Lecture 6: Emulating a faulty mesh

30th September and 1st October, 2008

In the last lecture, we talked about how to emulate one network on another by constructing embeddings with low load, congestion and dilation. In this lecture we apply this embedding technique to emulate a fault-free square mesh on a faulty square mesh. Up to certain fault parameters we can achieve embeddings with constant load and congestion and small dilation. This allows us to use the result of the previous lecture to say that our we have an emulation strategy with low slowdown. The material in this lecture is taken from [1].

6.1 Introduction: Mesh emulation

An embedding is used to map nodes of a fault free network to nodes of faulty network so that all the connections of the fault free network can be simulated in the faulty network. This enables us to operate on a faulty network as if it was fault free. The fault free network acts as the guest and the faulty network acts as the host in the embedding.

$$G \Longrightarrow G_f$$
$$Guest \Longrightarrow Host$$

We will consider both a random fault model and adversarial fault model and will prove that in either scenario, mesh is fault tolerant. We also show that a constant slowdown can be achieved in the network for limited fault probability in random fault model or limited no. of faults in adversarial faulty model. We will consider the special case of mesh networks.
In the random model, we assume that each processor fails independently with some probability $p < p\prime$, where $p\prime$ is a small constant. We refer to a mesh with such failures as a $p-$faulty mesh. In the worst-case model, we assume that an adversary chooses $k < n$ processors to fail. We will study the following *parameters*:

- *Congestion C*: Maximum number of guest network edges using same edge in the host network.

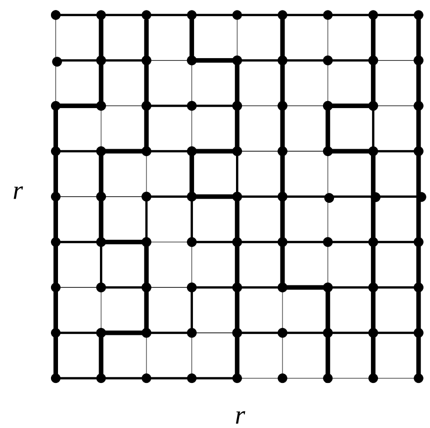

Figure 1: $r \times r$submesh of (α, r)gridlike mesh

- *Dilation D*: Maximum length of a path in the host network that maps to an edge in the guest network.
- *Load L*: Maximum number of nodes mapped to a single node of host Graph.

These parameters decide the slow-down factor for the embedding. One unit of time in the guest network can be simulated as $O(L+D+C)$ units of time in the simulated network. If a slowdown of at most r is required, then a component of size $\Theta(\frac{n}{r})$ must be present. If only one component had all the mapped nodes and less then n nodes are there in that component, then load would have exceeded r and hence slowdown would have been more than r.

Now, a faulty $n \times n$ mesh is said to be (α, r) gridlike if for every $r \times r$ submesh, there are at least $(1-\alpha)r$ vertex disjoint fault free paths connecting the left side to the right and least $(1-\alpha)r$ vertex disjoint fault free paths connecting the top side to the bottom, as shown in figure 1. Each of these paths must have length at most $2r$. Also, an $(\frac{1}{3}, r)$ grid like mesh is referred to as r-grid like mesh.

We observe that an (α, r) grid like mesh has at least $((1-\alpha)r)^2$ nodes which lie both on horizontal as well as vertical paths. Lets call these nodes as good points. In this way, $(\frac{1}{3}, r)$ grid like mesh has at least $\frac{4}{9}r^2$ good nodes in each $r \times r$ sub mesh.

Now, we will first prove that a fault free $n \times n$ mesh can be emulated on an r-grid like $n \times n$ mesh with $O(r)$ slowdown. Then we prove that in adversarial fault case, the mesh is grid like for a limited number of faults. Finally, we show that in randomized fault model, the mesh is (α, r) grid like with high probability. This leads us to show that a faulty mesh with limited fault probability can emulate a $n \times n$ fault free mesh with slow down $O(\log n)$ with high probability. This proves the applicability of embeddings in mesh structures with faults.

6.2 Embedding in r-grid like mesh

In this section, we will prove that an r-grid like mesh can emulate fault free mesh with $O(r)$ slow down.

Theorem 6.1 *An $n \times n$ mesh can be embedded into an $n \times n$ r-gridlike mesh with $O(1)$ load, $O(1)$ congestion and $O(r)$ dilation.*

Proof.
We emulate $\frac{n}{3} \times \frac{n}{3}$ fault free mesh and then increase load of each node by a factor of 9. The given $n \times n$ mesh can be divided into $r \times r$ non- overlapping sub meshes and the total number of such sub meshes possible are $O(n^2)$.

Now consider a single $r \times r$ sub mesh, it has at least $\frac{2}{3}r$ horizontal vertex

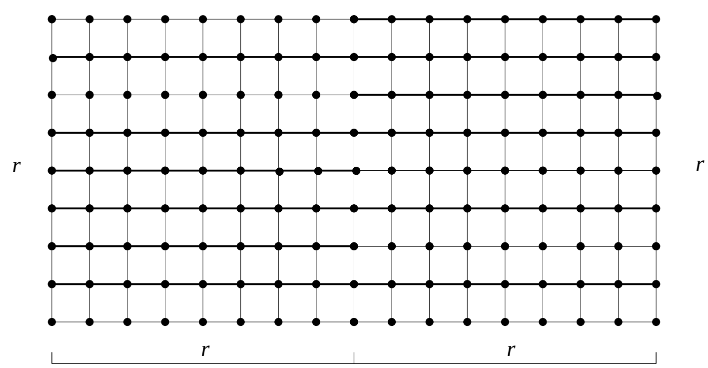

Figure 2: two adjacent $r \times r$ submeshes

disjoint path paths from left side to right side of the submesh as the given mesh is $r-$gridlike. Consider the $r \times r$ identical adjacent to this on right side, that submesh will also have $\frac{2}{3}r$ paths. So there will be at least $\frac{r}{3}$ vertex disjoint paths from nodes on left side of left submesh to the nodes on right side of right submesh. This set paths will exist between any two adjacent r x

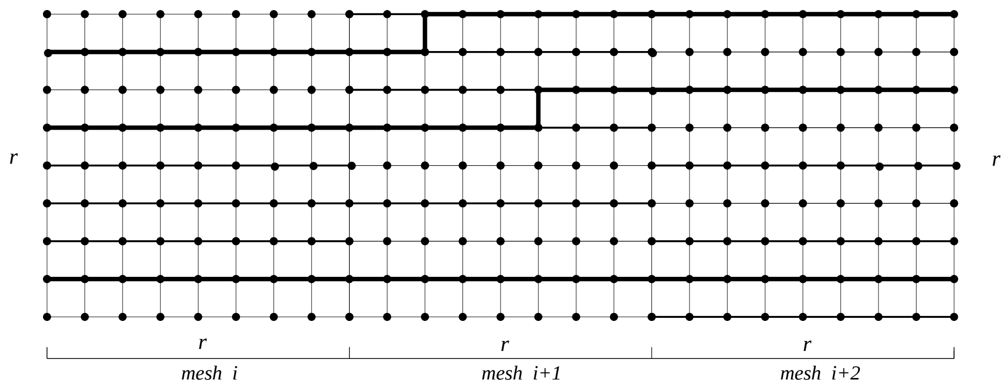

Figure 3: sequence of three $r \times r$ submeshes

r sub meshes and the length of these paths will be atmost 4r. Now consider a sequence of $r \times r$ submeshes from left side of $n \times n$ mesh to its right side and number the submeshes as $1, 2, 3, ..., \frac{n}{r}$. If we consider i, $i+1$ and $i+2$ meshes of this sequence and let A be the set of vertex disjoint paths from left side of mesh i to the right side of mesh $i+1$ and B be the set of vertex disjoint paths from right side of mesh $i+2$ to left side of mesh $i+1$. The mesh $i+1$ has $\frac{2}{3}r$ vertical vertex disjoint paths from top to bottom. With the help of these paths we can join the paths in set A with set B and can construct $\frac{r}{3}$ vertex disjoint paths from left side of mesh i to right side of mesh $i+2$, as shown in figure 3. The maximum length of the part of these newly constructed paths in mesh $i+1$ is $4r$, as vertex disjoint paths can be of length $2r$ in each horizontal and vertical directions. We can join $\frac{r}{3}$ nodes on left side of $n \times n$ to the right side of this mesh. The total length of these paths will be at most $4r \cdot \frac{n}{r} = 4n$.
Here, Highways refer to the sequences of $r \times r$ sub meshes and lanes refers to vertex disjoint paths from one side to its opposite side in an $n \times n$ mesh. On each highway, we have ensured $\frac{r}{3}$ lanes. We have $\frac{n}{r}$ highways in the mesh $r \times r$ mesh, so in all we have at least $\frac{n}{3}$ lanes in horizontal direction and also in vertical direction. These $\frac{n}{3}$ paths in each direction will ensure $\frac{n^2}{9}$ good nodes.
We can map the nodes of fault free mesh $\frac{n}{3} \times \frac{n}{3}$ to these good nodes. In this embedding we have ensured:

- *Load* $= 1$ (as node to node mapping is one-one)

- *Dilation* $= 8r$ (one edge of guest network may have to traverse 4r edges in adjacent submeshes)
- *Congestion* $= 2$ (an edge may be shared by at most one horizontal and one vertical path in host)

If we map 9 nodes of $n \times n$ to a single node of $\frac{n}{3} \times \frac{n}{3}$ mesh then we can emulate the $n \times n$ mesh using a r-gridlike mesh. Now the parameters will become:

- *Load* $= 9$
- *Dilation* $= 8r$
- *Congestion* $= 6$

In this way, we have obtained an embedding for $n \times n$ non faulty mesh on an r-gridlike faulty mesh with $O(r)$ slow down.

6.3 Adversarial Faults and (α, r) gridlike meshes

In this section, we prove that an $n \times n$ mesh is $(\alpha, \frac{k}{\alpha})$ gridlike even after k adversarial faults.

Lemma 6.1 *An $n \times n$ mesh with $k < n$ faults is $(\alpha, \frac{k}{\alpha})$ gridlike for any α such that $\frac{k}{n} < \alpha < 1$.*

Proof.
Consider a sub mesh of size $\frac{k}{\alpha} \times \frac{k}{\alpha}$ and assume that all the k faults are contained in this submesh. If we consider $\frac{k}{\alpha}$ rows and $\frac{k}{\alpha}$ columns of a fault free $\frac{k}{\alpha} \times \frac{k}{\alpha}$ mesh then it has $\frac{k}{\alpha}$ vertex disjoint paths in horizontal and vertical directions. Now, each faulty node can potentially remove a row and a column. So, k faults can remove k horizontal and k vertical paths from the fault free $\frac{k}{\alpha} \times \frac{k}{\alpha}$ submesh. As $\frac{k}{n} < \alpha < 1$, the submesh contains at least $\frac{k}{\alpha} - k$ horizontal and $\frac{k}{\alpha} - k$ vertical vertex disjoint paths. Hence, each $\frac{k}{\alpha} \times \frac{k}{\alpha}$ mesh has at least $\frac{k}{\alpha} - k = k(\frac{1}{\alpha} - 1) = \frac{k}{\alpha}(1 - \alpha)$ horizontal and vertical paths. Hence, the mesh is $(\alpha, \frac{k}{\alpha})$gridlike.

This proof helps us to appreciate the robustness of the grid like property.

6.4 Emulating in a Random Fault Model

In this section, we will figure out whether we can emulate a fault free model in a random faulty model. We have already shown that ideal mesh can be emulated in an (α, r) gridlike mesh with $O(r)$ slowdown under certain conditions. We now prove that a random fault model ensures that the mesh is (α, r) gridlike with high probability. Specifically, we prove that for a given value of α, $P(\alpha)$ exists such that for all $p < P(\alpha)$, a $p-$faulty mesh will be (α, r) gridlike with probability at least $1 - \frac{1}{n}$ where $r = O(\log_{\frac{1}{\alpha}} n)$. Now, the definition of (α, r) gridlike mesh requires $(1-\alpha)r$ vertex disjoint paths to be of length at most $2r$. We first establish a lower bound for the number of vertex disjoint paths of length at most $2r$ in an $r \times r$ mesh for given number of total number of vertex disjoint paths. Then, we use this lemma and Menger's Theorem to prove that a $p-$faulty mesh is (α, r) gridlike mesh with high probability.

Lets prove the following lemma.

Lemma 6.2 *If there are $(1 - \frac{\alpha}{2})r$ vertex disjoint paths from one side of an $r \times r$ mesh to the other side, then at least $(1-\alpha)r$ of them have length at most $2r$.*

Proof.
Let βr be the number of vertex disjoint paths , with length at least $2r$, out of the given $(1 - \frac{\alpha}{2})r$ vertex disjoint paths. We also observe that any path from one side to the other side in an $r \times r$ mesh will have length at least r. Now, the number of nodes involved in the vertex disjoint paths are at least $(1 - \frac{\alpha}{2}\beta)r \cdot r + 2\beta r \cdot r$. But the total number of nodes in an $r \times r$ mesh is at most r^2. Therefore, $(1 - \frac{\alpha}{2} - \beta)r^2 \leq r^2$. This implies that $1 - \frac{\alpha}{2} - \beta \leq 1$. Hence, $\beta \leq \frac{\alpha}{2}$.

So among the given set of $(1 - \frac{\alpha}{2})r$ vertex disjoint paths, at most $(\frac{\alpha}{2})r$ paths have length more than $2r$. Hence, at least $(1 - \frac{\alpha}{2})r - \frac{\alpha}{2}r$ paths out of the given set have length at most $2r$. Hence proved.

Theorem 6.2 *Given a constant $\alpha > 0$, there is a constant $p(\alpha) > 0$ such that a p-faulty $n \times n$ mesh is (α, r) gridlike with probability at least $1 - \frac{1}{n}$ where $r = O(\log_{\frac{1}{p}} n)$ and $p < p(\alpha)$.*

Proof. We prove that if the given faulty mesh is not (α, r) gridlike, then there will be some $r \times r$ submesh which will have less than $(1-\alpha)r$ vertex disjoint paths in any direction. The lemma proved above will then provide an upper limit on total number of vertex disjoint paths in the $r \times r$ mesh. Then an application of Menger's Theorem provides an upper bound on number of non-faulty vertices which is translated into probability that a mesh is not (α, r) gridlike.

Consider an $r \times r$ submesh in the given $p-$faulty mesh. Lets consider events P_1 and P_2 as given below:

P_1: There are $(1 - \frac{\alpha}{2})r$ vertex disjoint paths from one side of an $r \times r$ mesh to the other side.

P_2: There are at least $(1 - \alpha)r$ vertex disjoint paths from one side of an $r \times r$ mesh to the other side that have length at most $2r$.

Now, the above lemma states that P_1 implies P_2. By contraposition, $\neg P_2 \rightarrow \neg P_1$ which implies that $Pr(\neg P_2) \leq Pr(\neg P_1)$.

If the mesh is not $(\alpha, r)-$gridlike, then there are less than $(1 - \alpha)r$ vertex disjoint paths of length at most $2r$ from one side to another side. It is $\neg P_2$ and implies that there cannot be $(1 - \frac{\alpha}{2})r$ vertex disjoint paths of any length from one side to another side which is $\neg P_1$.

Now, we will apply Menger's Theorem. The result establishes relation between number of vertex disjoint paths and minimum number of nodes to be removed from the graph to disconnect it. Two sets A and B are said to be k-connected if removal of at most $k - 1$ vertices cannot disconnect them. The Menger's Theorem is as follows :

Theorem 6.3 *Two sets A and B are k-connected iff there are k vertex disjoint paths between A and B.*

Applying Menger's Theorem to the current problem, we observe that if a mesh is not (α, r)-gridlike, then there exists a set of at most $(1 - \frac{\alpha}{2})r$ vertices whose removal disconnects the two sides. This implies that there is some cut along which number of non-faulty vertices is less than $(1 - \frac{\alpha}{2})r$. In other words, there will be a set of fewer than $(1 - \frac{\alpha}{2})r$ non-faulty nodes whose removal will disconnect the mesh. That cut will be the bottleneck for the number of vertex disjoint paths. Any cut that disconnects left side from the right side of the $r \times r$ submesh will contain at least one node in each row. Also, we observe that any cut of a mesh contains a set of nodes that forms a simple path from the top to bottom. These two properties allow us to consider only vertical cuts allowing 45 degrees of freedom (If i^{th} vertex was chosen in row j then one out of $(i, i-1, i+1)th$ vertices can be considered to be part of the cut in $(j+1)th$ row).

Total number of such cuts $= r \cdot 3^{r-1}$ Lets try to bound the probability that such a vertical cut exists which has at most $(1-\frac{\alpha}{2})$non-faulty nodes.

Now, assume that $\frac{2ep}{\alpha} < 1$.
Let Pr(A cut is bad) $= p_1$. Then,

$$\begin{aligned}
p_1 &\leq \sum_{k=1}^{(1-\frac{\alpha}{2})r} \binom{r}{k} (1-p)^k p^{r-k} \\
&\leq \sum_{k=1}^{(1-\frac{\alpha}{2})r} \binom{r}{k} p^{r-k} \\
&\leq \sum_{k=1}^{(1-\frac{\alpha}{2})r} \left(\frac{epr}{r-k}\right)^{r-k} \\
&\leq \sum_{k=1}^{(1-\frac{\alpha}{2})r} \left(\frac{epr}{r-(1-\frac{\alpha}{2})r}\right)^{r-k} \\
&\leq \sum_{k=1}^{(1-\frac{\alpha}{2})r} \left(\frac{2ep}{\alpha}\right)^{r-k} \\
&\leq c \cdot \left(\frac{2ep}{\alpha}\right)^{r}
\end{aligned}$$

Here, the last step follows from the assumption made above.

Pr(There exists a bad cut) $\leq r3^r c(\frac{2ep}{\alpha})^r$
$= cr(\frac{6ep}{\alpha})^r$ where $\frac{6ep}{\alpha} < 1$ or $p < \frac{\alpha}{6e}$

The above discussion was based upon cut within an $r \times r$ submesh. Lets consider the complete $n \times n$ mesh. The number of $r \times r$ submeshes present in an $n \times n$ mesh is $O(n^2)$. So, Pr($n \times n$ mesh is not (α, r) grid like) = Pr(at least one $r \times r$ submesh in the given $n \times n$ mesh has a bad cut) i.e. at most (number of total $r \times r$ submeshes)(Pr(a given $r \times r$ submesh has a bad cut) which is at most $rn^2(\frac{6ep}{\alpha})^r$.

If this last term is $\leq \frac{1}{n}$, then the given mesh will be (α, r) grid like with probability at least $1 - \frac{1}{n}$. This condition is satisfied when $r = O(\log_{\frac{1}{p}} n)$. Hence the theorem has been proved.

6.5 Related Results

Theorem 6.4 *With probability at least* $1-\frac{1}{n}$*, a p-faulty* $n\times n$ *mesh with load O(1), congestion O(1) and dilation* $O(\log_{\frac{1}{p}} n)$ *i.e. with slow down* $O(\log_{\frac{1}{p}} n)$*.*

This result is obtained by combining Theorem 6.1 and Theorem 6.2. Theorem 6.2 shows that a p-faulty mesh is $(\frac{1}{3}, r)$-grid like with probability at least $1 - \frac{1}{n}$ and Theorem 6.1 shows that non-faulty $n \times n$ can be emulated on an r-gridlike mesh with $O(r)$ dilation.

Theorem 6.5 *An* r*-gridlike* $n\times n$ *mesh can emulate a fault free* $m\times m$ *mesh with* $O(r + (\frac{m}{n})^2)$ *slowdown.*

We first emulate $n \times n$ on the r-gridlike mesh using Theorem 6.1 and then map nodes of $m \times m$ mesh to the nodes of $n \times n$ non faulty mesh. This second mapping makes the total load as $O(\frac{m}{n})^2)$.Hence, in final embedding, load is $(\frac{m}{n})^2$, dilation is r and congestion is $\frac{m}{n}$.

References

[1] C. Kaklamanis, A.R. Karlin, F.T. Leighton, V. Milenkovic, P. Raghavan, S.B. Rao, and A. Tsantilas. Asymptotically tight bounds for computing with faulty arrays of processors. In *Proc. of the 22nd ACM Symp. on Theory of Computing (STOC)*, pages 285–296, 1990.