

Brief Announcement:

Papillon: Greedy Routing in Rings

Ittai Abraham¹ and Dahlia Malkhi² and Gurmeet Singh Manku³

¹ Hebrew University of Jerusalem, ittai@cs.huji.ac.il

² Hebrew University of Jerusalem and Microsoft Research, dalia@microsoft.com

³ Google, manku@google.com

Abstract. We construct the first n -node degree- d ring-based network with worst-case GREEDY routes of length $\Theta(\log n / \log d)$ hops.

We study GREEDY routing over uni-dimensional metrics⁴ defined over n nodes lying in a ring. GREEDY routing in graph (V, E) with distance function $\delta : V \times V \rightarrow \mathcal{R}^+$ entails the following decision: Given target node t , node u with neighbors $N(u)$ forwards a message to $v \in N(u)$ such that $\delta(v, t) = \min_{x \in N(u)} \delta(x, t)$. The distance metric we use over n nodes placed in a circle is the clockwise-distance between pairs of nodes (the full paper contains a similar study of the absolute-distance metric):

$$\delta_{clock}(u, v) = v - u \text{ if } v \geq u, n + v - u \text{ otherwise.}$$

Let Δ_{clock} denote the worst-case GREEDY route length with δ_{clock} . $\bar{\Delta}_{clock}$ denotes the average GREEDY route length over all-pairs.

Summary of results. We construct a family of network topologies, the *Papillon*⁵, in which GREEDY routes are asymptotically optimal. Papillon has GREEDY routes of length $\Delta_{clock} = \Theta(\log n / \log d)$ hops in the worst-case when each node has d out-going links. Papillon is the first construction that achieves asymptotically optimal worst-case GREEDY routes.

Upon further investigation, two properties of Papillon emerge: (a) GREEDY routing does not send messages along shortest paths, and (b) Edge congestion with GREEDY routing is not uniform – some edges are used more often than others. We exhibit the first property by identifying routing strategies that result in paths shorter than those achieved by GREEDY routing. In fact, one of these strategies guarantees uniform edge-congestion.

The Papillon network. The intuitive idea of the Papillon construction is as follows. Nodes are arranged in m levels, where $n = \kappa^m m$. A node at level i has links to nodes one level down (wrapping at zero). Level- i links cover nodes between

⁴ The principles of this work can be extended to higher dimensional spaces. We focus on one-dimension for simplicity.

⁵ Our constructions are variants of the well-known butterfly family, hence the name Papillon.

the current node, and distance roughly κ^i hops away, where hops are counted by the logical ordering of the nodes on the ring, and cut the distance to the target to roughly $\kappa^{(i-1)}$. In essence, Papillon is a geometric analogue of the bitonic butterfly network.

The reason that greedy routing works on Papillon is the following. So long as the current node's level i is too small, i.e., the distance to the target is more than κ^i , any link followed serves simply to go down the levels until we wrap at level zero and reach level $m-1$. When the current level i is large enough, i.e., the distance to the target is at most κ^i , then one of the level links cuts the distance to κ^{i-1} . In this manner, eventually the target is reached.

$\mathcal{P}_{clock}(\kappa, m)$ is a directed graph for any $\kappa, m \geq 1$. Let $n = \kappa^m m$. Let $\ell(u) \equiv (m-1) - (u \bmod m)$. Each node has κ links. For node u , these links are to nodes $(u+x) \bmod n$, where $x \in \{1 + im\kappa^{\ell(u)} \mid i \in [0, \kappa-1]\}$.

Theorem 1. For $\mathcal{P}_{clock}(\kappa, m)$, $\Delta_{clock} \leq 3m-2$ and $\bar{\Delta}_{clock} \leq 2m-1$.

Putting $d = \kappa$, we obtain that in an n -node, degree- d network, worst-case routes are $\Theta(\log n / \log d)$. Curiously, GREEDY routing is not along shortest paths.

Theorem 2. There exists a non-GREEDY strategy for \mathcal{P}_{clock} with routes of length at most $2m-1$. The average is at most $1.5m$. The strategy is congestion-free, with uniform load on edges for all-pairs communication.

Previous results. With $\Theta(\log n)$ out-going links per node, several graphs over n nodes in a circle support GREEDY routes with $\Theta(\log n)$ GREEDY hops. Deterministic graphs with this property include variants of Chord, and this is also the known lower bound on any uniform graph (i.e., whose link-distances are the same from every node) with distance function $\delta_{clockwise}$. Randomized graphs which trade degree- d with route lengths $\Theta((\log^2 n)/d)$ on average include randomized-Chord and Symphony, with a gap to the known lower bound on uniform randomized network of $\Omega(\frac{\log^2 n}{d \log \log n})$.

Papillon extends the above results by constructing a non-uniform graph with $\Delta_{clock} = \Theta(\log n / \log d)$ in an n -node degree- d network, which is asymptotically optimal. Previously, this tradeoff was achieved on butterfly-like networks only with non-greedy routing (Viceroy, Ulysses, and Mariposa) or by greedy routing assisted with look-ahead. For $d = o(\log n)$, this beats the lower bound on uniform, randomized greedy routing networks (and it meets it for $d = O(\log n)$). In the specific case of $d = \log n$, our greedy routing achieves $O(\log n / \log \log n)$ average route length.

Please see <http://arxiv.org/abs/cs.DC/0507034> for further details.