whitham theory: They replace instantaneous velocity adaption with momentum (an additional momentum equation for the fluid-dynamical models, the acceleration rules for the hopping model), which leads to the instability of the lamellar flow of high densities.

7.5 Random sequential vs. parallel update

As mentioned in Chapter 3, one usually assumes that random sequential updating is a better approximation of reality than parallel updating, and that one chooses parallel updating only for computational reasons. Why is for traffic flow parallel updating more realistic?

The problem is that for the random sequential update, the *fluctuations* of the speed are proportional to v_{max}, whereas in the parallel update, they are independent of v_{max}. In consequence, the *relative* fluctuations around $\overline{v}$ are constant for random sequential update, but decrease as $\propto 1/v_{max}$ for parallel update. This is what makes the cruise control limit relevant for higher velocities: Fluctuations at free driving are decreased.

In order to "save" the random sequential update, one would thus have to reduce the fluctuations around the mean velocity for free driving. This could, e.g., be achieved by moving smaller spatial steps and making the probability of moving proportional to the velocity of the car. But then, during one cycle a vehicle would have to be moved several times, which would make the simulation much slower. Noise-reduction techniques for surface growth models work similarly [81, 179].

7.6 Summary for traffic flow theory

As one sees, once one has an overview over connections between the different hopping models and to the fluid-dynamical models, there is a wide variety of traffic flow phenomena which could be readily explained. This will be, together with discussions with traffic engineering people and with data evaluations, the topic of further work. Here, I just want to summarize the—in my view—most important results which have come out of this research in the area of traffic flow:

- There is strong evidence from this work, from fluid-dynamical results, and from car-following simulations that, on a sufficiently small scale, stable traffic flow regimes are only: light traffic $\rho < \rho_c$, outflow traffic $\rho = \rho_c$, and dense traffic $\rho \approx 1$. All other traffic with $\rho > \rho_c$ undergoes a spontaneous phase separation into outflow regions of ρ_c and jammed regions with densities close to one. Many traffic measurements can be explained using this clear picture.

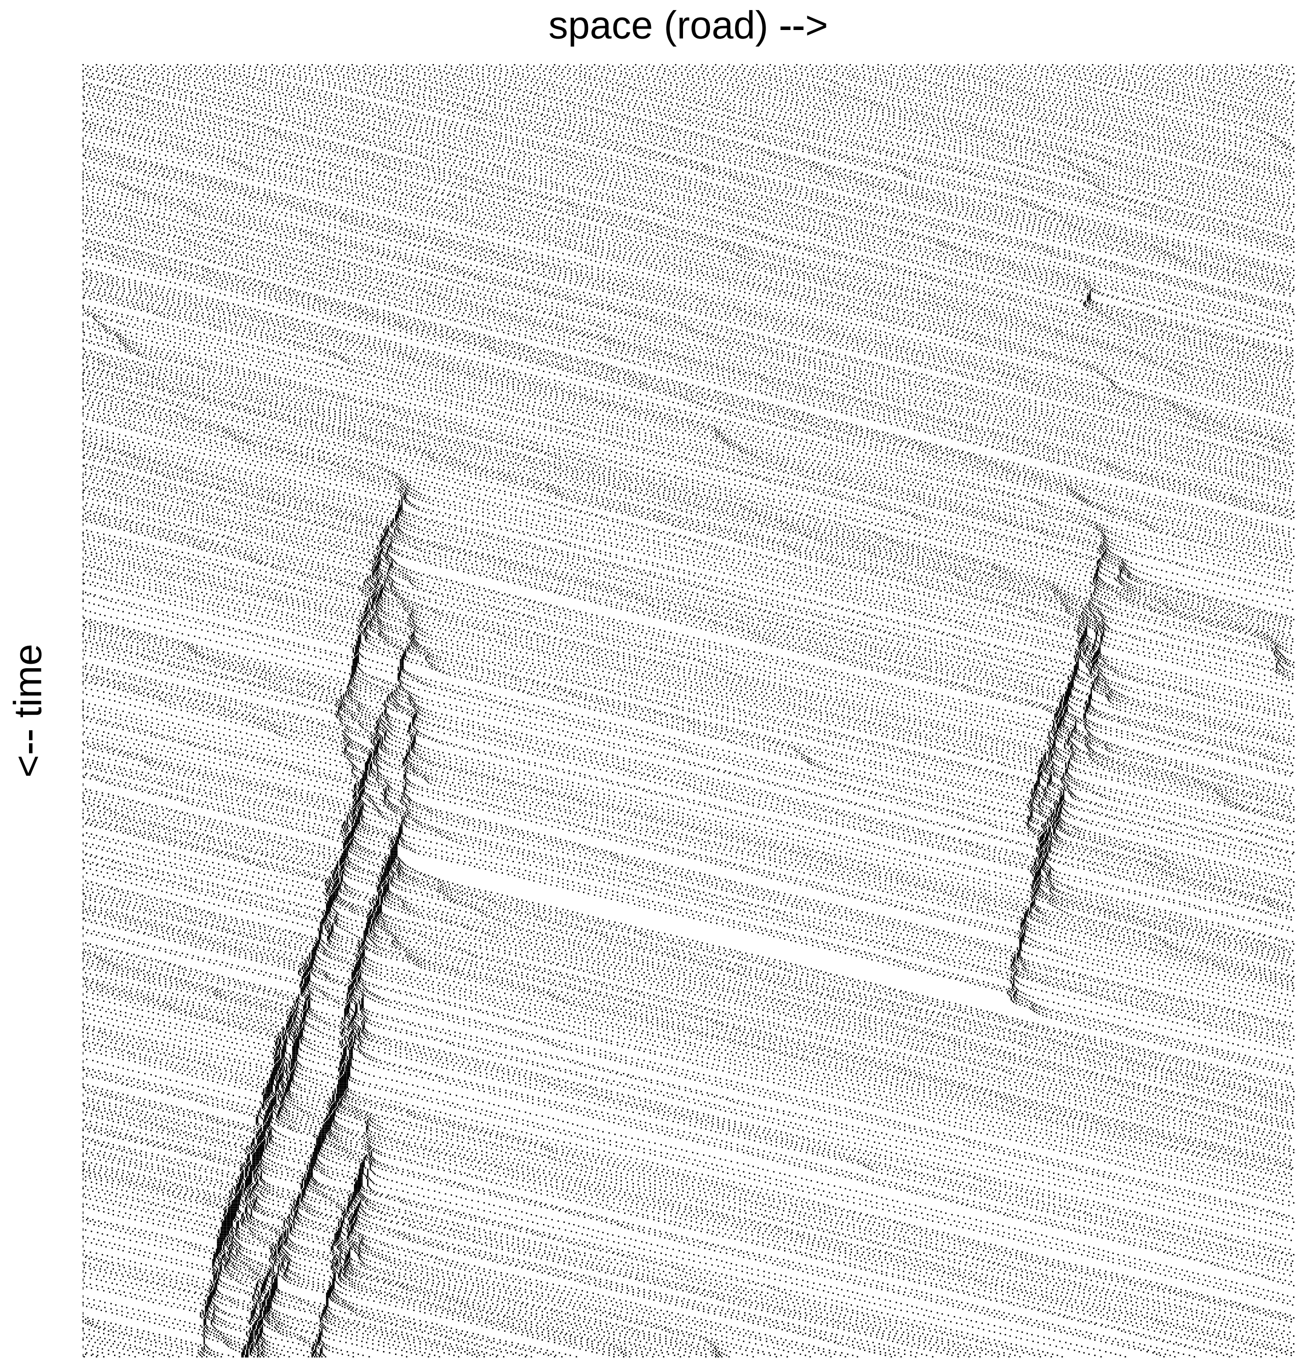

Figure 7.6: Space-time plot for parallel update, $v_{max} = 5$, $\rho = 0.09$ (i.e. slightly above $\rho(q_{max})$), starting from ordered initial conditions. The ordered state is meta-stable, i.e. "survives" for about 300 iterations until is spontaneously separates into jammed regions and into regions with $\rho = \rho(q_{max})$.

- Since all traffic above ρ_c eventually separates into the two regions explained above, the maximum flow is given by the outflow from a jam, and this is given by the acceleration rules. Faster acceleration at the outflow of a start-stop-wave greatly enhances the throughput.

This situation is described in the average by a *linear* Burgers equation. The fluctuations around this average picture (apart from the branching complications) are described by a random walk.

- The opposite picture would be that traffic flow is mostly constrained by disturbances "all over the place".

  This picture corresponds to the noisy nonlinear Burgers equation and the standard KPZ universality class.

- All results indicate that in reality there is a crossover between both descriptions, where the "cruise control limit"-description is valid on short scales, and the Burgers description on large scales. Both descriptions are thus correct, but for different length scales, and judging from traffic observations, the crossover length is larger than 10 km.

  The length and time scale of this crossover increase with *de*creasing p_{free}, the fluctuations at maximum speed.

- Traffic jams out of nowhere are caused by fluctuations at maximum speed. In the fluid-dynamical description, they are caused as instabilities in the equations for *dynamic* waves, but survive as *kinematic* waves [88, 80, 96]. The CA model supports this picture.

- Reducing the fluctuations at maximum speed does not increase maximum throughput very much. Therefore, installing technology for better controlling maximum speed will not be very useful without coupling it to the acceleration (see above).

  The reason why that does not help is that maximum flow is dominated by the outflow from a jam, and as long as it is not possible to keep traffic *completely* free of jams, it does not make a big difference if there are many or few jams in the system. However, reducing fluctuations at maximum speed will largely enhance throughput when used in bottlenecks of a shorter length: Then, it is feasible to keep the traffic along this length completely free of jams.

- Jam waves branch. Between the branches, one finds holes "of all sizes" (see Chapter 5). This makes the concept of queues difficult for state estimation problems: From a spatial configuration at a given time step alone, one cannot decide if two jams have the same dynamical origin or not—one needs the history.

The results about the universality classes mean that, to a large extent and in the way described here, all car following models will lead to the same macroscopic behavior. Since all major simplifications of the the CA traffic model lead either to losing the effect of spontaneous jam formation or to losing the lamellar flow regime, it is improbable that a much simpler model can be found which gives the same macroscopic behavior. On the other hand, using a more complicated model is only justified if one is interested in effects which go beyond the ones identified in this work. This is certainly the case for effects like the design of ramps or intersections, but seems—at least in first approximation—useless for most studies on the freeway network level. In other words: If one is satisfied with a numeric version (i.e. including diffusion and noise) of the Lighthill-Whitham (LW) model for certain questions but needs a microscopic version, then the approach presented here should be more than enough. It is even better than discrete LW in at least one respect: It correctly reproduces the instability above $\rho(q_{max})$. For example, whereas discrete LW will always predict the same average queue before a bottleneck, the particle model will allow both days without breakdowns and days with very severe breakdowns. This can make a difference when, e.g., only the queue from the severe breakdown reaches back into, say, another intersection.

Moreover, together with our implementation results on different supercomputers (Chapter 11) this means that, on current supercomputers, it will be difficult to find a faster microscopic traffic simulation model with the same phenomenology.