2024 European Control Conference (ECC)
June 25-28, 2024. Stockholm, Sweden

Entangled Gain and Phase Analysis for the Internet

Ding Zhang[1], Jiazheng Wang[2], Fan Zhang[2], Ioannis Lestas[3], Li Qiu[4]

***Abstract*— We conduct a local stability analysis of a class of Internet congestion control protocols, approaching the problem from a novel entangled gain and phase perspective. Our approach incorporates a recently revitalized quantitative MIMO phase concept, which we use to reexamine the local version of certain classical stability results. This phase information is entangled with the gain information through the concept of the Davis-Wielandt shell, providing an intuitive graphical interpretation and reducing the conservatism of stability conditions. Leveraging these tools, we derive decentralized stability conditions for the Internet protocols under consideration.**

I. Introduction

Transmission Control Protocol (TCP), one of the largest manmade feedback systems, is a crucial component of the modern Internet. They employ specific congestion control strategies to prevent or alleviate the congestion at network switches, thereby enabling low network delay and high throughput [1]. Despite the proliferation of numerous TCP protocols, their stability continues to be a concern.

A TCP network comprises many end-to-end connections. Each end-to-end connection, as illustrated in Fig. 1, consists of: (1) a source host, (2) a data flow generated by the source host, (3) a path, which is a sequence of consecutive communication links and network switches, through which the data flow is transmitted, and (4) a destination host for which the data flow is oriented. A typical TCP network deploys two groups of protocols, with one group implemented at the source hosts, and another implemented at the network switches. The host protocol adapts its packet sending rate to the aggregated congestion measure that it sees. Meanwhile, the switch protocol generates a certain form of congestion measure based on its arrival rate and/or queue length. The two groups of protocols, or dynamics in control terms, are interconnected according to the topology defined by the routing table, forming a large-scale TCP network.

The Network Utility Maximization (NUM) has served as a unified framework to interpret many folds of the equilibrium of the TCP dynamics [2], [3]. Roughly speaking, a large class of protocols can be regarded as distributed projected gradient descent algorithms, which solve certain types of convex optimization problem. Different protocols correspond to different utility functions that are maximized. The static properties of the network equilibrium, which coincides with the optimal solution, can be well understood from established optimization theory (see [4] for an expository treatment). When the network is delay-free, the global asymptotic stability of the network has been thoroughly investigated in [5], [6] under the passivity and dissipativity frameworks.

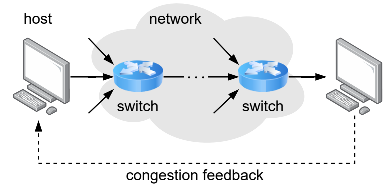

Fig. 1: An end-to-end connection in a TCP network

However, the ubiquitous network delays pose substantial challenge to global analysis while their impact on stability should not be underestimated [7]. In this context, we shall focus on local stability analysis of the TCP networks, which is sufficient for practical purposes. The aforementioned tools [5], [6] now become inadequate due to the severe phase lag at high frequency. In contrast, a purely gain-type condition, though capable of handling delays, may incur too much conservatism. Analyzing such systems requires tools beyond small gain and passivity. Many recent efforts have been dedicated to this direction. Lestas et al. [8], [9], [10] have studied some interesting spectrum-containing convex sets such as the S-hull and Davis-Wielandt shell. In particular, they reveal the equivalence between a class of quadratic separation and the separation of Davis-Wielandt shells (abbreviated as DW shell hereafter). Building on the equivalence, they propose several general stability conditions in the form of shell separation. The works [11], [12], [13] leverage a recently proposed quantitative MIMO phase measure [14], [15], [16]. They impose *either the gain or the phase condition* at each frequency to enforce generalized Nyquist criterion [17]. However, it is worth noting that the conditions proposed for power systems in [13] rely on the global topology information, which limits their applicability in network-related contexts. In [18], homotopy arguments are used to combine conditions arising from different network decompositions so as to develop broad classes of distributed stability conditions with a plug-and-play operation.

Expanding upon these works, this paper aims to contribute from the following aspects:

*This work was supported by Hong Kong Rsearch Grants Council under GRF 16201120. The work of D. Zhang was supported by Huawei-HKUST PhD Fellowship.

[1]D. Zhang is with Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, China dzhangbg@ust.hk

[2]J. Wang, F. Zhang are with Theory Lab, Huawei Hong Kong Research Center, Hong Kong, China {wang.jiazheng, zhang.fan2}@huawei.com

[3]I. Lestas is with the Department of Engineering, University of Cambridge, Trumpington Street, Cambridge, CB2 1PZ, United Kingdom icl20@cam.ac.uk

[4]L. Qiu is with Southern University of Science and Technology, Shen Zhen, China, and Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong, China eeqiu@ust.hk

- We exploit the loop structure of TCP networks, and precisely characterize the union of DW shells of a class of matrices. Based on this and the last item, we tailor the DW shell-based stability condition to a new distributed stability certificate for TCP networks.
- We emphasize graphically the connection between the gain, the phase, and the DW shell based analysis, both at the matrix level and in deriving stability conditions for TCP networks. This graphical understanding provides clear guidance for the stability vlidation and parameter tuning to achieve more robust network stability.

In the remaining sections of this paper, we first introduce the dynamical models that describe TCP networks and the MIMO phases in Section II. Then in Section III, we develop our main results centered around the DW shell, and we use these results to derive the local stability certificates for TCP networks in Section IV, whereas we conclude this paper in Section V. Due to the space constraint, the proofs will be included in the extended journal version of this paper.

Notation. Let $\imath = \sqrt{-1}$ be the imaginary unit. We denote by $\mathbb{C}$ and $\mathbb{R}$ the field of real and complex numbers, respectively. Then the positive reals and nonnegative reals are denoted by $\mathbb{R}_+$ and $\overline{\mathbb{R}}_+$. For an m-by-n matrix M, we use $M^\mathsf{T}, \overline{M}, M^\mathsf{H}$ to denote its transpose, conjugate, and conjugate transpose, respectively. The matrix $|M|$ refers to elementwise absolute value of M. A square matrix A is congruent to B, denoted by $A \simeq B$, if there exists nonsingular T such that $A = TBT^\mathsf{H}$; it is similar to B, denoted by $A \sim B$, if there exists nonsingular T such that $A = TBT^{-1}$. The spectrum of A is denoted by $\Lambda(A)$ and $\lambda_i(A)$ refers to a specific eigenvalue of A. The spectral raidus of A, denoted by $\rho(A)$, is the largest magnitude attained by elements of $\Lambda(A)$. The indicator $\mathbf{1}_\star$ takes the value 1 if $\star$ is true and the value 0 otherwise. When $|\cdot|$ is taken on a finite set, it refers to its cardinality.

II. Preliminaries

A. Flow Model and Linearization

Consider a TCP network comprising n hosts and m switches. Let $\mathcal{I} = \{1, \dots, n\}, \mathcal{J} = \{1, \dots, m\}$ denote the sets of indices of hosts and switches. We shall use indices i or j to refer to a host or a switch, respectively. The path of the data flow from host i is denoted by $\mathcal{J}_i$ while $\mathcal{I}_j$ denotes the set of hosts that the switch j serves.

The flow-level behavior of a TCP network can be described by a set of deterministic delay differential equations:

$$\chi_{g,i}\left(\frac{\mathrm{d}}{\mathrm{d}t}\right) x_i(t) = g_i(x_i(t), x_i(t-T_i), u_i(t)), \tag{1a}$$

$$\chi_{h,j}\left(\frac{\mathrm{d}}{\mathrm{d}t}\right) v_j(t) = h_j(v_j(t), y_j(t)), \tag{1b}$$

$$y_j(t) = \sum_{i \in \mathcal{I}_j} x_i(t - f_{ji}), \tag{1c}$$

$$u_i(t) = \sum_{j \in \mathcal{J}_i} v_j(t - b_{ij}), \tag{1d}$$

TABLE I: Typical congestion control protocols

Protocols	Congestion Measure	Local Type*		Ref.
		Host	Switch	
Reno/Droptail	Packet loss probability	type 1	type 0	[19]
RCP	Rate mismatch and queue size	type 0	type 1	[20]
BIC, CUBIC	Packet loss probability	type 1	type 0	[21]
High-throughput TCP with smoothed rate feedback	Packet marking probability	type 1	type 1	[22] [4, p. 90]

* Refer to Table II and its context for the definition.

for $i = 1, \dots, n$ and $j = 1, \dots, m$. In (1), both $g_i(\cdot)$ and $h_j(\cdot)$ are static nonlinearities. The $\frac{\mathrm{d}}{\mathrm{d}t}$ is the time derivative operator and $\chi_{\cdot,\cdot}(\frac{\mathrm{d}}{\mathrm{d}t})$ is a polynomial of $\frac{\mathrm{d}}{\mathrm{d}t}$. In this paper, we restrict this polynomial to have either degree 0 or 1. This already encompasses a large group of existing protocols, as listed by Table I. Signals x_i, y_j refer to the sending rate of host i and the arrival rate of switch j while v_j, u_i refer to the congestion measure at switch j and the aggregated congestion measure at host i. The scalar-valued signals will sometimes be stacked into vector-valued signals x, y, u, v. The delays f_{ji}, b_{ij} are the forward delay from host i to switch j and the backward delay from switch j to host i, respectively. Referring to [5, A. I] and [4, Sec. 3], we made the following assumptions:

- the congestion measure is additive during the aggregation, as already indicated by (1d);
- the equilibrium $(\hat{x}, \hat{y}, \hat{u}, \hat{v})$ of (1) is unique;
- the round-trip delay for each host i is a constant, i.e., $f_{ji} + b_{ij} = T_i$ for all $j \in \mathcal{J}_i$.

Then we linearize (1a) and (1b) around the unique equilibrium, which leads us to the local protocol dynamics of two types listed in Table II. As mentioned, we only consider TCP networks with host and switch protocols being either type-0 or type-1 in this paper.

TABLE II: Types of local dynamics

Types	0	1	
Local Dynamics	k	$\dfrac{k}{s+p}$	$k > 0, p > 0$

The g_i, h_j are now overloaded to denote transfer functions of linearized host and switch dynamics. We collect them into two diagonal transfer matrices: $G(s) := \operatorname{diag}\{g_i(s)\}$ and $H(s) := \operatorname{diag}\{h_j(s)\}$. Due to the practical compatibility constraint, the types of all hosts should be the same, and so are the types of all switches. However, their parameters may exhibit heterogeneity. The Laplace transforms of (1c) and (1d) are collected into a forward routing delay matrix $R_f(s)$ with $[R_f(s)]_{ji} = e^{-sf_{ji}} \mathbf{1}_{j \in \mathcal{J}_i}$ and a backward routing delay matrix $R_b(s)$ with $[R_b(s)]_{ij} = e^{-sb_{ij}} \mathbf{1}_{i \in \mathcal{I}_j}$. Let $R_0 = R_f(0)$, it follows from the definition that $R_b(0) = R_0^\mathsf{T}$. The matrix R_0 carries pure topology information of the network, and it necessarily has full row rank due to the uniqueness

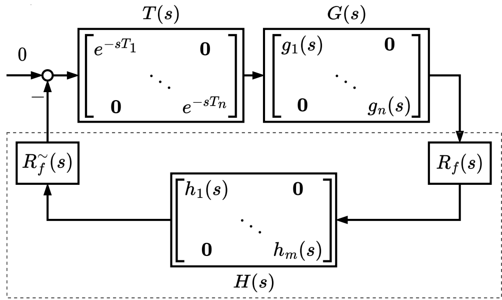

Fig. 2: System diagrams of linearized congestion control protocols around its equilibrium (negative feedback)

assumption on the network equilibrium. Under our previous constant round-trip delay assumption, we can factorize $R_b(s)$ as $R_b(s) = T(s)R_f(-s)^{\mathsf{T}}$ where $T(s) := \operatorname{diag}\{e^{-sT_i}\}$. The linearized TCP network is now depicted by Fig. 2, and we are interested in the closed-loop stability of this system, a.k.a., the local stability of the original nonlinear network. Before diving into this, we shall first introduce the MIMO phases [15], [14].

B. MIMO Phases

MIMO phases boil down to matrix phases at each frequency. The definition of the matrix phases relies on the concept of (angular) numerical range (see [23, Ch. 1] for an introductory overview). Let $L \in \mathbb{C}^{n\times n}$ be a square matrix, its numerical range $\mathbf{W}(L)$, defined as $\mathbf{W}(L) := \{z^{\mathsf{H}}Lz \mid z \in \mathbb{C}^n, z^{\mathsf{H}}z = 1\}$, is a convex and compact set in the complex plane. Meanwhile, by removing the unit norm constraint on the nonzero vector z, we obtain the *smallest* cone $\mathbf{W}'(L)$, namely its angular numerical range, that contains $\mathbf{W}(L)$, i.e., $\mathbf{W}'(L) := \{z^{\mathsf{H}}Lz \mid z \in \mathbb{C}^n, z \neq 0\}$. Note that the numerical range of matrix L contains its spectrum $\Lambda(L)$. If L is normal, then $\mathbf{W}(L)$ is exactly the convex hull of $\Lambda(L)$.

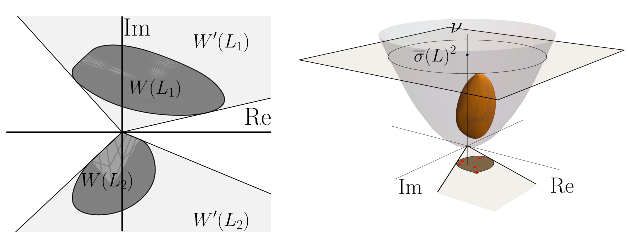

(a) Numerical range and angular numerical range: L_1 is sectorial, L_2 is quasisectorial but not sectorial

(b) DW shell of a matrix

Fig. 3: 2D and 3D convex sets associated with a matrix

Define the following classes of matrices:

Definition 1: A square matrix L is said to be quasisectorial if the closure of its $\mathbf{W}'(L)$ is pointed. Furthermore, it is said to be sectorial if $0 \notin \mathbf{W}(L)$.

Remark 1: Note that a convex cone is said to be *pointed* if it does not contain nontrivial subspaces. We also remark that a sectorial matrix is always nonsingular due to Def. 1 and the spectrum containment property of its numerical range.

A quasisectorial matrix L of rank r must be congruent to a direct sum of $\operatorname{diag}\{e^{\imath\alpha_1}, \ldots, e^{\imath\alpha_r}\}$ and a zero block of size $(n-r)$. Albeit being multi-valued, arguments $\alpha_1, \ldots, \alpha_r$ can always be selected such that their values lie whithin a closed interval of length less than π. We refer to any such choice of values as *a set of phases* for the matrix L. In each set of phases, we denote by $\underline{\phi}(L), \overline{\phi}(L)$ the smallest and the largest phases of L. It is important to note that matrix phases are unique modulo 2π and remain invariant under congruence. For system analysis, we should choose phases in a way such that the phase responses of the system are continuous functions of frequencies.

Geometrically, it is worth noting that the rays of $\overline{\phi}(L)$ and $\underline{\phi}(L)$ are the extreme rays of $\mathbf{W}'(L)$, as shown in Fig. 3a. The following simplified version of [14, Lem. 2.4] lays the foundation for the phase-relevant analysis.

Lemma 1: Let $L, M \in \mathbb{C}^{n\times n}$. Assume that L is quasisectorial and M is sectorial, then the product LM has r nonzero eigenvalues where $r = \operatorname{rank}(L)$. The arguments of these nonzero eigenvalues can be chosen such that

$$\underline{\phi}(L) + \underline{\phi}(M) \leq \angle\lambda_i(LM) \leq \overline{\phi}(L) + \overline{\phi}(M) \tag{2}$$

for $i = 1, \ldots, r$.

It follows from Lem. 1 that we can guarantee the nonsingularity of $I + LM$ by imposing the *small phase* condition: $\underline{\phi}(L) + \underline{\phi}(M) > -\pi$ and $\overline{\phi}(L) + \overline{\phi}(M) < \pi$.

III. Gain and Phase Entanglement via Davis-Wielandt Shell

A complex scalar z has its polar form $|z|e^{\imath\angle z}$, which elegantly encompassed both its gain and phase. Now that singular values of a matrix L carry its gain information and numerical range carries its phase information, how can we entangle both information? Our approach to answering this question is via the DW shell, which is a 3D convex set defined as follows:

$$\begin{aligned}\mathbf{DW}(L) := \mathbf{co}\{&(z^{\mathsf{H}}Lz, z^{\mathsf{H}}L^{\mathsf{H}}Lz) \mid \\ &z \in \mathbb{C}^n, z^{\mathsf{H}}z = 1\} \subseteq \mathbb{C} \times \overline{\mathbb{R}}_+\end{aligned} \tag{3}$$

where $\mathbf{co}(\cdot)$ denotes the convex hull operation. Fig. 3b shows the DW shell of a sectorial matrix. Note that the projection of $\mathbf{DW}(L)$ onto the complex plane is exactly $\mathbf{W}(L)$, while its third coordinate takes value in $[\underline{\sigma}(L)^2, \overline{\sigma}(L)^2]$. Meanwhile, all DW shells are lower bounded by the paraboloid $\{(s, |s|^2) \mid s \in \mathbb{C}\}$. A point $(s, z) \in \mathbf{DW}(L)$ is on the paraboloid iff $s \in \Lambda(L)$ [24]. In our context, the most important implication comes from the separation of DW shells of two matrices:

Lemma 2 ([25, Thm. 2.1]): Let $L, M \in \mathbb{C}^{n\times n}$, then $L + V^{\mathsf{H}}MV$ is nonsingular for all unitary V if and only if $\mathbf{DW}(-L) \cap \mathbf{DW}(M) = \emptyset$.

Lem. 2 establishes the equivalence between DW shell separation and the nonsingularity of the so-called *unitary orbit* around matrices L and M. This is useful when dealing with unitarily invariant[1] uncertainties. The analysis of closed-loop stability with respect to these uncertainties requires the characterization of the union of DW shells of the uncertainty set to be as precise as possible.

Recall the diagram in Fig. 2, where we break the loop at the host input. At each frequency, the lower part $R_f^{\mathsf{H}} H R_f$, though not necessarily normal, exhibits the special structure resembling unitoid matrices[2]. To handle matrices with such structure, we characterize the union of DW shells of a special class of matrices as follows:

Proposition 1: Let $M \in \mathbb{C}^{n\times n}$ be normal, and r be a positive integer, then

$$\begin{aligned}&\bigcup_{\Gamma\in\mathbb{C}^{n\times r},\|\Gamma\|_2\le 1} \mathbf{DW}(\Gamma^{\mathsf{H}} M\Gamma)\\ &= \bigcup_{s\in\mathbf{co}\,(\{0\}\cup\Lambda(M))} \left\{(s,t)\mid |s|^2\le t\le \max_{(s,l)\in\mathcal{P}(M)} l\right\}\end{aligned} \tag{4}$$

where

$$\mathcal{P}(M)=\begin{cases}\{(\lambda,|\lambda|^2)\mid \lambda\in\mathbf{co}\,(\{0\}\cup\Lambda(M))\}, & r=1,\\ \mathbf{co}\,\left(\{(\lambda,|\lambda|^2)\mid\lambda\in\{0\}\cup\Lambda\,(M)\}\right), & r>1.\end{cases}$$

Remark 2: When $r=1$, $\Gamma^{\mathsf{H}} M\Gamma$ is just a scalar, and $\mathcal{P}(M)$ is obtained by lifting each point in the convex hull of the origin and the spectrum of M to the paraboloid bound. When $r\ge 2$, $\mathcal{P}(M)$ is obtained by first lifting the origin and the spectrum of M to the paraboloid bound and then taking the convex hull. Subsequently, the DW shell union in Prop. 1 can be obtained by vertically sweeping $\mathcal{P}(M)$ down to the paraboloid bound.

By considering Prop. 1 and Lem. 2 jointly, we have

Theorem 2: Let $L\in\mathbb{C}^{n\times n}, M\in\mathbb{C}^{m\times m}$ be normal matrices, the following statements are equivalent:

1) $L+\Gamma^{\mathsf{H}} M\Gamma$ is nonsingular for all Γ such that $\|\Gamma\|_2\le 1$;
2) There exists a closed disc $\mathbf{D}$ such that $\{0\}\cup\Lambda\,(M)\subseteq\mathbf{D}$ and $\Lambda\,(-L)\subseteq\mathbb{C}\backslash\mathbf{D}$.

IV. Local Stability Analysis for TCP Networks

We are now all set to study the local stability of TCP networks using tools from the previous section. All forthcoming conditions rely on the generalized Nyquist criterion [17].

Here, the TCP network under consideration is locally stable if the eigenloci of the loop transfer function $GTR_f^{\sim} HR_f$ do not encircle $-1+\imath 0$ point[3]. This condition holds if

$$\det(\mu I + GTR_f^{\mathsf{H}} HR_f)\ne 0 \tag{5}$$

for all $\omega\in[0,\infty], \mu\in[1,\infty]$. Since G,H are either type-0 or type-1, $GTR_f^{\mathsf{H}} HR_f$ is bounded for all frequencies, and (5) always holds for all $\omega\in[0,\infty]$ when $\mu=\infty$. Meanwhile, the loop transfer function $GTR_f^{\mathsf{H}} HR_f$ has non-negative eigenvalues at $\omega=0$ or ∞, which implies that (5) holds for all $\mu\in[1,\infty]$ at these frequencies. Therefore, by letting $\zeta=1/\mu$, (5) can be rewritten as

$$\det((GT)^{-1}+\zeta R_f^{\mathsf{H}} HR_f)\ne 0 \tag{6}$$

for all $\omega\in(0,\infty),\zeta\in(0,1]$. Next, we define a diagonal matrix $D:=\operatorname{diag}\left\{\sum_{j\in\mathcal{J}_i}|\mathcal{I}_j||h_j|; i=1,\ldots,n\right\}$, and let $\tilde{R}_f=|H|^{1/2}R_fD^{-1/2}$. Then (6) is equivalent to

$$\det((DGT)^{-1}+\zeta\tilde{R}_f^{\mathsf{H}}\operatorname{diag}\left\{e^{\imath\angle h_j}; j=1,\ldots,m\right\}\tilde{R}_f)\ne 0$$

for all $\omega\in(0,\infty),\zeta\in(0,1]$. Note that $\|\tilde{R}_f\|_2\le 1$ since

$$\begin{aligned}\|\tilde{R}_f\|_2 &= \rho(\begin{bmatrix}0 & |H|^{1/2}R_fD^{-1/2}\\ D^{-1/2}R_f^{\mathsf{H}}|H|^{1/2} & 0\end{bmatrix})\\ &= \rho(\begin{bmatrix}0 & \operatorname{diag}\{|\mathcal{I}_j|\}^{-1}R_f\\ D^{-1}R_f^{\mathsf{H}}|H|\operatorname{diag}\{|\mathcal{I}_j|\} & 0\end{bmatrix})\\ &\le \|\begin{bmatrix}0 & \operatorname{diag}\{|\mathcal{I}_j|\}^{-1}R_f\\ D^{-1}R_f^{\mathsf{H}}|H|\operatorname{diag}\{|\mathcal{I}_j|\} & 0\end{bmatrix}\|_\infty\le 1.\end{aligned}$$

Except the unit gain bound, we assume that other information of $\tilde{R}_f$ is unknown. The parameter ζ can be absorbed into this uncertain part as well. Then the TCP network is stable if

$$\det((DGT)^{-1}+\Gamma^{\mathsf{H}}\operatorname{diag}\left\{e^{\imath\angle h_j}; j=1,\ldots,m\right\}\Gamma)\ne 0 \tag{7}$$

for all $\omega\in(0,\infty),\|\Gamma\|_2\le 1$.

A. Small Gain, Small Phase, and Crossover Frequencies

Observe that $\|\Gamma^{\mathsf{H}}\operatorname{diag}\left\{e^{\imath\angle h_j}; j=1,\ldots,m\right\}\Gamma\|_2\le 1$, and it suffices to guarantee (7) at ω by imposing the small gain condition $\|DGT(\imath\omega)\|_2<1$, or more explicitly, $g_i(i\omega)\sum_{j\in\mathcal{J}_i}|\mathcal{I}_j||h_j(\imath\omega)|<1$ for all i. Note that, for the hosts considered in this paper, $\|DGT\|_2$ is monotonically decreasing w.r.t. ω, and we define a gain crossover frequency

$$\omega_1 := \inf\{\omega\in\mathbb{R}_+ \mid |g_i(\imath\omega)|\sum_{j\in\mathcal{J}_i}|\mathcal{I}_j||h_j(\imath\omega)|<1\,\forall i\}.$$

The network complies to the small gain condition over the high frequency range $(\omega_1,\infty]$.

From the phase perspective, note that

$$\begin{aligned}&\mathbf{cl}\,\left(\mathbf{W}'(\Gamma^*\operatorname{diag}\left\{e^{\imath\angle h_j}; j=1,\ldots,m\right\}\Gamma)\right)\\ &\qquad\subseteq\mathbf{cl}\,\left(\mathbf{W}'(\operatorname{diag}\left\{e^{\imath\angle h_j}; j=1,\ldots,m\right\})\right)\end{aligned}$$

where $\mathbf{cl}\,(\cdot)$ is the closure operation. Hence the matrix $\Gamma^{\mathsf{H}}\operatorname{diag}\left\{e^{\imath\angle h_j(\imath\omega)}; j=1,\ldots,m\right\}\Gamma$ is frequencywise quasi-sectorial. Its phases are zero for all $\omega\in(0,\infty)$ if H is type-0. If H is type-1, its phases at $\omega\in(0,\infty)$ are contained in the interval $[-\arctan\frac{\omega}{\underline{p}_h}, -\arctan\frac{\omega}{\overline{p}_h}]\subseteq(-\frac{\pi}{2},\frac{\pi}{2})$ where $\underline{p}_h,\overline{p}_h$ denote the smallest and the largest absolute values of switch poles. The phase responses of DG are similar. When ignoring the delays, i.e., $T=I$, it is obvious that the small phase condition is met for all $\omega\in(0,\infty)$ regardless of the types of host and switch protocols. This resembles the passivity-based results. Yet the round-trip delay T introduces severe phase lag to the loop at high frequency, we define a phase crossover frequency

$$\omega_{-\pi} := \sup\{\omega\in\mathbb{R}_+ \mid \underline{\phi}\,(GT(\imath\omega))+\underline{\phi}\,(H(\imath\omega))>-\pi\}.$$

[1]By "unitarily invariant", we mean the uncertainty set is invariant under unitary congruence.

[2]Matrices that are diagonalizable via congruence.

[3]For notational simplicity, we will suppress the dependency on the generalized frequency s or frequency $\imath\omega$, which is clear from the context.

The network satisfies the small phase condition and holds infinitely large gain margin over $[0, \omega_{-\pi})$.

Requiring that $\omega_{-\pi} > \omega_1$ suffices to ensure the stability of the network. Such condition is essentially the same as the mixed gain and phase conditions proposed in [12], [13]. The DW shell-based gain and phase entanglement enables us to handle the cases where $\omega_{-\pi} \leq \omega_1$.

B. Entangled Gain and Phase Condition

Applying Thm. 2 to (7) gives the following proposition:

Proposition 3: A TCP network is locally stable if for each $\omega \in (0, \infty)$, there exists a closed disc such that $\{0\} \cup \{e^{i\angle h_j(\imath\omega)} \mid j = 1, \ldots, m\}$ is contained in the disc while $\{-(g_i(\imath\omega)e^{-\imath\omega T_i} \sum_{j \in \mathcal{J}_i} |\mathcal{I}_j||h_j(\imath\omega)|)^{-1} \mid i = 1, \ldots, n\}$ is outside the same disc.

Under the mapping $-1/z$, the stability condition above can be restated as $g_i(\imath\omega)e^{-\imath\omega T_i} \sum_{j \in \mathcal{J}_i} |\mathcal{I}_j||h_j(\imath\omega)|, i = 1, \ldots, n$ lie inside a uniform open disc that is contained in the feasible region specified by the switch protocols. For networks adopting type-0 switches, this feasible region is the complex plane excluding the real interval $[-\infty, -1]$, as depicted in the left plot of Fig. 4. If type-1 switches are used instead, an approximation of this feasible region is

$$\mathbb{C} \backslash \{z \mid |z| \geq 1, \angle z \in [-\pi - \overline{\phi}(H(\imath\omega)), -\pi - \underline{\phi}(H(\imath\omega))]\},$$

as shown in the right plot of Fig. 4. Note that this approximation introduces no extra conservatism considering the monotonicity of gain and phase functions of scaled hosts.

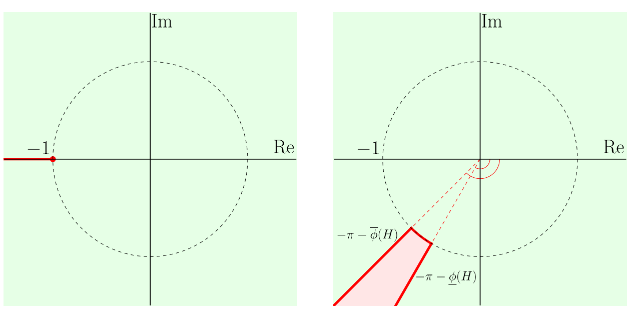

Fig. 4: Feasible regions (shaded in green) for the host part DGT when the switches are type-0 (left) or type-1 (right)

These feasible regions for hosts can be viewed as the union of feasible regions given by the small gain condition and the small phase condition. In other words, our proposed condition requires that each scaled host with round-trip delay satisfies either the small gain or the small phase condition. Additionally, these hosts should be bounded inside a uniform disc contained in the feasible region. This disc trivially exists if hosts all comply to the small gain condition, or all comply to the small phase condition. Hence, we just need to examine the condition in Prop. 3 over the finite frequency range $[\omega_{-\pi}, \omega_1]$, provided these two frequencies can be estimated.

We also remark that the worst-case phases among all switches are used to construct the feasible region for hosts. This indeed leads to potential conservatism. Nevertheless, our currently proposed condition is suitable for cases where switches exhibiting similar phase responses. Clearly, type-0 switches fall into this category. For type-1 switches, their phase responses being similar means that their poles should be close to each other. Switch poles are usually linked to switch capacities and/or weights in the low-pass filter. The former suggests the condition is more effective for the stability analysis of local area networks or a single layer in a hierarchical network. Meanwhile, the weights are usually set within a suggested range and will not be too diverse.

Next, we demonstrate the efficacy of our proposed condition via a numerical experiment.

Example 1 (High-throughput/Scalable TCP): Consider a network implementing the high-throughput TCP with a smoothed rate feedback as described in [4, Sec. 6.1.3] and [26]. The high-throughput TCP adopts a more aggressive host control strategy than the conventional Reno protocol so as to shorten the recovery time on the detection of packet loss and increase the network throughput. It has the following flow-level description:

$$\begin{aligned} \text{host:} \quad & \dot{x}_i(t) = \kappa_i x_i(t - T_i)(w_i - x_i(t)u_i(t)), \\ \text{switch:} \quad & \epsilon_j \dot{z}_j(t) = -z_j(t) + y_j(t), \\ & v_j(t) = (z_j(t)/c_j)^{B_j}, \end{aligned}$$

in which the switch protocol cascades a low-pass filter with a static M/M/1/B packet loss probability function, and c_j, B_j refer to the capacity and the buffer size of switch j. Meanwhile $\kappa_i, w_i, \epsilon_j$ are tunable network parameters. Given the equilibrium $\hat{x}_i, \hat{u}_i, \hat{y}_j$, the local dynamics of network hosts and switches have the forms $g_i(s) = \frac{\kappa_i w_i \hat{x}_i/\hat{u}_i}{s + \kappa_i \hat{x}_i \hat{u}_i}$ and $h_j(s) = \frac{k_j/\epsilon_j}{s + 1/\epsilon_j}$ where $k_j = \frac{\mathrm{d}v_j}{\mathrm{d}z_j}(\hat{y}_j)$. Now consider a high-throughput TCP network with 6 hosts severed by 4 switches. Referring to the network literature, e.g., [26], the packet size is set to 1000 bytes. For the hosts, the parameter w_i is set to 1 for all $i = 1, \ldots, 6$ while κ_i is set to 0.15 for $i = 1, 2, 5, 6$ and 0.3 for $i = 3, 4$. Their constant round trip times are assumed to be $15, 10, 30, 27, 15$ and $30\,\text{ms}$, respectively. For switches, we set ϵ_j to be $1.1, 1.1, 1.4$ and 1.5 for $j = 1, 2, 3$ and 4. Their capacities (and buffer sizes) are assumed to be 1000 Mbps (100 KB), 800 Mbps (300 KB), 800 Mbps (200 KB) and 1000 Mbps (100 KB). The hosts and switches are interconnected over the network topology defined by the following routing matrix:

$$R_0 = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{bmatrix}.$$

The forward delays and backward delays are chosen randomly under the constraints of positivity and constant round trip time. Their values are not critical as the stability criteria discussed herein are robust against them. We implemented the nonlinear flow-level dynamics of the network using MATLAB&Simulink and simulated the system for this specific network configuration. The decentralized condition in [4, Thm. 6.5] stipulates $(\kappa_i w_i T_i \sum_{j \in \mathcal{J}_i} k_j \hat{y}_j)/\hat{u}_i < 1$ for all

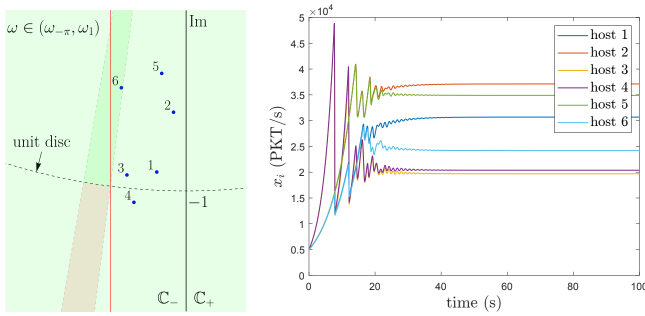

Fig. 5: Left: the frequency snapshot of scaled hosts at some ω in the interval $(\omega_{-\pi}, \omega_1)$; Right: host rate output of the flow-level simulation

$i = 1, \ldots, 6$ to ensure network stability. However, for our case, the left hand sides for all hosts are evaluated to be

$$0.4382 \quad 0.2475 \quad \mathbf{1.4473} \quad \mathbf{1.3197} \quad 0.3625 \quad 0.9767,$$

showing the third and the fourth hosts violate the condition. Furthermore, we observed that $\omega_{-\pi} = 5.2058 < \omega_1 = 5.3155$. For each $\omega \in [\omega_{-\pi}, \omega_1]$, as depicted in the left plot of Fig. 5, the conditions as derived in [13] cannot be met since the scaled hosts fail to uniformly satisfy the small gain or small phase condition. In contrast, our proposed condition is met. In particular, the scaled hosts fall inside the right half plane defined by the vertical line with real part being $\cos(-\pi - \underline{\phi}(H(\imath\omega)))$. This leads us to conclude that our network is locally stable. The simulation result as shown in the right plot of Fig. 5 verified our conclusion.

The above example shows that our proposed condition outperforms the existing ones for certain network configuration. In addition, it is not hard to verify that, at a fixed frequency, the condition in [4, Thm. 6.5] ensures that Eqa. (6) holds. Hence, our proposed condition can be naturally integrated with the existing one in a frequencywise way, i.e., one just has to ensure one of the two conditions to hold at each positive frequency so as to certificate the network stability.

V. CONCLUSIONS

We have introduced a decentralized, entangled-gain-and-phase condition that ensures the local stability of a specific class of TCP networks. The condition leverages both the gain and the phase information at each frequency, thus becoming less conservative than those obtained by frequencywise application of either small gain or small phase conditions. In obtaining the condition, we have studied the properties of DW shells, which integrate the concepts of the MIMO gains and phases. In particular, we have precisely characterized the union of DW shells of a class of matrices, whose structure echoes the loop structure of typical TCP networks. The non-conservatism of the proposed condition was validated through a numerical example of a High-throughput TCP network incorporating a smoothed rate feedback.

REFERENCES

[1] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-down Approach*, 8th ed. Hoboken: Pearson, 2021.

[2] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, no. 3, pp. 237–252, Mar. 1998.

[3] S. H. Low, F. Paganini, and J. C. Doyle, "Internet congestion control," *IEEE Control Systems Magazine*, vol. 22, no. 1, pp. 28–43, Feb. 2002.

[4] R. Srikant, *The Mathematics of Internet Congestion Control*, ser. Systems & Control: Foundations & Applications. Birkhäuser Basel, 2004.

[5] J. T. Wen and M. Arcak, "A unifying passivity framework for network flow control," *IEEE Transactions on Automatic Control*, vol. 49, no. 2, pp. 162–174, Feb. 2004.

[6] M. Arcak, "Compositional Design and Verification of Large-Scale Systems Using Dissipativity Theory: Determining Stability and Performance from Subsystem Properties and Interconnection Structures," *IEEE Control Systems Magazine*, vol. 42, no. 2, pp. 51–62, Apr. 2022.

[7] R. Johari and D. K. H. Tan, "End-to-End Congestion Control for the Internet: Delays and Stability," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 818–832, Dec. 2001.

[8] I. Lestas and G. Vinnicombe, "Heterogeneity and scalability in group agreement protocols: Beyond small gain and passivity approaches," *Automatica*, vol. 46, no. 7, pp. 1141–1151, July 2010.

[9] I. Lestas, "On network stability, graph separation, interconnection structure and convex shells," in *2011 50th IEEE Conference on Decision and Control and European Control Conference*, Dec. 2011, pp. 4257–4263.

[10] ——, "Large Scale Heterogeneous Networks, the Davis-Wielandt Shell, and Graph Separation," *SIAM Journal on Control and Optimization*, vol. 50, no. 4, pp. 1753–1774, Jan. 2012.

[11] D. Zhao, W. Chen, and L. Qiu, "When Small Gain Meets Small Phase," *arXiv:2201.06041 [eess.SY]*, Jan. 2022.

[12] L. Woolcock and R. Schmid, "Mixed Gain/Phase Robustness Criterion for Structured Perturbations with an application to Power System Stability," *IEEE Control Systems Letters*, pp. 1–1, 2023.

[13] L. Huang, D. Wang, X. Wang, H. Xin, P. Ju, K. H. Johansson, and F. Dörfler, "Gain and Phase: Decentralized Stability Conditions for Power Electronics-Dominated Power Systems," *IEEE Transactions on Power Systems*, pp. 1–16, 2024.

[14] W. Chen, D. Wang, S. Z. Khong, and L. Qiu, "A phase theory of MIMO LTI systems," *SIAM Journal on Control and Optimization*, 2024 (in press).

[15] D. Wang, W. Chen, S. Z. Khong, and L. Qiu, "On the phases of a complex matrix," *Linear Algebra and its Applications*, vol. 593, pp. 152–179, May 2020.

[16] X. Mao, W. Chen, and L. Qiu, "Phases of discrete-time LTI multivariable systems," *Automatica*, vol. 142, p. 110311, Aug. 2022.

[17] C. A. Desoer and Y. T. Wang, "On the generalized Nyquist stability criterion," *IEEE Transactions on Automatic Control*, vol. 25, no. 2, pp. 187–196, Apr. 1980.

[18] K. Laib, J. Watson, Y. Ojo, and I. Lestas, "Decentralized stability conditions for DC microgrids: Beyond passivity approaches," *Automatica*, vol. 149, p. 110705, Mar. 2023.

[19] V. Jacobson, "Congestion avoidance and control," *ACM SIGCOMM Computer Communication Review*, vol. 18, no. 4, pp. 314–329, Aug. 1988.

[20] N. Dukkipati, N. McKeown, and A. G. Fraser, "RCP-AC: Congestion Control to Make Flows Complete Quickly in Any Environment," in *IEEE INFOCOM 2006*, Apr. 2006, pp. 1–5.

[21] L. Xu, H. K., and R. Injong, "Binary increase congestion control (BIC) for fast long-distance networks," in *IEEE INFOCOM 2004*, vol. 4, Mar. 2004, pp. 2514–2524.

[22] G. Vinnicombe, "On the Stability of Networks Operating TCP-Like Congestion Control," *IFAC Proceedings Volumes*, vol. 35, no. 1, pp. 217–222, Jan. 2002.

[23] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*. Cambridge University Press, June 1994.

[24] C. K. Li, Y. T. Poon, and N. S. Sze, "Davis-Wielandt shells of operators," *Operators and Matrices*, no. 3, pp. 341–355, 2008.

[25] ——, "Eigenvalues of the Sum of Matrices from Unitary Similarity Orbits," *SIAM Journal on Matrix Analysis and Applications*, vol. 30, no. 2, pp. 560–581, Jan. 2008.

[26] T. Kelly, "Scalable TCP: Improving performance in highspeed wide area networks," *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 2, pp. 83–91, Apr. 2003.