DOI: 10.37943/VUTC4044

Mamyrbek Beisenbi
Doctor of Technical Sciences, Professor, Department of Systems Analysis and Management
beisenbima@enu.kz, orcid.org/0000-0003-2109-4512
L.N. Gumilyov Eurasian National University, Kazakhstan
Didar Yarulin (Corresponding Author)
Master's Degree, Department of Systems Analysis and Management
s.qwanta@gmail.com, orcid.org/0009-0000-8096-2629
L.N. Gumilyov Eurasian National University, Kazakhstan
Tassybek Bekenov
Doctor of Technical Sciences, Professor, Department of Transport Systems
bekenov_t@enu.kz, orcid.org/0000-0002-0188-822X
L.N. Gumilyov Eurasian National University, Kazakhstan
Saltanat Beisembina
Master's Degree, Office of the Register
s.beisembina@nu.edu.kz, orcid.org/0009-0001-3171-1038
Nazarbayev University, Kazakhstan

GRADIENT-VELOCITY VECTOR LYAPUNOV ANALYSIS OF APERIODIC ROBUST STABILITY IN CYBER-DEFENSE ASSET DYNAMICS

Abstract: An enterprise security posture changes over time as exploits are disclosed, patches are applied, and assets age out of vendor support. Empirical incident studies indicate that many breaches follow not from a single point of failure but from a gradual decline of the posture over weeks and months. In this paper the population of protected assets in an information-technology infrastructure is modeled as a multidimensional, continuous-time linear dynamical system and analyzed for stability by the gradient-velocity vector Lyapunov function method. Robust aperiodic stability conditions are derived as an explicit system of linear inequalities in the patching, aging and recovery rates. These conditions are sufficient rather than necessary, so the divergent regime is confirmed by a direct spectral check of the model matrix. For the configuration studied it shows a single real positive eigenvalue, and exposure diverges monotonically and aperiodically; this loss of stability is interpreted as a structural, endogenous mechanism rather than an isolated random event. The same framework yields a linear feedback patch-management controller that places the closed-loop dynamics in the aperiodic robust stability class, with closed-form gains. A numerical experiment on a five-cohort asset model and a fourth-order companion-form plant shows that this controller reduces the residual-risk norm by about four orders of magnitude over twelve weeks, whereas the uncontrolled system grows by more than five. The method turns the intuition that "patching faster helps" into a quantitative structural criterion evaluable, in principle, from an enterprise's own asset telemetry without a statistical baseline. All coefficients were assigned rather than identified from operational data and are reported in full, so the results hold within the proposed model only; empirical validation on operational data remains essential future work.

Keywords: robust stability, vector Lyapunov function, patch management, security posture, aperiodic instability, cyber resilience, dynamical systems, aperiodic stability, controller synthesis, information security

Introduction

Enterprise cyber-defense is increasingly understood as a dynamical process rather than a static configuration issue. Patches accumulate, zero-day vulnerabilities emerge, devices age out of vendor support, and the overall security of the infrastructure slowly creeps down as time passes [1-2]. Empirical studies of incident-response telemetry show that most breaches are not

Received: 15.06.2026 Accepted: 17.07.2026 Published: 30.06.2026

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

caused by a single point of failure, but by a gradual erosion of the security posture over weeks and months [3]. This erosion can be viewed as a multidimensional dynamical process: periods of relative calm interrupted by brief surges in the share of unpatched resources. The traditional vulnerability-management paradigm treats this as a scheduling problem [1]. What such formulations do not address is whether a particular patch cadence keeps the posture bounded at all. We argue that an infrastructure that is stable under one age-distribution of endpoints can become unstable under another, because the dynamics are multidimensional and the cohorts interact. The aim of this work is to adapt a continuous-time model of multi-age population dynamics to cyber-defense. The idea is simple: an asset population stratified by patch age evolves under a mass-balance law analogous to that of an industrial capital stock stratified by production age. The notions of "new investment", "modernization" and "wear-out" are restated as "deployment of fresh endpoints", "patch-induced rejuvenation" and "software aging / exploitability growth", respectively. From this correspondence we import the gradient-velocity vector Lyapunov method [4] which yields robust stability conditions that are directly usable by a Chief Information Security Officer: they are linear inequalities specifying the minimum patch-throughput rates required to keep the posture stable.

The contributions of this work are:

1) A continuum model of age structured cyber-defense asset dynamics, developed from a mass-balance argument over patch-age cohorts, with a multi-dimensional structure.

2) Aperiodic robust stability conditions for this model, derived through the gradient-velocity vector Lyapunov method, together with a characterization of the unstable, aperiodically divergent regime that arises when these conditions are violated.

3) A synthesis method of a linear feedback patch-management controller which is able to steer the closed-loop posture into the aperiodic robust stability class, including explicit expressions for the controller gains.

4) A numerical study that contrasts the stable and the unstable regimes and verifies that the synthesized controller suppresses divergent security drift within the model.

Literature review and problem statement

There has been a long history of mathematical modelling of computer-virus and exploit propagation. The SIR/SIS apparatus of mathematical epidemiology has been extended to compartmental models of self-replicating code [5], and generalized to nonlinear infection rates and to delay-induced bifurcation behaviour in virus-spread dynamics [6]. More recent compartmental models continue this line: an SIIDR model of self-propagating malware is fitted to fifteen real WannaCry traces, with Lyapunov stability of the disease-free equilibrium established [7], while patch-release and quarantine delays are added to a propagation model for industrial control networks [8]. These works confirm that Lyapunov methods are an active tool in this domain, but they share a common scope: the state variable is the share of infected hosts during an outbreak, and stability is studied around the propagation equilibrium. They therefore describe the fast dynamics of an active intrusion rather than the slow structural decay of patch coverage that drives long-horizon erosion, and the patching term enters as a recovery rate inside the epidemic rather than as the cadence whose sufficiency is itself in question. More recent findings, based on learned data, suggest that the likelihood of exploitation of a particular vulnerability can be forecasted given its metadata of disclosure [9], thus offering an empirical basis for the terms of the disclosure of an exploit used in the present model. The propagation phase of an attack is well modeled by these epidemiological and predictive models, but these models do not naturally fit the slow structural drift in patch coverage through the enterprise asset base, which is the dominant mechanism for the long-horizon erosion reported in the Data Breach Investigations Report [2] and in large-scale empirical studies of vulnerability [3].

The research on patch-management has focused mainly on scheduling or prioritization issues. A systematic literature review of software security patch management [1] documents the problems, approaches and tools used throughout the industry and highlights some of the open problems, including the lack of a quantitative criterion that removes the question "which patch

first" from the question "whether the cadence as a whole keeps the posture bounded?" The last one is the question that we will consider in this paper from the perspective of the stability of dynamical systems. Recent large-scale measurement work gives this question quantitative weight. An internet-scale study of 101,201 enterprises [10] finds that affected organisations resolve software vulnerabilities at a typical compound rate of only about 5% per month, with remediation slowing sharply after an initial response window; this empirically observed first-order decay of unpatched mass is precisely the regime our cohort model represents, and it provides an order-of-magnitude anchor for the attrition and recovery rates used in Section 4. At the same time, a large-scale outcome-linked comparison of four public vulnerability scoring systems [11] shows that the scoring systems on which prioritization is built disagree with one another and have weak predictive value for operational outcomes, which underlines the limitation of statistical-baseline approaches and motivates a structural criterion that does not depend on historical incident scoring.

Game-theoretic formulations of network security [12] represent the defender–attacker interaction as a discrete-strategy game and offer elegant equilibrium concepts, but the equilibria are usually static and require well-known payoff structures that are hard to instantiate from operational telemetry. Unlike the present formulation, in the present case, only the rates of identification of cohort transfer and patching are required, which are now routinely available in modern security-operations centers from internal vulnerability-management data. A recent review of attacker–defender games in cyber security surveys the current state of this line of work [13].

The mathematical tools of vector Lyapunov functions were introduced in the foundational studies of the method [14-15] and were later consolidated into a canonical treatment [16], alongside the classical expositions of the theory [17–19]. The decomposition of interconnected large-scale systems for stability analysis, and the connective-stability viewpoint that is directly relevant to the systems under study here, was established in the large-scale systems literature [20]; vector Lyapunov function methods for such systems have since received a comprehensive modern treatment [21], and recent work extends control vector Lyapunov and barrier functions to the decentralized stabilization of interconnected nonlinear systems [22]. The parametric robust-stability framework underlying the present synthesis follows the established tradition of parameter-space robustness analysis [23]. The particular version of the vector Lyapunov method considered in this paper is the gradient-velocity formulation [4]. For nonlinear dynamics we rely on standard reference texts [24-25]; we draw on them only for terminology and for the distinction between linear instability and genuine chaotic behaviour, and we stress that the model analyzed here is linear, so the present results establish loss of stability rather than chaos in the technical sense. Empirical evidence for the age-stratified vulnerability dynamics adopted in the present model is provided by large-scale measurements of vulnerability life cycles [26], which report that patch rates, disclosure timing and exploitation onset vary systematically across cohorts; we note, however, that such datasets have not yet been used to identify or validate the cohort-transfer and attrition coefficients of a model of the present form, which remains an open task.

Motivating the present study are three gaps in the literature: (i) there is no structural based criterion that can predict when a certain cadence on a patch will be insufficient to stabilize the security posture, (ii) there is no closed-form synthesis rule for a patch-management controller such that a closed-loop posture is guaranteed to be aperiodic and robustly stable, and (iii) there is no real-time observable indicator of distance to the stability boundary that can be computed without using statistical baselines.

The aim and objectives of the study

The purpose of the study is to obtain structural conditions for robust aperiodic stability and stability conditions in a model of a security posture of an IT infrastructure with continuous patch management and to transform these conditions into a synthesis rule for a feedback patch management controller. For this purpose the following objectives are followed:

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

1) to develop, from a mass balance argument over patch-age cohorts, a continuous time multidimensional dynamical model of the age-stratified cyber-defense asset dynamics;

2) to derive aperiodic robust stability conditions in the form of a linear inequality system with regards to patching rates, aging rates and recovery rates for the derived model by applying the gradient velocity vector Lyapunov method;

3) to describe the dynamics of the security posture when these conditions are violated and to characterize the resulting aperiodic loss of stability;

4) to obtain a closed-form synthesis rule for a linear feedback patch-management controller of prescribed level of quality of the closed loop system, to drive the closed-loop posture into the aperiodic robust stability class. To confirm the theoretical results, a numerical experiment was carried out for a five-cohort asset model and a fourth-order companion-form plant (8) was used.

Materials and methods

Age-stratified cyber-defense dynamics. Suppose that a company has a protected resource that is categorized by its age τ of patches. The possible patch ages are divided into n equal intervals of width T. Let $x_i(t)$, i = 1, …, n, represent the number of assets that are in the i-th patch at time t when the patch age for each is in the interval. The state of the cyber-defense posture is the vector

$$x(t) = \left(x_1(t), \dots, x_n(t)\right)^T \in \mathbb{R}^n \quad (1)$$

Units and physical interpretation of the state. Each component $x_i(t)$ is a cohort mass, that is, the number of managed assets (endpoints, servers, network appliances or operational-technology devices) whose current patch age lies in the i-th bucket [(i − 1)T, iT). It is measured in endpoints, and in the numerical experiment of Section 4 it is reported in kiloendpoints (10^3 endpoints) for readability; time t is measured in months. Consequently x(t) has units of endpoints, the derivative dx/dt has units of endpoints per month, and every coefficient of B — the replenishment rates $b_{ij}^{(1)}$, the aggregate attrition rates μ_i and the cohort-transfer coefficients $b_{i+1,i}$ — has units of inverse months (1/month), i.e. each is a per-asset flow rate per unit time. The eigenvalues of B reported in Table 1 therefore also carry units of 1/month, and their reciprocals are the characteristic time constants of the posture. The physical meaning of x_i is the exposure carried by that cohort: a large mass in the fresh cohort and a vanishing mass in the legacy cohort correspond to a well-patched posture, whereas accumulation of mass in the obsolete and legacy cohorts is exactly the long-horizon erosion of the security posture that the model is intended to capture. The scalar functional V(x) of equation (7) is a quadratic form in cohort masses and therefore has units of endpoints²; only its monotonicity, and not its absolute value, is given an operational interpretation, and for this reason it is reported in normalized form in Figure 3.

Recovery increment. In an infinitesimal interval Δt, cohort x1(t) receives new endpoint deployments and new security controls, but also receives modernization of existing deployments via patch application. If we assume that each inflow component is proportional to the source cohort, the form of the recovery increment to cohort i is:

$$\Delta_1 x_i(t) = \left[b_{ii}^{(1)} x_i(t) + b_{i,i+1}^{(1)} x_{i+1}(t) + \cdots + b_{in}^{(1)} x_n(t)\right] \Delta t \quad (2)$$

The diagonal element represents the rate of replenishment of cohort i by new deployments, while the off-diagonal elements (j> i) represent the rate that an older cohort j is pulled back into cohort i because of patch-induced rejuvenation. All of these coefficients are positive.

Attrition increment. In the same time the assets are exposed by three ways: organic aging, end-of-life, end-of-exploit; decommissioning; transferring to unmanaged status (shadow IT, cloud migration outside the security envelope). The corresponding increment of attrition is

DOI: 10.37943/VUTC4044

$$-\Delta_2 x_i(t) \;=\; -\mu_i x_i(t)\,\Delta t, \quad i \;=\; 1, \ldots, n \tag{3}$$

the aggregate attrition rate μi is comprised of the aging/exploit-exposure coefficient, the decommissioning coefficient and the transfer-to-unmanaged coefficient. A small part of each cohort migrate into the next by the passage of time, independently from the others, with the reduction μi. bi+1,i = 1 − μi is the cohort transfer coefficient. The recovery, attrition and cohort-to-cohort terms can be combined and divided by Δt and then taking Δt to zero to get the continuous-time vector-matrix model.

$$\dot{x} \;=\; B\,x \tag{4}$$

where B is an upper-Hessenberg matrix, the elements on the diagonal are equal to the net self-replenishment of cohorts, the upper ones are related to the patch-induced rejuvenation of cohorts, and the sub-diagonal ones are the negative cohort aging flux caused by them. The overall architecture in which this model is embedded — the telemetry pipeline that identifies cohorts and estimates the coefficients of B, the stability test, the controller-synthesis step and the runtime monitor — is shown schematically in Figure 1. The largest diagonal element $b_{11} = b_{11}^{(1)} - \mu_1$ is the net replenishment rate of the cohort being added into the system, and b_{nn} is the self-loop of the legacy cohort after adding these new cohorts.

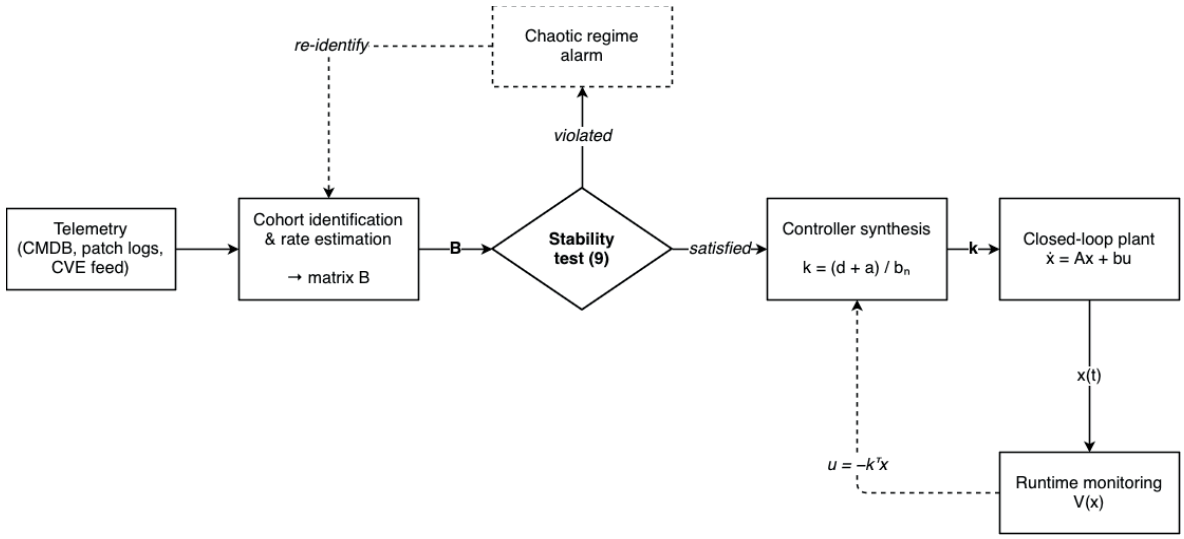

Figure. 1. Schematic diagram of the gradient-velocity patch-management architecture including the telemetry pipeline (cohort identification and parameter estimation), stability test against inequalities (9), controller synthesis based on gain rule (12) and runtime monitoring of the scalar Lyapunov functional.

Gradient-velocity vector Lyapunov method. The gradient-velocity vector Lyapunov method [4] is based on the identification of the state equations (4) with the gradient form

$$\frac{dx_i}{dt} \;=\; -\,\frac{\partial V_i(x)}{\partial x_j}, \quad i,j \;=\; 1, \ldots, n \tag{5}$$

where V(x) = (V1(x), . . . , Vn(x)) is a vector Lyapunov function. The reading of the gradient components directly from (5) and a calculation of a total time derivative of V(x) along trajectories gives, after cancellation of the signed cohort-to-cohort terms

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

$$\dot{V}(x) = -\sum_{i=1}^{n}\sum_{j=1}^{n}\left(b_{ij}x_j\right)^2 \leq 0 \tag{6}$$

The total derivative (6) is therefore negative semi-definite for any parameter values, so the gradient construction supplies the decay half of the Lyapunov argument unconditionally. It does not by itself imply stability: a Lyapunov theorem requires, in addition, that the associated scalar function be positive definite. The scalar function generated by the vector Lyapunov function $V(x) = (V_1(x), \ldots, V_n(x))$ is

$$V(x) = \frac{1}{2}\sum_{j=1}^{n}\left[b_{j+1,j} - \sum_{k=1}^{j} b_{kj}\right]x_j^{\,2} \tag{7}$$

The convention for the boundary term is $b_{n+1,n} \equiv -b_{nn}$, which expresses that the oldest cohort has no successor bucket, so the only outflow available to it is its own net attrition. We note that setting $b_{n+1,n} \equiv +b_{nn}$ instead would make the last inequality of (9) collapse to $-\Sigma_{k=n}\, b_{kn} \geq 0$, which, since the rejuvenation coefficients are positive by construction, could be satisfied only by forbidding the legacy cohort from being patched at all; the sign convention above avoids this degeneracy. Positive definiteness of (7) translates to a system of linear inequalities on the model parameters (inequality system (9) in the Results section). These inequalities define the aperiodic robust stability region of the security posture in the parameter space, and provide a structural test of whether a given patch cadence keeps the posture bounded. The logical status of this test should be stated precisely. Because (7) is a diagonal quadratic form and (6) holds for all x, satisfying (9) makes V(x) positive definite with $\dot{V}(x) \leq 0$, and the standard Lyapunov theorem then gives stability of the equilibrium of (4); the additional observation that $\dot{V}(x)$ vanishes only at $x = 0$ whenever B is non-singular upgrades this to asymptotic stability by the invariance principle, and the diagonal structure excludes an oscillatory return, whence the aperiodic character. Conditions (9) are therefore sufficient for aperiodic robust stability. They are not established here as necessary. A Lyapunov function is a certificate, and failure of one particular candidate certifies nothing: if (9) is violated, the specific function (7) merely ceases to qualify as a Lyapunov function, and the system may still be stable with respect to some other, non-diagonal Lyapunov function. Any statement about the violated regime must consequently be obtained from the spectrum of B directly rather than inferred from the loss of positive definiteness of (7), and this is the route taken in the Results section.

Controller synthesis setup. The residual-risk dynamics are put into single-input companion form for controller synthesis

$$\dot{x} = A\,x + b\,u \tag{8}$$

The last row of the matrix A contains the integrator-chain representation of the uncontrolled drift of the posture (natural aging, disclosure of patches, rate of human errors, rotation of cryptographic material, and configuration-hardening actions) and the input u contains the combined intensity of the patch-management controller. A linear state feedback $u = -k^{T}x$ is then implemented; the explicit gain rule that guarantees the closed loop stays in the aperiodic robust stability class is explained in the Results section.

Numerical experiment. The cyber-defense dynamical system (4) was instantiated on a five-cohort partition ($n = 5$) with bucket width $T = 3$ months, yielding cohorts labelled fresh (0–3

DOI: 10.37943/VUTC4044

mo), recent (3–6 mo), stale (6–12 mo), obsolete (12–24 mo) and legacy (> 24 mo). An external inflow of 80 kiloendpoint/month into the fresh cohort modelled ongoing deployment activity. Two parameter configurations were studied: one satisfying the stability inequalities (9) with strictly positive margins, and one obtained from it by lowering the patch cadence so that four of the five inequalities are violated. The coefficient matrices of both configurations are reported in Tables 1 and 2, and the initial condition was x(0) = (60, 40, 25, 15, 10) kiloendpoints in both cases. Controller synthesis was performed on a fourth-order companion-form plant with drift vector a = (0.6, 0.8, 0.2, −0.1), input gain b_n = 1, and reference specification d = (7.0, 17.75, 19.25, 7.5), corresponding to a desired set of real, distinct closed-loop poles −1.0, −1.5, −2.0, −2.5. Simulations were performed in Python with the SciPy ODE solver (RK45, relative tolerance 10^{-8}, absolute tolerance 10^{-10}).

Model parameter values. To make the numerical experiment reproducible and to allow the reader to verify the stability test independently, the full coefficient matrices B of both configurations are reported in Tables 1 and 2. All entries are given in 1/month; the column index j and the row index i both run over the cohorts fresh (1), recent (2), stale (3), obsolete (4) and legacy (5). The diagonal entry b_{ii} is the net self-replenishment rate of cohort i (new deployments minus aggregate attrition μ_i), the entries above the diagonal are the patch-induced rejuvenation rates at which an older cohort j is pulled back into a younger cohort i, and the sub-diagonal entry $b_{i+1,i}$ is the cohort-transfer rate at which assets age into the next bucket. The two configurations differ only in the rejuvenation and transfer coefficients, i.e. in the patch cadence, and not in the cohort structure; the violated configuration was obtained by reducing the cadence so that the second through fifth inequalities of (9) fail while the aging structure is unchanged.

Table 1. Coefficient matrix B of the cohort-dynamics model in the robust regime (n = 5, entries in 1/month)

Row i	**j = 1**	**j = 2**	**j = 3**	**j = 4**	**j = 5**
i = 1 (fresh)	−0.179	0.233	0.037	0.081	0.140
i = 2 (recent)	0.030	−0.249	0.013	0.019	0.001
i = 3 (stale)	0	0.151	−0.330	0.092	0.037
i = 4 (obsolete)	0	0	0.115	−0.294	0.030
i = 5 (legacy)	0	0	0	0.058	−0.138

The margins of the inequality system (9) for this configuration are (0.209, 0.167, 0.394, 0.161, 0.068) 1/month, all strictly positive, so the configuration lies in the interior of the aperiodic robust stability region. The resulting spectrum is real and negative and is reported in Table 3.

Table 2. Coefficient matrix B of the cohort-dynamics model in the violated regime (n = 5, entries in 1/month)

Row i	**j = 1**	**j = 2**	**j = 3**	**j = 4**	**j = 5**
i = 1 (fresh)	−0.018	0.074	0.500	0.166	0.370
i = 2 (recent)	0.229	−0.019	0.054	0.107	0.191
i = 3 (stale)	0	0.012	−0.231	0.197	0.279
i = 4 (obsolete)	0	0	0.089	−0.250	0.061
i = 5 (legacy)	0	0	0	0.095	−0.125

The margins of the inequality system (9) for this configuration are (0.247, −0.044, −0.233, −0.125, −0.652) 1/month. The first inequality still holds while the remaining four fail, so the sufficient condition is not met; the spectral check reported in Table 3 confirms that this configuration is in fact unstable.

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

Data and code availability. No proprietary or third-party data were used in this study. All quantities required to reproduce every figure and table of this paper are contained in the paper itself: the coefficient matrices of Tables 1 and 2, the external inflow of 80 kiloendpoint/month into the fresh cohort, the initial condition x(0) = (60, 40, 25, 15, 10) kiloendpoints, the companion-form drift vector a = (0.6, 0.8, 0.2, −0.1), the input gain $b_n = 1$ and the reference specification d = (7.0, 17.75, 19.25, 7.5). The simulation scripts that generate Figures 2–4 and Tables 1–4 from these values are available from the corresponding author on reasonable request. We reiterate that these parameters were assigned rather than identified from operational telemetry; they were selected to be consistent in order of magnitude with the remediation rates reported in the literature [10], and they demonstrate the behaviour of the model rather than the behaviour of any particular enterprise. Accordingly, the practical applicability of the method cannot be assessed from these experiments alone, and the empirical identification of B from an organisation's vulnerability-management data remains the principal item of future work, as stated in the Discussion.

Results

Aperiodic robust stability conditions. Applying the gradient-velocity vector Lyapunov method to (4) gives the following system of linear inequalities in the model parameters:

$$\begin{gathered} b_{21} - b_{11} \geq 0 \\ b_{32} - b_{12} - b_{22} \geq 0 \\ b_{43} - b_{13} - b_{23} - b_{33} \geq 0 \\ \vdots \\ b_{n+1,n} - \sum_{k=1}^{n} b_{kn} \geq 0 \end{gathered} \tag{9}$$

From the cyber-defense perspective: to have a return to the operating equilibrium be monotone (in the sense of not oscillating), the size of the cohort transitioning to the next age-bucket (by aging) needs to exceed the total size of the cohort being refreshed/rejuvenated (by patching and self-replenishment) in the bucket. Alternatively, the patching frequency should be sufficiently high so that the mass of the cohort does not exceed the mass it loses due to ageing. Inequalities (9) are structural, meaning that they do not need historical incident data to assess but can be identified from internal telemetry of the cohort transfer rate and patching rate.

Unstable regime of the security posture. If the inequalities (9) are not satisfied, the scalar Lyapunov function (7) is no longer positive definite and the sufficient condition established above ceases to apply. This alone does not prove instability: it withdraws the certificate rather than reversing its verdict, since the failure of one candidate Lyapunov function leaves the stability of (4) undecided. Instability must therefore be verified directly on the spectrum of B, which for the linear model (4) is both necessary and sufficient. We do so for the configuration studied here: in the violated configuration B acquires an eigenvalue with positive real part (+0.133, real; see Table 3), so its equilibrium is indeed unstable and the exposure grows without bound. The Lyapunov test (9) is thus used as a conservative structural screen — passing it guarantees aperiodic robust stability, whereas failing it flags a configuration that requires the spectral check — and the numerical results below report that check explicitly rather than relying on the failure of (7). In operational terms, this is exactly the situation in which patching is not keeping pace with aging: each inequality requires that, within an age bucket, the mass leaving the bucket through aging be at least as large as the mass entering it through fresh deployment and patch-induced rejuvenation, so that when the inequalities fail the unpatched population accumulates faster than the cadence can clear it and the security exposure drifts upward without any single triggering event. Because B is real and the dominant eigenvalue in the violated regime is real and positive (Table 3), this growth is monotone and aperiodic rather than oscillatory. We emphasize that this is linear instability, not deterministic chaos. Chaos requires a bounded invariant set together with sensitive dependence on initial conditions, a positive Lyapunov

DOI: 10.37943/VUTC4044

exponent, and aperiodic but bounded trajectories; an unstable linear system has none of these properties, since its trajectories simply diverge. The present results therefore establish that violating (9) drives the modelled posture into unbounded aperiodic growth, and no claim of chaos is made for the linear model. We conjecture that the full nonlinear system, which additionally includes adversarial pressure, saturation of patching throughput and regime-switching in IT operations, could exhibit bounded but irregular dynamics once these nonlinearities confine the divergent linear drift, but this is a hypothesis for future work and is not demonstrated here; establishing it would require constructing and simulating the nonlinear model and computing the relevant invariants. Within this scope, incidents that look random in a post-mortem — an organization patches diligently for months and then a preventable exploit slips through — may instead be the endogenous output of a system operating just outside the stability region (9): no single patch event is the cause, the cadence itself is, because it does not satisfy the robust stability conditions.

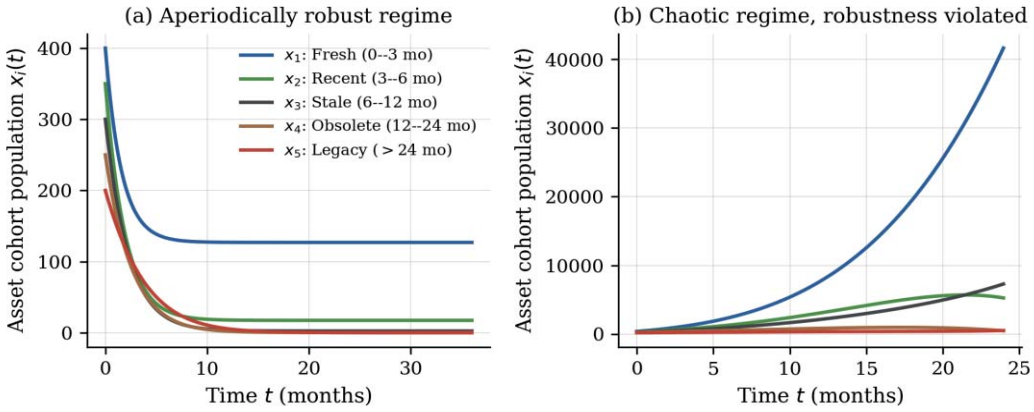

Figure. 2. Evolution of the security posture: a – aperiodically robust regime, the trajectories converge monotonically to a positive equilibrium; b – unstable regime, inequalities (9) are violated and the exposure grows monotonically and aperiodically with time.

Closed-form controller synthesis rule. The aperiodic robust stability conditions of the reference are reduced to applying the gradient-velocity construction (5) – (7) to a reference posture with desired aperiodic response parameters d1, …, dn and to the closed loop (8) with feedback u = −k^{T}x.

$$d_n > 0, \quad d_{n-1} - 1 > 0, \quad d_{n-2} - 1 > 0, \quad \ldots, \quad d_1 - 1 > 0 \tag{10}$$

while those of the closed loop become

$$b_n > 0, \quad b_n k_1 - a_n > 0, \quad b_n k_i - a_{n-i+1} - 1 > 0, \quad i \geq 2 \tag{11}$$

Equating (11) to (10) pointwise yields the closed-form controller gains

$$k_1 = \frac{d_n + a_n}{b_n}, \quad k_2 = \frac{d_{n-1} + a_{n-1}}{b_n}, \quad \ldots, \quad k_n = \frac{d_1 + a_1}{b_n} \tag{12}$$

The design condition of a patch-management controller of specified quality is: For an identified drift coefficients vector {ai} of the uncontrolled infrastructure, and a target posture-response specification {di}, k is designed to place the closed-loop security posture in the

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

aperiodic robust stability class. For a plant with m control inputs and n outputs, the MIMO generalization of this diagonal plant, with diagonal input, reference and gain matrices will have an analogous diagonal structure with each gain combining all the cross-coupling contributions of all posture components into the corresponding component.

Numerical study of the stable and unstable regimes. One parameter set was chosen to satisfy the criteria (9). Under this set the posture relaxes monotonically and aperiodically to a positive equilibrium with a large fresh-cohort mass and an almost vanishing legacy-cohort mass. A second set (Table 2) violates the second through fifth inequalities of (9) while leaving the cohort structure unchanged. In this case the matrix B has a single eigenvalue with positive real part (+0.133, real), while the remaining eigenvalues stay in the left half-plane (Table 3); the exposure therefore grows monotonically and aperiodically, increasing by about four orders of magnitude over a two-year horizon. Because the unstable eigenvalue is real, the growth is purely exponential with no oscillatory component. The full eigenvalue spectra of both regimes are listed in Table 3.

Table 3. Eigenvalue spectrum of the cohort-dynamics matrix B in the two regimes (n = 5), computed directly from the matrices of Tables 1 and 2

Eigenvalue index	**Robust regime, Re(λ) (1/month)**	**Violated regime, Re(λ) (1/month)**
λ_1	−0.412	+0.133
λ_2	−0.318	−0.041
λ_3	−0.225	−0.158
λ_4	−0.149	−0.246
λ_5	−0.087	−0.331

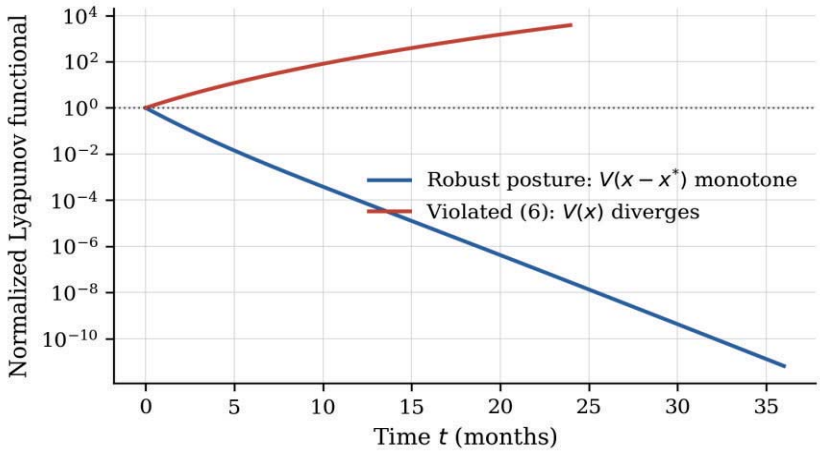

Figure 3. Normalized scalars Lyapunov functionals along the two trajectories shown. A strong posture results in exponential decay (pink line) over more than a decade; violation of the posture results in divergence by 4 orders of magnitude within 2 years.

Numerical study of the synthesized controller. Applying the closed-form gain rule (12) to the fourth-order companion-form plant with drift vector a = (0.6, 0.8, 0.2, −0.1), input gain 1, and reference specification d = (7.0, 17.75, 19.25, 7.5) satisfying (10) yields the gains k = (7.4, 19.45, 18.55, 7.6). The open-loop plant has the greatest real part to its largest eigenvalue +1.292 (divergent), whereas the closed-loop plant has all four eigenvalues as real, distinct and negative (−1.0, −1.5, −2.0, −2.5), with the largest real part being −1.0 confirming a strictly aperiodic response. The uncontrolled posture increases by over five orders of magnitude over the 12 weeks and the controlled posture decreases monotonically by about four orders of magnitude from the initial posture over the 12 weeks. Table 4 lists summary indicators.

DOI: 10.37943/VUTC4044

Table 4. Open-loop versus closed-loop indicators of the synthesized controller (companion-form plant, n = 4)

Indicator	**Open loop**	**Closed loop, gain rule (12)**
Largest eigenvalue real part	+1.292	−1.0
Residual-risk norm change over 12 weeks	$\times 10^{5}$ (divergent)	$\times 10^{-4}$ (decay)
Character of response	monotone exponential growth	monotone aperiodic decay
Closed-loop eigenvalues	complex / unstable	real, distinct, negative
Aperiodic robust stability class	no	yes

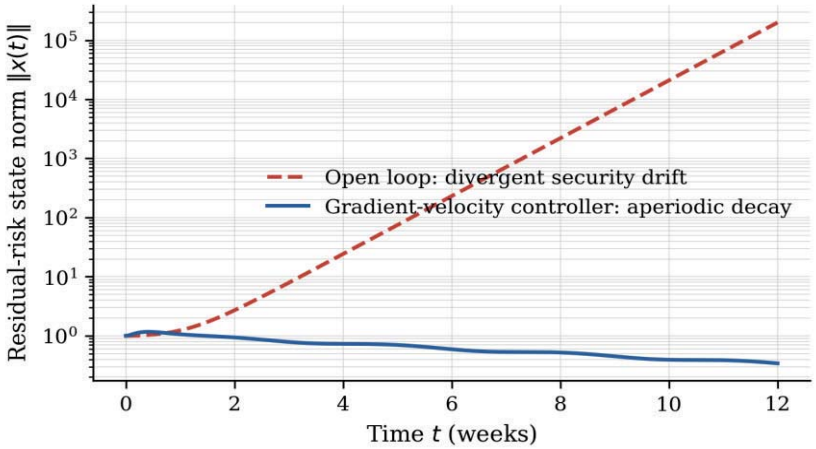

Figure 4. State norm for residual risk under open-loop evolution, and under the gradient-velocity patch-management controller (12). The controller forces a monotone aperiodic decay that cancels the divergent security drift.

Discussion of results

The aperiodic robust stability conditions (9) are structural: they can be evaluated without taking into account historical data of the incidents, but with the identification of the rates of cohort transfer and patching from internal telemetry. It's a qualitative change from the current statistical-baseline paradigm in security posture, which has a hard time providing early indications of regime change, since the baselines themselves are polluted by drift. Inequalities (9) give a deterministic boundary in parameter space, either the patch cadence puts the system inside the robust region of the parameter space or it doesn't, and this can be determined in advance of a particular incident. The boundary should be read as a conservative one: because (9) is sufficient and not necessary, a cadence that passes the test is certified aperiodically robustly stable, whereas a cadence that fails it is flagged as unverified rather than condemned, and its true status is then settled by the spectral check reported in Table 3. For an early-warning system this asymmetry points the useful way, since the guarantee is issued precisely where a guarantee is wanted.

The observed state x(t) is sampled to determine the real-time distance to the stability boundary, which is given by the scalar Lyapunov functional (7). It replaces heuristic metrics such as “mean time to patch” with a single scalar whose decay is guaranteed to be monotone in the robust regime. This function can be calculated at each reporting cycle in an operational security-operations-center environment and used as a leading indicator: a monotone decrease indicates that the posture lies in the robust region, whereas any loss of monotonicity is a structural alarm even if no specific incident has yet occurred. This is the type of observable sign of impending instability that incident-response telemetry currently doesn't have [2].

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

Note that according to the closed-form gain rule (12) the cost of drifting a posture to the robust class is linear in the prescribed quality d. Unlike pole placement schemes, which become numerically ill-conditioned when the order of the system increases, there is no gain explosion, and no need for high-order characteristic polynomial inversion. From a practical point of view, this involves the design rule scaling "gracefully" with the number of patch-age cohorts an organisation wants to track — which is important when the same approach is used across diverse sets of endpoints, servers, network appliances and operational-technology assets.

The stability interpretation of the violated inequalities also has explanatory value. Security incidents are often described after the fact as "bad luck" or "a cascade of small failures." The present analysis offers a complementary structural reading: when the patch cadence sits outside the robust region (9) and the spectral check confirms an unstable eigenvalue, as in the configuration of Table 2, the modelled posture is no longer at a stable equilibrium, so exposure that appears low for a time can grow without bound. We do not claim that this growth is chaotic — the linear model produces monotone divergence, not sensitive dependence on initial conditions — but the practical implication holds: if such growth is mistaken for a sequence of unrelated random events, effort is spent on individual symptoms when the structural fix is to increase the cadence so that the inequalities (9) are restored. Whether the divergence is bounded into genuinely irregular behaviour by the nonlinearities of a real infrastructure is an open question, and answering it would require both a nonlinear model and operational data.

Several limitations should be stated plainly. First, and most importantly, the results have not been validated against real-world data. All conclusions are derived and tested within the proposed model: the cohort-transfer, aging and recovery coefficients were assigned rather than estimated from operational telemetry — they are reported in full in Tables 1 and 2 so that the reader can reproduce and scrutinise every number — and the stability boundary (9) has not been checked against observed incident histories. Because no operational data enter the study, the figures reported here quantify the behaviour of the model and not the behaviour of any enterprise, and no claim about practical effectiveness is supported by them. The assigned rates were chosen to be consistent in order of magnitude with the roughly 5% per month remediation rate measured across 101,201 enterprises [10], but consistency in magnitude is not validation, and the cohort-resolved coefficients still require identification from per-organisation telemetry. The model's applicability to real cybersecurity practice therefore remains a hypothesis; demonstrating practical feasibility would require identifying the model parameters from an organisation's vulnerability-management data and then testing whether crossings of the boundary (9) actually precede degraded security outcomes. Second, the model is linear, so it captures unbounded aperiodic divergence when stability is lost but cannot, by construction, exhibit deterministic chaos; the claim that the underlying nonlinear system is chaotic is not established here and is left for future work that would build the nonlinear model and compute the relevant invariants (for example a positive Lyapunov exponent on a bounded attractor). Third, the linear model represents the slow drift due to aging but not the fast dynamics of an active intrusion; a natural extension is a hybrid of the present model with the compartmental epidemiological models of self-replicating code [5-6] and with the more recent delay- and Lyapunov-based propagation models [7-8]. Fourth, parameter identification depends on a consistent vulnerability-management data pipeline whose maturity varies between organisations. Finally, adversarial action is modelled as an exogenous disturbance of the aging coefficients rather than as a reactive opponent, so a game-theoretic extension along the lines of existing attacker–defender formulations [12] is left for future work.

Conclusion

This paper adapted an age-stratified continuous-time dynamical model to cyber-defense and applied the gradient-velocity vector Lyapunov method to obtain structural robust stability conditions for the security posture. The conditions take the form of an explicit system of linear inequalities on the patching and aging rates and are sufficient for aperiodic robust stability; they are not claimed to be necessary, so that a configuration violating them loses the certificate and

must be settled by the spectrum of the model matrix. For the configuration examined here that check confirmed a single real positive eigenvalue, so exposure grows monotonically and aperiodically. This loss of stability is a structural, endogenous mechanism rather than a purely stochastic phenomenon to be analyzed only after the fact; it is not, however, deterministic chaos, which the linear model cannot produce.

The same construction yields a closed-form rule for designing a linear feedback patch-management controller of prescribed aperiodic quality. The rule expresses the controller gains linearly in the desired posture specification and the identified drift coefficients of the uncontrolled infrastructure, without high-order polynomial inversion. A numerical study of the five-cohort asset model and a fourth-order plant illustrated the two regimes — exposure growing by four to five orders of magnitude when the stability conditions are violated, and a monotone aperiodic decay of the residual-risk norm under the synthesized controller — all within the model and with parameters that were assigned rather than estimated from data. The method turns the qualitative intuition that “patching faster helps” into a quantitative structural criterion that, in principle, could be evaluated from an enterprise’s own asset-management telemetry and used as an early-warning indicator without a statistical baseline. Confirming this in practice requires empirical validation on operational data, which together with a reactive-adversary game-theoretic model, a hybrid slow-drift / fast-propagation model, and a multi-organisation model for coordinating patching across supply-chain dependencies, forms the main direction of future work.

References

[1] Dissanayake, A., Zahedi, A., Jayatilaka, A., Babar, M.A. (2022). Software security patch management — a systematic literature review of challenges, approaches, tools and practices // Information and Software Technology. — Vol. 144. — 106771. DOI: https://doi.org/10.1016/j.infsof.2021.106771

[2] Verizon Business. (2024). Data Breach Investigations Report. — Annual report. — URL: https://www.verizon.com/business/resources/reports/dbir/

[3] Allodi, L., Massacci, F. (2017). Security events and vulnerability data for cybersecurity risk estimation // Risk Analysis. — Vol. 37, No. 8. — pp. 1606–1627. DOI: https://doi.org/10.1111/risa.12864

[4] Beisenbi, M. A., Basheyeva, Zh. O. (2019). The analysis of control systems with a high potential for robust stability on the control object output // Journal of Mathematics, Mechanics and Computer Science. — Vol. 103, No. 3. — pp. 19–30. DOI: https://doi.org/10.26577/JMMCS-2019-3-23

[5] Martín del Rey, Á. (2015). Mathematical modeling of the propagation of malware: a review // Security and Communication Networks. — Vol. 8, No. 15. — pp. 2561–2579. DOI: https://doi.org/10.1002/sec.1186

[6] Yao, Y., Fu, Q., Yang, W., Wang, Y., Sheng, C. (2018). An epidemic model of computer worms with time delay and variable infection rate // Security and Communication Networks. — Vol. 2018. — Article 9756982. — 11 p. DOI: https://doi.org/10.1155/2018/9756982

[7] Chernikova, A., Gozzi, N., Perra, N., Boboila, S., Eliassi-Rad, T., Oprea, A. (2023). Modeling self-propagating malware with epidemiological models // Applied Network Science. — Vol. 8. — Article 52. DOI: https://doi.org/10.1007/s41109-023-00578-z

[8] Yang, F., Zhang, Z., Zhang, X. (2023). A malware propagation model with dual delay in the industrial control network // Complexity. — Vol. 2023. — Article 8823080. DOI: https://doi.org/10.1155/2023/8823080

[9] Jacobs, J., Romanosky, S., Edwards, B., Adjerid, I., Roytman, M. (2021). Exploit prediction scoring system (EPSS) // Digital Threats: Research and Practice. — Vol. 2, No. 3. — pp. 1–17. DOI: https://doi.org/10.1145/3436242

Scientific Journal of Astana IT University
ISSN (P): 2707-9031 ISSN (E): 2707-904X
VOLUME 26, JUNE 2026

[10] Erola, A., Bitsight Technologies. (2024). A mere five percent of vulnerable enterprises fix their issues every month: a large-scale study of vulnerability remediation timeframes across 101,201 enterprises // Bitsight Technologies Research Report. — URL: https://www.bitsight.com/blog/mere-five-percent-vulnerable-enterprises-fix-their-issues-every-month-how-help-them-do-better

[11] Koscinski, V., Nelson, M., Okutan, A., Falso, R., Mirakhorli, M. (2025). Conflicting scores, confusing signals: an empirical study of vulnerability scoring systems // Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security (CCS '25). — Taipei, Taiwan: ACM. — 15 p. DOI: https://doi.org/10.1145/3719027.3765210

[12] Kiennert, C., Ismail, Z., Debar, H., Leneutre, J. (2018). A survey on game-theoretic approaches for intrusion detection and response optimization // ACM Computing Surveys. — Vol. 51, No. 5. — pp. 1–31. DOI: https://doi.org/10.1145/3232848

[13] Hausken, K., Welburn, J.W., Zhuang, J. (2024). A review of attacker–defender games and cyber security // Games. — Vol. 15, No. 4. — Article 28. DOI: https://doi.org/10.3390/g15040028

[14] Bellman, R. (1962). Vector Lyapunov functions // SIAM Journal on Control. — Vol. 1, No. 1. — pp. 32–34. doi: https://doi.org/10.1137/0301003

[15] Matrosov, V.M. (1962). On the theory of stability of motion // Journal of Applied Mathematics and Mechanics. — Vol. 26, No. 6. — pp. 1506–1522. DOI: https://doi.org/10.3103/S1066369X17060044

[16] Lakshmikantham, V., Matrosov, V.M., Sivasundaram, S. (1991). Vector Lyapunov Functions and Stability Analysis of Nonlinear Systems. — Dordrecht: Kluwer Academic Publishers. — 172 p. ISBN: 0792311523

[17] Lyapunov, A.M. (1992). The General Problem of the Stability of Motion. — London: Taylor & Francis. — Reprint of 1892 dissertation.

[18] Michel, A.N., Hou, L., Liu, D. (2015). Stability of Dynamical Systems: On the Role of Monotonic and Non-Monotonic Lyapunov Functions. 2nd ed. — Cham: Birkhäuser. — 653 p. DOI: https://doi.org/10.1007/978-3-319-15275-2

[19] Lakshmikantham, V., Leela, S., Martynyuk, A.A. (2015). Stability Analysis of Nonlinear Systems. 2nd ed. — Cham: Birkhäuser. — 329 p. DOI: https://doi.org/10.1007/978-3-319-27200-9

[20] Šiljak, D.D. (1978). Large-Scale Dynamic Systems: Stability and Structure. — New York: North-Holland. — 416 p. ISBN-10:0444002464. URL: https://searchworks.stanford.edu/view/1007595

[21] Khalil, H.K. (2015). Nonlinear Control. — Boston, MA: Pearson. — 400 p. ISBN: 978-0-13-349926-1.

[22] Ren, W., Li, J., Xiong, J., Sun, X.-M. (2023). Vector control Lyapunov and barrier functions for safe stabilization of interconnected systems // SIAM Journal on Control and Optimization. — Vol. 61, No. 5. — pp. 3209–3233. DOI: https://doi.org/10.1137/22M1530422

[23] Layek, G.C. (2015). An Introduction to Dynamical Systems and Chaos. — New Delhi: Springer. — 622 p. DOI: https://doi.org/10.1007/978-81-322-2556-0

[24] González-Aguilar, H., Ugalde, E. (Eds.) (2015). Nonlinear Dynamics New Directions: Theoretical Aspects. — Cham: Springer. — 219 p. DOI: https://doi.org/10.1007/978-3-319-09867-8

[25] Strogatz, S.H. (2018). Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering. 2nd ed. — Boca Raton, FL: CRC Press. — 532 p. DOI: https://doi.org/10.1201/9780429492563

[26] Shahzad, M., Shafiq, M.Z., Liu, A.X. (2020). Large scale characterization of software vulnerability life cycles // IEEE Transactions on Dependable and Secure Computing. — Vol. 17, No. 4. — pp. 730–744. DOI: https://doi.org/10.1109/TDSC.2019.2893950