



Graph-theoretic constraints on vesicle traffic networks

SOMYA MANI¹, KESAV KRISHNAN² and MUKUND THATTAI^{3*} 

¹Center for Soft and Living Matter, Institute for Basic Science, Ulsan, South Korea

²Indian Statistical Institute, Bengaluru, India

³Simons Centre for the Study of Living Machines, NCBS-TIFR, Bengaluru, India

*Corresponding author (Email, thattai@ncbs.res.in)

MS received 18 September 2021; accepted 24 December 2021

Eukaryotic cells use small membrane-enclosed vesicles to transport molecular cargo between intracellular compartments. Interactions between molecules on vesicles and compartments determine the source and target compartment of each vesicle type. The set of compartment and vesicle types in a cell define the nodes and edges of a transport graph known as the vesicle traffic network. The transmembrane SNARE proteins that regulate vesicle fusion to target compartments travel in cycles through the transport graph, but the paths they follow must be tightly regulated to avoid aberrant vesicle fusion. Here we use graph-theoretic ideas to understand how such molecular constraints place constraints on the structure of the transport graph. We identify edge connectivity (the minimum number of edges that must be removed to disconnect a graph) as a key determinant that separates allowed and disallowed types of transport graphs. As we increase the flexibility of molecular regulation, the required edge connectivity decreases, so more types of vesicle transport graphs are allowed. These results can be used to aid the discovery of new modes of molecular regulation and new vesicle traffic pathways.

Keywords. SNARE recycling; graph connectivity; self-organization

1. Introduction

Membrane-enclosed intracellular compartments, such as the ER, Golgi apparatus, endosomes and vacuoles, are a defining feature of all eukaryotic cells. Small membrane-enclosed carriers called vesicles bud from specific source compartments and fuse to specific target compartments, thereby transporting molecular cargo. A cell contains many different types of compartments and vesicles, each characterized by a molecular composition comprising specific proteins and lipids (Bonifacino and Glick 2004; Cai *et al.* 2007). The source and target compartments of each vesicle type are determined by specific interactions between these molecules (Ramadas and Thattai 2013; Mani and Thattai 2016a). This defines a self-organized, directed transport graph known as the vesicle traffic network.

Biochemical constraints restrict how various molecules can move along the transport graph. Membrane-integral molecules such as lipids and transmembrane proteins, as well as luminal cargo, can move only via vesicles (Bonifacino and Glick 2004; Cai *et al.* 2007). Membrane-peripheral molecules, such Rab and Arf GTPases, can exchange between membranes and the cytosol (Stenmark 2009; Sztul *et al.* 2019). The transmembrane SNARE proteins, which are required for the fusion of vesicles to compartments, are particularly constrained (Jahn and Scheller 2006). The fusion of some vesicle type to its desired target compartment occurs only when a specific v-SNARE on the vesicle physically entwines with a specific cognate t-SNARE on the target. The same v-SNARE is then returned to the vesicle's source compartment, to drive another round of fusion of that vesicle type (Südhof and Rothman 2009). v-SNAREs must therefore travel in a cycle through the transport graph, via successive vesicle carriers, yet only participate in the fusion of one specific vesicle type along that path. Since even a

This article is part of the Topical Collection: Emergent dynamics of biological networks.

single copy of a SNARE protein has the capacity to drive vesicle fusion (Van Den Bogaart *et al.* 2010), the activity of v-SNAREs along such cyclic paths must be tightly regulated (Südhof and Rothman 2009). As a result, only certain types of transport graphs are allowed (Shukla *et al.* 2017; Bhattacharyya *et al.* 2019).

Here we rigorously demonstrate how constraints on local molecular interactions lead to constraints on the global structure of the transport graph. We show that a graph-theoretic feature known as edge connectivity (Bang-Jensen and Gutin 2008, Chapter 7) determines which vesicle transport graphs are allowed under various SNARE regulation scenarios. In previous work, we used computational model checking to examine all possible transport graphs containing up to ten compartments (Shukla *et al.* 2017; Bhattacharyya *et al.* 2019). We found that, as molecular regulation became more flexible, the transport graphs became less constrained and required lower edge connectivity. We now give theoretical explanations for these computational results. Our analysis serves as a bridge between biochemistry and cell biology, and can be used to assess the completeness of the vesicle traffic maps that have been experimentally determined for various eukaryotic species (Mani and Thattai 2016; Bhattacharyya *et al.* 2019).

2. Results

2.1 Edge connectivity for directed graphs

Our approach is based on the connectivity properties of directed graphs (Bang-Jensen and Gutin 2008, Chapter 7). We consistently use the term ‘edge’ to mean ‘directed edge’ (since all the graphs we consider are directed); and we consistently use the term ‘nodes’ to mean ‘vertices’ (to avoid any confusion with ‘vesicles’). A directed graph is said to be strongly connected if, for every pair of nodes x and y , there is a directed path from x to y and a directed path from y to x . A directed cycle is a directed path from a node to itself. A directed simple cycle is a directed cycle in which no node repeats except for the first and last nodes. A directed graph is k -edge-strongly connected if it remains strongly connected whenever fewer than k edges are removed. A graph that is k -edge-strongly connected is clearly k' -edge-strongly connected for every $k' \leq k$. The edge-strong connectivity of a directed graph is the largest k for which it is k -edge-strongly connected. A


graph cut partitions the nodes of a graph into two disjoint sets, and a k -cut is a graph cut across which exactly k edges cross. For a directed graph that is k -edge-strongly connected, every graph cut must be at least a $2k$ -cut, with at least k edges crossing in each direction. Graph cycles and graph connectivity are intimately related (Methods).

2.2 Composition and transport graphs

We use two distinct but related representations of a vesicle traffic network (figure 1A). The *composition* representation captures the biochemical view: it specifies which types of vesicles bud from and fuse to which types of compartments. The *transport* representation captures the cell-biological view: it specifies the vesicle-mediated transport fluxes that couple compartments at distinct cellular locations.

The composition graph is a directed bipartite graph $G_{CO} = \{N_C \cup N_V, E_B \cup E_F\}$. Each node $x \in N_C$ is a compartment type, each node $\mu \in N_V$ is a vesicle type. Budding edges $E_B \subseteq \{(x, \mu)\}$ show the source compartment types for each vesicle type, and fusion edges $E_F \subseteq \{(\mu, x)\}$ show the target compartment types for each vesicle type. Each vesicle type must have at least one source compartment type and one target compartment type. G_{CO} is termed *single-source* or *single-target* if every vesicle type has precisely one source or target compartment type, respectively. The plasma membrane is itself treated as a compartment, from which endocytic vesicles bud and to which secretory vesicles use. Our analysis does not consider extracellular vesicles such as exosomes (Raposo and Stoorvogel 2013).

The transport graph is a directed multigraph with loops $G_{TR} = \{N_C, E\}$. Each node $x \in N_C$ is a compartment. Each edge $(\mu, x, y) \in E$ is a tuple associated with a vesicle label $\mu \in N_V$, and a source and target compartment pair $(x, y) \in N_C^2$. We permit multiple edges with distinct labels between the same pair of compartments, corresponding to distinct vesicle types. For each composition graph $G_{CO} = \{N_C \cup N_V, E_B \cup E_F\}$ we can *generate* a corresponding unique transport graph $G_{TR} = \{N_C, E\}$ as follows (figure 1A). The set of nodes N_C of the transport graph is the same as the set of compartment nodes N_C of the composition graph. Further, the transport graph contains an edge $(\mu, x, y) \in E$ if and only if the composition graph contains the budding and fusion edges $(x, \mu) \in E_B$ and $(\mu, y) \in E_F$.



◀ **Figure 1.** Graph representations of vesicle traffic networks. (A) Large grey circles are compartment nodes, small white circles are vesicle nodes. Each edge e_k of the transport graph G_{TR} is associated with some vesicle type μ_j of the composition graph G_{CO} . (B) A_i and a_k are the amounts of molecule type A on the nodes x_i and edges e_k of G_{TR} . Across edge e_k , r_k is the rate of passage of vesicles per unit time, and $r_k a_k$ is the flux of molecule type A . In steady state, $dA_i/dt = 0$ for each membrane-integral molecule type on each compartment. Here, I is the incidence matrix of G_{TR} . (C) Fluxes obey $r_k a_k \geq 0$, so $[r_1 a_1, r_2 a_2, \dots]^T$ must be a non-negative linear combination of directed simple cycles. The columns of the matrix B show all the directed simple cycles of G_{TR} , and $c_v \geq 0$. (D) We can assign distinct membrane-integral molecules and distinct rates to each directed simple cycle of G_{TR} . In this example, every cycle has unit rate, and the $M = 4$ molecule types are assigned as follows: every cycle carries one unit of membrane (brown); and each distinct cycle carries one unit of a distinct membrane-integral molecule (blue, green, red). M -tuples on the transport graph show molecular fluxes for each molecule type ($r_k a_k$). M -tuples on the composition graph show vesicle compositions for each molecule type (a_k). Observe that, in this example, each edge of G_{TR} has a distinct membrane-integral composition. We can show that this occurs if and only if G_{TR} has no 2-cut (Methods).

2.3 Molecular compositions

The molecules that regulate vesicle budding and fusion can either be *membrane-integral*, and constrained to move via vesicles, or *membrane-peripheral* and able to exchange between membranes and the cytosol. In addition to these regulatory molecules, the interior of a vesicle also contains various types of cargo molecules. Cargo is moved from one organelle to another by inter-compartment vesicles, or exchanged with the extracellular environment by endocytic or secretory vesicles that bud from or fuse into the plasma membrane.

Let M be the number of distinct molecule types carried by compartments or vesicles. The molecular composition function $\theta : N_C \cup N_V \rightarrow \mathbb{R}^M$ labels each $x \in N_C$ or $\mu \in N_V$ of the composition graph G_{CO} with an M -tuple $\theta(x)$ or $\theta(\mu)$, giving the amounts of the M distinct molecule types on the corresponding compartment or vesicle. By definition, θ is injective over the domain N_C or N_V , so distinct vesicle nodes have distinct compositions, and distinct compartment nodes have distinct compositions. Since vesicle sizes

span a narrow range (Ratamero and Royle 2019), we assume that all vesicles carry precisely one unit of membrane.

We can map molecules from the composition graph to the corresponding transport graph. Let the composition graph G_{CO} have compartment nodes $N_C = \{x_i\}$ and vesicle nodes $N_V = \{\mu_j\}$. Let the corresponding transport graph G_{TR} have compartment nodes $N_C = \{x_i\}$ and edges $E = \{e_k\}$, where each edge e_k corresponds to some vesicle type $\mu_{j(k)}$ of the composition graph. On this transport graph, each compartment x_i is associated with the M -tuple $\theta(x_i)$ and each edge e_k is associated with the M -tuple $\theta(\mu_{j(k)})$. Suppose the l th entry of the M -tuple corresponds to molecule type A . For ease of notation we use upper-case or lower-case letters to represent the quantity of A on compartment x_i or edge e_k : $A_i \equiv \theta(x_i)_l$ or $a_k \equiv \theta(\mu_{j(k)})_l$. If A represents membrane amount, then $a_k = 1$ on every edge.

2.4 Molecular flows

We now introduce dynamics. The rate function $\rho : E \rightarrow \mathbb{R}^+$ labels each edge $e = (\mu, x, y) \in E$ of G_{TR}

with a value $r = \rho(e)$, giving the number per unit time at which the corresponding vesicle type passes along that edge. If $r_k = \rho(e_k)$ is the rate associated with edge e_k , the net flux of molecule A via edge e_k is $r_k a_k$. If we neglect synthesis and degradation, the quantity of any membrane-integral molecule A on compartment x_i obeys the differential equation $dA_i/dt = \sum_k I_{ik} r_k a_k$ (figure 1B). Here, I is the incidence matrix of the directed transport graph G_{TR} , with columns corresponding to self loops set to zero. Membrane-peripheral molecules can have additional exchange fluxes between membrane and cytoplasmic pools, and luminal cargo molecules can have additional exchange fluxes with the extracellular environment. Since every edge of G_{TR} carries membrane, and all membrane derives from the ER, there must be at least one directed edge crossing in each direction across any graph cut (if not, membrane would accumulate on one side of the cut). This means G_{TR} must be strongly connected.

For compartment compositions to be in *steady state*, each membrane-integral molecule must obey $\sum_k I_{ik} r_k a_k = 0$, $r_k a_k \geq 0$ (figure 1B). Such a flow is called a circulation, and is always a nonnegative linear combination of directed simple cycles (Bang-Jensen and Gutin 2008, Chapter 3), implying: $r_k = \sum_v B_{kv} \tilde{r}_v$ (balancing membrane) and $a_k = (\sum_v B_{kv} \tilde{r}_v \tilde{a}_v) / (\sum_v B_{kv} \tilde{r}_v)$ (balancing membrane-integral molecules) (figure 1D). Here, B is the matrix whose columns are the directed simple cycles of G_{TR} , and $\tilde{r}_v, \tilde{a}_v \geq 0$ are rates and molecule amounts associated with each cycle (figure 1C,D). If multiple edges of G_{TR} correspond to the same vesicle type of G_{CO} , there are additional constraints on \tilde{a}_v . Removing the steady state requirement allows more complex dynamics, explored elsewhere (Mani and Thattai 2016a).

2.5 Regulation of vesicle fusion

The molecules carried by compartments and vesicles are not passive cargo, they actively regulate vesicle budding and fusion. In this analysis we do not apply any constraints on vesicle budding: each compartment type can bud out an arbitrary number of vesicle types. To describe constraints on vesicle fusion, we define a fusion indicator function $\phi: \mathbb{R}^M \times \mathbb{R}^M \rightarrow \{0, 1\}$. Here, $\phi(\hat{\mu}, \hat{y}) = 1$ if and only if a vesicle with M -tuple $\hat{\mu}$ can fuse with a compartment with M -tuple \hat{y} . We say that the composition graph G_{CO} is *well-fused* when, for all vesicle-compartment pairs $\{(\mu, y) \in N_C \times N_V\}$, there is a fusion edge $(\mu, y) \in E_F$ if and only if

$\phi(\theta(\mu), \theta(y)) = 1$. That is, all allowed fusions occur and no disallowed fusions occur.

In physiological conditions, every fusion event requires the binding of a specific v-SNARE on a vesicle with a specific cognate t-SNARE on a target compartment (Jahn and Scheller 2006; Bhattacharyya et al. 2019). We use a simplified description of SNARE binding that captures the essential biological details; we have explored more complex scenarios in previous computational work (Shukla et al. 2017; Bhattacharyya et al. 2019). We assume that v-SNAREs and t-SNAREs come in cognate pairs, and that members of distinct cognate pairs do not cross-interact. Each t-SNARE is always active. Each v-SNARE can be in an active or inactive state, regulated by other molecules on the same vesicle. The binding of a single active v-SNARE molecule with a single active t-SNARE molecule results in vesicle fusion (Van Den Bogaart et al. 2010). The fusion indicator function $\phi(\hat{\mu}, \hat{x})$ can be constructed to encode all these details.

2.6 Allowed and disallowed transport graphs

We say a transport graph G_{TR} is *allowed* if there exists a composition graph G_{CO} , along with molecular compositions θ , rates ρ , and a fusion function ϕ , such that: (i) G_{TR} is generated from G_{CO} ; (ii) G_{TR} is in steady state; (iii) G_{CO} is well-fused and single-target. Recall: steady state means compartment compositions are constant over time; well-fused means the target compartment types for each vesicle type are determined by local molecular interactions; and single-target means that any vesicle type is guaranteed to deliver its cargo to a single target compartment type.

To check if a given transport graph G_{TR} is allowed, it is sufficient to search over composition graphs G_{CO} that are simultaneously single-source and single-target. (Any vesicle type in G_{CO} with multiple source compartment types can always be replaced by multiple compartment-specific vesicle types with the same fusion properties, by adding membrane-peripheral dummy molecules; this results in a new composition graph G'_{CO} that generates the same G_{TR} .) This means we can associate each edge in G_{TR} with a distinct vesicle type, and can therefore refer to edge compositions and vesicle compositions interchangeably. We can assign distinct membrane-peripheral molecules to each compartment and edge of G_{TR} , since membrane-peripheral molecules have no steady-state constraints. We can assign distinct membrane-integral molecules to each compartment and directed simple cycle of G_{TR} ,

and assign rates $\tilde{r}_v > 0$ to each directed simple cycle. This assignment ensures G_{TR} is in steady state.

We now deal with the well-fused requirement. Without loss of generality, we can associate each edge e in G_{TR} with a distinct v-SNARE V_e responsible for its fusion. We assign V_e to some directed simple cycle containing e , and its cognate t-SNARE T_e to e 's target compartment. (Note that we may need to carefully choose the cycle for each V_e ; multiple v-SNAREs may be assigned to the same cycle, and multiple t-SNAREs may be assigned to the same target compartment.) Consider v-SNARE V_e responsible for the fusion of edge e . Each edge of the directed simple cycle followed by V_e has a distinct target compartment. By construction, only the target compartment of edge e carries the cognate t-SNARE T_e . The system will be well-fused if and only if, for each edge e in G_{TR} , V_e is active only on edge e . To check this, we consider four classes of SNARE regulation, ranging from more to less flexible, which serve to illustrate our main ideas. We have explored additional types of fusion regulation in previous computational work (Shukla *et al.* 2017; Bhat-tacharyya *et al.* 2019).

- *Membrane-peripheral regulation:* v-SNARE activity is a function of the membrane-peripheral molecules on a vesicle. Each edge in G_{TR} has a distinct membrane-peripheral composition, so we can always ensure that V_e is active only on edge e . Every allowed G_{TR} is strongly connected, and every strongly connected G_{TR} is allowed.
- *Membrane-integral regulation:* v-SNARE activity is a combinatorial function of the membrane-integral molecules on a vesicle. We can show (Methods) that each edge of a strongly connected graph belongs to a distinct set of directed simple cycles if and only if the graph does not have a 2-cut. If G_{TR} has a 2-cut, the two edges e_1 and e_2 crossing it will have identical membrane-integral compositions, so any v-SNARE active on one will be active on the other. If G_{TR} has no 2-cut, each edge will have a distinct membrane-integral composition, so we can always ensure that V_e is active only on edge e . Every allowed G_{TR} is strongly connected with no 2-cut, and every strongly connected G_{TR} with no 2-cut is allowed.
- *SNARE inhibition:* Each v-SNARE has a stoichiometric membrane-integral inhibitor; a v-SNARE is active on a vesicle whenever the amount of inhibitor is below the amount of v-SNARE, and is inactive otherwise (figure 2, left panels). Consider v-SNARE V_e responsible for the fusion of

edge $e = (\mu, x, y)$, and let I_e be its cognate inhibitor. To ensure that V_e is active only on edge e of the directed simple cycle followed by V_e , I_e must be above threshold on every edge except e . For I_e to be in steady state, we must then assign it to a path from x to y that does not contain e . Suppose G_{TR} is not 2-edge-strongly connected. Then there is some edge e which, when removed, leaves no other path from x to y , so I_e will not be in steady state. Suppose G_{TR} is 2-edge-strongly connected. Then there is always an x to y path when edge e is removed which, when combined with the y to x path followed by V_e , forms a directed simple cycle. We can assign I_e to this cycle. Every allowed G_{TR} is 2-edge-strongly connected. Every 2-edge-strongly connected G_{TR} is allowed.

- *SNARE activation:* Each v-SNARE has a stoichiometric membrane-integral activator; a v-SNARE is active on a vesicle whenever the amount of activator is above zero, and is inactive otherwise (figure 2, right panels). Consider v-SNARE V_e responsible for the fusion of edge $e = (\mu, x, y)$, and let A_e be its cognate activator. To ensure that V_e is active only on edge e of the directed simple cycle followed by V_e , A_e must be present only on edge e . Suppose G_{TR} is not 2-edge-strongly connected. Then there is a graph cut that contains a single edge e' crossing in one direction. Let $e = (\mu, x, y)$ be one of the edges crossing in the opposite direction. Then all paths from y to x must pass through e' , so e' contains both V_e and A_e . This violates the well-fused condition. Suppose G_{TR} is 2-edge-strongly connected. Then, by Menger's theorem (Bang-Jensen and Gutin 2008, Chapter 7), we can always find two edge-independent paths from y to x . These paths define two cycles that have only edge e in common. We can assign V_e to one cycle and A_e to the other. Every allowed G_{TR} is 2-edge-strongly connected. Every 2-edge-strongly connected G_{TR} is allowed.

2.7 Vesicles with multiple target compartments

If we permit multi-target composition graphs in addition to single-target ones in the definition of 'allowed', then the conditions listed above are sufficient but not necessary for a transport graph to be allowed. Every allowed transport graph is still strongly connected, all strongly connected transport graphs are still allowed under membrane-peripheral regulation, and some (but




Figure 2. SNARE inhibition and activation. Solid edges are single-step paths, wavy edges are edge-disjoint paths of arbitrary length. Edge colours represent the presence of a v-SNARE V and its cognate inhibitor I or activator A . We consider two types of SNARE regulation. **(A,C,E,G)** Inhibition: Free V is active, the VI complex is inactive. **(B,D,F,H)** Activation: Free V is inactive, the VA complex is active. **(A,B)** Sub-graphs of the transport graph showing connections between two compartments. V is responsible for fusion of the solid edge. V , along with I or A , must follow the designated paths to ensure that V is active only on the solid edge and nowhere else. We term these the I -motif or the A -motif. **(C,D)** The transport graph from figure 1. Edges are coloured to show a collection of flows consistent with the I -motif or A -motif. **(E,F)** Blue edges can be made to fit the I -motif. Teal edges can be made to fit the A -motif. Starred edges do not fit either motif. Therefore, this graph is not allowed under either inhibition or activation scenarios. **(G,H)** With the addition of two edges to the transport graph, every edge can be made to fit the I -motif or A -motif. Each edge will be associated with a distinct V , regulated by a cognate I or A . **(I)** We can show that G_{TR} is allowed under SNARE inhibition or activation scenarios if and only if it is 2-edge-strongly connected.

3. Discussion

Vesicle traffic is self-organized: molecular interactions set up the flows, and the flows define which molecular interactions can occur. It is this feedback that distinguishes vesicle traffic networks from other well-studied flow networks that arise in engineering contexts. We have previously studied vesicle traffic as a dynamical system (Ramadas and Thattai 2013; Mani and Thattai 2016a). Since the functions that govern vesicle budding and fusion are highly nonlinear, these systems can exhibit complex dynamics when started from arbitrary initial conditions. However, once the system has reached a steady state, the resulting molecular fluxes obey linear constraints. This property allows us to decouple the analysis of molecular rules from the analysis of flows, and see if they are consistent with one another: assuming some scenario of molecular regulation, we can check if some experimentally determined transport graph has the required minimum edge connectivity (table 1). If not, then either the regulation mechanism is more complex than assumed, or some transport pathway has not yet been characterized (Bhattacharyya et al. 2019). The constraints we have uncovered here explain why vesicle traffic networks are highly distinct from random networks with similar degree distributions (Mani and Thattai 2016a).

not all) strongly connected transport graphs with 2-cuts are allowed under the other regulation scenarios. Under membrane-integral or SNARE activation/inhibition scenarios, if the transport graph has a 2-cut then some vesicle type must have multiple target compartment types (figure 3). A special case of this arises if the 2-cut is across back-and-forth edges between a pair of compartments (figure 3A). The vesicle type corresponding to each edge would then fuse to its own source compartment type as well as to its desired target compartment type, creating a self loop on the transport graph. We term this the ‘weak single-target’ case: every cargo molecule loaded onto a vesicle may return multiple times to its source compartment, but is eventually delivered to the desired target compartment. Cells may have evolved specific mechanisms to prevent such ‘back-fusion’ of vesicles (Kamena and Spang 2004).




Figure 3. Multi-target vesicle traffic. Some, but not all, transport graphs with 2-cuts may be allowed if we remove the single-target requirement. If the transport graph has a 2-cut, the two edges (red, blue) crossing the cut have identical membrane-integral compositions. Under membrane-integral or SNARE activation/inhibition regulation scenarios, the vesicle types corresponding to each edge must fuse to the target compartment types of both edges. (A, B) We can map the additional fusion edges to existing edges of the given transport graph, so it is allowed. Multiple edges of the transport graph correspond to a single vesicle type of the composition graph. (C) We cannot map the additional fusion edges to any existing edges of the given transport graph, so it is not allowed.

Table 1. Necessary and sufficient conditions for allowed transport graphs

SNARE regulation class	Allowed transport graphs
Membrane-peripheral	Strongly connected
Membrane-integral	Strongly connected, no 2-cuts
SNARE inhibition	2-edge-strongly connected
SNARE activation	2-edge-strongly connected

The central physiological constraint is the need to regulate the fusogenic v-SNAREs as they move through the vesicle traffic network. Our analysis explores a spectrum of potential SNARE regulatory mechanisms that can be mapped to known molecular players. Likely candidates for stoichiometric v-SNARE inhibitors are the cognate t-SNAREs themselves (Jahn and Scheller 2006; Bhattacharyya *et al.* 2019): data

across multiple systems suggest that certain v-SNAREs are transported between compartments as inactive v-t complexes (Pryor *et al.* 2008; Schäfer *et al.* 2012; Kent *et al.* 2012; Karnahl *et al.* 2017). SNARE activity is directly regulated by several other types of molecules, including SM proteins (Baker *et al.* 2015), tethers (Dubuke and Munson 2016), and covalent modification enzymes (Conibear and Davis 2010; Greaves and Chamberlain 2011). These regulators can be recruited to specific vesicle types by various membrane-integral proteins and lipids, or by membrane-peripheral molecules such as Rab GTPases and their effectors (Cai *et al.* 2007; Stenmark 2009). Specific categories of signalling lipids, which can be covalently modified by enzymes or transported between distinct membranes by lipid transfer proteins, can also play the role of membrane-peripheral regulators (Cockcroft and Raghu 2018).

Many types of vesicles are carried between compartments by molecular motors on cytoskeletal tracks, introducing a spatial component to vesicle traffic (Stenmark 2009). This means only certain types of vesicles and compartments come into proximity, relaxing the constraints on SNARE regulation. However, while cytoskeletal transport of vesicles is likely to play a role in increasing efficiency of cargo delivery, the specificity of vesicle fusion mainly arises from local molecular interactions (Cai *et al.* 2007). For example, intra-Golgi vesicular transport is similarly regulated whether Golgi compartments are spatially stacked or dispersed (Mani and Thattai 2016a). It is therefore likely that the constraints we describe are broadly applicable, even in species with strong spatial organization of vesicle traffic. However, real vesicle traffic systems have constraints beyond those discussed here. For example, vesicle budding is subject to various restrictions on cargo selection; and the number and variety of molecule types that regulate vesicle budding and fusion are limited. Such constraints may increase the required edge connectivity of the transport graph, for the same class of SNARE regulation.

The assignment of SNARE regulators, inhibitors, or activators to cycles of the transport graph amounts to a coincidence detection strategy: a specific set of molecules that drive vesicle fusion only come together in a specific vesicle type. This problem has the flavor of the recreational mathematical puzzles in which a boatman must assist various types of passengers and cargo across a river, while ensuring that the wrong combination never cross together (Schwartz 1961). Our problem, like the river crossing puzzle, is amenable to a

graph-theoretic analysis: based only on very general assumptions about vesicles and compartments, independent of quantitative biochemical details, we have uncovered rigorous and surprising constraints that govern vesicle traffic networks across species and cellular contexts.

4. Methods

Theorem 1 *Each edge of a strongly connected graph belongs to a distinct set of directed simple cycles if and only if the graph does not have a 2-cut.*

Proof Let G be a strongly connected graph. Suppose G has a 2-cut. This must consist of two edges e_1 and e_2 crossing in opposite directions. Every cycle containing e_1 passes through e_2 and vice versa, therefore e_1 and e_2 belong to the same set of directed simple cycles. Conversely, suppose two edges $e_1 = (x_1, y_1)$ and $e_2 = (x_2, y_2)$ of G belong to the same set of directed simple cycles. Let P be the set of nodes on all directed y_2 -to- x_1 paths in $G \setminus \{e_2\}$, and let Q be the set of nodes on all directed y_1 -to- x_2 paths in $G \setminus \{e_1\}$. Every node of G is in $P \cup Q$, and in the graph $G \setminus \{e_1, e_2\}$ there are no directed paths from P to Q and no directed paths from Q to P (otherwise we would be able to construct directed simple cycles containing e_1 but not e_2 and vice versa). Therefore P and Q define a 2-cut of G across which only edges e_1 and e_2 cross. \square

Acknowledgements

We thank Arnab Bhattacharyya for pointing out the importance of edge connectivity to the vesicle traffic problem. We thank Gregory Gutin, Anjali Jaiman, Ramya Purkanti and Ankit Shukla for useful discussions. SM acknowledges funding from the Institute for Basic Science, South Korea, Project Code IBS-R020-D1. MT acknowledges funding from the Simons Foundation (287975) which supports collaborative activities of the Simons Centre for the Study of Living Machines.

References

Baker RW, Jeffrey PD, Zick M, Phillips BP, Wickner WT, and Hughson FM 2015 A direct role for the Sec1/Munc18-family protein Vps33 as a template for SNARE assembly. *Science* **349** 1111–1114

Bang-Jensen J and Gutin GZ 2008 *Digraphs: theory, algorithms and applications* (Springer)

Bhattacharyya A, Gupta A, Kuppasamy L, Mani S, Shukla A, Srivas M and Thattai M 2019 A formal methods approach to predicting new features of the eukaryotic vesicle traffic system. *Acta Inform.* **58** 1–37

Bonifacino JS and Glick BS 2004 The mechanisms of vesicle budding and fusion. *Cell* **116** 153–166

Cai H, Reinisch K, and Ferro-Novick S 2007 Coats, tethers, Rabs, and SNAREs work together to mediate the intracellular destination of a transport vesicle. *Dev. Cell* **12** 671–682

Cockcroft S and Raghu P 2018 Phospholipid transport protein function at organelle contact sites. *Curr. Opin. Cell Biol.* **53** 52–60

Conibear E and Davis NG 2010 Palmitoylation and depalmitoylation dynamics at a glance. *J. Cell Sci.* **123** 4007–4010

Dubuke ML and Munson M 2016 The secret life of tethers: the role of tethering factors in SNARE complex regulation. *Front. Cell Dev. Biol.* **4** 42

Greaves J and Chamberlain LH 2011 Differential palmitoylation regulates intracellular patterning of SNAP25. *J. Cell Sci.* **124** 1351–1360

Jahn R and Scheller RH 2006 SNAREs—engines for membrane fusion. *Nat. Rev. Mol. Cell Biol.* **7** 631–643

Kamena F and Spang A 2004 Tip20p prohibits back-fusion of COPII vesicles with the endoplasmic reticulum. *Science* **304** 286–289

Karnahl M, Park M, Mayer U, Hiller U, and Jürgens G 2017 ER assembly of SNARE complexes mediating formation of partitioning membrane in Arabidopsis cytokinesis. *eLife* **6** e25327

Kent HM, Evans PR, Schäfer IB, Gray SR, Sanderson CM, Luzio JP, Peden AA, and Owen DJ 2012 Structural basis of the intracellular sorting of the SNARE VAMP7 by the AP3 adaptor complex. *Dev. Cell* **22** 979–988

Mani S and Thattai M 2016a Stacking the odds for Golgi cisternal maturation. *eLife* **5** e16231

Mani S and Thattai M 2016b Wine glasses and hourglasses: Non-adaptive complexity of vesicle traffic in microbial eukaryotes. *Mol. Biochem. Parasitol.* **209** 58–63

Pryor PR, Jackson L, Gray SR, Edeling MA, Thompson A, Sanderson CM, Evans PR, Owen DJ, and Luzio JP 2008 Molecular basis for the sorting of the SNARE VAMP7 into endocytic clathrin-coated vesicles by the ArfGAP Hrb. *Cell* **134** 817–827

Ramadas R and Thattai M 2013 New organelles by gene duplication in a biophysical model of eukaryote endomembrane evolution. *Biophys. J.* **104** 2553–2563

Raposo G and Stoorvogel W 2013 Extracellular vesicles: Exosomes, microvesicles, and friends. *J. Cell Biol.* **200** 373–383

Ratamero EM and Royle SJ 2019 Calculating the maximum capacity of intracellular transport vesicles. *bioRxiv*. <https://doi.org/10.1101/555813>

- Schäfer IB, Hesketh GG, Bright NA, Gray SR, Pryor PR, Evans PR, Luzio JP, and Owen DJ 2012 The binding of Varp to VAMP7 traps VAMP7 in a closed, fusogenically inactive conformation. *Nat. Struct. Mol. Biol.* **19** 1300–1309
- Schwartz BL 1961 An analytic method for the “difficult crossing” puzzles. *Math. Mag.* **34** 187
- Shukla A, Bhattacharyya A, Kuppusamy L, Srivas M, and Thattai M 2017 Discovering vesicle traffic network constraints by model checking. *PloS One* **12** e0180692
- Stenmark H 2009 Rab GTPases as coordinators of vesicle traffic. *Nat. Rev. Mol. Cell Biol.* **10** 513–525
- Südhof TC and Rothman JE 2009 Membrane fusion: grappling with SNARE and SM proteins. *Science* **323** 474–477
- Sztul E, Chen P-W, Casanova JE, *et al.* 2019 ARF GTPases and their GEFs and GAPs: concepts and challenges. *Mol. Biol. Cell* **30** 1249–1271
- Van Den Bogaart G, Holt MG, Bunt G, Riedel D, Wouters FS, and Jahn R 2010 One SNARE complex is sufficient for membrane fusion. *Nat. Struct. Mol. Biol.* **17** 358–364

Corresponding editor: MOHIT KUMAR JOLLY