# REF: Resilience Evaluation Framework

Bruno Sousa
CISUC, University of Coimbra
Polo II, Pinhal de Marrocos
3030-290, Coimbra, Portugal
Email: bmsousa@dei.uc.pt

Kostas Pentikousis
Huawei Technologies
European Research Center
Carnotstrasse 4, 10587 Berlin, Germany
Email: k.pentikousis@huawei.com

Marilia Curado
CISUC, University of Coimbra
Polo II, Pinhal de Marrocos
3030-290, Coimbra, Portugal
Email: marilia@dei.uc.pt

***Abstract*—Resilience is one of the key goals of multihoming, as multiple interfaces/addresses on multiaccess nodes can be used to increase fault tolerance. Future networks will be empowered with multiaccess but, at the same time, applications and protocols must incorporate mechanisms to optimize the use of numerous interfaces/addresses. Until now, the resilience of multihoming support in various protocols has been assessed in an ad-hoc manner. This paper introduces the Resilience Evaluation Framework (REF), which provides for an objective evaluation of the resilience capacity of a protocol. As an example evaluation, we use REF to study the resilience of SCTP multihoming. We employ OMNET++ simulations to demonstrate the suitability of REF for such evaluations and compare it with the Quality of Resilience (QoR) schema. We show that REF is suitable for general protocol resilience evaluations and defines a more realistic evaluation framework than QoR.**

**Keywords – Multihoming, Quality of Resilience, Computer network management, reliability, performance and SCTP.**

## I. Introduction

New communication possibilities can be explored as network nodes become multiaccess capable. Resilience can be achieved if, besides adding multiple access technologies in each device, applications and protocols are equipped with efficient mechanisms to deal with the diversity of interfaces/addresses. For instance, the efficient mechanisms can go beyond the traditional primary-backup model and use all the paths simultaneously.

So far, the evaluation of the resilience support for a given protocol does not have a standard mechanism which enables the choice of a protocol for future networks based on the effective resilience support. The Resilience Evaluation Framework (REF) is proposed in this paper allows for the comparative study of the resilience capacity of different protocols. Commonly, researchers define specific metrics to perform the evaluation of protocols/technologies (see Section II). The problem with this type of approach is that metrics and protocol evaluation methodologies are closely tied to the research problem. Thus, they are not extensible to generic scenarios. REF addresses this limitation by defining objective measures of resilience that can be compared on an equal footing between different protocols.

We show the applicability of REF in this paper through a case study of the Stream Control Transport Protocol (SCTP) [1]. SCTP includes failure detection of the primary path and recovery switching to backup paths. In addition, the 1:1 protection model of SCTP can be enhanced to a 1+1 protection model with the Concurrent Multipath Transfer (CMT) extension [2]. Further, we compare REF with the Quality of Resilience (QoR) proposal [3] (originally tailored for MPLS networks) since both include availability and recovery efficiency metrics.

The remainder of this paper is organized as follows: Section II identifies the most relevant related work within the SCTP and resilience evaluation areas. Section III introduces our Resilience Evaluation Framework (REF) and Section IV presents our evaluation methodology and comparative simulation results. Finally, Section V concludes the paper.

## II. Related work

This section relates REF, on the one hand, with previous work on evaluating resilience in general, and, on the other, with SCTP-specific resilience evaluation studies.

### A. Resilience

Resilience has been evaluated in various ways and for different protocols. Both the Resilience-Differentiated Quality of Service (RD-QoS) framework [4] and Quality of Resilience (QoR) [3], [5] assess the resilience support of Multi Protocol Label Switching (MPLS). However RD-QoS does not assess the recovery cost and only includes a time analysis based on the ITU-T M.495 model [6]. QoR combines QoS metrics (e.g. packet loss, delay) with resilience metrics (e.g. steady-state availability, mean downtime). Nonetheless, the metrics rely on several parameters that affect the results and the formulation is tied to MPLS, lacking a broader applicability.

Recovery efficiency, as well as the protection model supported (e.g., 1+1 or 1:N) is addressed in [7]. Nevertheless, the evaluation relies on non deterministic methods, which depend on the application requirements. Other proposals evaluate resilience solely based on the availability criteria [8], [9], or are limited to the exploration of the protection mechanisms [10]. In short, previous work does not provide a complete scheme to assess resilience, taking into consideration availability, as well as protection and recovery factors. Within this perspective, the QoR framework is the most complete, although it focuses on MPLS. QoR also employs histograms that can be mapped to

978-1-4244-7286-4/10/$26.00 

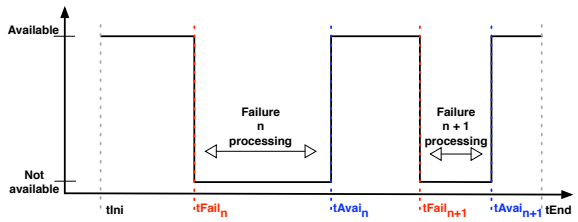


Fig. 1: Availability Model according states of ITU-T Rec. E.800 [15].

user satisfaction, nevertheless the effective resilience support of a protocol is the aim of REF.

### B. SCTP evaluation

Previous SCTP evaluation work on the evaluation of SCTP, such as [11] and [12] determine the best configuration for the failover mechanism of SCTP to transport different types of data. Thus, the evaluation methodology can not be employed as a generic tool to assess resilience.

Some SCTP evaluation proposals are tied to specific applications, as the case of ECHO [13], that is suited for VoIP applications over SCTP, while others concentrate efforts on the robustness of SCTP, as in [14] to avoid spurious failures.

The analysis of related work has clearly shown the lack of a framework to evaluate the resilience capabilities of SCTP in an objective and application/protocol-independent way.

## III. REF - Resilience Evaluation Framework

The Resilience Evaluation Framework (REF) aims at evaluating protocol resilience support in an objective way, without relying on application requirements. REF relies on the ITU-T M.495 [6] model to determine recovery performance, and on the ITU-T E.800 availability model [15] to determine availability. The final resilience assessment is done by following Def. 1

*Definition 1:* **- Resilience** is a mechanism to assure service robustness, by ensuring that resources are re-established in case of failures [7]. This re-establishment is possible due to protection (actions before failure) and/or restoration schemes (actions after failure).

REF introduces the term *ipath* to designate an interface, a path or a link, in order to be as generic as possible. The mathematical formulation of REF is summarized in Table I for different cases. The column titled *Base* presents the generic specification, while columns titled *1:1* and *1+1* correspond to the respective protection models.

REF considers the following assumptions in its formulation:

- All *time* variables are expressed in *milliseconds* (ms);
- All *message size* variables are expressed in *bytes*;
- The *capacity* of *ipath* is in *byte/s* and is constant;
- Percentage calculated values rely on the range of $[0, 1]$;
- Failures have a *min* of $\{0\}$ and a *max* of $\{n\}$;
- A node has a *min* of $\{2\}$ and a *max* of $\{z\}$ *ipaths*;
- A node has *min* of $\{1\}$ and *max* of $\{bk\}$ backup *ipaths*;

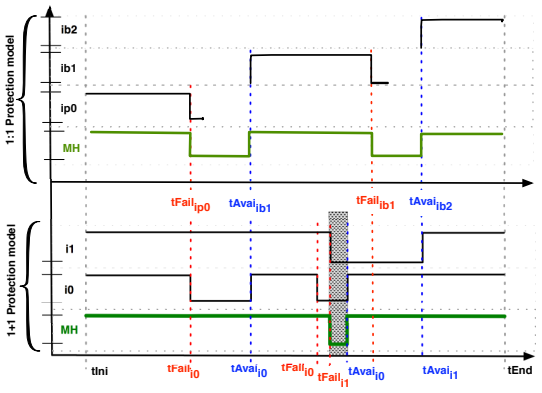


Fig. 2: Availability with **1:1** and **1+1** Protection Models.

- Cost functions rely on $\beta_O$ and $\beta_R$ that are empirically determined;

$R_{MH}$- Resilience is a function of $Av$- Availability, $Rc$- Recovery, which are determined based on the protection scheme supported, as shown in Eq. 1.

$$R_{MH} = Av \times Rc \tag{1}$$

Fig. 1 depicts an availability model on which the service has two states, available $(state = 1)$ or unavailable $(state = 0)$, according to ITU-T E.800 . After a failure instant, $tFail_n$, the procedures for failure processing are undertaken (see [7]). Whilst failure processing mechanisms can be handled at different layers, REF only considers the processes at the layer of the evaluated protocol.

### A. Availability ($Av$)

Availability $Av$ is the ratio of the Mean Up Time (MUT) over the total time (MUT+MDT); see also [5], [7]. Fig. 1 illustrates a scenario with two failures, on which MUT corresponds to the moments where $available = 1$ and can be formulated for generic cases with $n$ failures according to Eq. T.1 On a non failure situation, $MUT = tEnd - tIni$, since $tFail_n = tAvai_n = 0$. The Mean Down Time (MDT) considers the $available = 0$ moments, therefore for $n$ failures, it is determined according to Eq. T.3.

A key aspect in REF is the evaluation of availability in the context of end-host multihoming, which can have multiple interfaces. To the best of our knowledge, previous availability approaches only take into account the availability of the overall service [5] or of interfaces/paths (in REF designated as *ipath*) in isolation [8], [9], without taking into consideration the role of each *ipath*. The role of an *ipath* dictates whether it acts as a primary *ipath* or a backup *ipath*. This role is associated with the protection model, with the *1:1*, *1+1* and *1:N* models being the most generic from a multihoming perspective [7]. The *1:1* model states that the backup *ipaths* are only employed in the failure of the primary *ipath*, while the *1+1* allows the simultaneous use of primary and backup *ipaths*, as illustrated in Fig. 2. In the *1:N* model *N* backup

TABLE I: REF Mathematical Formulation

| | Base | 1:1 Case | 1+1 Case |
|---|---|---|---|
| **MUT** | $(tFail_1 - tIni) + \sum_{i=2}^{n}(tFail_i - tAvai_{i-1}) + (tEnd - tAvai_n)$ **T.1** | $(tFail_0 - tIni) + \sum_{t=1}^{bk-1}(tFail_t - tAvai_t) + (tEnd - tAvai_{bk})$ **T.2** | |
| **MDT** | $MDT = \sum_{i=1}^{n}(tAvai_i - tFail_i)$ **T.3** | | $MDT_{1+1} = MDTit_{1+1}$ with, $it_{1+1} = m\{tAvai_0, tAvai_1\} - M\{tFail_0, tFail_1\}$ **T.4** |
| **ERc** | $ERc_t = \left(1 + \frac{\sum_{i=1}^{n} tRc_{t,i}}{MDT_t}\right)^{-1}$ **T.5** | $\left(1 + \frac{tRc_0 + \sum_{t=1}^{bk}(tRc_t)}{MDT_0 + \sum_{t=1}^{bk} MDT_t}\right)^{-1}$ **T.6** | $\left(1 + \frac{\sum_{i=1}^{n} tRc_{i_{1+1}}}{\sum_{i=1}^{n} MDTit_i}\right)^{-1}$ **T.7** |
| **LRc** | $LRc = \frac{C_c \sum_{i=1}^{n}(teRS_{c,i} - tsFD_{c,i})}{C_p(tEnd - tStart)}$ **T.8** | $\frac{C_0 \cdot tRc_0 + \sum_{t=1}^{bk}(C_t \cdot tRc_t)}{C_0(tEnd - tStart)}$ **T.9** | $\frac{simRc_{1+1}}{\sum_{j=1}^{z} parLRc_j}$ **T.10** $= \frac{\frac{(C_0+C_1)\sum_{i=1}^{n}\left(tRc_{i_{1+1}}\right)}{(C_0+C_1)(tEnd-tStart)}}{\sum_{j=1}^{z}\cdot\frac{(C_c)\sum_{j=1}^{n_c}(tRS_j - tFD_j)}{C_c\ (tEnd-tStart)}}$ |
| **allSig** | $allSig = nM \cdot \overline{MSi} + nT \cdot \overline{TDu}$ **T.11** | | |
| **sigRc** | | $nM_0 \cdot \overline{MSi_0} + nT_0 \cdot \overline{TDu_0} + \sum_{t=1}^{bk}\left(nM_t \cdot \overline{MSi_t} + nT_t \cdot \overline{TDu_t}\right)$ **T.12** | $\sum_{j=1}^{n_s}\left(nM_{i0(j)} \cdot \overline{MSi_{i0(j)}} + nT_{i0(i)} \cdot \overline{TDu_{i0(j)}} + nM_{i1(j)} \cdot \overline{MSi_{i1(j)}} nT_{i1(j)} \cdot \overline{TDu_{i1(j)}}\right)$ **T.13** |
| **sigRt** | $sigRt_{(t)} = \frac{\sum_{i=1}^{n} sigRc_{t,i}}{allSig_t}$ **T.14** | | |
| **Xb** | $Xb_c = min\{{}^{C_c}/_{C_p}, {}^{C_p}/_{C_c}\}$ **T.15** | | $min\left\{\frac{C_{befFailure}}{C_{afterFailure}}, \frac{C_{afterFailure}}{C_{befFailure}}\right\}$ **T.16** |
| **Resto** | $Resto = \frac{nSuccessFulRecovers}{nFailures}$ **T.17** | | $\frac{nSuccessRec_{if0} + nSuccessRec_{if1}}{nFailure_{i0} + nFailure_{i1}}$ **T.18** |

*ipaths* protect one primary ipath. Cholda et al. [5] refer to the *M:N* model as a generic model, on which *N* backup *ipaths* protect *M* primary *ipaths*. In this paper we consider REF for the the *1:1* and *1+1* models only.

*1) 1:1 Protection Model:*
According to Fig. 2, the determination of $MUT_{1:1}$ and $MDT_{1:1}$ can be based on Eq. T.2 and T.3. $MUT_{1:1}$ considers the availability of all *ipaths* in a sequential mode, starting with the primary and following the respective backup *ipaths*.

*2) 1+1 Protection Model:*
In the *1+1* model, *i0* - the primary *ipath* is used simultaneously with *i1* - the backup ipath, as illustrated in Fig. 2. MUT corresponds to the union of MUTs for each *ipath*. An *OR* boolean logic can be employed to determine MUT, by including all the moments on which, at least one *ipaths* is available. The probability of downtime is lower, as if one *ipath* fails, another assures the service delivery. Thus, the downtime is the intersection of MDT on both interfaces, which is calculated based on the difference between the minimum $m$ available time and the maximum $M$ failure time, as given in Eq. T.4

### B. Recovery (Rc)

Recovery encompasses the actions necessary to return to a normal state. For such, different processes may occur within a recovery scheme, namely, Failure Detection (FD), Failure Notification (FN), determination of new paths, Recovery Switching (RS) and, finally, Restoration to the initial service levels [5].

According to Fig. 3, the recovery time $tRc$ is determined as follows: $tRc = \sum_{i=1}^{5} T_i$. In a simplistic approach, the recovery time can be determined based on the end time of recovery ($teRS$) and the start time of Failure Detection ($tsFD$), $tRc = teRS - tsFD$. Summing the time of the different processes, the recovery time is $tRc = tFD + tFN + tRS$.

To assess the performance of recovery schemes, different metrics/factors must be evaluated: $ERc$ - recovery time efficiency; $LRc$ - the recovery impact, which can correspond to the traffic that is affected by recovery schemes; $ORc$ - the recovery overhead, i.e. the cost of recovery in terms of signalling; and $QRc$ - the quality provided by recovery, which measures whether operation returns to the same conditions as before the failure occurred. $QRc$ is determined by restorability and backup link quality [5], [17]. Restorability indicates the percentage of failed *ipaths* that can be recovered [18]. As with availability, the operation of the recovery scheme depends on

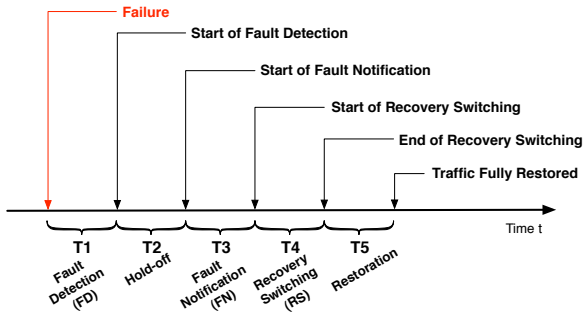


Fig. 3: Recovery model of ITU-T M.495 with slight adaptations from [16].

the protection and/or restoration model adopted. While Protection models have one of more backup *ipaths* pre-established before failures, Restoration models establish the backup *ipath* on failure events. The former can attain better performance in terms of recovery time (no need of signalling to establish an *ipath*), but with the drawback of having a higher cost in terms of resources, since an *ipath* is dedicated for recovery [4], [5].

The recovery efficiency $ERc$ of *ipath* $t$ corresponds to the ratio between the time to recover from all failures $\{i \cdots n\}$ and the mean downtime, as given in Eq. T.5. REF introduces this metric to differentiate protocols that have optimized mechanisms to provide fast recovery.

The recovery impact $LRc$ is determined based on the affected traffic (considered lost or prone to restransmission) during the recovery process, as depicted in Eq. T.8. REF considers the relation between the capacity of the current *ipath*, $C_c$ and the capacity of the primary path $C_p$, as opposed to [5] which considers only the capacity of the primary path. By considering the $C_c/C_p$ ratio it is possible to determine the affected traffic based on the recovery time but also on the capacities of *ipath*.

The recovery overhead $ORc$ represents the signalling cost of the recovery operations, see Eq. 2. This compound metric establishes the difference between the recovery models, as protection models have backup links pre-established, while restoration models need to establish them based on signalling. $ORc$ expresses a cost function, where the interest is, on one hand to minimize signalling ratio (e.g., signalling distributed along the service lifetime), and on the other hand to maximize signalling ratio diversity (e.g., improve load sharing).

$$ORc = \beta_O sigRt + (1 - \beta_O) sigDv \tag{2}$$

In REF, the determination of signalling considers the approaches that are based on message signalling and approaches that employ timers to trigger recovery actions. REF considers average values for $MSi$ - message size and $TDu$ - timeout durations, as these can vary in the measure interval. The recovery overhead is determined with $allSig$ - the overall signalling, $sigRt$ - the signalling ratio and $sigDv$ - the signalling ratio diversity metrics. $allSig$ includes all the signalling in the interval $(tEnd - tStart)$ and is calculated according to Eq. T.11. Overall signalling includes the $MSi$ with $nM$- total number of messages transmitted, and $nT$- number of timeouts, with different durations $TDu$. The signalling performed during recovery ($sigRc_i$) from failure $i$ is calculated by the signalling overhead during $tRc_i$ - the time of recovery, employing the same logic for the overall signalling, but only for the instant $tRc_i$. The signalling ratio, $sigRt_{(t)}$, for *ipath* $t$, establishes the relation between the signalling during recovery from possible $n$ failures and the overall signalling, as depicted in Eq. T.14. Higher values indicate that the signalling overhead is concentrated in the recovery processes. The signalling ratio diversity ,$sigDv$, assesses how signalling is balanced among all *ipaths* and can be calculated based on the relation of the minimum ratio and the maximum ratio of all *ipaths*, $sigDv = min\{sigRt_{(t)}\}/max\{sigRt_{(t)}\}$. If $sigDv \rightarrow 1$ the signalling load is distributed more equally between the *ipaths*. The recovery overhead metric in REF is clearly distinct from previous proposals [5], [17], which do not consider the signalling overhead, or consider it in a simplistic manner without any diversity analysis.

The quality provided by recovery, $QRc$, is a relation between $Xb$ - the quality of backup *ipaths* and $Resto$ - the restorability. The $Xb$ factor is determined based on $C_c$ and $C_p$ the capacity of current and primary *ipath*, respectively, as given in Eq. T.15, following the same calculation as in [5]. REF adds $Resto$ - restorability, which accounts for the ratio of failed connections successfully recovered, as depicted in Eq. T.17. Restorability in REF allows to assess to what extent the recovery is performed. $QRc$ is based on the average quality of backup *ipaths*, as each *ipath* has its own quality factor and on the restorability ratio, $QRc = \overline{Xb} \times Resto$.

Recovery, $Rc$, is determined according to Eq. 3, where the interest is, on one hand, to minimize the affected traffic and the overhead of recovery procedures, and on the other, to maximize the quality provided by recovery and the recovery efficiency.

$$Rc = \beta_R(LRc \times ORc) + (1 - \beta_R)(ERc \times QRc) \tag{3}$$

This concludes the brief introduction to the general formulation of REF. The following subsections present the specification of REF according to the protection model supported. REF considers the protection model for the recovery performance assessment, a clear advance from previous work such as [5], [17].

*1) 1:1 Protection model:*
Recovery efficiency, $ERc_{1:1}$, considers the time of recovery of all *ipaths* (primary and backups) and their respective MDT, see Eq. T.6. Recovery impact, $LRc_{1:1}$, assesses the affected traffic during recovery of the primary *ipath* and the recovery of the respective backup *ipaths*, in relation to the theoretical traffic that could be transmitted, if no failures had occurred, during the time service, as illustrated in Eq. T.9. Finally, recovery overhead, $ORc_{1:1}$, is determined according to Eq. 2. Nevertheless, the signalling during recovery is determined for

the instant of recovery and includes the messages or timers that are associated with the primary *ipath* and backup *ipaths*, as given in Eq. T.12.

*2) 1+1 Protection model:*

In the *1+1* protection model, the recovery processes only occur when both interfaces are down simultaneously, as shown in Fig. 2. The time for *ipath* to recover from failure $i$ ($tRc_{i_{1+1}}$) depends on the minimum time of recovery and on the maximum time of failure detection from one of the *ipaths* {i0,i1}, following the logic depicted in Eq. T.4.

The recovery efficiency, $ERc_{1+1}$, is determined based on the recovery time from $n$ possible failures and on the respective simultaneous downtime as per Eq. T.7.

The affected traffic, $LRc_{1+1}$, must consider the affected traffic in two distinct cases: a) *Simultaneous*, ${}_{sim}LRc_{1+1}$ - on which there is no service since both *ipaths* are down simultaneously; b) *Partial* - on which failures only affect one *ipath*. The affected traffic, $LRc_{1+1}$, corresponds to the relation between the affected traffic in the simultaneous case and the sum of the affected traffic in the partial cases, see Eq. T.10. Within $z$ *ipaths*, the simultaneous cases assume that traffic is forwarded simultaneously on different *ipaths*, thus the capacity is considered the sum of all affected *ipaths*. In the partial cases, the traffic impact is considered in isolation for each failed *ipath*.

Although, partial failures can affect the traffic, the service is only disrupted on simultaneous failures, therefore the signalling performed during recovery, in the $ORc_{1+1}$-recovery overhead determination, only considers the recovery performed for $n_s$ simultaneous failures, as depicted in Eq. T.13. The overall signalling includes all the signalling performed during the time service in all *ipaths*.

In the *1+1* cases, the primary and backup *ipath* are used simultaneously, therefore there is no real notion for backup *ipath*. In this context, the backup link quality corresponds to the minimum quality level that is achieved on a failure event, as Eq. T.16 shows. The restorability $Resto_{1+1}$ considers the number of successful recovery performed for each *ipath* in relation to the number of failures in the respective *ipath*, as given in Eq. T.18.

The quality provided by recovery, $QRc_{1+1}$, in the *1+1* protection model is calculated based on the possible $n$ failures and on the *ipath* where failures occur.

$$QRc_{1+1} = \begin{cases} 1 \cdot Resto_{1+1} & \text{if n = 0,} \\ Xb_{i0} \cdot Resto_{1+1} & \text{if } C_0 > C_1 \text{ and } n_{i0} \geq 1, \\ Xb_{i1} \cdot Resto_{1+1} & \text{if } C_1 \geq C_0 \text{ and } n_{i1} \geq 1. \end{cases} \quad (4)$$

REF can be employed to assess the resilience of any given protocol as long as the the protection model ($1:1$ or $1+1$) is taken into consideration. For instance, SCTP is under the $1:1$, while MPLS can be under the $1+1$, if considering load balancing characteristics.

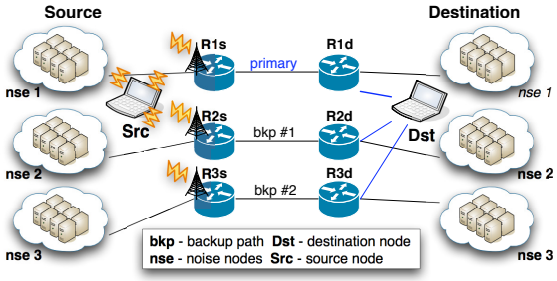


Fig. 4: Simulation Scenario.

| Parameter | RFC4960 (*Std*) | Optimized (*Opt*) |
|---|---|---|
| PMR | 5 | 3 |
| RTOmin | 1000 (ms) | 20 (ms) |
| RTOmax | 60000 (ms) | 60000 (ms) |
| SACK delay | 200 (ms) | 20 (ms) |

TABLE II: SCTP failover parameters.

## IV. Evaluation and Results

This section presents our comparative evaluation of SCTP resilience using REF and QoR. Fig. 4 illustrates the simulation scenario, which includes multihomed nodes with a primary *ipath* and two backup *ipaths*. The scenario includes nodes that introduce "background" traffic that causes congestion (based on bursts) in the respective *ipaths*. We use this scenario to evaluate SCTP resilience in the presence of failures occurring in the primary and backup *ipaths* and consider different types of data traffic. Different networks are considered on the source and destination sides. Moreover the source node moves linearly and starts connected to all wireless access routers (in order to include all the configured addresses during the association phase of SCTP). The scenario is modeled in OMNET++ simulator [19] using the SCTP extension [20].

The source node performs two handovers, from primary to bkp #01 and from bkp #01 to bkp #02 *ipaths*, respectively. In addition, we take $\beta_O = \beta_R = 0.05$, an empirical value based on experience with the simulation scenario.

The evaluation considers both SCTP failover parameters and the application in use (both include the sets for the data and VoIP applications). The SCTP failover parameters are configured according to RFC 4960 [1] while the optimized configurations come from [11], [12], as listed in Table II. Other configurable parameters of SCTP, such as Association Max Retrans and RTOinit follow the values recommended in RFC 4960.

We consider both VoIP and data applications for our evaluation. VoIP traffic is based on the G.723.1 [21] codec employing a bit rate of 6.3kbps. In addition, we configure SCTP to deliver all DATA chunks received immediately to the upper layer (i.e. unordered). FTP is chosen for our data application. We consider these types of applications due to the diversity of their requirements. Each test comprises 20 runs and a simulation

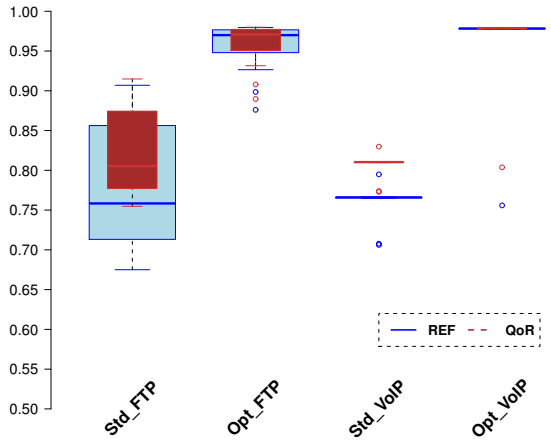


Fig. 5: SCTP availability, as measured by REF ($Av$) and QoR ($Q_A$).

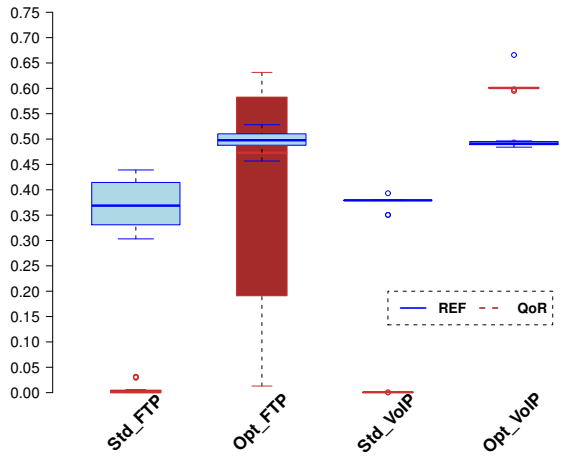


Fig. 6: SCTP resilience as measured by REF and QoR.

time of 300s.

We report the simulation results using boxplots, due to the richer information they present in terms of statistics (median, quartiles, maximum and minimum).

### A. Availability Results

The $Av$- Availability parameter of REF can be compared with $QA$- Quality of Availability of QoR. REF is similarly to QoR, as it determines an higher degree of availability of SCTP for optimized cases (with decreased SACK interval), as shown in Fig 5. Moreover, REF and QoR point out the fact that the availability in SCTP relies solely on its failover parameters and not on the sets of applications. It should be underlined that, differently from QoR, availability in REF relies on the protection model and it is independent of theoretical factors.

### B. Resilience

Fig. 6 illustrates the resilience of SCTP as measured by REF and QoR. In contrast with the results on availability, we see a divergence between REF- and QoR-measured resilience values. QoR reports that SCTP with standard configuration as per RFC 4960 has virtually no resilience. REF, due to the protection model, reports a resilience value of $0.37$ for both data and VoIP applications. Both QoR and REF show that SCTP resilience improves with the optimizations proposed in [11], [12]. In particular, REF reports a resilience value of nearly $0.50$, on median. QoR, on the other hand, reports a median resilience value of $0.47$ for FTP traffic and $0.6$ for VoIP traffic. We also note that the spread of the QoR resilience values is considerably larger than for REF. To sum up, the QoR SCTP resilience values for the standard configuration point to the schemas lack of generality. REF, on the other hand, provides a more realistic resilience evaluation framework by taking into consideration that the *1:1* protection model of SCTP.

## V. Conclusion and Next Steps

This paper introduced REF, our proposal for generic resilience evaluation framework. We compared REF with QoR and showed that its evaluation results for the case of SCTP are more realistic. In this study we used REF to assess the resilience capacity of SCTP. Although further evaluation work is in progress, we claim that REF can assess objectively and in an application independent manner the resilience of other protocols as well. Other protocols can easily be assessed without any modification to REF due to its generality. Such assessments require only that the respective protection model supported is taken into consideration. We plan to use REF to evaluate the resilience of a range of multihoming solutions in the near term.

REF allows us to argue that SCTP is a transport protocol that supports resilience, although such capacity depends on the failover parameters of SCTP. SCTP is resilient within standard configurations, although with a poor performance, and not without any support ($= 0$), as reported by QoR. This result puts in evidence the granularity of REF.

A key aspect in REF is the independence of application characteristics to assess the resilience performance. REF is an important tool to assess the resilience support of a protocol, and its main novelty relies in the integration of multihoming capabilities in the evaluation, as well as, the independence of topologies and protocols. In addition, REF can be employed to define comparison points when assessing the resilience support since it can be used for any protocol without any modification.

We conclude that REF is capable of objectively assessing the resilience support of a protocol in a straightforward manner, without relying on theoretical parameters. Being REF a comparison tool we aim to employ it in the evaluation of resilience capacity of protocols in a testbed.

## ACKNOWLEDGMENT

The first author would like to acknowledge the support of the PhD grant SFRH/BD/61256/2009 from Ministério da Ciência, Tecnologia e Ensino Superior, FCT, Portugal.

## REFERENCES

[1] R. Stewart (ed.), "Stream Control Transmission Protocol," IETF Request for Comments: 4960, September 2007.

[2] J. Iyengar, P. Amer, and R. Stewart, "Concurrent Multipath Transfer Using SCTP Multihoming Over Independent End-to-End Paths," *Networking, IEEE/ACM Transactions on*, vol. 14, no. 5, pp. 951–964, 2006.

[3] P. Cholda, A. Jajszczyk, and K. Wajda, "A unified quality of recovery (QoR) measure," *Int. J. Commun. Syst.*, vol. 21, no. 5, pp. 525–548, 2008.

[4] A. Autenrieth, "Differentiated Resilience in IP-Based Multilayer Transport Networks," Ph.D. dissertation, Technische Universitat Munchen, Munchen, April 2003.

[5] P. Cholda, J. Tapolcai, T. Cinkler, K. Wajda, and A. Jajszczyk, "Quality of Resilience as a Network Reliability Characterization Tool," *Network, IEEE*, vol. 23, no. 2, pp. 11–19, 2009.

[6] ITU-T, "Recommendation M.495, Transmission Restoration and Transmission Route Diversity: Terminology and General Principles," 1993.

[7] M. Pioro and D. Medhi, *Routing, Flow and Capacity Design in Communication and Computer Networks*. Elsevier, July 2004.

[8] "A measurement-based analysis of multihoming," in *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, Karlsruhe, Germany.

[9] S. Huang, Y. Xu, and L. Zhang, "A Path Diversity Metric for End-to-End Network," in *Dependable Computing, 2007. PRDC 2007. 13th Pacific Rim International Symposium on*, 2007, pp. 115–122.

[10] M. Al-Kuwaiti, N. Kyriakopoulos, and S. Hussein, "A Comparative Analysis of Network Dependability, Fault-tolerance, Reliability, Security, and Survivability," *Communications Surveys & Tutorials, IEEE*, vol. 11, no. 2, pp. 106–124, 2009.

[11] J. Eklund, K. Grinnemo, S. Baucke, and A. Brunstrom, "Tuning SCTP Failover for Carrier Grade Telephony Signaling," *Computer Networks*, August 2009.

[12] L. Budzisz, R. Ferrús, A. Brunstrom, K. Grinnemo, R. Fracchia, G. Galante, and F. Casadevall, "Towards Transport-layer Mobility: Evolution of SCTP Multihoming," *Computer Communications*, vol. 31, no. 5, pp. 980–998, March 2008.

[13] J. Fitzpatrick, S. Murphy, M. Atiquzzaman, and J. Murphy, "Using Cross-Layer Metrics to Improve the Performance of End-to-End Handover Mechanisms," *Computer Communications*, vol. In Press, Accepted Manuscript, June 2009.

[14] Adrian Gauch and Yoshifumi Nishida, "SCTP Profiling Framework for Multi-homed Environment," in *Proceedings of the 1st International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking*. Los Angeles, California: ACM, 2006, pp. 54–56.

[15] ITU-T, "ITU-T Recommendation E.800, Series E: Overall Network Operation, Telephone Service, Service Operation and Human Factors - Quality of telecommunications services: concepts, models, objectives and dependability planning - Terms and definitions related to the quality of telecommunication services ," 9 2008.

[16] Achim Autenrieth, "Recovery time analysis of differentiated resilience in MPLS," in *Design of Reliable Communication Networks, 2003. (DRCN 2003). Fourth International Workshop on*, 2003, pp. 333–340.

[17] Eusebi Calle and Jose L. Marzo and Anna Urra, "Protection performance components in MPLS networks," *Computer Communications*, vol. 27, no. 12, pp. 1220–1228, July 2004.

[18] D. Griffith, K. Sriram, S. Klink, and N. Golmie, "Optimal Mixtures of Different Types of Recovery Schemes in Optical Networks." [Online]. Available: http://www.antd.nist.gov/pubs/paper1.pdf

[19] "OMNeT++ - discrete event simulation environment," http://www.omnetpp.org/, 2009.

[20] I. Rungeler, M. Tuxen, and E. P. Rathgeb, "Integration of SCTP in the OMNeT++ simulation environment," in *Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems \& workshops*. Marseille, France: ICST, 2008, pp. 1–8.

[21] ITU-T, "ITU-T Recommendation G.723.1, Dual Rate Speech Coder For Multimedia Communications Transmitting at 5.3 and 6.3 kbit/s ," 3 1996.