



HAL
open science

High-Quality Fault Resiliency in Fat Trees

John Gliksberg, Antoine Capra, Alexandre Louvet, Pedro Javier Garcia,
Devan Sohier

► **To cite this version:**

John Gliksberg, Antoine Capra, Alexandre Louvet, Pedro Javier Garcia, Devan Sohier. High-Quality Fault Resiliency in Fat Trees. IEEE Micro, 2020, 40 (1), pp.44-49. 10.1109/MM.2019.2949978 . hal-03861715

HAL Id: hal-03861715

https:

//hal-universite-paris-saclay.archives-ouvertes.fr/hal-03861715

Submitted on 21 Nov 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution| 4.0 International License

High-Quality Fault Resiliency in Fat-Trees

John Gliksberg^{*†‡}, Antoine Capra[†], Alexandre Louvet[†], Pedro Javier García[‡], and Devan Sohier^{*}

Abstract—Coupling regular topologies with optimised routing algorithms is key in pushing the performance of interconnection networks of supercomputers. In this paper we present Dmodc, a fast deterministic routing algorithm for Parallel Generalised Fat-Trees (PGFTs) which minimises congestion risk even under massive network degradation caused by equipment failure. Dmodc computes forwarding tables with a closed-form arithmetic formula by relying on a fast preprocessing phase. This allows complete re-routing of networks with tens of thousands of nodes in less than a second. In turn, this greatly helps centralised fabric management react to faults with high-quality routing tables and no impact to running applications in current and future very large-scale HPC clusters.

Index Terms—HPC, routing, fat-tree, fault-resiliency

1 INTRODUCTION

A MAJORITY of current leading network topologies for High Performance Computing (HPC) clusters are fat-tree variants. (The five most powerful clusters of the June 2019 Top500 list [1] had fat-tree topologies.) These networks have some form of static routing tables computed by a centralised routing engine and uploaded to all switches.

It is sufficient for fat-tree-specific routing algorithms to be minimal to guarantee deadlock-free routing, and the regular nature of their target topology class should simplify load-balancing strategies. PGFTs [2] describe all regular fat-trees for which there is at most one downward switch-path from any switch to any node (as shown in Figure 1). In this article we refer to fat-trees as PGFTs. The oblivious routing algorithm for non-degraded PGFTs (Dmodk [2], see Section 2) uses this property and their connection logic to provide load balance through an arithmetic rule.

Due to the sheer amount of equipment in current and future supercomputers, hardware failures are to be expected [3] (especially in optical links [4] typically used in higher levels of fat-trees) and should not hinder running applications as far as possible. The fabric manager can react to equipment failures that do not break graph connectivity by uploading updated routing tables. In order to do this it requires a fault-resilient routing algorithm capable of rapid re-routing. The challenge is to provide these characteristics while maintaining high-quality static load balance.

Some of the research regarding oblivious fault-resilient routing focuses on techniques that explicitly target degradations to regular fat-trees [5] [6]; there are several re-routing strategies for these techniques. OpenSM’s UPDN [7] and Ftree [8] routing engines can also be applied from scratch to a degraded fat-tree. PQFT [5] is similar, though it requires a complete list of faults. The combination of Dmodk + Ftrnd_diff [9] available in BXI FM [10] is applied in an offline/online manner (with an iterative list of network changes and an up-to-date view of the network), the goal being fast reaction to faults with minimal routing changes.




Fig. 1. $PGFT(3; 2,2,3; 1,2,2; 1,2,1)$ with leaf switches shown in grey.

Fabriscale [11] also provides fast centralised re-routing of fat-trees, by precomputing alternative routes. A short summary of the limits of the existing approaches is provided in a previous conference paper [12].

The approach that we propose to meet that challenge is to apply the closed-form arithmetic formula of Dmodk while relaxing the topological constraint. For that purpose, we compute shortest paths explicitly rather than relying on an addressing scheme, and we balance load according to locally propagated information rather than relying on level-wide constants. These two goals are addressed together during preprocessing and will be the focus of this article, whose main contribution is the detailed algorithm description in Section 3.

2 DMODK

The Dmodk routing algorithm and corresponding PGFT topology are described in detail in [2]. The algorithm relies on a criterion (not shown here) to determine whether a destination d must be routed within the down ports and, if so, which one. Otherwise, an arithmetic formula defines the up port (with index p) to select:

$$p = \left\lfloor d / \prod_{k=1}^l w_k \right\rfloor \bmod (w_{l+1} p_{l+1})$$

The level-wide constants (or *arities*) w_l and p_l respectively denote the numbers of uplinks and of interlinks of all switches at level l . With this formula, each destination’s routes are coalesced as early as possible, and routes to different destinations are spread out as much as possible, thus

* Li-Parad, UVSQ, Versailles, {john.gliksberg,devan.sohier}@uvsq.fr
 † Atos, Bruyères-le-Châtel, {antoine.capra,alexandre.louvet}@atos.net
 ‡ RAAP, UCLM, Albacete, pedrojavier.garcia@uclm.es

minimising collisions between independent traffic. These closed-form steps rely on a given organisation of addresses of switches and indexing of their ports. Dmodk is a very low complexity and perfectly parallel routing algorithm for PGFTs, but it is not applicable to degraded PGFTs or irregular fat-trees.

3 DMODC DESCRIPTION

The idea behind the fault-resilient algorithm that we propose is to rely on local information while using the same closed-form arithmetic formula as Dmodk. The c in Dmodc refers to the neighbouring switches explicitly determined to be *closer* to the destination among which paths are chosen. The aim is fast centralised computation of routing tables for degraded PGFTs, providing optimal or well-balanced deterministic routes even under heavy fabric degradation. The algorithm begins with a preprocessing phase (that can be multi-threaded) followed by a parallel computation phase. Links are assumed to be bi-directional; notations used in the expressions hereafter are defined in Table 1.

TABLE 1
Notations Used in Expressions.

S	is the set of switches
L	is the set of leaf switches ($L \subset S$)
N	is the set of nodes
E	is the set of edges
λ_n	is the (only) leaf switch connected to node n ($\lambda_n \in L$)
$\vec{}, \overleftarrow{}$	respectively denote down and up links, according to rank
G_s	is the ordered list of port groups of switch s
Ω_g	is the switch connected to port group g
$\#$	denotes cardinality
<i>Dmodc-specific notations:</i>	
$c_{s,l}$	is the cost of switch s to leaf switch l
Π_s	is the divider of switch s

3.1 Basic Preprocessing

For ranking, levels and link directions are determined according to leaf switches being equivalent to the lowest level. Groups of ports linked to the same switch are prepared and sorted by globally unique identifier (GUID) to help with same-destination route coalescing.

3.2 Cost

We define the cost $c_{s,l}$ of a switch s to a leaf switch l to be the minimum number of hops between each other under up-down restrictions according to rank, as defined in Procedure 1 and illustrated in Figure 2. This later allows us to determine valid paths by exploring neighbouring switches and comparing costs. That exploration could be done here to prepare sets of output ports, but it's better to leave it for later since each set is only used once (see Subsection 3.4). Other all-pairs shortest paths methods could be substituted here.




Fig. 2. Example sequence of cost propagation steps in a degraded part of a network. Costs to the bottom-right switch are shown in switches. At each propagation step, the updated costs are in grey. Note that in steps 3–5, some propagations are interrupted due to the $c_{s,l} + 1 < c_{s',l}$ condition in the procedure. They could have been achieved with a simple $c_{s',l} = \infty$ condition instead; however this would have also interrupted the propagation of 2 in step 5. As a result, the long path on the left would not have been avoided. For PGFTs (degraded or not), such cases are actually impossible and the simple condition would suffice; but it would not guarantee shortest up-down paths in fat-tree-like topologies.

Procedure 1 Compute costs and dividers

```

for all  $s \in S$  do
  for all  $l \in L$  do
     $c_{s,l} \leftarrow \infty$ 
   $\Pi_s \leftarrow 1$ 
  for all  $l \in L$  do
     $c_{l,l} \leftarrow 0$ 
  for all  $s \in S$  sorted in ascending rank order do
     $\pi \leftarrow \Pi_s \times \#\{s' \vec{\phantom{s}} s\}$ 
    for all  $s' \vec{\phantom{s}} s$  do
      for all  $l \in L \mid c_{s,l} + 1 < c_{s',l}$  do
         $c_{s',l} \leftarrow c_{s,l} + 1$ 
      for all  $s' \vec{\phantom{s}} s \mid \Pi_{s'} < \pi$  do
         $\Pi_{s'} \leftarrow \pi$ 
  for all  $s \notin L$  sorted in descending rank order do
    for all  $s' \overleftarrow{\phantom{s}} s$  do
      for all  $l \in L \mid c_{s,l} + 1 < c_{s',l}$  do
         $c_{s',l} \leftarrow c_{s,l} + 1$ 

```

Thanks to the up-down restriction, the complexity of this procedure is in $\mathcal{O}(\#E\#L)$. This restriction is only for efficiency, it does not enforce deadlock-freedom. Some fat-tree-like topologies would result in up-down-up-down paths (if such shortcuts appear in neighbouring switches), since path selection does not distinguish up and down neighbours. Avoiding this requires a slightly different method: an extra integer must be stored, similar to cost but only for downpaths. More detail can be found in Section 3.4.

In our partially parallel implementation, each worker thread obtains a block of switches to propagate with one barrier per level upwards, then downwards.

3.3 Divider

Dmodc is based on the same arithmetic formula as Dmodk. Prior to the modulo operation, it begins with an integer division by the product of $\#\{s' \in S \mid s' \vec{} s\}$ (the upward arity




Fig. 3. Example sequence of divider propagation steps in a degraded part of a network. Dividers are shown in switches. At each propagation step, the updated dividers are in grey. Note that in step 2, the first upswitch is not updated because $\pi = 2 \times 2 \leq 6$. Even though there are multiple degradations in the considered case, all top switches end up with the divider that they would have had in the complete network.

of s) of switches at each lower level. This value represents the number of consecutive destinations to route through the same port. It is multiplied when going up levels to mirror the number of consecutive choices by switches below before each switch is chosen again. To reflect the actual state of the network (in which switches of the same level may have different arities), only local information must be considered; in turn, this operation is based on the products of up-to-date counts of upswitches (switches connected above), as defined in Procedure 1. Each downpath corresponds to a potential divider value, and we choose to keep only the maximum (as illustrated in Figure 3). The underlying motivation is to generate the same values as in the non-degraded PGFT, as long as the topological subgroup is not systematically degraded. The complexity of this part of the procedure is in $\mathcal{O}(\#E)$.

3.4 Routes Computation

The deterministic output port $p_{s,d}$ and alternative output ports $P_{s,d}$ of every switch s for every destination $d \in N$ (not directly linked to s) are selected with a closed-form formula based on the results previously determined. First, port groups leading *closer* to λ_d are selected in (1) (without taking ranking into account), setting corresponding alternative output ports in (2):

$$C_{s,\lambda_d} \leftarrow \{g \in G_s \mid c_{\Omega_{g,\lambda_d}} < c_{s,\lambda_d}\} \quad (1)$$

$$P_{s,d} \leftarrow \{p \in g \mid g \in C_{s,\lambda_d}\} \quad (2)$$

Selected port groups C are stored in an array (ordered by GUID of their remote switch), also represented by C : individual groups are accessed with indices $i \in [0, \#C_{s,\lambda_d}[$ using the $C_{s,\lambda_d}[i]$ notation. From this, the output port group is chosen in (3) and the port within that group in (4):

$$g_{s,d} \leftarrow C_{s,\lambda_d} \left[\left\lfloor \frac{d}{\Pi_s} \right\rfloor \bmod \#C_{s,\lambda_d} \right] \quad (3)$$

$$p_{s,d} \leftarrow g_{s,d} \left[\left\lfloor \frac{d}{\Pi_s \times \#C_{s,\lambda_d}} \right\rfloor \bmod \#g_{s,d} \right] \quad (4)$$

Routes are computed in a loop over leaves so that $C_{s,\lambda}$ is determined only once for all nodes connected to λ (with $P_{s,d}$




Fig. 4. Example route computation with s in grey, $\Pi_s = 4$, and $d = 20$. Costs to λ_{20} are shown in switches. Indices are ordered from left to right. The top-right group is chosen as $g_{s,20}$ because $\lfloor 20/4 \rfloor \bmod 2 = 1$, and the right port in $g_{s,20}$ is chosen as $p_{s,20}$ because $\lfloor 20/(4 \times 2) \rfloor \bmod 3 = 2$.

also unchanging $\forall d \mid \exists \lambda_d$). Figure 4 illustrates assignments (1), (2), (3), and (4).

The cost variant for up-down restriction described in 3.2 requires (1) to compare c values for upswitches and the downpath cost value for downswitches.

4 RESULTS

The algorithm was implemented in the fabric management software for Atos's Bull eXascale Interconnect (BXI). The same code has been used for validation, simulation, and in production.

4.1 Validity

Routing is valid for degraded PGFTs if and only if the cost of every leaf switch to every other leaf switch is finite: this reflects every node pair having an up-down path. Our implementation includes a pass through all leaf switch pairs to verify this condition. The up-down path restriction is sufficient to guarantee deadlock-freedom within degraded PGFTs [6].

4.2 Runtime

Our C99 implementation had computation of cost, divider, and routes spread over POSIX threads fetching work with a switch-level granularity. Figure 5 reports complete algorithm execution time alongside OpenSM (version 3.3.21) routing times (measured by adding timers in the source code) running on the same machine. For clusters ranging up to many tens of thousands of nodes, Dmodc provides fast enough re-routing for a centralised fabric manager to react to faults before applications are interrupted.

4.3 Quality

The routing algorithm was tested for quality by generating randomly degraded networks, computing corresponding routing tables, and then determining maximum congestion risk for multiple communication patterns. This study is available in a previously published extended abstract [12]. The results are comparable or better than the other available algorithms across the studied range of degradations.




Fig. 5. Algorithm runtime on a 2.50GHz Intel Xeon E5-2680 v3 for Real-Life Fat-Trees of varying sizes (in log-log scale; lower is better).

5 CONCLUSION

The simulation results in Section 4 show that Dmodc provides high-quality centralised fault-resilient routing for PGFTs at a fraction of the runtime of existing algorithms, without relying on partial re-routing. Dmodc is also applicable to fat-tree-like topologies (as mentioned in Figure 2) but with lower-quality load balancing. As defined here, no effort has been made to minimise the size of updates to be uploaded to switches throughout the fabric.

This algorithm is implemented inside BXI FM [10] and has been successfully deployed to an 8490 node PGFT production network in which it helps provide fault-resiliency even when faced with thousands of simultaneous changes.

ACKNOWLEDGEMENTS

This research has been undertaken under a cooperation between CEA and Atos. with the goal of co-designing extreme computing solutions. This research was partly funded by a grant of Programme des Investissements d’Avenir. This work has been jointly supported by the Spanish Ministry of Science, Innovation and Universities under the project RTI2018-098156-B-C52 and by JCCM under project SBPLY/17/180501/000498. BXI development was also part of ELCI, the French FSN (Fond pour la Société Numérique) cooperative project that associates academic and industrial partners to design and provide software components for new generations of HPC datacenters.

REFERENCES

- [1] E. Strohmaier, J. Dongarra, H. Simon, M. Meuer, and H. Meuer, “Top500, <https://www.top500.org/>,” 1993–2018.
- [2] E. Zahavi, “D-mod-k routing providing non-blocking traffic for shift permutations on real life fat trees,” *CCIT Report*, 2010.
- [3] J. Domke, T. Hoefler, and S. Matsuoka, “Fail-in-place network design: interaction between topology, routing algorithm and failures,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*. IEEE Press, 2014.
- [4] T. Connors, T. Groves, T. Quan, and S. Hemmert, “Simulation framework for studying optical cable failures in dragonfly topologies,” in *International Parallel and Distributed Processing Symposium (IPDPSW)*. IEEE, 2019.

- [5] E. Zahavi, I. Keslassy, and A. Kolodny, “Quasi fat trees for hpc clouds and their fault-resilient closed-form routing,” in *22nd Annual Symposium on High-Performance Interconnects (HOTI)*. IEEE, 2014.
- [6] J.-N. Quintin and P. Vigneras, “Transitively deadlock-free routing algorithms,” in *2nd International Workshop on High-Performance Interconnection Networks in the Exascale and Big-Data Era (HiPINEB)*. IEEE, 2016.
- [7] OpenSM, *Current OpenSM Routing & UPDN Routing Algorithm*, <https://github.com/linux-rdma/opensm>, retrieved 2007.
- [8] E. Zahavi, G. Johnson, D. J. Kerbyson, and M. Lang, “Optimized InfiniBand™ fat-tree routing for shift all-to-all communication patterns,” *Concurrency and Computation: Practice and Experience*, 2010.
- [9] P. Vigneras and J.-N. Quintin, “The bxi routing architecture for exascale supercomputer,” *The Journal of Supercomputing*, 2016.
- [10] “Bull exascale interconnect, <https://atos.net/en/products/high-performance-computing-hpc/bxi-bull-exascale-interconnect>.”
- [11] J. C. Villanueva, T. Skeie, and S.-A. Reinemo, “Routing and fault-tolerance capabilities of the fabric-scale fm compared to opensm,” Tech. rep. July, Tech. Rep., 2015.
- [12] J. Gliksberg, A. C. Capra, A. Louvet, P. J. García, and D. Sohier, “High-quality fault-resiliency in fat-tree networks (extended abstract),” in *26th Symposium on High-Performance Interconnects (HOTI)*. IEEE, 2019.

John Gliksberg MSc in Computer Science for HPC, BSc in Mathematics and Physics, is a PhD student at Versailles Saint-Quentin-en-Yvelines University (UVSQ), France; in an international cotutelle with Castilla-La Mancha University (UCLM), Spain; and in an industrial thesis contract with Atos, France. His research interests include routing algorithms for HPC interconnects and he has developed several routing algorithms used in BXI-based clusters.

Antoine Capra PhD in Computer Science, is a lead technical developer in two BXI projects at Atos, France. His thesis work focused on virtualisation in HPC context. His research interests include routing algorithms for HPC interconnects and optimisations to the MPI software layer. He has contributed to several routing algorithms used in BXI-based clusters.

Alexandre Louvet Lead technical developer in multiple BXI projects at Atos, France. His research interests include storage, algorithmic, and low-level optimisations for HPC. He has re-architected the BXI fabric management software, developed several routing algorithms used in BXI-based clusters and improved numerous aspects of BXI’s stack.

Pedro Javier García PhD in Computer Science, is an Associate Professor of Computer Architecture and Technology at Castilla-La Mancha University (UCLM), Spain. His research focuses on high-performance interconnection networks. He has published around 60 refereed papers in journals and conferences. He has guided 2 doctoral theses. He has coordinated 3 research projects funded by the Spanish and Castilla-La Mancha Governments. He has coordinated 4 R&D agreements between UCLM and different companies.

Devan Sohier PhD in Computer Science, is a full professor at Versailles Saint-Quentin-en-Yvelines University (UVSQ), within the LI-ParAD lab. He has previously been an associate professor at the University of Reims and at UVSQ. His primary research interest is distributed algorithms (in particular probabilistic distributed algorithms in stochastic environments), with applications to HPC. He has advised 3 PhD students, and is currently advising 2 others.