Spatial Coupling and the Threshold Saturation Phenomenon

Shrinivas Kudekar
Qualcomm Research

Ruediger Urbanke
EPFL

ISIT 2013 ISTANBUL

QUALCOMM®

July 7th, 2013

ECOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

The latest version of these slides (Keynote and PDF) can be found at
https://ipg.epfl.ch/doku.php?id=en:publications:scc_tutorial

Part III: Threshold saturation - proof by pictures

Threshold Saturation via Spatial Coupling: Why Convolutional LDPC Ensembles Perform so well over the BEC

Wave-Like Solutions of General One-Dimensional Spatially Coupled Systems

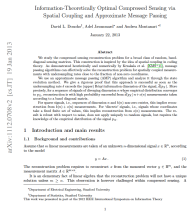
Information-Theoretically Optimal Compressed Sensing via Spatial Coupling and Approximate Message Passing

A Simple Proof of Threshold Saturation for Coupled Scalar Recursions

A Potential Theory of General Spatially-Coupled Systems via a Continuum Approximation

Let us now see how the threshold saturation phenomenon can be proved. We will only discuss the case of the BEC. In this relatively simple case we can argue mostly in a graphical way and the area theorem has a particularly nice graphical interpretation. The proof for general channels goes along similar lines but is technically more difficult.

Even more generally, currently there are proofs of the threshold saturation phenomenon known for the following cases.
1. Sparse graph codes and transmission over any BMS channel.
2. Any system whose state (for the uncoupled system) is a scalar or a vector.
3. Compressive sensing.

We cover in the following slides the methods for 1 and 2. The proofs for compressive sensing involve many other components (in particular a discussion of the approximate message-passing algorithm) which go far beyond what we can cover here. For further details we point the reader to http://arxiv.org/abs/1112.0708.

For the BEC there are currently three approaches for the proof. They of course share some important features but are nevertheless somewhat different.

Maxwell Construction: The Hidden Bridge between Iterative and Maximum a Posteriori Decoding

Threshold Saturation via Spatial Coupling: Why Convolutional LDPC Ensembles Perform so well over the BEC

Proof via Maxwell Construction

We begin with the proof of threshold saturation via the Maxwell construction using the EXIT functions. Historically speaking, this is the first proof that spatially coupled codes achieving capacity under BP decoding when transmitting over the BEC. Later on the same approach led to the proof that spatially coupled codes universally achieving capacity under BP decoding over the whole class of BMS channels. We limit our discussion here to the BEC. In the last part of this tutorial we will touch upon the general case. The details of the proof for the BEC can be found in http://arxiv.org/pdf/1001.1826.pdf, whereas the general case is described in http://arxiv.org/pdf/1004.3742.pdf.

Proof via Maxwell Construction

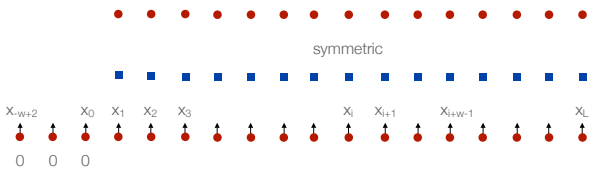

the x (erasure fraction) is a function of the position i

$$x_i = \epsilon\Big(1 - \frac{1}{w}\sum_{j=0}^{w-1}\big(1 - \frac{1}{w}\sum_{k=0}^{w-1} x_{i+j-k}\big)^{d_r-1}\Big)^{d_l-1} \qquad x_i = 0 \quad \forall\, i \notin [1, L]$$

Starting at what channel parameter does this recursion have non-trivial fixed points?

The answer is of course ϵ^{Area} as we have seen in Part I.

The starting point for the EXIT function analysis is the investigation of FPs of the (coupled) DE equations. As shown in the slides above, the DE equations for spatially coupled ensembles are multi-dimensional instead of scalar as for the uncoupled case. Recall that we are interested in finding the largest channel parameter ϵ so that the recursion, when started with the all-one vector inside the range [1, L], converges to the all-zero vector. We want to prove that this parameter is ϵ^{Area}, the area threshold.

Proof via Maxwell Construction

A word of warning: In some of the figures the constellation goes from 1 to L and in some other ones the indices are chosen symmetrical around zero. This index shift has clearly no bearing on the behavior of the system and we hope that it does not cause confusions.

The proof has three parts. The first ingredient in the proof is to show the existence of a special FP of the coupled DE equations as shown in
the slides. The second ingredient is to prove that any such FP must have a channel parameter which is very close to the area threshold.
The final part consists in proving that this implies that if we transmit below the area threshold that the DE recursions converge to the all-zero constellation.

Let us discuss each of these statements in more detail.

Proof via Maxwell Construction

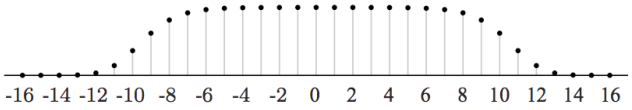

A word of warning: In some of the figures the constellation goes from 1 to L and in some other ones the indices are chosen symmetrical around zero. This index shift has clearly no bearing on the behavior of the system and we hope that it does not cause confusions.

The proof has three parts. The first ingredient in the proof is to show the existence of a special FP of the coupled DE equations as shown in
the slides. The second ingredient is to prove that any such FP must have a channel parameter which is very close to the area threshold.
The final part consists in proving that this implies that if we transmit below the area threshold that the DE recursions converge to the all-zero constellation.

Let us discuss each of these statements in more detail.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

A word of warning: In some of the figures the constellation goes from 1 to L and in some other ones the indices are chosen symmetrical around zero. This index shift has clearly no bearing on the behavior of the system and we hope that it does not cause confusions.

The proof has three parts. The first ingredient in the proof is to show the existence of a special FP of the coupled DE equations as shown in
the slides. The second ingredient is to prove that any such FP must have a channel parameter which is very close to the area threshold.
The final part consists in proving that this implies that if we transmit below the area threshold that the DE recursions converge to the all-zero constellation.

Let us discuss each of these statements in more detail.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

Saturation:

show that any “such” FP must have a channel parameter very close to the area threshold

A word of warning: In some of the figures the constellation goes from 1 to L and in some other ones the indices are chosen symmetrical around zero. This index shift has clearly no bearing on the behavior of the system and we hope that it does not cause confusions.

The proof has three parts. The first ingredient in the proof is to show the existence of a special FP of the coupled DE equations as shown in
the slides. The second ingredient is to prove that any such FP must have a channel parameter which is very close to the area threshold.
The final part consists in proving that this implies that if we transmit below the area threshold that the DE recursions converge to the all-zero constellation.

Let us discuss each of these statements in more detail.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

Saturation:

show that any “such” FP must have a channel parameter very close to the area threshold

Convergence:

show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

A word of warning: In some of the figures the constellation goes from 1 to L and in some other ones the indices are chosen symmetrical around zero. This index shift has clearly no bearing on the behavior of the system and we hope that it does not cause confusions.

The proof has three parts. The first ingredient in the proof is to show the existence of a special FP of the coupled DE equations as shown in
the slides. The second ingredient is to prove that any such FP must have a channel parameter which is very close to the area threshold.
The final part consists in proving that this implies that if we transmit below the area threshold that the DE recursions converge to the all-zero constellation.

Let us discuss each of these statements in more detail.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

Let us discuss how the existence of such a special FP can be proved. Let us first describe what we are looking for. Recall what we mean by a FP. We call a vector x whose components are the erasure fractions at the various indices a *constellation*. Recall that we have a special boundary condition. At all the indices outside the constellation we assume that the corresponding x value is 0, i.e., we have perfect knowledge. A constellation x which when inserted into the DE equations results again in x is called a FP of DE.

The special FP whose existence is established has the following properties: the FP is unimodal; further, towards the boundary the x_i values are very close to 0 whereas towards the middle the constellation takes on an essentially constant value; this value is equal to the FP value of the uncoupled system for the same channel parameter; finally, the transition between these two regions is “quick”; this means that the number of positions where the x_i value is in in the range [delta, 1-delta] x_{middle} is O(w), where the delta value only enters in the constant implied in the O(w) value and this constant is finite for all strictly positive delta values.

Thus, the special FP “couples” the zero value to the largest value of FP of DE (of the uncoupled system). The first part of the proof is to show the existence of such an FP of coupled DE equations.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

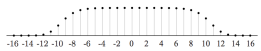

Let us discuss how the existence of such a special FP can be proved. Let us first describe what we are looking for. Recall what we mean by a FP. We call a vector x whose components are the erasure fractions at the various indices a *constellation*. Recall that we have a special boundary condition. At all the indices outside the constellation we assume that the corresponding x value is 0, i.e., we have perfect knowledge. A constellation x which when inserted into the DE equations results again in x is called a FP of DE.

The special FP whose existence is established has the following properties: the FP is unimodal; further, towards the boundary the x_i values are very close to 0 whereas towards the middle the constellation takes on an essentially constant value; this value is equal to the FP value of the uncoupled system for the same channel parameter; finally, the transition between these two regions is “quick”; this means that the number of positions where the x_i value is in in the range [delta, 1-delta] x_{middle} is O(w), where the delta value only enters in the constant implied in the O(w) value and this constant is finite for all strictly positive delta values.

Thus, the special FP “couples” the zero value to the largest value of FP of DE (of the uncoupled system). The first part of the proof is to show the existence of such an FP of coupled DE equations.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

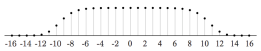

unimodal shape

Saturday, July 13, 13 8

Let us discuss how the existence of such a special FP can be proved. Let us first describe what we are looking for. Recall what we mean by a FP. We call a vector x whose components are the erasure fractions at the various indices a *constellation*. Recall that we have a special boundary condition. At all the indices outside the constellation we assume that the corresponding x value is 0, i.e., we have perfect knowledge. A constellation x which when inserted into the DE equations results again in x is called a FP of DE.

The special FP whose existence is established has the following properties: the FP is unimodal; further, towards the boundary the x_i values are very close to 0 whereas towards the middle the constellation takes on an essentially constant value; this value is equal to the FP value of the uncoupled system for the same channel parameter; finally, the transition between these two regions is “quick”; this means that the number of positions where the x_i value is in in the range [delta, 1-delta] x_{middle} is O(w), where the delta value only enters in the constant implied in the O(w) value and this constant is finite for all strictly positive delta values.

Thus, the special FP “couples” the zero value to the largest value of FP of DE (of the uncoupled system). The first part of the proof is to show the existence of such an FP of coupled DE equations.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

many values very close to 0

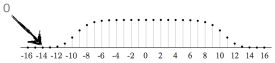

unimodal shape

Saturday, July 13, 13 8

Let us discuss how the existence of such a special FP can be proved. Let us first describe what we are looking for. Recall what we mean by a FP. We call a vector x whose components are the erasure fractions at the various indices a *constellation*. Recall that we have a special boundary condition. At all the indices outside the constellation we assume that the corresponding x value is 0, i.e., we have perfect knowledge. A constellation x which when inserted into the DE equations results again in x is called a FP of DE.

The special FP whose existence is established has the following properties: the FP is unimodal; further, towards the boundary the x_i values are very close to 0 whereas towards the middle the constellation takes on an essentially constant value; this value is equal to the FP value of the uncoupled system for the same channel parameter; finally, the transition between these two regions is “quick”; this means that the number of positions where the x_i value is in in the range [delta, 1-delta] x_{middle} is O(w), where the delta value only enters in the constant implied in the O(w) value and this constant is finite for all strictly positive delta values.

Thus, the special FP “couples” the zero value to the largest value of FP of DE (of the uncoupled system). The first part of the proof is to show the existence of such an FP of coupled DE equations.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

flat towards the middle

many values very close to 0

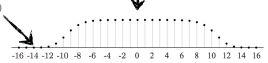

unimodal shape

Saturday, July 13, 13 8

Let us discuss how the existence of such a special FP can be proved. Let us first describe what we are looking for. Recall what we mean by a FP. We call a vector x whose components are the erasure fractions at the various indices a *constellation*. Recall that we have a special boundary condition. At all the indices outside the constellation we assume that the corresponding x value is 0, i.e., we have perfect knowledge. A constellation x which when inserted into the DE equations results again in x is called a FP of DE.

The special FP whose existence is established has the following properties: the FP is unimodal; further, towards the boundary the x_i values are very close to 0 whereas towards the middle the constellation takes on an essentially constant value; this value is equal to the FP value of the uncoupled system for the same channel parameter; finally, the transition between these two regions is “quick”; this means that the number of positions where the x_i value is in in the range [delta, 1-delta] x_{middle} is O(w), where the delta value only enters in the constant implied in the O(w) value and this constant is finite for all strictly positive delta values.

Thus, the special FP “couples” the zero value to the largest value of FP of DE (of the uncoupled system). The first part of the proof is to show the existence of such an FP of coupled DE equations.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

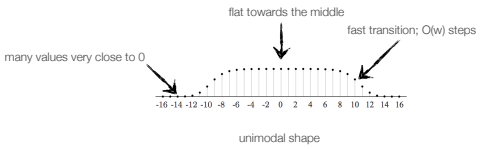

Let us discuss how the existence of such a special FP can be proved. Let us first describe what we are looking for. Recall what we mean by a FP. We call a vector x whose components are the erasure fractions at the various indices a *constellation*. Recall that we have a special boundary condition. At all the indices outside the constellation we assume that the corresponding x value is 0, i.e., we have perfect knowledge. A constellation x which when inserted into the DE equations results again in x is called a FP of DE.

The special FP whose existence is established has the following properties: the FP is unimodal; further, towards the boundary the x_i values are very close to 0 whereas towards the middle the constellation takes on an essentially constant value; this value is equal to the FP value of the uncoupled system for the same channel parameter; finally, the transition between these two regions is “quick”; this means that the number of positions where the x_i value is in in the range [delta, 1-delta] x_{middle} is O(w), where the delta value only enters in the constant implied in the O(w) value and this constant is finite for all strictly positive delta values.

Thus, the special FP “couples” the zero value to the largest value of FP of DE (of the uncoupled system). The first part of the proof is to show the existence of such an FP of coupled DE equations.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

So how do we show the existence? The basic idea is to show the existence of such an FP *which has an average erasure fraction that is neither too large nor too small*. If such a FP exists, then if the chain is sufficiently long, there cannot be too many sections with an erasure fraction close to zero, since the average will not be the desired entropy. Similarly, there cannot be too many sections with erasure value being too large since the average will then be larger than the desired entropy. As a further simplification we can consider *one-sided* such FPs, i.e., we only consider one half of the constellation. To show the existence of such an FP we consider a sub-space $(x_{-L}, ..., x_0)$ of $[0,1]^{L+1}$, where L is the length of the chain, which is ordered and such that the $x_{-L} < ... < x_0$. Furthermore, the average erasure probability of the constellation is equal to some desired value. Then, we define a map which is applied on this space. This map is an essentially the DE map but where the channel is kept “free.” This means, we first apply the DE operator (without the effect of the channel) to a constellation and then we choose the ``appropriate” channel parameter so as to fulfill the average erasure constraint. We now use a standard FP theorem due to Cauty to show that such a map must have at least one FP.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

consider the density evolution map with a “free” channel parameter for increasing one-sided constellations 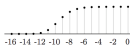

So how do we show the existence? The basic idea is to show the existence of such an FP *which has an average erasure fraction that is neither too large nor too small*. If such a FP exists, then if the chain is sufficiently long, there cannot be too many sections with an erasure fraction close to zero, since the average will not be the desired entropy. Similarly, there cannot be too many sections with erasure value being too large since the average will then be larger than the desired entropy. As a further simplification we can consider *one-sided* such FPs, i.e., we only consider one half of the constellation. To show the existence of such an FP we consider a sub-space $(x_{-L}, ..., x_0)$ of $[0,1]^{L+1}$, where L is the length of the chain, which is ordered and such that the $x_{-L} < ... < x_0$. Furthermore, the average erasure probability of the constellation is equal to some desired value. Then, we define a map which is applied on this space. This map is an essentially the DE map but where the channel is kept “free.” This means, we first apply the DE operator (without the effect of the channel) to a constellation and then we choose the ``appropriate” channel parameter so as to fulfill the average erasure constraint. We now use a standard FP theorem due to Cauty to show that such a map must have at least one FP.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

consider the density evolution map with a “free” channel parameter for increasing one-sided constellations 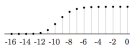

Use Schauder FP Theorem (Cauty):
every continuous map f from a convex compact subset of a topological vector space to itself has a FP

So how do we show the existence? The basic idea is to show the existence of such an FP *which has an average erasure fraction that is neither too large nor too small*. If such a FP exists, then if the chain is sufficiently long, there cannot be too many sections with an erasure fraction close to zero, since the average will not be the desired entropy. Similarly, there cannot be too many sections with erasure value being too large since the average will then be larger than the desired entropy. As a further simplification we can consider *one-sided* such FPs, i.e., we only consider one half of the constellation. To show the existence of such an FP we consider a sub-space $(x_{-L}, ..., x_0)$ of $[0,1]^{L+1}$, where L is the length of the chain, which is ordered and such that the $x_{-L} < ... < x_0$. Furthermore, the average erasure probability of the constellation is equal to some desired value. Then, we define a map which is applied on this space. This map is an essentially the DE map but where the channel is kept “free.” This means, we first apply the DE operator (without the effect of the channel) to a constellation and then we choose the ``appropriate” channel parameter so as to fulfill the average erasure constraint. We now use a standard FP theorem due to Cauty to show that such a map must have at least one FP.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

consider the density evolution map with a “free” channel parameter for increasing one-sided constellations

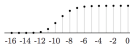

Use Schauder FP Theorem (Cauty):
every continuous map f from a convex compact subset of a topological vector space to itself has a FP

the “FP” here means therefore both the constellation as well as the channel parameter

So how do we show the existence? The basic idea is to show the existence of such an FP *which has an average erasure fraction that is neither too large nor too small*. If such a FP exists, then if the chain is sufficiently long, there cannot be too many sections with an erasure fraction close to zero, since the average will not be the desired entropy. Similarly, there cannot be too many sections with erasure value being too large since the average will then be larger than the desired entropy. As a further simplification we can consider *one-sided* such FPs, i.e., we only consider one half of the constellation. To show the existence of such an FP we consider a sub-space $(x_{-L}, \ldots, x_0)$ of $[0,1]^{L+1}$, where L is the length of the chain, which is ordered and such that the $x_{-L} < \ldots < x_0$. Furthermore, the average erasure probability of the constellation is equal to some desired value. Then, we define a map which is applied on this space. This map is an essentially the DE map but where the channel is kept “free.” This means, we first apply the DE operator (without the effect of the channel) to a constellation and then we choose the ``appropriate” channel parameter so as to fulfill the average erasure constraint. We now use a standard FP theorem due to Cauty to show that such a map must have at least one FP.

Proof via Maxwell Construction

Existence:

show that “such” a FP exists
“such”: small at the boundary, fast transition, not too small and flat in the middle

Saturation:

show that any “such” FP must have a channel parameter very close to the area threshold

Convergence:

show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to $\varepsilon^{\text{Area}}$. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the next few slides and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

Proof via Maxwell Construction

Saturday, July 13, 13 10

Saturday, July 13, 13 10

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to $\varepsilon^{\text{Area}}$. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the next few slides and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

Proof via Maxwell Construction

Saturday, July 13, 13 10

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to $\varepsilon^{\text{Area}}$. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the next few slides and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

Saturation

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to ε^{Area}. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the movie and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

The movie in this slide shows a dynamic look at the EXIT curve for the coupled code ensemble. Each point on the EXIT curve corresponds to a FP of DE. As the movie proceeds we see a point on the EXIT curve and we see the corresponding FP on the left. This picture serves as the guiding principle to show that the special FP must have a channel value close to the area threshold. To see this, focus on the vertical part of the EXIT function. Notice that as the point on the EXIT function drops vertically, the constellation on the left “moves” to the right. Indeed, what we show is that once we construct the special FP, obtained from the first part, any constellation got by moving the special FP inside, is also an FP of coupled DE having the **same** channel value as the special FP. This is seen from the picture when the constellation moves inside, the channel value in the EXIT function remains constant, hence the vertical drop. Furthermore, since the flat part of the FP is equal to the FP of forward DE of the underlying uncoupled ensemble, the EXIT function of the coupled ensemble follows the BP EXIT curve of the uncoupled ensemble till the channel value becomes equal to the channel value of the special FP. Finally, once the constellation has moved sufficiently inside, there is not much entropy remaining, and one can safely ignore this part for area analysis. Now, since the area under the EXIT curve is equal to the design rate of the code (irrespective of the ensemble under consideration) and the EXIT curve follows the BP EXIT curve of the uncoupled ensemble, it must be that the channel value of the special FP is equal to the area threshold of the uncoupled ensemble. The last assertion follows from the analysis presented in part II of this tutorial. Amazing, isn’t it :-)!

Saturation

Saturday, July 13, 13 11

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to ε^{Area}. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the movie and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

The movie in this slide shows a dynamic look at the EXIT curve for the coupled code ensemble. Each point on the EXIT curve corresponds to a FP of DE. As the movie proceeds we see a point on the EXIT curve and we see the corresponding FP on the left. This picture serves as the guiding principle to show that the special FP must have a channel value close to the area threshold. To see this, focus on the vertical part of the EXIT function. Notice that as the point on the EXIT function drops vertically, the constellation on the left “moves” to the right. Indeed, what we show is that once we construct the special FP, obtained from the first part, any constellation got by moving the special FP inside, is also an FP of coupled DE having the **same** channel value as the special FP. This is seen from the picture when the constellation moves inside, the channel value in the EXIT function remains constant, hence the vertical drop. Furthermore, since the flat part of the FP is equal to the FP of forward DE of the underlying uncoupled ensemble, the EXIT function of the coupled ensemble follows the BP EXIT curve of the uncoupled ensemble till the channel value becomes equal to the channel value of the special FP. Finally, once the constellation has moved sufficiently inside, there is not much entropy remaining, and one can safely ignore this part for area analysis. Now, since the area under the EXIT curve is equal to the design rate of the code (irrespective of the ensemble under consideration) and the EXIT curve follows the BP EXIT curve of the uncoupled ensemble, it must be that the channel value of the special FP is equal to the area threshold of the uncoupled ensemble. The last assertion follows from the analysis presented in part II of this tutorial. Amazing, isn’t it :-)!

Saturation

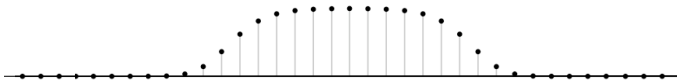

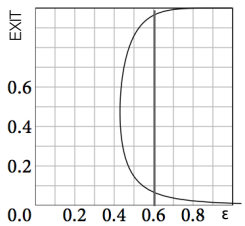

Saturday, July 13, 13 11

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to ε^{Area}. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the movie and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

The movie in this slide shows a dynamic look at the EXIT curve for the coupled code ensemble. Each point on the EXIT curve corresponds to a FP of DE. As the movie proceeds we see a point on the EXIT curve and we see the corresponding FP on the left. This picture serves as the guiding principle to show that the special FP must have a channel value close to the area threshold. To see this, focus on the vertical part of the EXIT function. Notice that as the point on the EXIT function drops vertically, the constellation on the left “moves” to the right. Indeed, what we show is that once we construct the special FP, obtained from the first part, any constellation got by moving the special FP inside, is also an FP of coupled DE having the **same** channel value as the special FP. This is seen from the picture when the constellation moves inside, the channel value in the EXIT function remains constant, hence the vertical drop. Furthermore, since the flat part of the FP is equal to the FP of forward DE of the underlying uncoupled ensemble, the EXIT function of the coupled ensemble follows the BP EXIT curve of the uncoupled ensemble till the channel value becomes equal to the channel value of the special FP. Finally, once the constellation has moved sufficiently inside, there is not much entropy remaining, and one can safely ignore this part for area analysis. Now, since the area under the EXIT curve is equal to the design rate of the code (irrespective of the ensemble under consideration) and the EXIT curve follows the BP EXIT curve of the uncoupled ensemble, it must be that the channel value of the special FP is equal to the area threshold of the uncoupled ensemble. The last assertion follows from the analysis presented in part II of this tutorial. Amazing, isn’t it :-)!

Saturation

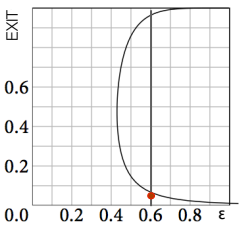

Saturday, July 13, 13 11

Now we want to show that IF we have such a special FP, then its channel parameter must be very close to ε^{Area}. The idea of showing this is the following. Given such a FP we construct a whole family “almost” FPs but with various EXIT parameters. We do this by essentially cutting out as little or as much from the middle part as we want. This gives us a whole family as shown in the movie and this family gives rise to an EXIT curve. We can then show that the area under this EXIT curve is very close to the rate of the code and also the part of the EXIT curve where we simply cut out the middle part must correspond to a vertical line. So we conclude that the line must be exactly where the Maxwell construction puts this line and so the position of the line is the area threshold.

The movie in this slide shows a dynamic look at the EXIT curve for the coupled code ensemble. Each point on the EXIT curve corresponds to a FP of DE. As the movie proceeds we see a point on the EXIT curve and we see the corresponding FP on the left. This picture serves as the guiding principle to show that the special FP must have a channel value close to the area threshold. To see this, focus on the vertical part of the EXIT function. Notice that as the point on the EXIT function drops vertically, the constellation on the left “moves” to the right. Indeed, what we show is that once we construct the special FP, obtained from the first part, any constellation got by moving the special FP inside, is also an FP of coupled DE having the **same** channel value as the special FP. This is seen from the picture when the constellation moves inside, the channel value in the EXIT function remains constant, hence the vertical drop. Furthermore, since the flat part of the FP is equal to the FP of forward DE of the underlying uncoupled ensemble, the EXIT function of the coupled ensemble follows the BP EXIT curve of the uncoupled ensemble till the channel value becomes equal to the channel value of the special FP. Finally, once the constellation has moved sufficiently inside, there is not much entropy remaining, and one can safely ignore this part for area analysis. Now, since the area under the EXIT curve is equal to the design rate of the code (irrespective of the ensemble under consideration) and the EXIT curve follows the BP EXIT curve of the uncoupled ensemble, it must be that the channel value of the special FP is equal to the area threshold of the uncoupled ensemble. The last assertion follows from the analysis presented in part II of this tutorial. Amazing, isn’t it :-)!

Saturation

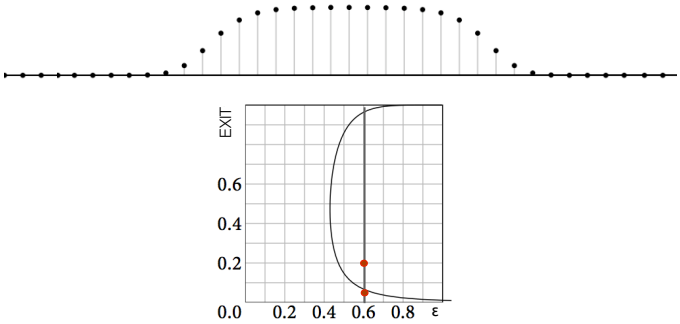

Saturation

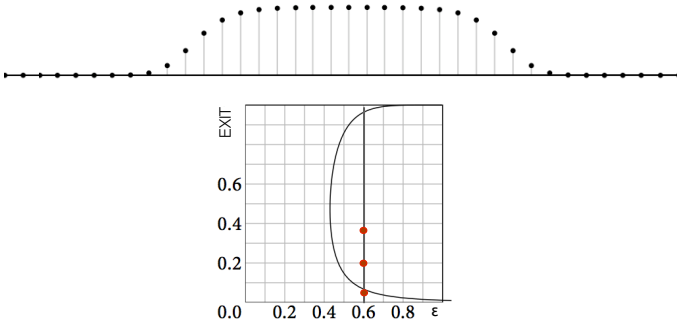

Saturation

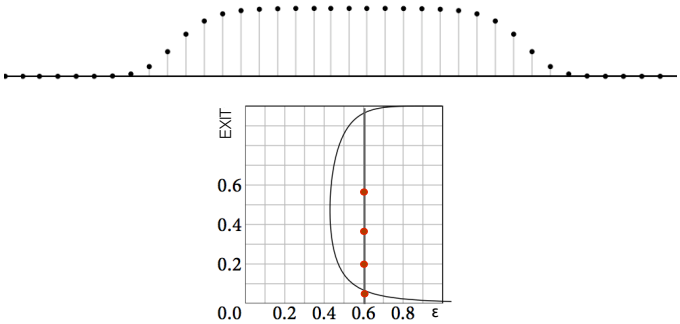

Saturation

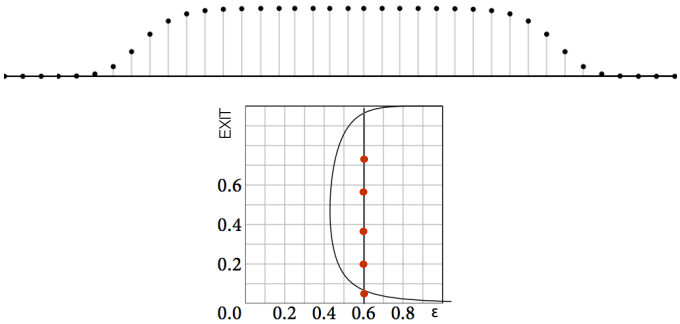

Saturation

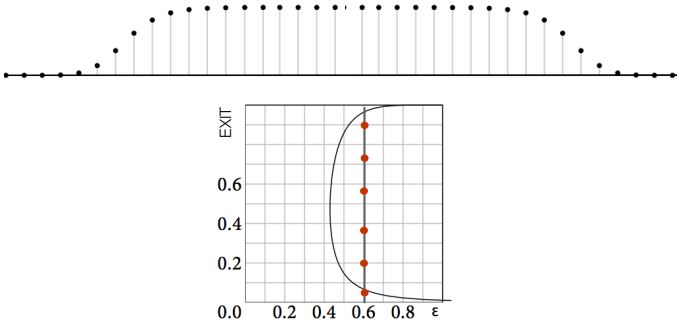

Saturation

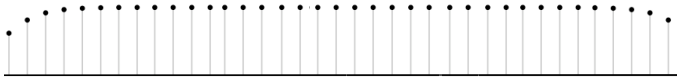

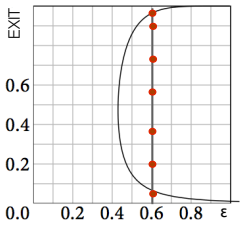

Saturation

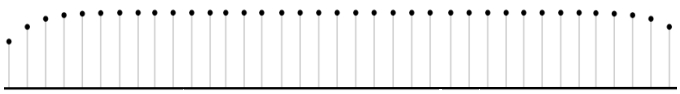

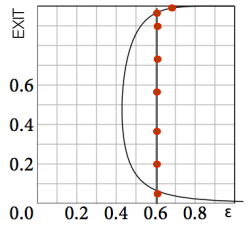

Saturation

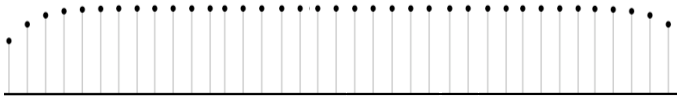

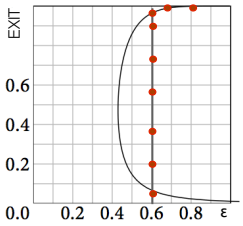

Saturation

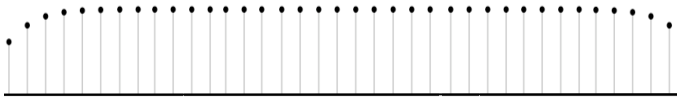

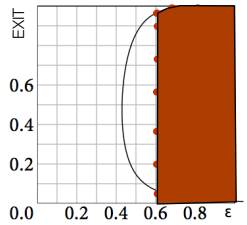

Saturation

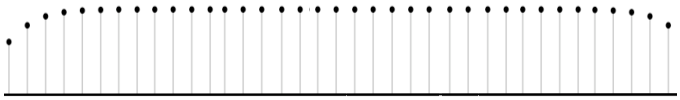

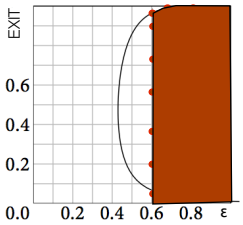

but any such approx. EXIT curve must have area close to the rate

$\rightarrow \epsilon \approx \epsilon^{\text{Area}}$

Saturation

The movie in this slide shows a dynamic look at the EXIT curve for the coupled code ensemble. Each point on the EXIT curve corresponds to a FP of DE. As the movie proceeds we see a point on the EXIT curve and we see the corresponding FP on the left. This picture serves as the guiding principle to show that the special FP must have a channel value close to the area threshold. To see this, focus on the vertical part of the EXIT function. Notice that as the point on the EXIT function drops vertically, the constellation on the left "moves" to the right. Indeed, what we show is that once we construct the special FP, obtained from the first part, any constellation got by moving the special FP inside, is also an FP of coupled DE having the **same** channel value as the special FP. This is seen from the picture when the constellation moves inside, the channel value in the EXIT function remains constant, hence the vertical drop. Furthermore, since the flat part of the FP is equal to the FP of forward DE of the underlying uncoupled ensemble, the EXIT function of the coupled ensemble follows the BP EXIT curve of the uncoupled ensemble till the channel value becomes equal to the channel value of the special FP. Finally, once the constellation has moved sufficiently inside, there is not much entropy remaining, and one can safely ignore this part for area analysis. Now, since the area under the EXIT curve is equal to the design rate of the code (irrespective of the ensemble under consideration) and the EXIT curve follows the BP EXIT curve of the uncoupled ensemble, it must be that the channel value of the special FP is equal to the area threshold of the uncoupled ensemble. The last assertion follows from the analysis presented in part II of this tutorial. Amazing, isn't it :-)!

Saturation

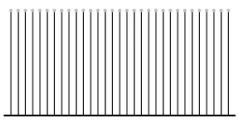

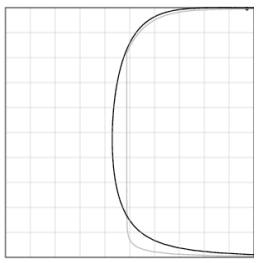

The movie in this slide shows a dynamic look at the EXIT curve for the coupled code ensemble. Each point on the EXIT curve corresponds to a FP of DE. As the movie proceeds we see a point on the EXIT curve and we see the corresponding FP on the left. This picture serves as the guiding principle to show that the special FP must have a channel value close to the area threshold. To see this, focus on the vertical part of the EXIT function. Notice that as the point on the EXIT function drops vertically, the constellation on the left "moves" to the right. Indeed, what we show is that once we construct the special FP, obtained from the first part, any constellation got by moving the special FP inside, is also an FP of coupled DE having the **same** channel value as the special FP. This is seen from the picture when the constellation moves inside, the channel value in the EXIT function remains constant, hence the vertical drop. Furthermore, since the flat part of the FP is equal to the FP of forward DE of the underlying uncoupled ensemble, the EXIT function of the coupled ensemble follows the BP EXIT curve of the uncoupled ensemble till the channel value becomes equal to the channel value of the special FP. Finally, once the constellation has moved sufficiently inside, there is not much entropy remaining, and one can safely ignore this part for area analysis. Now, since the area under the EXIT curve is equal to the design rate of the code (irrespective of the ensemble under consideration) and the EXIT curve follows the BP EXIT curve of the uncoupled ensemble, it must be that the channel value of the special FP is equal to the area threshold of the uncoupled ensemble. The last assertion follows from the analysis presented in part II of this tutorial. Amazing, isn't it :-)!

Proof via Maxwell Construction

Existence: 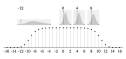 show that “such” a FP exists “such”: small at the boundary, fast transition, not too small and flat in the middle

Saturation: show that any “such” FP must have a channel parameter very close to the area threshold

Convergence: show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

Proof via Maxwell Construction

Convergence: show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

Saturday, July 13, 13 22

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

Proof via Maxwell Construction

Convergence: show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

pick $\varepsilon < \varepsilon^{Area}$ and run DE starting with all-one constellation
assume we get stuck in non-trivial FP

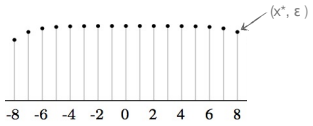

Saturday, July 13, 13 22

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

Proof via Maxwell Construction

Convergence: show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

pick $\varepsilon < \varepsilon^{Area}$ and run DE starting with all-one constellation
assume we get stuck in non-trivial FP

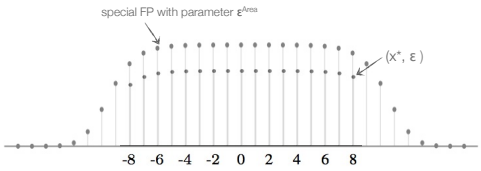

Saturday, July 13, 13 22

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

Proof via Maxwell Construction

Convergence:

show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

pick $\varepsilon < \varepsilon^{Area}$ and run DE starting with all-one constellation
assume we get stuck in non-trivial FP

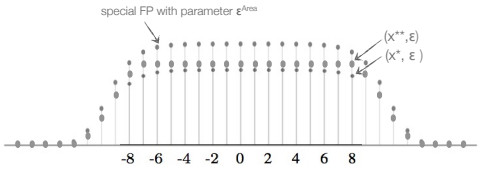

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

Proof via Maxwell Construction

Convergence:

show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

pick $\varepsilon < \varepsilon^{Area}$ and run DE starting with all-one constellation
assume we get stuck in non-trivial FP

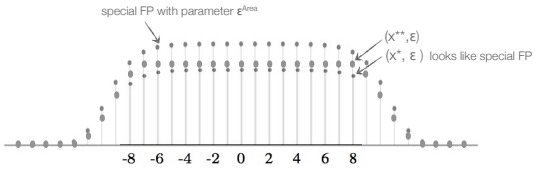

Saturday, July 13, 13 22

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

Proof via Maxwell Construction

Convergence:

show that this implies convergence to perfect decoding FP for a channel parameter below the area threshold

pick $\varepsilon < \varepsilon^{Area}$ and run DE starting with all-one constellation
assume we get stuck in non-trivial FP

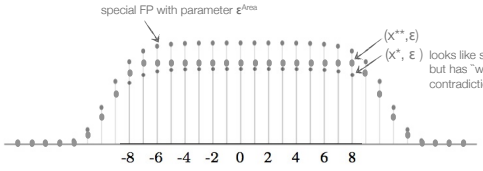

Saturday, July 13, 13 22

It remains to show that this implies that if we transmit at a channel parameter just a little bit below the area threshold then DE will converge to the all-zero constellation.

To see this, consider the following experiment. Take a constellation of size lets say L. Initialize it with the all-one vector inside the constellation and 0 outside and take a channel parameter $\varepsilon < \varepsilon^{Area}$. Run DE. DE produces a sequence of monotonically decreasing (point-wise) constellations which are bounded by 0 (again point-wise) from below. So DE must converge. Call the limit x^*. Assume that x^* is not the all-zero FP but a non-trivial one. Note that it is clear that at each point in the constellation the value of the FP can be no larger than the FP we would get for the uncoupled case if we used the channel ε.

Consider now this special FP whose existence we proved previously and which has an associated channel parameter very close to ε^{Area}. If we take such a FP over a large enough range then this FP will be point-wise strictly larger than x^*. Therefore if we apply DE to this special FP but with parameter ε instead of ε^{Area} then again DE will converge and it must converge to something non-trivial and no smaller than x^*. Further, this FP, call it x^{**}, fulfills all the conditions which we asked our special FP to fulfill. But we just saw in the previous slides that such a special FP can only exist when the channel parameter is very close to ε^{Area}, a contradiction.

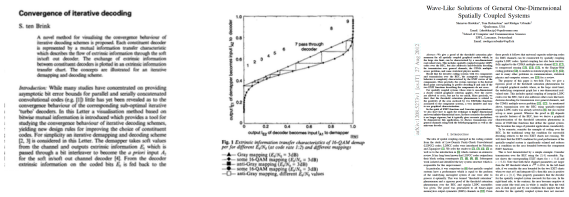

Proof via EXIT Charts

The title "Proof via EXIT Charts" is a little bit of misnomer. Even though the criterion is in terms of a balance of EXIT charts, the proof is in fact shares some similarities with the potential function approach, but it has its own twist.

Take Continuum Limit

Let us quickly explain. The first step of the proof consists of considering an appropriately chosen continuous version of this problem.

For the random coupled ensemble, we introduced the window, w, over which the random connections are made. This discrete system is hard to analyze. Instead, one can consider the limit when w goes to infinity. Of course, in this case the length of the chain has to go to infinity as well. If we increase the length w and scale the horizontal axis by the same proportion then in the limit we can think of the constellation as a continuous curve rather than a set of spikes. To this corresponds a DE equation where instead of taking sums to average we integrate over a finite support window.

Take Continuum Limit

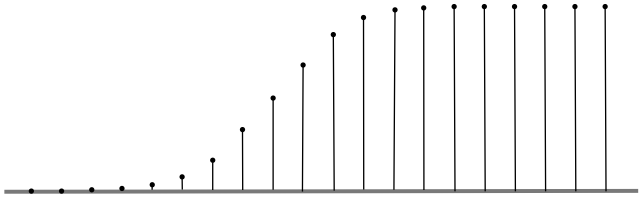

$$x_i = \epsilon\Big(1 - \frac{1}{w}\sum_{j=0}^{w-1}(1 - \frac{1}{w}\sum_{k=0}^{w-1} x_{i+j-k})^{d_r-1}\Big)^{d_l-1} \qquad x_i = 0 \quad \forall\ i \notin [1, L]$$

Let us quickly explain. The first step of the proof consists of considering an appropriately chosen continuous version of this problem.

For the random coupled ensemble, we introduced the window, w, over which the random connections are made. This discrete system is hard to analyze. Instead, one can consider the limit when w goes to infinity. Of course, in this case the length of the chain has to go to infinity as well. If we increase the length w and scale the horizontal axis by the same proportion then in the limit we can think of the constellation as a continuous curve rather than a set of spikes. To this corresponds a DE equation where instead of taking sums to average we integrate over a finite support window.

Take Continuum Limit

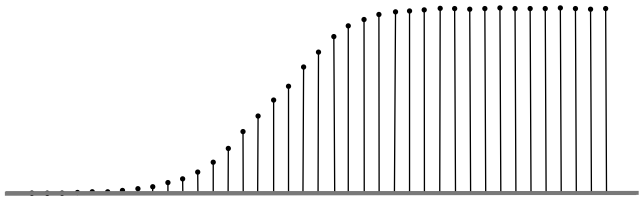

$$x_i = \epsilon\Big(1 - \frac{1}{w}\sum_{j=0}^{w-1}(1 - \frac{1}{w}\sum_{k=0}^{w-1} x_{i+j-k})^{d_r-1}\Big)^{d_l-1} \qquad x_i = 0 \quad \forall\ i \notin [1, L]$$

double the size of w and
scale horizontal axis

Let us quickly explain. The first step of the proof consists of considering an appropriately chosen continuous version of this problem.

For the random coupled ensemble, we introduced the window, w, over which the random connections are made. This discrete system is hard to analyze. Instead, one can consider the limit when w goes to infinity. Of course, in this case the length of the chain has to go to infinity as well. If we increase the length w and scale the horizontal axis by the same proportion then in the limit we can think of the constellation as a continuous curve rather than a set of spikes. To this corresponds a DE equation where instead of taking sums to average we integrate over a finite support window.

Continuum Limit

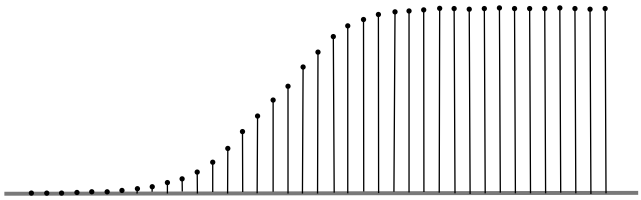

The next step consists in analyzing this continuous system. It is convenient to think of systems of infinite length, i.e., the horizontal axis extends from -∞ to +∞. Further, instead of consider a two-sides constellation as we usually do, we consider a *one-sided* constellation, i.e., we only look at the "left" part of the constellation.

Note that this approach does not only work for coding but can be applied in the general setting of systems whose state is one-dimensional. We will discuss this in more detail in part IV of this tutorial.

Continuum Limit

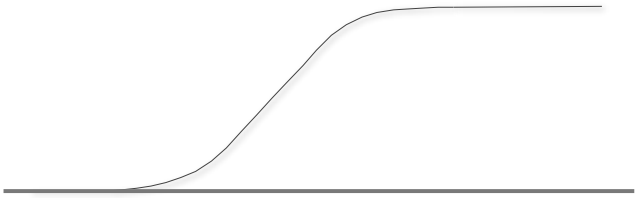

The next step consists in analyzing this continuous system. It is convenient to think of systems of infinite length, i.e., the horizontal axis extends from -∞ to +∞. Further, instead of consider a two-sides constellation as we usually do, we consider a *one-sided* constellation, i.e., we only look at the ``left" part of the constellation.

Note that this approach does not only work for coding but can be applied in the general setting of systems whose state is one-dimensional. We will discuss this in more detail in part IV of this tutorial.

Continuum Limit

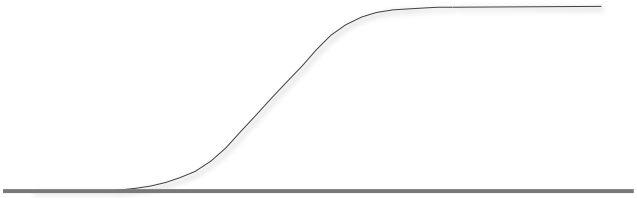

State of the system given by the *continuous* function f(x)

$\omega(x)$ is the smoothing kernel and has finite support

$$f(x) = \epsilon((g \otimes \omega)(x))^{d_l - 1} \qquad g(x) = 1 - (1 - (f \otimes \omega)(x))^{d_r - 1}$$

Saturday, July 13, 13 25

The next step consists in analyzing this continuous system. It is convenient to think of systems of infinite length, i.e., the horizontal axis extends from -∞ to +∞. Further, instead of consider a two-sides constellation as we usually do, we consider a *one-sided* constellation, i.e., we only look at the ``left" part of the constellation.

Note that this approach does not only work for coding but can be applied in the general setting of systems whose state is one-dimensional. We will discuss this in more detail in part IV of this tutorial.

Solutions of Continuum DE: Right-Moving Wave

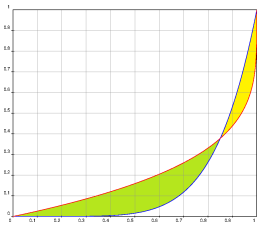

For the continuous system it is then proved that depending on the balance of the areas in the EXIT chart picture of the uncoupled system, one has three different scenarios. (We omit here the trivial case where the curves do not overlap at all
since in this case it is easy to show that we will decode.)

Consider first the scenario where the channel parameter is below the area threshold (but above the BP threshold of the uncoupled ensemble). This means that in the picture on the right the curves do overlap but only little and the area on the left is larger than the area on the right. In this case one can show that there does not exist a FP of DE but there exists a one-sided constellation x so that if we apply DE to this constellation we get x back but shifted ``to the right." Given that our one-sided constellation represents the "left" part of an actual constellation, saying that the wave is moving to the right means that the decoder is working and in each step decodes a further part of the constellation. The shift which we see in each iteration corresponds to the decoding speed and so tells us how many iterations we will need. To summarize, below the area threshold we get a decoding wave which moves to the right and the decoder works.

Solutions of Continuum DE: Right-Moving Wave

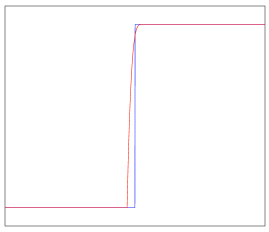

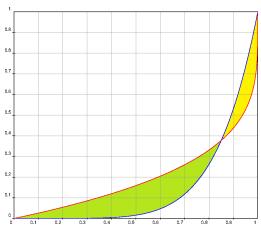

Saturday, July 13, 13 26

For the continuous system it is then proved that depending on the balance of the areas in the EXIT chart picture of the uncoupled system, one has three different scenarios. (We omit here the trivial case where the curves do not overlap at all
since in this case it is easy to show that we will decode.)

Consider first the scenario where the channel parameter is below the area threshold (but above the BP threshold of the uncoupled ensemble). This means that in the picture on the right the curves do overlap but only little and the area on the left is larger than the area on the right. In this case one can show that there does not exist a FP of DE but there exists a one-sided constellation x so that if we apply DE to this constellation we get x back but shifted ``to the right." Given that our one-sided constellation represents the "left" part of an actual constellation, saying that the wave is moving to the right means that the decoder is working and in each step decodes a further part of the constellation. The shift which we see in each iteration corresponds to the decoding speed and so tells us how many iterations we will need. To summarize, below the area threshold we get a decoding wave which moves to the right and the decoder works.

Solutions of Continuum DE: Stationary Wave

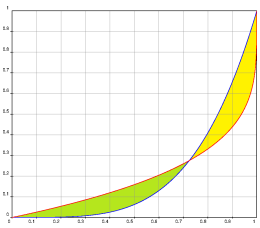

Assume next that the areas are exactly in balance, this means that we are transmitting exactly at the area threshold. In this case one can prove that the continuous version of DE has a FP as shown above. This FP can be thought of as a stationary wave or a wave with zero speed.

Solutions of Continuum DE: Stationary Wave

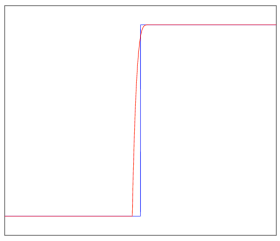

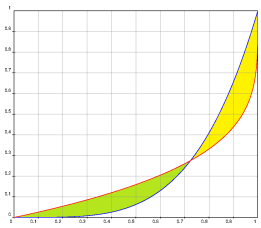

Saturday, July 13, 13 27

Assume next that the areas are exactly in balance, this means that we are transmitting exactly at the area threshold. In this case one can prove that the continuous version of DE has a FP as shown above. This FP can be thought of as a stationary wave or a wave with zero speed.

Solutions of Continuum DE: Left-Moving Wave

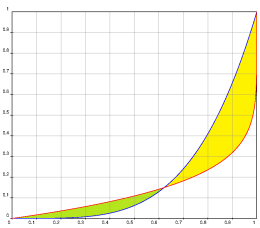

Finally, consider the case where we are transmitting above the area threshold. In the picture on the right this means that the curves overlap so much so that the area on the left is smaller than the area on the right.

For this case one can then show that there does not exist a non-trivial FP but only a continuous constellation x, so that after one round of DE we get the same constellation back but shifted ``to the left." So if we continue to run DE then in effect we get a wave which is moving ``to the left." According to our interpretation this means that the decoder does not work.

Solutions of Continuum DE: Left-Moving Wave

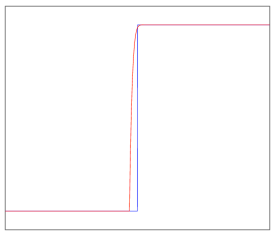

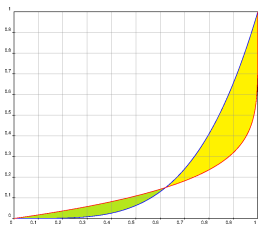

Saturday, July 13, 13 28

Finally, consider the case where we are transmitting above the area threshold. In the picture on the right this means that the curves overlap so much so that the area on the left is smaller than the area on the right.

For this case one can then show that there does not exist a non-trivial FP but only a continuous constellation x, so that after one round of DE we get the same constellation back but shifted ``to the left." So if we continue to run DE then in effect we get a wave which is moving ``to the left." According to our interpretation this means that the decoder does not work.

Proof via EXIT Charts

In a final step one needs to reconnect the continous system to the actual discrete system and show that if the w is not too small then the behavior of the discrete system is well predicted by the behavior of the continous system.

Proof via EXIT Charts

Show that the discrete system behaves similar to the continuous system as long as w is chosen not too small. In other words, the difference between the two decays with w.

In a final step one needs to reconnect the continous system to the actual discrete system and show that if the w is not too small then the behavior of the discrete system is well predicted by the behavior of the continous system.

Information-Theoretically Optimal Compressed Sensing via Spatial Coupling and Approximate Message Passing

A Simple Proof of Threshold Saturation for Coupled Scalar Recursions

A Potential Theory of General Spatially-Coupled Systems via a Continuum Approximation

Proof via Potential Functions

Uncoupled Potential Function

$$
\begin{aligned}
U(x,\epsilon) = & \int_0^x (z - \epsilon\lambda(1-\rho(1-z)))\rho'(1-z)\mathrm{d}z \\
= & \; x(1-\rho(1-x)) - R(1-x) - \epsilon\Lambda(1-\rho(1-x))
\end{aligned}
$$

$U'(x,\epsilon) \geq 0$

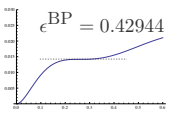

$U(x,\epsilon) \geq 0$

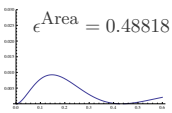

Recall the definition of the potential function for an uncoupled system. Recall that the BP threshold for the uncoupled system was defined as the supremum of all channel parameters so that the potential function is strictly increasing, i.e., the derivative of the potential function must be strictly positive, whereas the area threshold is defined as the supremum of all channel parameters so that the potential function itself is strictly positive.

Rewriting of DE equations

$$x_i^{(\ell+1)} = \frac{1}{w}\sum_{k=0}^{w-1} f\left(\frac{1}{w}\sum_{j=0}^{w-1} g(x_{i+j-k}^{(\ell)}); \varepsilon_{i-k}\right)$$

$$\boldsymbol{x}^{(\ell+1)} = \boldsymbol{A}_2^{\mathsf{T}} \boldsymbol{f}(\boldsymbol{A}_2 \boldsymbol{g}(\boldsymbol{x}^{(\ell)}); \varepsilon)$$

$$\boldsymbol{A}_2 = \frac{1}{w}\underbrace{\begin{bmatrix} \overbrace{1 & 1 & \cdots & 1}^{w} & 0 & 0 & \cdots & 0 \\ 0 & 1 & 1 & \cdots & 1 & 0 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & 1 & 1 & \cdots & 1 & 0 \\ 0 & \cdots & 0 & 0 & 1 & 1 & \cdots & 1 \end{bmatrix}}_{2L+w} \Big\} 2L+1$$

First, following the lead of Yedla, Jian, Nguyen, and Pfister, we rewrite the DE equations in vector form. Note that in above formulas we consider a *one-sided* system of size 2L+1 and the matrix A encodes the connections so that a multiplication with the matrix A respectively its transpose A^T we perform in effect the averaging operation at the check and variable nodes. Note that there is also a subtle difference of the above DE equations compared to the equations we have written so far. In the above equations the x quantities represent the *averages* of the quantities which we called x previously. Clearly, both of these quantities form a state of the system. Why do we switch notation here? It turns out that taking the averages is considerably more convenient when we introduce the potential of the coupled system.

Coupled Potential Function

$$\begin{aligned} U(x,\epsilon) = & \int_0^x (z - \epsilon\lambda(1-\rho(1-z)))\rho'(1-z)\mathrm{d}z \\ = & \; x(1-\rho(1-x)) - R(1-x) - \epsilon\Lambda(1-\rho(1-x)) \end{aligned}$$

$$\begin{aligned} U(\boldsymbol{x};\varepsilon) &= \int_{\mathcal{C}} \boldsymbol{g}'(\boldsymbol{z})(\boldsymbol{z} - \boldsymbol{A}^{\mathsf{T}} \boldsymbol{f}(\boldsymbol{A}\boldsymbol{g}(\boldsymbol{z});\varepsilon)) \cdot \mathrm{d}\boldsymbol{z} \\ &= \boldsymbol{g}(\boldsymbol{x})^{\mathsf{T}}\boldsymbol{x} - G(\boldsymbol{x}) - F(\boldsymbol{A}\boldsymbol{g}(\boldsymbol{x});\varepsilon), \end{aligned}$$

$$\boldsymbol{g}'(\boldsymbol{x}) = \mathrm{diag}([g'(x_i)]) \quad G(\boldsymbol{x}) = \int_{\mathcal{C}} \boldsymbol{g}(\boldsymbol{z}) \cdot \mathrm{d}\boldsymbol{z} \quad F(\boldsymbol{x};\varepsilon) = \int_{\mathcal{C}} \boldsymbol{f}(\boldsymbol{z};\varepsilon) \cdot \mathrm{d}\boldsymbol{z}$$

The potential function of the coupled system can now be introduced in analogy to the uncoupled system. Note that formally the two look very much the same, except that now we have to deal with vector valued functions.

Main Claim

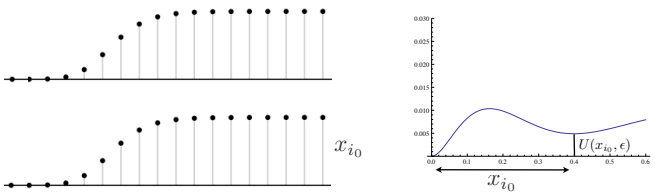

$$U(\boldsymbol{S}\boldsymbol{x};\varepsilon) - U(\boldsymbol{x};\varepsilon) = -U(x_{i_0};\varepsilon)$$

The main claim is the following. Assume we are given a one-sided increasing constellation and we shift the constellation by one position to the right. Then the difference in the potential functions of the shifted and unshifted constellation is equal to minus the uncoupled potential function of the rightmost position. So if we are transmitting above the area threshold this value will be strictly negative. Now think of this difference as a derivative. It is then clear that the wave will move to the right since in this direction the energy is lowered!

To make this precise one uses a Taylor series expansion of U(Sx; eps) around U(x; eps) up to second order and
shows that it is dominated (for w not too small) by the constant and linear term.

Main Claim

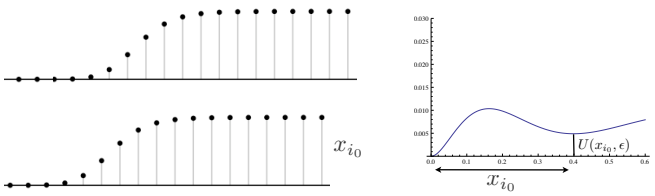

$$U(\boldsymbol{S}\boldsymbol{x};\varepsilon) - U(\boldsymbol{x};\varepsilon) = -U(x_{i_0};\varepsilon)$$

The main claim is the following. Assume we are given a one-sided increasing constellation and we shift the constellation by one position to the right. Then the difference in the potential functions of the shifted and unshifted constellation is equal to minus the uncoupled potential function of the rightmost position. So if we are transmitting above the area threshold this value will be strictly negative. Now think of this difference as a derivative. It is then clear that the wave will move to the right since in this direction the energy is lowered!

To make this precise one uses a Taylor series expansion of U(Sx; eps) around U(x; eps) up to second order and
shows that it is dominated (for w not too small) by the constant and linear term.

Three sides of the same coin :-)

At the area threshold:

Special FP exists, with long tail, quick transition and large flat part

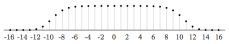

Operational Interpretation:

Special FP is unstable under BP and collapses to zero below the area threshold

At the area threshold:

Stationary wave exists

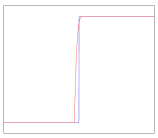

Operational Interpretation:

Below the area threshold, we get a traveling wave moving to the right implying successful BP decoding

At the area threshold:

Potential function has zero gradient

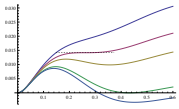

Operational Interpretation:

Below the area threshold, potential energy is strictly decreasing implying convergence to perfect decoding

To summarize: the three different proof techniques have essentially the same foundations. Something happens in the coupled DE system at the area threshold of the uncoupled system: in the first proof technique, we show the existence of a special FP; in the second proof technique we show the existence of a stationary wave and the third proof technique it is shown that the potential function of the coupled system has zero gradient. To complete the proof in each technique an operational interpretation is then shown: in the first proof technique, the special FP is unstable and at any channel value below the area threshold, the application of the BP decoder collapses the special FP to zero; in the second proof technique it was shown that the wave travels to the right for channel values below the area threshold resulting in successful BP decoding and lastly in the third proof technique it is shown that below the area threshold, the potential energy is always strictly decreasing.