Breaking the VLB Barrier

Randomized Routing with a Single Spraying Hop

Tegan Wilson
Daniel Amir
Nitika Saran
Bobby Kleinberg
Vishal Shrivastav
Hakim Weatherspoon

STOC 2024

Does Reconfigurable Networking Really Need Spraying?

Routing paths in reconfigurable networks are composed of ***direct hops*** and ***spraying hops***.

Direct: one step closer to the destination.

Spraying: random step to balance load.

Spraying uses bandwidth inefficiently.

Can we avoid or reduce it?

Image credit: Denish C *(License: CC BY-NC-ND 2.0)*

The Oblivious Reconfigurable Network Abstraction

- N network nodes.
- Discrete time slots 0, 1, 2, …
- Nodes can send to/receive from one neighbor per time slot.
  - *d>1 neighbors reduces to this*
- Connection schedule = sequence of permutations $\sigma_0, \sigma_1, \sigma_2, \ldots$
  - *At time t, node i can send to* $\sigma_t(i)$.

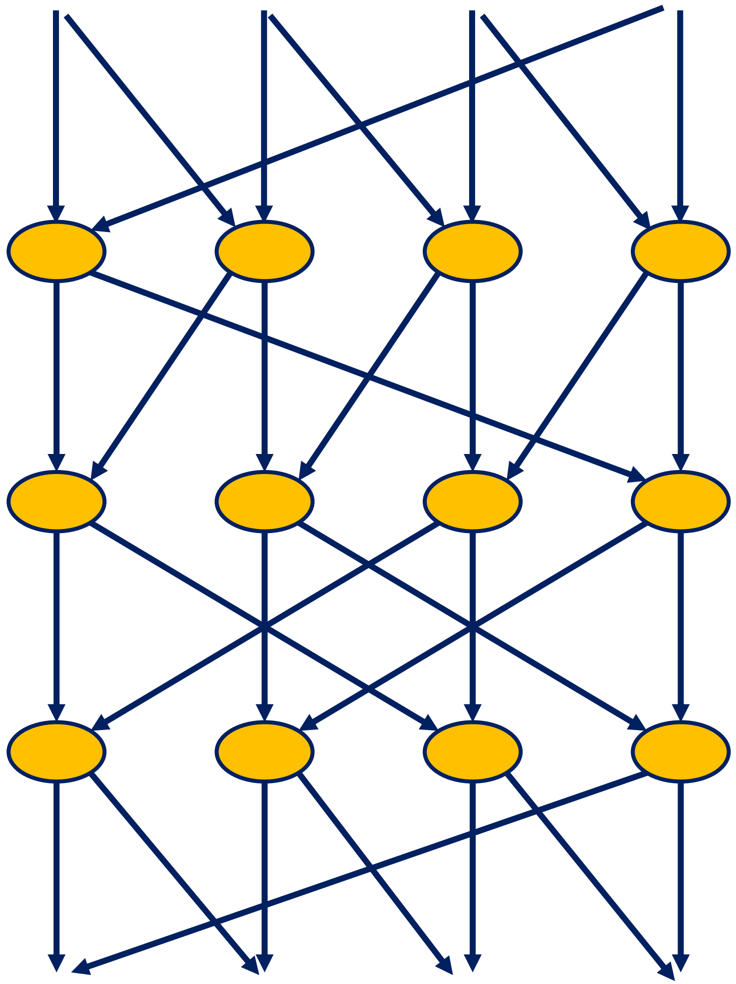

The Oblivious Reconfigurable Network Abstraction

- N network nodes.
- Discrete time slots 0, 1, 2, …
- Nodes can send to/receive from one neighbor per time slot.
  - *d>1 neighbors reduces to this*
- Connection schedule = sequence of permutations $\sigma_0, \sigma_1, \sigma_2, \ldots$
  - *At time t, node i can send to $\sigma_t(i)$.*

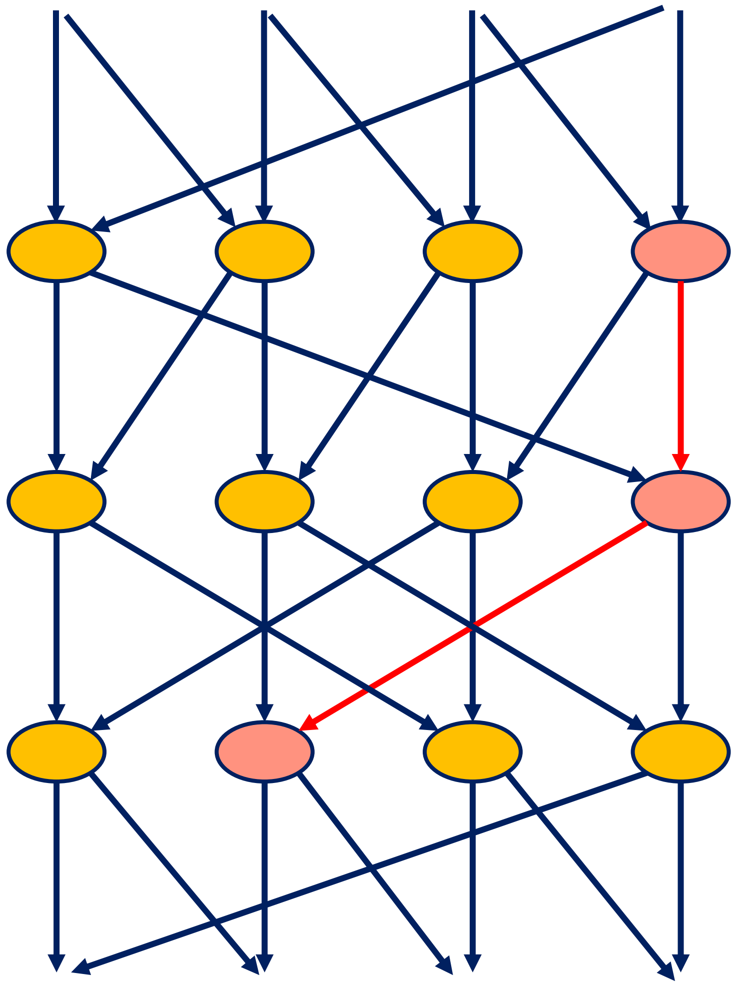

Routing Schemes

- **Routing Scheme** = probability distribution over paths for each (source, destination, starting time).
- ***ε*-balanced:** ∀ permutation demand matrix, $\frac{\text{max load}}{\text{avg load}} \leq 1 + \varepsilon$.
  - Here "load" means "expected flow on a link in a specific time slot".
- **Latency (L):** number of time slots from source to destination.
- **Hop count (h):** number of physical links traversed.

Valiant Load Balancing

- Used in RotorNet, Shoal, Sirius, Shale.
- Spraying hops = shortest path from source to random intermediate node.
- Direct hops = shortest path from intermediate to destination.
- Converts routing that is ε-balanced for uniform demands to ε-balanced for all permutation demands.
- **BUT…** results in **2x latency** and **2x hop-count**.
- **Is this doubling really necessary?**
  Yes, in some other network design settings.
  (Keslassy, Chang, McKeown, Lee 2005; Babaioff & Chuang 2007; AWSWKA 2022)

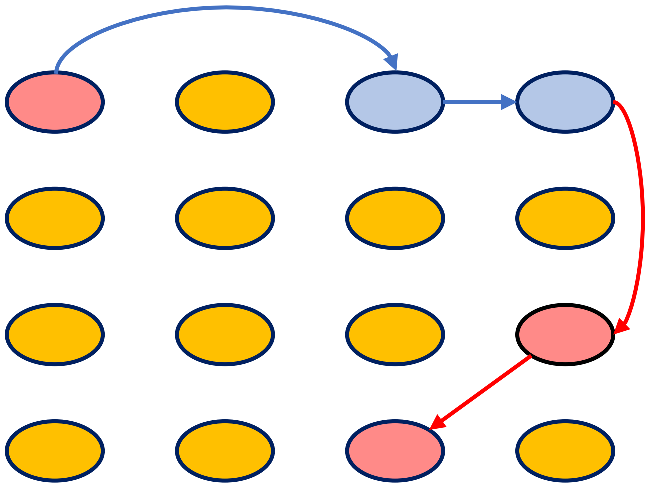

Why RotorNet/Shoal/Sirius Need VLB

- Suppose demand is a random permutation matrix, Π.
- Routing from source a to destination $\Pi(a)$ on a direct hop is only possible if $\sigma_t(a) = \Pi(a)$, which happens with probability $1/N$.
- So, average hop-count $\gtrsim 2$ is unavoidable.
- Impossibility is because the connection schedule is demand oblivious, holds even if routing is demand-aware.
- *Analogue of this reasoning for multi-hop routing?*

The "Speed of Light" Barrier: $L \geq \frac{g}{e} N^{1/g}$

- From a given source, how many nodes reachable in latency L and hop-count g?
- At most

$$\binom{L}{0} + \binom{L}{1} + \cdots + \binom{L}{g} \leq \left(\frac{eL}{g}\right)^g$$

- So, point-to-point reachability in $\leq g$ hops requires $L \geq \frac{g}{e} N^{1/g}$.
- Shale at level g achieves $L = 2gN^{1/g}$, $\boldsymbol{h = 2g}$

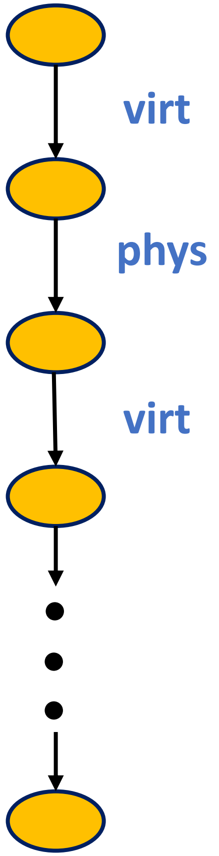

Necessity of Additive Stretch 1

- The link from a to $\sigma_t(a)$ belongs to a g-hop path only if $\Pi(a)$ is reachable from $\sigma_t(a)$ in latency $L-1$ and hop-count $g-1$.
- The number of such destinations is $\binom{L-1}{0} + \cdots + \binom{L-1}{g-1} \le \left(\frac{eL}{g}\right)^{g-1}$.
- When $L = O\left(gN^{1/g}\right)$ this is $O\left(N^{1-1/g}\right)$.
- Pr(direct hop leaving source a at time t) = $O\left(N^{-1/g}\right)$.
- So, average hop count $h \gtrsim g+1$ unavoidable when $L = O\left(gN^{1/g}\right)$.
- Again, this impossibility holds whenever connection schedule is demand oblivious, even if routing is demand aware.

Sufficiency of Additive Stretch 1

Theorem (WASKSW'24): For all $\varepsilon > 0, g \geq 1$, for infinitely many N, there exists a probability distribution on N-node ORN designs with:

- max latency $L = O\left(gN^{1/g}\right)$
- max hop count $h = g + 1$
- ∀ permutation demands, Pr(not ε-balanced) < $\frac{1}{N^{100}}$

TL;DR: If you're OK with negligible probability of violating load balance, **a single spraying hop is all you need!** VLB is overkill.

Contrast with AWSWKA'22: if load balance holds with probability 1, then $\boldsymbol{h \gtrsim 2g}$.

Key Idea #1: Shale in vector spaces

- Identify nodes with vectors in $\mathbb{F}_p^g$.
- Group time slots into phases (round robins) of length *p-1.*
- For $v \in \mathbb{F}_p^g$, a v-phase starting at time $t+1$ uses $\sigma_{t+s}(a) = a + sv$.
- For basis $B \subset \mathbb{F}_p^g$, a B-epoch consists of v-phases for each $v \in B$.
- Shale's connection schedule is made up of B-epochs where B = standard basis.

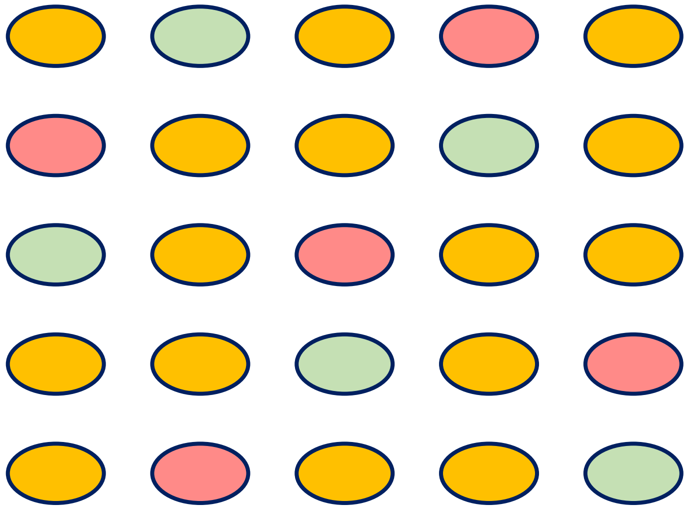

Key Idea #2: Constellations

- A basis defines a unique direct path.
  Instead **use an overcomplete set of vectors for redundancy**.
- **Constellation** = set of vectors in $\mathbb{F}_p^g$, any g of which form a basis.
- E.g. Vandermonde vectors $v = (1, x, x^2, \dots, x^{g-1})$.
- Connection schedule: sequence of v-phases for v ranging over a constellation of size $C(g+1)$, where $C = O(\log n)$.
- Routing scheme:
  - **randomly sample $g + 1$ phases**, one from each block of length C.
  - Take **at most 1 hop in each selected phase** (and exactly one in first and last).

Key Idea #3: Global Random Shuffle

- Choose the bijection $\{\text{nodes}\} \leftrightarrow \mathbb{F}_p^g$ uniformly at random.
- Limits an adversary’s ability to correlate the demand matrices with the connection schedule and routing scheme.

Some intuitions:

- Unlike in VLB, location after spraying is not uniformly random.
- Start of direct hop sequence correlates with destination!
- To still guarantee that load is ε-balanced by direct hops, incorporate two extra sources of randomness: timing of phases + global shuffle.

Summary: Given a latency bound of $\tilde{O}\left(gN^{1/g}\right)$ for integer g

Goal	Average Hop Count	Congestion	
Full Network Connectivity (lower bound)	g	–	Naïve counting
Uniform Multicommodity Flow	g	g	[AWSWKA'22]
Oblivious Routing (prob. 1)	$2g$	$2g$	[AWSWKA'22] (uses VLB)
Oblivious Routing (w.h.p.)	$g + 1$	$g + 1 + \delta$ $\forall \delta > 0$	**This work**
Demand-Aware Routing (prob. 1)	$g + 1$	$g + 1 + \delta$ $\forall \delta > 0$	**This work**

Two Closing Thoughts

1. Replacing VLB with a single spraying hop seems useful in general. Maybe we haven’t found the killer app yet.

2. A key to breaking the VLB barrier: **randomizing network topology independently of traffic demands.** ... *an unexpected side benefit of reconfigurability.*