2013 American Control Conference (ACC)
Washington, DC, USA, June 17-19, 2013

Invalidation of Dynamic Network Models

Diego Ruiz[1] and Jorge Finke[2]

Abstract— **Models of discrete event systems combine ideas from control theory and computer science to represent the evolution of distributed processes. We formalize a notion of the invalidation of models presumed to describe dynamics on networks, and introduce an algorithm to evaluate a class of event–driven processes that evolve close to an invariant and stable state. The algorithm returns the value true, if according to the proposed notion of invalidation, the evolution of empirical observations is inconsistent with the stability properties of the model. To illustrate the approach, we represent a generic decision–making process in which the marginal utility of allocating agents to particular nodes rests on the well–known concept in economy theory of the law of diminishing returns.**

I. INTRODUCTION

Building network–based applications can help us understand complex phenomena arising from distributed processes [1], [2]. It is envisioned that applications that leverage the cues and services of massive decision–makers, often referred to as agents, replace more vulnerable and costly centralized systems. Discrete event systems provide an abstract representation of the pattern (sequence) of events underlying the tangled economies on which these applications generally operate. The ability to predict the evolution of an event–driven process depends on the extent to which a particular representation (model) describes the dynamic behavior of its empirical counterpart [3].

Checking for consistency with the data forms an important step in addressing the invalidation problem for dynamic models [4]–[6]. The work in [4] introduces an invalidation technique for continuous–time models that rests on functions reminiscent of Lyapunov functions called barrier certificates. Certificates evaluate whether a model, and its feasible set of parameters, are consistent in the sense that the evolution of the data does not contradict the possible state trajectories generated by the set of differential equations (i.e., certificates divide the state space into regions that trajectories starting from a given set of initial states cannot reach). While finding a proper certificate is in general an NP–hard problem, the technique in [4] addresses the case where the vector fields are polynomials, and the initial, final, and parameter sets can be described by polynomial equalities and inequalities (semi–algebraic sets). For this particular case, the sum of square (SOS) decomposition provides a tractable computational relaxation [7]. Using SOS decomposition, the work in [8] introduces semi–definite programs to approximate SOS decomposition and evaluate dynamics represented by difference equations [9], [10]. For discrete event systems, however, because state trajectories evolve at asynchronous instants of time, the construction of certificates cannot derive from SOS decomposition.

This paper proposes a notion of invalidation for these types of systems. In particular, we present sufficient conditions that would make a model inconsistent with an available set of data. As in [4], constructing a function that serves as a certificate rests on the stability properties of the model. Based on the identified conditions, we propose an algorithm to evaluate models presumed to describe dynamics on networks.

To illustrate the approach we formulate a generic allocation problem as a distributed optimization process where agents approach resource sites (i.e., choose destination nodes) that maximize their marginal gains according to the law of diminishing returns. The outcome of the matching process of agents to resources resembles a solution concept called the ideal free distribution (IFD) where the underlying decision–making mechanism not only depends on the quality of resources but also on the number of agents allocating their efforts to the same resource [11]. We compare the dynamics underlying of the IFD with data on the balancing of units of load in parallel computing [12]. While the distribution of load reaches the same equilibrium as the generic allocation model, our algorithm shows that the stability properties of the model are inconsistent with the evolution of the data. In other words, the proposed notion of invalidation rejects the agents' strategy of diminishing returns as a plausible mechanism behind the balancing of load.

The remainder of this work is organized as follows: Section II introduces a generic allocation problem used to illustrate a dynamic network process. We model the allocation process in Section III and present analytical results that guarantee that the IFD is stable (Proposition 1). Section IV formalizes the notion of invalidation for discrete event systems (Definition 1) and identifies conditions that evaluate a model using both the level sets (Theorem 1) and gradient vectors of the Lyapunov function (Corollary 1). The application of Theorem 1 and Corollary 1 to the allocation model in Section III is presented in Section V. Section VI draws some conclusions.

II. AN ALLOCATION PROBLEM

The generic problem of the optimal allocation of agents across a network can be found across various disciplines, including sociology, economics, engineering, and ecology. When competing for a set of resources, the concept of the IFD captures an optimal allocation that depends not only on

[1]D. Ruiz is with the Department of Mathematics, Universidad del Cauca, Popayán, Colombia. Email: dfruiz@unicauca.edu.co

[2]J. Finke is with the Department of Electrical Engineering and Computer Science, Pontificia Universidad Javeriana, Cali, Colombia. Email: finke@ieee.org

the number of resources available at each node but also on the number of agents sharing the same location. The word "ideal" refers to the assumption that agents have perfect sensing capabilities for simultaneously determining "utility" in each of a finite number of nodes (assumed to be a correlate of the quality of the available resource). The word "free" indicates that agents can move at no cost and instantaneously from any node directly to any other node at any time. Under the above conditions, the allocation process takes place on a fully connected network. The IFD is an invariant network state where no agent can increase its marginal gain by unilateral deviations from one strategy to another; hence at the IFD all agents achieve equal marginal gains (i.e., the IFD is in fact a Pareto optimum [13], [14]).

To capture agent dynamics under less restrictive conditions, consider an undirected network $\mathcal{G} = (\mathcal{N}, \mathcal{A})$, where $\mathcal{N} = \{1, \ldots, n+1\}$ represents the set of $n+1$ nodes and $\mathcal{A}$ represents the set of edges as 2–element subsets of $\mathcal{N}$. If $\{i, j\} \in \mathcal{A}$, we say that nodes $i, j \in \mathcal{N}$ are adjacent. The neighborhood $\mathcal{N}_i$ of node i is defined as the set of all its adjacent nodes, i.e., $\mathcal{N}_i := \{j \in \mathcal{N} : \{i, j\} \in \mathcal{A}\}$. The network carries a fixed number of agents $q > 0$ (e.g., load) to be distributed across a set of nodes (e.g., a set of interconnected computers). Agents can be transferred from one node to another at no cost but only according to $\mathcal{A}$.

Note that modeling the agent dynamics on a network $\mathcal{G}$ relaxes to some extent the free assumption. To represent dynamic behavior based on the law of diminishing returns, let x_i denote the number of agents, u_i represent a utility function associated to node i, and consider the following assumptions on the network.

(A1) The network $\mathcal{G}$ is connected, without self–loops or parallel edges.

(A2) Any node $i \in \mathcal{N}$ has a strictly increasing utility function $u_i(x_i)$ that is continuously differentiable and strictly concave on $\mathcal{R} = [0, q]$. The derivative of u_i with respect to its argument $s_i(x_i) := \frac{du_i}{dx_i}$ (i.e., its marginal utility function) is Lipschitz continuous on $\mathcal{R}$.

Let $x, y \in \mathcal{R}$, $x \neq y$. Note that because s_i satisfies the Lipschitz condition, there exists a positive constant K_i such that under the usual metric on $\mathcal{R}$ we have

$$\frac{|s_i(x) - s_i(y)|}{|x - y|} \leq K_i \tag{1}$$

Furthermore, since s_i is strictly decreasing in $\mathcal{R}$ note that

$$\frac{s_i(x) - s_i(y)}{x - y} < 0 \tag{2}$$

The allocation process rests on marginal utility functions of the form

$$s_i(x_i) := \frac{w_i}{x_i + 1} \tag{3}$$

where $w_i > 0$ characterizes the quality of the resources available at node i (e.g., load processing rate) [15]. To guarantee that all nodes have a positive number of agents at the IFD, we assume that

$$q + n + 1 > \max_{i \in \mathcal{N}} \left\{ \frac{1}{w_i} \right\} \sum_{i \in \mathcal{N}} w_i \tag{4}$$

According to (3), ignoring the constant of integration, the utility function associated to node i is given by

$$u_i(x_i) := \int s_i(x_i) dx_i = w_i \ln(x_i + 1)$$

The total utility function u is defined as

$$u: \begin{array}{ccc} \mathcal{R}^{n+1} & \rightarrow & \mathbb{R}_0^+ \\ (x_1, \ldots, x_{n+1}) & \mapsto & \sum\limits_{i \in \mathcal{N}} u_i(x_i) \end{array}$$

where $\mathbb{R}_0^+$ denotes the set of positive real numbers including zero. Next, we formalize the dynamics of allocating agents under the premise that the first agent to be allocated to node i yields a greater increase in $u_i(x_i)$ than the second and subsequent agents.

III. A DISCRETE EVENT MODEL

Let Δ_q be the simplex of all $n+1$ tuples on $\mathcal{R}^{n+1}$ constrained by the total number of available agents

$$\Delta_q := \left\{ (x_1, \ldots, x_{n+1}) \in \mathcal{R}^{n+1} : \sum_{i \in \mathcal{N}} x_i = q \right\}$$

Because it requires n states to represent the dynamics on $\mathcal{G}$, we define the set $\mathcal{X}$ as the projection of the simplex Δ_q on the $n-$dimensional space $x_1 \ldots x_n$, so that

$$\mathcal{X} := \{(x_1, \ldots, x_n) \in \mathcal{R}^n : (x_1, \ldots, x_n, x_{n+1}) \in \Delta_q\} \tag{5}$$

Denote $\mathbf{x} = (x_1, \ldots, x_n)$ and let $\mathcal{X}^*$ be the set defined by

$$\mathcal{X}^* = \{\mathbf{x} \in \mathcal{X} : s_i(x_i) = s_j(x_j), \forall i, j \in \mathcal{N}\} \tag{6}$$

Because of Assumption (A2), for a fixed q that satisfies (4), there exists a unique state $\mathbf{x}^* = (x_1^*, \ldots, x_n^*) \in \mathcal{X}^*$ such that $u(x_1^*, \ldots, x_{n+1}^*)$ is maximized if and only if $s_i(x_i^*) = s_j(x_j^*)$ for all $i, j \in \mathcal{N}$ [16]. If every agent has the same impact on any node (i.e., $w_i = w_j$ for all $i, j \in \mathcal{N}$), the IFD is achieved when there are the same number of agents at each node and

$$\mathcal{X}^* := \{\mathbf{x} \in \mathcal{X} : x_i = x_j, \forall i, j \in \mathcal{N}\} \tag{7}$$

When $\mathbf{x}^* \in \mathcal{X}^*$, we denote henceforth the number of agents at any node as x^* (instead of x_i^*).

The dynamics are governed by a discrete event system represented by $\mathcal{S} = (\mathcal{X}, \mathcal{G}, \mathcal{E}, g, f_e)$, where $\mathcal{E}$ denotes the set of possible events that drive the processes on $\mathcal{G}$. We denote an event at time index $k \in \mathbb{N}$ as $e(k)$, where $\mathbb{N}$ denotes the set of all nonnegative integer numbers. For $\mathbf{x}(k) \in \mathcal{X}$, we say that an event $e(k)$ is active, if it belongs to an activation function g (i.e., $e(k) \in g(\mathbf{x}(k))$. Note that there is just one state at any k, but there could be many active events. If at state $\mathbf{x}(k)$ an active event $e(k)$ occurs, the transition function f_e generates the state $\mathbf{x}(k+1)$ defined by $\mathbf{x}(k+1) := f_{e(k)}(\mathbf{x}(k))$. Furthermore, if there exists a deadlock at k, the only active event is the null event e^0, where $f_{e^0}(\mathbf{x}(k)) = \mathbf{x}(k)$. For a fixed k,

$E_k = \{e(0), e(1), \ldots, e(k-1)\}$ defines the finite sequence of k events and $E = \{e(k), e(k+1), \ldots\}$ the infinite sequence of events after and including k. The concatenation of the sequences E_k and E defines the infinite sequence of events, that is, $E_k E = \{e(0), \ldots, e(k-1), e(k), \ldots\}$. By definition, E_0 represents the empty sequence of events. Let $\mathbf{E}$ denote the set of all event trajectories in $\mathcal{S}$.

The set $\mathcal{X}_0$ denotes the initial states of the system. If $\mathbf{x}(0) \in \mathcal{X}_0$, $\mathbf{E}(\mathbf{x}(0))$ is the set of all event trajectories starting from $\mathbf{x}(0)$. Starting from $\mathbf{x}(0)$ and applying the sequence E_k of k events such that $E_k E \in \mathbf{E}(\mathbf{x}(0))$ the system reaches the state $\mathbf{x}(k)$ defined by the function

$$\begin{array}{rccc} \mathcal{F}: & \mathcal{X}_0 \times \mathbf{E} \times \mathbb{N} & \rightarrow & \mathcal{X} \\ & (\mathbf{x}(0), E_k, k) & \mapsto & \mathbf{x}(k) \end{array}$$

Note that $\mathcal{F}(\mathbf{x}(0), E_k, k)$ defines the evolution of the system.

To define decision–making strategies that guarantee any state trajectory starting from the set of initial states $\mathcal{X}_0 \subseteq \mathcal{X}$ converges to the invariant set $\mathcal{X}^*$ (i.e., $\mathcal{F}(\mathbf{x}(0), E_k, k) \rightarrow \mathcal{X}^*$ for all $\mathbf{x}(0) \in \mathcal{X}_0$ as $k \rightarrow \infty$), we now specify $\mathcal{E}$ and $g(\mathbf{x}(k))$.

Let

$$\mathcal{M}_i := \{j \in \mathcal{N}_i : x_j > x_i\} \tag{8}$$

be the set of neighboring nodes with a number of agents above that of node i. Let e_{ji} represent the reallocation of α_{ji} agents from node $j \in \mathcal{M}_i$ to node i, and

$$\alpha_{ji} = \frac{x_j - x_i}{d_i + 1}$$

where d_i represents the degree of node i (i.e., the cardinality of the neighborhood $\mathcal{N}_i$). Note that $\alpha_{ji} \in \mathbb{R}^+$. Let $\mathcal{E}_\alpha = \{e_{ji}\}$ be the set of all possible reallocations of agents on the network $\mathcal{G}$. The set of events is given by the power set of $\mathcal{E}_\alpha$ without the empty set, i.e., $\mathcal{E} = \mathcal{P}(\mathcal{E}_\alpha) \setminus \{\emptyset\}$. This definition of $\mathcal{E}$ ensures that there always exists an event that can occur. Note that an event $e(k) \in \mathcal{E}$ is a set in which each element represents the reallocation of agents from some node $j \in \mathcal{N}$ to neighboring nodes.

An event $e(k)$ is active at time index k (denoted by $e(k) \in g(\mathbf{x}(k))$) only if for all $e_{ji} \in e(k)$, $j \in \mathcal{M}_i$. If an active event $e(k)$ occurs with $e_{ji} \in e(k)$, we consider $\alpha_{ji'} = 0$ for all $e_{ji'} \in e(k)$. The allocation strategy ensures that the number of agents at node i does not exceed the local maximum in the neighborhood $\mathcal{N}_i$. Note that the proposed scheme does not require that we know $\mathcal{X}^*$ explicitly.

If the event $e(k) \in g(\mathbf{x}(k))$ occurs, the transition between states is defined as

$$x_i(k+1) = x_i(k) - \sum_{\{j : e_{ij} \in e(k)\}} \alpha_{ij} + \sum_{\{j : e_{ji} \in e(k)\}} \alpha_{ji} \tag{9}$$

In other words, the number of agents at any node $i \in \mathcal{N}$ at $k+1$ is equal to the number of agents at k minus the number of agents that reallocate to neighboring nodes, plus the number of agents reaching node i from neighboring nodes in $\mathcal{M}_i$. Note that if $\mathbf{x}(k) \in \mathcal{X}^*$ then the state $\mathbf{x}(k+1) \in \mathcal{X}^*$ (i.e., $\mathcal{X}^*$ is invariant because $\mathcal{M}_i = \emptyset$ for all $i \in \mathcal{N}$).

The following proposition characterizes the stability properties of $\mathcal{X}^*$.

Proposition 1: The invariant set $\mathcal{X}^*$ defined in (7) has a region of asymptotic stability equal to $\mathcal{X}$.

The proof of Proposition 1 is presented in the supplement to this article available at http://www.jfinke.org/public_html/publications.html. There we use

$$\mathcal{V}(\mathbf{x}) := \max\{x_1, \ldots, x_n, q - \textstyle\sum_{i=1}^{n} x_i\} - x^* \tag{10}$$

as a Lyapunov function to show that the IFD (i.e., the set in (7)) is an invariant and stable network state of the model. Next, Section IV presents a notion of invalidation and introduces an algorithm to evaluate whether the evolution of empirical data is consistent with the stability properties. In Section V we come back to (10) to evaluate the dynamics underlying the IFD.

IV. INVALIDATION

To formalize the notion of invalidation of a model, we need to introduce the following notation. Let ρ be the metric on $\mathcal{X}$ and $\mathcal{V}$ a continuous Lyapunov function for an invariant state $\mathcal{X}^*$. Define a level set of $\mathcal{V}$ corresponding to a real value c as $\mathcal{L}_c := \{(x_1, \ldots, x_n) \in \mathcal{X} : \mathcal{V}(x_1, \ldots, x_n) = c\}$. Moreover, let $\{\hat{\mathbf{x}}\} := \{\hat{\mathbf{x}}(1), \ldots, \hat{\mathbf{x}}(m)\}$ be a finite data trajectory of m points in the domain of $\mathcal{V}$. We say that the system $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$ if the data trajectory $\{\hat{\mathbf{x}}\}$ is inconsistent with the stability properties of $\mathcal{X}^*$. Formally, the invalidation of a discrete event system is given by Definition 1.

Definition 1: A discrete event system $\mathcal{S}$ with an invariant and stable set $\mathcal{X}^*$ and a set of initial states $\mathcal{X}_0$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$ if there exist two neighborhoods of $\mathcal{X}^*$, $\mathcal{U}_1 \subseteq \mathcal{U}_2$, and two time indices, $k' < k'' \leq m$, such that the following conditions are satisfied

(D1) $\mathcal{F}(\mathbf{x}(0), E_k, k) \in \mathcal{U}_2$, for all $\mathbf{x}(0) \in \mathcal{U}_1$, all $k \geq 0$ and all E_k, $E_k E \in \mathbf{E}(\mathbf{x}(0))$; and

(D2) $\hat{\mathbf{x}}(k') \in \mathcal{U}_1 \cap \mathcal{X}_0$ and $\rho(\hat{\mathbf{x}}(k''), \mathcal{U}_2) > 0$.

The left plot of Figure 1 illustrates that any motion of the system $\mathcal{S}$ starting from an initial state $\mathbf{x}(0) \in \mathcal{U}_1$ (in particular $\hat{\mathbf{x}}(k')$) must remain in the neighborhood $\mathcal{U}_2$ (according to Condition (D1)). The right plot depicts how a finite data trajectory $\{\hat{\mathbf{x}}\}$ with $\hat{\mathbf{x}}(k'')$ such that $\rho(\hat{\mathbf{x}}(k''), \mathcal{U}_2) > 0$ for some k'' contradicts the stability properties of the model (according to Condition (D2)). The following theorem presents sufficient conditions that would invalidate a discrete event system $\mathcal{S}$ based on Definition 1.

Theorem 1: Let $\mathcal{S}$ be a discrete event system with an invariant and stable set $\mathcal{X}^*$ and a set of initial states $\mathcal{X}_0$. If there exist a real function $\mathcal{V} : \mathcal{X} \rightarrow \mathbb{R}_0^+$, and a finite data trajectory $\{\hat{\mathbf{x}}\}$ with m points such that

(C1) $\mathcal{V}$ is a continuous Lyapunov function that guarantees the stability of $\mathcal{X}^*$; and

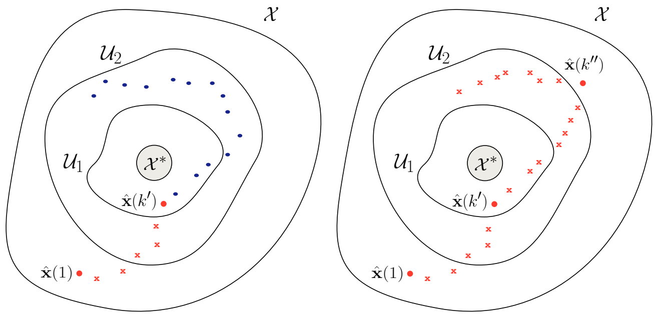

Fig. 1. Illustration of the notion of invalidation. Condition (D1) (left plot) states that any trajectory starting from any initial state $\mathbf{x}(0) \in \mathcal{U}_1$ (in particular from $\hat{\mathbf{x}}(k')$) remains in the neighborhood $\mathcal{U}_2$; Condition (D2) (right plot) requires that the data trajectory $\{\hat{\mathbf{x}}\}$ satisfies $\rho(\hat{\mathbf{x}}(k''), \mathcal{U}_2) > 0$ for some $\hat{\mathbf{x}}(k'')$.

(C2) There exist $k' < k'' \leq m$ such that $\mathcal{V}(\hat{\mathbf{x}}(k'')) > \mathcal{V}(\hat{\mathbf{x}}(k'))$, with $\hat{\mathbf{x}}(k') \in \mathcal{X}_0$

then $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$.

Proof: Suppose there exists a function $\mathcal{V}$ satisfying (C1) and (C2). Let $c_1 = \mathcal{V}(\hat{\mathbf{x}}(k'))$, $c_2 = \mathcal{V}(\hat{\mathbf{x}}(k''))$, and $a \in \mathbb{R}$ be a constant such that $c_1 < a < c_2$. Define the closed sets $\Omega_{c_1} := \{\mathbf{x} \in \mathcal{X} : \mathcal{V}(\mathbf{x}) \leq c_1\}$, $\Omega_a := \{\mathbf{x} \in \mathcal{X} : \mathcal{V}(\mathbf{x}) \leq a\}$, $\Omega_{c_2} := \{\mathbf{x} \in \mathcal{X} : \mathcal{V}(\mathbf{x}) \leq c_2\}$. Note that $\Omega_{c_1} \subset \Omega_a \subset \Omega_{c_2}$. Define $\mathcal{L}_a := \{\mathbf{x} \in \mathcal{X} : \mathcal{V}(\mathbf{x}) = a\}$ and $\mathcal{L}_{c_2} := \{\mathbf{x} \in \mathcal{X} : \mathcal{V}(\mathbf{x}) = c_2\}$ as the level sets of $\mathcal{V}$ corresponding to the values a and c_2, respectively. Because $\mathcal{L}_a$ and $\mathcal{L}_{c_2}$ are closed sets, and $\Omega_{c_1} \cap \mathcal{L}_a = \emptyset$, $\Omega_a \cap \mathcal{L}_{c_2} = \emptyset$, there exist two neighborhoods of $\mathcal{X}^*$, $\mathcal{U}_1$ and $\mathcal{U}_2$, such that $\Omega_{c_1} \subset \mathcal{U}_1 \subset \Omega_a$ and $\Omega_a \subset \mathcal{U}_2 \subset \Omega_{c_2}$. Let $\mathbf{x}(0)$ be any initial state in $\mathcal{U}_1$. According to (C1) we know that $\mathcal{V}$ is non–increasing along all possible motions of $\mathcal{S}$ starting at $\mathbf{x}(0)$. In other words

$$\mathcal{V}(\mathbf{x}(0)) \geq \mathcal{V}(\mathcal{F}(\mathbf{x}(0), E_k, k)) \tag{11}$$

for all $k > 0$ and all E_k such that $E_k E \in \mathbf{E}(\mathbf{x}(0))$. By (11) we have $\mathcal{V}(\mathcal{F}(\mathbf{x}(0), E_k, k)) < a$ which implies that $\mathcal{F}(\mathbf{x}(0), E_k, k) \in \mathcal{U}_2$. Condition (D1) is satisfied with the established neighborhoods $\mathcal{U}_1$ and $\mathcal{U}_2$.

By construction and (C2), $\hat{\mathbf{x}}(k') \in \mathcal{U}_1$ and $\rho(\hat{\mathbf{x}}(k''), \mathcal{U}_2) > 0$ which implies that Condition (D2) is satisfied. Therefore, there exist two neighborhoods of $\mathcal{X}^*$, $\mathcal{U}_1$ and $\mathcal{U}_2$, and two time indices, $k' < k'' \leq m$ such that Conditions (D1) and (D2) are satisfied. The discrete event system $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$. ■

Remark: Theorem 1 allows us, given an available set of data, to narrow down the search for valid models. Note that checking for consistency may rest on different stability properties. Whether stronger types of stability are appropriate depends on the application at hand.

Next, let $\nabla\mathcal{V}(\mathbf{x})$ denote the gradient vector of $\mathcal{V}$ at $\mathbf{x} \in \mathcal{X}$. Algorithm 1 presents a procedure to invalidate a model using gradient vectors associated to the Lyapunov function $\mathcal{V}$ under the following assumption.

(A3) The level sets of $\mathcal{V}$ partition the state space into a finite number of regions such that each region R_i has a constant gradient vector and $\bigcap R_i = \mathcal{X}^*$.

Algorithm 1 Invalidation procedure

Requires: A Lyapunov function $\mathcal{V}$ and a finite trajectory $\{\hat{\mathbf{x}}\}$ of m data points
Returns: true if $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$ or false otherwise

```
Define k ← 1 and invalidation ← 0
Suppose that x̂(k) ∈ R_i, x̂(k+1) ∈ R_j
while invalidation ≠ 1 and k ≠ m do
   Find the gradient ∇V(x̂(k))
   Define the vector v ← x̂(k+1) − x̂(k)
   Find the angle θ between v and ∇V(x̂(k))
   if i = j and 0 ≤ θ < π/2 then
      invalidation ← 1
   else
      Find the intersection point (P) of v with the level
      set L_V(x̂(k))
      Let d1 ← ρ(x̂(k), P) and d2 ← ‖v‖
      if d1 < d2 then
         invalidation ← 1
      end if
   end if
   k ← k + 1
end while
if invalidation = 1 then
   return true
else
   return false
end if
```

Corollary 1: Let $\mathcal{S}$ be a discrete event system with a set of initial states $\mathcal{X}_0$. Let $\mathcal{V} : \mathcal{X} \to \mathbb{R}_0^+$ be a Lyapunov function that satisfies (A3) and proves the stability of the invariant set $\mathcal{X}^*$. If Algorithm 1 returns the value true, then $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$.

Proof: Let $\{\hat{\mathbf{x}}\}$ be a finite data trajectory with m points contained in the domain of $\mathcal{V}$. Let $\hat{\mathbf{x}}(k)$, $\hat{\mathbf{x}}(k+1)$ be two consecutive data for some $k \in \{1, \ldots, m-1\}$ such that $\hat{\mathbf{x}}(k+1) \neq \hat{\mathbf{x}}(k)$. Suppose that $\hat{\mathbf{x}}(k+1) \in R_j$ and $\hat{\mathbf{x}}(k) \in R_i$. Define the vector $v := \hat{\mathbf{x}}(k+1) - \hat{\mathbf{x}}(k)$ and let $\theta \in [0, \pi]$ be the angle between the vectors v and $\nabla\mathcal{V}(\hat{\mathbf{x}}(k))$. Suppose Algorithm 1 returns the value true and consider the following two cases.

If $i = j$, then the gradient vector does not leave region R_i. For $0 \leq \theta < \pi/2$ the dot product $v \cdot \nabla\mathcal{V}(\hat{\mathbf{x}}(k))$ is greater than zero and the orthogonal projection of $\nabla\mathcal{V}(\hat{\mathbf{x}}(k))$ on the vector v has the same orientation as v. Because the gradient vector $\nabla\mathcal{V}(\hat{\mathbf{x}}(k))$ points in the direction in which $\mathcal{V}$ grows the fastest between $\hat{\mathbf{x}}(k+1)$ and $\hat{\mathbf{x}}(k)$, it causes the function $\mathcal{V}$ to increase in the direction of v. As a result $\mathcal{V}(\hat{\mathbf{x}}(k+1)) > \mathcal{V}(\hat{\mathbf{x}}(k))$, and according Theorem 1, $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$.

If $i \neq j$ or $\pi/2 \leq \theta < \pi$, define $d_1 := \rho(\hat{\mathbf{x}}(k), \mathbf{P})$ and $d_2 := \|v\|$, where $\mathbf{P}$ is the intersection point of v (or any vector that lies on v) with the level set $\mathcal{L}_{\mathcal{V}(\hat{\mathbf{x}}(k))}$. Note that if $d_1 < d_2$, then $\mathcal{V}(\hat{\mathbf{x}}(k+1)) > \mathcal{V}(\hat{\mathbf{x}}(k))$, because $\mathbf{P}$ lies on v and $\mathcal{L}_{\mathcal{V}(\hat{\mathbf{x}}(k))} = \mathcal{L}_{\mathcal{V}(\mathbf{P})}$. According Theorem 1, the discrete event system $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$. ■

V. APPLICATION

This section applies Definition 1 to the generic allocation model presented in Section III. For the function $\mathcal{V}$ in (10) the region R_i is given by

$$R_i := \{(x_1, \ldots, x_n) \in \mathcal{X} : \max\{x_1, \ldots, x_{n+1}\} = x_i\} \tag{12}$$

Because $\bigcup_{i \in \mathcal{N}} R_i = \mathcal{X}$, the level sets of $\mathcal{V}$ partition the state space into a finite number of regions. Furthermore, $(x_1, \ldots, x_n) \in \bigcap_{i \in \mathcal{N}} R_i$ if and only if $\max\{x_1, \ldots, x_{n+1}\} = x_i$ for all $i \in \mathcal{N}$ implying $x_1 = \cdots = x_{n+1}$ $(\bigcap_{i \in \mathcal{N}} R_i = \mathcal{X}^*)$. Now, the gradient vector of $\mathcal{V}$ satisfies:

1) Suppose $\max\{x_1, \ldots, x_{n+1}\} = x_i,\ i \neq n+1$. Then $\mathcal{V}(x_1, \ldots, x_n) = x_i - x^*$ and $(x_1, \ldots, x_n) \in R_i$. Thus $\nabla\mathcal{V}(x_1, \ldots, x_n) = \xi_i$, where ξ_i

$$\xi_i := \begin{cases} 1 & \text{at the } ith \text{ position} \\ 0 & \text{otherwise} \end{cases}$$

is an n–dimensional vector.

2) Suppose $\max\{x_1, \ldots, x_{n+1}\} = x_{n+1}$. Then $\mathcal{V}(x_1, \ldots, x_n) = x_{n+1} - x^*$ and $(x_1, \ldots, x_n) \in R_{n+1}$. Since $x_{n+1} = q - x_1 - \cdots - x_n$ we have $\mathcal{V}(x_1, \ldots, x_n) = q - (x_1 + \cdots + x_n) - x^*$ and $\nabla\mathcal{V}(x_1, \ldots, x_n)$ is the n–dimensional vector $-(1, \ldots, 1)$.

Note that when $n = 2$

$$\nabla\mathcal{V}(x_1, x_2) = \begin{cases} (1,0) & \text{if} \quad (x_1, x_2) \in R_1 \\ (0,1) & \text{if} \quad (x_1, x_2) \in R_2 \\ -(1,1) & \text{if} \quad (x_1, x_2) \in R_3 \end{cases}$$

Thus, because each region has a constant gradient vector, the model satisfies Assumption (A3).

To evaluate to what extent the generic allocation model may capture the dynamics of a network of computers balancing load, consider the data presented in [12], resulting from a balancing algorithm that allocates 12000 agents (load units) across three nodes. Table I shows the empirical initial distribution of load across New Mexico, Taipei, and Frankfurt. The simplex Δ_q on which the dynamics evolve is

$$\Delta_q = \{(x_1, x_2, x_3) \in \mathcal{R}^3 : x_1 + x_2 + x_3 = 12000\}$$

Because, given x_1 and x_2, the number of agents at node 3

	Node 1	Node 2	Node 3
Location	New Mexico	Taipei	Frankfurt
Initial load distribution	6000	4000	2000

TABLE I

INITIAL DISTRIBUTION OF LOAD (AGENTS) ACROSS NODES (COMPUTERS). DATA FROM [12].

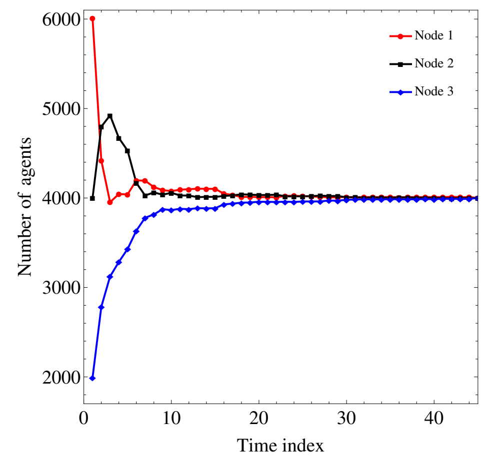

Fig. 2. Evolution of 12000 agents across 3 nodes with $\hat{\mathbf{x}}(1) = (6000, 4000)$.

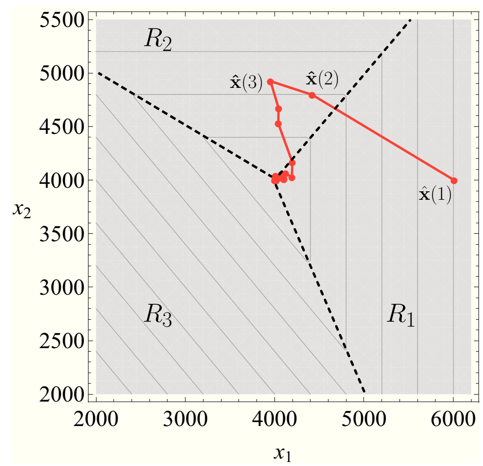

Fig. 3. Invalidation of $\mathcal{S}$. The triple $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$ invalidates the model at $k = 3$.

is known, the set of states is

$$\mathcal{X} = \{(x_1, x_2) \in \mathcal{R}^2 : (x_1, x_2, x_3) \in \Delta_q\}$$

The invariant set $\mathcal{X}^* = \{\mathbf{x} \in \mathcal{X} : x_1 = x_2 = 4000\}$ represents the IFD. Because $\mathcal{X}^*$ is asymptotically stable in the large, $\mathcal{X}_0 = \mathcal{X} = \mathcal{R}^2$.

Figure 2 shows the data trajectory $\{\hat{\mathbf{x}}\} = \{\hat{\mathbf{x}}(1), \ldots, \hat{\mathbf{x}}(45)\}$ of the evolution of load at each node across the network. There are several peaks when evaluating the Lyapunov function along $\{\hat{\mathbf{x}}\}$. In particular, the data points corresponding to $\hat{\mathbf{x}}(2) = (4419, 4797)$ and $\hat{\mathbf{x}}(3) = (3953, 4924)$, yield $\mathcal{V}(\hat{\mathbf{x}}(2)) = 797 < \mathcal{V}(\hat{\mathbf{x}}(3)) = 924$ and thus $\mathcal{V}(\hat{\mathbf{x}}(3)) - \mathcal{V}(\hat{\mathbf{x}}(2)) > 0$. The discrete event system $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$ based on Theorem 1. Figure 3 shows that while the trajectory $\{\hat{\mathbf{x}}\}$ converges to the IFD ($\mathcal{V}$ converges to zero), the trajectory jump from the level set $\mathcal{L}_{797}$ (at $k = 2$) to the level set $\mathcal{L}_{924}$ (at $k = 3$), invalidates the model.

Next, consider invalidating the model by applying Algorithm 1.

Example: At $k = 1$ we have $\hat{\mathbf{x}}(1) = (6000, 4000)$, and the value of variable *invalidation* in Algorithm 1 is zero. At $k = 2$, $\hat{\mathbf{x}}(2) = (4419, 4797)$. The procedure completes the following steps.

1) Find the gradient $\nabla\mathcal{V}(\hat{\mathbf{x}}(2))$. Because $\hat{\mathbf{x}}(2) \in R_2$, then $\nabla\mathcal{V}(\hat{\mathbf{x}}(2)) = (0, 1)$.
2) Define the vector $v := \hat{\mathbf{x}}(3) - \hat{\mathbf{x}}(2)$. Since $\hat{\mathbf{x}}(3) = (3953, 4924)$, then $v = (-466, 127)$.
3) Find the angle θ between the vectors v and $\nabla\mathcal{V}(\hat{\mathbf{x}}(2))$. Because $\cos\theta = \frac{\nabla\mathcal{V}(\hat{\mathbf{x}}(2))\cdot v}{\|\nabla\mathcal{V}(\hat{\mathbf{x}}(2))\|\|v\|}$, $\theta = 0,41\pi$. Thus, $0 \leq \theta < \dfrac{\pi}{2}$, Algorithm 1 returns the value true, and $\mathcal{S}$ is invalidated by $\{\{\hat{\mathbf{x}}\}, \mathcal{X}^*, \mathcal{X}_0\}$.

VI. CONCLUSION

Our work formalizes a notion of invalidation of dynamic network processes modelled using a class of discrete event systems with an invariant and stable network state (Definition 1). A discrete event system is invalidated if the evolution of empirical observations is inconsistent with the stability properties of the model. Based on Definition 1 we identify sufficient conditions on the Lyapunov function that invalidate the model (Theorem 1). Using Theorem 1, we introduce an algorithm that applies the gradient vector of the Lyapunov function to evaluate whether the stability properties are inconsistent with the data (Corollary 1).

To illustrate the proposed approach we formulate a generic allocation problem of agents across nodes in which the marginal utility of allocating agents rests on the law of diminishing returns. The optimal distribution is achieved when the same number of agents are distributed across each node. We prove that this network state is stable.

To evaluate the validity of the allocation model to represent the dynamics of load balancing algorithms we use data on the balancing of load across an empirical computer network. While the allocation model reaches the same equilibrium as the balancing of load, applying Theorem 1 or Corollary 1 will show that the stability properties of the model are inconsistent with the data.

VII. ACKNOWLEDGMENTS

The first author acknowledges the support of the Department of Mathematics at Universidad del Cauca.

REFERENCES

[1] A. Leon-Garcia and I. Widjaja, *Communication Networks*. McGraw-Hill, 2003.

[2] A. Grabowski, "Interpersonal interactions and human dynamics in a large social network," *Physica A: Statistical Mechanics and its Applications*, vol. 385, no. 1, pp. 363 – 369, 2007.

[3] W. Rand, D. G. Brown, S. E. Page, R. Riolo, L. E. Fernandez, and M. Zellner, "Statistical validation of spatial patterns in agent-based models," in *Proceedings of Agent-Based Simulation*, Montpellier, France, 2003.

[4] S. Prajna, "Barrier certificates for nonlinear model validation," *Automatica*, vol. 42, no. 1, pp. 117 – 126, 2006.

[5] S. Prajna and A. Jadbabaie, "Safety verification of hybrid systems using barrier certificates," in *Proceedings of Hybrid Systems: Computation and Control*, pp. 477 – 492, Philadelphia, PA, 2004.

[6] S. Prajna and A. Rantzer, "On the necessity of barrier certificates," in *Proceedings of the 16th IFAC World Congress*, Prague, Czech Republic, 2005.

[7] S. Prajna, A. Papachristodoulou, P. Seiler, and P. A. Parrilo, "SOSTOOLS and its control applications," in *Positive Polynomials in Control, Lecture Notes in Control and Information Sciences*, vol. 312, pp. 273–292, Springer, 2005.

[8] J. Anderson and A. Papachristodoulou, "On validation and invalidation of biological models," *BMC Bioinformatics*, vol. 10, no. 1, pp. 132 – 144, 2009.

[9] S. Prajna, A. Papachristodoulou, P. Seiler, and P. A. Parrilo, *SOSTOOLS: Sum of squares optimization toolbox for MATLAB*, 2004.

[10] J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones," *Optimization Methods and Software*, vol. 11, no. 4–5, pp. 625–653, 1999.

[11] S. D. Fretwell and H. L. Lucas, "On territorial behavior and other factors influencing habitat distribution in birds," *Acta Biotheoretica*, vol. 19, no. 1, pp. 16–36, 1969.

[12] J. Chiasson, Z. T. Z. Tang, J. Ghanem, C. T. Abdallah, J. D. Birdwell, M. M. Hayat, and H. Jerez, "The effect of time delays on the stability of load balancing algorithms for parallel computations," *IEEE Transactions on Control Systems Technology*, vol. 13, no. 6, pp. 932–942, 2005.

[13] R. S. Cantrell, C. Cosner, D. L. Deangelis, and V. Padron, "The ideal free distribution as an evolutionarily stable strategy," *Journal of Biological Dynamics*, vol. 1, no. 3, pp. 249–271, 2007.

[14] R. Cressman and V. Křivan, "The ideal free distribution as an evolutionarily stable state in density-dependent population games," *Oikos*, vol. 119, no. 8, pp. 1231–1242, 2010.

[15] M. Sokolowski, F. Tonneau, and E. Freixa Baqué, "The ideal free distribution in humans: An experimental test," *Psychonomic Bulletin Review*, vol. 6, no. 1, pp. 157–161, 1999.

[16] B. J. Moore, J. Finke, and K. M. Passino, "Optimal allocation of heterogeneous resources in cooperative control scenarios," *Automatica*, vol. 45, no. 3, pp. 711 – 715, 2009.