Fault tolerance and network integrity measures: the case of computer-based systems

Peter Andras, Olusola Idowu, Panayiotis Periorelis

School of Computing Science

University of Newcastle

NE1 7RU, Newcastle upon Tyne, UK

Short title: Network integrity measures of fault tolerance

Abstract

Fault tolerance is a key aspect of the dependability of complex computer-based systems. Fault tolerance may be difficult to measure directly in complex real world systems, and we propose here to measure it in terms of integrity preservation of the system under the assumption of a particular fault occurrence distribution. We measure the integrity preservation ability of the system by measuring the change of structural integrity of the graph representing the system while it is exposed to random node removal according to the assumed fault distribution. We show how to use such measures to measure the integrity reservation of computer-based systems and in this way indirectly their fault tolerance. We discuss the application of the proposed method in the context of a real world example, the Linux operating system. The results indicate that integrity preservation metrics can serve as an appropriate measure of fault tolerance of complex computer-based systems.

Keywords: complex computer-based systems, measurement-based evaluation, network analysis, network model, probabilistic modelling and evaluation, operating systems

Corresponding author:

Peter Andras

School of Computing Science, University of Newcastle

Newcastle upon Tyne, NE1 7RU, United Kingdom

Tel. +44-191-2227946; Fax. +44-191-2228232; E-mail: peter.andras@ncl.ac.uk

1. Introduction

The concept of fault tolerance of complex computer-based systems, and in particular of computers and software, emerged very early in the 1950s [1]. It was recognized that unexpected faults may emerge in computer-based systems, and that effective dealing with such faults it is critical for highly dependable systems. Fault tolerance is a key measure of the dependability of computer-based systems [1,2], dependability being defined as reliability, availability, safety, security, survivability, and maintainability of a system [3].

Generally systems can be perceived as a set of units that are interconnected by their actions and behaviours [4]. Computer-based systems can be seen as systems with units which can be computer hardware, software, humans, and possibly a variety of other machines and human artefacts containing sensors and actuators. The interconnecting actions and behaviours of these units can take the form of data entry to the computers, data communications between hardware components, data interchange and processing by software components, and display or communication of data to actuators.

An interesting issue is how to measure the fault tolerance of a computer-based system. Systematic mathematical analysis of fault tolerance of models of computer-based systems started in 1960s [1]. Typically fault tolerance is evaluated by full probabilistic analysis of the system, by calculating measures such as mean time to failure and mean time to repair under the assumption of a fault occurrence scenario (e.g., identical and independent fault occurrence distribution for each system component) [5-8].

One stream of fault tolerance research is focused on the analysis of graphs that represent computer-based systems [9-11]. These works assume a fault occurrence scenario in the graph (e.g., node failure or edge failure) and measure the probability of connectedness [12, 13] or of having flow capacity above a given limit [14, 15] of the graph as a proxy measure for the fault tolerance of the system represented by the graph. The main drawback of these methods is that they are very computationally intensive and in many cases they are restricted to a narrow range of particular graph topologies [16-18].

An alternative way to analyse the robustness of systems is to use structural graph analysis methods that reveal vulnerable components and the sensitivity to structural damage of the system [19]. These methods assess the integrity of the system and the change of integrity measures after structural damage to the system in terms of structural measures, such as diameter, average minimum path length or average clustering coefficient. The underlying theoretical assumption is that system structural integrity implies functional integrity of the system [20, 21]. This is supported by practical examples, which show that structural integrity and functional integrity of systems are strongly correlated [19, 22]. Consequently, the analysis of the structural integrity of the graph representing a system by appropriate structural measures can provide indicator measures of the functional integrity of the system.

We propose in this paper the use of structural graph analysis methods to measure the integrity of computer-based systems. We measure the likely structural damage as an approximation of likely functional damage due to the presence of faults. In this way we can assess the fault tolerance of the system by measuring the likely change of structural integrity of the system.

The rest of the paper is structured as follows. Section 2 discusses system integrity measures. In Section 3 we analyse the link between fault tolerance and integrity measures. Section 4 presents an example of the application of the proposed methodology to the assessment of fault tolerance of computer-based systems. Finally, in section 5 we draw some conclusions of the paper.

2. System Integrity

Systems are sets of component units interconnected by their interactions [4]. Component units interact by their behaviour modifying the state of the units participating in such interactions. In a stronger sense we may consider systems as only those sets of interacting component units, in which the interactions between components depend primarily on earlier interactions between system components [23]. We should also point out that system components may also

interact with other units, which are not part of the system. Such interactions constitute the system's interaction with its environment.

The integrity of a system can be defined in functional terms as the system's ability to perform the full range of system behaviours [21]. The system behaviours are possible patterns of behaviours of its component units [1]. Some of these behaviours may have an effect on the system's environment, while others may cause only a change of the internal behaviour of the system.

Measuring functional integrity directly may be difficult, as the full range of possible system behaviour may not be known [21]. A way to approximate the functional integrity of a system is to measure its structural integrity [20]. In practical cases of living cells [22], nervous systems of animals [24], and technological systems [19] it has been shown that their functional integrity correlates strongly with their structural integrity. Measuring structural integrity is much simpler than measuring functional integrity, in the sense that it requires only the measurement of the existence of components and interactions between components, disregarding the actual functional semantics of interactions and interaction patterns.

Structural integrity measures of systems are based on the measurement of the structural integrity of the underlying graph structure of the system, which is made of nodes representing system units, and edges or arcs representing undirected or directed interactions between system units. (We consider undirected graphs only in what follows). A graph representation of a complex system of nodes and interactions is shown in Figure 1.

Simple measures of structural integrity of graphs include the diameter, the average minimum path length and the average clustering coefficient of the graph. The diameter is defined as the largest of the minimal path lengths between nodes of the graph:

$$D(G) = \max\{l(i, j) \mid i, j \in V\}, G = \{V, E\} \quad (1)$$

where V is the set of nodes and E is the set of edges of the graph, and l(i,j) is the minimal length of a path between the nodes i and j. The average minimum path length is defined as the average length of minimal paths between all pairs of nodes of the graph:

$$\mu(=\frac{1}{|V|^2}\sum_{i,j\in V} l(i,j), G=\{V,E\} \tag{2}$$

where we use the same notations as above. The clustering coefficient of a node is the proportion between the number of existing edges between the neighbours of the node and the number of all possible edges:

$$c(i)=\frac{2\cdot|\{(i,j)\,|\,(i,j)\in E\}|}{|\{j\,|\,(i,j)\in E\}|\cdot(|\{j\,|\,(i,j)\in E\}|-1)}, i\in V, G=\{V,E\} \tag{3}$$

The average clustering coefficient of the graph is the algebraic average of the clustering coefficients of all nodes:

$$\eta(G)=\frac{1}{|V|}\cdot\sum_{i\in V} c(i) \tag{4}$$

We note that the above measures evaluate somewhat different aspects of the graph integrity; none of them provides a comprehensive evaluation of the graph integrity. In order to be on the safe side in practical applications the best practice is to use such a set of simple integrity measures and evaluate the graph integrity using the resulting set of integrity measure values (i.e., by considering a vector of integrity measure values). In particular, if we need a single value measure of the graph integrity on the basis of a vector of integrity measures, the safest is to take the value indicating the greatest amount of integrity loss.

Other more sophisticated measures of graph integrity include the calculation of coefficients of the graph's characteristic polynomial, and eigenvalues of the graph's adjacency matrix. These methods can provide a full picture of the graph's integrity and in principle capture all its aspects. The disadvantage of these methods is that they are computationally very expensive, and the calculation of the required numbers may be impractical for very large graphs representing complex systems. The above introduced simple integrity measures are well correlated with the more general measures. The largest eigenvalue of the adjacency matrix is related to the density of the edges, the second eigenvalue is related to the conductivity within the network [25]. The second coefficient of the characteristic polynomial is related to the

number of edges, while the third coefficient is twice the number of triangles in the network [26].

An important issue regarding the use of graph integrity measures to assess the integrity of systems is that of how to actually measure the system components and their interactions. One approach can be to consider the design of the system, if this is available. (For technological systems this might often be the case.) However, this approach can lead one to fall into the trap of showing the robustness of the designed system and not of the actual system. We believe that the right approach is to measure the existing components of the real system and their existing interactions in order to assess the integrity and robustness of the actual system. However, we recognize that in some practical cases such measurements might prove to be difficult (e.g., monitoring of human – computer interactions), limiting the applicability of the structural graph analysis based assessment of system integrity evaluation.

In the case of computer-based systems we typically have a set of non-computer related units (e.g., humans, sensors), a set of hardware units making the computer hardware part of the system, and usually a very large set of software modules, constituting the units of the system. In some cases we ignore the non-computer related units and even the hardware part of the system and we focus our attention exclusively on the system made of software units. The interactions between software units take the form of data transactions between them, which can be measured by appropriate monitoring of the system [27].

3. Fault tolerance and integrity preservation

Faults are unexpected behaviours of system components. Faults in computer-based systems may have a number of origins; they can be classified as design faults, physical faults and interaction faults [3]. Faults cause errors in the system, which are deviations from the expected behaviour of the system. Errors in computer-based systems may stay latent, until they are detected, when they cause abnormal behaviour at the interface of the system with its

environment [1]. Errors cause failures of the system, when the system is unable to perform its function correctly [3].

Faults in the system may occur at various places. An important feature of faults is their occurrence within the system and their distribution at these places within the system. In many cases we may suppose that the faults may appear at any system unit according to the same occurrence distribution, no unit being more susceptible for producing faulty interactions than others [5]. In some cases we may also use the hypothesis that the likelihood of faulty interactions is proportional with the likelihood of the unit being involved in interactions. In other cases the fault distribution may follow some peculiar well defined distribution, such as in the case of faults induced by malicious logic (e.g., attacks by hackers). The types of fault distributions determine the fault occurrence environment of the system.

Fault tolerance is the ability of the system to maintain its functionality in the presence of active faults [1]. Fault tolerance is typically achieved by error detection, recovery and fault handling [3]. Fault tolerance of computer-based systems depends on the fault occurrence environment of the system (e.g., in presence of naturally occurring faults the system may prove sufficiently fault tolerant, while in the presence of targeted attack by hackers, it may prove fault sensitive).

In general the measure of fault tolerance of the system is a relative measure, which shows to what extent the system preserves its functionality in a certain fault occurrence environment [1]. To assess the fault tolerant nature of a computer-based system we need to assess the level of functionality of the system within the considered fault occurrence environment. In other words we need to evaluate the functionality preservation ability of the system. Usually some probabilistic approach is used to evaluate fault tolerance measures such as mean time to failure or mean time to repair [1]. These methods take into consideration the whole system resulting computationally very intensive analyses in case of large systems [6, 9]. To perform such exhaustive evaluations may prove difficult in practice, as monitoring and assessing all aspects of the functionality of the system and performing all the required calculations may be extremely time and resource consuming [21]. Alternative methods were proposed recently,

involving game theoretic approaches [10], formal languages inspired analysis [28] and structural network analysis approaches [19].

We adopt the structural network analysis approach, which has the key advantage that it implies a relatively low computational load for the evaluation of large systems. We measure the fault tolerance of the system by evaluating the ability of the system to preserve its integrity. The measure of integrity preservation is calculated by using system integrity measures based on a structural graph analysis of the graph representing the system. As structural integrity is strongly correlated with functional integrity, the structural integrity preservation measure provides a proxy measure of the functional integrity preservation measure of the system. Consequently, we can use the structural integrity measures introduced in the previous section to measure the change of the integrity of the system in a given fault occurrence environment.

To measure the effects of faults on the integrity of the system, we simulate the faults by sampling the fault occurrence distributions and then evaluating the integrity measures of the system in the presence of simulated faults. The presence of faults causes the elimination from the graph of the system of edges between nodes or of nodes of the graph. These changes happen according to the fault occurrence distributions and have the effect that the integrity measures of the system graph are modified. The expected changes in terms of integrity measures may be calculated analytically in the case of small systems or can be evaluated by numerical simulations in the case of large and complex systems. The expected changes associated with a fault occurrence environment characterise the system's integrity preservation ability and are used as an approximate measure of the fault tolerance of the system.

To show how to use the calculated integrity measures to assess the fault tolerant nature of a system we consider below a toy example. Let us consider a software system of 1000 units of which corresponding graph representation is shown in Figure 2. The system's structural integrity measures are the following:

- diameter: D(S)=19

- average minimum path length: $\mu(S)=3.58151$
- average clustering coefficient: $\eta(S)=0.022702$

We consider a fault occurrence environment in which the faults occur with equal uniform probability (p=0.15) at each unit of the system and each fault temporarily knocks out the system unit where it occurs. To evaluate the fault tolerance of the system we perform a numerical simulation of the fault occurrences, and evaluate the integrity measures of the system for each simulation. After the simulations we calculate the average values and variances of the system integrity measures. We chose to run 20 simulations in order to get reliable estimates of mean values (the variance of the mean value calculated from n measurements is σ/square-root(n), where σ is the variance of the calculated values). The calculations after the simulations led to the values:

- diameter: $\overline{D}(S) = 23.05, \sigma(D(S)) = 4.1355$
- average minimum path length: $\overline{\mu}(S) = 3.69947, \sigma(\mu(S)) = 0.03978$
- average clustering coefficient: $\overline{\eta}(S) = 0.023049, \sigma(\eta(S)) = 0.001438$

To evaluate the integrity preservation ability of the system we calculate first, whether the average values of system integrity measures after the simulation of faults differ significantly or not from the corresponding values calculated for the fully functional system. Next we calculate the normalized distance of the pre-damage and post-damage integrity measure values, which together with their attached statistical significance levels characterize the fault tolerance of the system. In order to be on the safe side, we choose the worst measure (i.e., the largest and most significant distance) to be the numerical evaluation of the fault tolerance of the system. In the case of the above system the normalized distances (z-score, i.e., the distance measured between the mean value and original value in units equal to the standard deviation – $(v_{original} - m)/\sigma_m$) and statistical significance levels (statistical significance levels show how likely is that the original value is the same as the estimated mean value after damage, low p-value indicates that the likelihood of them being the same is very low, or in other words the two values differ significantly) are listed in Table 1.

In the case of the above toy example we have shown how to apply in principle the proposed structural graph analysis based integrity evaluation methods to assess the fault tolerance of a computer-based system. The data shown in the last column of Table 1 shows the values of the likelihoods that original value of the integrity measure is the same as its value after the damage. The results indicate that under the above described fault occurrence environment assumption the system suffers significant damage ($p<0.01$) in terms of diameter and average shortest path length, the amount of the latter damage being more significant than the former. Considering the most significant damage (i.e., the damage in terms of average path length, $p=7.29 \times 10^{-14}$), we conclude that under the considered fault occurrence assumption the system represented by the graph suffers very significant structural and functional damage, and consequently has low fault tolerance.

4. Applications

Linux is one of the most popular operating systems, which is due to a good extent to its open source based development. It is commonly claimed that Linux is more reliable and secure than many other operating systems. An immediate question is how fault tolerant is Linux actually.

We analysed the network structure of the Linux under typical running conditions with a set of usual programs running. To perform the analysis we considered the calls between the classes present in the Linux kernel (version 2.4.19). We found 6815 classes and 19909 calls between them, by parsing the source code of the classes. The interaction network of the classes (see Figure 3) was then analysed in terms of structural network analysis.

Analysing the connectivity distribution of the processes we found that the distribution follows a power law distribution similar to the case of the Internet [19] (see Figure 4). This indicates that among the Linux classes there are relatively few very highly connected classes (which call and are called by many other classes) and many others with relatively few connections. This implies that similarly to the Internet [19] the Linux is very robust and fault tolerant if

faults happen randomly following a uniform fault distribution over the processes (run-time representation of classes), while it should be very vulnerable and fault sensitive if faults are distributed such that they affect mostly the most highly-connected processes.

We performed an analysis of the Linux class network to evaluate the effects of faults on integrity measures. We simulated a scenario with uniform random distribution with $p = 0.05$ probability of faults at each node. We also performed a simulated a scenario when the likelihood of a node being faulty was proportional with the connectivity of the node. The analysis results are shown in Table 2 and Table 3.

The results show that as we expected Linux is remarkably fault tolerant in a fault occurrence environment characterised by uniform fault distribution, while it is significantly more fault sensitive in the case of a fault distribution centred on the mostly linked processes. This suggests that indeed the common belief about the reliability and fault tolerance of Linux is well founded in case of random uniformly distributed errors, but also highlights that Linux is also a vulnerable system in case of well designed malicious attacks.

5. Conclusions

Fault tolerance is measurable aspect of the dependability of computer-based systems. Direct measurement of fault tolerance of large real world systems poses considerable problems, considering that most existing work is focused on exhaustive analytical evaluation of relatively simple model systems [9-11]. An approach to find a proxy measure for the fault tolerance of large systems is to measure their structural integrity preservation under the assumption of a fault occurrence environment. This measure is based on the assumption that functional integrity is strongly correlated with structural integrity [20, 21], which is supported by experimental analysis of various real world complex systems [19, 22].

Integrity and integrity preservation of large systems can be measured by structural graph analysis of their graph representation, where nodes represent system units and edges represent interactions between system units. Simpler (e.g., average minimum path length) or more

complex (e.g., eigenvalues of the adjacency matrix) measures can be used to estimate the structural integrity preservation of the system while exposed to faults occurring in accordance with an assumed fault distribution. Using these measures we can evaluate the likely amount of loss of integrity (or integrity damage) and the statistical significance of this loss.

A toy example and a real world example were presented to show the application of this fault tolerance measurement approach. These examples show that indeed the proposed methods can be applied effectively and lead to meaningful conclusions about the analysed systems. In the case of the real world example (Linux) the analysis indicates that the system is very fault tolerant under the assumption of uniform fault distribution, but that, not surprisingly, it is very vulnerable under the assumption of a fault distribution driven by properly targeted malicious interventions. However, the methods discussed here enable the extent of this effect to be analyzed in useful detail.

We believe that using relatively simple network integrity measures can simplify considerably the effective analysis of fault tolerance of large real world computer-based systems. Although these methods do not provide an exact measure of fault tolerance they provide good approximations of the actual measure. Such approximate measures can be used to rapidly determine the effects of a variety of fault occurrence environments, allowing the designers and developers of large systems to prepare appropriate defence and repair strategies to support the dependability of their system effectively and efficiently.

Acknowledgement: The authors thank for very helpful comments from Brian Randell.

References

1. Lee, PA and Anderson, T: Fault Tolerance. Principles and Practice. Wien: Springer-Verlag, 1990, 2nd ed.

2. Laprie, JC (ed): Dependability: basic concepts and terminology in English, French, German, Italian and Japanese. Wien, Springer-Verlag, 1992.

3. Aviziensis, A, Laprie, JC, Randell, B: Fundamental concepts of dependability. Newcastle University Report no. CS-TR-739 2001.

4. von Bertalanffy, L: General System Theory: foundations, development, applications. Harmondsworth, Penguin, 1973.

5. Amari, SV: Generic rules to evaluate system-failure frequency. IEEE Transactions on Reliability 2000; 49: 85-87.

6. Chang, Y-R, Amari, SV, Kuo, S-Y: Computing system failure frequencies and reliability importance measures using OBDD. IEEE Transactions on Computers 2004; 53: 54-68.

7. Ou, Y and Dugan, JB: Approximate sensitivity analysis for acyclic Markov reliability models. IEEE Transactions on Reliability 2003; 52: 220-230.

8. Scherrer, C, and Steininger, A: Dealing with dormant faults in an embedded fault-tolerant computer system. IEEE Transactions on Reliability 2003; 52: 512-522.

9. Billinton, R, Jonnavithula, S: Calculation of duration, frequency, and availability indexes in complex networks. IEEE Transactions on Reliability 1999; 48: 25-30.

10. Bell, MGH: The use of game theory to measure the vulnerability of stochastic networks. IEEE Transactions on Reliability 2003; 52: 63-68.

11. Cheng, Y, He, Z: Bounds on the reliability of distributed systems with unreliable nodes & links. IEEE Transactions on Reliability 2004; 53: 205-215.

12. Beichelt, F and Tittmann, P: A generalized reduction method for the connectedness probability of stochastic networks. IEEE Transactions on Reliability 1991; 40: 198-204.

13. Elperin, T Gertsbakh, I Lomonosov, M: Estimation of network reliability using graph evolution models. IEEE Transactions on Reliability 1991; 40: 572-581.

14. Chan, Y, Yim, E, Marsh, A: Exact and approximate improvement to the throughput of a stochastic network. IEEE Transactions on Reliability 1997; 46: 473-486.

15. Kishimoto, W: Reliable flow with failures in a network . IEEE Transactions on Reliability, 1997; 46: 308-315.

16. Al-Sadi, J, Day, K, Ould-Khaoua, M: Unsafety vectors: a new fault-tolerant routing for the binary n-cube. Journal of Systems Architecture 2002; 47: 783-793.

17. Goerdt, A: Random regular graphs with edge faults: expansion through cores. Theoretical Computer Science 2001; 264: 91-125.

18. Goerdt, A and Molloy, M: Analysis of edge deletion processes on faulty random regular graphs. Theoretical Computer Science 2003; 297: 241-260.

19. Albert R, Jeong H, Barabási A-L: Diameter of the World Wide Web. Nature 1999; 401: 130-131.

20. Andrews, JD, Beeson, S: Birnbaum's measure of component importance for noncoherent systems. IEEE Transactions on Reliability 2003; 52: 213-219.

21. Ferrandi, F, Fummi, F, Pravadelli, G, Sciutto, D: Identification of design errors through functional testing. IEEE Transactions on Reliability 2003; 52: 400-412.

22. Jeong, H, Mason, SP, Barabasi, A-L, Oltvai, ZN: Lethality and centrality in protein networks. Nature 2001; 411: 41-42.
23. Charlton, B, Andras, P: The nature and function of management - a perspective from systems theory. Philosophy of Management 2004; 3: 3-16.
24. Scannell, JW, Blackmore, C, Young, MP: Analysis of connectivity in the cat cerebral-cortex. Journal of Neuroscience 1995; 15: 1463-1483.
25. Farkas, IJ, Derenyi I, Barabasi A-L, Vicsek T: Spectra of Real-World Graphs: Beyond the Semi-Circle Law. arXiv:cond-mat/0102335 2001.
26. Biggs, N: Algebraic Graph Theory. Cambridge, Cambridge University Press, 1994.
27. Periorellis, P, Idowu, OC, Lynden, SJ, Young, MP, Andras, P: Dealing with complex networks of protein interactions: A security measure. In Proceedings of 9th IEEE International Conference on Engineering of Complex Systems (ICECCS), Bellini, P, Bohner, SA, Steffen, B (eds) pp 29-36, IEEE Computer Society 2004.
28. Phoha, VV, Nadgar, AU, Ray, A, Phoha, S: Supervisory control of software systems. IEEE Transactions on Computers 2004; 53: 1187-1199.

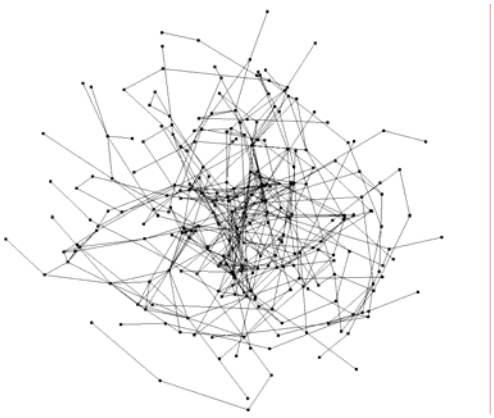

Figure 1. A graph representation of a complex systems, in which system units are represented as nodes and their interactions as edges of the graph.

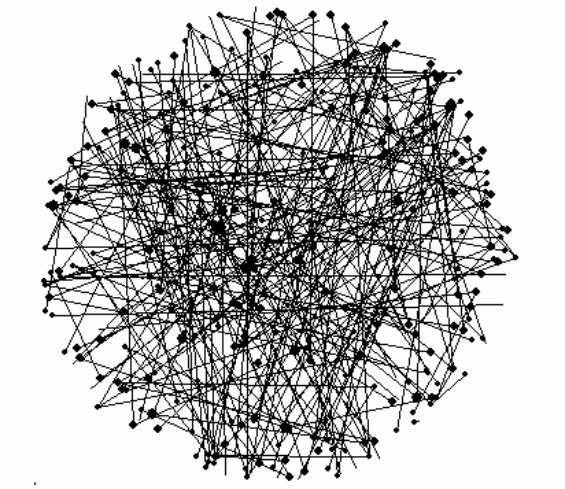

Figure 2. The graph representation of the model system with 1000 nodes. The size of the nodes indicates the number of connections of the node. Only the subset of more connected nodes and the subset of connections between these nodes are displayed to keep the figure comprehensible.

Table 1. Summary of integrity measures of the system before and after damage, including the z-score for the original values considering the mean and variance of the after damage values (z-score = (original – damage mean)/(damage variance / square-root(20))), and the statistical significance level of the difference between the original values and the mean values calculated after the damage. The p values above 0.1 are omitted.

Integrity measure	Original	Damage mean	Damage variance	z-score	p-value
Diameter	19	23.05	4.1355	4.37964	0.000012
AvShortPath	3.58151	3.29378	0.03978	13.25906	7.29×10^{-14}
AvClusCoef	0.02270	0.02304	0.00143	1.07918	-

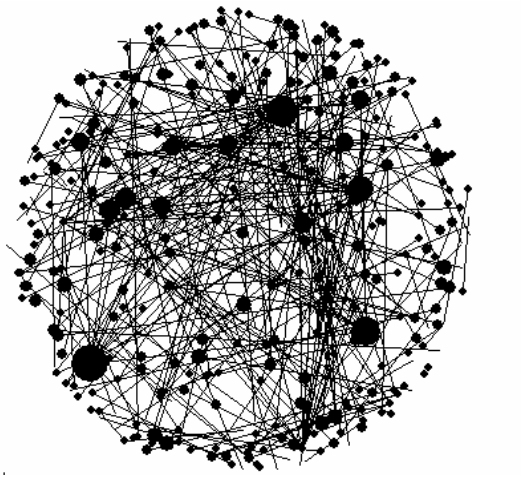

Figure 3. The graph representation of the Linux. The size of the nodes indicates the number of connections of the node. Only the subset of more connected nodes and the subset of connections between these nodes are displayed to keep the figure comprehensible.

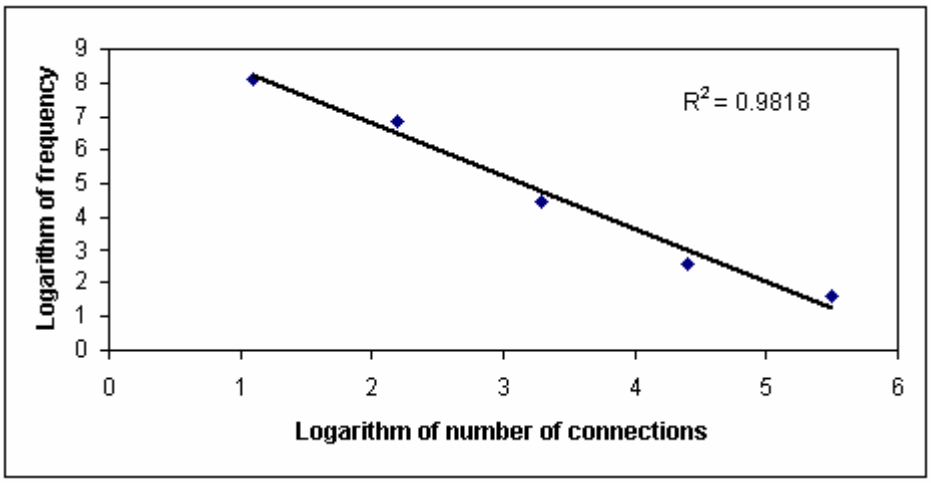

Figure 4. The relationship between the number of connections and the number of nodes with connectivity in a given range. The horizontal axis shows the logarithm of the number of connections, and the vertical axis the logarithm of the number (frequency) of connections with connectivity value in the corresponding range (the considered connectivity value ranges were, 0-3, 4-9, 10-27, 28-81, and 82-243) The figure shows the log-log graph, indicating clearly the power law distribution of connectivity values (i.e., the graph is linear, indicating a connectivity frequency relationship of the form $f = \lambda \cdot c^{a}$, where a is the slope of the line fitted to the data points and λ is a parameter given by the horizontal intersection point of the line)

Table 2. Summary of integrity measures of Linux before and after random damage, including the z-score for the original values considering the mean and variance of the after damage values (z-score = (original – damage mean)/(damage variance / square-root(20))), and the statistical significance level of the difference between the original values and the mean values calculated after the damage. The p values above 0.1 are omitted.

Integrity measure	Original	Damage mean	Damage variance	z-score	p-value
Diameter	44	38.85	4.869	4.729	2.25×10^{-6}
AvShortPath	12.01	11.50	0.740	3.077	0.00209
AvClusCoef	0.133	0.133	0.007	0.350	-

Table 3. Summary of integrity measures of Linux before and after targeted damage, including the z-score for the original values considering the mean and variance of the after damage values (z-score = (original – damage mean)/(damage variance / square-root(20))), and the statistical significance level of the difference between the original values and the mean values calculated after the damage. The p values above 0.1 are omitted.

Integrity measure	Original	Damage mean	Damage variance	z-score	p-value
Diameter	44	35.45	5.062	7.552	1.41×10^{-13}
AvShortPath	12.01	10.47	1.444	4.766	1.87×10^{-6}
AvClusCoef	0.133	0.132	0.005	1.018	-