

# Fault Tolerance Logical Network Properties of Irregular Graphs

ICA3PP'12, Fukuoka, Japan

Christophe Cérin ♣ , **Camille Coti** ♣ , Michel Koskas ♠

♣ Université Paris 13

♠ Institut national de la recherche agronomique

*Sept 6th 2012*

Cloud computing is using the Internet to interconnect resources

- Connecting together HUGE amounts of computing and storage resources
- Designed to be resilient

Resilient?

- In March 2011 Armenia was disconnected from the Internet by a 75yo Georgian woman who stole a wire for copper ("the hacker with the shovel")
- In August Wikipedia was disconnected from the Internet because a data center in Florida was disconnected by an accidental cable cut



Cloud computing is using the Internet to interconnect resources

- Connecting together HUGE amounts of computing and storage resources
- Designed to be resilient

### Resilient?

- In March 2011 Armenia was disconnected from the Internet by a 75yo Georgian woman who stole a wire for copper ("the hacker with the shovel")
- In August Wikipedia was disconnected from the Internet because a data center in Florida was disconnected by an accidental cable cut



Cloud computing is using the Internet to interconnect resources

- Connecting together HUGE amounts of computing and storage resources
- Designed to be resilient

Resilient?

- In March 2011 Armenia was disconnected from the Internet by a 75yo Georgian woman who stole a wire for copper ("the hacker with the shovel")
- In August Wikipedia was disconnected from the Internet because a data center in Florida was disconnected by an accidental cable cut



How can we characterize a network topology wrt expectations for parallel applications?

- **Diameter** : longest shortest path between any two vertices from the graph.  
In other words: maximum number of hops made by a message to reach its destination.
- **Node connectivity** : minimum number of vertices that must be removed to disconnect the graph.  
In other words: how many nodes can fail before we can expect connectivity to be lost.
- **Link connectivity** : minimum number of links that must be removed to disconnect the graph.  
In other words: how many links (cables...) can be broken before we can expect connectivity to be lost.
- **Fault diameter** : diameter of the graph, given the maximum number of failure before the graph becomes bipartite ( $\kappa - 1$ , if  $\kappa$  is the node connectivity)  
In other words: how the diameter evolves in degraded conditions.

How can we characterize a network topology wrt expectations for parallel applications?

- **Diameter** : longest shortest path between any two vertices from the graph.  
In other words: maximum number of hops made by a message to reach its destination.
- **Node connectivity** : minimum number of vertices that must be removed to disconnect the graph.  
In other words: how many nodes can fail before we can expect connectivity to be lost.
- **Link connectivity** : minimum number of links that must be removed to disconnect the graph.  
In other words: how many links (cables...) can be broken before we can expect connectivity to be lost.
- **Fault diameter** : diameter of the graph, given the maximum number of failure before the graph becomes bipartite ( $\kappa - 1$ , if  $\kappa$  is the node connectivity)  
In other words: how the diameter evolves in degraded conditions.

How can we characterize a network topology wrt expectations for parallel applications?

- **Diameter** : longest shortest path between any two vertices from the graph.  
In other words: maximum number of hops made by a message to reach its destination.
- **Node connectivity** : minimum number of vertices that must be removed to disconnect the graph.  
In other words: how many nodes can fail before we can expect connectivity to be lost.
- **Link connectivity** : minimum number of links that must be removed to disconnect the graph.  
In other words: how many links (cables...) can be broken before we can expect connectivity to be lost.
- **Fault diameter** : diameter of the graph, given the maximum number of failure before the graph becomes bipartite ( $\kappa - 1$ , if  $\kappa$  is the node connectivity)  
In other words: how the diameter evolves in degraded conditions.

How can we characterize a network topology wrt expectations for parallel applications?

- **Diameter** : longest shortest path between any two vertices from the graph.  
In other words: maximum number of hops made by a message to reach its destination.
- **Node connectivity** : minimum number of vertices that must be removed to disconnect the graph.  
In other words: how many nodes can fail before we can expect connectivity to be lost.
- **Link connectivity** : minimum number of links that must be removed to disconnect the graph.  
In other words: how many links (cables...) can be broken before we can expect connectivity to be lost.
- **Fault diameter** : diameter of the graph, given the maximum number of failure before the graph becomes bipartite ( $\kappa - 1$ , if  $\kappa$  is the node connectivity)  
In other words: how the diameter evolves in degraded conditions.



Regular graphs have some symmetries

- Some properties can be extracted to simplify the computation of these metrics
- e.g., to compute the diameter

Here we are talking about **irregular graphs**

- No such property to simplify the computation

Large-scale graphs

- Large number of nodes!
- The complexity of the algorithms matters a lot

This algorithm answers two major questions:

- Is our set of vertices/nodes connected or disconnected?
- How many connected components do we have?

Algorithm: based of the Breadth First Search (BFS) algorithm.

- 1: Start with the first non visited vertex
- 2: Visit its connected component
- 3: Restart until there is no more vertex to visit (*i.e.*,all vertices have been visited)
- 4: The number of times we do step 3 = number of connected components

If all the vertices are visited during the first pass of the algorithm: the graph is connected.

This algorithm answers two major questions:

- Is our set of vertices/nodes connected or disconnected?
- How many connected components do we have?

Algorithm: based of the Breadth First Search (BFS) algorithm.

- 1: Start with the first non visited vertex
- 2: Visit its connected component
- 3: Restart until there is no more vertex to visit (*i.e.*,all vertices have been visited)
- 4: The number of times we do step 3 = number of connected components

If all the vertices are visited during the first pass of the algorithm: the graph is connected.


Naive approach:

```
1: for  $i = 1$  to  $n$  do  
2:   for NumChoice = 1 to MaxChoiceVertices do  
3:     choose  $i$  vertices among  $n$ ;  
4:     cancel the vertices;  
5:     if graph is not connected then  
6:       return  $i$   
7:     end if  
8:   end for  
9: end for
```

Problem: exponential complexity!

Dichotomous approach:

- Remove  $n/2$  vertices
- Check if the graph is still connected
- Remove more or less vertices depending on the previous try




**Require:** Low = 1 ; Up = MAX.

**Require:** (MAX = NbVertices or MinDegree).

```
1: set  $K = \log n$ 
2: while Up - Low > 1 do
3:   m = (Low + Up) / 2
4:   repeat
5:     Remove m vertices (or edges)
       randomly
6:     if graph not connected then
7:       Up = m
8:     end if
9:   until K times
10:  if all graphs are connected then
11:    Low = m
12:  end if
13: end while
14: return Low + 1
```

Dichotomous approach:

- Remove  $n/2$  vertices
- Check if the graph is still connected
- Remove more or less vertices depending on the previous try



**Require:** Low = 1 ; Up = MAX.

**Require:** (MAX = NbVertices or MinDegree).

```
1: set  $K = \log n$ 
2: while  $Up - Low > 1$  do
3:    $m = (Low + Up) / 2$ 
4:   repeat
5:     Remove  $m$  vertices (or edges)
       randomly
6:     if graph not connected then
7:        $Up = m$ 
8:     end if
9:   until  $K$  times
10:  if all graphs are connected then
11:     $Low = m$ 
12:  end if
13: end while
14: return  $Low + 1$ 
```

Doing an extensive enumeration of all the vertices of the graph would be too expensive

- Randomized algorithm

```
1: set Diameter = 0
2: repeat
3:   Select randomly a vertex, name it 'current vertex'
4:   and mark it as visited. Set 'Current diameter' to 0
5:   while current vertices have non visited neighbors do
6:     a) Compute the non visited neighbors of current vertices
7:     b) Replace the current vertices by their non visited vertices
8:     c) Add 1 to 'Current diameter'
9:   end while
10:  if 'Current diameter' > 'Diameter' then
11:    'Diameter' = 'Current diameter'
12:  end if
13: until 'some' vertices have been visited
14: return 'Diameter'
```

Several benchmark files (see our paper for all 4 of them)

Topology of the Internet: obtained from a campaign of traceroute calls and dedicated tools (WebGraph, eDonkey, MetroSec).

- web: 1 719 037 vertices and 11 095 298 edges
- p2p: 5 792 297 vertices and 142 038 401 edges
- web: 39 459 925 web pages (vertices) and 783 027 125 links (edges)
- ip: 2 250 498 vertices and 19 394 216 edges

Characteristic of the Internet: end-users have a degree of 1.



## Comparison with bounds given by Magnien et al

Table: Comparison between estimated diameters

	Magnien et al's results tlb - dslb - hdtub - rtub - tub	Our results
inet	29-31-34-34-38	25 (24)
p2p	8-9-10-10-10	8 (7)
web	26-32-33-33-34	22 (23)
ip	9-9-9-9-10	8 (7)

Table: Metrics for Fault Tolerance ( $\delta_{min} > 1$ )

	Link co.	Node co.	Fault diameter
inet	2	18	24
p2p	2	1054	7
web	2	36	23
ip	2	391	7

- Computing metrics such as diameter and node connectivity of large-scale, irregular graphs is too expensive to be done extensively
- Probabilistic algorithms to approach these values
- Our tool can be used for any graph, no assumption on the topology

Conclusions on the resilience of the Internet

- Quite resilient (was the initial goal)
- End users are still subject to disconnections