

QUALITATIVE PROPERTIES OF MATHEMATICAL MODEL FOR DATA FLOW

CORY D. HAUCK*

Computational and Applied Mathematics Group
Oak Ridge National Laboratory
1 Bethel Valley Road, Bldg. 5700, Oak Ridge, TN 37831-6164, USA
and
Department of Mathematics
University of Tennessee
227 Ayres Hall, 1403 Circle Drive, Knoxville TN 37996-1320, USA

MICHAEL HERTY AND GIUSEPPE VISCONTI*

Institut für Geometrie und Praktische Mathematik (IGPM)
RWTH Aachen University
Templergraben 55, 52062 Aachen, Germany

(Communicated by Shi Jin)

ABSTRACT. In this paper, properties of a recently proposed mathematical model for data flow in large-scale asynchronous computer systems are analyzed. In particular, the existence of special weak solutions based on propagating fronts is established. Qualitative properties of these solutions are investigated, both theoretically and numerically.

1. Introduction. The increasing number and diversity of processing units in modern supercomputers presents a significant challenge to the understanding of how data is globally distributed in these systems. This issue is critically important in emerging computer architectures for which the cost of data movement is becoming a critical driver for system design, affecting hardware and software choices and even basic algorithms [12, 17, 18]. At the same time, these supercomputers are becoming increasingly difficult to model. Indeed, simulating the detailed state of a supercomputer requires an even larger computer, thereby making such approaches prohibitive for the largest systems.

With the help of new mathematical models, a more efficient utilization of available processing units might be possible. Such models are necessarily coarse-grained approximations, used to understand the gross behavior of the system in order to

2020 *Mathematics Subject Classification.* 35L65, 35L40.

Key words and phrases. Conservation laws, predictive performance, weak solutions, data flow, front propagation.

The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for the United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

* Corresponding author: Giuseppe Visconti.

make design, control, and optimization more tractable. In this spirit, a coarse-grained mathematical model has been recently derived in [5] to describe the movement of processed data in an extreme-scale computer system. This model focuses on a idealized setting, in which a lattice of processors perform operations in parallel and communicate with nearest neighbors. Although very simple, the model allows for local variations in processor speed, due to differences in hardware performance or local problem complexity, as well as asynchronously operation, meaning that processors manipulate and share data with their neighbors as they are able, as opposed to a staged process in which processors communicate only after they have all completed a step of assigned work.

The model proposed in [5] takes the form of a PDE for the density of data at a particular processor location and stage in the global computation. It is derived formally from a microscopic ODE description. The notion of coarse-graining discrete and semi-discrete model of networks and related systems to derive PDE models is common mathematical strategy that is used in many applications, including vehicular [4, 11, 16] and pedestrian traffic flow [1, 6, 8, 15], supply chains [2, 3], bacterial movement [14] and machine learning applications [7].

In the current work, we establish some qualitative properties of the coarse-grained PDE model in [5]. We focus specifically on the phenomena of front propagation and present simulation results that highlight the theoretical findings. In the context of the current application, these fronts describe the movement of data through the different stages of a computational task. The reason for this analysis is two-fold. First, it was observed in [5] that, after a transient phase, many initial conditions eventually relax to a front whose shape depends on local variations in the processor speed. Second, these fronts can be used to characterize three important features for application: the first time at which some portion of the data reaches a state of completion, the rate at which the rest of the data reaches this state, and the time at which all data reaches this stage. Moreover, the relatively simple form of these fronts enables the investigation of optimization and control strategies based on some application-specific metrics. As a first step in this direction, we explore how variations in the processor speed affect front behavior.

2. Mathematical description of the model. In this section, we briefly summarize the microscopic ODE model from [5] and the macroscopic model that is obtained from it. The microscopic model describes a computer system as a lattice of processors, the dimension of which can be arbitrary, but finite. For application purposes, only one-, two-, and three-dimensional lattices are realistic, and here we focus the one-dimensional case for simplicity. It is assumed in [5] that each processor performs the same task, which is broken into discrete stages. The rate at which a given processor moves data from stage to stage depends not only on its intrinsic processing rate, but also on the availability of usable data in the processor and its neighbors.

Mathematically, the amount of data at time t that sits in stage $k \in \{1, \dots, k_{\max}\}$ of processor $i \in \{1, \dots, i_{\max}\}$ is given by $q_{i,k}(t)$. In the absence of any throttling, the flow of data between stages is given by a (processor dependent) rate $a_i(t) \geq 0$. However the actual flow may be reduced due to a lack of available data in a given processor or its nearest neighbors. Throttling due to a lack of usable data in the nearest neighbors is characterized by the parameter $\eta := k_{\max}/i_{\max} > 0$. As η increases the global effect of local slowdowns increases.

For each $i \in \{1, \dots, i_{\max}\}$ and $k \in \{1, \dots, k_{\max}\}$, the evolution of $q_{i,k}$ is modeled by the ordinary differential equation

$$\dot{q}_{i,k}(t) = f_{i,k-1}(t) - f_{i,k}(t), \quad q_{i,k}(0) = q_{i,k}^0, \quad f_{i,0}(t) = f_i^{\text{in}}(t), \quad (1)$$

where $q_{i,k}^0$ is the (known) initial amount of data present in each processor i at each stage k and $f_i^{\text{in}}(t)$ is the (known) inflow of unprocessed (or raw) data at processor i and time t . For $k \in \{1, \dots, k_{\max}\}$, $f_{i,k}$ is the flow of data in processor i from stage k to $k+1$. It is given by the product of the maximum processing rate a_i and the composition of two *throttling functions* v_1 and v_2 ; that is,

$$f_{i,k} = a_i v_1 \left(v_2(q_{i,k}, Q_{i+1,k} - Q_{i,k} + q_{i,k}, Q_{i-1,k} - Q_{i,k} + q_{i,k}); q_* \right), \quad (2)$$

where $Q_{i,k}(t)$ is the total amount of data in processor i that by time t has traversed the first $k-1$ stages of the computation:

$$Q_{i,k}(t) = \sum_{j=k}^{k_{\max}} q_{i,j}(t) + \int_0^t f_{i,k_{\max}}(s) ds, \quad (3)$$

and $Q_{i\pm 1,k} - Q_{i,k} + q_{i,k}$ is the amount of usable data available from processor $i \pm 1$. The throttling functions are given by

$$v_1(q; q_*) = \max \left\{ 0, \min \left\{ 1, \frac{q}{q_*} \right\} \right\}, \quad (4)$$

$$v_2(q, \Delta_+, \Delta_-) = \min \{q_{i,k}, \max\{\Delta_+, 0\}, \max\{\Delta_-, 0\}\}. \quad (5)$$

The function v_1 is a linear ramp, sometimes referred to as roof-line model [19]. The steepness of the ramp is determined by the parameter $q_* > 0$, the value at which the processor is fully on. The function v_2 determines the amount of data ready to process based on what is available locally and from nearest neighbors. In particular, if not enough data is available from neighboring processors, then $f_{i,k} < a_i$.¹

The macroscopic model in [5] is a continuum approximation for (1) that is derived in the limit of infinitely many processors ($i_{\max} \rightarrow \infty$) and stages ($k_{\max} \rightarrow \infty$), subject to that constraint that η is a fixed, positive, finite constant. This continuum model is a PDE for a function ρ , where $\rho(t, x, z)$ is interpreted as the local density of data at time t , processor location x , and stage completion variable z . It takes the form

$$\partial_t \rho(t, x, z) + \partial_z \Phi(\rho(t, x, z), \partial_x P(t, x, z)) = 0, \quad (t, x, z) \in \mathbb{R}^+ \times \mathbb{T} \times (0, 1), \quad (6a)$$

$$\rho(0, x, z) = \rho_0(x, z), \quad (x, z) \in \mathbb{T} \times (0, 1), \quad (6b)$$

$$\Phi(\rho(t, x, 0), \partial_x P(t, x, 0)) = F^{\text{in}}(t, x), \quad (t, x) \in \mathbb{R}^+ \times (0, 1), \quad (6c)$$

where $\mathbb{R}^+ = (0, \infty)$; \mathbb{T} is the torus parameterized by $x \in [0, 1]$; F^{in} is a scaled, continuous version of f^{in} ; P and Φ are given by

$$P(t, x, z) = \int_z^1 \rho(t, x, y) dy + \int_0^t \Phi(\rho(s, x, 1), \partial_x P(s, x, 1)) ds, \quad (7a)$$

$$\Phi(\rho, \sigma) = \alpha w_1(w_2(\rho, \sigma, -\sigma)); \quad (7b)$$

¹In [5], there is another parameter $\beta \in [0, 1]$ in the definition of v_2 that measures the effective strength of coupling between processors. Here we assume $\beta = 1$ and do not consider it any further.

and w_1 and w_2 are scaled version of v_1 and v_2 :

$$w_1(u) = \min \left\{ \max \left\{ \frac{u}{\rho_*}, 0 \right\}, 1 \right\} \quad (8a)$$

$$w_2(u, v_1, v_2) = \min \left\{ u, \max\{u + \eta v_1, 0\}, \max\{u + \eta v_2, 0\} \right\}. \quad (8b)$$

The model (6) is a conservation law that describes the movement of data entering the system at $z = 0$ and exiting at $z = 1$. The flux Φ is the composition of the two scaled throttling functions: First w_2 determines the amount of data available based on the local density and the density of left and right neighbors. Once w_2 is specified, w_1 computes the effective processing speed assuming a linear ramp, where the parameter $\rho_* > 0$ is the minimum data density needed to operate at full capacity. The parameter η controls the degree of throttling due to the lack of available data in neighboring processors.

At first glance, (6) appears to include a recursive definition of the flux function Φ due to equation (7a). However, after integrating (6a) over $(s, z) \in [0, t] \times [0, 1]$ and applying (6c) and (6b), the result can substituted into (7a) to find the following formula for P :

$$P(t, x, z) = \int_0^t F^{\text{in}}(s, x) ds + \int_0^1 \rho_0(x, z) dz - \int_0^z \rho(t, x, z) dz. \quad (9)$$

Alternatively, since $\partial_z P(t, x, z) = -\rho(t, x, z)$, (6a) can be rewritten as a closed Hamilton-Jacobi equation for P :

$$\partial_t P - \Phi(-\partial_z P, \partial_x P) = 0. \quad (10)$$

The reformulation of (6a) in terms of (10) is the basis for the analysis and simulation performed in [5]. In the current work, we proceed by analyzing (6) directly in terms of ρ .

3. Qualitative properties of the mathematical model. In this section, we investigate qualitative properties of (6). We begin by simplifying the expression for Φ in (7b) using the auxiliary function

$$W_2(s) = \min\{\bar{w}_2(s), \bar{w}_2(-s)\}, \quad (11)$$

where

$$\bar{w}_2(s) = \min\{1, \max\{1 + \eta s, 0\}\}; \quad (12)$$

(see Figure 1). In terms of w_1 and W_2 , Φ takes the form

$$\Phi(\rho, \sigma) = \begin{cases} \alpha w_1 \left(\rho W_2 \left(\frac{\sigma}{\rho} \right) \right), & \rho \neq 0 \\ 0, & \rho = 0. \end{cases} \quad (13)$$

Next for simplicity, we artificially extend the domain in z to $[0, \infty)$ and assume initial data with compact support. This assumption allows us to ignore the contribution of the outflow in (7a). Under this assumption, the definition of P is modified as follows:

$$\textbf{Assumption : } P(t, x, z) = \int_z^\infty \rho(t, x, y) dy \quad \forall t \geq 0, x \in \mathbb{T}, z \in [0, \infty). \quad (14)$$

Hence, in the following we consider the model (6) on the extended domain

$$\mathbf{D} := \mathbb{R}^+ \times \mathbb{T} \times \mathbb{R}^+ \quad (15)$$

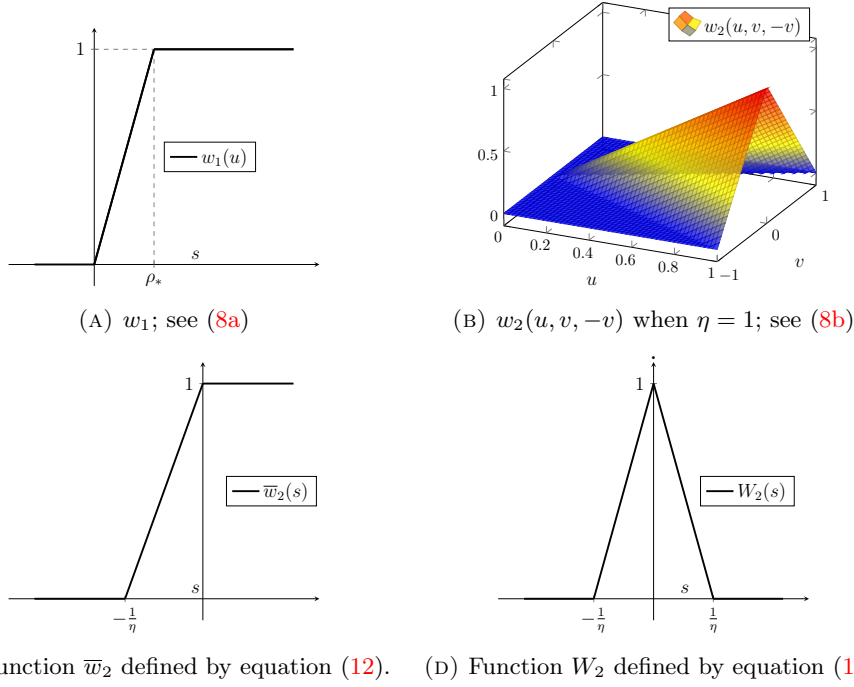


FIGURE 1. Shape of the throttling functions used in the definition of the flux Φ in (13).

and introduce the auxiliary variable

$$\sigma(t, x, z) := \partial_x \int_z^\infty \rho(t, x, y) dy, \quad (t, x, z) \in \mathbf{D}. \quad (16)$$

The resulting system for $(\rho, \sigma): \mathbf{D}^2 \rightarrow \mathbb{R}^2$ is then

$$\partial_t \rho + \partial_z \Phi(\rho, \sigma) = 0, \quad (t, x, z) \in \mathbf{D}, \quad (17a)$$

$$\partial_x \rho + \partial_z \sigma = 0, \quad (t, x, z) \in \mathbf{D}, \quad (17b)$$

$$\rho(0, x, z) = \rho_0(x, z), \quad (x, z) \in \mathbb{T} \times \mathbb{R}^+, \quad (17c)$$

$$\rho(t, x, 0) = \rho_b(t, x), \quad \sigma(t, x, 0) = \sigma_b(t, x), \quad (t, x) \in \mathbb{R}^+ \times \mathbb{T}. \quad (17d)$$

The relation of the quantities (ρ, σ) to the original model (6), modified by the previous Assumption, is as follows:

$$\sigma_b(t, x) = \partial_x \int_0^\infty \rho(t, x, y) dy, \quad (18a)$$

$$\Phi(\rho_b(t, x), \sigma_b(t, x)) = F^{\text{in}}(t, x), \quad (18b)$$

$$\lim_{z \rightarrow \infty} \sigma(t, x, z) = 0. \quad (18c)$$

Definition 3.1 (Weak solution). Given initial data $\rho_0 \in L^\infty(\mathbb{T} \times \mathbb{R}^+)$ and boundary data $(\rho_b, \sigma_b) \in L^\infty(\mathbb{R}^+ \times \mathbb{T})^2$, we call $(\rho, \sigma) \in [L^\infty(\mathbf{D})]^2$ a weak solution of (17) if

for all smooth compactly supported functions $\varphi : \mathbf{D} \rightarrow \mathbb{R}$ and $\psi : \mathbf{D} \rightarrow \mathbb{R}$,

$$\begin{aligned} \int_{\mathbf{D}} \rho \partial_t \varphi + \Phi(\rho, \sigma) \partial_z \varphi \, dz \, dx \, dt + \int_{\mathbb{T} \times \mathbb{R}^+} \rho_0(x, z) \varphi(0, x, z) \, dz \, dx \\ + \int_{\mathbb{R}^+ \times \mathbb{T}} \Phi(\rho_b(t, x), \sigma_b(t, x)) \varphi(t, x, 0) \, dx \, dt = 0, \end{aligned} \quad (19a)$$

$$\int_{\mathbf{D}} \rho \partial_x \psi + \sigma \partial_z \psi \, dz \, dx \, dt + \int_{\mathbb{R}^+ \times \mathbb{T}} \sigma_b(t, x) \psi(t, x, 0) \, dx \, dt = 0. \quad (19b)$$

We do not know whether a solution in the sense of Definition 3.1 exists for (17) with general initial and boundary data or whether such a solution is unique and depends continuously on the data.² However, for the particular initial and boundary data given in (28), we establish the existence of special solutions below.

The weak formulation (19) gives rise to a Rankine–Hugoniot jump condition [10]. Consider the surface

$$S := \{(t, x, z) \in \mathbf{D} : z = \zeta(t, x)\} \quad (20)$$

for a differentiable function $\zeta : \mathbb{R}_0^+ \times \mathbb{T} \rightarrow \mathbb{R}_0^+$ and assume that the functions ρ and σ satisfy (17) point-wise in the interior of $\mathbf{D} \setminus S$. Then standard arguments (see for example [13, Section 11.1.1]) can be used to establish that

$$\Phi(\rho_\ell, \sigma_\ell) - \Phi(\rho_r, \sigma_r) = \partial_t \zeta(t, x) (\rho_\ell - \rho_r), \quad (21)$$

where

$$\rho_\ell = \rho(t, x, \zeta(t, x)^-), \quad \sigma_\ell = \sigma(t, x, \zeta(t, x)^-), \quad (22)$$

$$\rho_r = \rho(t, x, \zeta(t, x)^+), \quad \sigma_r = \sigma(t, x, \zeta(t, x)^+). \quad (23)$$

3.1. Special solutions for front propagation. In this section, we investigate solutions to (17) that take the form of fronts; see Figure 2. As discussed in the introduction, such solutions are important from the point of view of applications.

3.1.1. Existence of Solutions. In the case of a single front, the initial and boundary conditions for ρ take the form

$$\rho_0(x, z) = rH(\zeta_0(x) - z) \quad \text{and} \quad \rho_b(t, x) = r, \quad (24)$$

where H denotes the Heaviside function and $r > 0$ is a positive constant. In addition, we enforce the consistency of σ_b as prescribed in (18a):

$$\sigma_b(t, x) = r \partial_x \zeta(t, x). \quad (25)$$

Theorem 3.2. *Let $\rho_*, \eta > 0$ be positive constants and let $\alpha = \alpha(t, x) \in C^1(\mathbb{R}^+ \times \mathbb{T})$ be a strictly positive function. Given initial and boundary conditions of the form (24) and (25), where*

$$0 < r < \rho_* \quad (26)$$

and $\zeta_0 \in C^1(\mathbb{T})$ is non-negative, assume that there is a C^1 solution $\zeta : \mathbb{R}_0^+ \times \mathbb{T} \rightarrow \mathbb{R}^+$ that satisfies

$$\partial_t \zeta(t, x) = \frac{\alpha(t, x)}{\rho_*} W_2(\partial_x \zeta(t, x)), \quad \zeta(0, x) = \zeta_0(x). \quad (27)$$

²In [5], a continuous vanishing viscosity solution is established using known results from the mathematical literature on Hamilton–Jacobi equations. However, translating such a result to an L^1 -theory for ρ would require P to be absolutely continuous. While possible, this additional regularity has not yet been determined.




FIGURE 2. Left: Contour plot of the density $\rho(t, x, z)$ for a fixed time t . Processed data with constant density $r > 0$ is depicted in blue up to a stage of completion $z = \zeta(t, x)$. Zero data (grey) is prescribed for completion stages $z > \zeta(t, x)$, c.f. equation (28). Right: Similar plot of a density with constant values r_1 and r_2 and regions separated by functions $\zeta_1(t, x)$ and $\zeta_2(t, x)$.

Then there pair (ρ, σ) given by

$$\rho(t, x, z) = rH(\zeta(t, x) - z), \quad \sigma(t, x, z) = rH(\zeta(t, x) - z)\partial_x \zeta(t, x) \quad (28)$$

is a weak solution for (17) in the sense of Definition 3.1 with initial and boundary conditions given by (24) and (25).

Before proving this result, some remarks are in order.

1. The condition in (26) guarantees that $\frac{r}{\rho_*} W_2\left(\frac{\sigma}{r}\right) < 1$. In this case the flux function Φ (c.f. (13)) is given by the simplified formula

$$\Phi(\rho, \sigma) = \begin{cases} \frac{\alpha\rho}{\rho_*} W_2\left(\frac{\sigma}{\rho}\right), & \rho \neq 0, \\ 0, & \rho = 0. \end{cases} \quad (29)$$

This formula and the Rankine–Hugoniot condition in (21) together form the key components of the existence proof.

2. While the condition $r < \rho_*$ is sufficient to obtain (29), it is not necessary. Indeed $W_2(\sigma/\rho) = 0$ for σ sufficiently large; in particular, (29) holds at any (x, t) such that $\partial_x \zeta(t, x) > \eta^{-1}$.
3. It is currently open as to whether equation (27) has a C^1 solution. However, since α, ρ_* , and W_2 are non-negative, it is clear that any such solution will be non-decreasing. This fact is consistent with the notion that data is always processed toward completion.
4. A result similar to Theorem 3.2 can be obtained for initial data separated by a finite number of non-intersecting fronts. Consider for example the case of two fronts (see Figure 2-b):

$$\rho_0(x, z) = r_1 H(\zeta_{1,0}(x) - z) + r_2 [H(\zeta_{2,0}(x) - z) - H(\zeta_{1,0}(x) - z)] \quad (30)$$

with positive constants $r_1 \neq r_2$ and $r_1 < \rho_*$ and $r_2 < \rho_*$. If $0 < \zeta_{1,0}(x) < \zeta_{2,0}(x)$ for all $x \in \mathbb{T}$, then there is a weak solution (ρ, σ) in the sense of

Definition 3.1 of the form

$$\rho(t, x, z) = (r_1 - r_2)H(\zeta_1(t, x) - z) + r_2H(\zeta_2(t, x) - z) \quad (31a)$$

$$\sigma(t, x, z) = (r_1 - r_2)\partial_x \zeta_1(t, x)H(\zeta_1(t, x) - z) \quad (31b)$$

$$+ r_2\partial_x \zeta_2(t, x)H(\zeta_2(t, x) - z), \quad (31c)$$

provided $\zeta_1(0, x) = \zeta_{1,0}(x)$, $\zeta_2(0, x) = \zeta_{2,0}(x)$, and

$$\partial_t \zeta_1(t, x) \quad (32a)$$

$$= \frac{\alpha(t, x)}{\rho_*(r_2 - r_1)} \left[r_2 W_2(\partial_x \zeta_2(t, x)) - r_1 W_2\left(\left(\frac{r_1 - r_2}{r_1}\right)\partial_x \zeta_1(t, x) + \frac{r_2}{r_1}\partial_x \zeta_2(t, x)\right) \right],$$

$$\partial_t \zeta_2(t, x) = \frac{\alpha(t, x)}{\rho_*} W_2(\partial_x \zeta_2(t, x)). \quad (32b)$$

Proof of Theorem 3.2. With the C^1 assumption on ζ , (ρ, σ) given by equation (28) clearly belongs to $L^\infty(\mathbf{D})^2$. By construction the given solution (28) fulfills boundary and initial condition pointwise and hence also in weak form. It therefore remains to verify that (ρ, σ) is a weak solution in the interior of \mathbf{D} .

We first verify that (ρ, σ) fulfills (19b). For ψ smooth and compactly supported in the interior of \mathbf{D} , we need to show that:

$$\int_{\mathbf{D}} [rH(\zeta(t, x) - z)\partial_x \psi(t, x, z) + rH(\zeta(t, x) - z)\partial_x \zeta(t, x)\partial_z \psi(t, x, z)] dz dx dt = 0. \quad (33)$$

By definition of the Heaviside function,

$$\begin{aligned} & \int_{\mathbf{D}} [rH(\zeta(t, x) - z)\partial_x \psi(t, x, z) + rH(\zeta(t, x) - z)\partial_x \zeta(t, x)\partial_z \psi(t, x, z)] dz dx dt \\ &= r \int_{\mathbb{R}_0^+ \times \mathbb{T}} \int_0^{\zeta(t, x)} [\partial_x \psi(t, x, z) + \partial_x \zeta(t, x)\partial_z \psi(t, x, z)] dz dx dt \\ &= r \int_{\mathbb{R}_0^+ \times \mathbb{T}} \left[\int_0^{\zeta(t, x)} \partial_x \psi(t, x, z) dz + \partial_x \zeta(t, x)\psi(t, x, \zeta(t, x)) \right] dx dt \\ &= r \int_{\mathbb{R}_0^+ \times \mathbb{T}} \frac{\partial}{\partial x} \left(\int_0^{\zeta(t, x)} \psi(t, x, z) dz \right) dx dt = 0, \end{aligned} \quad (34)$$

where the integral in the last line vanishes due to periodicity of the domain with respect to x .

We next verify that (ρ, σ) fulfills (19a). Let $S = \{(t, x, z) \in \mathbf{D} : z = \zeta(t, x)\} \subset \mathbf{D}$ be the surface of the front described by ζ , and let $\mathbf{S}_+ := \{(t, x, z) \in \mathbf{D} : z > \zeta(t, x)\}$ and $\mathbf{S}_- := \{(t, x, z) \in \mathbf{D} : z < \zeta(t, x)\}$. For $z < \zeta(t, x)$ the explicit form of (ρ, σ) implies the following relation

$$\frac{\sigma(t, x, z)}{\rho(t, x, z)} = \partial_x \zeta(t, x). \quad (35)$$

Hence according to (29),

$$\Phi(\rho(t, x, z), \sigma(t, x, z)) = \begin{cases} \frac{\alpha(t, x)r}{\rho_*} W_2(\partial_x \zeta(t, x)) & (t, x, z) \in \mathbf{S}_- \\ 0 & (t, x, z) \in \mathbf{S}_+. \end{cases} \quad (36)$$

Let φ be any smooth, compactly supported function in the interior of \mathbf{D} . In \mathbf{S}_\pm , $\Phi(\rho, \sigma)$ and ρ are constant, and in \mathbf{S}_+ , they both vanish. Hence,

$$\begin{aligned} \int_{\mathbf{D}} \rho \partial_t \varphi + \Phi(\rho, \sigma) \partial_z \varphi \, dz \, dx \, dt &= \int_{\mathbf{S}_-} r \partial_t \varphi + \frac{\alpha r}{\rho_*} W_2(\partial_x \zeta) \partial_z \varphi \, dz \, dx \, dt \\ &= \int_{\mathbb{T}} \int_{\{(t, z): z = \zeta(t, x)\}} \varphi \left(-r \partial_t \zeta + \frac{\alpha r}{\rho_*} W_2(\partial_x \zeta) \right) \frac{1}{\|(1, \partial_t \zeta)\|} dA(t, z) dx, \end{aligned} \quad (37)$$

where $dA(t, z)$ denotes the surface measure on the curve $\{(t, z) : z = \zeta(t, x)\}$. The final expression above vanishes due to the definition of ζ given in (27). This finishes the proof. \square

3.1.2. Properties of fronts and their associated weak solutions. It is the equation for the front ζ in (27) that essentially determines the behavior of the special solutions (28). We assume again that $0 < r < \rho_*$. With the definition of W_2 in (11), this equation takes the form

$$\partial_t \zeta(t, x) = \frac{\alpha(t, x)}{\rho_*} \times \begin{cases} 1 - \eta |\partial_x \zeta(t, x)|, & |\partial_x \zeta(t, x)| < \eta^{-1}, \\ 0, & |\partial_x \zeta(t, x)| \geq \eta^{-1}. \end{cases} \quad (38)$$

In general, we do not know if there exists a solution to (38), due to the discontinuous flux. However, we may consider particular solutions related to special initial conditions $\zeta_0(x)$. Those special solutions are used as test cases for numerical simulations in the next section.

- If for all $(t, x) \in \mathbb{R}^+ \times \mathbb{T}$, $|\partial_x \zeta(t, x)| < \eta^{-1}$ and $\alpha(t, x) = \alpha$ is constant, then formally $\partial_x \zeta(t, x)$ fulfills the conservation law

$$\partial_t(\partial_x \zeta)(t, x) + \frac{\alpha \eta}{\rho_*} \partial_x |\partial_x \zeta(t, x)| = 0, \quad (\partial_x \zeta)(0, x) = \partial_x \zeta_0(x) \quad (39)$$

with flux function $f(u) = C|u|$ and $C = \frac{\alpha \eta}{\rho_*}$. It is known that there exists a weak (entropic) solution to (39) as long as $\partial_x \zeta_0(x)$ is in L^∞ , has locally bounded variation, and is continuous from the left [9, Proposition 3.1]. Moreover, if $\partial_x \zeta_0(x)$ is piecewise constant, this solution is given by piecewise constant states separated by a finite number of traveling discontinuities [9, Lemma 3.2]. We investigate such piece-wise solutions in Examples 2 and 3 in the numerical results of Section 4.

- If α is constant, and if $0 \leq \partial_x \zeta_0(x) < \eta^{-1}$ locally in x , then for sufficiently small t ,

$$\zeta(t, x) = \zeta_0 \left(x - \eta \frac{\alpha}{\rho_*} t \right) + \frac{\alpha}{\rho_*} t. \quad (40)$$

Similarly if $-\eta^{-1} < \partial_x \zeta_0(x) \leq 0$ locally, then for sufficiently small t ,

$$\zeta(t, x) = \zeta_0 \left(x + \eta \frac{\alpha}{\rho_*} t \right) + \frac{\alpha}{\rho_*} t. \quad (41)$$

We use these formulas to make comparisons with numerical results in Examples 2-4 in Section 4.

- If locally $|\partial_x \zeta_0(x)| \geq \eta^{-1}$, then $\zeta(t, x) = \zeta_0(x)$ for t sufficiently small. In this case the model predicts that all processors are stalled. However, the condition on ζ_0 cannot hold globally due to the periodicity assumption in x . We explore the behavior of the model for this scenario in Example 5 in Section 4.

4. Numerical scheme and computational simulations. In this section, we construct a numerical discretization for the system. Results of numerical simulations using this discretization are presented to illustrate the behavior of special front-type solutions discussed in the previous section.

4.1. Numerical discretization. Our strategy for approximating (17) is based on the relaxed system of equations

$$\partial_t \rho + \partial_z \Phi(\rho, \sigma) = 0, \quad (42a)$$

$$\varepsilon \partial_t \sigma + \partial_x \rho + \partial_z \sigma = 0, \quad (42b)$$

where $\varepsilon > 0$ is a small parameter. Formally, (17a) and (17b) are obtained from (42) in the limit as $\varepsilon \rightarrow 0$. This limit is considered later after a suitable discretization of (42).

Before discretizing (42), we investigate hyperbolicity in the simplified case that $\alpha(t, x) > 0$ is constant.³ In this case, (42) takes the non-conservative form

$$\partial_t U + B^x(U) \partial_x U + B^z(U) \partial_z U = 0, \quad (43)$$

where $U = (\rho, \sigma)$ and

$$B^x(U) = \begin{pmatrix} 0 & 0 \\ \frac{1}{\varepsilon} & 0 \end{pmatrix}, \quad B^z(U) = \begin{pmatrix} \partial_\rho \Phi & \partial_\sigma \Phi \\ 0 & \frac{1}{\varepsilon} \end{pmatrix}. \quad (44)$$

Under the assumption (26) and provided that $\rho > 0$, the simplified form of Φ in (29) implies that

$$\partial_\rho \Phi(\rho, \sigma) = \frac{\alpha}{\rho_*} (W_2(s) - s W_2'(s)), \quad \partial_\sigma \Phi(\rho, \sigma) = \frac{\alpha}{\rho_*} W_2'(s), \quad (45)$$

where $s = \sigma/\rho$.

The system (43) is hyperbolic if for any $\xi = (\xi_1, \xi_2) \in \mathbb{R}^2$ the matrix $B(U, \xi) = \xi_1 B^x(U) + \xi_2 B^z(U)$ is diagonalizable in the field of real numbers [13]. Because the function W_2 is not smooth, $B(U, \xi)$ is not differentiable when $s = 0$ or $s = \pm 1/\eta$. However, away from these points,

$$B(U, \xi) = \begin{cases} \begin{pmatrix} 0 & 0 \\ \frac{\xi_1}{\varepsilon} & \frac{\xi_2}{\varepsilon} \end{pmatrix}, & |s| > \frac{1}{\eta}, \\ \left[\begin{pmatrix} \xi_2 \frac{\alpha}{\rho_*} & -\xi_2 \operatorname{sgn}(s) \eta \frac{\alpha}{\rho_*} \\ \frac{\xi_1}{\varepsilon} & \frac{\xi_2}{\varepsilon} \end{pmatrix} \right], & 0 < |s| < \frac{1}{\eta}, \end{cases} \quad (46)$$

and the eigenvalues of $B(U, \xi)$ are

$$(\lambda_1, \lambda_2) = \begin{cases} \left(0, \frac{\xi_2}{\varepsilon} \right), & |s| > \frac{1}{\eta}, \\ (\lambda_1^*, \lambda_2^*), & 0 < |s| < \frac{1}{\eta}, \end{cases} \quad (47)$$

³Hyperbolicity will not change if α is space or time-dependent, provided it is sufficiently smooth.




FIGURE 3. Domain of hyperbolicity of system (42). In both cases, the non-hyperbolic region (gray) is bounded by the ξ_1 -axis and by $\xi_2 = \pm C\xi_1$ with slope $C = \frac{4\epsilon\alpha\rho_*\eta}{(\epsilon\alpha - \rho_*)^2} \rightarrow 0$ as $\epsilon \rightarrow 0^+$.

where $(\lambda_1^*, \lambda_2^*)$ are the roots of the characteristic polynomial

$$p(\lambda) = \lambda^2 - \xi_2 \left(\frac{\alpha}{\rho_*} + \frac{1}{\epsilon} \right) \lambda + \frac{\xi_2 \alpha}{\epsilon \rho_*} (\xi_2 + \text{sgn}(s) \xi_1 \eta). \quad (48)$$

For $|s| > \eta^{-1}$, the system is strictly hyperbolic. For $0 < |s| < \eta^{-1}$, the polynomial p in (48) has discriminant

$$d_\epsilon(\xi_1, \xi_2) = \left(\frac{\xi_2}{\epsilon \rho_*} \right)^2 \left[(\epsilon\alpha - \rho_*)^2 - 4 \text{sgn}(s) \epsilon \alpha \rho_* \eta \frac{\xi_1}{\xi_2} \right] \quad (49)$$

which, for $\xi_2 \neq 0$, is negative (so that λ_1^* and λ_2^* have non-zero imaginary component) if and only if

$$\text{sgn}(s) \frac{\xi_1}{\xi_2} > \frac{(\epsilon\alpha - \rho_*)^2}{4\epsilon\alpha\rho_*\eta}. \quad (50)$$

For any $\epsilon > 0$, the non-hyperbolic region $\{(\xi_1, \xi_2) : d_\epsilon(\xi_1, \xi_2) < 0\}$ is nontrivial; see Figure 3. However, for any fixed (ξ_1, ξ_2) , there exists an $\epsilon^*(\xi_1, \xi_2)$ small enough that $d_\epsilon(\xi_1, \xi_2) \geq 0$ for all $\epsilon < \epsilon^*(\xi_1, \xi_2)$. Hence as $\epsilon \rightarrow 0^+$, the non-hyperbolic region vanishes. The numerical method below is constructed by discretizing the relaxation system in (42) and then formally setting $\epsilon = 0$. This fact partially justifies the use of (42) even though it is not everywhere hyperbolic when $\epsilon > 0$.

To derive a numerical scheme for (42), we first discretize in time using a first-order method. We are eventually interested in the $\epsilon \rightarrow 0$ limit of the fully discretized scheme; hence the equation for ρ in (42a) is treated explicitly, but the equation for σ in (42b) is treated implicitly. Let $\rho^n(x, z) \simeq \rho(n\Delta t, x, z)$ and $\sigma^n(x, z) \simeq \sigma(n\Delta t, x, z)$ be the discrete approximations of ρ and σ , respectively, at time $n\Delta t$. We set

$$\rho^{n+1} = \rho^n - \Delta t \partial_z \Phi(\rho^n, \sigma^n), \quad (51a)$$

$$\sigma^{n+1} = \sigma^n - \frac{\Delta t}{\epsilon} (\partial_x \rho^{n+1} + \partial_z \sigma^{n+1}). \quad (51b)$$

To discretize (51) in x and z , we introduce a bounded computational domain $(x, z) \in [0, 1] \times [0, 1]$, which is divided into $N^x \times N^z$ uniform cells of size $\Delta x \times \Delta z$. The cell centers are denoted by (x_i, z_j) , for $i \in \{1, \dots, N^x\}$ and $j \in \{1, \dots, N^z\}$; and the cell edges are denoted by $x_{i+\frac{1}{2}} = (i + \frac{1}{2})\Delta x$ and $z_{j+\frac{1}{2}} = (j + \frac{1}{2})\Delta z$, respectively, for $i \in \{0, \dots, N^x\}$ and $j \in \{0, \dots, N^z\}$. The approximate cell averages of ρ^n and σ^n on the cell centered at (x_i, z_j) are denoted by

$$\bar{R}_{ij}^n \simeq \frac{1}{\Delta x \Delta z} \int_{x_{i-1/2}}^{x_{i+1/2}} \int_{z_{j-1/2}}^{z_{j+1/2}} \rho^n(x, z) dx dz \quad (52)$$

and

$$\bar{S}_{ij}^n \simeq \frac{1}{\Delta x \Delta z} \int_{x_{i-1/2}}^{x_{i+1/2}} \int_{z_{j-1/2}}^{z_{j+1/2}} \sigma^n(x, z) dx dz, \quad (53)$$

respectively.

We use a Lax-Friedrichs approximation of $\partial_z \Phi$ in (51a) and in (51b) a centered discretization for $\partial_x \rho$ and a one-sided discretization of $\partial_z \sigma$. The resulting fully discrete scheme is

$$\bar{R}_{ij}^{n+1} = \bar{R}_{ij}^n - \frac{\Delta t}{\Delta z} (F_{i,j+1/2}^n - F_{i,j-1/2}^n) + \frac{1}{4} (\bar{R}_{i+1,j}^n - 2\bar{R}_{ij}^n + \bar{R}_{i-1,j}^n), \quad (54a)$$

$$\bar{S}_{ij}^{n+1} = \bar{S}_{ij}^n - \frac{\Delta t}{2\epsilon \Delta x} (\bar{R}_{i+1,j}^{n+1} - \bar{R}_{i-1,j}^{n+1}) - \frac{\Delta t}{\epsilon \Delta z} (\bar{S}_{i,j+1}^{n+1} - \bar{S}_{ij}^{n+1}), \quad (54b)$$

where

$$F_{i,j+\frac{1}{2}}^n = \frac{1}{2} (\Phi(\bar{R}_{i,j+1}^n, \bar{S}_{i,j+1}^n) + \Phi(\bar{R}_{ij}^n, \bar{S}_{ij}^n) - a(\bar{R}_{i,j+1}^n - \bar{R}_{ij}^n)). \quad (55)$$

The monotonicity of the numerical flux $F_{i,j+\frac{1}{2}}^n$ is guaranteed provided that

$$\max_{(R,S)} \left\| \partial_\rho \Phi(R, S) \right\| =: a \leq \frac{\Delta z}{\Delta t}. \quad (56)$$

This condition motivates a dynamic choice of Δt : At each time level t^n we set $\Delta t = \frac{\Delta z}{a^n}$ for

$$a^n = \max_{(i,j): 1 \leq i \leq N^x, 1 \leq j \leq N^z} \left\| \partial_\rho \Phi(\bar{R}_{ij}^n, \bar{S}_{ij}^n) \right\|. \quad (57)$$

The right-biased stencil in the discretization of $\partial_z \sigma$ in equation (54b) warrants some discussion. Indeed, given that the boundary condition for σ is given at $z = 0$, it seems more natural to use a left-biased stencil. However, such an approach does not allow the boundary condition to be computed implicitly via the consistency relation in (18a). However, if ρ has compact support, then (18c) holds. As long as the support of the solution does not reach the boundary of the computational domain at $z = 1$, this condition can be enforced there, in an implicit fashion. This approach is consistent with the fact that σ at a given z_* is determined entirely by ρ at values of $z > z_*$. Moreover, numerical calculations confirm that enforcing the boundary condition in this way yields stable results.

The final scheme used in Section 4.2 is given by the formal limit of (54) in the limit $\epsilon \rightarrow 0$:

$$\bar{R}_{ij}^{n+1} - \bar{R}_{ij}^n = -\frac{\Delta t}{\Delta z} [F_{i,j+1/2}^n - F_{i,j-1/2}^n] + \frac{1}{4} (\bar{R}_{i+1,j}^n - 2\bar{R}_{ij}^n + \bar{R}_{i-1,j}^n), \quad (58a)$$

$$\bar{S}_{ij}^{n+1} = \frac{\Delta z}{2\Delta x} (\bar{R}_{i+1,j}^{n+1} - \bar{R}_{i-1,j}^{n+1}) + \bar{S}_{i,j+1}^{n+1}. \quad (58b)$$

This scheme above must be accompanied by boundary and initial conditions. The initial values \bar{R}_{ij}^0 are obtained by integration of the initial data ρ_0 :

$$\bar{R}_{ij}^0 = \frac{1}{\Delta x \Delta z} \int_{x_{i-1/2}}^{x_{i+1/2}} \int_{z_{j-1/2}}^{z_{j+1/2}} \rho_0(x, y) dx dz. \quad (59)$$

The cell averages \bar{S}_{ij}^0 are obtained by applying a midpoint rule to the consistency relation (18a):

$$\bar{S}_{ij}^0 = \frac{1}{\Delta x} \sum_{\ell=j}^{N^z} \rho_0(x_{i+1/2}, z_j) - \rho_0(x_{i-1/2}, z_j). \quad (60)$$

For \bar{S}_{ij}^n , zero boundary conditions are prescribed at $j = N^z$ consistent with the right-based stencil in the discretization of $\partial_z \sigma$:

$$\bar{S}_{i, N^z}^n = 0. \quad (61)$$

For \bar{R}_{ij}^n , boundary conditions at $z = 0$ are computed from $\rho_b(t)$. For the special solutions discussed in the previous section $\rho_b = r$ and hence

$$\bar{R}_{i,1}^n = r. \quad (62)$$

Due to the central discretization of $\partial_z \rho$, boundary conditions for \bar{R}_{ij}^n at $j = N^z$ need to be prescribed. Since in the simulation no data reaches this boundary we implement Neumann boundary conditions

$$\bar{R}_{i, N^z}^n = \bar{R}_{i, N^z-1}^n. \quad (63)$$

4.2. Simulation results.

Example 1: validation with the microscopic model. We consider a test problem that has been used to validate the macroscopic model (17) and the numerical method (58) by comparing with a simulation of the microscopic model introduced in [5]. For this problem, $\rho_* = 1.0$ and $\eta = 1.0$. The initial and boundary conditions are given by

$$\rho_0(x, z) = 1.5 (\sin(2\pi z))^6 \chi_{[0, 0.5]}(z), \quad \rho_b(t, x) = 0, \quad \sigma_b(t, x) = 0, \quad (64)$$

and the processor speed is

$$\alpha(x) = 1 - 0.4(\sin(\pi x))^2. \quad (65)$$

Both models are simulated up to a final time $T_{\text{fin}} = 0.5$.

The (x, z) domain for the macroscopic model is discretized using $N^x = N^z = 800$ cells, and simulated with the algorithm in (58). In order to be consistent with the choice of η , the microscopic model uses $i_{\text{max}} = 800$ processors and $k_{\text{max}} = 800$ stages. The microscopic model is integrated in time using the explicit Euler method.

In Figure 4 we show the initial density and the final density profiles for both models at time T_{fin} . We observe qualitative agreement between the microscopic and the macroscopic solutions. In both cases, the data around $x = 0.5$ is processed more slowly than the data in the rest of the domain, due to a slower processing rate α there.

In the remaining examples of this subsection, we explore the behavior of front-type solutions. Unless otherwise stated the following parameters are used in all simulations:

$$r = 0.5, \quad \rho_* = 0.8, \quad \alpha = 0.1, \quad T_{\text{fin}} = 2.0. \quad (66)$$

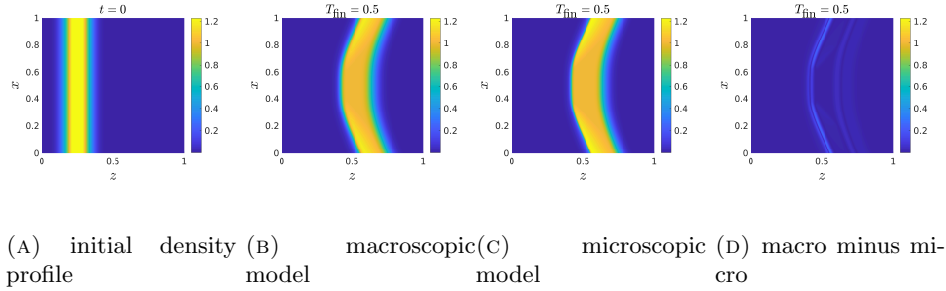


FIGURE 4. Density profiles for Example 1.

The value of T_{fin} ensures that data does not reach $z = 1$; in particular the assumption in (14) holds. These parameter choices, along with the initial and boundary conditions ensure that

$$0 \leq \rho(t, x, z) < \rho_*, \quad \text{for all } (t, x, z) \in \mathbf{D}, \quad (67)$$

in which case the flux Φ is given by the simplified formula in (29).

Example 2: front with constant profile. We consider the case of an initial constant front, namely

$$\zeta_0(x) = \zeta_0 = 0.2, \quad \forall x \in \mathbb{T}, \quad (68)$$

which provides a trivial example for a piece-wise constant solution $\partial_x \zeta$ to (39). Moreover, the evolution of this front is given explicitly in (40) and (41):

$$\zeta(t, x) = \zeta_0 + \frac{\alpha}{\rho_*} t. \quad (69)$$

In this example η is set to 0.5 although the solution in (69) is independent of this choice.

Numerical results are shown in Figure 5. Since the processing rate α is constant the front moves with the same speed for each $x \in [0, 1)$, as expected. Moreover, the analytical and the numerical solutions provide the same position of the front at final time.

Example 3: V-shaped front with small η . We consider an initial condition characterized by a V-shape front under the condition that $\eta|\partial_x \zeta| < 1$; cf. (38). This condition models the physical situation in which processors are not initially at the same stage of completion. Such a situation may be realized if a local group of processors is, at some previous time, slower than the rest [5]. Specifically, we let

$$\zeta_0(x) = (1 - 2z_0)|x - 0.5| + z_0 \quad (70)$$

where $z_0 = \zeta_0(0.5) < 0.5$ is a fixed parameter.

Based on the theoretical findings in [9] for (39), we expect that the front at later times will be piecewise linear. Moreover, the formulas in (40) and (41) can be used to generate the local solution analytically. Indeed, under the assumption $\eta|\partial_x \zeta| < 1$, a small calculation shows that the local, short time solution away from $x = 0.5$ is given by

$$\zeta_{\text{loc}}(t, x) = \zeta_0(x) + \frac{\alpha}{\rho_*} \left(1 - \eta|\partial_x \zeta_0|\right) t. \quad (71)$$

where $|\partial_x \zeta_0| = (1 - 2z_0)$. However at $x = 0.5$, Theorem 3.2 does not apply, since the front is not in differentiable there; hence (40) and (41) are not valid. We conjecture

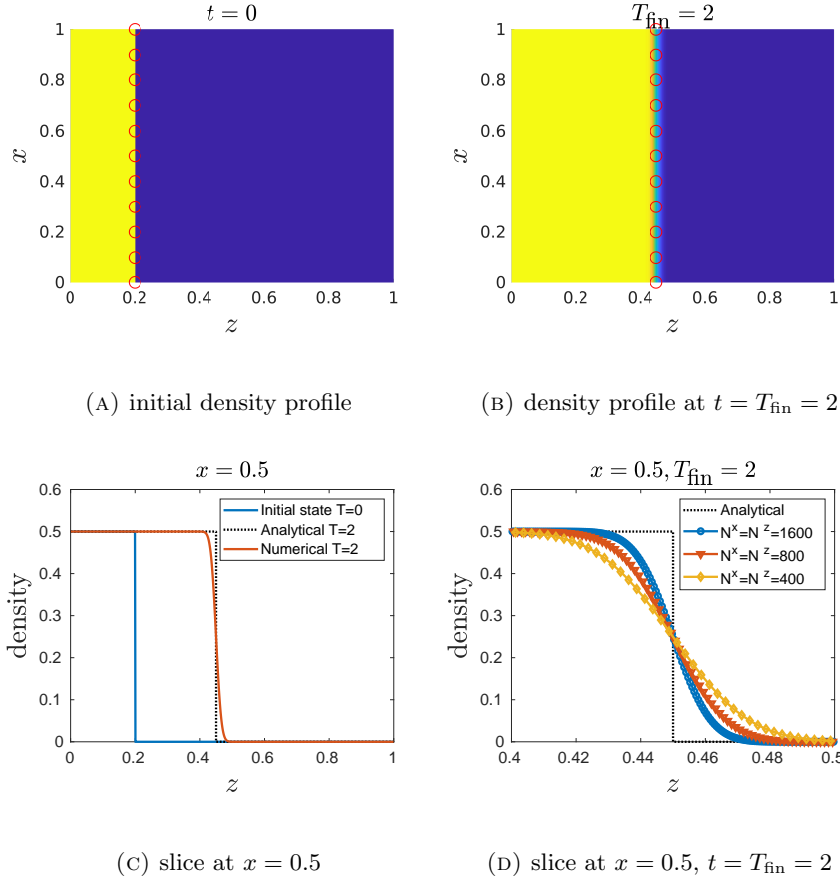


FIGURE 5. Density profiles based on the a simulation of the constant front solution (69) with the scheme in (58). The simulations in panels (a)-(c) use $N^x = N^z = 800$ computational cells, while in panel (d), numerical solutions are computed at different refinement levels. In the top row, the analytical front is marked by circles.

that at this point, the front moves forward at full speed α/ρ_* and, as it catches up with neighboring points in the front, these points also move forward at full speed. This means that the global front solution is given by

$$\zeta(t, x) = \max \left\{ \zeta_{\text{loc}}(t, x), z_0 + \frac{\alpha}{\rho_*} t, \right\}. \quad (72)$$

In Figure 6 we show numerical results for this test problem using the algorithm in (58). We use a grid with $N^x = N^z = 800$ computational cells. We set $z_0 = 0.2$ (so that $|\partial_x \zeta_0(x)| = 0.6$) and let $\eta = 1/1.1$. It is clear from these plots that the analytic solution (72) and the numerical solution provide the same position of the front.

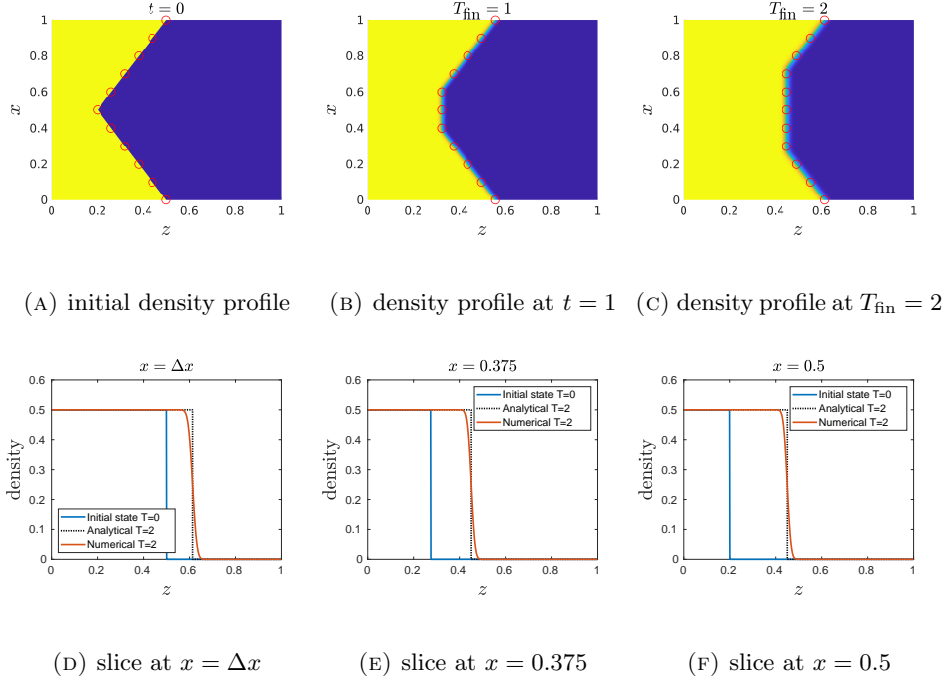


FIGURE 6. Density profiles based on the simulation of the V-shaped initial front (70) with the scheme in (58) using $N^x = N^z = 800$. The analytic solution is computed using (72). In the top row, the analytical front is marked by circles.

Example 4: Smooth front with small η . We repeat the test problem above, except now the initial front is given by a smooth function:

$$\zeta_0(x) = \frac{1}{4} \cos(2\pi x) + 0.3. \quad (73)$$

In order to fulfill the condition $\eta|\partial_x \zeta| < 1$ we choose $\eta^{-1} = \max_{x \in \mathbb{T}} \zeta_0(x) + 1 = \pi/2 + 1$.

In Figure 7 we show numerical results for this test problem using the algorithm in (58). We use a grid with $N^x = N^z = 800$ computational cells and simulate the solution to a time horizon $T_{\text{fin}} = 2$. As in the previous example, (72) can be used to evaluate the location of the front analytically. The results in Figure 7 demonstrate that the numerical and analytical solutions agree even when $\partial_x \zeta_0$ is not piece-wise constant.

Example 5: V-shaped front with large η . We consider again the initial condition (70) characterized by a V-shape front. The model parameters are the same as in Example 3, except that η is now large enough to ensure that $\eta|\partial_x \zeta| > 1$. In particular $|\partial_x \zeta| = 0.6$, but $\eta = 10$. In this case, we expect stalling for those processors $x \neq 0.5$; see (38). Again at $x = 0.5$ the solution is not described by Theorem 3.2, since there the front is not C^1 . We expect that the front at this point will proceed at the maximum speed α/ρ_* and that the stalled points away front $x = 0.5$ will also begin to move at this rate after they have been overtaken. In other words, we

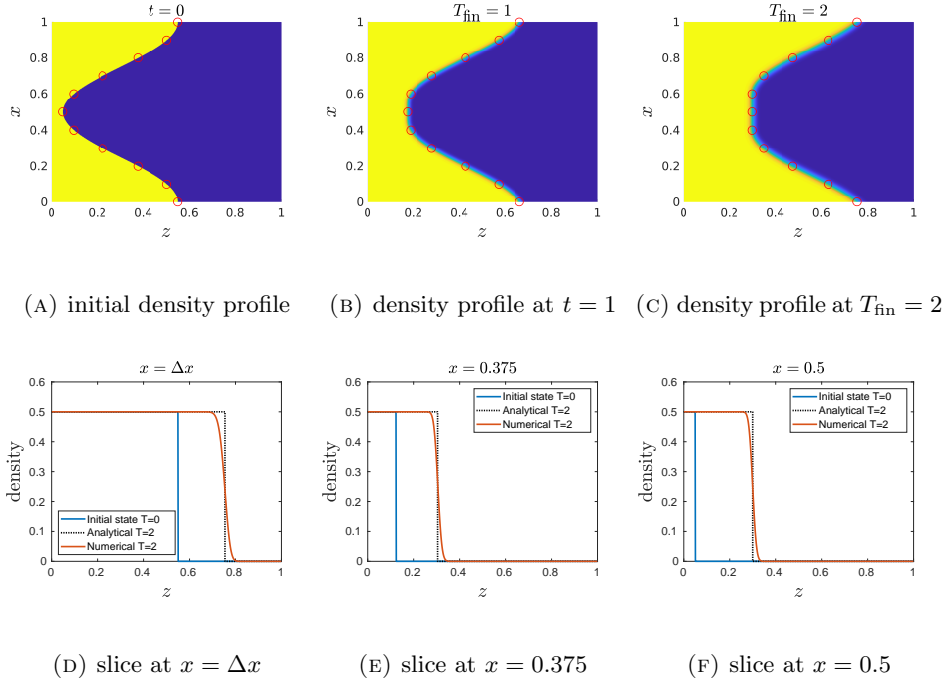


FIGURE 7. Density profiles based on the simulation of the V-shaped initial front in (70), using the scheme in (58) with $N^x = N^z = 800$. The analytic solution is computed using (72). In the top row, the analytical front is marked by circles.

conjecture that the global formula for the solution front is

$$\zeta(t, x) = \max \left\{ \zeta_0(x), z_0 + \frac{\alpha}{\rho_*} t \right\}. \quad (74)$$

In Figure 8 we show numerical results for a grid with $N^x = N^z = 800$ computational cells and a time horizon $T_{\text{fin}} = 2$. These results agree with the analytical solution in (74).

4.3. Variations in the processor rate. The discrete model (1) can be used to predict how variations in the processor rate affect computer performance, and to understand whether these rates can be controlled to achieve a prescribed objective, such as faster throughput. For the continuum model in (17), such a control is specified via the function α . In this section, we briefly explore the effects of adapting α with the intent of a more extensive study in later work.

In order to normalize the control action, we assume that for all t ,

$$\int_{\mathbb{T}} \alpha(t, x) dx = \bar{\alpha}. \quad (75)$$

Since α is non-negative, so too is $\bar{\alpha}$; we assume further that $\bar{\alpha} > 0$. Then several choices for α can be envisioned. The simplest choice is to set α equal to a constant. In this case, the normalization in (75) implies that $\alpha = \bar{\alpha}$. For a front-type solution,

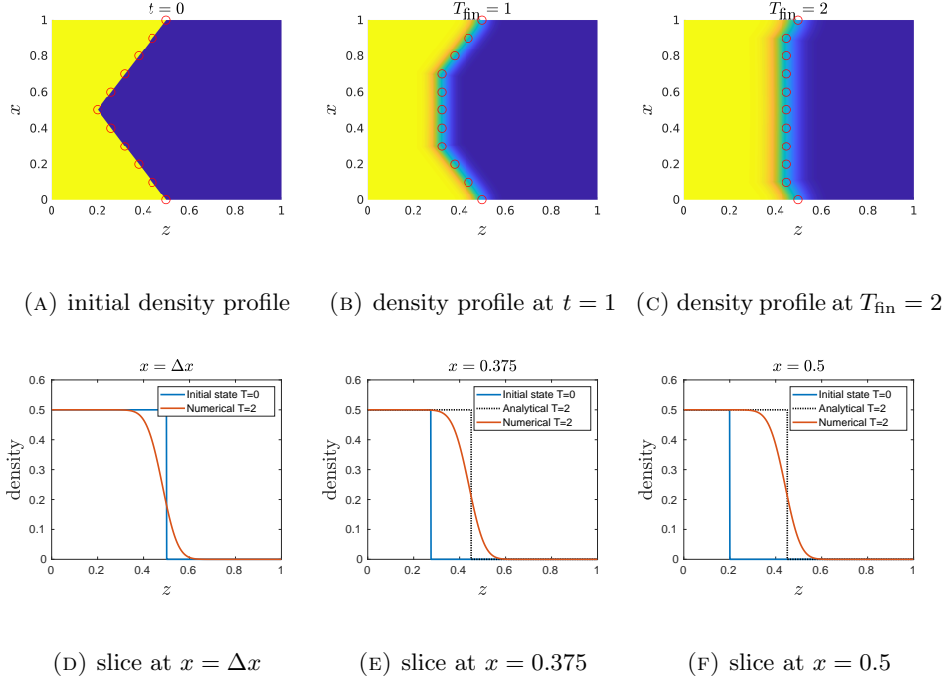


FIGURE 8. Density profiles based on the simulation of the V-shaped initial front (70) with the scheme in (58) using $N^x = N^z = 800$. The analytical solution is computed using (74). In the top row, the analytical front is marked by circles.

this simple choice will be compared with a policy that depends on the initial data

$$\rho_0(x, z) = rH(\zeta_0(x) - z), \quad r < \rho_*, \quad \text{and} \quad \zeta_0(x) > 0, \quad (76)$$

Since the slope $\partial_x \zeta(t, x)$ controls the throttling due to neighboring processors, the condition $\partial_x \zeta(t, x) = 0$ is desirable. We therefore propose a choice for α that gives priority to processors that are trailing in the initial configuration. Specifically, we consider

$$\alpha(x) = C_\alpha \rho_* (\zeta_{\max} - \zeta_0(x)), \quad (77)$$

where $\zeta_{\max} > \max_{x \in \mathbb{T}} \zeta_0(x)$ is a fixed constant and C_α , which depends on ζ_0 , is chosen so that (75) holds. This choice of α will increase the speed of trailing processors

In order to compare the temporal and spatial evolution of ρ based on the policy (77) versus the constant $\alpha = \bar{\alpha}$ policy, we compute some quantities of interest

$$\begin{aligned} \omega_1(t, z) &:= \int_0^t \int_{\mathbb{T}} \Phi(\rho(s, x, z), \sigma(s, x, z)) dx ds, \\ \omega_2(t, z) &:= \int_{\mathbb{T}} \Phi(\rho(t, x, z), \sigma(t, x, z)) dx, \\ \omega_3(t, z) &:= \int_0^t \int_{\mathbb{T}} \rho(s, x, z) dx ds. \end{aligned} \quad (78)$$




FIGURE 9. Density profiles for two different policies choices of α with the scheme in (58) using $N^x = N^z = 800$.

The quantity ω_1 measures the cumulative outflow at stage of completion z up to time t ; the quantity ω_2 measures the current outflow at stage of completion z ; and the quantity ω_3 measures the cumulative part density at stage of completion z . The first two quantities are indicators of the processed data whereas the last indicator shows the load of the processors, which may be related to the consumption of energy during computation. In each case, larger values of ω_i are preferable.

For our numerical tests, we set $T_{\text{fin}} = 1$ and $\bar{\alpha} = 0.5$. The initial front $\zeta_0(x)$ is given as in (70) with $z_0 = 0.1$, in which case $C_\alpha \approx 0.5/0.575$. The quantities of interest are evaluated for $z = 0.5$ and $z = 0.75$. Figure 9 shows the initial density profile and the profiles at $T_{\text{fin}} = 1$ for each of the two policies. For the policy based on equation (77), the control has been active long enough to push the processors that were initially trailing ahead of the others. This means the control has been on for too long and should have been modified or turned off at an earlier time. While the results of the two policies appear similar, the quantities of interest $w_i(t)$, $i = 1, 2, 3$ displayed in Figure 10 demonstrate the differences. In particular, with respect to w_1 , the policy in (77) outperforms the constant α policy.

5. Summary. In this paper, we have analyzed a recently derived mathematical model for the evolution of processed data in large-scale, asynchronous computers [5]. After the introduction of an auxiliary variable, the model is expressed as a system of partial differential equations. It is possible to prove existence of discontinuous solutions to this system for a particular class of initial and boundary conditions. These solution describes the flow of information as a series of propagating fronts. A numerical scheme for the reformulated system has also been designed based on a relaxation approximation. Numerical simulations based on this scheme demonstrate qualitative agreement with theoretical findings.

We have also briefly explored the effects of local processor speed on quantities of interest predicted by the model. In future work, we intend to investigate more extensively control mechanisms for optimizing important objectives related to performance of the large-scale computers. Further it is planned to validate the model based on experimental data.

Acknowledgments. MH and GV thank the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) for the financial support through 20021702/GRK2326, 333849990/IRTG-2379, HE5386/14,15,18-1,19-1,22-1 and under Germany's Excellence Strategy EXC-2023 Internet of Production 390621612. The funding through HIDSS-004 is acknowledged. This work has been supported also by the

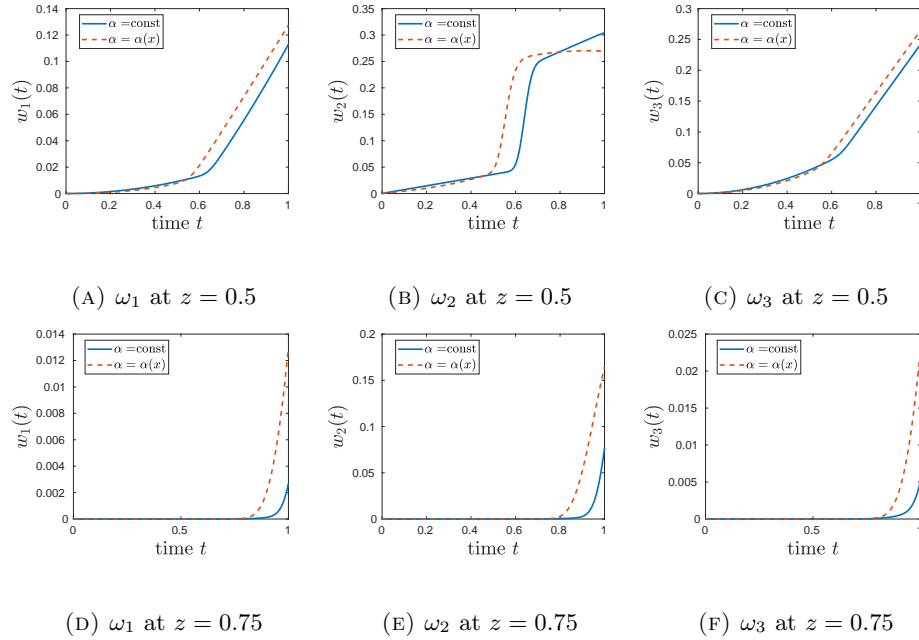


FIGURE 10. Comparison of the quantities of interest $w_i, i = 1, 2, 3$ given in (78).

US National Science Foundation, RNMS (KI-Net) grant 11-07444, and the U.S. Department of Energy, Office of Advanced Scientific Computing Research. The work of CDH was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. De-AC05-00OR22725.

REFERENCES

- [1] C. Appert-Rolland, P. Degond and S. Motsch, [Two-way multi-lane traffic model for pedestrians in corridors](#), *Netw. Heterog. Media*, **6** (2011), 351–381.
- [2] D. Armbruster, P. Degond and C. Ringhofer, [A model for the dynamics of large queuing networks and supply chains](#), *SIAM J. Appl. Math.*, **66** (2006), 896–920.
- [3] D. Armbruster, S. Göttlich and M. Herty, [A scalar conservation law with discontinuous flux for supply chains with finite buffers](#), *SIAM J. Appl. Math.*, **71** (2011), 1070–1087.
- [4] A. Aw, A. Klar, T. Materne and M. Rascle, [Derivation of continuum traffic flow models from microscopic follow-the-leader models](#), *SIAM J. Appl. Math.*, **63** (2002), 259–278.
- [5] R. C. Barnard, K. Huang and C. Hauck, [A mathematical model of asynchronous data flow in parallel computers](#), *IMA Journal of Applied Mathematics*, **85** (2020), 865–891.
- [6] N. Bellomo and C. Dogbé, [On the modelling crowd dynamics from scaling to hyperbolic macroscopic models](#), *Mathematical Models and Methods in Applied Sciences*, **18** (2008), 1317–1345.
- [7] H. Chan, M. J. Cherukara, B. Narayanan, T. D. Loeffler, C. Benmore, S. K. Gray and S. K. R. S. Sankaranarayanan, [Machine learning coarse grained models for water](#), *Nature Communications*, **10** (2019), 379.
- [8] A. Chertock, A. Kurganov, A. Polizzi and I. Timofeyev, [Pedestrian flow models with slowdown interactions](#), *Mathematical Models and Methods in Applied Sciences*, **24** (2014), 249–275.
- [9] C. M. Dafermos, [Polygonal approximations of solutions of the initial value problem for a conservation law](#), *J. Math. Anal. Appl.*, **38** (1972), 33–41.

- [10] C. M. Dafermos, *Hyperbolic Conservation Laws in Continuum Physics*, vol. 325 of Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences], 2nd edition, Springer-Verlag, Berlin, 2005.
- [11] M. Di Francesco and M. D. Rosini, [Rigorous derivation of nonlinear scalar conservation laws from follow-the-leader type models via many particle limit](#), *Archive for Rational Mechanics and Analysis*, **217** (2015), 831–871.
- [12] J. Dongarra, J. Hittinger, J. Bell, L. Chacón, R. Falgout, M. Heroux, P. Hovland, E. Ng, C. Webster and S. Wild, *Applied Mathematics Research for Exascale Computing*, Technical report, U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research Program, 2014.
- [13] L. C. Evans, *Partial Differential Equations*, American Mathematical Society, 2010.
- [14] A. Galante and D. Levy, [Modeling selective local interactions with memory](#), *Physica D: Nonlinear Phenomena*, **260** (2013), 176–190.
- [15] D. Helbing, A fluid dynamic model for the movement of pedestrians, *Complex Systems*, **6** (1992), 391–415.
- [16] H. Holden and N. H. Risebro, [Models for dense multilane vehicular traffic](#), *SIAM J. Math. Anal.*, **51** (2019), 3694–3713.
- [17] C. Murray et al., *Basic Research Needs for Microelectronics: Report of the Office of Science Workshop on Basic Research Needs for Microelectronics*, Technical report, USDOE Office of Science (SC)(United States), 2018.
- [18] J. S. Vetter et al., [Extreme Heterogeneity 2018-Productive Computational Science in the Era of Extreme Heterogeneity: Report for DOE ASCR Workshop on Extreme Heterogeneity](#), Technical report, USDOE Office of Science (SC), Washington, DC (United States), 2018.
- [19] S. Williams, A. Waterman and D. Patterson, [Roofline: An insightful visual performance model for multicore architectures](#), *Communications of the ACM*, **52** (2009), 65–76.

Received November 2020; revised May 2021.

E-mail address: hauckc@ornl.gov

E-mail address: herty@igpm.rwth-aachen.de

E-mail address: visconti@igpm.rwth-aachen.de