

Threshold Attendance under Soft-Crash Model: TAG Protocol and Markovian Analysis

Jerzy Konorski

Faculty of Electronics, Telecommunications and Informatics
Gdansk University of Technology
Gdansk, Poland
jekon@eti.pg.edu.pl

Abstract—A realistic and systematic network evaluation should subsume an availability model and a failure model. We combine a "hard availability" model we call *threshold attendance*, whereby a certain minimum number of network elements must be present at any time, with a *soft-crash* failure model, whereby after experiencing a failure, a network element is still able to function correctly for a while in an emergency mode at a risk of a major crash. A *Threshold Attendance Guarantee* (TAG) protocol, earlier studied from a security viewpoint, is deployed to ensure threshold attendance while controlling the duration of the emergency mode. We study the network under TAG using an "isolated" node-type Markovian analysis, offering insights into the tradeoffs between some relevant availability and reliability characteristics, and showing a simple model-free way to account for a positive correlation between the network elements' behavior that can reflect mass disasters.

Keywords—*availability; reliability; soft-crash; threshold attendance; Markovian model; isolated node*

I. INTRODUCTION

Contemporary network systems are subject to tightening performance and availability requirements, hence research into their efficiency in the presence of failures and failure recovery has recently attracted attention. The underlying network model should contain, among others, (a) an availability model that specifies the conditions for the network service to be perceived as available; they can have a form of a global consensus, such as the correct functioning at any time of at least one of predefined subsets of network elements, and (b) a reliability model of a network element, including stochastic lifetime and intermittency properties. Existing models vary from crash (fail-stop) to omission/timing to Byzantine failures ([1]).

In this paper we offer some preliminary insights by considering a simplified "hard availability" model we call *threshold attendance*, whereby a certain minimum number of network elements must function correctly at any time. This is motivated, e.g., by similar requirements in a smart electric grid [2], where some minimum electric load is necessary to keep the grid stable, or in cloud computing environments, where at any given time a certain number of servers must be running. Accordingly, a *Threshold Attendance Guarantee* (TAG) protocol, first studied in [3] from a

security viewpoint, is deployed. Maintaining a global consensus in a failure-prone distributed system, of which threshold attendance is a special case, is generally treated in [4]. A *soft-crash* failure model is moreover introduced, whereby after experiencing a failure, a network element is still able to function correctly for a while in an emergency mode (e.g., with reduced power supply, functionality or security) at a risk of a major crash. Unless this happens, the element only goes absent when authorized to do so.

The paper contains the following contributions. First, we study the network operation under TAG via approximate "isolated" node-type Markovian analysis to state some relationships and tradeoffs between relevant availability and reliability characteristics. Such analyses have been used in the network literature, dating back to [5], [6], and [7], but to the author's knowledge not under global consensus conditions and soft-crash model. Rather than carrying out a full performability analysis in the sense of [8], with a reward function related to a quantified risk of enough network elements incurring a major crash during the emergency mode to violate threshold attendance, we show how the duration of the emergency mode, hence the risk of a major crash, is controlled by TAG configuration. Second, using a Pólya urn scheme, we show a simple model-free way to extend the analysis so that the effects of positive correlation between the network elements' behavior can be assessed, e.g., to reflect mass disasters.

II. NETWORK OPERATION

A. Network Model

The network consists of N interconnected nodes. Each node controls a device whose operation is equally vital to the perception of the network service availability. The device in particular can be the hardware/software architecture of the node itself, or an external object, e.g., a piece of machinery as in sensor and actuator networks [9] or an appliance presenting some electric load in electric grid systems. The device is unreliable and can be either in the UP or DOWN state. For the network "hard availability" we state a consensus-type threshold attendance requirement that at least T of the N devices be in the UP state at all times, where $0 < T \leq N$, as ensured by the deployed TAG protocol. In the adopted soft-crash device reliability model, transitions to the DOWN state

Work supported by the Statutory Fund of the Faculty of Electronics, Telecommunications and Informatics, Gdansk University of Technology.

occur only when allowed by the controlling node and never spontaneously (provided that a major crash does not occur). The term *node* refers to a TAG agent functionality. A designated node moreover acts as the *Availability Supervisor* responsible for authorization of transitions to the DOWN state.

The soft-crash model illustrated in Fig. 1 permits a device in the UP state report a failure at any time. In the ensuing emergency mode the device continues to function in the UP state for a while, albeit at a risk of a major crash. This risk is not quantified here except that can be plausibly assumed to grow over time; rather, we seek to evaluate the duration of the emergency mode with a view of keeping it small enough as to almost never incurring a major crash. Thus only when authorization from the Availability Supervisor is received by controlling node does the device enter the DOWN state, whereupon repair/maintenance processes can start. When they are complete, according to the device's maintainability characteristics, the device returns to the UP state.




Fig. 1. The soft-crash model of a device.

B. TAG Protocol

A TAG message exchange is triggered by a node whose controlled device reports a failure and ends with the node allowing the device to enter the DOWN state. Hence, there may be multiple simultaneous message exchanges in progress. In a simplified view (Fig. 2; more details follow in Sec. III), the node can be in one of three states: ENABLED, REQUEST SENT, or DISABLED; in the former two, the controlled device remains in the UP state, and in the third one is allowed to enter the DOWN state. In response to the controlled device reporting a failure, the node attempts to obtain authorization to allow the device to enter the DOWN state, where repair/maintenance processes prevent a major crash. The node sends to the Availability Supervisor a *Disable Request* message and enters the REQUEST SENT state awaiting a *Disable Authorization* message. In the meantime, the Availability Supervisor sends *Endorsement Request* messages to nodes that recently have sent *Endorsement Ready* messages (thereby notifying that their controlled devices have returned to the UP state) and starts collecting the *Endorsement* messages from them. Due to the logical FIFO message channels, if at least T such messages are collected, the Availability Supervisor can be certain that sending a *Disable Authorization* message to the node triggering the message exchange will not violate the threshold attendance requirement: any of the T endorsers, in order to allow its controlled device to enter the DOWN state, will have to collect another T endorsements [3]. A *Disable Authorization* message then terminates the message exchange; the controlled device is allowed to enter the DOWN state and repair/maintenance can start. When the device returns to the UP state, the node sends an

Endorsement Ready message to the Availability Supervisor and subsequently can act as an endorser.

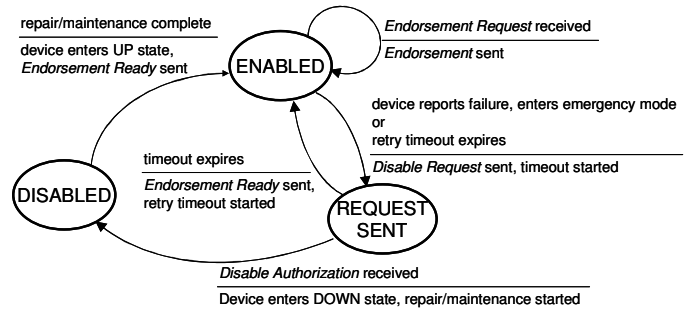


Fig. 2. Node states and state transitions under TAG.

While in the REQUEST SENT state, the node has to occasionally (after a timeout) return to the ENABLED state, where it can act as an endorser. This it does out of concern that indefinitely waiting for a *Disable Authorization* message while not sending *Endorsement* messages might indefinitely suspend other ongoing message exchanges and so preclude a correct termination of its own. Having returned to the ENABLED state, the node may send a *Disable Request* message and enter the REQUEST SENT state again after a retry timeout.

III. MARKOV MODEL OF AN ISOLATED NODE

We wish to quantify the long-term proportion of time a generic ("isolated") node is in a given TAG state, assuming that the other nodes exhibit statistically similar behavior. The idea is to identify which parameters of the "isolated" node's model are governed by its internal dynamics driven by the TAG protocol and the controlled device's reliability and maintainability, and which are influenced by the other nodes' operation. The latter give rise to a fixed-point relationship that captures the interactions of the "isolated" node with the other nodes, and permits to adjust the model to account for the presence of the rest of the network.¹

To simplify the analysis we consider discrete-time dynamics, with time split into slices of a fixed length Δ , assumed small compared to the time span of the conducted analysis; for convenience we let $\Delta = 1$. Finally, we assume that any TAG-related event can occur within a time slice with a fixed probability. This also pertains to the toggling between the UP and DOWN states of the devices that the network nodes control. Hence, the relevant uninterrupted sojourn times are geometrically distributed.

Under the above assumptions, TAG execution at a generic "isolated" network node follows a Markov chain with six distinguished states and the state transitions as in Fig. 3. Depicted are the TAG states along with their labels and stationary probabilities denoted by x_i , $i = 0, \dots, 5$. Compared to Fig. 2, additional states are introduced to grasp the dynamic TAG operation in a distributed

¹ For moderate N the joint state space can be analyzed instead, though at larger N it does not seem feasible or promise much analytical insight.

environment. The transition probabilities are discussed in the following subsection. We assume all the model parameters are selected so as to ensure ergodicity of the Markov chain.

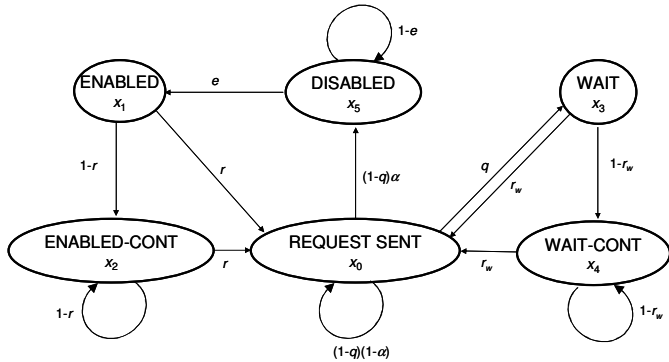


Fig. 3. States and state transitions of "isolated" node's Markov chain.

A. "Isolated" Node States and State Transitions

Compared with Fig. 2, the considered Markov model introduces two "continuation" states: ENABLED-CONT and WAIT-CONT. These states are distinguished to represent a situation where the node whose device is currently in the UP state was in the ENABLED or WAIT state in the previous time slice. It is an important distinction since, from the Availability Supervisor's perspective, only nodes in the ENABLED-CONT or WAIT-CONT state have sent *Endorsement Ready* messages and so are perceived as potential endorsers, i.e., able to send an *Endorsement* message in response to an *Endorsement Request* message.

Below we briefly explain the meaning of the states and introduce probabilities related to the state transitions.

- In the ENABLED state, the controlled device is in the UP state and can report a failure within any forthcoming time slice with a probability r , measuring the reliability of the device. A *Disable Request* message is then sent to the Availability Supervisor and the REQUEST SENT state is entered. Recall that in our soft-crash model, reporting a failure triggers an emergency mode during which the device continues to function correctly for a while, albeit at a risk of a major crash. The remaining probability $1 - r$ governs the transition to ENABLED-CONT, a state at which the node is perceived by the Availability Supervisor as a potential endorser, as it has sent an *Endorsement* message.
- In the DISABLED state, the controlled device is in the DOWN state, and repair/maintenance processes are in progress. These can bring about a return to the UP (hence, ENABLED) state within any forthcoming time slice with a probability e . This probability measures the device's maintainability—the effectiveness of the repair/maintenance processes. In such a case the node sends to the Availability Supervisor an *Endorsement Ready* message to notify it is able to act as an endorser.

- Upon entering the REQUEST SENT state, the node awaits a *Disable Authorization* message from the Availability Supervisor for a timeout duration that can end within any forthcoming time slice with a probability q . The *Disable Authorization* message arrives if the Availability Supervisor has sent *Endorsement Request* messages to nodes it perceives as potential endorsers and collected T or more *Endorsement* messages. This happens with the probability α of there being among the other $N - 1$ nodes at least T nodes currently in the ENABLED-CONT or WAIT-CONT states. The node then enters the DISABLED state. Note that $q = 0$ implies infinite mean timeout duration and gives rise to deadlocks (where the Availability Supervisor waits indefinitely before the necessary number of *Endorsement* messages are collected, since the other nodes are in the REQUEST SENT state).
- The WAIT state is entered after the timeout expires and no *Disable Authorization* message has arrived (this message is later ignored). The node sends to the Availability Supervisor an *Endorsement Ready* message. It may then retry sending a *Disable Request* message after a retry timeout expires, which happens within any forthcoming time slice with a probability r_w (thus $1/r_w$ is the mean retry timeout duration) and causes a return to the REQUEST SENT state. Note that with $r_w = 0$ a *Disable Authorization* message may never arrive. The WAIT-CONT state is distinguished for reasons similar as in the case of the ENABLED-CONT state. In the WAIT and WAIT-CONT states the controlled device is in the UP state.

Regarding the above model we make several comments:

- The reciprocals of the probabilities r , e , and r_w are the mean uninterrupted sojourn times in the ENABLED-CONT, DISABLED, and WAIT-CONT states, respectively, whereas the reciprocal of q is the mean timeout duration.
- q , and r_w are TAG protocol parameters, whereas N , T , r , and e are external parameters characterizing the network size and requirements, and controlled devices' reliability and maintainability.
- It is reasonable to assume that $r_w \geq r$, i.e., *Disable Request* message sending retries are on average more likely within a time slice than its first sending.
- α is the only parameter that couples TAG operation at our analyzed node with that at the other nodes.

B. Stationary Probabilities and Related Characteristics

Let $\mathbf{x} = (x_0, x_1, x_2, x_3, x_4, x_5)$ be the stationary probability distribution of the considered Markov chain as indicated in Fig. 3. Since ergodicity is ensured, this distribution yields the desired statistics of state occupancy times. We have the following Chapman-Kolmogorov equations, subject to the normalization constraint $x_0 + \dots + x_5 = 1$:

$$\begin{aligned}
x_0 &= (1-q)(1-\alpha)x_0 + rx_1 + rx_2 + r_w x_3 + r_w x_4, \\
x_1 &= ex_5, \quad x_2 = (1-r)x_1 + (1-r)x_2, \\
x_3 &= qx_0, \quad x_4 = (1-r_w)x_3 + (1-r_w)x_4, \\
x_5 &= (1-q)\alpha x_0 + (1-e)x_5,
\end{aligned} \tag{1}$$

Given the parameters e , r , q , r_w , and α , the system (1) admits a unique solution, which can be found analytically. If we define $M = (1-q)\alpha(1/r + 1/e) + 1 + q/r_w$, it becomes:

$$\begin{aligned}
Mx_0 &= 1, \quad Mx_1 = (1-q)\alpha, \quad Mx_2 = (1-q)\alpha(1/r - 1), \\
Mx_3 &= q, \quad Mx_4 = q(1/r_w - 1), \quad Mx_5 = (1-q)\alpha/e.
\end{aligned} \tag{2}$$

The stationary probabilities (2) measure a few interesting characteristics of TAG operation, namely:

- x_5 —the proportion of time where the controlled device is allowed to stay in the DOWN state,
- $y = x_2 + x_4$ —the proportion of time where the node can act as an endorser, i.e., send a valid *Endorsement* message (note it cannot be done in the ENABLED or WAIT states),
- $w = x_0 + x_3 + x_4$ —the proportion of time where the node is waiting for a *Disable Authorization* message, hence the total risk of a major device crash during an emergency mode,
- $z = r(x_1 + x_2) + r_w(x_3 + x_4)$ —the mean frequency of sending a *Disable Request* message, the main communication cost of the TAG protocol.
- $v = (1-q)\alpha x_0$ —the mean frequency of transitions to the DISABLED state,
- $\phi = z/v - 1$ —the communication overhead of the TAG protocol, i.e., the mean number of excess *Disable Request* messages per one transition to the DISABLED state (1 is subtracted as one message is a necessary minimum).

Another characteristic of interest is the mean time u to a transition to the DISABLED state, starting immediately after a transition to the REQUEST SENT state. This measures the "unfulfilled request" time—the mean duration of the device operation in an emergency mode, where it is especially vulnerable to a major crash and so a prompt transition to the DISABLED state is desirable. Looking at possible transitions out of the REQUEST SENT state in Fig. 3 one derives a first passage-type equation in u :

$$u = (1-q)\alpha + (1-q)(1-\alpha)(u+1) + q(u+1+1/r_w). \tag{3}$$

On the other hand, u is composed of successive timeouts in the REQUEST SENT state and uninterrupted sojourn times in the WAIT and WAIT-CONT states. Hence,

$$u = \frac{1+q/r_w}{(1-q)\alpha} = \left(\frac{1}{q} + \frac{1}{r_w} \right) \phi. \tag{4}$$

C. Tradeoffs between Selected Characteristics

Intuitively, small values of x_5 should be postulated as this implies a small proportion of device DOWN time. Also, w and u should be kept small to minimize the risk of a major device crash. these are conflicting postulates: the shorter the node has to wait for a *Disable Authorization* message, the more frequently it enters the DISABLED state. Hence, there is a tradeoff between x_5 and w or u . How can it be quantified and can it be controlled by the TAG protocol configuration? After some algebra, (2) yields the following tradeoffs:

$$(e/r + 1)x_5 + w = 1, \quad (eu + e/r + 1)x_5 = 1. \tag{5}$$

Note that neither the TAG protocol configuration (q and r_w) nor the network size N or threshold attendance T influence the tradeoffs (5), the only relevant parameters being r and e , both related to the device reliability and maintainability.

Another tradeoff occurs between x_5 and the communication overhead ϕ . Here, too, small values of both characteristics are desirable. However, they are impossible to be achieved simultaneously because of the following relationship derived from (2):

$$(e(1/q + 1/r_w)\phi + e/r + 1)x_5 = 1. \tag{6}$$

Now q and r_w (through the sum of their reciprocals) are relevant for the tradeoff, yet N or T are not, similarly as in (5). The tradeoff (6) should be considered for q and r_w ensuring stable network operation, as discussed later.

IV. "ISOLATED" NODE-TYPE NETWORK-WIDE ANALYSIS

We notice that the probability α is the only element that couples the analyzed "isolated" node with the rest of the network. On the one hand, α is an input parameter of the presented Markov chain model and on the other hand can be calculated based on the solution (2). This leads to a fixed-point equation for α .

A. Use of Single-Node Stationary Probabilities

To account for the other nodes one proceeds as follows:

- 1) select an $\alpha \in [0, 1)$,
- 2) with e , q , r , and r_w treated as fixed, calculate $\mathbf{x}(\alpha)$ based on (2) and the selected α ,
- 3) calculate $y(\alpha) = x_2(\alpha) + x_4(\alpha)$,
- 4) using the definition of α and assuming inter-node statistical independence, state the fixed-point equation for α .



$$\alpha = \sum_{k=T}^{N-1} \binom{N-1}{k} y(\alpha)^k (1-y(\alpha))^{N-1-k}, \quad (7)$$

- 5) if the selected α equals the right-hand side of (7) up to a tolerable error then stop, otherwise go to step 1),

The above amounts to solving a nonlinear equation with a single unknown, of the form $\alpha = f(\alpha)$, where the function $f(\cdot)$ is expressed as the right-hand side of (7). Alternatively, (7) can be stated using the function $g(y) = \sum_{k=T}^{N-1} \binom{N-1}{k} y^k (1-y)^{N-1-k}$ as

$$\alpha = g(y(\alpha)), \quad (8)$$

B. Solution Existence and Uniqueness

The nonlinear equation (7) admits at least one solution $\alpha^* \in [0, 1)$, since $y(\alpha)$ is continuous in that interval and on its two edges takes the values

$$y(0) = \frac{q(1/r_w - 1)}{1 + q/r_w}, \quad y(1) = \frac{(1-q)(1/r - 1) + q(1/r_w - 1)}{(1-q)(1/r + 1/e) + 1 + q/r_w}. \quad (9)$$

Therefore, $f(\alpha)$ is continuous too, with $0 < f(0), f(1) < 1$. To examine the uniqueness of the solution note that $y(\alpha)$ is hyperbolic and so monotonous in $\alpha \in [0, 1)$. It is also nonincreasing iff $y(0) \geq y(1)$. After some algebra the latter condition becomes

$$q \geq \xi / (1/e + 1 - \xi), \quad (10)$$

where $\xi = (1/r - 1)/(1/r_w - 1)$. As explained below, $g(\cdot)$ can be viewed as a probability distribution function, hence nondecreasing for the whole range of $y \in [0, 1]$. Therefore, if (10) holds then $f(\alpha)$ is continuous and nonincreasing in $\alpha \in [0, 1)$ and so the solution of (7) is unique. Note that when $r_w = r$ (i.e., $\xi = 1$) condition (10) simplifies to $q \geq e$ and stipulates that the mean timeout duration in the REQUEST SENT state does not exceed the mean uninterrupted sojourn time in the DISABLED state.

If (10) does not hold, i.e., if $y(0) < y(1)$, one may notice that the function $y(\alpha)$ is concave in $\alpha \in [0, 1]$. To find the behavior of $f(\alpha)$ we use a representation $g(y) = Pr[U_{(T,N-1)} < y]$, where $U_{(T,N-1)}$ is the T^{th} order statistic among the set of $N-1$ independent random variables uniformly distributed on $[0,1)$. It is well-known that $U_{(T,N-1)}$ has the $Beta_{T,N-T}$ probability distribution with mean $\tau = T/N$ [10]. (Hence, $g(y) = I_{T,N-T}(y)$, where $I_{T,N-T}(\cdot)$ is the so-called regularized incomplete beta function.)

Second, since $N > 1$ and $T > 1$, this probability distribution is unimodal with a maximum at $y = \tau$, hence $g(y)$ is convex for $y < \tau$ and concave for $y > \tau$, with an inflection point at $y = \tau$. Therefore,

if $y(0) \geq \tau$ then the function $f(\alpha) = g(y(\alpha))$ is concave (as a superposition of two concave functions), as well as nondecreasing and with $f(0) > 0$. As such, it only admits a unique solution of (7). The condition $y(0) \geq \tau$ is equivalent of

$$q \geq \tau / ((1-\tau)/r_w - 1). \quad (11)$$

For large N and T , $U_{(T,N-1)}$ becomes approximately Gaussian, with mean τ . This stems either from a Taylor series expansion of $Beta_{T,N-T}$ around τ or from the fact that $U_{(T,N-1)} = U_{(1,N-1)} + (U_{(2,N-1)} - U_{(1,N-1)}) + \dots + (U_{(T,N-1)} - U_{(T-1,N-1)})$, where for large N the summands become almost statistically independent with mean $1/N$ [10] and the central limit theorem applies. As a result, $g(\cdot)$ can be approximated using the error function $erf(\cdot)$.

Conditions (10) and (11) are sufficient though not necessary for the uniqueness of the solution of (7). When neither of them holds, one may resort to numerical calculation of $f(\alpha)$. Parameter configurations yielding multiple solutions of (7) are possible. E.g., Fig. 4 illustrates the case with two solutions and the dependence of the solutions on q for very small q . The Markov chain analyzed in Sec. III then ceases to be ergodic. Fig. 3 shows that the reason is that the Markov chain is not irreducible—assuming $r_w > 0$, the WAIT and WAIT-CONT states become transient and the TAG operation becomes unstable. Mathematically, the only solution becomes $\alpha^* = 0$, at which $x_0 = 1$ and $x_i = 0$ for $i \neq 0$ (i.e., REQUEST SENT is an absorbing state and the other states are transient).




Fig. 4. Illustration of graphically solving (7) (top) and the solution values α^* (bottom) for very small q ($N = 100$, $T = 40$, $e = 0.04$, $r = 0.05$, $r_w = 0.06$).

An example region of stable TAG operation in Fig. 5 illustrates that the smaller the value of r_w , the better the Markov model reflects true TAG operation, i.e., stability for all $q > 0$.

V. MODELING STATISTICAL INTER-NODE DEPENDENCIES

The function $g(\cdot)$ defined in (8) reflects the assumption that the Markov chains produced by TAG operation at different nodes are statistically independent; such an assumption is characteristic of existing "isolated"-node network analyses. Various forms of statistical inter-node dependence might be incorporated into the model by generalizing the function $g(\cdot)$; note that the resulting analysis is model-free as it does not rely on assumed nature of

interactions between specific parameters at different nodes. Especially useful are inter-node dependence models introducing positive correlation between the nodes' ability to act as endorsers, for in this way one can reflect network-wide phenomena that tend to spread in time and space, such as mass disasters.




Fig. 5. Region of stable TAG operation ($N = 100$, $T = 40$, $e = 0.04$, $r = 0.05$).

One possibility is adopting a (slightly generalized) variation of the Pólya urn scheme [11]. In the classical formulation, balls of two colors are successively drawn from an urn and each drawn ball is replaced along with a number of same-color balls. In our generalization, neither initial nor replacement numbers of balls need to be integers. A thought experiment consists here in generating, for all the nodes except the "isolated" one being analyzed via the Markov chain model, a sequence of binary random variables $(\underline{X}_1, \dots, \underline{X}_{N-1})$, where $\underline{X}_j = 1$ represents the current ability of j^{th} node to act as an endorser, and $\underline{X}_j = 0$ represents its current inability to do so. The random draws of \underline{X}_j (i.e., 0s or 1s) is arranged so that the probability of drawing 1 for the j^{th} node is:

$$Pr[\underline{X}_j = 1] = (y + n_j \delta) / (1 + (j-1) \delta), \quad (12)$$

where $j = 1, \dots, N-1$, n_j is the number of 1s obtained in the previous draws, and $\delta \geq 0$ is a continuous-value parameter controlling the strength of the inter-node dependence (analogous to the number of same-color balls added to the urn after successive drawings in the classical scheme). In particular, by taking $\delta = 0$ one reverts to the statistical inter-node independence assumption. Under the Pólya urn scheme, the function $g(\cdot)$ is generalized as follows:

$$g(y) = \sum_{k=T}^{N-1} Pr[\sum_{j=1}^{N-1} \underline{X}_j = k], \quad (13)$$

where the constituent probabilities can be shown to be

$$Pr[\sum_{j=1}^{N-1} \underline{X}_j = k] = \binom{N-1}{k} \times \frac{\prod_{n=0}^{k-1} (y + n\delta) \prod_{n=0}^{N-2-k} (1 - y + n\delta)}{\prod_{n=1}^{N-2} (1 + n\delta)}. \quad (14)$$

Numerical calculation shows that the solution value α^* of (7) for a generalized function $g(\cdot)$ with $\delta > 0$ is typically smaller than that with $\delta = 0$; Fig. 6 shows an example. This quantifies the adverse impact of positive-correlation phenomena like mass disasters. With $\delta > 0$, the condition (11) remains in force, since except when T approaches N , $g(y)$ has been observed to have a single inflection point around $y = \tau$ or to be concave for all $y \in [0, 1]$. If (11) does not hold, stability is nevertheless observed for practically all $q > 0$, i.e., much more frequently than with $\delta = 0$.




Fig. 6. Comparison of solutions of (7) for statistical inter-node independence and for Pólya urn scheme ($N = 100$, $T = 40$, $e = 0.04$, $r = 0.05$, $r_w = 0.06$).

VI. NUMERICAL EXPERIMENTS

Due to the space limit we only present several numerical insights to illustrate the potential of the "isolated"-node approach, leaving more extensive results to a future publication. In the notation of Sec. III.B, we are interested in x_5 , the proportion of time in the DISABLED state, and u , the mean "unfulfilled request" time given by (4). For convenience, the latter is normalized to $1/r + 1/e$, a TAG independent quantity signifying the mean length of a device's UP?DOWN cycle; the resulting characteristic is denoted u_{norm} . With N , T , e , and r fixed, these two characteristics are controlled by the TAG configuration (q, r_w) . Recall that small values of either of them are desirable.

Fig. 7 compares x_5 obtained for various (q, r_w) TAG configurations under statistical inter-node independence ($\delta = 0$) and with a positive correlation-type inter-node dependence ($\delta = 0.5$). Each dot represents a pair of respective x_5 values corresponding to some TAG configuration (only configurations ensuring stable TAG operation are depicted). Regardless of the fixed parameters, a number of TAG configurations produce a slightly better characteristic under positive correlation (dots below the diagonal), however, for other configurations the characteristic worsens considerably under positive correlation. A tentative conclusion is that the former configurations capitalize on the nodes' positive correlation of the ability to act as endorsers, whereas the latter configurations are more prone to the effects of mass disasters.

Fig. 8 offers analogous insights regarding u_{norm} : here, depending on the fixed parameters, positive inter-node correlation may either invariably worsen the characteristic or produce various comparisons depending on the TAG configuration.




Fig. 7. Mean uninterrupted sojourn time in the ENABLED state; *top*: $N = 100$, $T = 40$, $e = 0.04$, $r = 0.05$, *bottom*: $N = 20$, $T = 10$, $e = 0.3$, $r = 0.2$.




Fig. 8. Mean "unfulfilled request" time; *top*: $N = 100$, $T = 40$, $e = 0.04$, $r = 0.05$, *bottom*: $N = 20$, $T = 10$, $e = 0.3$, $r = 0.2$.

VII. CONCLUSION

We have investigated the fulfillment of global consensus-type conditions in networked structures under a soft-crash reliability model subject to a "hard availability" threshold attendance re-

quirement. In this model, network nodes control the UP and DOWN states of devices that, after reporting failure, can still continue functioning in an emergency mode, albeit at a risk of a major crash. The consensus-type condition stipulates that a minimum threshold number of devices be in the UP state at any time. The underlying TAG protocol has a device reporting failure wait until permitted to enter the DOWN state and initiate appropriate repair/maintenance processes, and involves occasional control message exchange among the nodes and the Availability Supervisor. Using an "isolated" node-type Markovian analysis we have quantified the tradeoffs between relevant characteristics, such as the mean time until a device is allowed to enter the DOWN state (the risk of a major crash) and the necessary communication overhead. Accounting for a positive correlation between the nodes' behavior via a Pólya urn model has also been discussed, offering preliminary insights into the impact of network-wide phenomena that tend to spread, e.g., mass disasters.

Future work is planned to include, among others, subtler consensus-type conditions such as minimum attendance of specific node subsets or routing over disjoint paths, selfish refusal to send *Endorsement* messages, and detailed stochastic models of major crashes following failure reporting. Also, quantifying the emergency mode duration-dependent risk of a major crash should lead to a full-scale performability analysis. Simulative validation of the "isolated" node-type analysis in some environments is envisaged too. (Promising preliminary results have been obtained for a wireless sensor and actuator network environment with the data sink acting as Availability Supervisor.)

REFERENCES

- [1] A.S. Tanenbaum and M. van Steen. Distributed Systems: Principles and Paradigms, 2nd ed. Pearson Prentice-Hall, 2007.
- [2] W. Wang, Yi Xu, and M. Khanna, "A survey on the communication architectures in smart grid", *Comp. Networks*, vol. 55, 2011, pp. 3604–29.
- [3] J. Konorski and A. Makutunowicz, "Guaranteeing threshold attendance of W/WSAN nodes in a reverted security paradigm," *Proc. IEEE ICOIN*, Phuket, Thailand, Feb. 2014.
- [4] J. Turek and D. Shasha, "The many faces of consensus in distributed systems", *Computer*, vol. 25, no. 6, 1992, pp. 8–17.
- [5] G. B. Agnew and J. W. Mark, "Performance modeling for communication networks at a switching node," *IEEE Trans. Comm.*, vol. COM-32, 1984, pp. 902–910.
- [6] N. Shaham, "Stochastic models for multihop packet radio networks, in H. Takagi, ed. *Stochastic Analysis of Computer and Communication Systems*. North-Holland: Elsevier Science Publishers, 1990.
- [7] G. Bianchi, "Performance analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE J. on Selected Areas in Commun.*, vol. 18, no. 3, 2000, pp. 535–547.
- [8] J.F. Meyer, "On Evaluating the Performability of Degradable Computing Systems," *IEEE Trans. Computers*, vol. 29, no. 8, pp. 720-731, Aug. 1980.
- [9] R. Verdone, D. Dardari, G. Mazzini, and A. Conti. *Wireless Sensor and Actuator Networks: Technologies, Analysis and Design*. Academic Press, 2008.
- [10] W. Feller. *An Introduction to Probability Theory and its Applications*. New York: J. Wiley and Sons. 1966.
- [11] H. M. Mahmoud. *Pólya Urn Models*, Chapman and Hall–CRC, 2009.