Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Context-Aware Stability-Preserving State Transfer for Structurally Reconfigurable Sdr Receivers

Serhii Liventsev

slivencev@gmail.com

KPI https://orcid.org/0000-0002-9768-8133

Research Article

Keywords: software-defined radio, structural reconfiguration, heterogeneous state spaces, context-aware state transfer, state projection, multiple Lyapunov functions, directional dwell time, transient bit-error rate, QPSK, 16-QAM

Posted Date: September 1st, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10868870/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

CONTEXT-AWARE STABILITY-PRESERVING STATE TRANSFER FOR STRUCTURALLY RECONFIGURABLE SDR RECEIVERS

Serhii Liventsev (0000-0002-9768-8133)

Department of Telecommunication Systems, Igor Sikorsky Kyiv Polytechnic Institute,
Kyiv, Ukraine, Industrialnyi Lane 2, Campus 30, 03056,
s.liventsev@kpi.ua

Abstract. Structural reconfiguration of a software-defined radio (SDR) receiver changes not only its modulation and processing parameters but also the dimension and physical interpretation of its internal state. Direct state resetting or dimension-agnostic copying can therefore produce discontinuities that degrade synchronization and bit-error rate (BER) after switching. This paper develops a context-aware state-transfer procedure for heterogeneous QPSK and 16-QAM receiver models. Common phase, frequency-offset, and filter states are preserved exactly, whereas configuration-dependent demapper-metric biases are reconstructed using a regularized projection conditioned on measured signal-to-noise ratio and its increment. The identified operator is constrained through its projection-energy gain, and the standard multiple-Lyapunov dwell-time argument is applied to determine admissible directional switching intervals. A reproducible paired Monte Carlo study with 200 runs in four channel and interference scenarios evaluates mapping accuracy, transient BER, goodput, switching frequency, and stability-related quantities using direct baseband bit decisions. Relative to context-free ridge regression, contextual projection reduces independent mapping RMSE by 6.20% and the QPSK-to-16-QAM projection-energy gain by 25.82%. Stability-constrained identification further reduces this gain to 1.2449 while increasing mapping RMSE by only 0.91%. Compared with symmetric 4/4 dwell time, the reliability-oriented 3/1 directional profile reduces mean BER by 2.86–7.09% and QPSK-to-16-QAM transient BER by 10.46–27.87%, with a 4.17–6.40% increase in switching frequency. Online recursive adaptation provides no additional benefit under the considered observation model. The results demonstrate that heterogeneous SDR reconfiguration should jointly coordinate context-dependent state transfer, projection-energy constraints, and directional dwell time.

Keywords. software-defined radio; structural reconfiguration; heterogeneous state spaces; context-aware state transfer; state projection; multiple Lyapunov functions; directional dwell time; transient bit-error rate; QPSK; 16-QAM

1. Introduction

Software-defined radio (SDR) permits software-controlled reconfiguration of modulation, coding, synchronization, and filtering. In SDR-RAN and O-RAN, supervisory applications may select admissible configurations from channel and performance measurements [1–3, 6], but initialization of the new signal-processing chain remains a local receiver operation.

Configuration-selection models usually represent reconfiguration as an instantaneous parameter change. This abstraction is incomplete when source and target receivers have different state dimensions: QPSK and 16-QAM may share synchronization coordinates but require different demapper states.

Switched-system theory provides reset maps, multiple Lyapunov functions, and dwell-time conditions [13–18], but it does not identify an SDR-specific transfer operator. Such an operator must preserve common physical states, reconstruct configuration-dependent coordinates, and bound the jump between Lyapunov metrics.

The problem addressed in this paper is therefore formulated as follows: given stable SDR configurations with heterogeneous state spaces, construct a causal inter-configuration state-transfer operator, limit its induced Lyapunov-energy gain, and coordinate it with a directional dwell-time rule.

The study considers five-dimensional QPSK and seven-dimensional 16-QAM receiver models. Phase, frequency-offset, and filter states are copied exactly; demapper-metric biases are reconstructed by a regularized affine map conditioned on measured SNR and its increment.

The identified state Jacobian is coupled to mode-dependent Lyapunov metrics through a projection-energy gain, which is penalized during identification and inserted into the standard dwell-time condition. Direction-dependent gains then motivate separate forward and reverse dwell intervals.

A paired Monte Carlo experiment with 200 runs in four scenarios uses direct baseband bit decisions. Context reduces mapping RMSE by 6.20%, the stability penalty limits the forward gain to 1.2449 with a 0.91% RMSE increase, and the 3/1 profile reduces mean BER by 2.86–7.09% relative to symmetric 4/4 dwell time.

The contribution is limited to the following application-specific results:

1. A physically interpretable contextual state-transfer procedure is formulated for SDR configurations with different state dimensions. It preserves synchronization coordinates exactly and reconstructs target demapper states using only causal channel observations.
2. The projection is identified under an explicit Lyapunov-energy penalty. This reduces the forward energy gain without materially degrading independent mapping accuracy.
3. The standard multiple-Lyapunov dwell-time argument is numerically instantiated using the identified directional gains. The resulting asymmetry motivates separate forward and reverse dwell intervals.

4. A paired four-scenario experiment separates projection accuracy from dwell-time effects and identifies 2/1 and 3/1 profiles for throughput- and BER-oriented operation, respectively.

The scope is an application-specific state-transfer and switching procedure; no new general dwell-time theorem, globally optimal controller, or standards-compliant xApp is claimed.

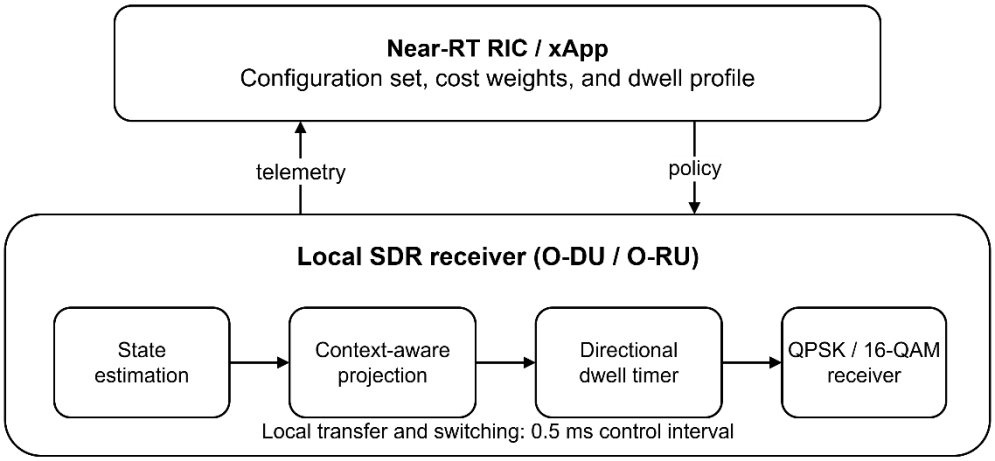


Fig. 1. Hierarchical integration of state transfer into an SDR/O-RAN control architecture.

Section 2 positions the research gap. Sections 3–5 define the receiver, projection, and stability condition. Sections 6–7 present the controller and evaluation; Sections 8–9 discuss applicability and conclude the paper.

2. Related Work and Research Gap

2.1 Adaptive radio and O-RAN control

O-RAN and adaptive-radio studies establish hierarchical control of modulation, coding, power, spectrum, and processing configurations [1–6]. They determine which configuration should be selected but generally omit the receiver state immediately before and after a structural transition.

Here, the Near-RT RIC is only a supervisory integration point; sample-level projection and state initialization remain local receiver functions.

2.2 State transfer and adaptation between models

Bumpless-transfer methods reduce transients when controllers are exchanged [10–12], but they do not directly identify mappings between modulation-dependent receiver states of unequal dimension.

Transfer learning and domain adaptation reuse representations under distribution change [7–9]. They motivate contextual and deployment-shift testing, whereas the present problem additionally requires exact physical invariants and a bounded Lyapunov-energy jump.

2.3 Synchronization and receiver-state modelling

Classical synchronization models justify including phase, frequency-offset, and filter states [19, 20]. These coordinates retain their meaning across QPSK and 16-QAM, whereas the number and interpretation of demapper metrics change with the constellation.

State-aware radio adaptation is also motivated by contemporary RAN automation [1–3], but the cited architectures do not provide a reproducible projection between explicit heterogeneous receiver-state dimensions.

2.4 Switched- and hybrid-system stability

Multiple Lyapunov functions, reset maps, and dwell-time conditions combine within-mode contraction with switching-induced growth [13–18]. Their general formulations assume that the jump map or its gain is known or bounded.

For heterogeneous SDR receivers, this gain depends on target-state initialization: resetting, truncation, and learned projection yield different multipliers under identical fixed-mode dynamics.

2.5 Research gap and positioning

The unresolved problem is to transfer state between SDR configurations of different dimensions while preserving common coordinates, exploiting causal context, limiting the Lyapunov-energy jump, and supporting an implementable switching rule.

The controlled QPSK/16-QAM study separates three claims:

- projection accuracy concerns reconstruction of the target state;
- projection-energy gain concerns compatibility with the Lyapunov metrics;

- BER and goodput concern the operational consequences of the dwell-time policy.

This separation is essential because lower state RMSE need not yield a proportional BER reduction, while lower BER can require more switching or reduced goodput.

3. Physically Interpretable Heterogeneous SDR Model

This section defines the heterogeneous receiver states, their mode-dependent dynamics, the observable channel context, and the direct relationship between state-transfer error and bit decisions. Unlike the previous model, communication degradation is not represented by an empirical state-to-SNR penalty. Instead, synchronization and demapper-state errors act directly on the received baseband samples and decision metrics.

3.1 Configuration set and switching signal

The SDR receiver operates in one of two modulation configurations,

$$\Theta = \{\theta_Q, \theta_{16}\}. \qquad (3.1)$$

In Equation (3.1), $\theta_Q$denotes the QPSK receiver configuration and θ_{16}denotes the 16-QAM configuration. These modes differ in both spectral efficiency and internal state dimension.

The active configuration at sampling instant $k$is specified by

$$\sigma_k \in \Theta. \qquad (3.2)$$

In Equation (3.2), $\sigma_k$is the discrete switching signal. A transition $\sigma_{k-1} = \theta_i \to \sigma_k = \theta_j$activates the inter-configuration state-transfer operator introduced in Section 4.

The state estimator and local switching controller use a sampling interval of $T_s = 0.5$ms. This interval represents the local SDR control loop rather than the execution period of a Near-Real-Time RAN Intelligent Controller.

3.2 Physically interpretable receiver states

The QPSK receiver state is defined as

$$x_{Q,k} = [\phi_k \quad \omega_k \quad z_k \quad b_{I,s,k} \quad b_{Q,s,k}]^T \in \mathbb{R}^5. \qquad (3.3)$$

In Equation (3.3), $\phi_k$is the residual carrier-phase error; $\omega_k$is the residual normalized frequency offset; $z_k$is the synchronization or loop-filter state; $b_{I,s,k}$and $b_{Q,s,k}$are residual biases of the in-phase and quadrature sign-decision metrics. The subscript $s$denotes a sign decision.

The 16-QAM receiver requires two additional decision thresholds and is represented by

$$x_{16,k} = [\phi_k \quad \omega_k \quad z_k \quad b_{I,s,k} \quad b_{Q,s,k} \quad b_{I,a,k} \quad b_{Q,a,k}]^T \in \mathbb{R}^7. \qquad (3.4)$$

In Equation (3.4), $b_{I,a,k}$and $b_{Q,a,k}$are residual biases of the in-phase and quadrature amplitude-decision metrics; the subscript $a$denotes discrimination between the inner and outer 16-QAM levels. The first two demapper coordinates therefore retain the physical meaning of the QPSK sign metrics, whereas the additional coordinates are specific to 16-QAM.

The common synchronization-related state is

$$x_{\mathrm{inv},k} = [\phi_k \quad \omega_k \quad z_k]^T. \qquad (3.5)$$

In Equation (3.5), $x_{\mathrm{inv},k}$defines the configuration-invariant subspace. These coordinates are copied exactly during reconfiguration because their physical interpretation does not change with the modulation order.

Table 1. Physical interpretation of the configuration-dependent receiver states

State	Configuration	Physical meaning	Unit
ϕ	QPSK, 16-QAM	Residual carrier-phase error	rad
ω	QPSK, 16-QAM	Residual normalized frequency offset	rad per normalized symbol interval
z	QPSK, 16-QAM	Loop-filter or matched-filter gain state	dimensionless
$b_{I,s}$	QPSK, 16-QAM	Bias of the in-phase sign metric	dimensionless
$b_{Q,s}$	QPSK, 16-QAM	Bias of the quadrature sign metric	dimensionless
$b_{I,a}$	16-QAM	Bias of the in-phase amplitude threshold	dimensionless
$b_{Q,a}$	16-QAM	Bias of the quadrature amplitude threshold	dimensionless

This definition replaces the ambiguous variables L_1–L_4. It also provides a direct engineering interpretation of projection error: an incorrectly transferred state shifts a specific demapper decision boundary.

3.3 Mode-dependent state dynamics

For configuration $i \in \{Q, 16\}$, the receiver-state error is

$$e_{i,k} = x_{i,k} - \bar{x}_{i,k}. \qquad (3.6)$$

In Equation (3.6), $x_{i,k}$is the estimated receiver state and $\bar{x}_{i,k}$is the configuration-consistent reference state. The vector $e_{i,k}$determines both the Lyapunov energy used in Section 5 and the residual distortion applied to the baseband receiver.

Between structural transitions, the error evolves according to

$$e_{i,k+1} = A_i e_{i,k} + \xi_{i,k}. \qquad (3.7)$$

In Equation (3.7), $A_i$is the mode-dependent state matrix and $\xi_{i,k}$is a zero-mean stochastic disturbance representing estimation noise and unmodelled receiver dynamics.

The matrices used in the reproducible simulator are

$$A_Q = \begin{bmatrix} 0.74 & 0.04 & 0 & 0 & 0 \\ 0 & 0.78 & 0.03 & 0 & 0 \\ 0 & 0 & 0.68 & 0 & 0 \\ 0 & 0 & 0 & 0.58 & 0.05 \\ 0 & 0 & 0 & 0.02 & 0.60 \end{bmatrix}, \qquad (3.8)$$

$$A_{16} = \begin{bmatrix} 0.78 & 0.05 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.82 & 0.04 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.72 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.64 & 0.03 & 0 & 0 \\ 0 & 0 & 0 & 0.02 & 0.66 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.70 & 0.03 \\ 0 & 0 & 0 & 0 & 0 & 0.02 & 0.72 \end{bmatrix}. \qquad (3.9)$$

In Equations (3.8) and (3.9), the first three rows describe residual synchronization dynamics, while the remaining rows describe relaxation of the demapper-metric biases. All eigenvalues lie inside the unit disk. Consequently, each fixed receiver configuration is asymptotically stable; the additional stability problem arises from state jumps at structural transitions.

The disturbance standard deviations are 0.006, 0.018, and $0.008$for the phase, frequency, and filter coordinates, respectively. Each demapper-bias coordinate has a disturbance standard deviation of 0.04. These values are fixed for all compared state-transfer methods.

3.4 Channel, interference, and observable context

For scenario s, the unjammed signal-to-noise ratio evolves as

$$g_k = \alpha_s g_{k-1} + (1 - \alpha_s)\bar{\gamma}_s + \nu_{s,k},\ \nu_{s,k} \sim \mathcal{N}\left(0, \sigma_{\nu,s}^2\right). \qquad (3.10)$$

In Equation (3.10), $g_k$is the background SNR in decibels; $\alpha_s$is its temporal correlation coefficient; $\bar{\gamma}_s$is the scenario-dependent mean SNR; and $\nu_{s,k}$is the innovation process. Scenario-specific values are reported in the experimental section.

The interference indicator is represented by a two-state Markov process,

$$\Pr\ (j_k = 1 \mid j_{k-1} = 0) = p_{\text{on},s},\ \Pr\ (j_k = 0 \mid j_{k-1} = 1) = p_{\text{off},s}. \qquad (3.11)$$

In Equation (3.11), $j_k = 1$denotes active interference; $p_{\text{on},s}$and $p_{\text{off},s}$determine the duration and recurrence of interference intervals.

The effective SNR and its controller-side estimate are

$$\gamma_k = \text{clip}(g_k - J_s j_k, 0{,}15), \quad \hat{\gamma}_k = \text{clip}\left(\gamma_k + \epsilon_{s,k}, 0{,}15\right)$$

where:

$$\epsilon_{s,k} \sim \mathcal{N}\left(0, \sigma_{\epsilon,s}^2\right).$$

In Equation (3.12), $J_s$is the interference attenuation in decibels; $\epsilon_{s,k}$is the SNR-measurement error; and $\text{clip}(\cdot, 0, 15)$restricts the simulated SNR to the interval from 0 to 15 dB. Only $\hat{\gamma}_k$, rather than the true γ_k, is available to the switching controller and contextual projection.

The causal context vector is

$$c_k = \begin{bmatrix} (\hat{\gamma}_k - 7.5)/4 \\ (\hat{\gamma}_k - \hat{\gamma}_{k-1})/2 \end{bmatrix}. \qquad (3.13)$$

In Equation (3.13), the first component represents normalized measured channel quality, while the second represents its local increment. Both quantities are available before a QPSK-to-16-QAM transition and therefore do not introduce future-information leakage.

3.5 Direct baseband communication model

The analytical additive-white-Gaussian-noise BER curves are retained only as reference characteristics:

$$P_b^Q(\gamma) = Q\left(\sqrt{2 \cdot 10^{\gamma/10}}\right),\ P_b^{16}(\gamma) = \frac{3}{4} Q\left(\sqrt{0.8 \cdot 10^{\gamma/10}}\right). \qquad (3.14)$$

In Equation (3.14), P_b^Q and P_b^{16} are the reference QPSK and 16-QAM BER functions, respectively, and $Q(\cdot)$ is the Gaussian tail function. These expressions are used for controller-level quality prediction and analytical verification, but not for imposing a state-dependent SNR penalty.

For transmitted bits $u_{\ell,m} \in \{0, 1\}$, the normalized QPSK and 16-QAM symbols are

$$s_{Q,\ell} = \frac{(1 - 2u_{\ell,1}) + \mathrm{j}(1 - 2u_{\ell,2})}{\sqrt{2}}, \qquad (3.15)$$

$$s_{16,\ell} = \frac{(1 - 2u_{\ell,1})(3 - 2u_{\ell,2}) + \mathrm{j}(1 - 2u_{\ell,3})(3 - 2u_{\ell,4})}{\sqrt{10}}. \qquad (3.16)$$

In Equations (3.15) and (3.16), ℓ is the symbol index and $u_{\ell,m}$ is the m-th transmitted bit. The normalization provides unit average symbol energy for both constellations.

The received baseband sample is generated as

$$r_{\ell,k} = a_z(e_{z,k}) s_{\sigma_k,\ell} \exp\left[\mathrm{j}(e_{\phi,k} + e_{\omega,k}\tau_\ell)\right] + \frac{n_{\ell,k}}{\sqrt{2m_{\sigma_k} 10^{\gamma_k/10}}}, \qquad (3.17)$$

where

$$a_z(e_{z,k}) = \max\left\{0.2,\ 1 - 0.35|e_{z,k}|\right\}. \qquad (3.18)$$

In Equations (3.17) and (3.18), $e_{\phi,k}$, $e_{\omega,k}$, and $e_{z,k}$ are the residual synchronization errors; τ_ℓ is the normalized symbol-time coordinate; $m_Q = 2$ and $m_{16} = 4$ are the numbers of bits per symbol; and $n_{\ell,k}$ is unit-variance circular complex Gaussian noise. Equation (3.17) makes phase, frequency, and filter-state errors act directly on the received signal.

For 16-QAM, the four decision metrics are

$$\begin{aligned} d_{I,s,\ell} &= \mathrm{Re}(r_{\ell,k}) + \kappa e_{b_{I,s},k}, \\ d_{Q,s,\ell} &= \mathrm{Im}(r_{\ell,k}) + \kappa e_{b_{Q,s},k}, \\ d_{I,a,\ell} &= |\mathrm{Re}(r_{\ell,k})| - \frac{2}{\sqrt{10}} + \kappa e_{b_{I,a},k},\ \kappa = 0.12. \\ d_{Q,a,\ell} &= |\mathrm{Im}(r_{\ell,k})| - \frac{2}{\sqrt{10}} + \kappa e_{b_{Q,a},k}, \end{aligned} \qquad (3.19)$$

In Equation (3.19), $d_{I,s,\ell}$ and $d_{Q,s,\ell}$ are sign-decision metrics; $d_{I,a,\ell}$ and $d_{Q,a,\ell}$ are inner-versus-outer amplitude metrics; and κ converts a normalized residual bias into a decision-boundary displacement. For QPSK, only the first two sign metrics are used.

The detected 16-QAM bits are

$$\hat{u}_{\ell,1} = 1\left[d_{I,s,\ell} < 0\right],\ \hat{u}_{\ell,2} = 1\left[d_{I,a,\ell} < 0\right],\ \hat{u}_{\ell,3} = 1\left[d_{Q,s,\ell} < 0\right],\ \hat{u}_{\ell,4} = 1\left[d_{Q,a,\ell} < 0\right]. \qquad (3.20)$$

In Equation (3.20), $1[\cdot]$ is the indicator function. The corresponding QPSK decisions use the sign metrics for the in-phase and quadrature bits.

The instantaneous simulated BER is obtained by direct bit counting,

$$\mathrm{BER}_k = \frac{1}{N_s m_{\sigma_k}} \sum_{\ell=1}^{N_s} \sum_{q=1}^{m_{\sigma_k}} 1\left[\hat{u}_{\ell,q} \neq u_{\ell,q}\right]. \qquad (3.21)$$

In Equation (3.21), $N_s = 128$ is the number of symbols processed per control interval and m_{σ_k} is the number of bits per symbol in the active configuration. State-transfer error therefore affects BER only through the synchronization distortion and decision-boundary displacement explicitly defined in Equations (3.17)–(3.20).

Thus, projection error affects BER through explicit synchronization distortion and shifted decision boundaries, without an empirical state-to-SNR penalty.

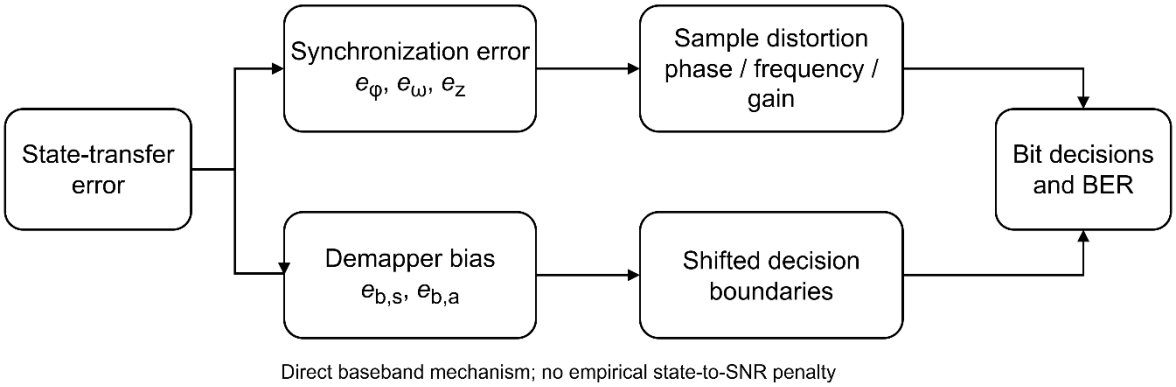


Fig. 2. Influence of state-transfer errors on baseband bit decisions.

4. Context-Aware Inter-Configuration State Transfer

The model in Section 3 separates configuration-invariant synchronization coordinates from modulation-dependent demapper states. This decomposition permits exact transfer of physically identical variables while reconstructing only the coordinates that change their meaning or dimension. The present section defines the causal context-aware projection and examines its sensitivity to the regularization coefficient.

4.1 State-transfer structure

Consider a structural transition from configuration $i$to configuration j. The target state immediately after switching is initialized according to

$$x_{j,k+1}^{\text{post}} = \mathcal{T}_{i\to j}\left(x_{i,k}^{\text{pre}}, c_k\right). \qquad (4.1)$$

In Equation (4.1), $x_{i,k}^{\text{pre}}$is the source state available at the end of interval k; $x_{j,k+1}^{\text{post}}$is the target state used at the beginning of interval $k+1$; $c_k$is the observable context vector from Equation (3.13); and $\mathcal{T}_{i\to j}$is the inter-configuration transfer operator. This timing convention prevents the use of future channel observations.

For the QPSK-to-16-QAM transition, the source state is decomposed as

$$x_{Q,k} = \begin{bmatrix} x_{\text{inv},k} \\ v_{Q,k} \end{bmatrix},\ v_{Q,k} = \begin{bmatrix} b_{I,s,k} \\ b_{Q,s,k} \end{bmatrix}. \qquad (4.2)$$

In Equation (4.2), $x_{\text{inv},k} \in \mathbb{R}^3$contains the phase, frequency, and filter coordinates, whereas $v_{Q,k} \in \mathbb{R}^2$contains the QPSK demapper-metric biases.

The target modulation-dependent state is

$$v_{16,k} = [b_{I,s,k} \quad b_{Q,s,k} \quad b_{I,a,k} \quad b_{Q,a,k}]^T \in \mathbb{R}^4. \qquad (4.3)$$

In Equation (4.3), the first two coordinates retain the QPSK sign-metric interpretation, while the last two coordinates represent the additional 16-QAM amplitude thresholds.

The proposed forward transfer is defined as

$$x_{16,k+1}^{\text{post}} = \begin{bmatrix} x_{\text{inv},k} \\ A_{Q16} v_{Q,k} + G_{Q16}\tilde{c}_k + b_{Q16} \end{bmatrix}. \qquad (4.4)$$

In Equation (4.4), $A_{Q16} \in \mathbb{R}^{4\times 2}$maps the source demapper state; $G_{Q16} \in \mathbb{R}^{4\times 2}$represents the observable channel-context contribution; $b_{Q16} \in \mathbb{R}^4$is an affine offset; and $\tilde{c}_k$is the standardized context vector.

For the QPSK-to-16-QAM transition, the modulation-dependent affine map requires 16 scalar multiplications and 16 additions: eight multiplications for the 4×2 state map, eight for the 4×2 context map, and four bias additions. The three invariant coordinates are copied without arithmetic. The reverse transition requires only structural extraction. Hence, the online complexity is constant, $O(1)$, and no optimization is performed during switching.

Context standardization is performed as

$$\tilde{c}_k = D_c^{-1}(c_k - \bar{c}),\ D_c = \text{diag}\ \left(s_{c,1}, s_{c,2}\right). \qquad (4.5)$$

In Equation (4.5), $\bar{c}$is the training-set context mean and $s_{c,1}$, $s_{c,2}$are the corresponding standard deviations. Standardization prevents the regression from being dominated by one context coordinate solely because of its numerical scale.

The homogeneous state Jacobian associated with Equation (4.4) is

$$\mathcal{P}_{Q16} = \begin{bmatrix} I_3 & 0_{3\times 2} \\ 0_{4\times 3} & A_{Q16} \end{bmatrix}. \qquad (4.6)$$

In Equation (4.6), I_3preserves the invariant synchronization subspace exactly. The matrix $\mathcal{P}_{Q16}$is used in the projection-energy analysis of Section 5, while the context and affine terms are treated as bounded transfer residuals.

For the reverse transition, the QPSK state is obtained by structural truncation,

$$x_{Q,k+1}^{\text{post}} = \begin{bmatrix} I_3 & 0_{3\times 4} & \\ 0_{2\times 3} & I_2 & 0_{2\times 2} \end{bmatrix} x_{16,k}^{\text{pre}}. \qquad (4.7)$$

In Equation (4.7), the invariant coordinates and the two sign-metric biases are retained, whereas the 16-QAM-specific amplitude-threshold states are discarded. This exact structural rule replaces the numerically cluttered reverse matrix with coefficients of order 10^{-6}.

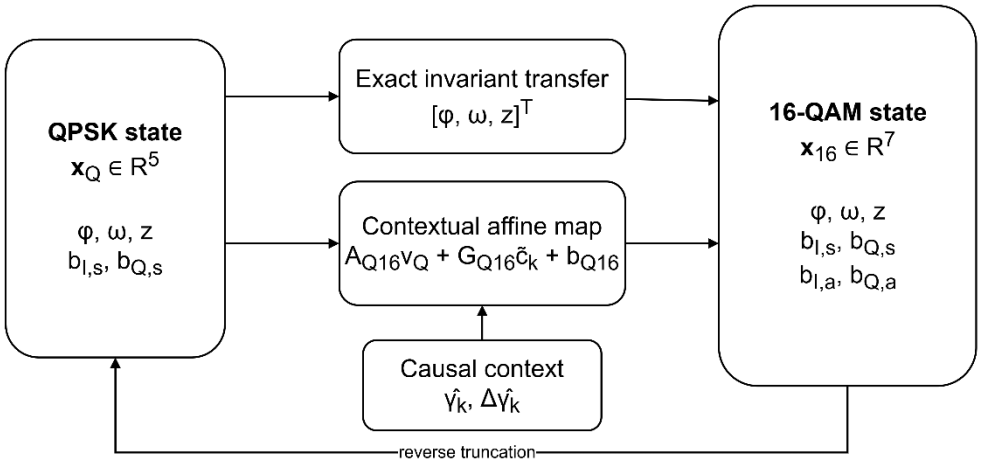


Fig. 3. Context-aware state transfer between heterogeneous QPSK and 16-QAM receivers.

4.2 Causal identification of the contextual map

The training data contain causal transition tuples

$$\mathcal{D}_{\text{tr}} = \left\{ \left(v_{Q,k}^{(r)}, c_k^{(r)}, v_{16,k+1}^{(r)} \right) \right\}_{r=1}^{N_{\text{tr}}}. \qquad (4.8)$$

In Equation (4.8), $r$indexes the training tuple and $N_{\text{tr}} = 5000$. Each source state and context observation is taken at interval k, whereas the target reference state corresponds to interval $k+1$. The identification procedure therefore reproduces the timing of an actual reconfiguration.

The augmented input vector is

$$z_k = \begin{bmatrix} v_{Q,k}^T & \tilde{c}_k^T & 1 \end{bmatrix}^T \in \mathbb{R}^5. \qquad (4.9)$$

In Equation (4.9), the first two coordinates represent the source demapper state, the next two represent the standardized context, and the final constant implements the affine offset.

Let

$$Z = [z_1 \quad z_2 \quad \cdots \quad z_{N_{\text{tr}}}],\ Y_{16} = [v_{16,2} \quad v_{16,3} \quad \cdots \quad v_{16,N_{\text{tr}}+1}]. \qquad (4.10)$$

In Equation (4.10), $Z \in \mathbb{R}^{5\times N_{\text{tr}}}$is the augmented source matrix and $Y_{16} \in \mathbb{R}^{4\times N_{\text{tr}}}$is the causal target-state matrix.

The coefficient matrix is identified from

$$W_{Q16} = \underset{W}{\text{argmin}} [\| Y_{16} - WZ \|_F^2 + \lambda \| W\Gamma \|_F^2]$$

where

$$W_{Q16} = [A_{Q16} \quad G_{Q16} \quad b_{Q16}],\ \Gamma = \text{diag}(1, 1, 1, 1, 0).$$

In Equation (4.11), $\|\cdot\|_F$is the Frobenius norm; $\lambda$is the Tikhonov regularization coefficient; and Γexcludes the affine offset from regularization. The objective penalizes excessive sensitivity to both source-state and context perturbations without shrinking the intercept.

The closed-form solution is

$$W_{Q16} = Y_{16} Z^T (ZZ^T + \lambda \Gamma)^{-1}. \qquad (4.12)$$

In Equation (4.12), all quantities are computed directly from the declared training tuples. The resulting transfer operator is therefore deterministic for the fixed data-generation seeds and regularization coefficient.

For $\lambda = 10^{-3}$, the identified state-dependent block is

$$A_{Q16} = \begin{bmatrix} 0.877890 & -0.004691 \\ 0.012398 & 0.906497 \\ 0.597972 & 0.211298 \\ -0.167143 & 0.666926 \end{bmatrix}. \qquad (4.13)$$

In Equation (4.13), the first two rows reconstruct the sign-metric biases, while the last two rows initialize the 16-QAM-specific amplitude-threshold biases.

The identified context block and affine offset are

$$G_{Q16} = \begin{bmatrix} 0.002150 & -0.000886 \\ -0.004269 & -0.001935 \\ 0.128298 & -0.011595 \\ -0.152606 & -0.004633 \end{bmatrix}, \; b_{Q16} = \begin{bmatrix} 0.029038 \\ -0.008651 \\ 0.118311 \\ -0.097815 \end{bmatrix}. \qquad (4.14)$$

In Equation (4.14), the largest context coefficients occur in the amplitude-threshold rows. This is consistent with the greater sensitivity of the additional 16-QAM decision levels to measured channel quality.

4.3 Independent validation

Projection accuracy is evaluated on $N_{\text{test}} = 20\,000$causal transition tuples that are independent of the training sample. For target state $\bar{x}_{16,k+1}$, the mapping error is measured as

$$\text{RMSE}_{\text{map}} = \sqrt{\frac{1}{7} \left\| x_{16,k+1}^{\text{post}} - \bar{x}_{16,k+1} \right\|_2^2}. \qquad (4.15)$$

In Equation (4.15), $x_{16,k+1}^{\text{post}}$is the projected state and $\bar{x}_{16,k+1}$is the paired configuration-consistent target state. The invariant and modulation-dependent coordinates are both included.

Table 2. Independent validation of QPSK-to-16-QAM state transfer

Transfer method	**Mean RMSE**	**Median RMSE**	**95th-percentile RMSE**	**Soft-state RMSE**
Copy-and-pad	0.28421	0.27086	0.48297	0.35618
Context-free ridge projection	0.18247	0.16628	0.33978	0.21023
Context-aware projection	**0.17117**	**0.15386**	**0.33574**	**0.19285**

Relative to context-free ridge regression, the contextual map reduces mean mapping RMSE by 6.20% and soft-state RMSE by 8.27%. Relative to copy-and-pad transfer, the reductions are 39.77% and 45.86%, respectively. The improvement is therefore concentrated in the modulation-dependent coordinates rather than produced by the exactly copied invariant block.

The state Jacobian of the contextual projection also reduces the QPSK-to-16-QAM projection-energy gain from 2.15596 to 1.59922, corresponding to a 25.82% reduction. This quantity is not used to claim stability at this stage; its formal interpretation is provided in Section 5.

4.4 Selection and sensitivity of the regularization coefficient

The regularization coefficient was examined over the interval

$$\lambda \in \{10^{-4}, 10^{-3}, 10^{-2}\}. \qquad (4.16)$$

In Equation (4.16), the tested interval spans one order of magnitude on either side of the selected value. The training and independent test samples were held fixed for all three cases.

Table 3. Sensitivity of the contextual projection to

λ	**Mean RMSE**	**95th-percentile RMSE**	**Soft-state RMSE**	μ_{Q16}
10^{-4}	0.17116710	0.33574430	0.19284845	1.59921957
10^{-3}	**0.17116711**	**0.33574419**	**0.19284846**	**1.59921610**

10^{-2}	0.17116715	0.33574315	0.19284853	1.59918133

Changing $\lambda$from $10^{-4}$to 10^{-2}changes mean RMSE by less than 3×10^{-5}%and $\mu_{Q16}$by approximately 0.0024%. The identified map therefore lies on a broad numerical plateau rather than at a narrowly tuned optimum.

No L-curve criterion was used because the independent validation error and projection-energy gain were effectively invariant across the tested decade. The value $\lambda = 10^{-3}$was selected as the central positive regularization level and retained for all subsequent experiments. This choice avoids an unregularized inverse while introducing no measurable accuracy penalty.

The observed plateau shows that the contextual gain is not an artifact of a narrowly tuned regularization value.

5. Stability-Constrained Projection and Dwell-Time Condition

The contextual map identified in Section 4 minimizes state-transfer error but does not directly limit the Lyapunov-energy jump produced by structural reconfiguration. This section incorporates the projection-energy gain into the identification criterion and applies the standard multiple-Lyapunov dwell-time argument to the resulting operator.

5.1 Mode-dependent Lyapunov metrics

For each fixed configuration $i \in \{Q, 16\}$, consider the quadratic Lyapunov function

$$V_i(e) = e^T P_i e,\ P_i \succ 0, \qquad (5.1)$$

where $P_i$is obtained from the discrete Lyapunov equation

$$A_i^T P_i A_i - P_i = -I. \qquad (5.2)$$

In Equations (5.1) and (5.2), $e$is the mode-dependent state error and $A_i$is the corresponding state matrix from Section 3. This is a standard construction for stable discrete-time systems.

The one-step contraction factor is

$$\rho_i = \lambda_{\max}\left(P_i^{-1/2} A_i^T P_i A_i P_i^{-1/2}\right). \qquad (5.3)$$

In Equation (5.3), $\lambda_{\max}(\cdot)$denotes the largest eigenvalue. The computed values are $\rho_Q = 0.6244$and $\rho_{16} = 0.6971$, confirming contraction within each fixed configuration.

5.2 Projection-energy gain

For a structural transition $i \to j$, the homogeneous projection-energy gain is defined as

$$\mu_{ij} = \lambda_{\max}\left(P_i^{-1/2} \mathcal{P}_{ij}^T P_j \mathcal{P}_{ij} P_i^{-1/2}\right). \qquad (5.4)$$

In Equation (5.4), $\mathcal{P}_{ij}$is the state Jacobian of the transfer operator. The quantity $\mu_{ij}$is the maximum multiplication of Lyapunov energy caused by the homogeneous part of the identified projection.

The complete transfer error includes the affine context-dependent residual,

$$e_j^+ = \mathcal{P}_{ij} e_i^- + r_{ij,k}, \qquad (5.5)$$

where $e_i^-$and $e_j^+$are the pre- and post-switch errors, respectively, and $r_{ij,k}$contains context-dependent reconstruction error and reference-state mismatch.

Applying Young's inequality with $\varepsilon > 0$gives

$$V_j\left(e_j^+\right) \le (1+\varepsilon)\mu_{ij} V_i(e_i^-) + (1+\varepsilon^{-1}) r_{ij,k}^T P_j r_{ij,k}. \qquad (5.6)$$

In Equation (5.6), the first term determines the multiplicative energy jump, while the second produces a bounded residual contribution. The analysis therefore establishes practical boundedness under stochastic mismatch rather than convergence exactly to zero.

5.3 Stability-constrained identification

Only the state-dependent block A_{Q16}affects the homogeneous gain in Equation (5.4). It is identified by minimizing

$$\mathcal{L}_\beta = \frac{1}{N_{\text{tr}}} \sum_{k=1}^{N_{\text{tr}}} \left\| v_{16,k+1} - A_{Q16} v_{Q,k} - G_{Q16} \tilde{c}_k - b_{Q16} \right\|_2^2 + \beta \mu_{Q16}. \qquad (5.7)$$

In Equation (5.7), the first term is the causal reconstruction loss defined in Section 4; $\mu_{Q16}$is computed from Equation (5.4); and $\beta \ge 0$controls the trade-off between projection accuracy and the induced energy jump.

The selected model minimizes μ_{Q16}subject to an independent mean-RMSE increase not exceeding 2% relative to the unconstrained contextual map,

$$\beta^* = \underset{\beta \geq 0}{\operatorname{argmin}} \mu_{Q16}(\beta) \quad \text{subject to} \quad \mathrm{RMSE}_{map}(\beta) \leq 1.02\mathrm{RMSE}_{map}(0)$$

In Equation (5.8), the constraint prevents a numerically favorable energy gain from being obtained through an excessive loss of transfer accuracy.

Table 4. Accuracy–stability trade-off for the contextual projection

β	μ_{Q16}	**Mean mapping RMSE**	**Calculated dwell bound, samples**
0	1.5992	0.17117	1.436
0.001	1.5071	0.17129	1.272
0.003	1.3487	0.17189	0.964
0.010	**1.2449**	**0.17273**	**0.742**

The selected value $\beta^* = 0.01$reduces $\mu_{Q16}$from 1.5992 to 1.2449, or by 22.16%, while increasing independent mapping RMSE by 0.91%. Increasing βbeyond 0.01 does not reduce μ_{Q16}further because the exactly preserved invariant block imposes a lower gain floor.

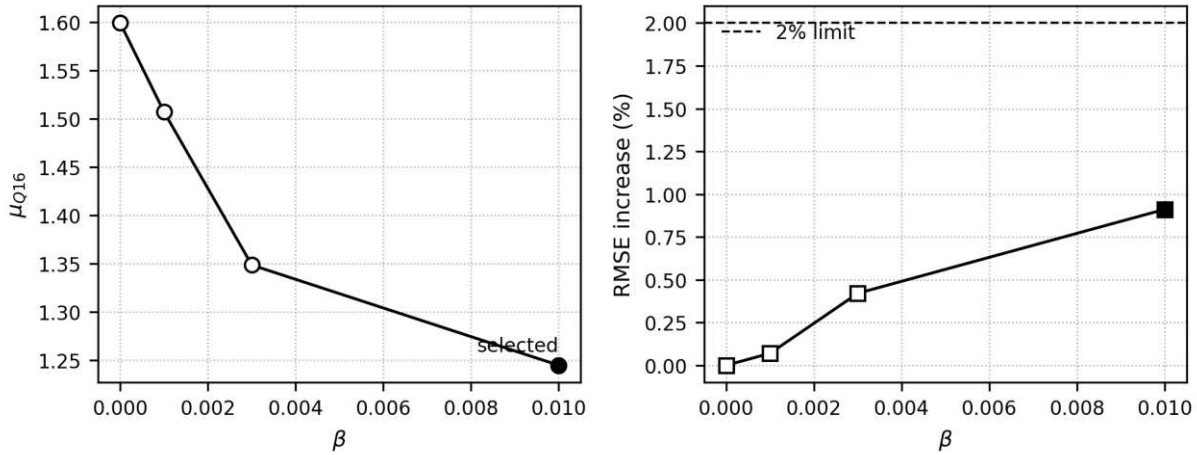


Fig. 4. Accuracy–stability trade-off of the contextual projection.

For the reverse structural truncation, the computed gain is

$$\mu_{16Q} = 0.9079. \qquad (5.9)$$

Equation (5.9) shows that the 16-QAM-to-QPSK transition does not amplify the homogeneous Lyapunov energy. This directional asymmetry motivates different dwell times for the forward and reverse transitions.

5.4 Application of the standard dwell-time argument

Proposition 1 (Application of the multiple-Lyapunov dwell-time argument). Consider the identified projection and mode-dependent Lyapunov metrics in Equations (5.1)–(5.4). Specializing the standard dwell-time result [13, 15–18] to the computed projection-energy gain gives the following sufficient homogeneous contraction condition between successive switches:

$$N_{d,ij} > \frac{\ln\left[(1+\varepsilon)\mu_{ij}\right]}{-\ln \rho_{\max}},\ \rho_{\max} = \max_i \rho_i. \qquad (5.10)$$

In Equation (5.10), $N_{d,ij}$ is the minimum number of samples between transitions in direction $i \to j$; ε is the Young-inequality parameter from Equation (5.6); and $\rho_{\max} = 0.6971$. Equation (5.10) is an application of the established multiple-Lyapunov result, not a new general stability theorem.

Equation (5.10) is obtained by specializing the standard multiple-Lyapunov dwell-time inequality for mode-dependent Lyapunov functions and bounded switching gains [13, 15–18].

Over the dwell interval, within-mode contraction and the subsequent projection produce a factor below unity whenever Equation (5.10) holds.

For $\varepsilon = 0.05$and the selected forward projection,

$$N_{d,Q16} > 0.742,\ N_{d,16Q} > -0.130. \qquad (5.11)$$

In Equation (5.11), the negative reverse bound indicates that the homogeneous reverse map is already non-expansive. The smallest implementable integer dwell profile is therefore 1/1 sample. However, this theoretical value does not account for finite reconstruction residuals, stochastic disturbances, or repeated decisions near the switching threshold.

For this reason, Section 6 evaluates the directional profiles 1/1, 2/1, and 3/1 empirically. The profiles 2/1 and 3/1 retain a margin above the sufficient homogeneous bound while permitting faster recovery from 16-QAM to QPSK than a symmetric dwell-time rule.

6. Directional Dwell-Time Reconfiguration Control

The bound in Section 5 is direction dependent because the QPSK-to-16-QAM projection is expansive in the Lyapunov metric, whereas the reverse truncation is non-expansive. A symmetric dwell-time rule would therefore delay recovery to QPSK without providing an equivalent stability benefit. This section introduces a directional admissibility rule while retaining the same one-step configuration cost for all evaluated modes.

6.1 Configuration cost and switching rule

For candidate configuration $i \in \Theta$, the controller evaluates

$$J_i(k) = w_1 \min\left\{\frac{P_b^i(\hat{\gamma}_k)}{P_{\mathrm{ref}}}, 3\right\} + w_2 D_i + w_3 E_i + w_4 C_{\mathrm{sw}}(i, \sigma_k). \qquad (6.1)$$

In Equation (6.1), $P_b^i(\hat{\gamma}_k)$ is the analytical BER estimate from Equation (3.14); $P_{\mathrm{ref}} = 0.05$; D_i and E_i are normalized delay and energy costs; and C_{sw} equals one for a requested structural transition and zero otherwise. The weights are $w = [0.45, 0.25, 0.20, 0.10]$.

The next configuration is selected as

$$\sigma_{k+1} = \arg\min_{i \in \Theta_{adm}(k)} J_i(k)$$

In Equation (6.2), $\Theta_{\mathrm{adm}}(k)$contains only configurations permitted by the directional dwell-time rule. The controller is a one-step policy; no claim of finite-horizon or global optimality is made.

A candidate transition is executed only if

$$J_{\sigma_k}(k) - J_i(k) \geq \delta_J, \ \delta_J = 0.03. \qquad (6.3)$$

In Equation (6.3), $\delta_J$is the hysteresis threshold used to suppress low-value oscillatory switching.

6.2 Directional dwell-time profiles

Let $q_{ij}(k)$denote the number of samples elapsed since the last structural transition. A requested transition $i \to j$is admissible when

$$q_{ij}(k) \geq N_{d,ij}. \qquad (6.4)$$

In Equation (6.4), $N_{d,ij}$is specified separately for QPSK-to-16-QAM and 16-QAM-to-QPSK transitions.

Five modes are used to separate the effects of dwell time and state projection.

Table 5. Experimental reconfiguration modes

Mode	Forward/reverse dwell	State transfer	Experimental purpose
D0	4/4	Stability-constrained contextual projection	Symmetric reference
D1	1/1	Stability-constrained contextual projection	Minimum theoretical profile
D2	2/1	Stability-constrained contextual projection	Throughput-oriented profile
D3	3/1	Stability-constrained contextual projection	BER-oriented profile
D4	3/1	Context-free ridge projection	Projection ablation for D3

All modes use the same cost, hysteresis, channel observations, and random realizations. Consequently, D0–D3 isolate the effect of directional dwell time, while D3–D4 isolate the effect of the contextual stability-constrained projection.

7. Experimental Evaluation

7.1 Scenarios and statistical protocol

The evaluation uses four channel and interference scenarios.

Table 6. Scenario parameters

Scenario	SNR correlation	Mean SNR, dB	Innovation std., dB	p_{on}	p_{off}	Interference loss, dB	Measurement std., dB	Projection shift

S0 nominal	0.96	10.0	0.70	0.035	0.18	5	0.35	0
S1 fast	0.85	10.0	1.20	0.050	0.20	5	0.50	0
S2 persistent	0.94	10.0	0.80	0.080	0.08	7	0.45	0
S3 shifted	0.92	9.5	1.00	0.060	0.12	6	0.80	0.20

Each experiment contains 200 paired Monte Carlo runs of 400 control samples, corresponding to 200 ms per run. Each control interval contains 128 directly demodulated symbols. Within one run, all modes use identical SNR, interference, transmitted bits, receiver noise, and state disturbances.

Normalized goodput is calculated as

$$G = \frac{1}{K}\sum_{k=1}^{K} \eta_{\sigma_k}(1 - \mathrm{BER}_k),\ \eta_Q = 1,\ \eta_{16} = 2. \qquad (7.1)$$

In Equation (7.1), $K = 400$and $\eta_i$represents relative uncoded spectral efficiency. The metric is a normalized comparison index rather than an absolute physical-layer bitrate.

Transient metrics are averaged over the first four samples after each structural transition. Paired t-tests and Wilcoxon signed-rank tests are applied to run-level differences. Benjamini–Hochberg correction is performed separately for each scenario and contrast. The achieved power for the primary mean-BER and transient-BER comparisons ranges from 0.81 to greater than 0.99 at the 0.05 significance level.

7.2 Effect of directional dwell time

Table 7 reports the symmetric reference D0 and the BER-oriented directional profile D3.

Table 7. Mean performance of symmetric and directional dwell-time profiles

Scenario	Mode	Mean BER	Peak BER	Normalized goodput	Switches per 200 ms	Q→16 transient BER
S0	D0, 4/4	0.001081	0.042627	1.734211	12.47	0.002709
	D3, 3/1	**0.001050**	**0.040986**	**1.736688**	12.99	**0.002295**
S1	D0, 4/4	0.001489	0.052959	1.728366	21.39	0.003291
	D3, 3/1	**0.001432**	**0.052373**	**1.728806**	22.76	**0.002947**
S2	D0, 4/4	0.004606	0.073711	**1.441104**	22.63	0.005963
	D3, 3/1	**0.004280**	**0.072061**	1.431654	23.70	**0.004302**
S3	D0, 4/4	0.003165	0.068066	**1.540036**	23.34	0.004983
	D3, 3/1	**0.002953**	**0.065576**	1.533965	24.51	**0.003914**

Reducing only the reverse dwell time allows an earlier transition from 16-QAM to the more robust QPSK mode. As a result, D3 reduces mean BER in all four scenarios. The advantage is largest under persistent interference and distribution shift.

Table 8. Paired D3–D0 comparison

Scenario	Mean BER change	BH-adjusted p	d_z	Q→16 transient BER change	BH-adjusted p	d_z
S0	−2.86%	0.0075	−0.216	−15.27%	0.0109	−0.201
S1	−3.86%	4.23×10^{-5}	−0.309	−10.46%	1.87×10^{-4}	−0.279
S2	−7.09%	2.82×10^{-35}	−1.105	−27.87%	2.67×10^{-27}	−0.911
S3	−6.72%	7.56×10^{-26}	−0.888	−21.44%	1.07×10^{-19}	−0.727

The BER reductions are accompanied by 4.17–6.40% more structural transitions. Goodput changes by $+0.143\%$, $+0.025\%$, -0.656%, and -0.394%in S0–S3, respectively. Therefore, D3 should be interpreted as a BER-oriented profile rather than a universally superior operating point.

The QPSK-to-16-QAM transient state RMSE increases by 1.96–4.37% relative to D0. This does not contradict the BER improvement: the directional rule changes the timing and number of transitions, while the projection operator remains unchanged. The result shows that improved configuration selection can reduce communication errors even when the average state mismatch immediately after switching does not decrease.

The absolute synchronization-loss fractions remain below $4.7 \times 10^{-4}$for every scenario and profile. D3 increases this fraction by at most 1.5×10^{-4}, which is small in absolute terms but prevents a claim that the 3/1 profile improves every reliability indicator.

7.3 Throughput-oriented 2/1 profile

The D2 profile provides a different operating point. Relative to symmetric 4/4 dwell time, D2 changes normalized goodput by $+0.449\%$in S0, $+0.628\%$in S1, -0.327%in S2, and $+0.080\%$in S3. Mean BER changes by -2.06%, -1.15%, -6.14%, and -5.67%, respectively.

These results support two engineering profiles:

- 2/1 for nominal and rapidly varying scenarios where normalized goodput is prioritized;
- 3/1 for persistent interference and distribution shift where lower mean and transient BER is prioritized.

This recommendation is scenario dependent and is not yet an automatic online profile-selection algorithm.

7.4 Contribution of the contextual projection

Under the same 3/1 dwell-time profile, D3 is compared with the context-free ridge mode D4. The contextual stability-constrained map reduces QPSK-to-16-QAM transient state RMSE by 1.44–2.47% and transient soft-state RMSE by 3.03–6.12%. Mean BER differences remain below 0.37%, and transient BER differences remain below 1.48%.

Accordingly, the contextual projection produces a clear improvement in state-space consistency, but its direct BER effect is small under the adopted demapper sensitivity. The larger BER improvements in Table 8 are attributable primarily to directional switching rather than to a change in the projection matrix. This distinction prevents the effects of state transfer and switching policy from being conflated.

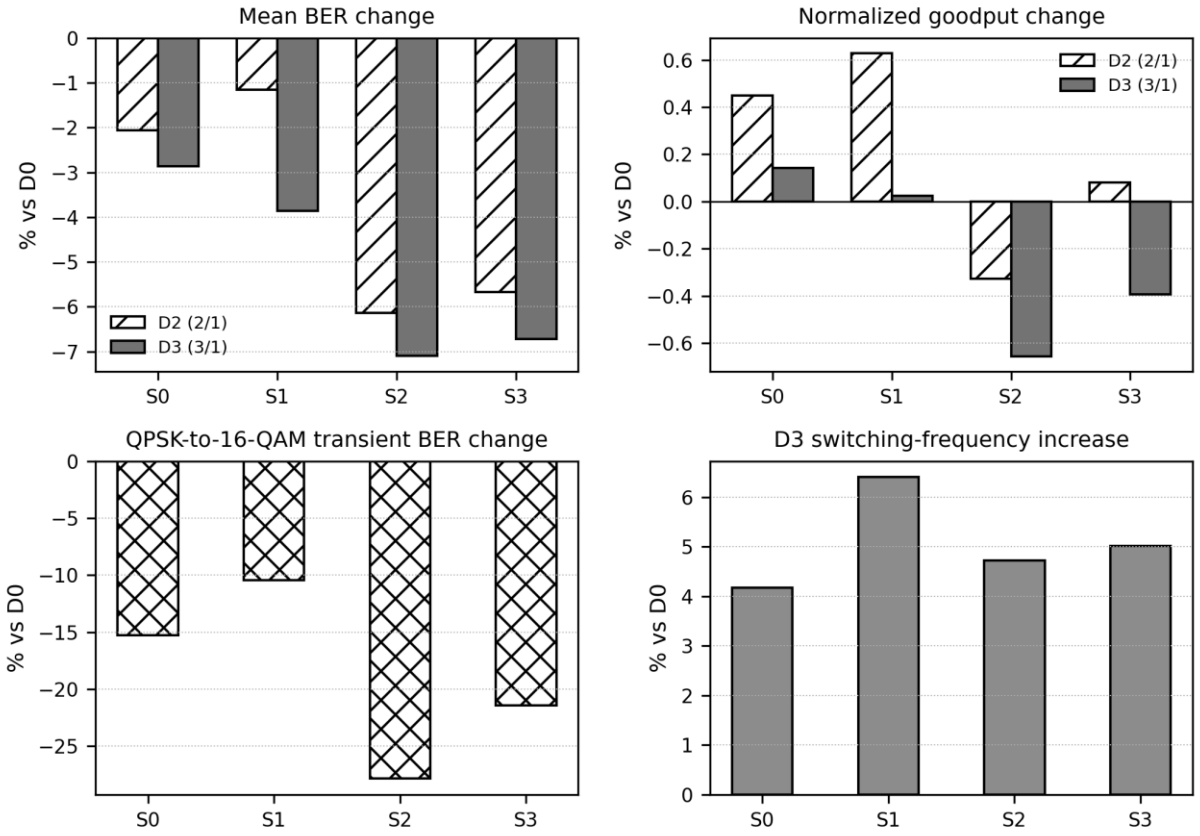


Fig. 5. Performance of directional dwell-time profiles relative to the symmetric reference.

Event-aligned recovery analysis shows that contextual transfer improves internal-state consistency beyond the immediate four-sample interval. Relative to context-free ridge transfer, state RMSE decreases by 1.54–2.54% and soft-state RMSE by 3.26–6.34% in samples 1–4, with smaller but still significant reductions of 0.37–0.56% and 0.83–1.31% in samples 5–10. The differences vanish in samples 11–20. No corresponding BER reduction is observed; BER differences become negligible after the first window. Thus, contextual transfer accelerates state-space recovery over approximately ten control samples, but the adopted demapper is insufficiently sensitive for this improvement to produce a communication-level gain.

8. Discussion

8.1 Interpretation of the validated contribution

The validation separates contextual reconstruction, energy-constrained identification, and directional dwell-time control. Context reduces independent mapping RMSE by 6.20% and the unconstrained forward gain by 25.82%; the stability penalty further reduces the gain to 1.2449 for a 0.91% RMSE increase.

Communication-level gains arise mainly from directional switching: D3 reduces mean BER by 2.86–7.09% and transient BER by 10.46–27.87%, but increases switching by 4.17–6.40% and lowers goodput in S2–S3. The result is therefore a controllable reliability–throughput trade-off.

Proposition 1 applies the established multiple-Lyapunov argument [13, 15–18]; the application-specific result is the projection-derived directional gain and its use in selecting dwell profiles.

The validated components serve different purposes. Contextual conditioning improves state-space consistency, whereas the stability penalty reduces the certified worst-case projection-energy gain without producing a separate significant BER improvement under the fixed 3/1 policy. Communication-level gains arise primarily from directional dwell-time control. Consequently, the projection should not be interpreted as a standalone BER optimizer or as a necessary enabler of the 3/1 profile. Its role is to preserve physically meaningful receiver state and provide a quantified switching-energy margin, while the dwell policy determines the observed BER–goodput trade-off.

8.2 O-RAN integration and timescale separation

The 0.5 ms interval belongs to the local SDR loop, not to E2 or a Near-RT RIC cycle. State estimation, projection, dwell enforcement, and structural switching must execute locally because the new receiver requires its internal state immediately.

In an O-RAN deployment, the functions are separated into two timescales:

1. The local SDR or O-DU/O-RU implementation estimates the receiver state, applies the projection, enforces the minimum dwell timer, and performs the structural transition.
2. A Near-RT RIC xApp may select the admissible configuration set, update cost weights, or choose between the 2/1 and 3/1 profiles using aggregated measurements reported over E2.

The internal five- or seven-dimensional state is not transported through E2. Only aggregated SNR, configuration, switching, transient-error, and dwell-profile indicators are exposed to supervisory control.

No new E2 service model is defined; O-RAN is treated as an integration context rather than a standards-compliant xApp implementation.

8.3 Why online adaptation was not retained

A causal recursive least-squares extension was evaluated under S3 with a noisy delayed target estimate and an energy-gain constraint. It increased transient RMSE by 8.83–14.69% and transient BER by 4.75–14.22%.

A safety gate either rarely activated or delayed valid 16-QAM transitions without measurable reliability benefit.

The fixed projection is therefore retained. Online adaptation would require persistent excitation, a validated target observer, uncertainty-aware acceptance, and a predictive risk criterion.

8.4 Physical interpretation and Bayesian relation

Replacing abstract soft states with sign- and amplitude-metric biases provides a direct link between transferred-state error and bit decisions, eliminating the previous empirical SNR penalty.

Amplitude-threshold states were approximately 1.7 times more influential than sign-metric states, but BER-weighted identification changed BER by less than 0.063%. Physical interpretation improves model validity without materially changing performance in the tested range.

Equation (4.12) is a maximum-a-posteriori coefficient estimate under Gaussian errors and priors, not an optimal Bayesian state estimator, which would require a full joint dynamical model.

8.5 Applicability and limitations

The results are conditional on stable linear within-mode dynamics, offline paired calibration states, QPSK/16-QAM dimensions of five and seven, and the declared baseband impairment model.

The simulator omits coding, timing acquisition, multipath estimation, AGC, and hardware impairments; 200-run Monte Carlo evidence does not replace hardware validation. With M configurations, pairwise identification may require M(M-1) maps, motivating a shared latent state or reference hub.

Extension to higher-order QAM does not alter the invariant synchronization subspace but increases the number of modulation-dependent threshold states. A direct pairwise implementation for Mconfigurations may require up to $M(M-1)$directed maps. A scalable alternative is to transfer all configurations through a common reference state or a shared latent demapper representation, reducing identification to $2M$directed maps. Verification for 64-QAM, 256-QAM, and OFDM receivers remains future work.

8.6 Practical implication

Implementation requires one small affine operation for the forward transfer, structural truncation in reverse, and a directional counter; no online optimization is required.

The 2/1 profile is recommended when normalized throughput is prioritized and 3/1 when BER under persistent interference or distribution shift is prioritized. Profile selection belongs to the slower supervisory timescale.

9. Conclusion

A context-aware state-transfer procedure was developed for structural reconfiguration between five-dimensional QPSK and seven-dimensional 16-QAM receiver models. The method preserves three physically invariant synchronization coordinates, reconstructs modulation-dependent demapper states from causal SNR context, and constrains the transfer through its Lyapunov-energy gain.

Independent validation showed a 6.20% mapping-RMSE reduction relative to context-free ridge regression. Stability-constrained identification reduced the forward gain to 1.2449 with a 0.91% RMSE increase. Applying the standard multiple-Lyapunov dwell-time argument revealed directional asymmetry and supported separate forward and reverse dwell intervals.

In four paired Monte Carlo scenarios, the 3/1 profile reduced mean BER by 2.86–7.09% and QPSK-to-16-QAM transient BER by 10.46–27.87% relative to symmetric 4/4 dwell time, at the cost of 4.17–6.40% more switching. The 2/1 profile provided the preferable throughput-oriented operating point. Hardware validation, coded and multipath reception, nonlinear acquisition, and scalable transfer among larger configuration libraries remain open tasks.

References

[1] O-RAN Alliance. O-RAN Architecture Description. Version 17.0, 2026. Available: https://www.o-ran.org/specifications.

[2] O-RAN Alliance. WG3: Near-Real-Time RAN Intelligent Controller and E2 Interface Workgroup. 2026. Available: https://www.o-ran.org/technical-groups/wg3.

[3] 3GPP. TR 38.843: Study on Artificial Intelligence (AI)/Machine Learning (ML) for NR Air Interface. Release 18, 2023. Available: https://portal.3gpp.org/desktopmodules/WorkItem/WorkItemDetails.aspx?workitemId=940084.

[4] Mitola, J. “The software radio architecture.” IEEE Communications Magazine, vol. 33, no. 5, pp. 26–38, 1995. DOI: 10.1109/35.393001.

[5] Jondral, F. K. “Software-defined radio—Basics and evolution to cognitive radio.” EURASIP Journal on Wireless Communications and Networking, vol. 2005, pp. 275–283, 2005. DOI: 10.1155/WCN.2005.275.

[6] Checko, A., Christiansen, H. L., Yan, Y., Scolari, L., Kardaras, G., Berger, M. S., and Dittmann, L. “Cloud RAN for mobile networks—A technology overview.” IEEE Communications Surveys & Tutorials, vol. 17, no. 1, pp. 405–426, 2015. DOI: 10.1109/COMST.2014.2355255.

[7] Pan, S. J., and Yang, Q. “A survey on transfer learning.” IEEE Transactions on Knowledge and Data Engineering, vol. 22, no. 10, pp. 1345–1359, 2010. DOI: 10.1109/TKDE.2009.191.

[8] Taylor, M. E., and Stone, P. “Transfer learning for reinforcement learning domains: A survey.” Journal of Machine Learning Research, vol. 10, pp. 1633–1685, 2009.

[9] Pan, S. J., Tsang, I. W., Kwok, J. T., and Yang, Q. “Domain adaptation via transfer component analysis.” IEEE Transactions on Neural Networks, vol. 22, no. 2, pp. 199–210, 2011. DOI: 10.1109/TNN.2010.2091281.

[10] Hanus, R., Kinnaert, M., and Henrotte, J.-L. “Conditioning technique, a general anti-windup and bumpless transfer method.” Automatica, vol. 23, no. 6, pp. 729–739, 1987. DOI: 10.1016/0005-1098(87)90029-X.

[11] Graebe, S. F., and Ahlén, A. L. B. “Dynamic transfer among alternative controllers and its relation to antiwindup controller design.” IEEE Transactions on Control Systems Technology, vol. 4, no. 1, pp. 92–99, 1996. DOI: 10.1109/87.481772.

[12] Zaccarian, L., Nešić, D., and Teel, A. R. “A common framework for anti-windup, bumpless transfer and reliable designs.” Automatica, vol. 38, no. 10, pp. 1735–1744, 2002. DOI: 10.1016/S0005-1098(02)00072-9.

[13] Branicky, M. S. “Multiple Lyapunov functions and other analysis tools for switched and hybrid systems.” IEEE Transactions on Automatic Control, vol. 43, no. 4, pp. 475–482, 1998. DOI: 10.1109/9.664150.

[14] Liberzon, D., and Morse, A. S. “Basic problems in stability and design of switched systems.” IEEE Control Systems Magazine, vol. 19, no. 5, pp. 59–70, 1999. DOI: 10.1109/37.793443.

[15] Hespanha, J. P., and Morse, A. S. “Stability of switched systems with average dwell-time.” In Proceedings of the 38th IEEE Conference on Decision and Control, pp. 2655–2660, 1999. DOI: 10.1109/CDC.1999.831330.

[16] Liberzon, D. Switching in Systems and Control. Boston: Birkhäuser, 2003. DOI: 10.1007/978-1-4612-0017-8.

[17] Chatterjee, D., and Liberzon, D. “Stability analysis of deterministic and stochastic switched systems via a comparison principle and multiple Lyapunov functions.” SIAM Journal on Control and Optimization, vol. 45, no. 1, pp. 174–206, 2006. DOI: 10.1137/040619429.

[18] Lin, H., and Antsaklis, P. J. “Stability and stabilizability of switched linear systems: A survey of recent results.” IEEE Transactions on Automatic Control, vol. 54, no. 2, pp. 308–322, 2009. DOI: 10.1109/TAC.2008.2012009.

[19] Gardner, F. M. Phaselock Techniques, 3rd ed. Hoboken, NJ: Wiley, 2005. DOI: 10.1002/0471732699.

[20] Meyr, H., Moeneclaey, M., and Fechtel, S. A. Digital Communication Receivers: Synchronization, Channel Estimation, and Signal Processing. New York: Wiley, 1998. DOI: 10.1002/0471200573.

[21] Popovski, P., Nielsen, J. J., Stefanović, Č., de Carvalho, E., Strom, E., Trillingsgaard, K. F., Bana, A.-S., Kim, D. M., Kotaba, R., Park, J., and Sørensen, R. B. “Wireless access for ultra-reliable low-latency communication: Principles and building blocks.” IEEE Network, vol. 32, no. 2, pp. 16–23, 2018. DOI: 10.1109/MNET.2018.1700258.