On the Reliability and Availability of Systems Tolerant to Stealth Intrusion

Luís T. A. N. Brandão*† and Alysson Bessani†

*Electrical & Computer Engineering Department, Carnegie Mellon University – Pittsburgh, USA
†LaSIGE, Informatics Department, Faculty of Sciences of the University of Lisbon – Lisbon, Portugal
Lbrandao@{andrew.cmu.edu,di.fc.ul.pt}, bessani@di.fc.ul.pt

***Abstract*— This paper considers the estimation of reliability and availability of intrusion-tolerant systems subject to non-detectable intrusions. Our motivation comes from the observation that typical intrusion tolerance techniques may in certain circumstances worsen the non-functional properties they were meant to improve, such as dependability. We start by modeling attacks as adversarial efforts capable of affecting the intrusion rate probability of components of the system. Then, we analyze several configurations of intrusion-tolerant replication and proactive rejuvenation, to find which ones lead to security enhancements. We consider different attack and rejuvenation models and take into account the mission time of the overall system and the expected time to intrusion of its components. In doing so, we identify thresholds that distinguish between improvement and degradation. We compare the effects of replication and rejuvenation and highlight their complementarity, showing improvements of resilience not attainable with any of the techniques alone, but possible only as a synergy of their combination. We advocate the need for thorougher system models, by showing fundamental vulnerabilities arising from incomplete specifications.**

***Keywords*-reliability; availability; resilience; security; dependability; intrusion and attack models; intrusion tolerance; replication; rejuvenation; assumptions;**

I. Introduction

The design of *dependable and secure distributed systems* usually considers the use of enhancing techniques, such as replication [1] and rejuvenation [2], in order to cope with faults and intrusions. In this paper we highlight attack scenarios for which the *reliability* and *availability* of dependable systems capable of tolerating intrusions (a.k.a., intrusion-tolerant systems [3]) might be decreased, when compared to non-enhanced systems. We show how over-simplified system models, with incomplete specifications, may leave room for intrinsic vulnerabilities.

From a *reliability theory* [4] standpoint, fault tolerance has been extensively studied as a common approach to deal with fail-prone components. A straightforward intuition is that techniques of redundancy in space (*replication*) and time (*rejuvenation*) usually make a system more dependable. *Replication* enables a system to withstand the failure of some nodes (a.k.a., replicas or components) up to a certain fault tolerance threshold, e.g., f out-of n; *rejuvenation* (a.k.a., *repair* or *recovery*) allows malfunctioning nodes to be restored to a healthy state.

In the context of *malicious* attacks, intrusion tolerance techniques [3] go beyond traditional fault tolerance. Besides enabling dependable systems to cope with crashes and (typically random) abnormal behaviors, it also tolerates undetected *intrusions*, where parts of the system become under control of an adversary. Intrusion tolerance explicitly aims to preclude such intrusions from implying global security failures, e.g., loss of confidentiality. One could believe that techniques used for traditional fault tolerance imply the same improvements for intrusion-tolerant systems. However, different requirements usually imply different levels of sophistication and thus different characteristics. For example, the ratio (f/n), between the threshold of tolerable intrusions (f) and the degree of replication (n), usually decreases. Also, implementation of rejuvenations must become resilient to nodes that might have already been stealthily intruded. Because of these differences, the security-enhancement being sought may be jeopardized, if the estimation of *reliability* $(\mathcal{R})$ or *availability* $(\mathcal{A})$ are neglected.

It is often argued that one of the pitfalls of *replication* and *rejuvenation* techniques is their inability to cope with *common-mode failures*. We delimit our scope by not addressing the ways in which *diversity* (in space or time) can be adequately implemented to circumvent such problem (see examples in [5]). Nevertheless, we point out that, *even* assuming a probabilistic independence of intrusion between nodes, dependability properties may still be brought down in intrusion-tolerant systems. We consider that existing practical scenarios (as we shall exemplify) fit well our models.

This paper aims to promote thorougher system model specifications for dependable systems, in a way that allows the measurement of the dependability properties whose improvement is being attempted. We exemplify variations of *reliability* and *availability* brought upon by different *models of attack*, ratios of *fault threshold over replication degree* and *rejuvenation strategies*.

We summarize the contributions of this paper as follows:

1) we propose an *intrusion model* that is directly dependent of the *adversarial effort* for intruding nodes and compare results for different instantiations of attack;
2) we identify scenarios where *intrusion-tolerant replication* intrinsically decreases *reliability* and *availability* of a system under attack;

3) we find configurations towards *reliability* and *availability improvement goals*, for finite, unbounded and infinite mission times, identifying in particular a synergy between *replication* and *rejuvenation*.

The remainder of the paper is organized as follows: Section II introduces our preliminary system model; Section III illustrates analytic and quantitative results, focusing on *reliability*, and formalizes a notion of *resilience-improvement*; Section IV extends the system model to consider *rejuvenations* and presents results for systems where this technique is employed, now focusing on *availability*; Section V describes some related work; Section VI concludes the paper with some final remarks. In adittion, an Appendix presents mathematical details that sustain the results of the paper.

II. Preliminary System Model

Definition 1. An *intrusion-tolerant replicated* $\langle n, f\rangle$ system, with $0 < f < n$, is a system made up of n nodes, correct while the simultaneous number of *intruded* nodes does not exceed f. $\langle 1, 0\rangle$ is called the *reference system* – one that fails when its single node is intruded.

With "intruded" we intend a meaning more general than *faulty*. For example, we want to include cases where a node might continue to execute correctly, despite already under the control of a *malicious* adversary. Such control may be as subtle as being able to decide at any time to interfere with the service running on the node.

Initially we are interested in comparing characteristics of $\langle n, f\rangle$ with those of $\langle 1, 0\rangle$, when the former is built as an architectural enhancement of the later, using intrusion-tolerant replication. Many implementations fit this model. For example: $f = n - 1$ for some synchronous crash fault-tolerant (*Crash FT*) protocols [1]; $f = \lfloor (n-1)/2 \rfloor$ for some *Byzantine fault-tolerant* (BFT) systems with synchrony [1] or using trusted components (e.g., [6]); $f = \lfloor (n-1)/3 \rfloor$ for general BFT systems (e.g., [7], [8]).

Definition 2. The *mission time* (MT) of a system is the uninterrupted interval of time during which the system is intended to be correct. MT may be finite and known, finite but unknown, or infinite.

Definition 3. The *reliability* ($\mathcal{R}$) of $\langle n, f\rangle$ is the probability that the system will never fail during its MT.

Definition 4. The *availability* ($\mathcal{A}$) of $\langle n, f\rangle$ is the probability that the system is not failed at an instant of time randomly chosen, uniformly from the MT period.

Definition 5. A dependability property (e.g., $\mathcal{R}$ or $\mathcal{A}$) of a $\langle n, f\rangle$ system is said to be *desirable* if it is *better* than in $\langle 1, 0\rangle$. For example, if $\mathcal{R}_{n,f} > \mathcal{R}_{1,0}$, then $\langle n, f\rangle$ is said to have desirable $\mathcal{R}$.

Assumption 1 (Intrusion model). *The system has a* $\langle n, f\rangle$ *architecture, with state represented by a vector* $\vec{\phi}$ *of length* n*. The state of each node* j*, with* $j \in \{1, ..., n\}$*, is given by* $\phi_j \in \{0, 1\}$*, with 0 for* healthy *and 1 for* intruded*. Each node starts in a* healthy state *and transitions probabilistically to an* intruded state*, according to an* intrusion rate probability *(IRP) density* $\lambda_j(t)$*, at each instant* (t)*, directly proportional to an* intrusion adversarial effort *(IAE) exerted on the node. The proportionality ratio IRP/IAE is the same for all nodes.*

The distinction between IRP and IAE allows avoiding the specification in advance of the behavior of the attacker. The proportionality relation implies that all nodes have the same probability of being intruded when subjected to the same IAE, even though an attacker may still choose to attack different nodes with different variations of effort. In our simple examples, we shall omit the proportionality constant and use $\lambda_j(t)$ indistinctively to specify IRP and IAE. Following a conservative estimation of reliability, it is considered that global *failure* of a $\langle n, f\rangle$ system occurs at the first instant of time in which more than f nodes are in *intruded* state. We now proceed with two *alternative* attack models, both of practical interest.

Assumption 2 (Attack models). *The system will be attacked in one of the following manners:*

- **Parallel Attack ($\|$)** – *The IAE is equal on all healthy nodes and has constant intensity* (λ)*.*
- **Sequential Attack ($\therefore$)** – *The IAE targets one healthy node at a time, with constant intensity* (λ)

Formally (and omitting t): a $\|$-attack satisfies $(\forall j)$ $(\lambda_j \equiv \lambda \times (1 - \phi_j))$; a $\therefore$-attack satisfies $(\exists j : \phi_j = 0) \Leftrightarrow (\exists j)\,[(\lambda_j \times (1 - \phi_j) = \lambda) \wedge (\forall j' \neq j)\,(\lambda_{j'} = 0)]$.

These two assumptions do not consider cases of exploitation of common vulnerabilities capable of leading to immediate simultaneous intrusion. Our implicit assumption is that *replication* considers intentional *diversity* to cope with some vectors of attack [5], from which systems usually try to protect themselves. Nonetheless, we intend to show that **even with independence of intrusions** *reliability* and *availability* can still be lowered by the use of *replication*. It is also worth mentioning a commonly overlooked fact: although *independence* is better than a possibility of *simultaneous collective intrusion*, it is not an optimal situation. As noted in [9], "better than independence can actually be attained".

Actually, there are two orthogonal axes of dependence: one is respective to intrusions (our model is indeed of independence, because the ratio IRP/IAE is a constant, which implies that intrusions do not become easier with time nor after other intrusions); another is respective to architectural aspects of attack (e.g., in the $\|$-attack model nodes are attacked independently of others, whereas in the $\therefore$ there is a *good* dependence in that each node under attack protects all remaining healthy nodes from being attacked).

At first glance, it may seem that the different attack models are only a matter of an attacker's choice, i.e., choosing $\|$

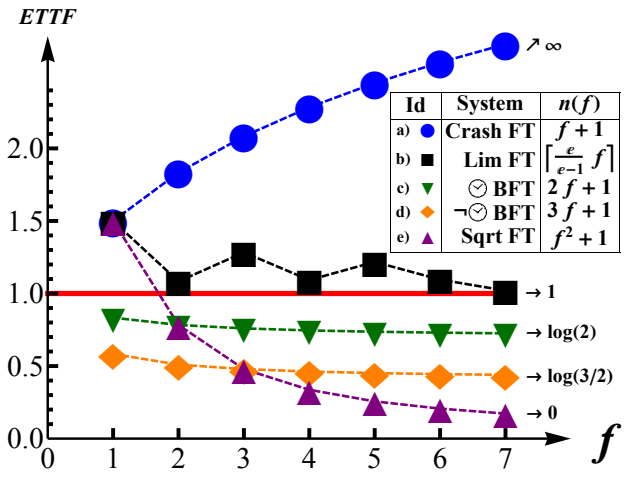

Figure 1. ETTF for *parallel* attack model. The horizontal (red) line highlights the reference ETTF of a single-node system if $\mu_{1,0} = 1$. The value to the right of each curve indicates the limit ETTF as $f \to \infty$. The vertical axis actually measures $b_{n,f} = \mu_{n,f}/\mu_{1,0}$, but for $\mu_{1,0} = 1$ it is equal to the ETTF of the respective $\langle n, f\rangle$ system.

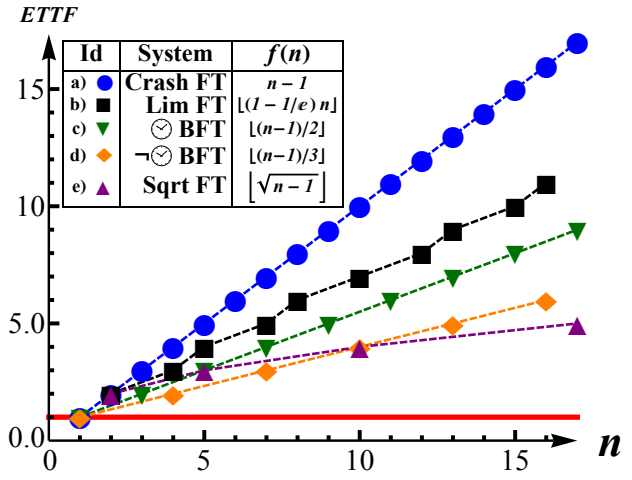

Figure 2. ETTF for *sequential* attack model. For each system, the horizontal coordinate of the plotted points is the minimal n that allows a certain threshold f. Note some awkwardnesss in system *Lim FT* (*b*, ■), where the ratio f/n is not monotonic for the sequence of plotted points (enabling f from 0 to 10) – note in $f(n)$ the division by a non-integer.

versus $\therefore$. attack and/or choosing *dispersed* (lower λ) versus *focused* (higher λ) *effort*. However, these options might be constrained, for example if the system's architecture does not expose itself to $\|$ attacks. Also, for some types of attack, the malicious goal of stealthiness may limit the *effort* on each node (e.g., too many incorrect password attempts per day may trigger some alarm). In such cases, a $\|$-attack on n nodes might not be replaceable by a $\therefore$-attack with a focused *effort* n times higher in a single node at a time, because otherwise it would be detected.

Some examples consistent with our assumptions:

- A set of nodes, each protected with a random-one-time-password, for a $\|$-attack using random password attempts, with equal frequency in all nodes.
- A set of software nodes, each diversified with *instruction set randomization* [9], for a buffer-overflow $\|$-attack. Here, the vulnerability might not be eliminated, but the way to exploit it becomes obfuscated – the code injection leading to intrusion varies per node.
- A set of nodes geographically dispersed, for a social-engineering $\therefore$-attack, requiring human presence and performed by a single person.

III. Time, Reliability, Resilience

In this section we try to determine the *reliability* ($\mathcal{R}$) of $\langle n, f\rangle$ systems, under both models of attack, considering different perspectives:

1) Which $\langle n, f\rangle$ systems have a *desirable expected time to failure* (ETTF)?
2) For which *mission time* (MT) does a $\langle n, f\rangle$ system have a *desirable* $\mathcal{R}$?
3) Given a MT, a goal of $\mathcal{R}$ and a functional relation between the *replication degree* n and the *intrusion tolerance threshold* f, how to adjust f or n?
4) How to define goals of $\mathcal{R}$*-improvement* and how to achieve them?

A. Expected Time to Failure

For $\langle 1, 0\rangle$ the parallel ($\|$) and sequential ($\therefore$) models of attack are equivalent. In both, the probability that the single node becomes intruded follows an exponential distribution (see Appendix). Here, the ETTF ($\mu_{1,0} = 1/\lambda$) is the inverse of the node's *intrusion rate probability* (IRP) (λ).

The ETTF is a metric often used for a quick intuition about the reliance of a system. Also, the MT of a system is often defined as a function of its ETTF. Thus, we now determine the circumstances in which the ETTF increases or decreases with the number of nodes. Let $\mu_{n,f}$ stand for the ETTF of a $\langle n, f\rangle$ system. We want to know if the ratio $b_{n,f} = \mu_{n,f}/\mu_{1,0}$ is higher or lower than 1.

To be practical, we shall group systems by functional relations $n(f)$, e.g., $n = f+1$ or $n = 2f+1$, typically characterizing different types of protocol and types of tolerated faults. We shall use suggestive names, such as Crash and (synchronous or asynchronous) BFT, to label such groups, but the analysis will be based not on the type of faults, but only on the relation between n and f.

For ***parallel*** attacks, the ratio is $b_{n,f} = H_n - H_{n-(f+1)}$, as deduced in [10], with $H_n = \sum_{i=1}^{n} i^{-1}$ being the *harmonic-number* function. Figure 1 shows curves for several cases. In the extreme of higher ETTFs is the type of system (●) that works correctly while at least one node is healthy ($f = n - 1$), having $\mu_{n,n-1} = \mu_{1,0} \times H_n$. The system maintains a desirable ETTF while $\mu_{n,f} > \mu_{1,0}$.

When the *intrusion tolerance threshold* ratio f/n decreases below a certain limit, the system eventually transitions to an undesirable ETTF. The *Lim FT* curve (■) illustrates, for several values f, the limit case of desirable ETTF. Asymptotically (in the limit $n \to \infty$), the transition occurs for $f/n = (1 - 1/e) \approx 0.63$, with $e \approx 2.718$ being Euler's number. For lower f/n ratios, the global ETTF decreases while the threshold f increases, as seen in curves with $f = (n-1)/2$ (▼) and $f = (n-1)/3$ (♦), as typically used in intrusion-tolerant systems, namely of type *Byzantine fault-tolerant* (BFT). Though decreasing, for these cases the ETTF still converges to a positive value. For example, with $f = (n-1)/3$ the ETTF tends to $\log(3/2) \approx 40.5\%$ of $\mu_{1,0}$. In a further extreme, when the ratio f/n itself converges to 0 while increasing f, the ETTF also converges to 0, as shown with the *Sqrt FT* curve (▲), with $n = f^2 + 1$. The extreme case happens without intrusion tolerance, i.e., $f = 0$ (not shown in the figure), implying a global ETTF of $\mu_{n,0} = \mu_{1,0}/n$, with faster convergence to 0 as n increases.

For ***sequential*** attacks, the ETTF is much higher, with $b_{n,f} = f + 1$ (also deduced in [10]), if λ is fixed when varying n. Each node has an *expected time to intrusion* of $\mu_{1,0}$, once it starts being attacked. The higher increase of ETTF with f is now the result of a (*good*) dependence between the *intrusion adversarial effort* (IAE) on different nodes. Intuitively, a node being attacked draws all the attention from the attacker and thus, while *healthy*, it protects the other nodes from being attacked. Figure 2 highlights the ETTF in function of n, for different $\langle f, n \rangle$ systems. Note that if this graphic was plotted in function of f, all curves would superpose, as $\mu_{n,f}$ is now a pure function of f.

In conclusion, the differences in types of attack ($\|$ versus $\therefore$), may make the difference between improving or worsening the ETTF of a system, when *augmenting* its configuration from $\langle 1, 0 \rangle$ to $\langle n, f \rangle$. This should bring to attention the importance of considering architectural aspects, that may limit the types of attack, when deciding on how to achieve *intrusion tolerance*.

B. Reliability per Mission Time

The ETTF is a useful metric, but there is no fundamental reason for it to be the desired MT. Thus, we now consider a more dynamic perspective and analyze the *reliability* ($\mathcal{R}$) for different MT values. We are interested in knowing *what are the MT for which intrusion tolerance does not worsen the system's reliability, when compared to that of* $\langle 1, 0 \rangle$. This information is important when one wants to define an adequate MT given a $\langle n, f \rangle$ system, or, vice-versa, select the best $\langle n, f \rangle$ system given a predetermined MT.

We shall use symbol τ to express time with an implicit unit of $\mu_{1,0}$ (the *expected time to intrusion* of a node under attack). In Appendix we include the explicit mathematical formulas for $\mathcal{R}$, in both attack models, involving hypergeometric functions.

Table I
RELIABILITY ($\mathcal{R}$) UNDER PARALLEL ATTACK

System Type	*n*	*f*	$\mathcal{R}$				
			$\tau = 0.2$	$\tau = 0.5$	$\tau = 1$	$\tau = 2$	$\tau = 5$
Reference	1	0	0.819	0.607	0 368	0.135	0.00674
No FT	2	0	0.670	0.368	0 135	0.0183	0.0000454
Crash FT	2	1	*0.967*	*0.845*	*0.600*	*0.252*	*0.0134*
⍉ BFT	3	1	*0.913*	*0.657*	0 306	0.0500	0.000136
Crash FT	3	2	*0.994*	*0.939*	*0.747*	*0.354*	*0.0201*
¬⍉ BFT	4	1	*0.847*	0.487	0 144	0.00891	1.22×10^{-6}
Lim FT	4	2	*0.979*	*0.828*	*0.469*	0.0911	0.000270
Crash FT	4	3	*0.999*	*0.976*	*0.840*	*0.441*	*0.0267*
⍉ BFT	5	2	*0.955*	*0.694*	0 264	0.0200	3.03×10^{-6}
¬⍉ BFT	7	2	*0.883*	0.434	0.0684	0.000751	2.88×10^{-10}
⍉ BFT	7	3	*0.976*	*0.723*	0 230	0.00834	7.10×10^{-8}
Lim FT	7	4	*0.997*	*0.910*	*0.509*	0.0568	0.0000105

Highlighted in blue, italic and slightly larger font-size are the cases with *desirable* $\mathcal{R}$, if compared with a $\langle 1, 0 \rangle$ system with the same MT.

Parallel Attack Model. Table I and Figure 3 show the variation of $\mathcal{R}_{n,f}(\tau)$ for several pairs $\langle n, f \rangle$. When *little* time has passed, an intrusion-tolerant system with $f > 0$ has desirable $\mathcal{R}$, because it is not yet likely that *many* nodes have been intruded. As time passes, more nodes are likely to have been intruded and thus a low ratio f/n may imply lower $\mathcal{R}$. In Figure 3 we show solutions (τ_{max}) of the MT for which $\mathcal{R}$ transitions from *desirable* to *undesirable*. In other words, $[0, \tau_{max}]$ is the interval for which $\mathcal{R}_{n,f}(\tau) \geq \mathcal{R}_{1,0}(\tau)$.

For example, consider a context that requires $n = 3f + 1$ and for which each node under attack has an estimated *expected time to intrusion* of 1 year. From Table I, we see that, when compared to $\langle 1, 0 \rangle$, a system $\langle 4, 1 \rangle$ has *desirable* $\mathcal{R}$ for $\tau = 0.2$, i.e. a MT of 2.4 months, because $(\mathcal{R}_{4,1}(0.2) > \mathcal{R}_{1,0}(0.2))$. However, for $\tau = 0.5$, i.e., a MT of 6 months, the respective $\mathcal{R}$ is *undesirable*, because $(\mathcal{R}_{4,1}(0.5) < \mathcal{R}_{1,0}(0.5))$. In Figure3 we determine $\tau = 0.264$ as the transition value for $\langle 4, 1 \rangle$.

This example illustrates why *intrusion tolerance* is not on its own aligned with dependability. In this case (parallel attacks and $n = 3f + 1$), the specification of MT (or, more precisely, $MT/\mu_{1,0}$) is needed to determine if intrusion tolerance brings an advantage or a disadvantage.

On a more global look to Table I and Figure 3, we note that different functional relations between n and f imply different MT-ranges of desirable $\mathcal{R}$:

1) any MT – e.g., the *Crash FT* curve is higher than the *Reference* curve for any positive MT;
2) MT up to some $\tau_{max} > 1$ – e.g., the *Lim FT* curve (with $f \geq 2$) intersects the *Reference* curve for $\tau > 1$;
3) MT up to some $\tau_{max} < 1$ – e.g., the BFT curves (with $f > 0$) intersect the *Reference* curve for $\tau < 1$;
4) never – e.g., system $\langle 2, 0 \rangle$ in Table I has $\mathcal{R}$ lower than the *Reference*, for any positive MT.

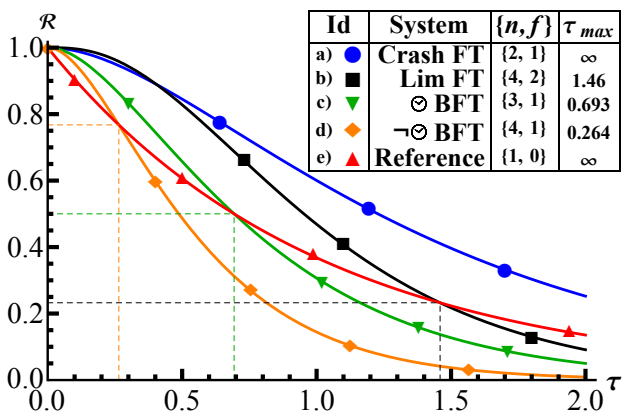

Figure 3. Reliability ($\mathcal{R}$) as a function of *mission time* (MT), for *parallel* attacks. Horizontal axis is $\tau = MT/\mu_{1,0}$; vertical axis is $\mathcal{R}$ (higher is better). Each curve represents $\mathcal{R}_{n,f}(\tau)$ for some pair $\langle n, f\rangle$. τ_{max} is the value satisfying $\tau \in [0, \tau_{max}] \Leftrightarrow R_{n,f}(\tau) \geq R_{1,0}(\tau)$.

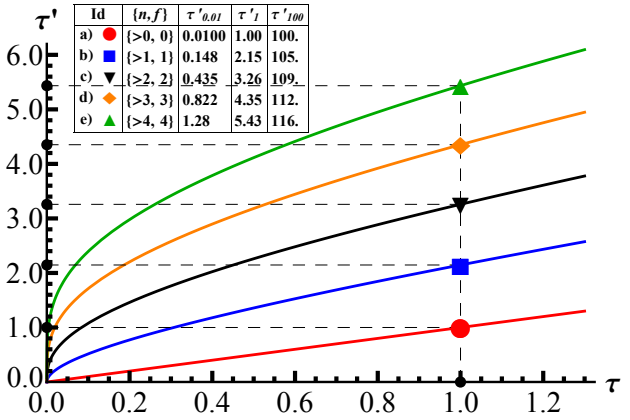

Figure 4. *Mission time* (MT) for same *reliability* ($\mathcal{R}$), for *sequential* attacks. Each curve is associated with an *intrusion tolerance threshold* f, being independent of n. Each point $\langle \tau, \tau' \rangle$ means that, for a MT τ in $\langle 1, 0\rangle$, the respective $\langle n, f\rangle$ system has the same $\mathcal{R}$ when its MT is τ'. Curve *a* (●), with $f = 0$, is the identity τ=τ'.

Sequential Attack Model. In this model, the time required to intrude more than f nodes is independent of the total number of nodes (n). For any MT, the $\mathcal{R}$ always grows with the intrusion-tolerance threshold f. Formally, $(\tau > 0 \wedge f > f' \geq 0) \Rightarrow \mathcal{R}_{n,f}(t) > \mathcal{R}_{n,f'}(t)$. Still, for any $\langle n, f\rangle$, $\mathcal{R}$ still converges to 0, as time increases. Due to space constraints we omit a table with numerical values of $\mathcal{R}$.

For the $\therefore$-attack model, a graphic equivalent to the one in Figure 3 would have no curve intersections, thus, we proceed directly to a new perspective with Figure 4, showing how, for a fixed $\mathcal{R}$, an increase of f allows an increase of MT. This graphic allows the determination of the adequate f, whenever wanting to increase the MT while maintaining the $\mathcal{R}$ of the overall system. Note that near $\tau = 1$ the multiplicative factor of MT-improvement is approximately linear with f, but for smaller values of τ the MT-improvement is much higher. For example, if having a $\langle 1, 0\rangle$ system being used for a MT $\tau = 0.01$, then a threshold $f = 4$ yields the same $\mathcal{R}$ for a new MT of $\tau' = 1.28$, i.e. 128× bigger. However, if $\mathcal{R}$ is instead compared with that of $\langle 1, 0\rangle$ for MT $\tau = 1$, then the replicated system can only be used for $\tau' = 5.43$. The analytic solution for $\mathcal{R}^{\therefore}_{1,0}(\tau) = \mathcal{R}^{\therefore}_{n,f}(\tau')$, in order of τ', is presented in Appendix, Equation 9.

C. *Transition times for* resilience

It is easy to understand what it means to increase the MT by a multiplicative factor. However, with $\mathcal{R}$ (a probability), the scale is not linear and thus it may not be meaningful to ask for a linear improvement (e.g., double the $\mathcal{R}$). To deal with this, Equation 1 defines a new scale, to which we suggestively call *resilience* (ρ), increasing linearly with the number of bits with which the $\mathcal{R}$ approximates to 1.

$$\rho_{n,f}(\tau) = -\log_2(1 - \mathcal{R}_{n,f}(\tau)) \quad (1)$$

Note the resemblance with a common notion of *security* of cryptographic schemes [11], in which a system has k bits of security if the best attack known to *break* it requires a *search* of complexity $O(2^k)$. Note also that, if using a logarithmic base 10, this would resemble the typical measure of "nines" for *availability* ($\mathcal{A}$). We now find the values of MT for which the ρ of $\langle n, f\rangle$ is c-times higher than that of $\langle 1, 0\rangle$. By solving $\rho_{n,f} \geq c \times \rho_{1,0}$, in order of τ, we get:

$$\mathcal{R}_{n,f}(\tau) \geq 1 - (1 - \mathcal{R}_{1,0}(\tau))^c \quad (2)$$

Table II presents some numerical solutions for the limits of MT for which $\langle n, f\rangle$ systems should be designed for, when considering a desired of ρ-*improvement* factor c, for a $\|$ attack model. Some interesting facts:

- Every $\langle n, f\rangle$ system has a maximum resilience factor c that it can sustain. For $n = f + 1$, the ρ-improvement factors are valid either for any MT ($\tau_{max} = \infty$) or for none at all. For the other illustrated systems, improvement factors are valid only for a finite duration of time.
- For $n > 1$, $f = 0$ always implies lower ρ, i.e., $(\forall n > 1)(\forall c \geq 1)(\forall t > 0)(\rho_{n,0} < c \times \rho_{1,0})$.
- For any $f \geq 1$, some ρ-improvements (i.e., $c > 1$) can be obtained for a small MT. However, only large ratios f/n allow ρ-improvements for large values of MT.

As an example of interpretation, consider $\tau_{max} = 0.0319$, obtained in Table II for $c = 2$ and $\langle 7, 2\rangle$, a possible BFT system with configuration $n = 3f + 1$. From Equation 4, in Appendix, we calculate the *reliability* ($\mathcal{R}$) for such MT. For $\langle 1, 0\rangle$ we get $\mathcal{R}_{1,0}(0.0319) \approx 0.969\%$, corresponding to a *resilience* of $\rho \approx 5.0$. For $\langle 7, 2\rangle$ we get $\mathcal{R}_{7,2}(0.0319) \approx 0.999\%$, i.e., $\rho \approx 10.0$. Thus, we have $\rho_{7,2} \approx c \times \rho_{4,1}$, with $c = 2$, for $MT \approx 0.0319 \times \mu_{1,0}$. If, for example, $\mu_{1,0}$ is 1 year, then a $\langle 7, 2\rangle$ system doubles the ρ of a non-replicated system if it is used for only 11.4 days (0.0319×1 year).

Table II
ρ-IMPROVEMENT FOR PARALLEL ATTACKS

System Type	n	f	$\tau_{max} : \rho_{n,f} \geq c \times \rho_{1,0}$						
			$c=0.1$	$c=0.5$	$c=1$	$c=1.25$	$c=1.5$	$c=2$	$c=3$
Reference	1	0	∞	∞	∞	0	0	0	0
No FT	2	0	2.25	0.481	0	0	0	0	0
Crash FT	2	1	∞	∞	∞	∞	∞	∞	0
$\odot$ BFT	3	1	3.36	1.59	0.693	0.382	0.144	0	0
Crash FT	3	2	∞	∞	∞	∞	∞	∞	∞
$\neg\odot$ BFT	4	1	1.73	0.746	0.264	0.120	0.0306	0	0
Lim FT	4	2	4.06	2.33	1.46	1.15	0.871	0.405	0
Crash FT	4	3	∞	∞	∞	∞	∞	∞	∞
$\odot$ BFT	5	2	2.19	1.20	0.693	0.512	0.360	0.129	0
$\neg\odot$ BFT	7	2	1.14	0.579	0.296	0.201	0.128	0.0319	0
$\odot$ BFT	7	3	1.78	1.07	0.693	0.559	0.445	0.259	0.0313
Lim FT	7	4	2.82	1.86	1.36	1.18	1.03	0.761	0.337

Cells contain τ_{max}, the maximum τ for which $\rho_{n,f}(\tau) \geq c \times \rho_{1,0}(\tau)$. Highlighted in blue, italic and slightly larger font-size are the cases where the respective improvement factor is valid for τ up to at least 1.

Due to space constraints we omit the respective table for the ***sequential*** case, for which *resilience* is generally better.

IV. AVAILABILITY AND THE ROLE OF REJUVENATIONS

In this section we analyze the security augmentation brought upon by the use of *rejuvenation* [7], [12], [13]. Consistently with our model of intrusions and attacks (assumptions 1 and 2), we assume that the eventual intrusion of a node, at a given time, does not make easier the future intrusion of other nodes or of the same node after rejuvenation. As mentioned in Section I, this type of independence is usually achieved by the use of *diversity*, which can be effective for certain vectors of attack. However, within our scope, we keep agnostic to such implementations and simply assume they are effective for our purposes.

In the interest of space we have omitted from the previous section the analysis of *availability* ($\mathcal{A}$). Its focus is not on the first global failure (probability of never failing), but instead on the accumulated delivery of service (probability of not being failed at a random instant). When not considering *rejuvenations*, $\mathcal{A}$ can be deduced by integrating the *reliability* ($\mathcal{R}$) across time (Equation 8 in Appendix). Now that we consider rejuvenations, it becomes more pertinent to consider $\mathcal{A}$ (Equation 11 in Appendix). Both $\mathcal{R}$ and $\mathcal{A}$ increase with rejuvenations, because it becomes more difficult for an attack to succeed in surpassing the *intrusion tolerance threshold* f. However, $\mathcal{A}$ has the extra benefit of accounting also the moments of correctness obtained after a first global failure. Thus, $\mathcal{A}$ is positive even for an infinite *mission time* (MT). This is relevant whenever global failure is not considered a catastrophic event and the reestablishment of service is considered worthy.

A. System Model Extension

If we could detect attacks and/or intrusions, then a reactive-rejuvenation scheme could be implemented [14]. For example: a detected attack could be mitigated by rejuvenating components more frequently; a detected intrusion could be amended by immediately rejuvenating the respective node. However, in our context of stealthiness, we can rely only on proactive-rejuvenation schemes.

We now proceed with two models of *proactive*-rejuvenation: *parallel* ($\|$) and *sequential* ($\therefore$).

Assumption 3 (Periodic Rejuvenations). *In a $\langle n, f\rangle$ system, let $\Delta > 0$ be the time between (periodic) rejuvenations of the same node. Let $\delta < \Delta$ (with $\delta \geq 0$) be the smallest time between rejuvenations of different nodes. For all $j \in \{1, ..., n\}$, node j begins its i^{th}* rejuvenation *(with $i \in \mathbb{N}_1$), at instant $(j-1) \times \delta + i \times \Delta$ and completes it in a constant amount of time r, with $r = \delta \times k$, $r = \Delta \times (k/n)$ and $k \in \mathbb{N}_0$. Rejuvenations can be of two types: **parallel** ($\|$), if $\delta = 0$, or **sequential** ($\therefore$) otherwise. A system without rejuvenation is as a $\|$-rejuvenating system with $\Delta = \infty$.*

Some observations about Assumption 3:

- For $\|$-rejuvenations we have $r = k = 0$, since nodes rejuvenate simultaneously ($\delta = 0$) every Δ time units and $k \in \mathbb{N}_0$ is constant.
- The number of *offline* nodes (i.e., being rejuvenated) at any instant in time is either $k = 0$, if $r = 0$, or a positive integer $k = r/\delta$. Thus, the number of *online* nodes is constant: $n' = n - k$. Each node is *online* for durations $\Delta - r$, interleaved with *offline* durations r.
- n now accounts also with the k offline nodes.
- Parameters r, δ and Δ will be expressed in time units of $\mu_{1,0}$, using symbol τ, as was already done with MT.

The choice of the rejuvenation scheme might not be arbitrary. For example, if the system implements non-stop operations, the rejuvenation process might require transfer of state from online nodes to rejuvenating nodes, thus making a $\therefore$ scheme more appropriate than a $\|$ one. In such cases, parameters k and r are relevant in terms of implementation. Actually, inability to enforce a fixed bounded limit on r may result in security vulnerabilities for some protocols, as noted for example in [12]. However, for the purpose of estimating $\mathcal{R}$ or $\mathcal{A}$, a $\therefore$-rejuvenating system, with n nodes and r time per individual recovery, is equivalent to one with $n' = n - k$ nodes and instantaneous rejuvenation ($r' = 0$) of nodes. Thus, henceforth we shall use a reduced notation (in subscript) to distinguish the type of rejuvenation:

- $\langle \|, \Delta \rangle$: *parallel* ($\|$) rejuvenations with period Δ and assuming $\delta = 0$.
- $\langle \therefore, \delta \rangle$: *sequential* ($\therefore$) rejuvenations, with consecutive nodes being rejuvenated at instants separated by δ.

B. Parallel rejuvenation

On each *parallel* ($\|$) rejuvenation, a $\langle n, f\rangle$ system resets to a completely healthy state. Formally, $\sum_{j=1}^{n} \phi_j(\Delta \times i) = 0$, for $i \in \mathbb{N}_1$. Thus, the overall *reliability* $\left(\mathcal{R}_{n,f,\|,\Delta}\right)$ can

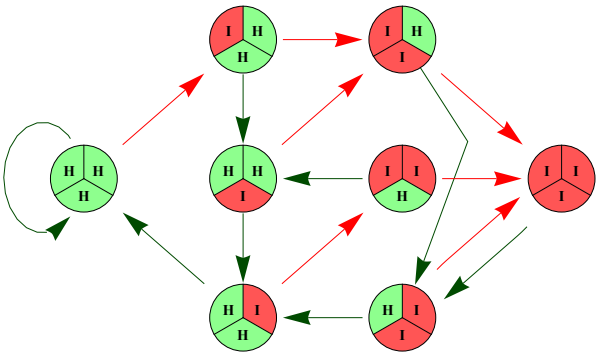

Figure 5. State diagram for *sequential* attack and *sequential* rejuvenation. Each circle represents a set of $n = 3$ nodes and their states: *healthy* (H) or *intruded* (I). A rejuvenation *heals* ($H \rightarrow H$ or $I \rightarrow H$) the right-upper triangle and then rotates the circle counter-clock-wise. An intrusion *intrudes* ($H \rightarrow I$) the healthy triangle further away from recovery.

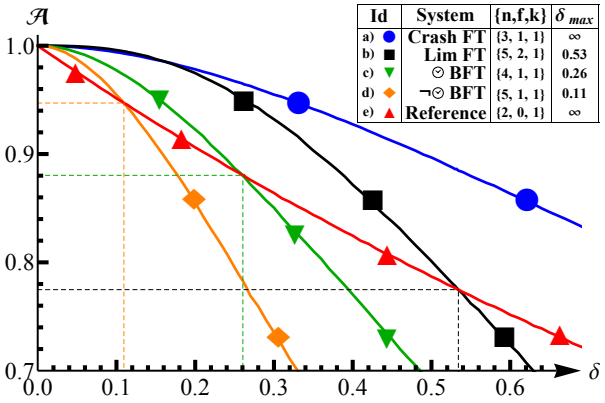

Figure 6. Availability ($\mathcal{A}$), for *sequential* attack and *sequential* rejuvenations. Horizontal axis measures δ – time offset between rejuvenations of different nodes. Vertical axis measures $\mathcal{A}$ – proportion of time for which the number of intruded nodes is at most f. For each $\langle n, f \rangle$ system, pairs $\langle \delta, \mathcal{A} \rangle$ were obtained, with δ spaced by intervals of at most 0.01 and $\mathcal{A}$ being an average of 10 probabilistic simulations of $\mathcal{A}\left(\tau = \delta \times 10^5\right)$.

be obtained (Equation 10 in Appendix) as a product of *reliabilities* ($\mathcal{R}_{n,f}$) in time-windows of width Δ or less.

Notably, for the reference system $\langle 1, 0 \rangle$ (or any other with $f = 0$), *rejuvenation* does not affect $\mathcal{R}$, because: (1) the *intrusion* of a node corresponds to the immediate failure of the system; and (2) the rejuvenation of a healthy node does not alter its *intrusion rate probability* (IRP). Consequently, if it is not possible to have fault tolerance, then a $\mathcal{R}$-improvement can only be obtained by using more reliable nodes. However, for $\langle n, f \rangle$ systems with $f > 0$, rejuvenation allows the healing of *intruded* nodes before the number of simultaneous intrusions exceeds f. This discussion shows that replication and rejuvenation have complementary roles:

- intrusion-tolerant replication, with $f > 0$, improves $\mathcal{R}$ for small MT, *but* it is prejudicial for large MT;
- rejuvenation cannot bring benefits before its first application, *but* it truncates the long-term degradation of the system, periodically bringing it back to a young stage.

By applying both techniques together, the $\mathcal{R}$-improvement might be valid even for an unbounded MT (finite but not known in advance). To achieve such overall improvement, for a $\langle n, f \rangle$ system, the period Δ should be less than the value of time (in Figure 3) for which $\langle n, f \rangle$ without rejuvenation has *undesirable* $\mathcal{R}$. In this way, for example even BFT systems under parallel-attack can have *desirable* $\mathcal{R}$ for unbounded MT. However, if $MT = \infty$, then the $\mathcal{R}$ of any $\langle n, f \rangle$ system is 0, whereas the $\mathcal{A}$ is still positive.

C. Sequential rejuvenation

In a *sequential* ($\therefore$) rejuvenation scheme, even though the rejuvenation instants are periodic, there is no periodic instant where the overall system state is deterministically reset. Thus, a strong-enough attacker may have a high probability of intruding nodes at a faster pace than their rejuvenation, consequently maintaining the number of intruded nodes above the threshold f for durations much longer than Δ. In a $\therefore$-attack-model, we assume an optimal *intrusion adversarial effort* (IAE) sequence, with the attacker always attempting to intrude the node which will remain un-rejuvenated for the longest time. Fig 5 shows the respective diagram.

The rules of probabilistic transition between states are easy to define and simulate. As an example, Figure 6 shows results for a $\therefore$-attack model, when varying δ (the offset between $\therefore$-rejuvenations). We consider cases with $k = 1$ and thus $\delta = r$. The graphic shows that different $\langle n, f \rangle$ systems have *desirable* $\mathcal{A}$ (i.e., higher than in $\langle 1, 0 \rangle$) for different offsets δ of rejuvenations: for any δ if $\langle 2, 1 \rangle$; only for $\delta \lesssim 0.26$ if $\langle 3, 1 \rangle$; only for $\delta \lesssim 0.1$ if $\langle 4, 1 \rangle$. Due to space constraints we skip the case of parallel rejuvenations.

D. A practical comparison of configurations

So far we have compared several $\langle n, f \rangle$ configurations, two models of attack, two models of rejuvenation and different perspectives of parameter selection. However, in real cases, further practical restrictions may condition the criteria for optimal configuration. We now proceed with an illustrative comparison-example. Consider that:

1) the underlying protocol requires $n = 2f + k + 1$, e.g., a typical synchronous or stateless BFT system with rejuvenation (e.g., [14]);
2) the number of nodes is bounded to $n \leq 4$;
3) the system will be attacked either sequentially with a focused IAE of $\lambda = 3$ per node, or in parallel with a dispersed IAE of $\lambda = 3/(n - k)$ per node.
4) the rate at which nodes can rejuvenate is proportional to the number of available off-line nodes, e.g., new (diversified) software replicas are generated using the computational resources of nodes that are not online.

With these restrictions, we want to answer the following question: *what is the configuration that enables a higher $\mathcal{A}$, for an infinite MT?* In answering this question, for different values of r (the time that a node takes to rejuvenate) we shall compare 5 different scenarios fitting the restrictions.

In order to make a *fair* comparison, for $\|$-rejuvenations we consider the existence of an offline *virtual node* (vk) in $\|$-rejuvenations, helping in the preparation of new replicas. This allows us to satisfy the restriction on the number of nodes: $n+vk = 4$. Additionally, since we would have $r \equiv 0$ for the $\|$-rejuvenation cases (Assumption 3), we consider that for the purposes of this comparison r is the time spent by the virtual node to prepare each other node's rejuvenation.

Among the 5 scenarios, one (the reference) has $f = 0$ and $k = 3$, while the other four have $f = 1$, contemplating the possible combinations of two types of attack ($\|$ and $\therefore$) with two types of rejuvenation ($\|$ and $\therefore$). We proceed with a compact notation to describe the base configurations:

- ***Single node*:** $\langle n, f, k, vk\rangle = \langle 4, 0, 3, 0\rangle$; $\langle rej, \delta, \Delta\rangle = \langle \therefore, r/3, (4/3)r\rangle$; $\lambda = 3$.
- ***$\therefore$-rej.*:** $\langle n, f, k, vk\rangle = \langle 4, 1, 1, 0\rangle$; $\langle rej, \delta, \Delta\rangle = \langle \therefore, r, 4r\rangle$; $\lambda = 1$ for $\|$-attack; $\lambda = 3$ for $\therefore$-attack.
- ***$\|$-rej.*:** $\langle n, f, k, vk\rangle = \langle 3, 1, 0, 1\rangle$; $\langle rej, \delta, \Delta\rangle = \langle \|, 0, 3r\rangle$; $\lambda = 1$ for $\|$-attack; $\lambda = 3$ for $\therefore$-attack.

Note that the *online* characteristics of the *single node* configuration are equivalent to those of a $\|$-rejuvenation scheme with: $\langle n, f, k, vk\rangle = \langle 1, 0, 0, 3\rangle$; $\langle rej, \delta, \Delta\rangle = \langle \|, 0, r/3\rangle$; $\lambda = 3$. For any case we have:

- at any given time, $n - k$ out-of n real nodes are online, k out-of n real nodes are rejuvenating, $k + vk$ nodes contribute to reduce δ;
- the global rejuvenation period is $\Delta = r \times n/(k + vk)$;
- the min time between rejuvenations of different nodes is $\delta = r/k$ for $\therefore$-rejuvenations, or $\delta = 0$ otherwise;
- the IAE exerted in each node under attack is $\lambda = 3/(n - k)$, which means $\sum_{j=1}^{n} \lambda_j(t) \leq 3$.

In Figure 7 we plot the availability of such systems, in function of parameter r (time required to recover each node). This figure shows interesting results:

1. For the same rejuvenation type, a focused $\therefore$-attack ($\lambda = 3$) is more effective than a dispersed $\|$-attack ($\lambda = 1$). This was expected, given that in a $\therefore$-attack an optimal sequence of intrusions is pursued and that in the $\|$-attack the sum of IAE decreases with the number of healthy nodes.

2. For $f = 1$, as r grows, $\therefore$-rejuvenations eventually yield a lower availability than $\|$ ones (see thumbnail in Figure 7). This was expected, as $\therefore$-rejuvenations cannot guarantee a periodic complete recovery. Thus, a fast enough intrusion of nodes may keep the system failed for a long time.

3. As r grows, the system with lowest intrusion tolerance threshold ($f = 0$) but higher rejuvenation rate eventually becomes more available than the alternatives. This means that, if single nodes cannot be rejuvenated quickly enough, then it is better to increase k than f.

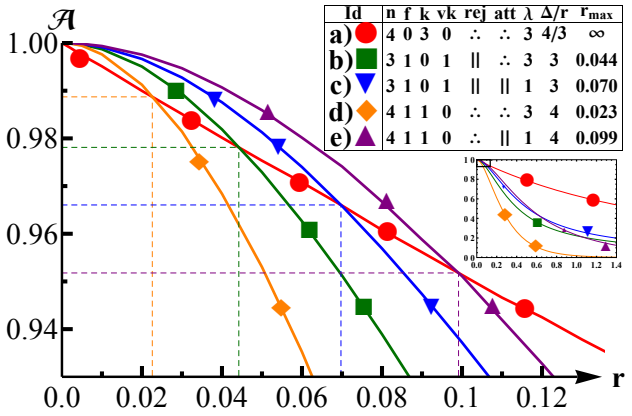

Figure 7. *Availability* ($\mathcal{A}$) of systems in function of *rejuvenation time* (r) per node. Legend: *rej* (*rejuvenation type*); *att* (*attack type*); $\|$ (parallel); $\therefore$ (sequential); r_{max} (maximum r for which $\mathcal{A}$ becomes less than that of the system with $f = 0$). The small rectangular frame on the center-right encloses a zoom-out thumbnail of the main graphic. For each curve, pairs $\langle r, \mathcal{A}\rangle$ were obtained with r spaced in intervals of at most 0.01 (in the main graphic) or 0.05 (in the thumbnail). For $\|$-*rej* cases, $\mathcal{A}$ was obtained from analytic formulas or numeric integration; for $\therefore$-*rej* cases each value $\mathcal{A}$ was obtained by averaging 10 probabilistic simulations of $\mathcal{A}(\tau = \delta \times 10^5)$.

V. Related Work

Intrusion Tolerance. Much research has been done on intrusion-tolerant protocols (e.g., [3], [7], [6], [8]). We do not focus on protocols, but instead on high level properties, such as the functional relation between replication degree n and intrusion tolerance threshold f. One of our main motivations is to show that intrusion tolerance is not necessarily aligned with reliability or availability. Such alignment depends on a set of parameters that must combine together in a way that gives rise to desirable dependability properties.

Reliability. Reliability has been widely studied [4], both in theory and practice. Many works consider detailed estimations of reliability. [15] (one out of many possible examples) studies a particular type of system and analyzes the probability that simultaneous faults actually lead to failure, thus distinguishing fatal from nonfatal faults. We instead follow a high level approach and, focusing on a context of malicious attack, base our estimates on simple and conservative modeling decisions: intrusions cannot be detected and any number of intrusions above the threshold implies immediate failure. Our analysis used some analytic results from [10].

Rejuvenations. The effect of rejuvenation schemes is the topic of previous research works. For example, [16], [14], [2] evaluate tradeoffs between pro-active and reactive recoveries. In a similar way, we compare different models of rejuvenation, but avoid reactive schemes, given our scope of stealth intrusions. The work in [17] mentions the infeasibility of enforcing a threshold of intrusions and considers proactive recovery as a possible mitigation. It also points out caveats

in asynchronous systems that depend on synchronous rejuvenations. In this paper we are not concerned with proving the feasibility of rejuvenations – we just assume their possibility. From there we study the configurations that provide an augmentation to reliability and availability.

Diversity. Much research has been done on the need of *diversity* in systems with rejuvenation (e.g., [18], [19], [20], [13]) and on how to avoid common modes of failure (e.g., [5]). We do not address the problem of node vulnerabilities, but are instead just concerned with the specification of intrusions as the result of direct attack efforts. We are interested in finding configurations that allow the best dependability properties. Nevertheless, we show that degradation is possible even when intrusion independence exists.

VI. Conclusions

In this paper we showed how some (often neglected) parameters play an important role in determining the *reliability* and *availability* of intrusion-tolerant systems. We focused on the impact of *mission time* (MT), rejuvenation strategy and attack model. Based on our analytical and simulation-based study we found four main insights that should be taken into account when designing dependable systems:

1) To assess the benefits of replication and rejuvenation, it is essential to specify the MT, or, more precisely, its relation with the *expected time to intrusion* of individual nodes. Its non-specification allows opportunity for undesired levels of *reliability* and/or *availability*. For example, intrusion-tolerant replication may be counter-productive in the long term if parallel attacks are in place and malicious stealth intrusions are expected. Even a simple distinction between finite, unbounded or infinite MT might help distinguishing configurations in respect to their dependability-enhancement.

2) The choice of rejuvenation type – *sequential* or *parallel* – is important for the overall reliability and availability of the system. For example, *sequential* rejuvenations, incapable of guarantying that the overall system is reset to a complete healthy state, allow a subtle time-window of attack not present in truly periodic *parallel* rejuvenations.

3) Replication and rejuvenation have complementary roles by improving the *reliability* and *availability* of systems for two opposite extremes of a mission timeline. For some configurations, reliability is benefited from the synergy of replication and rejuvenation, even for unbounded (but not infinite) mission times. This benefit can be expressed by the defined measure of *resilience*, which allows a linear expression of goals-of-improvement.

4) By specifying a relation between an effort of attack and the respective intrusion rate of nodes, it is possible to circumvent the problem of not being able to predict the behavior and power of a malicious adversary. In our examples, we considered that an "effort" exerts a proportional probabilistic rate of intrusion.

The study presented in this paper is a step in understanding how to use *intrusion tolerance* techniques to provide *tolerance to uncertainty of assumptions*, making it possible to design dependable systems that better withstand any instantiation of some hidden or unspecified parameters.

Acknowledgment

We thank the LADC'11 reviewers for their comments that helped us improve the paper. This research was partially supported by FCT (Fundação para a Ciência e a Tecnologia – Portuguese Foundation for Science and Technology) through the Carnegie Mellon Portugal Program under Grant SFRH/BD/33770/2009 (for the first author) and by FCT through project PTDC/EIA-IA/100581/2008 (REGENESYS) and the Multiannual (LaSIGE) program.

References

[1] F. B. Schneider, "Implementing fault-tolerant service using the state machine aproach: A tutorial," *ACM Computing Surveys*, vol. 22, no. 4, Dec 1990.

[2] Y. Huang, C. M. R. Kintala, N. Kolettis, and N. D. Fulton, "Software rejuvenation: Analysis, module and applications," in *Proceedings of 25th International Symposyum on Fault Tolerant Computing – FTCS-25*, Jun 1995.

[3] P. E. Veríssimo, N. F. Neves, and M. P. Correia, "Intrusion-tolerant architectures: Concepts and design," in *Architecting Dependable Systems*, ser. LNCS. Springer, 2003, vol. 2677.

[4] R. E. Barlow, *Mathematical Reliability Theory: From the Beginning to the Present Time*. World Scientific, Singapore, 2002, vol. 7.

[5] R. R. Obelheiro, A. N. Bessani, L. C. Lung, and M. Correia, "How practical are intrusion-tolerant distributed systems?" Dep. of Informatics, Univ. of Lisbon, DI-FCUL TR 06–15, 2006.

[6] M. Correia, N. F. Neves, and P. Veríssimo, "How to tolerate half less one Byzantine nodes in practical distributed systems," in *Proceedings of the 23rd IEEE Symposium on Reliable Distributed Systems - SRDS 2004*, Oct 2004.

[7] M. Castro and B. Liskov, "Practical byzantine fault tolerance and proactive recovery," *ACM Transactions on Computer Systems*, vol. 20, no. 4, Nov 2002.

[8] G. S. Veronese, M. Correia, A. N. Bessani, and L. C. Lung, "Spin one's wheels? Byzantine fault tolerance with a spinning primary," in *Proceedings of the 28th IEEE Symposium on Reliable Distributed Systems – SRDS'09*, Sep 2009.

[9] A. N. Sovarel, D. Evans, and N. Paul, "Where's the FEEB? the effectiveness of instruction set randomization," in *Proceedings of the 14th USENIX Security Symposium*, vol. 14, Aug 2005.

[10] K. S. Trivedi, *Probability and statistics with reliability, queuing and computer science applications*, 2nd ed. Chichester, UK: John Wiley and Sons Ltd., 2002.

[11] E. Barker, W. Barker, W. Burr, W. Polk, and M. Smid, "Recommendation for key management - part 1: General (revised)," NIST Special Publication 1/3, Mar 2007.

[12] P. Sousa, N. F. Neves, and P. Veríssimo, "How resilient are distributed f fault/intrusion-tolerant systems," in *Proceedings of the 2005 International Conference on Dependable Systems and Networks – DSN'05*, 2005.

[13] T. Roeder and F. B. Schneider, "Proactive obfuscation," *ACM Transactions on Computer Systems*, vol. 28, no. 2, 2010.

[14] P. Sousa, A. N. Bessani, M. Correia, N. F. Neves, and P. Verissimo, "Highly available intrusion-tolerant services with proactive-reactive recovery," *IEEE Transactions on Parallel and Distributed Systems*, vol. 21, no. 4, 2010.

[15] I. Koren and E. Shalev, "Reliability analysis of hybrid redundancy systems," *IEE Proceedings-E on Computers and Digital Techniques*, vol. 131, Jan 1984.

[16] A. Daidone, S. Chiaradonna, A. Bondavalli, and P. Veríssimo, "Analysis of a redundant architecture for critical infrastructure protection," in *Architecting Dependable Systems V*, ser. LNCS, vol. 5135. Springer, 2008.

[17] P. Sousa, N. F. Neves, and P. Verissimo, "Hidden problems of asynchronous proactive recovery," in *Proceedings of the 3rd workshop on on Hot Topics in System Dependability – HotDep'07*, 2007.

[18] S. Forrest, A. Somayaji, and D. H. Ackley, "Building diverse computer systems," in *Proceedings of the 6th Workshop on Hot Topics in Operating Systems – HotOS'97*, 1997.

[19] B. Littlewood and L. Strigini, "Redundancy and diversity in security," in *Proceedings of the 9th European Symposium on Research Computer Security – ESORICS'04*, ser. LNCS, vol. 3193. Springer, 2004.

[20] A. Bessani, A. Daidone, I. Gashi, R. Obelheiro, P. Sousa, and V. Stankovic, "Enhancing fault/intrusion tolerance through design and configuration diversity," in *Proceedings of the 3rd Workshop on Recent Advances on Intrusion-Tolerant Systems – WRAITS'09*, Jun 2009.

[21] F. W. Olver, D. W. Lozier, R. F. Boisvert, and C. W. Clark, *NIST Handbook of Mathematical Functions*. New York, NY, USA: Cambridge University Press, 2010.

[22] I. Wolfram Research, "Mathematica 7.0 for students," Champaign, Illinois, Feb 2009.

APPENDIX

From Assumption 1 and 2, we consider the case of a constant IAE and a proportionality between IAE and IRP. Thus, IRP is a constant λ and the intrusion of a node is modeled probabilistically with PDF $p_{1,0}^{(\lambda)}(t) = \lambda \times e^{-\lambda \times t}$, CDF $P_{1,0}^{(\lambda)}(t) = 1 - e^{-\lambda \times t}$ and ETTF $\mu_{1,0}^{(\lambda)} = 1/\lambda$. When possible we will omit λ. The overall *reliability* of $\langle n, f \rangle$ is $\mathcal{R}_{n,f}(t) = 1 - P_{n,f}(t)$, with $P_{n,f}(t)$ being the probability of global failure at instant t. When necessary, we shall distinguish the type of attack using superscripts ($\|$ or $\therefore$). When considering rejuvenations, we shall include subscripts with the respective symbols ($\|$ and Δ or $\therefore$ and δ).

Reliability ($\mathcal{R}$) with *parallel* ($\|$) attacks. The CDF of failure is in Equation 3. Calculating the sum, and subtracting it from 1, *reliability* becomes as in Equation 4, with ${}_2F_1$ being the *Hypergeometric2F1* function.

$$P_{n,f}^{\|}(t) = \sum_{i=f+1}^{n} \binom{n}{i} P_{1,0}(t)^i \times (1 - P_{1,0}(t))^{(n-i)} \quad (3)$$

$$\begin{aligned}\mathcal{R}_{n,f}^{\|}(t) = 1 - \left(e^{-\lambda t}\right)^{n-(f+1)} \left(1 - e^{-t\lambda}\right)^{f+1} \\ \times \binom{n}{f+1} \times_2 F_1\left(1, f+1-n; f+2; 1-e^{t\lambda}\right)\end{aligned} \quad (4)$$

Reliability ($\mathcal{R}$) with *sequential* ($\therefore$) attacks. The probability density $p_{n,f}^{\therefore}(t)$ that the $(f+1)$-th node is intruded exactly at instant t, is in Equation 5, with $p_0^{\therefore}(t) \equiv p_{1,0}(t)$. The global probability of failure $P_{n,f}^{\therefore}(t)$ is in Equation 6. The respective *reliability* is in Equation 7, with Q being the *Generalized Incomplete Regularized Gamma Function* [21], satisfying $Q(a, z_0, z_1) = \Gamma(a, s)/\Gamma(a)$ and $\Gamma(a, z_0, z_1) = \int_{t=z_0}^{z_1} t^{a-1} e^{-t} dt$.

$$p_{n,f}^{\therefore}(t) = \int_{t'=0}^{t} p_{n,f-1}^{\therefore}(t')\, p_{1,0}(t - t')\, dt' = \frac{(\lambda t)^f}{f!} \lambda e^{-\lambda t} \quad (5)$$

$$P_{n,f}^{\therefore}(t) = \int_{t'=0}^{t} p_{n,f}^{\therefore}(t')\, dt' \quad (6)$$

$$\mathcal{R}_{n,f}^{\therefore}(t) = 1 - \mathcal{P}_{n,f}^{\therefore}(t) = Q(f+1, \lambda t, \infty) \quad (7)$$

Availability ($\mathcal{A}$). $\mathcal{A}$ is the probability that the system is healthy at a random (uniformly selected) instant of time within the MT, as in Equation 8. Due to space constraints we do not expand the result of such integral for the two attack-models considered.

$$\mathcal{A}_{n,f}(t) = \frac{1}{t} \int_{t'=0}^{t} \mathcal{R}_{n,f}(t')\, dt' \quad (8)$$

MT for the same $\mathcal{R}$. For *sequential* attacks, solving $\mathcal{R}_{1,0}^{\therefore}(t) = \mathcal{R}_{n,f}^{\therefore}(t')$ in order of t' gives Equation 9, with $Q^{\langle 0,0,-1 \rangle}$ being the (3rd argument) inverse of $Q(a, z_0, z_1)$.

$$t' = Q^{\langle 0,0,-1 \rangle}\left(f+1, \infty, e^{-\lambda t}\right)/\lambda \quad (9)$$

***Parallel* ($\|$) rejuvenations.** Let $M = \lfloor \tau/\Delta \rfloor$ and $m = mod_{\Delta}\tau$. With $\|$-rejuvenations, the system is periodically restored to a completely healthy state. Thus, *reliability* (Equation 10) and *availability* (Equation 11) can be obtained in function of the formulas without rejuvenations, by partitioning the MT into windows of size Δ.

$$\mathcal{R}_{n,f,\|,\Delta}(\tau) = \mathcal{R}_{n,f}(\Delta)^M \mathcal{R}_{n,f}(m) \overset{\tau >> \Delta}{\approx} \mathcal{R}_{n,f}(\Delta)^{(\tau/\Delta)} \quad (10)$$

$$\begin{aligned}\mathcal{A}_{n,f,\|,\Delta}(\tau) = (1 - m/\tau) \times \mathcal{A}_{n,f}(\Delta) + \\ (m/\tau) \times \mathcal{A}_{n,f}(m) \overset{\tau >> \Delta}{\approx} \mathcal{A}_{n,f}(\Delta)\end{aligned} \quad (11)$$

Note: Software [22] was used to perform the simulations needed for Figure 6 and 7, plot all the graphics and tables and help deducing Equations 4, 5, 7 and 9.

