# All Pairs Shortest Path Performance Measurements

Jérôme Chiabaut chiabaut@nortel.com

March, 2008

## > Foreword

- The premise is that all pairs shortest path is too expensive and time consuming
- This presentation show calculations of all pairs shortest path for two moderate processors
- We feel All Pairs shortest path is feasible on current hardware here is some data.

# Platforms

| Platform | Core Duo | PowerQUICC |
|---|---|---|
| Processor | Xeon® Processor LV 2.0 GHz (Sossaman) | MPC8548E (PowerQUICC III) |
| Clock speed | 2.0 GHz | 1.0 GHz |
| L2 cache | 2 MB shared ECC | 512 kB ECC |
| Bus | 667 MHz FSB | 400 MHz |
| Memory controller | Intel E7520 (Lindenhurst) | N/A |
| Memory | 2 x 512 MB of 400 MHz ECC DDR2 LC3 | 1 x 1 GB of 533 MHz ECC DDR2 LC4 |
| Operating system | Fedora Core 5 – 2.6.15-1.2054_FC5smp | Linux mpc8548cds 2.6.11 |
| C Compiler | gcc version 4.1.0 20060304 | gcc version 3.4.3 20041021 |
| C flags | -O3 -fno-strict-aliasing -mtune=pentium-m | -O3 -fno-strict-aliasing |

# Topologies – Meshed Cores

> Topologies composed of a meshed core surrounded by dual-homed edges
  - Key parameters are the size of the mesh and the size and degree of meshiness of the core

> Lightly meshed large cores
  - Size/2 core routers with 2(log2(Size)-2) trunks and 2 lines each
  - Size/2 edge routers dual homed (2 links)
  - Size*Log2(Size)/2 point-to-point links
  - Size of core routers' neighbourhood of core nodes grows a log(Size)
  - Network diameter increases very slowly with size (4-6 typical)

> Fully meshed small cores
  - sqrt(2*Size) core routers with 2*sqrt(Size)-1 trunks and 2*sqrt(Size)-2 lines each
  - Size-sqrt(2*Size) edge routers dual homed (2 links)
  - 3*Size-5/2*sqrt(2*Size) point-to-point links
  - Size of core routers' neighbourhood of core nodes grows a sqrt(Size)
  - Network diameter is always 3

# Topologies – Ring Hierarchies

> Topologies composed of a primary ring surrounded by dual-homed secondary rings
  - Symmetric point-to-point links with metrics between 1 and 9
  - The instance running the algorithm is the node with the smallest ID in the primary ring
  - Key parameters are the size of the primary and secondary rings and the number of secondary rings which is assumed equal to the number of nodes on the primary ring

> Rings hierarchy
  - Primary ring with N nodes
  - N secondary rings with M nodes each
  - Two nodes in common between a secondary ring and the primary: each node of the primary ring is on two secondary rings
  - N(M-1) nodes and NM links in total
  - N and M vary between 5 and 15 (8 to 10 typical)
  - Very large network diameter (N/2+M-2)

# Topologies – Two-Tier Hierarchy

> Topologies composed of a small core surrounded by dual-homed metro nodes surrounded by dual-homed edge nodes
  - Symmetric point-to-point links with metrics between 1 and 9
  - The instance running the algorithm is the core node with the smallest ID
  - Key parameters are the size of the core, metro, and edge

> Two-Tier Hierarchy
  - N core nodes, fully meshed (except for N=2)
  - M metro nodes, dual-homed onto core nodes
  - L edge nodes, dual-homed onto metro nodes
  - L+M+N nodes and 2(L+M)+N(N-1)/2 links in total (2(L+M) for N=2)
  - L/M ~ M/N vary between 4 and 6
  - Small network diameter: 5 (4 for N=2)

# APSP Performance

| Topology | Nodes | Links | Neighbors | Diameter | Core Duo | PowerQUICC |
|---|---|---|---|---|---|---|
| Light mesh | 50 | 125 | 8 | 4 | 0.40 ms | 0.80 ms |
| Light mesh | 98 | 300 | 10 | 5 | 2.50 ms | 4.10 ms |
| Light mesh | 200 | 700 | 12 | 5 | 15.2 ms | 24.5 ms |
| Heavy mesh | 50 | 105 | 13 | 4 | 0.40 ms | 0.70 ms |
| Heavy mesh | 98 | 217 | 19 | 4 | 1.70 ms | 3.00 ms |
| Heavy mesh | 200 | 460 | 28 | 4 | 7.50 ms | 13.6 ms |
| Meshed core | 50 | 125 | 17 | 3 | 0.30 ms | 0.60 ms |
| Meshed core | 98 | 259 | 25 | 3 | 1.30 ms | 2.40 ms |
| Meshed core | 200 | 550 | 37 | 3 | 5.60 ms | 10.3 ms |
| Rings | 49 | 56 | 4 | 9 | 0.20 ms | 0.40 ms |
| Rings | 100 | 110 | 4 | 14 | 0.80 ms | 1.60 ms |
| Rings | 196 | 210 | 4 | 20 | 3.60 ms | 6.80 ms |
| Two-Tier | 52 | 100 | 10 | 4 | 0.40 ms | 0.60 ms |
| Two-Tier | 100 | 198 | 11 | 5 | 1.90 ms | 3.40 ms |
| Two-Tier | 203 | 413 | 14 | 5 | 7.80 ms | 13.5 ms |

# Backup Slides

# Light Mesh – 50 Switches

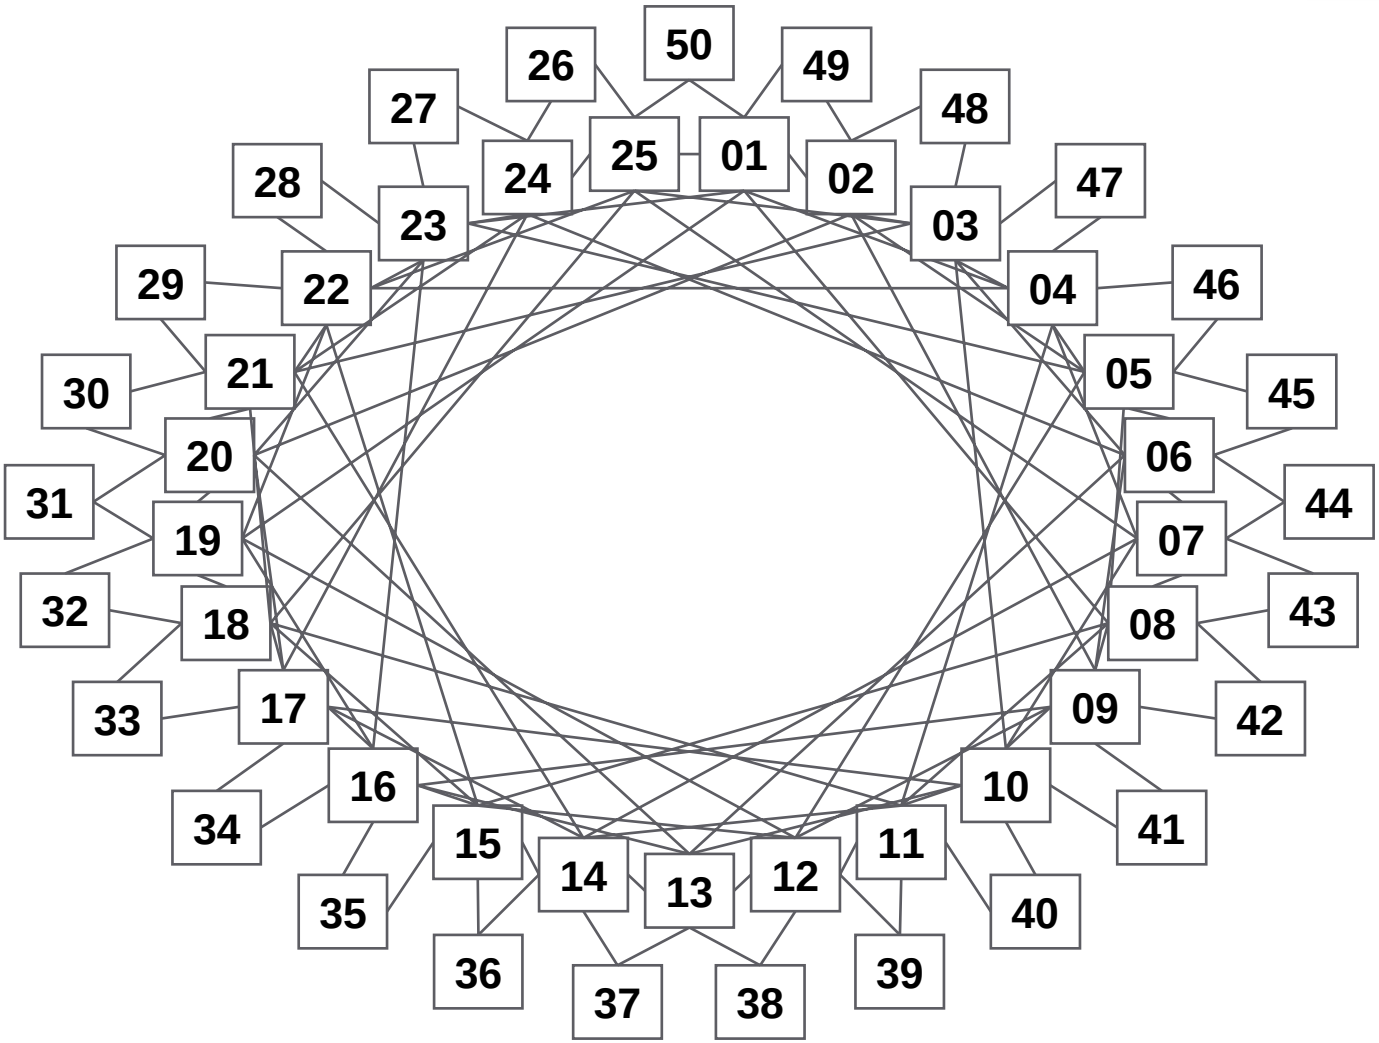

# Heavy Mesh – 50 Switches

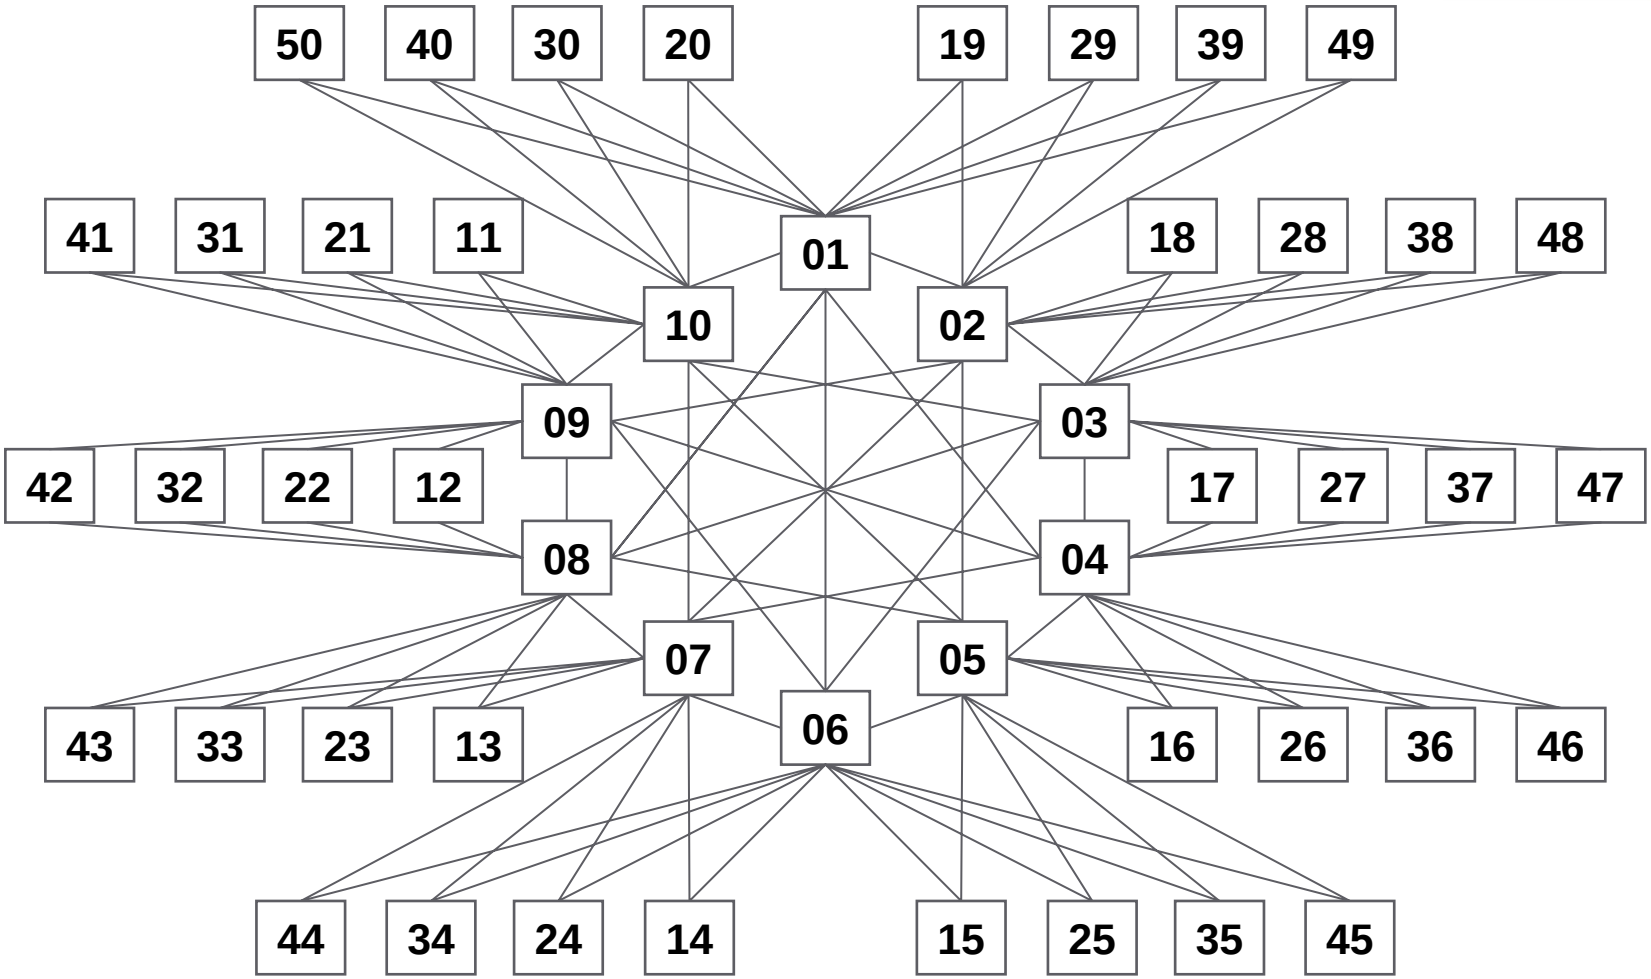

# Meshed Core – 50 Switches

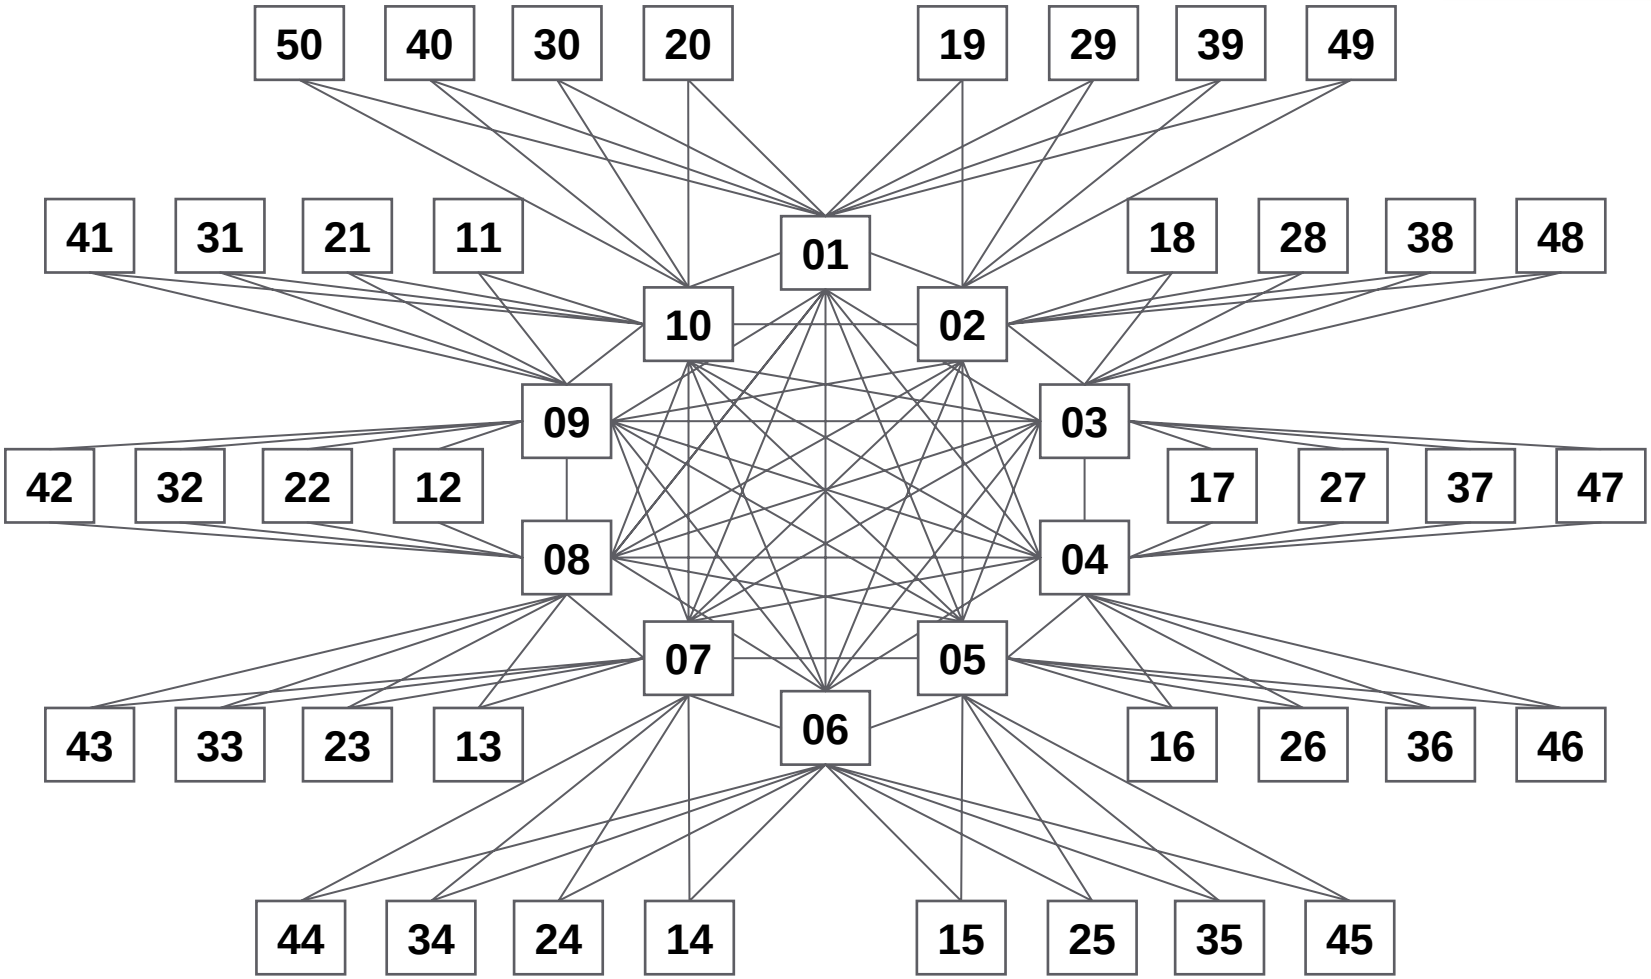

# Rings – 49 Switches

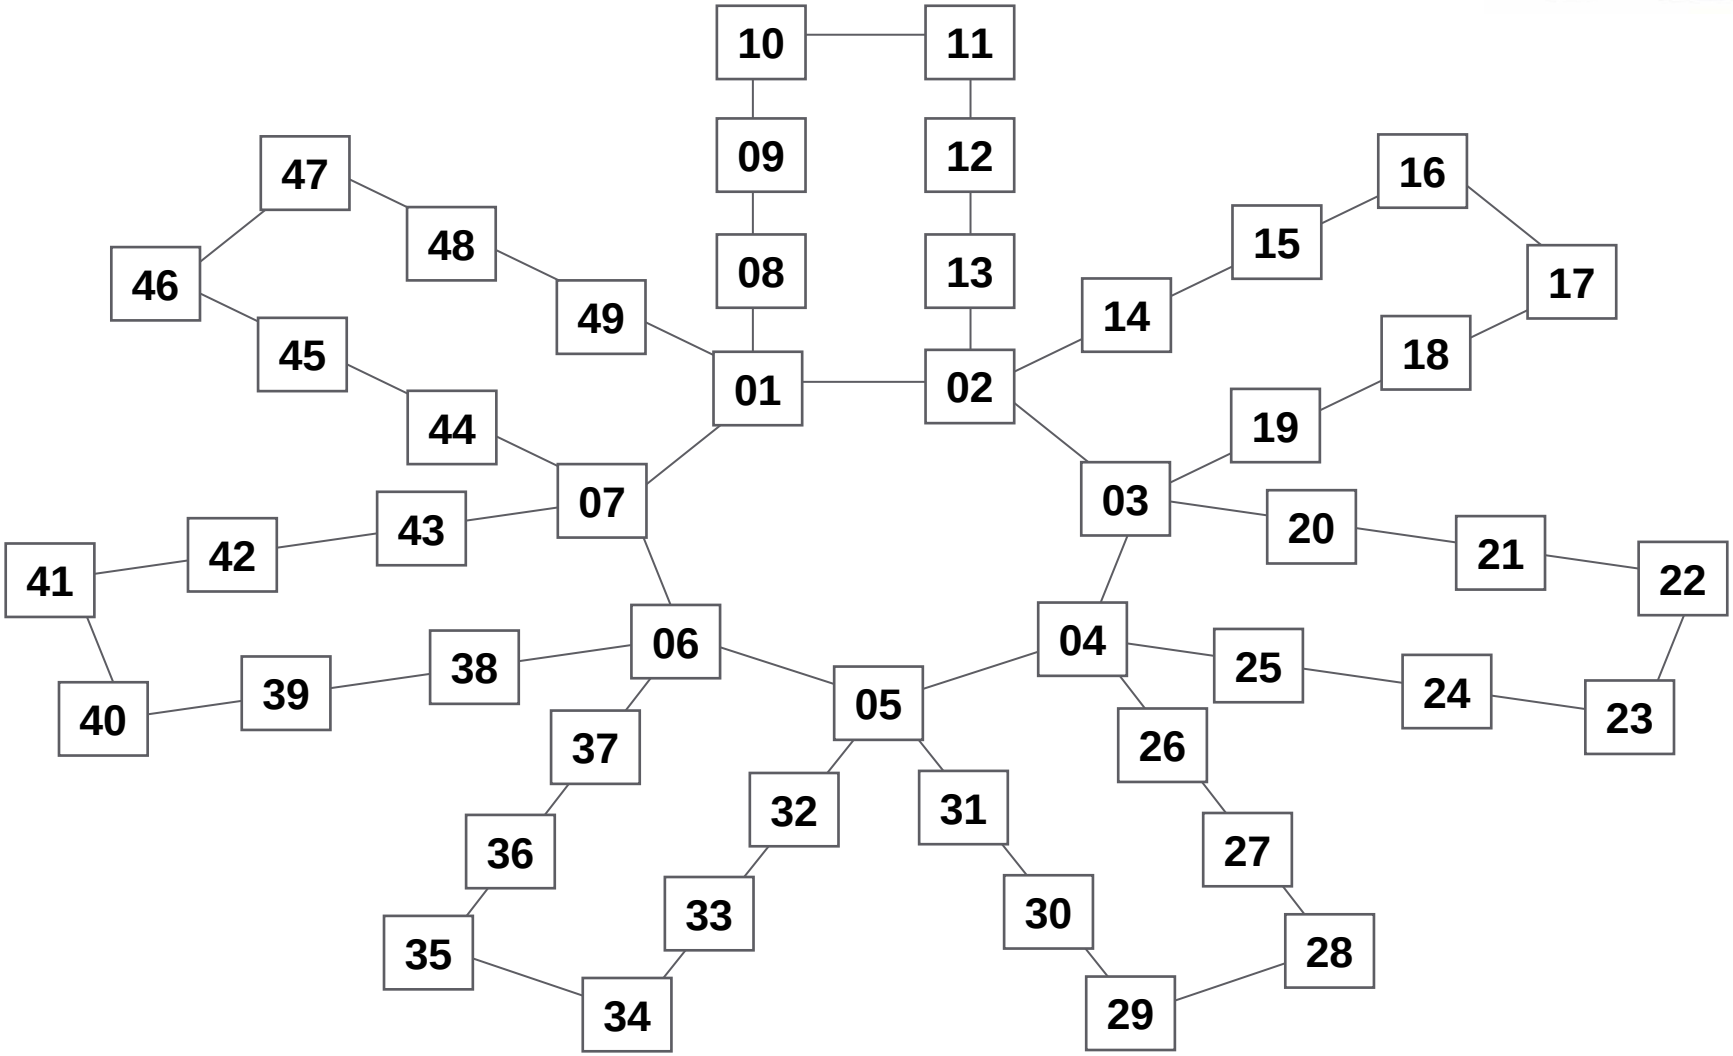

# Two-Tier – 52 Switches