Asynchronous Updates in Large Parallel Systems

Albert G. Greenberg
AT&T Research
albert@research.att.com

S. Shenker
Xerox PARC
shenker@parc.xerox.com

Alexander L. Stolyar
Motorola
stolyara@cig.mot.com

Abstract Lubachevsky [5] introduced a new parallel simulation technique intended for systems with limited interactions between their many components or sites. Each site has a local simulation time, and the states of the sites are updated asynchronously. This asynchronous updating appears to allow the simulation to achieve a high degree of parallelism, with very low overhead in processor synchronization. The key issue for this asynchronous updating technique is: how fast do the local times make progress in the large system limit? We show that in a simple K-random interaction model the local times progress at a rate $1/(K+1)$. More importantly, we find that the asymptotic distribution of local times is described by a *traveling wave* solution with exponentially decaying tails. In terms of the parallel simulation, though the interactions are local, a very high degree of global synchronization results, and this synchronization is succinctly described by the traveling wave solution. Moreover, we report on experiments that suggest that the traveling wave solution is *universal;* i.e., it holds in realistic scenarios (out of reach of our analysis) where interactions among sites are not random.

1 Introduction

Simulation is the most widely used and reliable tool for understanding the behavior of systems with many interacting components. Even if the interactions between the components are fairly simple and local in nature, and the states of the components are piecewise constant (as opposed to varying continuously), the dynamics of such systems are often beyond the reach of our current analytical techniques. Examples of such systems include large computer networks and certain models of interacting particles. In [5], Lubachevsky introduced an asynchronous updating technique for simulating such systems on large parallel machines.

A simplified version of the simulation technique can be described as follows. Consider a system with N components, or *sites.* Each site has its own local simulation time, which describes the simulation time up to which the current state of the site is valid. These local times are a function of the real time t, and we denote the local time of site i at real time t by $x_i(t)$. The state of each site is governed by a transition rule. The exact nature of the transition rule does not concern us here, but the nature of the dependencies introduced by these transition rules is of crucial importance. We can denote by $D_i(t)$ the set of sites whose states are relevant to the updating of site i at real time t.

The simulation method associates each site with a processor. The processor attempts to update the state of the site at random times, modeled here as a Poisson process. The rate of attempted updates per unit time is μ, which in this paper we set to be one. If an update attempt fails, then the site waits for the next randomly arriving update attempt. Site i can be updated at time t if and only if $x_i(t) < x_j(t)$ for all $j \in D_i(t)$. See Figure 1. The systems to which this method is applied are such that the state of each site, once updated, is guaranteed to remain unchanged for a period of time that can be computed at the time of updating. In applications, this period arises from the detailed model of the lags being simulated; examples include Ising models [5], Markovian networks of queues [8], and dynamic channel assignment schemes in wireless cellular systems [3]. After updating site i, the local time $x_i(t)$ can be incremented by the length of this static period, since the current site state will be valid at least until then. The length of the static period is a property of the system being simulated. We will assume that this period is an *iid* random variable governed by the probability den-

SIGMETRICS '96 5/96 PA, USA
© 1996 ACM 0-89791-793-6/96/0005...$3.50

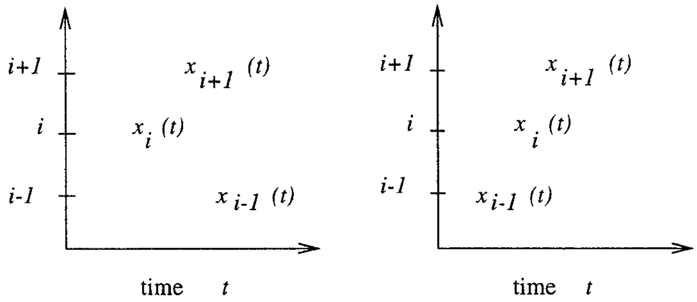

Figure 1: Example: Suppose that at time t, an update attempt at site i is dependent on sites $D_i(t) = \{i-1, i+1\}$. Site i can update if its local time ties or lags the local times of sites $D_i(t)$, as depicted on the left hand side. Otherwise, as depicted on the right hand side, site i cannot update.

sity $r(z)$, where we assume, without loss of generality, that the mean of this distribution is one: $\int_{-\infty}^{\infty} zr(z)dz$. Define $R(z) = \int_z^{\infty} r(u)du$. Although our results are applicable to a wider class of distributions, we will focus our attention on the case where $r(z) = e^{-z}$.

The key question about this simulation approach is, in the limit $N \to \infty$, do the local times $x_i(t)$ make sufficient progress. More specifically, this can be reduced to two separate questions. First, in the large system limit, do the local times progress at a nonzero speed so that $\lim_{t\to\infty} \frac{x_i(t)}{t} > 0$ holds for all i? Second, if a nonzero average speed is achieved, is the asymptotic distribution of local times reasonably tight? A tight distribution means that the system as a whole has made simulation progress. It is not hard to imagine the formation of long chains of dependencies, resulting in very few of the sites succeeding in update attempts, and in an asymptotic rate of progress that tends to 0 as N increases.

Sites in computer networks and in interacting particle systems often have local interactions, in the sense that a bounded subset of the sites is involved in each update. Our main result, presented in Section 2 but proven in Section 4, is that when each site interacts with K randomly chosen neighbors, the local times progress at a speed of $\frac{1}{K+1}$ for any $r(z)$. More importantly, the distribution of local times converges to a *traveling wave* solution with exponentially decaying tails. In Section 3, we present results drawn from extensive experiments treating more realistic interaction patterns, such as those arising in regular lattices. These experiments show that $O(1/K)$ growth in local times and traveling wave solutions are *universal*. In every case we consider that has bounded dependency sets $D_i(t)$, all initial conditions with a bounded distribution of initial times appear to converge to a traveling wave solution. In each case the trailing edge of this asymptotic traveling wave is exponentially tight, while the leading edge depends on the form of $r(z)$. We feel these results reveal the key structure of the asynchronous updating simulation technique, and demonstrates its scalability to large systems.

As this technique requires no rollbacks to correct temporary inconsistencies that arise during execution, it is known as a *conservative* technique. Work on the efficiency of conservative techniques relevant to our model includes [2] and [6]. A great deal of work has also been done on the efficiency of rollback techniques, such as Time Warp [4]. A recent survey of the state of the art in parallel and distributed simulation can be found in [7].

2 K-Random Interaction Model

The K-random interaction model chooses, for each update attempt, K sites at random to comprise the set $D_i(t)$. This set is rechosen for each update attempt, even if a previous update attempt has failed. We find it convenient, as a matter of convention, to assume that there is an additional ordering relationship that applies between two sites with identical local times, so that only one site prevents the other from updating. That is, if at any given time site i prevents site j from updating, then site j does not prevent site i from updating. However, the additional notation needed to describe this additional ordering relation is cumbersome, and is largely irrelevant to our treatment here, so we omit it. This additional ordering is superfluous when all sites have distinct local times. When the initial condition has the no-equal-local-times property, then with probability 1 this no-equal-local-times property continues to hold for all t.

Considering the limiting case of an infinite number of sites, we can define a function $f(x,t)$ to be the proportion of sites with local time less than or equal to x at time t. This function $f(\cdot,t)$ completely describes the state of the system at any time t.

Remark The formal limit transition to the case of an infinite number of sites is done in [1]. Namely, it is proven that a sequence of processes with increasing number of sites converges to a deterministic process described by equation (1) we introduce below.

The first question is: what is the average rate of progress of the simulation? That is, how fast does the average local time increase with real time. Recall that the average step size – the average increase in local times for a successful update – has been set to unity. The probability $p_i(t)$ that a given site i with local time x_i will be able to successfully update at time t (if an attempt to update is made) is given by $(1 - f(x_i(t), t))^K$. Let $p(t)$ be the average over all sites of the $p_i(t)$. This quantity is then given by:

$$p(t) = \int_0^1 (1 - f(x,t))^K df(x,t) = \frac{1}{K+1}$$

The average rate of progress, for any distribution $f(\cdot, t)$ and any function $r(z)$ with unit mean, is exactly $\frac{1}{K+1}$. One can motivate this result by considering a model where the sites are grouped into cliques of $K+1$ members and each site depends only on the other K sites in the clique. This is equivalent to an ensemble of fully interacting systems, each with $K+1$ sites. The average rate of progress in such systems is exactly $\frac{1}{K+1}$.

However, as we observed before, knowing the average rate of progress is not sufficient to conclude that this simulation technique is viable. At the end of the simulation we want to have essentially all sites to have made the same linear rate of progress. To ensure this, we must ask the second question: is the distribution of local times should be relatively tight? We therefore need to study the evolution of the distribution $f(x,t)$ in more detail. If we assume $f(\cdot, t)$ is differentiable (a condition we relax in Section 4) we can write the following evolution equation (with the notation that $f'(u,z) = \frac{\partial f}{\partial x}(u,t)$):

$$\frac{\partial f}{\partial t}(x,t) = -\int_{-\infty}^{x} (1 - f(u,t))^K f'(u,t) R(x-u) du \quad (1)$$

One can look for traveling wave solutions of this equation: $f(x,t) = \phi(x - vt)$ for some wave velocity v. Clearly, from the result about the average rate of progress, we must have $v = \frac{1}{K+1}$. For the case $r(z) = e^{-z}$, a family of traveling wave solutions is given by $\phi_\alpha(x) = 1 - (1 + e^{K(x-\alpha)})^{\frac{-1}{K}}$. These solutions have exponentially decreasing densities for large positive and negative divergences from the mean local time. If the system tends towards this solution then we are assured that the distribution of local times is sufficiently tight. The bulk of this paper is devoted to proving that a wide class of initial conditions all converge to the same traveling wave solution.

Because of its length, we delay the proof until Section 4. We first discuss the behavior of this scheme in more realistic scenarios.

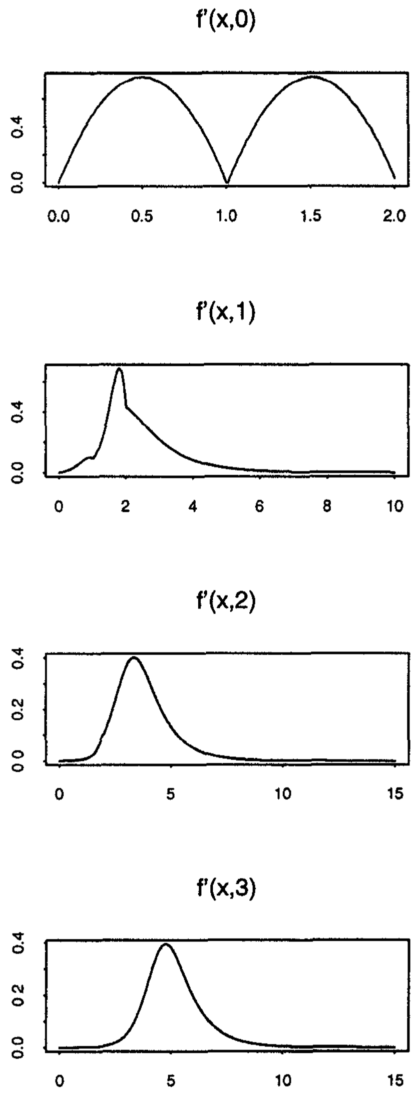

Figure 2: Rapid convergence to the traveling wave solution ($K = 2$). $f'(x,t) = \frac{\partial f}{\partial x}(x,t)$ is plotted along the vertical axis and x on the horizontal axis.

3 Experimental Results

In this Section, we begin with experimental results on the convergence of system to the traveling wave solution, and then explore more realistic models outside the reach of our analysis where the update dependencies adhere to a regular graph structure. The experiments show that the qualitative properties revealed by the analysis of the K-random model are universal.

Figure 2 illustrates the rapid convergence of the sys-

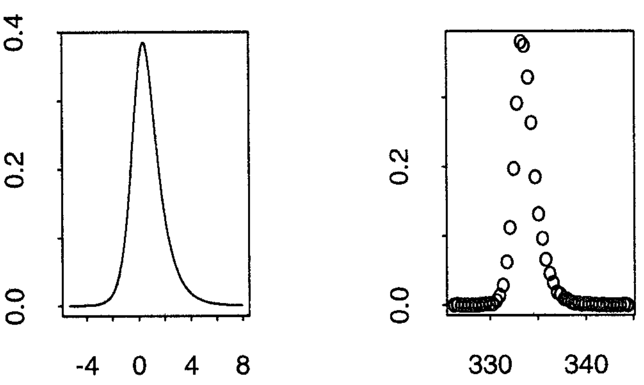

Figure 3: On the left hand side is a plot of the traveling wave $\phi'(x)$. On the right hand side is a snapshot of the local time density $f'(x,t)$ taken from a 10,000 site Monte Carlo simulation. $K = 2$.

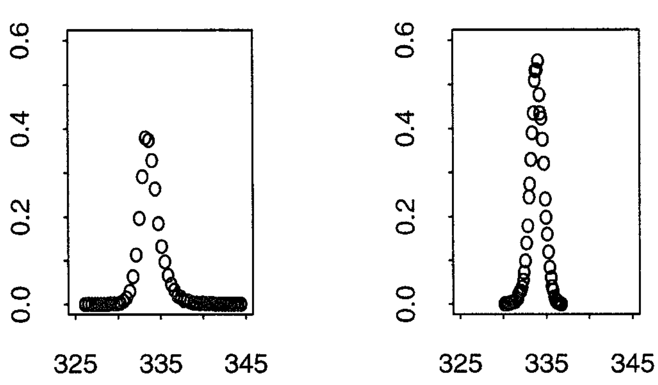

Figure 4: Local times taken from snapshots of the local time density $f'(x,t)$ taken from 10,000 site Monte Carlo simulations, with $K = 2$. The data on the left hand side is for $r(z) = exp(-z)$; the data on the right hand side is for $r(z) = 1/2$ for $z \in [0,2]$.

tem $f'(x,t)$ to the traveling wave solution $\phi'(x - t/(K+1))$. At time $t = 0$ (upper left plot), we assume the local time density $f'(x,0)$ is bimodal in $[0,2]$. However, by time $t = 4$, the density has already become unimodal, with a steep trailing edge and a leading edge that is beginning to look exponential. This trend continues so that we reach at time $t = 12$ a density $f'(x,12)$ that is indistinguishable from the traveling wave solution.

Figure 3 compares the traveling wave solution to the empirical density of local times drawn from a Monte Carlo simulation of the 10,000 site K-random neighbor model, after 1 million update attempts. The correspondence, illustrating the convergence of the finite model, to the infinite one described by (1) is apparent.

Figure 4 compares empirical densities of local times for the 10,000 site K-random model, with the static period either exponentially distributed ($r(z) = e^{-z}$, left plot) or uniformly distributed on $[0,2]$ ($r(z) = 1/2$ for $z \in [0,2]$, right plot). In both cases, we obtain traveling waves moving at the same rate. The Figure illustrates the rule that the more concentrated the distribution of the static period the sharper the wave obtained.

Last, we consider models where an event update attempts adhere to a graph structure. That is, the set of sites $D_i(t)$ relevant to an update attempt at site i are simply the neighbors of site i in a fixed graph. Figure 5 depicts results for three graphs, all having about 10,000 sites: a toroidally connected 2 dimensional mesh; a toroidally connected 3 dimensional mesh; and a butterfly graph, chosen for its logarithmic diameter, but fixed degree (number of neighbors of a given site) $K = 4$. The Figure illustrates that the empirical densities strongly resemble the traveling wave solutions for K-random models with K chosen as the degree of the graph. To first order the key determinant of behavior is the graph degree; in particular, the butterfly and the 2d mesh behave similarly. This is supported by the estimates of local time growth rates reported in Table 1. This is good news because the complete graph on N sites admits no parallelism (the local time growth rate is $1/N$), and the butterfly is typical of graphs that approach the connectivity of the complete graph as quickly as possible subject to having fixed degree.

Our data suggest that the qualitative properties of the K-random model are universal:

- The rate of growth of local time for a regular graph models of degree K is $O(1/K)$.
- For large systems the densities of local times converge to a traveling wave solution.

1d mesh K= 2	2d mesh K= 4	3d mesh K= 6	butterfly K= 4
0.21	0.094	0.058	0.085

Table 1: Estimated local time growth rates.

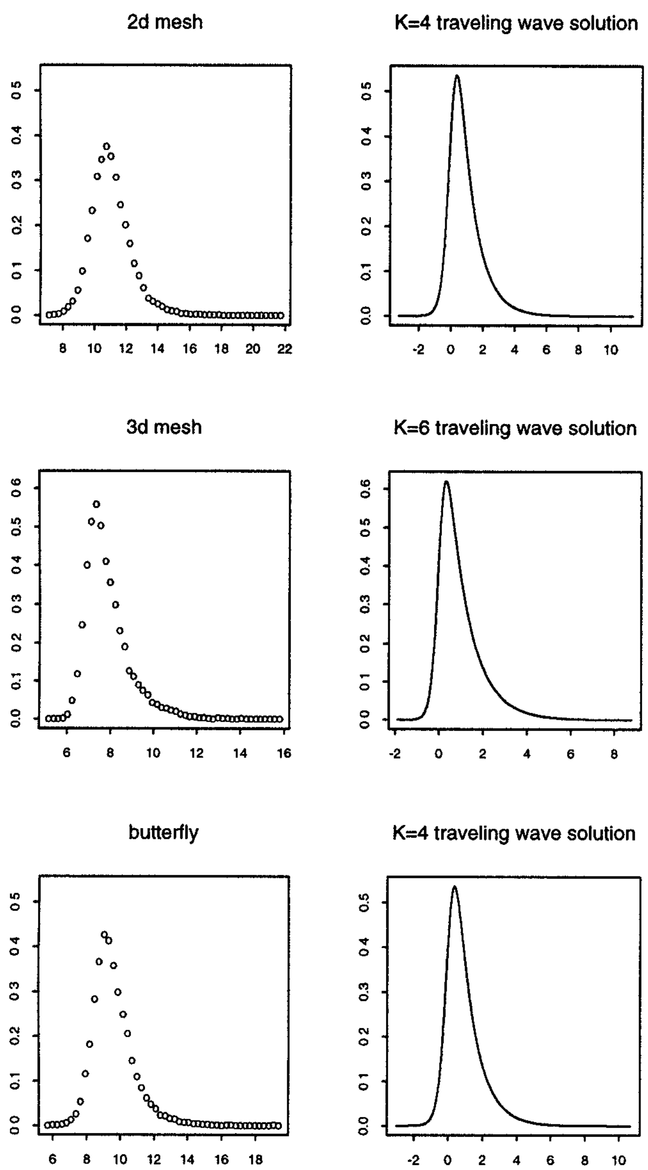

Figure 5: Impact of graph structure.

4 Convergence to the Traveling Wave Solution

We now return to the assertion made in Section 2 that essentially all initial conditions converge to a traveling wave solution. We first introduce some notation and then state this assertion more precisely, before presenting the proof.

4.1 Notation

For a probability distribution described by $f(x,t)$, we define $m(f(\cdot,t))$ to be the mean value of the distribution at time t:

$$m(f(\cdot,t)) \equiv \int_{-\infty}^{\infty} x df(x,t)$$

Similarly, define the inverse function

$$f^{-1}(y,t) \equiv \inf\{x | f(x,t) > y\}$$

for $0 \leq y < 1$.

We denote the L_1 norm of a function g by $||g||$:

$$||g|| \equiv ||g||_{L_1} \equiv \int_R |g(x)| dx$$

We will use the following notation to identify the positive and negative components of a function:

$$g^+ = \max(g,0), \quad g^- = -\min(g,0) \ .$$

We will also use a step function $\theta^{(a)}(x)$ given by:

$$\theta^{(a)}(x) \quad = \quad \begin{cases} 0, & x < a \\ 1, & x \geq a \end{cases}$$

Also, we set $v = \frac{1}{K+1}$, $q(y) = (1-y)^K$.

4.2 Statement of Results

As we described in Section 1, the function $f(x,t)$ describes the fraction of sites that have local times less than x at real time t. The dynamics of the simulation process is described by the following equation:

$$\frac{\partial f(x,t)}{\partial t} = h(x,t) \qquad (2)$$

where:

$$h(x,t) = -\int_0^{f(x,t)} dy \ q(y) R(x - f^{-1}(y,t)) \qquad (3)$$

Note that solutions to (1) are also solutions to this equation; the above formulation has the advantage that we need not assume differentiability of the distribution $f(\cdot,t)$. We will find it convenient to rewrite (3) as:

$$\begin{aligned} h(x,t) &= -\int_0^{f(x,t)} dy\ q(y) + \int_0^{f(x,t)} dy\ q(y) \\ &\cdot \int_{f^{-1}(y,t)}^{x} r(\xi - f^{-1}(y,t)) d\xi \qquad (4) \\ &\equiv -\int_0^1 dy\ q(y)\theta^{(f^{-1}(y,t))}(x) + \int_0^1 dy\ q(y) \\ &\cdot \int_0^\infty dz\ r(z)\theta^{(f^{-1}(y,t)+z)}(x) \qquad (5) \end{aligned}$$

We restrict ourselves to the set of *acceptable* initial conditions $f(x,0)$, which have the following properties:

(a) $f(x,0)$ is nondecreasing right continuous;

(b) $f(-\infty,0)=0,\ f(\infty,0)=1$;

(c) $f(x,0)$ has "integrable tails":

$$\int_{-\infty}^0 f(x,0)dx < \infty\ , \quad \int_0^\infty (1-f(x,0))dx < \infty$$

(d) $f(x,0)$ has zero mean:

$$m(f(\cdot,0)) = 0$$

Theorem 1 *If there exists a traveling wave solution*

$$\tilde{f}(x,t) = \phi(x - vt)$$

of equation (2) obeying initial conditions (a) - (d), then it is unique, the function $\phi(x)$ is continuous, and any other solution $f(x,t)$ of equation (2) satisfying (a) - (d) converges to that traveling wave solution as $t \to \infty$, both pointwise

$$f(x+vt,t) \to \tilde{f}(x+vt,t) \equiv \phi(x), \quad \forall x \in R \qquad (6)$$

and in the sense of L_1-distance

$$\|f(x+vt,t) - \phi(x)\| \to 0 \qquad (7)$$

Remarks:

1) As we mentioned in Section 2, for the case $r(z) = e^{-z}$, the traveling wave solution obeying (a) - (d) does exist and is given by $\phi(x-vt)$ with $\phi(x) = 1-(1+e^{K(x-\alpha)})^{\frac{-1}{K}}$ and α chosen so that $m(\phi)=0$.

2) Any traveling wave solution $\tilde{f}(x,t) = \phi(x-vt)$ defines a family of traveling wave solutions

$$\tilde{f}^{(a)}(x,t) \equiv \tilde{f}(x-a,t) = \phi(x-a-vt)$$

3) Since functions $f(x+vt,t)$ and $\phi(x)$ are nondecreasing, and $\phi(x)$ is continuous, the point-wise convergence (6) actually follows from the L_1-norm convergence (7).

We present the proof of this result in several phases. We first present an informal overview of the approach, followed by some preliminary results, and then we present the detailed proof.

4.3 Overview of Proof

4.3.1 Decreasing L_1 Distance

We will view a function $f(x,t)$ as a function $f(\cdot,t)$ of time t taking values in the set of nondecreasing right-continuous functions of x. The basic idea of the proof of Theorem 1 is to show that the L_1-distance between any two solutions $f_1(\cdot,t)$ and $f_2(\cdot,t)$ to equation (2), $\|f_1(\cdot,t)-f_2(\cdot,t)\|$, can only decrease in time. More precisely, we will show, that the derivative

$$\frac{\mathrm{d}}{\mathrm{dt}}\|f_1(\cdot,t)-f_2(\cdot,t)\| < 0$$

as long as $f_1 > f_2$ on a set of nonzero measure, and $f_1 < f_2$ on a set of nonzero measure. To show that we use the following argument. The function $h(\cdot,t)$ defined in (4) is the derivative of $f(\cdot,t)$ in L_1-space. (We will prove this later — equation (2) does not immediately imply this.) Formula (4) can be rewritten as follows:

$$h(x) = \int_0^1 \mathrm{dy}\ q(y) \int_0^\infty \mathrm{dz}\ r(z)[-\theta^{(f^{-1}(y))}(x)+\theta^{(f^{-1}(y)+z)}(x)]$$

Consider any two solutions $f_1(\cdot,t)$ and $f_2(\cdot,t)$ to equation (2). We can write

$$\frac{\mathrm{d}}{\mathrm{dt}}(f_1(\cdot,t)-f_2(\cdot,t)) = s(\cdot)\ ,$$

where

$$\begin{aligned} s(x) &= \int_0^1 \mathrm{dy}\ \ q(y) \int_0^\infty \mathrm{dz}\ \ r(z)\cdot s^{(y,z)}(x)\ , \\ s^{(y,z)}(\cdot) &= [-\theta^{(a_1)}(\cdot)+\theta^{(a_2)}(\cdot)] + [\theta^{(a_1+z)}(\cdot)-\theta^{(a_2+z)}(\cdot)]\ , \\ a_i &= f_i^{-1}(y),\ \ i=1,2\ . \end{aligned}$$

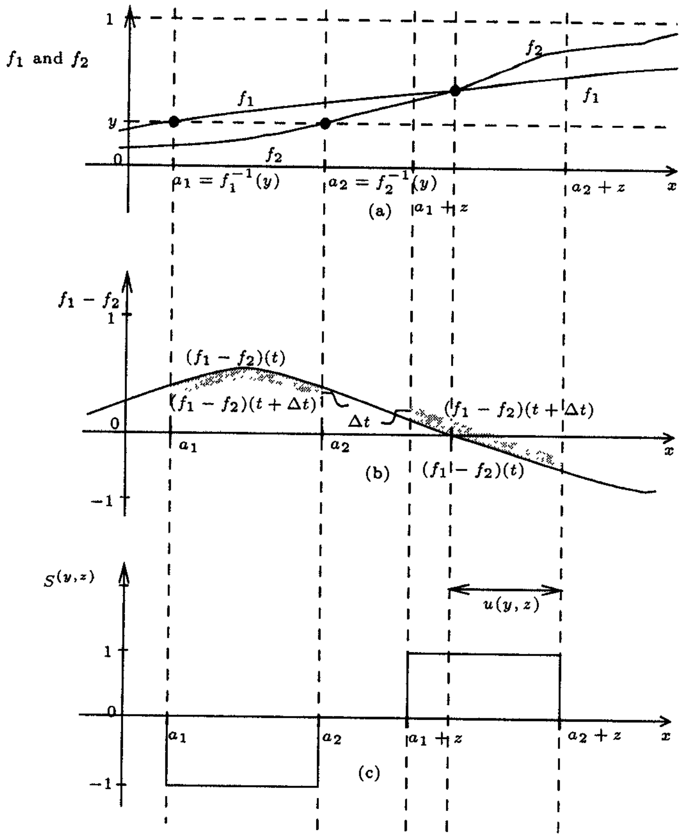

Figure 6:

(We will also denote, $s_l^{(y,z)} = [-h^{(a_1)} + h^{(a_2)}]$, $s_r^{(y,z)} = [h^{(a_1+z)} - h^{(a_2+z)}]$.) As we see, the derivative s of the difference $f_1 - f_2$ is a weighted sum of "primitive derivatives" $s^{(y,z)}$ defined for all $y \in (0,1)$ and $z \in (0,\infty)$. The form of function $s^{(y,z)}$ is illustrated in Figure 1.

For the case where

$$\frac{\mathrm{d}}{\mathrm{dt}}(f_1(\cdot,t) - f_2(\cdot,t)) = \sigma(\cdot)$$

let us denote

$$\begin{aligned} F_\sigma^+ &\equiv \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1(\cdot,t) - f_2(\cdot,t))^+\| \\ F_\sigma^- &\equiv \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1(\cdot,t) - f_2(\cdot,t))^-\| \end{aligned}$$

Let us find the expression for $F^+_{s(y,z)}$. For example, consider the case $a_1 = f_1^{-1}(y) < a_2 = f_2^{-1}(y)$, illustrated in Figure 1. We see, that

$$F^+_{s_l^{(y,z)}} = -(a_2 - a_1) \le 0$$

Indeed, by definition $a_1 = f_1^{-1}(y)$, so function $f_1 - f_2$ is positive everywhere in the segment $[a_1, a_2]$, and $s_l^{(y,z)} = -1$ in that segment.

We also see that

$$F^+_{s_r^{(y,z)}}(y,z) = (a_2 - a_1) - u(y,z)$$

where $u(y,z)$ is the length of the interval within $[a_1 + z, a_2 + z]$ where $f_1 - f_2$ is negative. Therefore,

$$F^+_{s(y,z)} = F^+_{s_l^{(y,z)}} + F^+_{s_r^{(y,z)}} = -u(y,z)$$

It is verified similarly that $F^-_{s(y,z)} = -u(y,z)$.

Considering the three different cases, $a_1 < a_2$, $a_1 = a_2$, and $a_1 > a_2$, we get the following result:

$$F^+_{s(y,z)} = F^-_{s(y,z)} = -u(y,z)$$

where nonnegative function $u(y,z)$ is defined as follows. Denote $c_1 = \min(a_1, a_2)$, $c_2 = \max(a_1, a_2)$. (Recall, that $a_i = f_i^{-1}(y)$). If $a_1 = a_2$, then $u(y,z) = 0$. If $a_1 \ne a_2$, then

$u(y,z) =$ {Measure of the set of those points in the segment $[c_1 + z, c_2 + z]$ in which $f_1 - f_2$ has the sign opposite to the sign of $f_1 - f_2$ in the segment $[c_1, c_2]$}.

Since the derivative s is a weighted sum of primitive derivatives $s^{(y,z)}$, it is natural to expect that the derivative

$$\frac{\mathrm{d}}{\mathrm{dt}}\|(f_1 - f_2)^+\| \equiv F_s^+$$

should be equal to the integral

$$-\int_0^1 \mathrm{dy}\ q(y) \int_0^\infty \mathrm{dz}\ r(z) \cdot u(y,z) \le 0$$

We will show (Theorem 2), that actually the following result holds:

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1 - f_2)^+\| &\le -\int_0^1 \mathrm{dy}\ q(y) \int_0^\infty \mathrm{dz}\ r(z) u(y,z) \\ &\le 0 \qquad (8) \\ \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1 - f_2)^+\| &= \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1 - f_2)^-\| \qquad (9) \end{aligned}$$

Then, obviously,

$$\begin{aligned} \frac{\mathrm{d}}{\mathrm{dt}}\|f_1 - f_2\| &= \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1 - f_2)^+\| + \frac{\mathrm{d}}{\mathrm{dt}}\|(f_1 - f_2)^-\| \\ &\le -2\int_0^1 \mathrm{dy}\ q(y) \int_0^\infty \mathrm{dz}\ r(z) u(y,z) \le 0\ . \end{aligned}$$

Therefore, $\|f_1 - f_2\|$, the L_1-distance between two solutions, has *a strictly negative derivative* as long as both sets (of argument x) where $f_1 - f_2$ is strictly positive and strictly negative have nonzero measures.

This, in turn, implies that if two solutions $f_1(\cdot, t)$ and $f_2(\cdot, t)$ have equal "mean value"

$$m(f_1(\cdot, t)) = m(f_2(\cdot, t))$$

then $\|f_1 - f_2\|$ has a strictly negative derivative as long as they *do not coincide*.

4.3.2 Convergence to a traveling wave

In the rest of the proof of Theorem 1 we consider the L_1-distance between an arbitrary solution $f(\cdot, t)$ with mean value

$$m(f(\cdot, t)) = vt$$

and the traveling wave solution $\tilde{f}(x, t) = \phi(x - vt)$ having equal mean value

$$m(\tilde{f}(\cdot, t)) = vt$$

In other words, we assume $f(\cdot, t) = f_1(x, t)$ and $\tilde{f}(\cdot, t) = f_2(\cdot, t)$ and use the properties (8) and (9) of $\|f_1 - f_2\|$. We consider the "time shifted" version of $f(x, t)$:

$$\hat{f}(x, t) = f(x + vt, t)$$

This is the function $f(\cdot, t)$ continually shifted to the left with speed v to keep its mean value constant:

$$m(\hat{f}(x + vt, t)) = 0, \quad \forall t$$

A similar shift applied to the traveling wave $\tilde{f}(x, t)$ makes the traveling wave time invariant

$$\tilde{f}(x + vt, t) \equiv \phi(x)$$

Thus, the problem is to prove that

$$\hat{f}(\cdot, t) \to \phi(x)$$

We prove the following sequence of statements.

(i) The family of functions $\hat{f}(\cdot, t)$, $t \geq 0$, viewed as probability distribution functions on the real axis, is relatively compact.

This is true, because otherwise $\|\hat{f}(\cdot, t) - \phi(\cdot)\|$ can be arbitrarily large, which is impossible.

This, in turn, implies, that the family of functions $\{\hat{f}(\cdot, t), t \geq 0\}$ has limiting functions as $t \to \infty$. Then we consider an arbitrary limiting function $\overline{f}(\cdot)$. It suffices to prove that $\overline{f}(\cdot) = \phi(\cdot)$. First, we show (Lemma 8), that

(ii) Any limiting function $\overline{f}(\cdot)$ must have a form $\phi^{(a)}(x) = \phi(x - a)$, i.e., a shifted version of $\phi(\cdot)$.

This is true, because otherwise we could find a "shifted version" $\phi^{(a)}$ such that the difference $\overline{f}(\cdot) - \phi^{(a)}(\cdot)$ is positive and negative on sets of nonzero measure. But, the function $\hat{f}(\cdot, t)$ gets infinitely close to $\overline{f}(\cdot)$ infinitely often. This would mean, that the derivative

$$\frac{\mathrm{d}}{\mathrm{d}t}\|\hat{f}(\cdot, t) - \phi^{(a)}(\cdot)\| < \eta < 0$$

is separated from zero by a negative constant η on a set of time instants t having infinite measure. This is impossible, because it would imply $\|\hat{f}(\cdot, t) - \phi^{(a)}(\cdot)\| \to -\infty$.

The last observation is (Lemma 9):

(iii) Any limiting function $\overline{f}(\cdot)$ is exactly equal to $\phi(x)$.

This is proved by a contradiction. If, for example, $\overline{f}(\cdot) = \phi^{(a)}(\cdot)$ for some $a < 0$, then

$$\lim_{t \to \infty} \|(\hat{f}(\cdot, t) - \phi(\cdot))^+\| \geq |a|$$

and therefore

$$\lim_{t \to \infty} \|(\hat{f}(\cdot, t) - \phi(\cdot))^-\| \geq |a|$$

Moreover, it is easy to see that

$$\lim_{t \to \infty} \|(\hat{f}(\cdot, t) - \phi^{(c)}(\cdot))^-\| \geq |a|$$

for arbitrarily large $c > 0$. But this is impossible, because we can always make $\|(\hat{f}(\cdot, 0) - \phi^{(c)}(\cdot))^-\|$ arbitrarily small by choosing a sufficiently large $c > 0$.

4.4 Preliminary Results

4.4.1 Properties of Solutions to (2)

Lemma 1 *(a) For any x and $t \geq 0$, $h(x, t) \leq 0$. If in addition $f(x, t) > 0$, then $h(x, t) < 0$.*

(b) $|h(x, t)| \leq \frac{1}{K+1}$

(c) For any fixed x, $f(x, t)$ is absolutely continuous on t.

(d) If $f(x_1, t_0) < f(x_2, t_0)$, $x_1 < x_2$, $t_0 \geq 0$, then $f(x_1, t) < f(x_2, t)$ for all $t \geq t_0$.

(e) If $f(x, 0) > 0$, then $f(x, t) > 0$, $\forall\ t \geq 0$.

(f) If $f(x_1,0) = f(x_2,0) > 0$, $x_1 < x_2$, then $h(x_1,0) < h(x_2,0)$.

(g) For any $t > 0$, the function $f(x,t)$ is strictly increasing in the interval $x_* < x < \infty$, where $x_* = -\infty$ if $f(x,0) > 0$, $\forall x$, and

$$x_* = \sup\{x \mid f(x,0) = 0\}$$

otherwise.

(h) For any $t > 0$, the inverse function $f^{-1}(y,t)$, $y \in (0,1)$, is continuous and nondecreasing.

(j) For any fixed x, the function $f(x,t) - f(x-,t)$ is strictly decreasing in t unless $f(x,t) = f(x-,t)$. Therefore the set of points x where $f(x,t)$ is not continuous remains unchanged in time.

Notice, that (j) immediately implies the continuity of the function $\phi(x)$ defining a traveling wave solution $\phi(x - vt)$.

Proof. Throughout the proof we use the following representation of the derivative $h(x)$ in (4):

$$h(x) = -h_l(x) + h_r(x)$$

where

$$h_l(x) = \int_0^{f(x)} dy\ q(y)$$

$$h_r(x) = \int_0^{f(x)} dy\ q(y) \int_{f^{-1}(y)}^{x} r(\xi - f^{-1}(y))d\xi$$

Obviously, both $h_l(x)$ and $h_r(x)$ are non-decreasing on x, and $h_l(x) \geq h_r(x)$.

(a) As mentioned above, $h_l(x) \geq h_r(x)$. If $f(x) > 0$, then $h_l(x) > h_r(x)$.

(b) $|h(x)| \leq h_l(x) \leq \int_0^1 q(y)dy = \frac{1}{K+1}$

(c) Follows from (b).

(d) Denote: $\Delta f = f(x_2) - f(x_1)$. Then

$$\frac{\mathrm{d}}{\mathrm{dt}}(\Delta f) = h(x_2) - h(x_1)$$

$$= -(h_l(x_2) - h_l(x_1)) + (h_r(x_2) - h_r(x_1))$$

$$\geq -(h_l(x_2) - h_l(x_1))$$

$$= -\int_{f(x_1)}^{f(x_2)} q(y)\ dy \geq -\int_0^{\Delta f} q(y)\ dy$$

We see that

$$\frac{\mathrm{d}}{\mathrm{dt}}(\Delta f) \geq -\Delta f + o(\Delta f)$$

which implies that Δf cannot attain 0 in finite time.

(e) Follows from (d).

(f) In this case $h_l(x_1) = h_l(x_2)$ and $h_r(x_1) < h_r(x_2)$.

(g) Follows from (d), (e), and (f).

(h) Follows from (g).

(j) It is verified directly, that if $f(x-,t) < f(x,t)$, then $\lim_{\epsilon\downarrow 0} h(x-\epsilon,t) < h(x,t)$. This implies the statement (j).

Lemma 2 *For any $t \geq 0$, $\|h(\cdot,t)\| = v$. For any $t \geq 0$ and $\tau \geq 0$,*

$$f(x,t+\tau) \leq f(x,t), \quad \forall x$$

and

$$\|f(\cdot,t+\tau) - f(\cdot,t)\| = \int_{-\infty}^{\infty} [f(x,t) - f(x,t+\tau)]dx = v\tau$$

Proof. Directly from the expression (4) we get:

$$\begin{aligned} -\|h(\cdot,t)\| &= \int_{-\infty}^{\infty} h(x,t)dx \\ &= \int_0^1 dy\ q(y) \\ &\quad \cdot \int_0^{\infty} dz\ r(z)\{f^{-1}(y) - (f^{-1}(y) + z)\} \\ &= -\int_0^1 dy\ q(y) \int_0^{\infty} dz\ r(z)z = -v \end{aligned}$$

Using Lemma 1(c) we can write

$$\begin{aligned} \int_{-\infty}^{\infty} [f(x,t) &- f(x,t+\tau)]dx \\ &= -\int_{-\infty}^{\infty} dx \int_t^{t+\tau} d\xi\ h(x,\xi) \\ &= -\int_t^{t+\tau} d\xi \int_{-\infty}^{\infty} h(x,\xi)\ dx = v\tau \end{aligned}$$

Lemma 3

$$f(\cdot,t+\Delta t) - f(\cdot,t) = h(\cdot,t)\Delta t + o(\Delta t) \qquad (10)$$

where the left and right sides of (10) are understood as L_1-valued functions of Δt .

Proof. According to Lemma 1(a) and Lemma 2, for any t, $h(\cdot,t)$ is non-positive function of x having norm v. Again, according to Lemma 1a) and Lemma 2, the same properties has the function

$$\hat{h}_{\Delta t}(\cdot,t) \equiv [f(\cdot,t+\Delta t) - f(\cdot,t)]/\Delta t$$

for any Δt. By definition of h we have convergence everywhere:

$$\lim_{\Delta t \downarrow 0} \hat{h}_{\Delta t}(x,t) = h(x,t), \ \ \forall x$$

This convergence everywhere along with the above properties of h and $\hat{h}_{\Delta t}$ easily imply the convergence $\hat{h}_{\Delta t} \to h$ in L_1.

4.4.2 Additional Preliminary Results

Consider two functions

$$g = \{g(x),\ x \in R\} \in L_1 \quad \text{and} \quad s = \{s(x),\ x \in R\} \in L_1$$

Denote:

$$\frac{d_s\|g^+\|}{dt} \equiv \lim_{\Delta t \downarrow 0} \frac{\|(g + s\Delta t + o(\Delta t))^+\| - \|g^+\|}{\Delta t} \tag{11}$$

and, similarly,

$$\frac{d_s\|g^-\|}{dt} = \lim_{\Delta t \downarrow 0} \frac{\|(g + s\Delta t + o(\Delta t))^-\| - \|g^-\|}{\Delta t} \tag{12}$$

The following Lemma 4 shows that the derivatives in (11) and (12) are well defined and gives an explicit expression for them.

Lemma 4 *Let $g, s \in L_1$. Then*

$$\frac{d_s\|g^+\|}{dt} = \int_R dx\ s(x)[I\{s>0\}\cdot I\{g \geq 0\} + I\{s<0\}I\{g>0\}]$$

and, respectively,

$$\frac{d_s\|g^-\|}{dt} = -\int_R dx\ s(x)[I\{s<0\}I\{g \leq 0\} + I\{s>0\}I\{g<0\}]$$

Proof. First, we notice that for any two functions $g, u \in L_1$,

$$\|(g+u)^+\| - \|g^+\| \leq \|u\|$$

Thus,

$$\|(g + s\Delta t + o(\Delta t))^+\| = \|(g + s\Delta t)^+\| + o(\Delta t)$$

and therefore

$$\begin{aligned} &\lim_{\Delta t \downarrow 0} \frac{\|(g + s\Delta t + o(\Delta t))^+\| - \|g^+\|}{\Delta t} \\ = &\lim_{\Delta t \downarrow 0} \frac{\|(g + s\Delta t)^+\| - \|g^+\|}{\Delta t} \\ = &\lim_{\Delta t \downarrow 0} \frac{1}{\Delta t} \int [(g + s\Delta t)^+ - g^+]\ dx \end{aligned}$$

Denote: $\Delta g_{\Delta t} \equiv (g + s\Delta t)^+ - g^+$. Then, breaking down the real axis into non-intersecting subsets according to the sign of $g(x)$ and $s(x)$ we can easily verify the following set of equalities:

$$\int_{s=0} \Delta g_{\Delta t} dx = 0, \ \ \forall \Delta t$$

$$\int_{s>0,g\geq 0} \Delta g_{\Delta t} dx = (\int_{s>0,g\geq 0} s(x)dx)\Delta t, \ \ \forall \Delta t$$

$$\int_{s>0,g<0} \Delta g_{\Delta t} dx = o(\Delta t)$$

$$\int_{s<0,g\leq 0} g_{\Delta t} dx = 0, \ \ \forall \Delta t$$

$$\int_{s<0,g>0} \Delta g_{\Delta t} dx = (\int_{s>0,g\geq 0} s(x)dx)\Delta t + o(\Delta t)$$

This completes the proof.

Lemma 5 *Let $g, s \in L_1$, and $s_n \xrightarrow{L_1} s$, $n \to \infty$. Then*

$$\frac{d_{s_n}\|g^+\|}{dt} \to \frac{d_s\|g^+\|}{dt}$$

and, respectively,

$$\frac{d_{s_n}\|g^-\|}{dt} \to \frac{d_s\|g^-\|}{dt}$$

Proof. The direct limiting transition in the explicit expressions for the derivatives given by Lemma 4 proves this lemma.

Lemma 6 *Let $g, s_1, s_2 \in L_1$, and*

$$\frac{d_{s_1}\|g^+\|}{dt} = \frac{d_{s_1}\|g^-\|}{dt}$$

$$\frac{d_{s_2}\|g^+\|}{dt} = \frac{d_{s_2}\|g^-\|}{dt}$$

Then

$$\frac{d_{s_1+s_2}\|g^+\|}{dt} = \frac{d_{s_1+s_2}\|g^-\|}{dt}$$

and

$$\frac{d_{s_1+s_2}\|g^+\|}{dt} \leq \frac{d_{s_1}\|g^+\|}{dt} + \frac{d_{s_2}\|g^+\|}{dt}$$

Proof. Easily proven using the explicit expressions of Lemma 4.

We are now prepared to present the detailed proof. This will be done in two steps. We first discuss the decrease in L_1-distance between two solutions, and then establish the convergence to the traveling wave solution.

4.5 Decreasing L_1 Distance

Consider two solutions $f_1(x,t)$ and $f_2(x,t)$ to the equation (2). Corresponding derivative functions we will denote $h_1(x,t)$ and $h_2(x,t)$. We are interested in the asymptotics of f_1 and f_2 when $t \to \infty$. Therefore, without loss of generality, we can assume that statements (g) and (h) of Lemma 1 hold for all t including $t = 0$.

Denote

$$\begin{aligned} g(x,t) &\equiv f_1(x,t) - f_2(x,t) \\ s(x,t) &\equiv h_1(x,t) - h_2(x,t) \end{aligned}$$

Let us fix $t = 0$. Denote

$$F_\sigma^+ = \frac{d_\sigma \|g^+(\cdot,0)\|}{dt}, \quad F_\sigma^- = \frac{d_\sigma \|g^-(\cdot,0)\|}{dt}$$

According to Lemmas 3 and 4,

$$F_s^+ = \frac{d_s \|g^+(\cdot,0)\|}{dt} = \frac{d \|g^+(\cdot,0)\|}{dt}$$

From the formula (4) we get

$$s(x) = \int_0^1 dy\ q(y) \int_0^\infty dz\ r(z) s^{(y,z)}(x)$$

where

$$\begin{aligned} s^{(y,z)} &= \left[-\theta^{(f_1^{-1}(y))} + \theta^{(f_2^{-1}(y))}\right] \\ &\quad + \left[\theta^{(f_1^{-1}(y)+z)} - \theta^{(f_2^{-1}(y)+z)}\right], \end{aligned}$$

for $0 < y < 1$, $z \geq 0$.

Lemma 7 *(a)* $F_{s^{(y,z)}}^+ = -u(y,z) =$

$$\begin{cases} -\displaystyle\int_{f_1^{-1}(y)+z}^{f_2^{-1}(y)+z} d\xi\ I\{f_1(\xi) < f_2(\xi)\}, & f_1^{-1}(y) \leq f_2^{-1}(y) \\ -\displaystyle\int_{f_2^{-1}(y)+z}^{f_1^{-1}(y)+z} d\xi\ I\{f_1(\xi) > f_2(\xi)\}, & f_1^{-1}(y) > f_2^{-1}(y) \end{cases}$$

(b) $F_{s^{(y,z)}}^+ = F_{s^{(y,z)}}^-$

Proof. Follows directly from the expression for the derivative given in Lemma 4. □

Obviously, $u(y,z) \geq 0$ for any y and z. As we assumed, statements (g) and (h) of Lemma 1 hold for all $t \geq 0$ including $t = 0$. Therefore, $f_i^{-1}(y)$, $i = 1,2$, are continuous nondecreasing functions. This and Lemma 7 imply that $u(y,z)$ is a continuous function of (y,z).

Theorem 2

$$\begin{aligned} F_s^+ &\leq -\int_0^1 q(y)dy \int_0^\infty r(z)dz\ u(y,z) \\ F_s^+ &= F_s^- \end{aligned}$$

Proof. For every $n = 1,2,3,\ldots$, let us define a function s_n as follows. Denote,

$$\begin{aligned} y_i &= \frac{i}{2^n}, \quad i = 0,1,\ldots,2^n \\ z_j &= \frac{j}{2^n}, \quad j = 0,1,2,\ldots,2^{2n}-1 \\ z_{2^{2n}} &\equiv \infty \\ \sigma_{ij} &= \begin{cases} s^{(y_i,z_j)}, & \text{if } i \geq 1 \\ s^{(y_1,z_j)}, & \text{if } i = 0 \end{cases} \end{aligned}$$

and, for $i = 0,1,\ldots,2^n-1$, $j = 0,1,\ldots,2^{2n}-1$,

$$\alpha_{ij} \equiv q(y_i)(y_{i+1}-y_i)\left(\int_{z_j}^{z_{j+1}} r(z)dz\right).$$

Let

$$s_n = \sum_{i=0}^{2^n-1} \sum_{j=0}^{2^{2n}-1} \alpha_{ij}\sigma_{ij}$$

It is easy to verify, that

$$s_n \xrightarrow{L_1} s$$

Lemma 5 implies that

$$F_{s_n}^+ \to F_s^+ \quad \text{and} \quad F_{s_n}^- \to F_s^- \tag{13}$$

It follows from Lemma 6 and Lemma 7 that

$$F_{s_n}^+ = F_{s_n}^-$$

which means, particularly, that

$$F_s^+ = F_s^- \leq 0$$

From Lemma 6 we get

$$F_{s_n}^+ \leq \sum_{i,j} F_{\sigma_{ij}}^+$$

where all $F_{\sigma_{ij}}^+$ are nonpositive. So,

$$\limsup_{n\to\infty} F_{s_n}^+ \leq \limsup \sum_{i=1}^{2^n-1} \sum_{j=0}^{2^{2n}-1} \alpha_{ij} F_{\sigma_{ij}}^+$$
$$= \limsup \left[-\int_0^1 q(y)dy \int_0^\infty r(z)dz \; u_n(y,z) \right] \quad (14)$$

where $u_n(y,z) =$

$$\begin{cases} u(y_i, z_j), & \text{if } \; y_i \leq y < y_{i+1}, \; z_j \leq z < z_{j+1}, \; i \geq 1 \\ 0, & \text{if } \; y_0 \leq y < y_1 \end{cases}$$

Continuity of the function $u(y,z)$ implies that

$$u_n(y,z) \to u(y,z)$$

for all $0 < y < 1$ and $0 \leq z < \infty$. Therefore, applying Fatou's lemma to the right-hand side of (14) we get

$$\limsup F_{s_n}^+ \leq -\int_0^1 q(y)dy \int_0^\infty r(z)dz \; u(y,z) \quad (15)$$

Convergence (13) and inequality (15) prove the theorem. □

4.6 Convergence to a Traveling Wave

Theorem 2 provides a key to the proof of our main result, Theorem 1. Indeed, it shows that as long as two solutions to (2) with the same "mean value" $m(f(\cdot,t))$, do not coincide, the L_1-distance between them is strictly decreasing. This, by the way, implies the uniqueness of the traveling wave solution (with fixed "mean value" at time 0).

Let us apply Theorem 2 to the functions

$$f_1(x,t) \equiv f(x,t)$$

and

$$f_2(x,t) \equiv \tilde{f}(x,t) \equiv \phi(x - vt)$$

Both functions are solutions to (2) with zero "mean value" at time 0:

$$m(f(\cdot,0)) = m(\tilde{f}(\cdot,t)) = 0 \quad (16)$$

Denote

$$\widehat{f}(x,t) \equiv f(x+vt,t)$$

Function $\widehat{f}(x,t)$ is a solution $f(x,t)$ being shifted to the left with speed v to keep its mean value equal to 0.

We have to prove that for any $x \in R$

$$\widehat{f}(x,t) \to \phi(x), \quad t \to \infty$$

Observe that $\widehat{f}(x,t)$ is a family of nondecreasing functions. We have

$$\|\widehat{f}(\cdot,t) - \phi(\cdot)\| < \infty$$

due to the "integrable tails" condition and the fact (Theorem 2) that $\|\widehat{f}(\cdot,t) - \phi(\cdot)\|$ is a nonincreasing function. This implies that if each function $\widehat{f}(x,t)$ is considered as a distribution function (in the probabilistic sense), then the family $\widehat{f}(\cdot,t)$ is relatively compact. This, in turn, implies that for any infinitely increasing sequence of time instants

$$t_1 < t_2 < \cdots < t_n < \cdots$$

there is a subsequence t_{n_k}, $k = 1, 2, \ldots$, such that

$$\widehat{f}(x, t_{n_k}) \to \overline{f}(x)$$

where $\bar{f}(x)$ is a nondecreasing right continuous function, $\bar{f}(-\infty) = 0$, $\bar{f}(\infty = 1$, and the convergence takes place in any point x where $\bar{f}(x)$ is continuous.

Lemma 8 *Any limiting function $\bar{f}(x)$ is a shifted function $\phi(x)$, i.e., there exist a constant a such that*

$$\bar{f}(x) = \phi^{(a)}(x) \equiv \phi(x-a) \quad (17)$$

Proof. Suppose, (17) is not true. Then we can always find a constant a such that $\bar{f}(x) - \phi^{(a)}(x)$ is positive for some x and negative for others. To be more definite, suppose there exist $x_1 < x_2$ such that

$$\bar{f}(x_1) > \phi^{(a)}(x_1), \quad \bar{f}(x_2) < \phi^{(a)}(x_2)$$

and $\bar{f}$ is continuous in points x_1 and x_2. This means, that there exist $\epsilon > 0$ and $\delta > 0$ such that the Lebesgue measure of time instants such that

$$\widehat{f}(x_1,t) > \phi^{(a)}(x_1) + \epsilon \quad (18)$$

and

$$\widehat{f}(x_2 - \delta, t) < \phi^{(a)}(x_2 - \delta) - \epsilon \quad (19)$$

is infinite. But it easily follows from Theorem 2 that if at time t both (18) and (19) are true, then the derivative F_s^+ is separated from 0 by a negative constant:

$$F_s^+ \leq \eta < 0$$

This would mean that

$$\|\widehat{f}(\cdot,t) - \phi^{(a)}(t)\| \downarrow -\infty$$

which is, of course, impossible. □

Lemma 9 *Any limiting function* $\bar{f}(x)$ *coincides with* $\phi(x)$*:* $\bar{f}(x) \equiv \phi(x) \equiv \phi^{(0)}(x)$.

Proof. Suppose, it's not true. Then there exists a limiting function

$$\bar{f}(x) = \phi^{(a)}(x)$$

where, for example, $a < 0$. This easily implies, that

$$\lim_{t\to\infty} \|(\widehat{f}(\cdot,t) - \phi(x))^{+}\| \geq |a|$$

Therefore,

$$\lim \|(\widehat{f}(\cdot,t) - \phi(x))^{-}\| \geq |a| \tag{20}$$

because $m(\widehat{f}(\cdot,t)) = m(\phi(\cdot)) \equiv 0$. From the fact that $\bar{f}(x)$ is a limiting function and inequality (20), we see that for any $\epsilon > 0$, there exists a sufficiently big $t > 0$ such that

$$\int_{1-\epsilon}^{1} (\widehat{f}^{-1}(y,t) - \phi^{-1}(y))^{+} dy \geq |a| > 0 \tag{21}$$

This means that for any $b > 0$,

$$\|(\widehat{f}(\cdot,t) - \phi^{(b)}(\cdot))^{-}\| \geq |a| \tag{22}$$

But inequality (22) is impossible, because we always can find a sufficiently big $b > 0$ such that

$$\|(\widehat{f}(\cdot,0) - \phi^{(b)}(\cdot))^{-}\| < |a|$$

□

The statement (6) of Theorem 1 has been proven. Suppose, (6) is true, but (7) is false. It may happen only if statement (21) is true for some positive number $|a|$. But this is impossible as it is shown in the proof of Lemma 9.

Theorem 1 has been proven.

References

[1] S. Yu. Popov A.G. Greenberg, V.A. Malyshev. Stochastic model of massively parallel computation. To appear in Markov Processes and Related Fields, 1996.

[2] R. E. Felderman and L. Kleinrock. Bounds and approximations for self-initiating distributed simulation without lookahead. *ACM Transactions on Modeling and Computer Simulation*, 1(4):386–406, 1991.

[3] A. G. Greenberg, B. D. Lubachevsky, D. M. Nicol, and P. E. Wright. Efficient massively parallel simulation of dynamic channel assignment schemes for wireless cellular communications. In *Proceedings of the 8th Workshop on Parallel and Distributed Simulations (PADS '94)*, 1994.

[4] D.R Jefferson. Virtual time. *ACM Trans. on Programming Languages and Systems*, 7(3):404–425, July 1985.

[5] B. D. Lubachevsky. Efficient distributed event-driven simulations of multiple-loop networks. In *ACM Sigmetrics conference on measurement and modelling of computer systems*, pages 12–21, May 1988.

[6] D. Nicol. Performance bounds on parallel self-initiating discrete-event simulations. *ACM Transactions on Modeling and Computer Simulation*, 1(1):24–50, 1991.

[7] D. M. Nicol and R. M. Fujimoto. Parallel simulation today. *Annals of Operations Research*, 1994.

[8] D. M. Nicol and P. Heidelberger. Parallel simulation of markovian queueing networks using adaptive uniformization. In *Proc. ACM SIGMETRICS conf. on measurement and modeling of comp. systems*, pages 135–145, May 1994.