THE BRANCHING DIFFUSION APPROXIMATION FOR A MODEL OF A SYNCHRONIZED QUEUEING NETWORK

N.G. Duffield
AT&T Laboratories,
Room 2C-323, 600 Mountain Avenue, Murray Hill, NJ 07974, USA.
E-mail: duffield@research.att.com

M. Kelbert
The European Business Management School, University College Swansea, UK, and International Institute for The Earthquake Prediction Theory and Mathematical Geophysics, The Russian Academy of Sciences, Moscow, Russia.
E-mail: m.kelbert@swansea.ac.uk

Yu.M. Suhov
Statistical Laboratory, DPMMS, University of Cambridge and St John's College, Cambridge, UK and Institute for Problems of Information Transmission, The Russian Academy of Sciences, Moscow, Russia.
E-mail: y.m.suhov@statslab.cam.ac.uk

We consider a model of a centralized network under arrival-synchronization constraints. The mean-field and diffusion approximation limits are performed and it is shown that the limiting waiting time process is described in terms of a branching diffusion. Relations with travelling-wave solutions to a non-linear PDE are also established.

1 Introduction

In this paper we consider a queueing network model under an arrival-synchronization constraint. The model leads to a higher-order Lindley equation [11]. The network has N input nodes $I_1, \ldots, I_N$ which are fed with independent Poisson flows $\xi_1, \ldots, \xi_N$ of intensity λ. All tasks arriving in the input flows ξ_j have i.i.d. service times with some distribution F. In addition, the tasks are provided with i.i.d. labels indicating a 'type". The labels are assumed to be independent on the lengths. The labels are pairs $(0, k)$ and $(1, k)$, $k = 1, \ldots, N$, the value $(0, k)$ being taken with probability $(1-p)/N$ and $(1, k)$ being taken with probability p/N.

We describe the network operation recursively as follows. A task arriving in flow ξ_j waits first until all tasks arrived earlier in this flow have left the network. Then at this time if the task has label $(0, k)$ it is serviced (for the duration of its service time), and leaves the network thereafter. If its label is $(1, k)$, the task also has to wait until all tasks with labels $(0, k)$ and $(1, k)$ which arrived earlier in any of the flows $\xi_{j'}$, $j' \neq j$, have left the network. Only after this further waiting period is the task serviced; then it leaves the network.

Tasks with label $(0,k)$ can be regarded as having being privileged in the sense that before being serviced they only wait for tasks with earlier arrival times in their *input* streams to leave the network. Let us call all other tasks unprivileged.

The network rule just described may be classified as an arrival-synchronization constraint. Models this sort arise naturally in parallel processing when a network deals with programs of several types which are processed in different nodes. Some types of program require information that becomes available only after the processing of all programs of a certain type (perhaps the same type) which were created earlier in all other nodes. Furthermore, programs created at a given node must be processed in the order of their creation. A review of properties of networks with synchronization constraints may be found in [4].

To be more specific, we give below a concrete example of the network that leads to this model. This example is related to a star-like network where, in addition to the input nodes $I_1, \ldots, I_N$, there is a central node, C, and a collection $O_1, \ldots, O_N$ of the output nodes.

One may think of the tasks as programmes being processed at the input nodes in which they are created. However, the results of processing the programs with labels $(0,k)$ and $(1,k)$ are collected in node O_k (we assume that the results are transmitted through the network without delay). A program with label $(0,k)$ starts processing as soon as all programs created at the same node earlier are terminated. A program with label $(1,k)$ has in addition to wait until all programs labelled $(0,k)$ or $(1,k)$ and created earlier, no matter in which node, are terminated. The running time of a program is, as before, equal to its service time.

A specific feature of the network under consideration is that some tasks with label $(1,k)$ have two predecessors. One of these predecessors is the previous task to arrive in the same flow ξ_j; the other is the previous task with the same label to arrive at any input node. (Sometimes these two predecessors are from the same flow ξ_j, or may even coincide). Each of these predecessors may have its own pair of predecessors, and so on. This creates a rather intricate picture, and all we can do is to establish sufficient conditions under which the network possesses stationary regimes and prove that these regimes are, in some sense, unique.

However, there are some natural approximation schemes available which allow us to simplify the probabilistic picture for both models. In particular, we study in this paper an asymptotic distribution of the random process representing the overall waiting times for the whole sequence of tasks arrived at a single node I_j (for definiteness, we consider below the input node I_1). More precisely, we perform two limiting procedures: (i) $N \to \infty$ and (ii) $\lambda \to \infty$, $F \Rightarrow \delta$ (δ is the Dirac measure) and $p \to 0$, with λ^2 Var s, $\lambda^{3/2}(\lambda^{-1} - \mathbf{E}s)$ and λp tending to finite (positive) limits. Here and below $\mathbf{E}$ and Var denote the expectation and the variance (we omit the indication of the probability distribution if there is no confusion about this). The first limiting procedure is known as the mean-field limit (see [3,7,11]); for the model under consideration it results in a specific decoupling phenomenon between the nodes which in turn leads to what was called in [11] a second-order Lindley equation. The second procedure represents the well-known diffusion approximation (see, e.g. [2], for the model under consideration. As a result of these procedures, we pass to a description in terms of branching diffusion processes where the independently moving Brownian particles occasionally give birth to the new generations which then proceed in a like manner.

Branching diffusion processes (in their simplest version) were studied in [17] in connection with the Kolmogorov–Petrovskii–Piskunov (KPP) equation [16], a non-linear PDE that plays an important role in various applications (see, e.g [10]]). This process has been intensively studied since (see, e.g. [5,6,9] and the references therein), but to our knowledge a branching diffusion process only recently appeared directly in a queueing network context (cf. [8]). Our result relates the virtual waiting time distribution function in an overloaded network with the travelling wave solution to the KPP equation.

2 The Models, Results and Main Probabilistic Constructions

2.1 *The models: a formal description.*

As was said, the random times $t_n^{(j)}$, $n = 0, \pm 1, \ldots$, of tasks' arrivals at the input node I_j, $j = 1, \ldots, N$, form a Poisson flow ξ_j of intensity λ (as usually, $t_0^{(j)}$ denotes the first point from ξ_j on $\mathbf{R}_+$, the non-negative half-axis). We can think of ξ_j as a marked random flow, with i.i.d. marks $(s_n^{(j)}, a_n^{(j)})$ where $s_n^{(j)}$ is the length and $a_n^{(j)}$ is the label of the n^{th} task arrived in ξ_j. The components $s_n^{(j)}$ and $a_n^{(j)}$ are also independent, and $s_n^{(j)}$ takes values in $\mathbf{R}_+$ according to the distribution F. The label $a_n^{(j)}$ takes the values $(0, k)$ and $(1, k)$, $k = 1, \ldots, N$, with probabilities $(1-p)/N$ and p/N, respectively.

We therefore can form, for each $k = 1, \ldots, N$, a flow η_k composed by the times of arrival of the tasks with labels $(0, k)$ and $(1, k)$ regardless of the input node of each task. Let us denote the points from flow η_k by $\bar{t}_m^{(k)}$, $m =$

$0, \pm 1, \pm 2 \ldots$. We associate elements of the flows as follows. With a task arriving in input flow ξ_j at time $t_n^{(j)}$ with labels $(0,k)$ and $(1,k)$, $k \geq 1$, one can associate an integer $m = m(n,j,k)$ indicating the place of epoch $t_n^{(j)}$ in the realization of the corresponding flow η_k. In other words,

$$t_n^{(j)} = \bar{t}_m^{(k)}, \qquad \text{provided that} \qquad a_n^{(j)} = (\cdot, k). \tag{1}$$

We can also 'transfer' to the flow η_k the service times of the corresponding tasks. The service time of a task that has appeared in flow η_k at time $\bar{t}_m^{(k)}$ is denoted by $\bar{s}_m^{(k)}$:

$$\bar{s}_m^{(k)} = s_n^{(j)} \quad \text{if} \quad \bar{t}_m^{(k)} = t_n^{(j)}. \tag{2}$$

This enables us to think of the η_k's as marked random flows. As follows from the description provided above, flows η_k are independent marked Poisson flows of intensity λ with i.i.d. marks. Of course, the η_k's are functions of the ξ_j's, in the sense that flows η_k are defined on the probability space Ω formed by the collections of the realizations of flows $\xi_1, \ldots, \xi_N$, with the probability distribution $\mathcal{P}_N$ that is the direct product of the distributions of flows ξ_j.

Denote by $w_n^{(j)}$ (respectively $\bar{w}_m^{(k)}$) the waiting time for the task that arrives at time $t_n^{(j)}$ in flow ξ_j (respectively, appears at time $\bar{t}_m^{(k)}$ in flow η_k). The network operating rule described in the introduction leads to the following recursive equations

$$w_n^{(j)} = \begin{cases} \max[0, w_{n-1}^{(j)} + s_{n-1}^{(j)} - (t_n^{(j)} - t_{n-1}^{(j)})], \\ \qquad\qquad\qquad \text{if } a_n^{(j)} = (0,k), 1 \leq k \leq N, \\ \max[0, w_{n-1}^{(j)} + s_{n-1}^{(j)} - (t_n^{(j)} - t_{n-1}^{(j)}), \\ \qquad \bar{w}_{m-1}^{(k)} + \bar{s}_{m-1}^{(k)} - (\bar{t}_m^{(k)} - \bar{t}_{m-1}^{(k)})], \\ \qquad\qquad\qquad \text{if } a_n^{(j)} = (1,k), 1 \leq k \leq N. \end{cases} \tag{3}$$

Our first result is about the existence of a stationary regime in the model under consideration. Sufficient condition that becomes asymptotically tight in the mean-field limit (see below) is given in Proposition 2.1. It is convenient to introduce a (2×2) matrix

$$\mathbf{A}(a) = \begin{pmatrix} \dfrac{\lambda(1-p)}{\lambda + a}\mathbf{E}e^{as} & \dfrac{2\lambda p}{\lambda + a}\mathbf{E}e^{as} \\ \dfrac{\lambda(1-p)}{\lambda + a}\mathbf{E}e^{as} & \dfrac{2\lambda p}{\lambda + a}\mathbf{E}e^{as} \end{pmatrix}. \tag{4}$$

Here $a > 0$ is a parameter. The matrix $\mathbf{A}(a)$ reflects a specific character of the random processes describing the stationary regime: it has a transparent meaning in the context of a random field on a Cayley tree-graph whose connection with the current problem we elaborate below. The norm of $A(a)$ is equal to

$$\frac{\lambda}{\lambda + a}(1+p)\mathbf{E}e^{as}. \tag{5}$$

Proposition 2.1 *For all N, under the condition*

$$\inf_{a>0} ||\mathbf{A}(a)|| < 1 \tag{6}$$

the system of equations (3) has, for $\mathcal{P}_N$-almost every array $\{t_n^{(j)}, s_n^{(j)}, a_n^{(j)}\}$, a solution that is unique in a class of the sequences $\{w_n^{(j)}\}$ satisfying the following condition of boundedness (in probability):

$$\lim_{y\to\infty} \sup_n \mathcal{P}_N\ (w_n^{(j)} > y) = 0. \tag{7}$$

If we concatenate the components $w_n^{(j)}$ to pairs $(s_n^{(j)}, a_n^{(j)})$ constituting the marks in flows ξ_j, the resulting family of extended flows $\bar{\xi}_j$, $1 \le j \le N$, is stationary in time.

2.2 The mean-field approximation: limiting stochastic equations.

The extended families $\{\bar{\xi}_j\}$ mentioned in Proposition 2.1 describe the working regimes in our network: any network characteristic may be, in principle, found from the joint probability distributions of flows $\bar{\xi}_1, \ldots, \bar{\xi}_N$, which we denote by $\bar{\mathcal{P}}$ ($\bar{\mathcal{P}}_N$). However, the trouble with distribution $\bar{\mathcal{P}}$ is that it contains a lot of dependences (it is not the product of the marginal distributions of flows $\bar{\xi}_j$) and does not provide any opportunity for a (non-trivial) explicit calculation. To simplify we take the limit $N \to \infty$ (see [11]). Under our assumptions, this limit is analogous to the mean-field limit in Statistical Physics in the sense that in this limit individual tasks wait for service in an independent background created by other tasks in the network. The limiting stationary distribution of $w^{(j)}$, for a fixed j (say $j = 1$), exists and is identified as a minimal solution to the following stochastic equations

$$w \simeq \max[0, w' + s' - l', \chi(w'' + s'' - l'')]. \tag{8}$$

Here $\simeq$ means equality in law, and the random variables w' and w'' in the RHS of (8) have the same distribution as w. We denote this distribution which solves (8) by Φ. We also denote by Φ the corresponding distribution function. Further on s' and s'' in (8) have the distribution F, l' and l'' are exponentially distributed with mean λ^{-1}, and χ takes value zero with probability $1-p$ and one with probability p. Finally, all random variables in the RHS of (8) are independent.

It turns out (see [13]) that, under the condition (6) the solution to the equations (8) exists, but is non-unique, however small λ or $\mathbf{E}s$. (Actually, condition $||\mathbf{A}(a)|| \leq 1$ is sufficient for this. It is also necessary (see below)). Under condition (6), there exists a linearly ordered continuum of solutions of (8), with different asymptotics of $\Phi(x)$ as $x \to \infty$. (Here we mean the standard stochastic order: for random variables u' and u'' with distribution functions Φ', Φ'' the partial order $u' \preceq u''$ means $\Phi'(x) \geq \Phi''(x)$, $x \in \mathbf{R}^1$). There exists a unique minimal solution, Φ^0, corresponding to the 'zero initial condition' that initial waiting times are zero:

$$w_0^{(j)} = 0, 1 \leq j \leq N. \tag{9}$$

If on the other hand $||A(a)|| > 1$ for all $a > 0$ then equation (8) has no solution in the set of proper distribution functions. However, equation (3) may be iterated from some initial condition (say the zero condition). This gives us a family of extended flows $\bar{\xi}_j^+$, $1 \leq j \leq N$, which are defined on the non-negative half-axis. The limit $N \to \infty$ may also be performed in this case.

Let us concentrate now on the structure of the sequence $w_n^{(j)}$ for a fixed value of j. We consider the case $j = 1$ and omit, for simplicity, index 1 writing ξ and $\bar{\xi}$ instead of ξ_1 and $\bar{\xi}_1$, ξ^+ and $\bar{\xi}^+$ instead of ξ_1^+ and $\bar{\xi}_1^+$, etc). Consider first the case where the conditions of Proposition 2.1 are fulfilled. As was said before, the extended flow $\bar{\xi}$, with marks $(s^{(j)}, a^{(j)})$, describes, for a given N, the stationary regime in the input node I_1. It is convenient to pass to the reduced labels of the tasks arrived, by setting $a_n^* = 0$, if $a_n = (0, k)$ and $a_n^* = 1$, if $a_n^* = (1, k)$, $1 \leq k \leq N$. The corresponding reduced flows with marks (s_n, a_n^*, w_n) may be denoted as ξ^* and $\bar{\xi}^*$, respectively. The new components of the mark a_n^* are (conditionally) i.i.d. and take their values 0 and 1 with probability $1 - p$ and p, respectively. Our results about the limit $N \to \infty$ for the model under consideration are summarized in the Theorem 1 following. Hereafter the notion of covergence we use is of a vague convergence of random processes, i.e. the convergence of the expectation values of bounded continuous local functions on the realization space of the process. The precise definition depends on the topology on this space, and we formally introduce it when we prove the corresponding theorem.

Theorem 1 *Under condition (6), the extended flow $\bar{\xi}^*$ converges, as $N \to \infty$, to a limiting flow Ξ^* which is characterised in an unique way by the following conditions:*

(i) The times t_n and the components of the mark s_n and a_n^*, $n = 0, \pm 1, \ldots$, are distributed as in flow ξ^*.

(ii) *Given a sequence* $\{(t_n, (s_n, a_n^*)), n = 0, \pm 1, \ldots\}$, *the sequence* $\{w_n, n = 0, \pm 1, \ldots\}$ *satisfies the recursive equations*

$$w_n = \begin{cases} \max\,[0, w_{n-1} + s_{n-1} - (t_n - t_{n-1})], & \text{if } a_n^* = 0, \\ \max\,[0, w_{n-1} + s_{n-1} - (t_n - t_{n-1}), \widetilde{w}_n + \widetilde{s}_n - L_n], & \\ & \text{if } a_n^* = 1, \end{cases} \tag{10}$$

where $\widetilde{w}_n, \widetilde{s}_n, L_n$, $n = 0, \pm 1, \ldots$, *are independent random variables,* $\widetilde{w}_n$ *has distribution* Φ^0, $\widetilde{s}_n$ *has distribution* F *and* L_n *is exponentially distributed with mean* λ^{-1}.

We can split the distribution Φ^0 of the minimal solution w to (8) into two parts: $\Phi^0 = (1-p)\Phi_1^0 + p\Phi_2^0$ where, with the notation of (8), Φ_1^0 is the distribution of the random variable $\max[0, w + s - l]$ and Φ_2^0 the distribution of $\max[0, w' + s' - l', w'' + s'' - l'']$. Then the marginal (conditional) distribution of the component of the mark w_n in the limiting flow in Theorem 1 is Φ_1^0 when $a_n^* = 0$ and Φ_2^0 when $a_n^* = 1$. (The precise formulation of this assertion requires the use of a Palm distribution. This will become more transparent when we pass to the diffusion approximation).

2.3 Random Processes on Random Trees

We now associate to solutions of (8) random processes on random trees. This will be used in proving the validity of mean-field and diffusion limits (see Theorems 2–4 below). Consider the Cayley tree rooted at an origin $\mathbf{O}$, and with branching ratio 2. Each edge e can be identified with a binary sequence, which also identifies its vertex furthest from the origin. (The origin is then identified with the null sequence). Denote by $\mathcal{L}^0$ the set of finite continuous paths starting at the origin: clearly each path $\mathbf{L} \in \mathcal{L}^0$ is identified with its final edge (or equivalently its final vertex) and hence is given again by a finite binary sequence. In a *random* tree on the other hand the branching ratio is random at each vertex, but maximally 2. We use various *equipments* of trees, these being the association of random variables (i.e. marks) with each of the edges of the tree.

We now associate with the solution to equation (8) a random process $\{w_e\}$ on the edges (or equivalently the vertices) of a random tree which possesses a natural stationarity property. The basic idea here is that each edge is associated with a task, and that adjacent edges above (or offspring) correspond to the predecessors of this task. Hence in this case each edge of the tree spawns one offspring with probability $1-p$ and two with probability p, independently of other edge. To each edge e we assign a pair of random variables l_e, s_e where s_e has distribution F and l_e has the exponential distribution with mean λ^{-1}.

The assumption of the (conditional) independence of the whole array l_e, s_e is of course preserved.

We specify the random variables $\{w_e\}$ to be related by a recursive system of equations similar to (3) above, namely

$$w_e = \begin{cases} \max[0, w_{\tilde{e}} + s_{\tilde{e}} - l_{\tilde{e}}] & \text{if } e \text{ has one offspring,} \\ \max[0, w_{\tilde{e}'} + s_{\tilde{e}'} - l_{\tilde{e}'}, w_{\tilde{e}''} + s_{\tilde{e}''} - l_{\tilde{e}''}], & \\ & \text{if } e \text{ has two offspring.} \end{cases} \tag{11}$$

Then minimal solution Φ^0 will be identified with the distribution of the quantity

$$\max[0, \sup_{\mathbf{L} \in \mathcal{L}^0} \sum_{e \in \mathbf{L}} (s_e - l_e)]. \tag{12}$$

In the random-tree picture, then the component Φ_1^0 gives the conditional distribution under the condition that the initial vertex **O** creates one offspring, and the component Φ_2^0 the conditional distribution under the condition that **O** creates two offspring.

2.4 The diffusion approximation.

The diffusion approximation we consider in this paper is related to the process of the virtual waiting time $v(\tau)$, $\tau \in \mathbf{R}$, generated by the queueing in our input node I, in the limiting extended flow Ξ^*:

$$v(\tau) = \max[0, w_{\widehat{n}} - (\tau - t_{\widehat{n}})] \tag{13}$$

where $\widehat{n}$ is defined by the requirement that $t_{\widehat{n}}$ is the first arrival time in flow ξ^* which precedes τ. Under the conditions of Theorem 1, this is a stationary, continuous-time regenerative process. The resulting process is complex (although parameters may be determined from Theorem 1). We obtain a simpler picture in terms of branching diffusion by making an approximation in which the intensity of the random flow $\bar{\xi}^*$ (or ξ^*) becomes large and the random variables s_n become negligibly small.

Let us first describe the limiting diffusion. Consider a family of Brownian particles, each particle having a type 1 or 2 and moving independently on a line. The drift coefficient of a particle of type i is μ_i and its diffusion coefficient is σ_i, $i = 1, 2$. After independent exponential lifetimes, with the mean values γ_1^{-1} and γ_2^{-1}, respectively, the particles die, each giving birth to two offspring, one of type 1 and one of type 2, The descendants then proceed independently according to their type. From [12] we have the following

Proposition 2.2 *Let X denote the supremum of the deviation, in the positive direction, of all particles taken over all times of particles' death. A sufficient condition for X to have a proper distribution is:*

$$\inf_{0<b<b^*}\left(\frac{\gamma_1}{\gamma_1-\mu_1 b-\frac{\sigma_1^2}{2}b^2}+\frac{\gamma_2}{\gamma_2-\mu_2 b-\frac{\sigma_2^2}{2}b^2}\right)<1, \tag{14}$$

where b^* is defined as the supremum of those values of $\widetilde{b}>0$ for which both denominators in the expression under the infimum are positive for any $b\in(0,\widetilde{b})$.

We shall obtain these diffusions by taking the following limit:

$$\lambda\to\infty,\ \lambda^2\mathrm{Var}_F\ s\to r\geq 0,\ \lambda^{3/2}(\mathbf{E}_F s-\lambda^{-1})\to\varphi\in\mathbf{R},\ \lambda p\to\gamma>0 \tag{15}$$

subject to the technical condition

$$\mathbf{E}_F s^{2+\alpha}\leq C, \tag{16}$$

with fixed constants $\alpha, C\in(0,\infty)$. In our model we will set $\mu_1=\mu_2=\varphi$, $\gamma_1=\gamma_2=\gamma$ and $\sigma_1=\sigma_2=\sigma:=(1+r)^{1/2}$, condition (14) becoming that

$$\varphi<0\ ,\qquad \varphi^2/\sigma^2>2\gamma\qquad . \tag{17}$$

We actually need not distinguish between the different types of particle due to the identify of their parameters.

We can represent the branching diffusion as a random process, now on a *non-random* rooted Cayley tree with fixed branching ratio 2. Roughly speaking, the correspondence with the random trees in the previous section is that the vertices of the fixed tree correspond to those vertices on the random tree where branching with ratio 2 actually takes place. We equip each edge e with an independent random pair (g_e, S_e), where g_e is distributed exponentially with mean γ^{-1}, while S_e is a Wiener trajectory on $[0, g_e)$ with drift coefficient φ and diffusion coefficient σ.

The distribution of the random quantity

$$\max\,[\,0,\ \sup_{\mathbf{L}\in\mathcal{L}^0}\sum_{e\in\mathbf{L}}S_e(g_e)\,] \tag{18}$$

may be identified with the distribution of the variable X of Proposition 2.2 (which justifies the use of the same symbol X in the notation). The distribution of X is denoted by $\bar{\Phi}^0$. As before, we have a pair of distributions, $\bar{\Phi}^0_1$ and $\bar{\Phi}^0_2$ according to whether we have one or two particles present initially.

Physically speaking, we perform in (18) an operation of 'glueing' the trajectories S_e along a path $\mathbf{L}$ and then compute the final displacement from the origin.

In fact we only need to know the second component, since if for example X_2 has distribution $\bar{\Phi}_2^0$, then

$$max[0, X_2 + S] \tag{19}$$

has distribution $\bar{\Phi}_1^0$, where S is the displacement of independent Wiener particle, with the drift coefficient φ and the diffusion coefficient σ, at the random stopping time that is exponentially distributed, with mean γ^{-1}.

As before, we can write stochastic equations analogous to (10) or, in the stationary form, to (8). The analogue of the second equation in (10) is

$$W_\beta = \max\,[\,0, W_{\beta'} + S(g_{e'}), W_{\beta''} + S(g_{e''})\,], \tag{20}$$

and the analogue of (8) is

$$W \simeq \max\,[\,0, W' + S(g'),\ W'' + S(g'')\,]; \tag{21}$$

we retained here the same rules of the notation and the conventions about the distributions as in (9) and (8). Under condition (17) the minimal solution $\bar{\Phi}_2^0$ to (21) is unique and is a proper probability distribution. However, there exists again a linearly ordered continuum of solutions distinguishable in their behavior at infinity (see [12]).

Our result on the diffusion approximation is as follows:

Theorem 2 *Suppose that λ, F and p are varied as indicated in (15), so that conditions (6), (16) and (17) hold true. Then, the random process $\{v(\tau), \tau \in \mathbf{R}^1\}$, converges to a stationary process, $\{V(\tau),\ \tau \in \mathbf{R}^1\}$. The distribution of $\{V(\tau),\ \tau \in \mathbf{R}^1\}$ is the marginal distribution of the second component in a pair $(\Theta, \{\widetilde{V}(\tau),\ \tau \in \mathbf{R}^1\})$, where Θ is a stationary Poisson process of intensity γ. The conditional distribution of process $\{\widetilde{V}(\tau),\ \tau \in \mathbf{R}^1\}$, given a realization $\{\tau_n^*, n \in \mathbf{Z}\}$, of process Θ, is characterized as follows.* **(a)** *The trajectories of the conditional process (which we denote by $\{\widetilde{V}(\tau|\{\tau_n^*, n \in \mathbf{Z}\}$) are continuous on $\mathbf{R}^1 \setminus \{\tau_n^*, n \in \mathbf{Z}\}$; on $\{\tau_n^*, n \in \mathbf{Z}\}$ they are right-continuous and have left limits;* **(b)** *The process $\{\widetilde{V}(\tau|\{\tau_n^*, n \in \mathbf{Z}\})\}$ possesses the following Markov cyclicity property: for any $j \in \mathbf{Z}$, the conditional distribution of the random variables $\widetilde{V}(\tau|\{\tau_n^*, n \in \mathbf{Z}\})$, $\tau \in (\tau_j^*, \tau_{j+1}^*)$, given the past $\widetilde{V}(\tau|\{\tau_n^*, n \in \mathbf{Z}\})$, $\tau < \tau_j^*$, depends on the value $\widetilde{V}(\tau_j^*|\{\tau_n^*, n \in \mathbf{Z}\})$ only, and is given by*

$$\widetilde{V}(\tau|\{\tau_n^*, n \in \mathbf{Z}\}) = \max\,[\sup_{\tau_j^* < t \leq \tau} S(\tau - t), \widetilde{V}(\tau_j^*|\{\tau_n^*, n \in \mathbf{Z}\}) + S(\tau - \tau_j^*)], \tag{22}$$

where $S(u)(= S^{(\tau^*)}(u))$, $u \geq 0$, is the time-homogeneous Wiener process with the drift coefficient φ and the diffusion coefficient σ, independent on $\widetilde{V}(\widetilde{\tau}|\{\tau_n^*, n \in \mathbf{Z}\})$, $\widetilde{\tau} \leq \tau_j^*$. c) *The values $\widetilde{V}(\tau_j^*|\{\tau_n^*, n \in \mathbf{Z}\})$, $j \in \mathbf{Z}$, at the initial points of the cycles are given by*

$$\widetilde{V}(\tau_j^*|\{\tau_n^*, n \in \mathbf{Z}\}) = \max\,[\widetilde{V}(\tau_j^* - 0|\{\tau_n^*, n \in \mathbf{Z}\}), Y_j], \tag{23}$$

where Y_j are i.i.d random variables, and each variable Y_j has distribution $\bar{\Phi}_2$ emerging in the model of the branching diffusion.

The distribution of $V(\tau)$ at a (non-random) time τ coincides with $\bar{\Phi}_1^0$.

2.5 The half-axis picture.

If we do not assume the condition (6) to hold, we are able to establish similar pictures on the non-negative half-axis $[0, +\infty)$. It is again convenient to pass to the reduced flows ξ^{+*} and $\bar{\xi}^{+*}$. Recall that initial condition (9) is supposed to be fulfilled. In order to describe the situation, we have to introduce 'cut-off' analogues of the constructions of the processes on trees used before.

We want to construct a 'piece' of a (random) tree which has been grown during a given time $\widetilde{t} > 0$. Take a realization of the tree and equip it with the pairs of the random variables (l_e, s_e) according to the rule described above. Given a collection of the values $\{l_e\}$, consider a part of the tree confined to those vertices β for which the sum $\sum_{\widetilde{e} \in \mathbf{L}_\beta} l_{\widetilde{e}} < \widetilde{t}$. Here $\mathbf{L}_\beta$ denotes the path from $\mathcal{L}^0$ leading to vertex β. Assign to the 'border' vertices the value $w_\beta = 0$ and compute the values w_β for the preceding vertices via equations (10). Equivalently, compute the maximum of the sums $\sum_{e \in \mathbf{L}}(s_e - l_e)$ over the paths $\mathbf{L} \in \mathcal{L}^0$ for which $\sum_{e \in \mathbf{L}} l_e < \widetilde{t}$. We obtain a pair of probability distributions $\Phi_1^0(\widetilde{t})$ and $\Phi_2^0(\widetilde{t})$ describing the conditional distributions of the random variable $w_{\mathbf{O}}$ under the condition that the origin $\mathbf{O}$ has, correspondingly, one or two offspring.

An analogue of Theorem 1 in the half-axis picture is:

Theorem 3 *The random flow $\bar{\xi}^{+*}$ converges, as $N \to \infty$, to a limiting flow $(\Xi^+)^*$ on the non-negative time half-axis $\mathbf{R}_+$. The distribution of flow $(\Xi^+)^*$ is characterized, in a unique way, by the following conditions:*

(i) the times t_n and the components of the mark s_n and a_n^*, $n = 0, \pm 1, \ldots$, are distributed as in flow ξ^+;

(ii) given a sequence $\{(t_n, (s_n, a_n^*)), n = 0, 1, \ldots\}$, sequence $\{w_n, n = 1, 2 \ldots\}$ satisfies recursive equations (10) where $\widetilde{w}_n, \widetilde{s}_n, L_n$, $n = 0, \pm 1, \ldots$, are again independent random variables, $\widetilde{w}_n$ now has distribution $\Phi_2^0(t_n)$ and

$\widetilde{s}_n$ and L_n have, as before, distribution F and the exponential distribution with mean λ^{-1}, respectively.

As in the case of Theorem 1, we can say that the marginal (conditional) distribution of the component of the mark w_n in the limiting flow in Theorem 3 coincides with $\Phi_1^0(t_n)$ when $a_n^* = 0$ and with $\Phi_2^0(t_n)$ when $a_n^* = 1$.

Theorem 4 (see below) deals with the process of the virtual waiting time $v^+(\tau)$, $\tau \in [0, \infty)$, generated by the queueing in the input node I under the zero initial condition (9). More precisely, we perform the diffusion approximation limit (15), now without assumptions (17). Again, in order to describe the situation, we have to consider the cut-off versions of the branching diffusion process described before. This may be done on the base of the construction of the corresponding cut-off non-random tree. As before, fix $\widetilde{t} > 0$. Take an infinite tree that starts with two edges and equip it with the pairs (g_e, S_e). Then, given a collection $\{g_e\}$, confine ourselves to those vertices β for which the sum $\sum_{\widetilde{e}\in \mathbf{L}_\beta} g_{\widetilde{e}}$ is $< \widetilde{t}$. Compute the value $W_{\mathbf{O}}$ via the cut-off version of formulas (20). We get a probability distribution $\bar{\Phi}_2^0(\widetilde{t})$. If we modify the construction allowing the tree to start with a single edge, we obtain distributions $\bar{\Phi}_1^0(\widetilde{t})$.

Theorem 4 *Suppose that λ, F and p are varied as indicated in (15). Then the random process $\{v_\tau^+, \tau \geq 0\}$ converges to a (non-stationary) regenerative process, $\{V^+(\tau),\ \tau \geq 0\}$, which is characterized by conditions* a), b) *and* c) *from Theorem 2 and the following condition.*

d′) *At the epoch of regeneration, τ^*, the process takes the value* max $[V^+(\tau^*-),\ Y^+(\tau^*)]$, *where $Y^+(\tau^*)$ is a random variable with distribution $\bar{\Phi}_2^0(\tau^*)$, which is independent on $V^+(\tau)$, $\tau < \tau^*$.*

e′) *The distribution of $V(\tau)$ at a (non-random) time t coincides with $\bar{\Phi}_1^0(\tau)$.*

2.6 A travelling wave bound.

Under the condition opposite to (6)

$$\inf_{a>0} ||\mathbf{A}(a)|| > 1, \tag{24}$$

the network does not possess a stationary regime in the mean-field limit. [An expected picture is that the non-overload condition in a finite network depends on N and is more 'liberal' than (6) (so that, for $N < \infty$, a network may have a unique stationary regime even when (6) is violated), but in the mean-field limit $N \to \infty$ this condition coincides with (6).] In terms of the diffusion approximation, it means that process $V^+(\tau)$ increases to infinity, in the sense that the distribution function $x \in \mathbf{R}_+ \mapsto \mathbf{Pr}\ (V^+(\tau) < x)$ vanishes, as $\tau \to \infty$.

It is interesting to analyze the speed of convergence of $\mathbf{Pr}(V^+(\tau) < x)$ to zero and it is conjectured that a good bound is provided by the travelling wave solution to a non-linear differential equation. This conjecture is based on a connection between the branching diffusion processes and non-linear partial differential equations of the so-called KPP (Kolmogorov-Petrovskii-Piskunov) type. Our model leads to the simplest model of branching diffusion with a single type Wiener process, and to a single KPP equation (see [16])

$$\frac{\partial}{\partial t}u(t,x) + \mu\frac{\partial}{\partial x}u(t,x) = \frac{\sigma^2\partial^2}{2\partial x^2}u(t,x) + \gamma u(t,x)(u(t,x)-1), t \geq 0, x \in \mathbf{R^1}. \tag{25}$$

The following types of solutions to (25) are of interest.

(a) Solutions of the Cauchy problem, with the initial date $u(t,x)]_{t=0} = \theta(x)$; where θ is the indicator function of the non-negative half-axis $\mathbf{R}_+$.

(b) Solutions that possess the travelling wave property: $u(t,x) = u^*(x - ct), x, t \in \mathbf{R}^1$, parameter $c \in \mathbf{R}^1$ is called the travelling wave velocity and the function u^* the travelling wave profile (with travelling wave velocity c).

The conjectured bound by the travelling wave is of the form $\mathbf{Pr}(V^+(t) < x) \leq u^0(x - m(t))$, where u^0 is the travelling wave profile corresponding to a minimal possible velocity $c^0 > 0$, and $\lim_{t\to\infty} t^{-1}m(t) = c^0$.

Proposition 2.3 *Under the condition*

$$\mu^2/\sigma^2 > 2\gamma \tag{26}$$

the solution $u(t,x)$ of the Cauchy problem with the initial date $\theta(x)$ satisfies the following bound: for $t > T(x)$

$$u(x,t) \leq u^0(x - m(t)),\ x \in \mathbf{R}^1. \tag{27}$$

Here, $u^0(x)$ is the (unique) travelling wave profile that corresponds to the minimal velocity $c^0 = 2^{1/2}\sigma + \mu$ and obey $u^0(0) = 1/2$ and

$$m(t) = c^0 t + logt + o(logt). \tag{28}$$

The function $u^0(x)$ is exponentially decreasing at $-\infty$:

$$\lim_{x\to-\infty} \frac{\log u^0(x)}{x} = \rho < 0, \tag{29}$$

and threshold $T(x)$ obeys $T(x) \leq \max[Ax, B]$, where $A, B > 0$ are constants.

For the proof of Proposition 2.3, see [17]. We prove the following result in section 5.

Theorem 5 *Under condition (26),then for $x > 0$ and $t > T(x)$*

$$\mathbf{Pr}(V^+(t) < x) \le u^0(x - m(t)). \tag{30}$$

The bound (30) illustrates the behaviour, as $t \to \infty$, of process $V^+(t)$ and hence the limiting behavior of original waiting time process $v^+(t)$, in the overloaded system (we mean here the diffusion approximation limit of $v^+(t)$). In particular, bound (30) implies that $\mathbf{Pr}(V^+(t) < x)$ decreases exponentially in t as $t \to \infty$.

3 Convergence to the mean-field limit

3.1 The basic construction: Proof of Proposition 2.1

Theorem 1 may be derived from Proposition 2.1 (or, more precisely, from an argument the proof of this proposition is based upon). Thus, we concentrate on the proof of Proposition 2.1 and complete the proof of Theorem 1 afterwards. The proof of Proposition 2.1 is split into two steps: (i) the construction of a family of extended flows $\bar{\xi}_j^{>t_0}, 1 \le j \le N$, on a time half-axis $t_0 + \mathbf{R}_+ = [t_0, \infty)$ with, say, the zero initial condition at time t_0, and (ii) the proof of the existence of the limit of family $\bar{\xi}^{>t_0}$ when $t_0 \to -\infty$. As noted before, the first step is merely a straightforward iteration of equation (3).

As to the second step, it is based on an innovation event method going back to works of Borovkov (see [1,2]). In our context, this method exploits a detailed analysis of a 'predecessor picture', for a given task, as this picture develops backwards in time. A basic component of the predecessor picture is a compatible family of oriented random graphs (called predecessor graphs), associated with the task arriving in flows ξ_j after time t_0. The predecessor graph associated with a given task indicates the numbers and the types of its subsequent predecessors (the number of outgoing edges from any vertex of the predecessor graph is always equal to 1 or 2).

Speaking of a given (or 'tagged') task, we mean a fixed pair (j_0, n_0) and pick the n_0-th successive task in flow $\xi_{j_0}^{>t_0} = \xi_{j_0}\big|_{t_0+\mathbf{R}_+}$, the restriction of the original input flow ξ_{j_0} to $[t_0, \infty)$. The construction of the predecessor graph begins with a 'root' which is denoted again by $\mathbf{O}$; the root is associated with a tagged task. The number of the edges going out of $\mathbf{O}$ equals one if the component of the mark $a_{n_0} = (0, k)$. In the case $a_{n_0} = (1, k)$, $k \ge 1$, the task, generally speaking, has two predecessors, one in each of two queues. If, by chance, these predecessors coincide, the corresponding edges are 'glued' together.

Thus, the vertices at the other ends of the edges are associated with the predecessors of the tagged task. The above procedure is then repeated for these vertices: the number of the edges going out of them are defined in the same fashion, depending on their marks, and glueing occurs when their predecessors coincide. The procedure is then again repeated for the 'next' predecessors, etc. As a result, we end up with an oriented graph $\Delta^{>t_0}_{(j_0,n_0)}$ which is obtained from the Cayley tree, of branching ratio two, by 'cutting' off some branches and 'glueing' some edges (both procedures being random). Clearly, the graph constructed can include some loops because different tasks may share the same predecessor in the past.

A similar procedure may be performed for any task from the restricted input flows $\xi_j^{>t_0}$, $1 \leq j \leq N$; by construction, the predecessor graphs arising for different tasks possess a natural compatibility property. For a fixed $t_0 > -\infty$, the predecessor graphs eventually 'stop' when we reach initial time point t_0. However, the construction may be carried on indefinitely, creating 'infinitely-extended' graphs. Given a tagged task (j_0, n_0), it is convenient to think of an infinite predecessor graph $\Delta_{(j_0,n_0)}$ and a family of its 'cut-off' subgraphs of 'depth' ℓ, $\ell = 1, 2, \ldots$. More precisely, we consider, for any $\ell = 1, 2, \ldots$, a family $\mathcal{M}_{(j_0,n_0)}(\ell)$ of the paths on $\Delta_{(j_0,n_0)}$ starting at $\mathbf{O}$ and having exactly ℓ edges. While speaking of a fixed tagged task, index (j_0, n_0) is systematically omitted.

Observe that for any tagged task (j_0, n_0),

$$\mathbf{E}_{\mathcal{P}_N^{>t_0}} \, |\mathcal{M}(\ell)| \, \leq (1+p)^{\ell}. \tag{31}$$

Here, $|\mathcal{M}(\ell)|$ is the cardinality of set $\mathcal{M}(\ell)$, and $\mathcal{P}_N^{>t_0}$ is the joint probability distribution of family $\{\xi_j^{>t_0}, 1 \leq j \leq N\}$; one can of course replace it with the joint distribution $\bar{\mathcal{P}}_N^{>t_0}$ of family $\{\bar{\xi}_j^{>t_0}, 1 \leq j \leq N\}$ (see above).

Pictorially, the predecessor graphs absorb the whole 'N-dependance' that is present in probability distribution $\bar{\mathcal{P}}_N^{>t_0}$. To determine the predecessor picture, we assign, to the edges e of the predecessor graph, random variables ξ_e of the form $\xi_e = s_e - T_e$. Random variables s_e and T_e are assumed to be independent, as for a single edge as for different edges. [More precisely, s_e and T_e are assumed to be conditionally independent, given a realization of the predecessor graph.]

Random variables s_e have distribution F and random variables T_e are distributed exponentially, with means λ^{-1}.

The object we end up with (a random graph $\Delta_{(j_0,n_0)}$ and an array of random variables ξ_e) is inferred from probability distributions $\bar{\mathcal{P}}_N^{>t_0}$ (or $\mathcal{P}_N^{>t_0}$), $t_0 > -\infty$, and defines what was called the predecessor picture. By construction, the predecessor pictures obtained for different tasks are compatible. Given a tagged task (j_0, n_0), it is convenient to consider a continuous family of sub-graphs of predecessor graph Δ, labeled by a parameter representing the difference $t_{n_0} - t_0$ between t_{n_0}, the time the tagged task arrived at the network and the initial time t_0. More precisely, in analogy with $\mathcal{M}(\ell)$ (see above), we consider, for any given $T > 0$, a (random) family $\mathcal{M}^{(T)}_{(j_0,n_0)}$ of the paths $\mathbf{L}$, on graph Δ, which start at root $\mathbf{O}$ and are extended up to time T backwards, counting from t_{n_0}. [It means that the sum $\sum\limits_{e\in\mathbf{L}} l_e \leq T$.]

Physically speaking, a predecessor picture characterizes a chain of events that determines the 'fate' of a given task in families of the extended flows $\{\bar{\xi}_j^{>t_0}\}$. For example, the waiting time of a tagged task (j_0, n_0) in family $\{\bar{\xi}_j^{>t_0}\}$ is identified with the random variable

$$\max\left[\,0, \sup_{\mathbf{L}\,\in\,\mathcal{M}^{(t_{n_0}-t_0)}} \; \sum_{e\,\in\,\mathbf{L}} \xi_e\,\right]. \tag{32}$$

Observe that the final part of the above construction (assigning random variables ξ_e to the edges of the predecessor graph) does not explicitly refer to value N, as soon as a realization of the predecessor graph is fixed.

Next, we introduce an innovation event

$$\mathbf{C} = \left\{\exists \ell_0 \in \mathbf{N}:\ \forall \mathbf{L} \in \mathcal{M}(\ell),\ \text{with}\ \ \ell \geq \ell_0,\ \sum_{e\in L} \xi_e \leq 0\right\}. \tag{33}$$

We will use a standard Borel–Cantelli Lemma to check that for any N

$$\mathcal{P}_N^{>t_0}(C) = 1. \tag{34}$$

This is based on the fact that the series

$$\sum_{\ell} \mathcal{P}_N^{>t_0}\left(\text{there exists}\ \ \mathbf{L} \in \mathcal{M}(\ell) : \sum_{e\in\mathbf{L}} \xi_e > 0\right). \tag{35}$$

converges uniformly in N.

The proof of this fact is straightforward: series (35) is majorized, term by term, by the series

$$\sum_{\ell} \mathbf{E}_{\mathcal{P}_N^{>t_0}} |\mathcal{M}(\ell)|\ \widehat{P}\left(\Sigma_\ell > 0\right). \tag{36}$$

Here $\widehat{P}$ is the joint probability distribution of the sequence of i.i.d. random variables of the form $\xi = s - t$, where s and t are independent and distributed as before (s has distribution F and t an exponential distribution of mean λ^{-1}). Furthermore, Σ_ℓ is the sum of ℓ random variables from the sequence.

By using Chernoff's inequality we conclude that, for any $a > 0$,

$$\widehat{P}\,(\Sigma_\ell > 0) \leq \left(\mathbf{E}\big(\exp\,(a\xi)\big)\right)^\ell = \Big(\frac{\lambda}{\lambda + a}\mathbf{E}\exp\,(as)\Big)^\ell. \tag{37}$$

Owing to condition (6) and bound (31), there exists $a > 0$ that the series (35) converges uniformly in N. By the Borel–Cantelli Lemma we conclude that (34) is valid. Moreover, the same argument shows that the supremum in (32) is attained at a path $\mathbf{L}$ that belongs to a set $\mathcal{M}^{(T_0)}_{(j_0,n_0)}$ (i.e. is extended backwards up to (random) time $T_0(= T_0(j_0, n_0))$, from point t_{n_0}), with T_0 being bounded in probability uniformly in N, $j_0 = 1, 2, \dots, N$, $n_0 = 1, 2, \dots$, and $t_0 \in \mathbf{R}$:

$$\lim_{y\to\infty}\ \sup_{N,j_0,n_0,t_0}\ \mathcal{P}^{>t_0}_N(T_0 > y) = 0. \tag{38}$$

[In fact, given a $k \geq 1$, a similar relation holds for the random variables $\bar{T}_0(j_0, n_0, k) = \max\Big[T(j_0, \widetilde{n}_0) : n_0 \leq \widetilde{n}_0 \leq n_0 + k\Big]$.] This allows us to control the 'memory' in extended flows $\bar{\xi}^{>t_0}_j$, $j = 1, \dots, N$, uniformly in N and perform the limit $t_0 \to -\infty$. As a result, we obtain a stationary family of extended flows $\{\xi_j\}$ that solves equations (3). Finally, a coupling argument (cf. [11]) implies the uniqueness of the solution obtained, in the class indicated in Proposition 2.1. This completes the proof of Proposition 2.1. ■

3.2 Proof of Theorem 1.

The uniformity in N occurring in estimates discussed in the previous section allows us to pass to the limit $N \to \infty$. Namely, by inspecting a 'cut-off' family $\mathcal{M}^{(T)}_{(j_0,n_0)}$ (or $\mathcal{M}_{(j_0,n_0)}(\ell)$) of paths on predecessor graph $\Delta_{(j_0,n_0)}$, for a given $T > 0$ (or $\ell \geq 1$), one can observe that as $N \to \infty$ the following probability vanishes: that among tasks that are met in the course of inspection, either

(i) at least two have predecessors which are unprivileged and have coincident labels; or

(ii) at least one serves as a predecessor for more than one task.

Physically speaking, what survives in the limit $N \to \infty$ corresponds to the picture described in Theorem 1: for any tagged task (j_0, n_0) and any $T >$

0, no branches in the T-extended piece of the original Cayley tree are cut and no vertices are glued. The aforementioned uniformity allows us to make this argument rigorous. In other words, all random variables emerging in the recurrence (3) in the limit $N \to \infty$ become independent. This guarantees the existence of the limiting flow Ξ^* satisfying conditions (10). ∎

4 Convergence to the diffusion limit

It is convenient to start with proving of Theorem 4. We will need minor technicalities.

4.1 Preliminaries.

The processes $\{v_\tau^+, \tau \geq 0\}$ and $\{V^+(\tau), \tau \geq 0\}$ live in the space D^+ of functions $\rho : [0, \infty) \to \mathbf{R}^1$, without the second-kind discontinuities. For definiteness, let us assume that the functions are right-continuous. We endow space D^+ with the standard Skorokhod topology and the standard nested family of the σ-algebras $\mathcal{D}(\tau)$, $\tau > 0$. Fix $\tau^0 > 0$ and a bounded $\mathcal{D}(\tau^0)$-measurable function $f : D^+ \to \mathbf{R}^1$ which is continuous in the above topology. Denoting by $\mathbf{P}^+$ the probability distribution of process $\{v_\tau^+\}$ and by $\mathbf{Q}^+$ that of process $\{V^+(\tau)\}$, we have to prove that, in the course of the limiting procedure (15),

$$\lim \mathbf{E}_{\mathbf{P}^+} f = \mathbf{E}_{\mathbf{Q}^+} f \,. \tag{39}$$

In the course of the proof, we use directly a construction of processes $\{v_\tau^+\}$ and $\{V^+(\tau)\}$ which is related to the non-random tree, more precisely, to 'equipments' of the tree with some random elements. To this end, we have to introduce probability distributions $\mathbf{P}_{\text{tree}}$ and $\mathbf{Q}_{\text{tree}}$ of the 'basic' equipped trees for these processes. [Actually, distribution $\mathbf{Q}_{\text{tree}}$ was in substance determined in Section 2; the construction provided below completes the formal side of the definition.] Distributions $\mathbf{P}_{\text{tree}}$ and $\mathbf{Q}_{\text{tree}}$ are defined on a σ-algebra $\mathcal{T}$ in the space T of the basic equipped tree realizations. A point Γ of space T is given by a tree-graph, together with a collection of positive numbers g_e assigned to its edges e. It is convenient to agree now that the origin $\mathbf{O}$ of the tree creates a single edge; any other vertex of the tree always creates two edges. The number g_e assigned to an edge e may be interpreted as the length of e.

The σ-algebra $\mathcal{T}$ is generated by the variables g_e. Both probability distributions $\mathbf{P}_{\text{tree}}$ and $\mathbf{Q}_{\text{tree}}$ correspond to the collection of i.i.d. random variables distributed exponentially, with means $(\lambda p)^{-1}$ (for $\mathbf{P}_{\text{tree}}$) and γ^{-1} (for $\mathbf{Q}_{\text{tree}}$).

The construction of process $\{v_\tau^+\}$ is in fact related to a refinement of a basic tree which may be called the 'barbed' tree. We replace the random variable g_e with an ordered set $\mathbf{l}_e = (l_e(1), \dots, l_e(n_e))$ consisting of a random number n_e of positive random 'bits'. The number of the bits is geometrically distributed, with parameter p, and, given n_e, their sizes $l_e(i)$, $i = 1, \dots, n_e$, are i.i.d. and have the exponential distribution with mean λ^{-1}. The assumption of the independence of pairs $(n_e, \mathbf{l}_e)$ for different edges is retained. The space of the barbed trees is denoted by T^*, the corresponding σ-algebra by $\mathcal{T}^*$ and the probability distribution arising by $\mathbf{P}^*_{\rm tree}$. The connection between $\mathbf{P}_{\rm tree}$ and $\mathbf{P}^*_{\rm tree}$ is established via the formula $g_e = \sum_i l_e(i)$. Pictorially, a barbed tree Γ^* may be associated with a basic tree where the length of each edge is partitioned into bits.

We construct the distributions of processes $\{v^+_\tau\}$, up to time τ^0, as follows. First, we perform an additional equipment of a barbed tree with the i.i.d. copies of the random variable s assigned to each partition point $z_e(1) = l_e(1), z_e(2) = l_e(1) + l_e(2), \dots, z_e(n_e) = l_e(1) + \dots + l_e(n_e)$. Given the whole array $\{l_e(i),\ s_e(i)\}$, we compute the value v^+_τ as follows. First, we extract an (infinite) path $\mathbf{L}_0$ from the origin $\mathbf{O}$ and truncate it at length τ^0 (we call this path basic). Truncation means that we fix an edge $e_0 = e_0(\tau^0)$ of $\mathbf{L}_0$ and a number i_0 $(= i_0(e_0))$ (or, equivalently, the corresponding partition point $z_{e_0}(i_0)$) such that

$$\bar{z}_{e_0}(1) = \tau^0 - \sum_{e \in \mathbf{L_0}: e \prec e_0} g_e - \sum_{i_0 < i \le n_{e_0}} l_{e_0}(i) > 0, \tag{40}$$

but

$$\bar{z}_{e_0}(1) - l_{e_0}(i_0) < 0. \tag{41}$$

Here $e \prec e_0$ means that edge e precedes e_0 on path $\mathbf{L}_0$ while moving from $\mathbf{O}$; in the case $i_0 = n_{e_0}$, the second sum in the LHS of (40) is omitted. Note the reversed direction of labelling the partition points on the edges of the barbed tree.

A similar truncation procedure is performed for any path from $\mathcal{L}^0$. It gives us a collection of pairs $(\bar{e}, i_0(\bar{e}))$ that, together with (e_0, i_0), forms a 'cut' of the barbed tree. We obtain a 'truncated' tree, from $\mathbf{O}$ to the cut, such that

$$G(\bar{e}, i_0(\bar{e})) = \sum_{e \in \mathbf{L}_{\bar{e}}: e \prec \bar{e}} g_e + \sum_{i > i_0(\bar{e})} l_{\bar{e}}(i) < \tau^0, \tag{42}$$

but

$$G(\bar{e}, i_0(\bar{e})) + l_{\bar{e}}(i_0(\bar{e})) > \tau^0. \tag{43}$$

Here, $\mathbf{L}_{\bar{e}}$ denotes the path from $\mathcal{L}^0$ that ends on $\bar{e}$.

The next step is to install the 'physical' clock along the (truncated) path $\mathbf{L}^0$. It means that we mark, on the time-interval $(0, \tau_0)$, points

$$\begin{aligned} \bar{z}_{e_0}(1), \quad \bar{z}_{e_0}(2) &= \bar{z}_{e_0}(1) + l_{e_0}(i_0+1), \dots \\ \bar{z}_{e_0}(n_{e_0} - i_0) &= \bar{z}_{e_0}(1) + l_{e_0}(i_0+1) + \ldots + l_{e_0}(n_{e_0}), \end{aligned} \tag{44}$$

then points

$$\begin{aligned} \bar{z}_{e_1}(1) &= \bar{z}_{e_0}(n_{e_0}) + l_{e_1}(1), \dots \\ \bar{z}_{e_1}(n_{e_1}) &= \bar{z}_{e_0}(n_{e_0}) + l_{e_1}(1) + \ldots + l_{e_1}(n_{e_1}), \end{aligned} \tag{45}$$

etc., until we reach (at time τ^0) the origin $\mathbf{O}$. Here and below e_1 denotes the edge preceding e_0 on the truncated main path $\mathbf{L}_0$, e_2 denotes the edge preceding e_1, etc. Points $\bar{z}_{e_1}(1)$, $\bar{z}_{e_2}(1), \dots$ are of a special significance: they correspond the time epochs when the (truncated) barbed tree branches. Note that the direction of the physical time is reversed with respect to the direction in which the tree develops.

4.2 Convergence on the half-axis.

We now proceed as follows. Given $\tau \in (0, \tau^0)$, we extract a part of the truncated barbed tree that contains a piece of the basic path corresponding to the time-interval $(0, \tau)$, together with all associated branches. Then set

$$v_\tau^+ = \max[0, \sup_{0<\tilde{\tau}\leq\tau} \max_{\mathbf{L}(\tilde{\tau},\tau)} \sum_{e,i} {}^{(0,\tau)} (\bar{s}_e(i) - \bar{l}_e(i))]. \tag{46}$$

The internal maximum $\max\limits_{\mathbf{L}(\tilde{\tau},\tau)}$ in (46) is taken over all paths on the truncated barbed tree which start at time τ (more precisely, at the corresponding point on the main path) and move backwards in the physical time, up to time $\tilde{\tau}$. Any such path $\mathbf{L}(0, \tau)$ is formed by: (i) a piece of the main path between τ and the first preceding point when another branch joined the main path, (ii) a collection of subsequent paths of the barbed tree (we mean by that the piece of a bit of the basic path which precedes, in the sense of the physical time-direction, point τ is appended (we call it the idle piece), as well as a piece of the border bit of a final path which starts immediately after the physical time zero). The sum $\sum\limits_{e,i}{}^{(0,\tau)}$ is over all bits that partition the corresponding path (including two pieces appended). The quantity $\bar{l}_e(i)$ is the length of the corresponding bit (or of its piece), and the quantity $\bar{s}_e(i)$ equals zero for the idle piece and the corresponding value $s_e(i)$ otherwise.

Formula (46) is of course equivalent to a recursive iteration of an operation analogous to that represented in (11).

The above construction determines process $\{v_\tau^+\}$ on the space of equipped barbed trees endowed with additional arrays of i.i.d. variables $s_e(i)$. In order not to make the exposition too much overloaded with the notation, we keep, for the corresponding probability space, the previous notation $(T^*, \mathcal{T}^*, \mathbf{P}_{\rm tree})$.

As to the distribution of process $\{V^+(\tau)\}$ (again up to time τ), it may be constructed over space $(T, \mathcal{T}, \mathbf{Q}_{\rm tree})$ in a similar way, by using an additional equipment of the basic (non-barbed) tree $\Gamma \in T$ with Wiener trajectories S_e assigned to the edges e of Γ. Trajectory S_e lives on the time-interval $(0, g_e)$ and has the drift coefficient φ and the diffusion coefficient σ, which are given by (15). The trajectories corresponding to the different edges are independent. As before, we use the same notation $(T, \mathcal{T}, \mathbf{Q}_{\rm tree})$ for the probability space that emerges after assigning Wiener trajectories S_e.

As before, we can perform the procedure of truncating the tree (with a fixed main path) and installing a physical clock, up to time τ^0. Here, the time-interval $(0, \tau^0)$ is partitioned by the points $\bar{z}_e$ when the main path is joined by other branches of the tree. Note that the trajectories S_e must 'live' in the reversed time which means that they take value zero at the time of junction with the main path (in the case of the main path itself, the trajectories vanish at the final epochs). Therefore, we make an agreement that $S_e(g_e) = 0$.

After this, we set

$$V^+(\tau) = \max\,[0, \sup_{0<\tilde{\tau}\leq\tau} \max_{\mathbf{L}(\tilde{\tau},\tau)} \sum_e {}^{(0,\tau)} \bar{S}_e(\bar{g}_e)\,]. \tag{47}$$

As before, the internal maximum $\max\limits_{\mathbf{L}(\tilde{\tau},\tau)}$ in (47) is taken over all paths on the truncated tree which start at time τ on the main path and move backwards in the physical time, up to time $\tilde{\tau}$. Any such path $\mathbf{L}(0, \tau)$ is again formed by a piece of the main path between time τ and the first preceding epoch when other branch joins the main path (we call this piece the τ-piece) and a collection of subsequent paths of the truncated tree (including a piece of the final path which starts immediately after the physical time zero). Summation $\sum_e {}^{(0,\tau)}$ is over all edges constituting the corresponding path (including two pieces appended). The quantity $\bar{g}_e$ is the length of edge e (or of its piece). Furthermore, the quantity $\bar{S}_e(\bar{g}_e)$ equals $S_e(0) - S_e(\bar{z}_e)$ for the τ-piece (here again e is the corresponding edge) and the value $S_e(0)$ otherwise.

For our purpose, it is convenient to work with conditional distributions $\mathbf{P}^+(\cdot|\mathcal{T})$ and $\mathbf{Q}^+(\cdot|\mathcal{T})$ that emerge from the construction provided. The distributions $\mathbf{P}^+(\cdot|\mathcal{T})$ and $\mathbf{Q}^+(\cdot|\mathcal{T})$ are determined for almost all $\Gamma \in T$ with respect to $\mathbf{P}_{\mathrm{tree}}$ and $\mathbf{Q}_{\mathrm{tree}}$. However, we can think of their natural variants determined everywhere on T.

The following lemma is a modification of the well-known results about the convergence to the diffusion limit (see e.g. [2], Ch.3 and the references therein, and [14], Sect. 4.4 and the references therein.)

Lemma 4.1 *In the limiting procedure (39), everywhere on T there exists the limit*

$$\lim \mathbf{E}_{\mathbf{P}^{+}(\cdot|\mathcal{T})} f = \mathbf{E}_{\mathbf{Q}^{+}(\cdot|\mathcal{T})} f. \tag{48}$$

Furthermore, for any given constant $\delta > 0$, the difference

$$|\mathbf{E}_{\mathbf{P}^{+}(\cdot|\mathcal{T})} f - \mathbf{E}_{\mathbf{Q}^{+}(\cdot|\mathcal{T})} f| \tag{49}$$

may be made uniformly small for any tree $\Gamma \in \mathcal{C}_\delta$ where $\mathcal{C}_\delta \subset T$ is the set of the trees for which, after the truncation by time τ^0, at most one branching point occurs on any sub-interval of $[0, \tau^0)$ of length δ and no branching point occurs on the intervals $[0, \delta)$ and $[\tau^0 - \delta, \tau^0)$.

The proof of Lemma 4.1 repeats that of the results quoted above and based on an appropriate invariance principle for random variable

$$\sum_{e,i}{}^{(0,\tau)} (\bar{s}_e(i) - \bar{l}_e(i)), \tag{50}$$

which yields the convergence to

$$\sum_{e}{}^{(0,\tau)} (\bar{S}_e(\bar{g}_e) - \bar{g}_e) \tag{51}$$

for the array of paths $\mathbf{L}(0, \tau]$. The convergence (48) is then a result of standard topological considerations, because, in the set-up of Theorem 4, $v^{+}(\tau)$ and $V^{+}(\tau)$ are local functionals of the array. It is here that condition (16) is used, to simplify the proof of the above invariance principle (this condition can certainly be weakened, at the expense of greater technical complexity). We omit the details: they are in essence repeated below, in the course of proving Theorem 2.

To finish the proof of (39), we write

$$\mathbf{E}_{\mathbf{P}^{+}} f = \mathbf{E}_{\mathbf{P}_{\mathrm{tree}}}[\mathbf{E}_{\mathbf{P}^{+}(\cdot|\mathcal{T})} f], \ \mathbf{E}_{\mathbf{Q}^{+}} f = \mathbf{E}_{\mathbf{Q}_{\mathrm{tree}}}[\mathbf{E}_{\mathbf{Q}^{+}(\cdot|\mathcal{T})} f], \tag{52}$$

then observe that the set $\mathcal{C}_\delta$ has both $\mathbf{P}_{\mathrm{tree}}$- and $\mathbf{Q}_{\mathrm{tree}}$-probabilities tending to one as $\delta \to 0$, and finally note that function $\mathbf{E}_{\mathbf{Q}^{+}(\cdot|\mathcal{T})} f$ is continuous in the Skorokhod topology.

4.3 Convergence on the whole axis.

Let us now pass to the proof of Theorem 2. Recall that under condition (17) the corresponding random variable X has a proper distribution, and hence the process $\{V(\tau),\ \tau \in \mathbf{R}^1\}$ is correctly defined.

Physically speaking, the difference between Theorems 2 and 4 is that in Theorem 2 we deal with (stationary) processes determined on the whole time-axis. In terms of the underlying Cayley tree picture, it means that we have to work with a random tree on the infinite half-axis. The point is, however, that the random variables v_τ and $V(\tau)$, $\tau \in \mathbf{R}^1$, depend essentially on a finite-depth 'piece' of the tree (see Lemma 4.2). This property allows us to perform a 'cut-off' procedure and reduce the problem to that of proving a convergence of truncated processes which are simply local functionals of a nice array of the external input processes.

The formal construction is as follows. The processes $\{v_\tau, \tau \in \mathbf{R}^1\}$ and $\{V(\tau), \tau \in \mathbf{R}^1\}$ live in the space D of functions $\mathbf{R}^1 \to \mathbf{R}^1$ without the second-kind discontinuities, which, as before, is endowed with the standard Skorokhod topology and the family of nested σ-algebras $\mathcal{D}(\tau_1, \tau_0)$, $\tau_1 < \tau_0$. Let us denote the distributions of processes $\{v_\tau,\ \tau \in \mathbf{R}^1\}$ and $\{V(\tau),\ \tau \in \mathbf{R}^1\}$ by $\mathbf{P}$ and $\mathbf{Q}$, respectively. As before, we need to relate processes $\{v_\tau\}$ and $\{V(\tau)\}$, more precisely, their restrictions $\{v_\tau,\ \tau \le \tau_0\}$, and $\{V(\tau),\ \tau \le \tau_0\}$, for a fixed $\tau_0 \in \mathbf{R}^1$, to the probability spaces $(T, \mathcal{T}^*, \mathbf{P}_{\mathrm{tree}})$ (more precisely, $(T^*, \mathcal{T}^*, \mathbf{P}^*_{\mathrm{tree}})$) and $(T, \mathcal{T}, \mathbf{Q}_{\mathrm{tree}})$ (although the construction formally depends on the choice of τ_0, there is an obvious consistency property that allows us to extend all results to the process determined on the whole time-axis). The technical details of the construction are similar to those above: we only point that, while installing a physical clock along the main path, we now work on the half-infinite axis $(-\infty, \tau^0)$, identifying time τ^0 with the origin of the tree $\mathbf{O}$ (as before, the origin emits a single edge). As a result, we do not have to care about the compatibility of a given particular construction with the direction of the physical time: all procedures develop backwards, from τ^0 to $-\infty$. The main difference is that the internal maxima in formulas (46) and (47) are now replaced by the suprema taken over all (finite) paths $\mathbf{L}(\tau)$ on an (infinite) tree Γ (or an infinite barbed tree Γ^*, in the case of process $\{v_\tau\}$):

$$v_\tau = \max\,[\ 0, \sup_{\mathbf{L}(\tau)} \sum_{e,i} {}^{(\tau)}(\bar{s}_e(i) - \bar{l}_e(i))\ ] \tag{53}$$

and

$$V(\tau) = \max\,[\ 0, \sup_{\mathbf{L}(\tau)} \sum_{e} {}^{(\tau)}\bar{S}_e(\bar{g}_e)\ ]. \tag{54}$$

However, as we said above, this dependence is essentially localized. More precisely, fix $\bar{\tau} < \tau_0$ and consider an auxiliary process $\{\widetilde{v}_{\bar{\tau},\tau},\ \bar{\tau} \le \tau \le \tau_0\}$ distributed as $\{v^+_{\tau-\bar{\tau}}\}$ and coupled with $\{v_\tau,\ \tau \le \tau_0\}$ in an obvious way so that we can use the symbol $\widehat{\mathbf{P}}^*_{\rm tree}$, having in mind the joint distribution of these two processes. Likewise, we consider an auxiliary process $\{\widetilde{V}(\bar{\tau},\tau),\ \bar{\tau} \le \tau \le \tau_0\}$ distributed as $\{V^+(\tau-\bar{\tau})\}$ and coupled with $\{V(\tau),\ \tau \le \tau_0\}$ in a similar way, allowing us to use the symbol $\widetilde{\mathbf{Q}}_{\rm tree}$ while working with the joint distribution of these two processes.

Lemma 4.2 *For any $\tau_1 \in (\bar{\tau}, \tau_0)$,*

$$\widehat{\mathbf{P}}^*_{\rm tree}(\exists \widehat{\tau} \in (\tau_1, \tau_0) \text{ such that } v_{\widehat{\tau}} \neq \widetilde{v}_{\bar{\tau},\widehat{\tau}}) < \psi(\tau_1 - \bar{\tau}) \tag{55}$$

where $\psi(u)$ monotonically tends to zero as $u \to \infty$. A similar bound holds for processes $\{\widetilde{V}(\bar{\tau},\tau)\}$ and $\{V(\tau)\}$.

Proof. Consider the case of processes $\{\widetilde{v}_{\bar{\tau},\tau}\}$ and $\{v_\tau\}$ only: the case of processes $\{\widetilde{V}(\bar{\tau},\tau)\}$ and $\{V(\tau)\}$ is completely similar.

By the construction, both processes 'live' on the same barbed tree, and they both are equal to suprema of sums of random variables along the (finite) paths on the barbed tree. In the case of process $\{v_\tau\}$ the paths are arbitrary and in the case of process $\{\widetilde{v}_{\bar{\tau},\tau}\}$ they should not propagate beyond the physical time $\bar{\tau}$. The assertion of Lemma 4.2 follows if we show that the aforementioned supremum, while taken over the paths of time-length $t > 0$, is positive with $\widehat{\mathbf{P}}^*_{\rm tree}$-probability decreasing to zero as t tends to infinity.

In fact, this probability does not exceed

$$\widehat{\mathbf{P}}^*_{\rm tree}(\exists \text{ a path } \mathbf{L} \in \mathcal{L}^0 \text{ of depth} \ge t^{1/3} : \sum_{\beta \in \mathbf{L}} (S_\beta - l_\beta) > 0) \tag{56}$$

$$+\widehat{\mathbf{P}}^*_{\rm tree}(\exists \text{ a path } \mathbf{L} \in \mathcal{L}^0 \text{ of length and depth } \le t^{1/3}). \tag{57}$$

By the depth of a path we mean here the number of its edges. The second term in (56) is bounded from above by $3^{t^{1/3}} \exp[-c\lambda p t^{2/3}]$ for some constant $c > 0$. The main work is with the first term: it is less than

$$\sum_{n \ge t^{1/3}} \sum_{\mathbf{L} \in \mathcal{L}^0 : d(\mathbf{L}) = n} \pi(\mathbf{L}) P_{\mathbf{L}}. \tag{58}$$

Here, $\pi(\mathbf{L})$ is the probability that a path $\mathbf{L}$ is present in the the tree and $P_{\mathbf{L}}$ is the probability that the sum $\sum_{\beta \in \mathbf{L}} (S_\beta - l_\beta) > 0$.

The further estimates are standard:

$$P_{\mathbf{L}} \le E_{\mathbf{L}}(a) = \prod_{\beta \in \mathbf{L}} E(\beta, a), \tag{59}$$

where $E_{\mathbf{L}}(a)$ is the expectation value of $\exp[a\sum_{\beta\in\mathbf{L}}(S_\beta - l_\beta)]$ and $E(\beta, a)$ is the expectation value of $\exp[a(S_\beta - l_\beta)]$. The last bound holds for any $a > 0$. In view of (31) a straightforward calculation leads to

$$\text{the RHS of (58)} \le \sum_{n\ge t^{1/3}} ((1+p)\inf_{a>0}\frac{\lambda}{(\lambda+a)}\mathbf{E}e^{as})^n. \tag{60}$$

Under conditions (6) the series (60) converges. This leads to the assertion of Lemma 4.2.

Lemma 4.2 allows us to finish quickly the proof of Theorem 2. In fact, let $f : D \to \mathbf{R}^1$ be a continuous bounded function which is measurable with respect to the σ-algebra $\mathcal{D}(\tau_1, \tau_0)$. Fix an arbitrary $\epsilon > 0$. By using Lemma 4.2, we can choose $\bar\tau < \tau_1$ so that the difference

$$|\mathbf{E}_{\mathbf{P}}f - \mathbf{E}_{\widetilde{\mathbf{P}}(\bar\tau)}f| < \epsilon, \tag{61}$$

and a similar bound holds for the expectations with respect to the measures $\mathbf{Q}$ and $\widetilde{\mathbf{Q}}(\bar\tau)$. Here $\widetilde{\mathbf{P}}(\bar\tau)$ and $\widetilde{\mathbf{Q}}(\bar\tau)$ are the distributions of processes $\{\widetilde v_{\bar\tau,\tau}\}$ and $\{\widetilde V(\bar\tau,\tau)\}$, respectively. Then, keeping $\bar\tau$ fixed, we perform the limiting procedure (15.) for processes $\{\widetilde v_{\bar\tau,\tau}\}$ obtaining that

$$\lim \mathbf{E}_{\widetilde{\mathbf{P}}(\bar\tau)}f = \mathbf{E}_{\widetilde{\mathbf{Q}}(\bar\tau)}f\,. \tag{62}$$

To prove (62), we just need to repeat the proof of Theorem 4. Since ϵ in (61) is arbitrary, we get from (61) the relation

$$\lim \mathbf{E}_{\mathbf{P}}f = \mathbf{E}_{\mathbf{Q}}f\,. \tag{63}$$

5 The travelling wave bound

The proof of Theorem 5 is based on the McKean formula (see [17]) for the solution $u(t;x)$ of the Cauchy problem equation (25), with initial date $\theta(x)$:

$$u(\tau;x) = \mathbf{Q}_{\rm tree}(Y_\tau < x), \tag{64}$$

where Y_τ, $\tau > 0$, is a random process over the probability space $(T, \mathcal{T}, \mathbf{Q}_{\rm tree})$ (with the additional equipment by Wiener trajectories S_e). Formally,

$$Y_\tau = \max_{\mathbf{L}(0,\tau)} \sum_e{}^{(0,\tau)}\ \left(\bar S_e(\bar g_e)\right)\,. \tag{65}$$

As follows from (47), with probability one, $Y_\tau \le V_\tau^+$ for all $\tau \ge 0$. Thus, the bound

$$u(\tau;x) \ge \mathbf{Q}^+(V_\tau^+ < x) \tag{66}$$

holds for all $x \in \mathbf{R}$ and $\tau \ge 0$. Using Proposition 2.3 completes the proof. ■

References

1. A.A. Borovkov, Convergence of the distributions of functionals of sequences and processes given on the entire axis. In: *Proceedings of Steklov Mathematical Institute, The USSR Academy of Sciences*, vol. **128**, 41–65, (1972).
2. A.A. Borovkov, *Asymptotical Methods in Queueing Theory.* (Springer-Verlag, Berlin, 1980)
3. F. Baccelli, F.I. Karpelevich, M.Ya. Kelbert, A.A. Pukhalskii, A.N. Rybko and Yu.M. Suhov, Mean-field limit for a class of queueing networks. *J. Stat. Phys.,* **66**, 803–825 (1992).
4. F. Baccelli, A. Makowski, Queueing models for systems with synchronization constraints. *Proceedings IEEE,* **77**, 138–161 (1989).
5. B. Chauvin and B. Rouault, KPP equation and supercritical branching Brownian motion in the subcritical speed area. Application to spatial trees. *Prob. Theory Rel. Fields,* **80**, 299–314 (1988).
6. B. Chauvin, B. Rouault and A. Wakolbinger, Growing conditioned trees. *Stoch. Proc. Appl.,* **39**, 117–130 (1991).
7. N.G. Duffield, Local mean-field Markov processes: an application to message switching networks. *Prob. Theory Rel. Fields,* **93**, 485–505 (1992).
8. N.G. Duffield, M.Ya. Kelbert and Yu.M. Suhov, A branching diffusion approximation for a class of queueing networks. In: *Stochastic Analysis and Applications* (I.M. Davies, A. Truman and K.D. Elworthy, Eds.) pp201–220 (World Scientific, Singapore, 1996).
9. F.M. Dekking and B. Host, Limit distributions for minimal displacement of branching random walk. *Prob. Theory Rel. Fields,* **90**, 403–426 (1991)
10. M.I. Freidlin, Limit theorems for large deviations and reaction-diffusion equations. *Ann. Prob.,* **13**, 639–675 (1985).
11. F.I. Karpelevich, M.Ya. Kelbert and Yu.M. Suhov, Higher-order Lindley equations. *Stoch. Proc. Appl.,* **53**, 65–96 (1994).
12. F.I. Karpelevich, M.Ya. Kelbert and Yu.M. Suhov, The branching diffusion, stochastic equations and travelling wave solutions to the equation of Kolmogorov–Petrovskii–Piskunov. In: *Cellular Automata and Cooperative Behaviour* (N. Boccara, R. Goles, S. Martinez and P. Ricco, Eds.) pp343–366 (Kluwer, Dordrecht, 1993).
13. F.I. Karpelevich, M.Ya. Kelbert and Yu.M. Suhov, Stochastic equations on random trees. In: *Cellular Automata and Cooperative Behaviour* (N. Boccara, R. Goles, S. Martinez and P. Ricco, Eds.) pp323–342 (Kluwer, Dordrecht, 1993).
14. M.Ya. Kelbert and Yu.M. Suhov, Mathematical theory of queueing net-

works. *Journ. of Soviet Mathematics,* **50**, 1527–1600 (1990).

15. M.Ya. Kelbert and Yu.M. Suhov, The Markov branching random walk and system of reaction diffusion (Kolmogorov-Petrovskii- Piskunov) equations. *Commun. Math. Phys.,* **167**, 607–634 (1995).
16. A. Kolmogorov, I. Petrovskii and N. Piscounov, Étude de l'équation de la diffusion avec croissance de la quantité de la matière et son application à un problème biologique. *Bull. Math. Moscow University,* **1**, 1–25 (1937)
17. H.P. McKean, Applications of Brownian motion to the equation of Kolmogorov - Petrovskii - Piscounov. *Comm. Pure Appl. Math.,* **28**, 323–331 (1975).