

# Recursive Lattice Search: Hierarchical Heavy Hitters Revisited

Kenjiro Cho

IIJ Research Laboratory

IMC'17 November 2, 2017

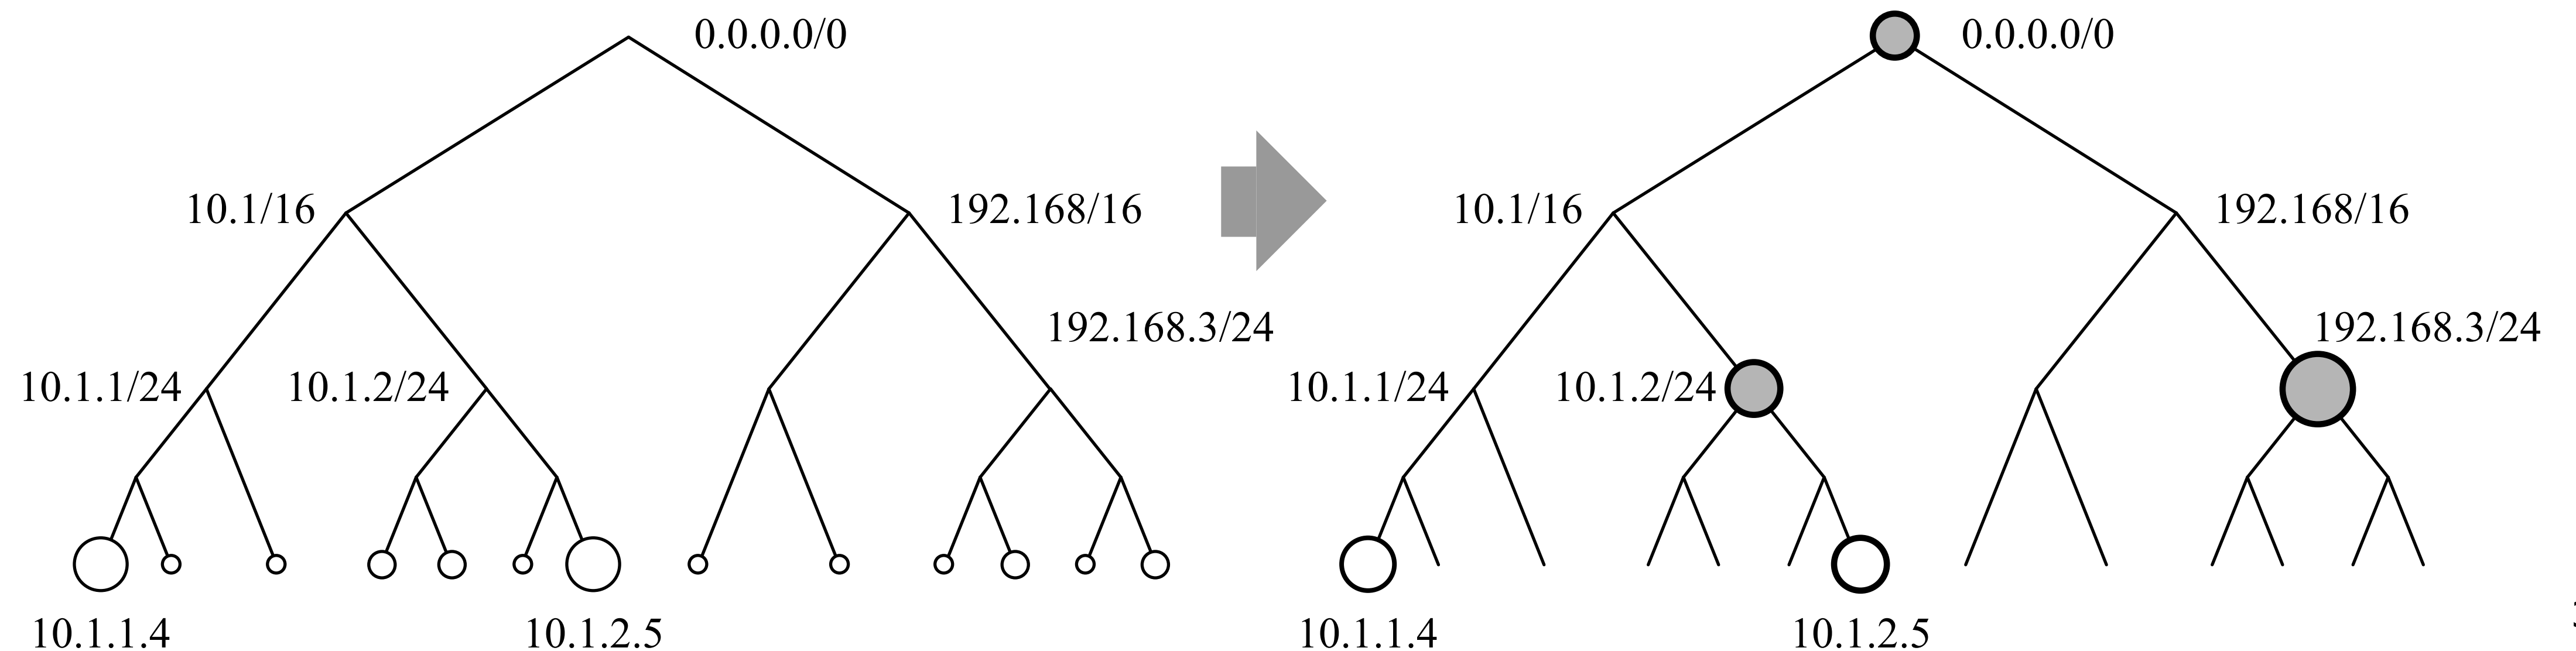


# Hierarchical Heavy Hitters (HHHs)

- identifying significant clusters across multiple planes
  - exploiting underlying hierarchical IP address structures
  - e.g., (src, dst) address pairs
    - $(1.2.3.4, *) \rightarrow$  one-to-many: e.g., scanning
    - $(*, 5.6.7.8) \rightarrow$  many-to-one: e.g., DDoS
    - $(1.2.3.0/24, 4.5.6.0/28) \rightarrow$  subnet-to-subnet
  - can be extended to higher dimensions (e.g., 5-tuple)
- powerful tool for traffic monitoring/anomaly detection

# Unidimensional HHH

- an HHH: an aggregate with count  $c \geq \varphi N$ 
  - $\varphi$ : threshold  $N$ : total input (e.g., packets or bytes)
- HHHs can be uniquely identified by depth-first tree traversal
  - aggregating small nodes until it exceeds the threshold



# Multi-dimensional HHH

- each node has multiple parents
  - many combinations for aggregation
  - much harder than one-dimension
- search space for 2-dimensional IPv4 addrs
  - $5 \times 5 = 25$  for bitwise aggregation
  - $33 \times 33 = 1089$  for bitwise aggregation

src: 1.2.3.4

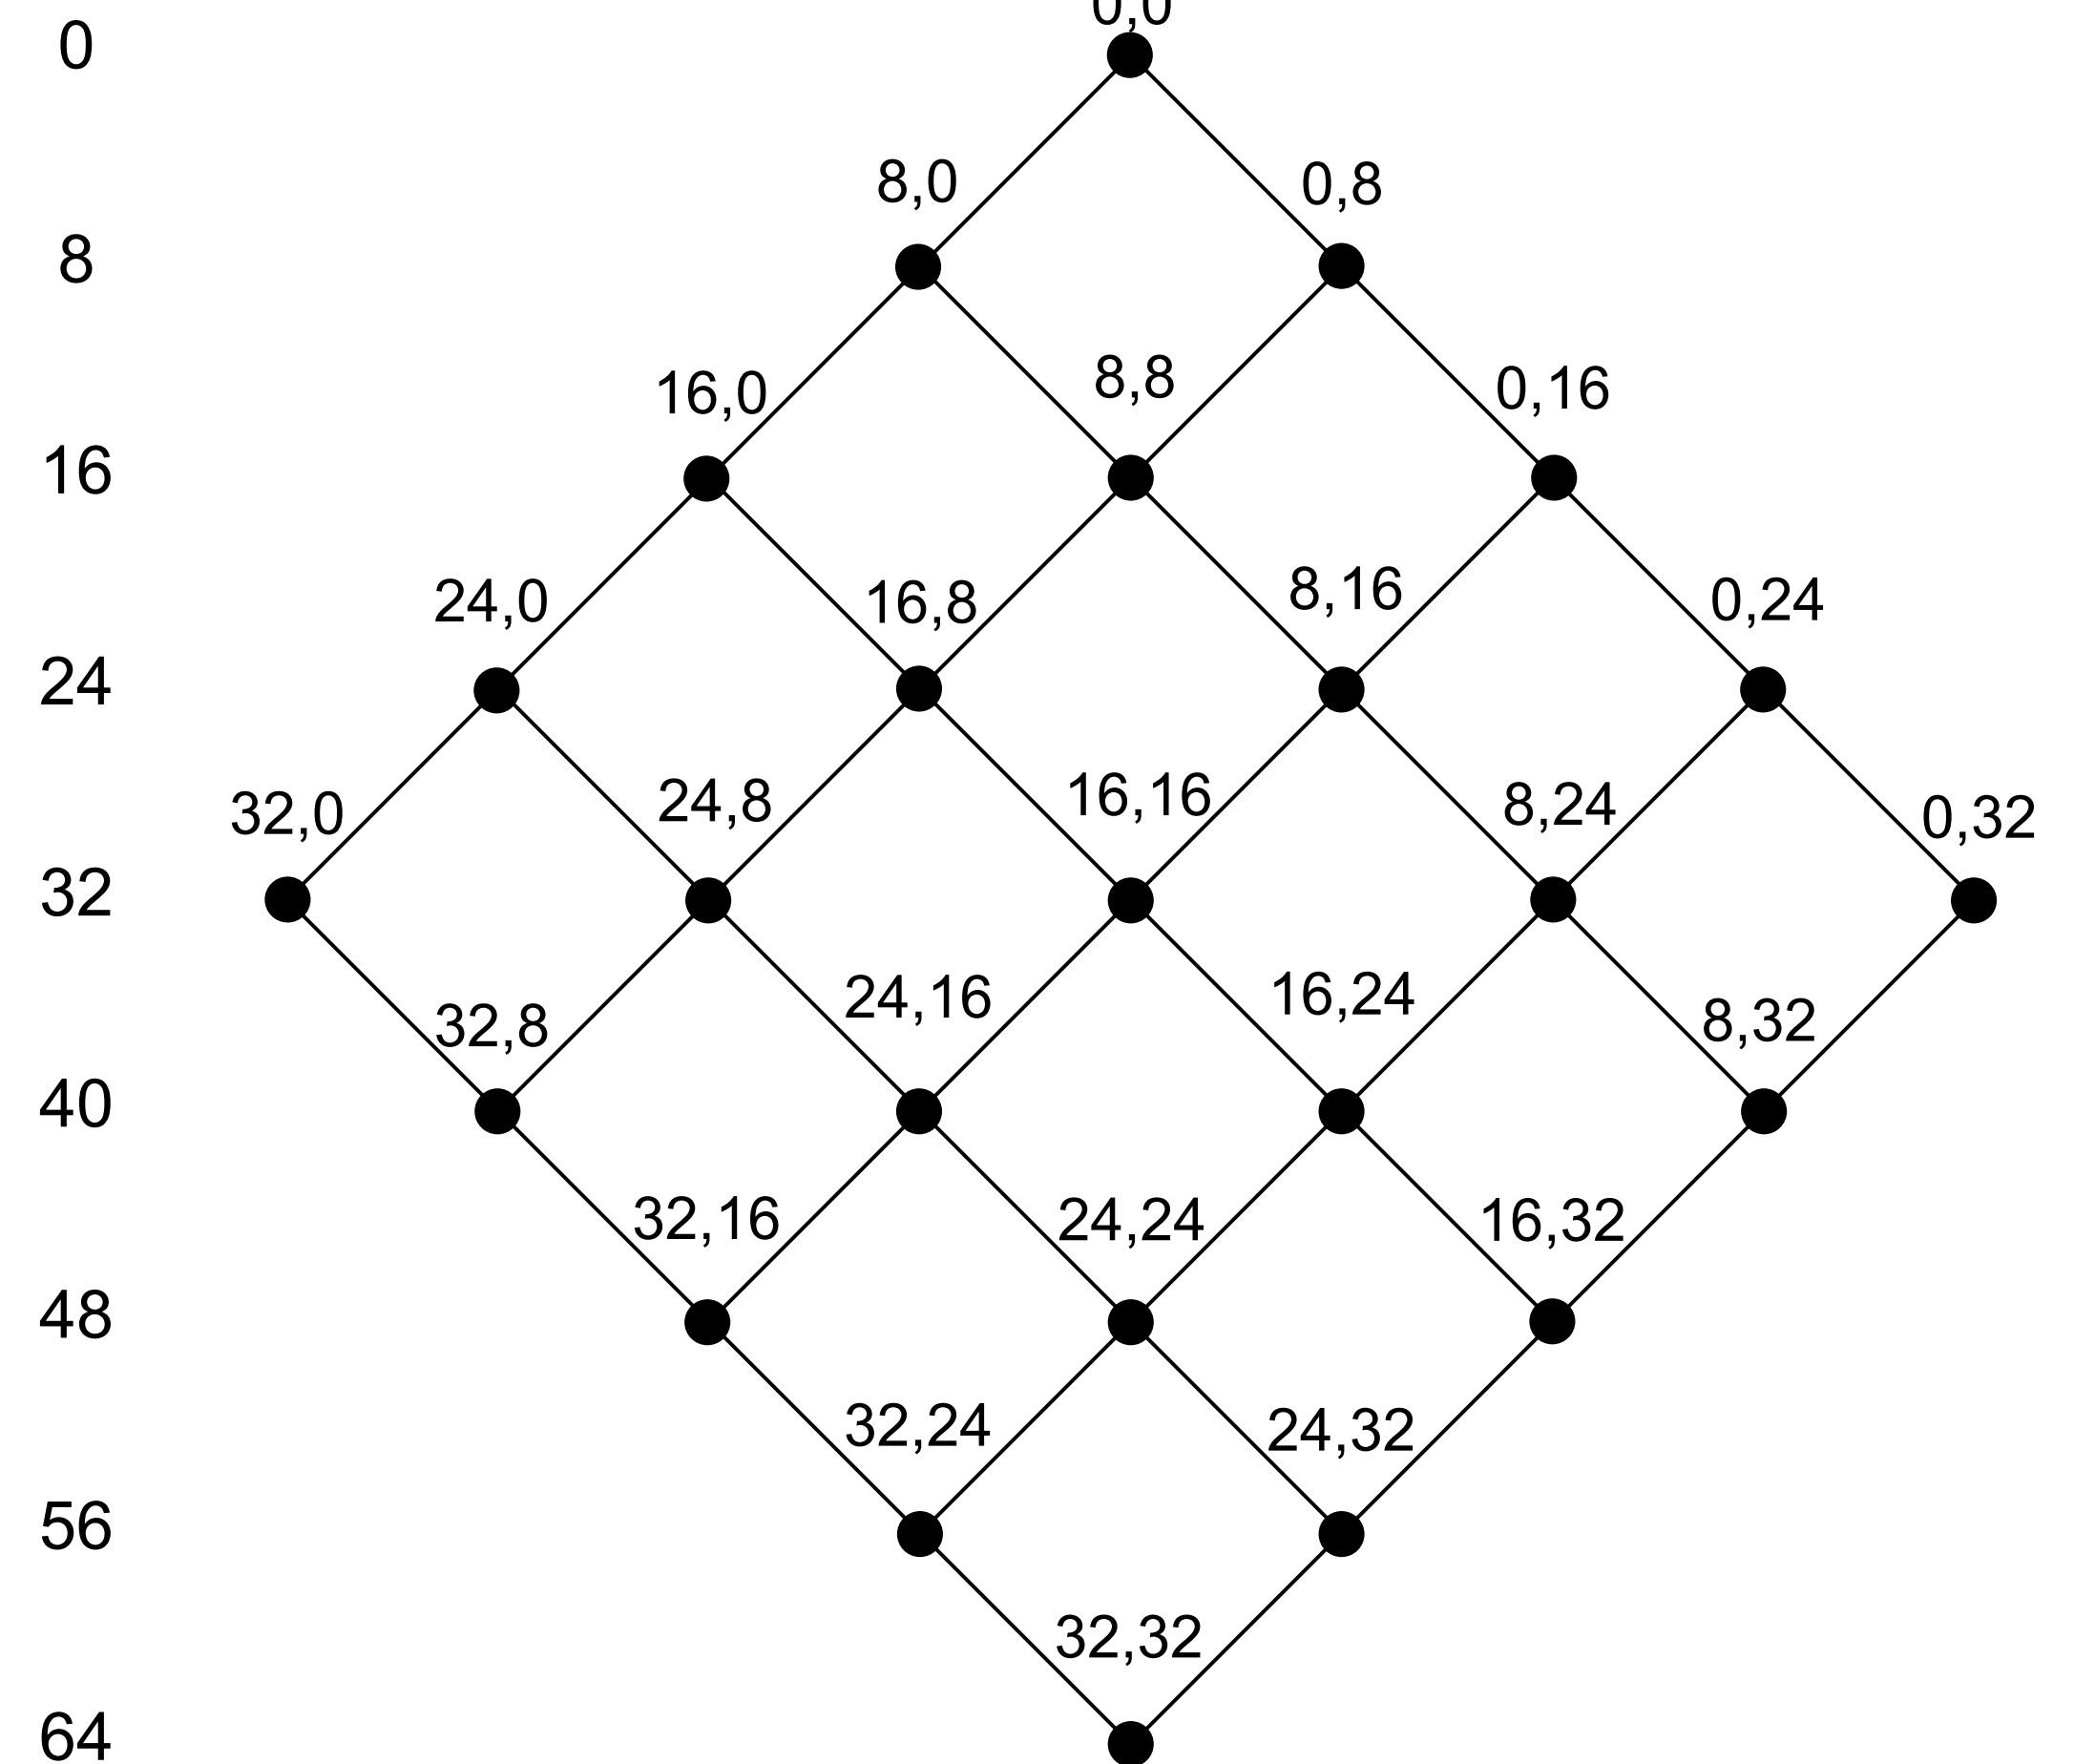
dst: 5.6.7.8

[1.2.3.4/32,5.6.0.0/16] [1.2.3.0/24,5.6.7.0/24] [1.2.0.0/16,1.2.3.4/32]

[1.2.3.4/32,5.6.7.0/24] [1.2.3.0/24,5.6.7.8/32]

[1.2.3.4/32,5.6.7.8/32]

sum of prefix lengths



Lattice for IPv4 prefix length pair with 8-bit granularity

# Challenges

- performance
  - bitwise aggregation is costly
- operational relevance
  - ordering: e.g., [32, \*] and [16, 16]
  - broad and redundant aggregates: (e.g., 128/4 and 128/2)
- re-aggregation
  - useful for interactive analysis (for zoom-in/out)

# Contributions

- new efficient HHH algorithm for bitwise aggregation
  - matches operational needs, supports re-aggregation
- open-source tool and open datasets
- more broadly, transforming the existing hard problem into a tractable one, by revisiting the commonly accepted definition

# Various HHH definitions

- discounted HHH ← we also employ this
  - exclude descendant HHHs' counts for concise outputs

$$c_i' = \sum_j c_j' \text{ where } \{j \in \text{child}(i) \mid c_j' < \varphi N\}$$

- rollup rules: how to aggregate counts to parent
  - overlap rule: allows double-counting to detect all possible HHHs
  - split rule: preserves counts ← we use a simple first-found split rule
- aggregation ordering
  - sum of prefix lengths ← we'll revisit this ordering

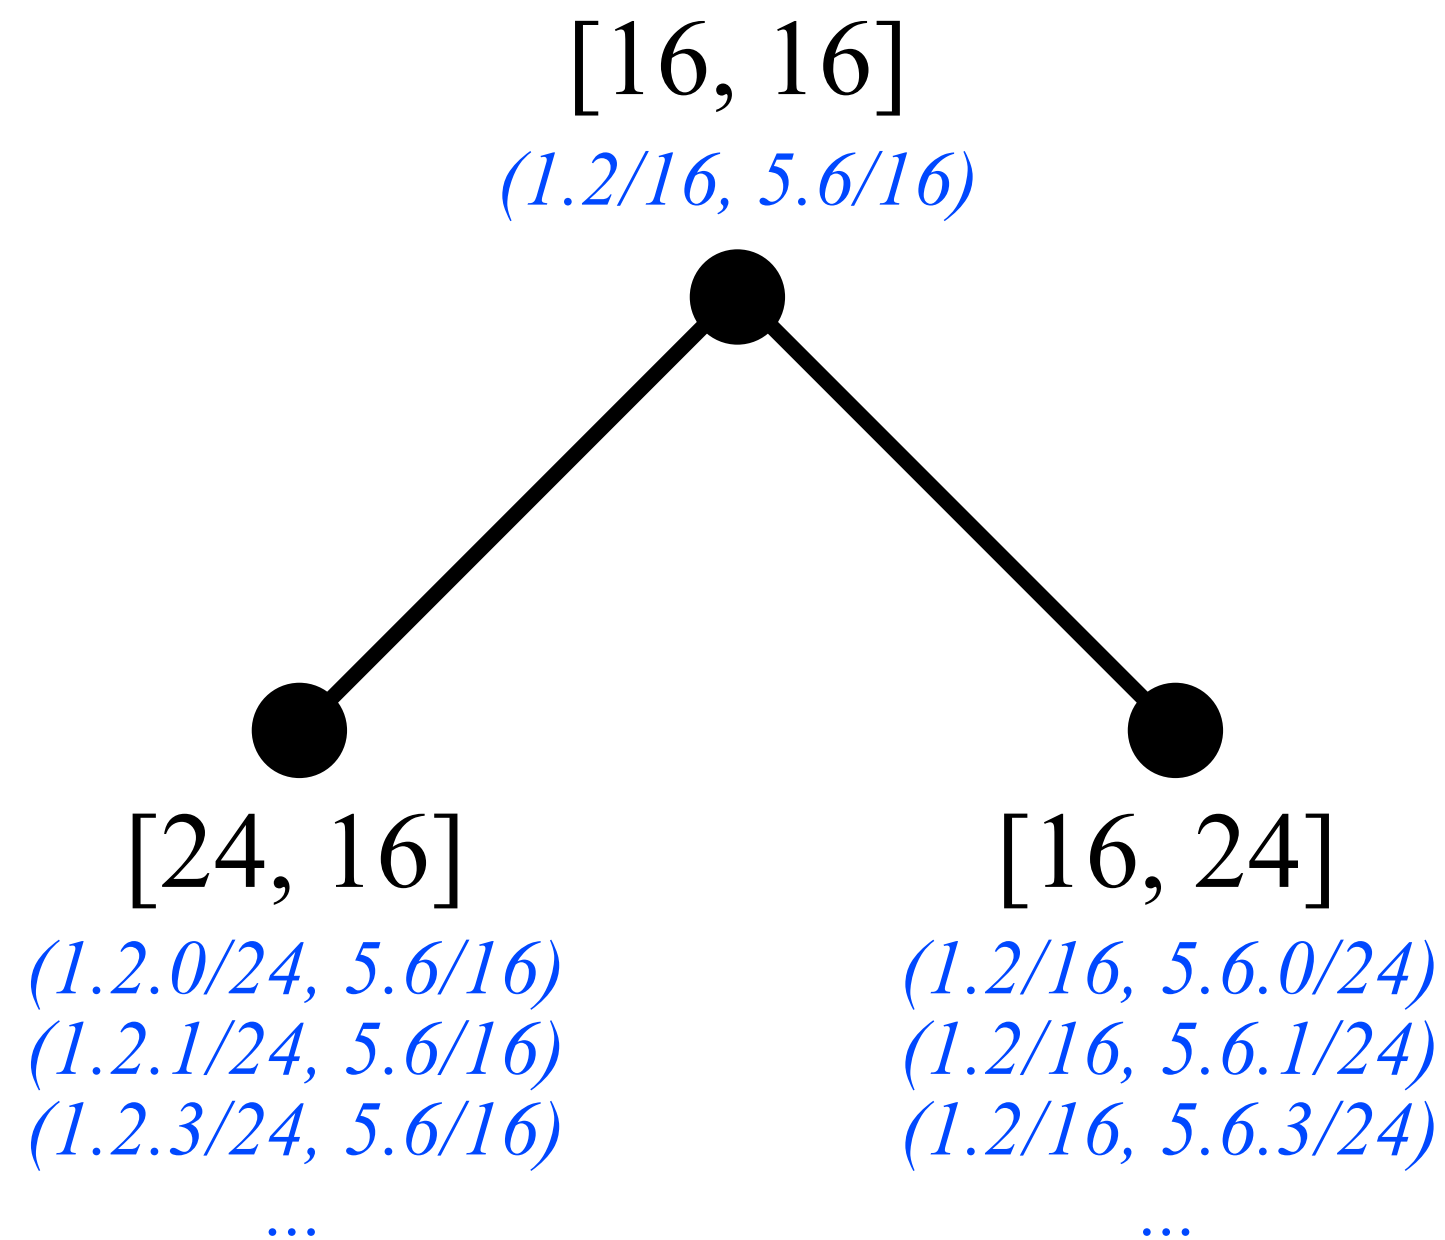
# Previous algorithms

- elaborate structures
  - cross-producting, grid-of-trie, rectangle-search
- theoretical analyses
  - streaming approximation algorithms w/ error bounds
- all the existing methods are bottom-up
  
- our algorithm: top-down, deterministic
  - no elaborate structure, no approximation, no parameter

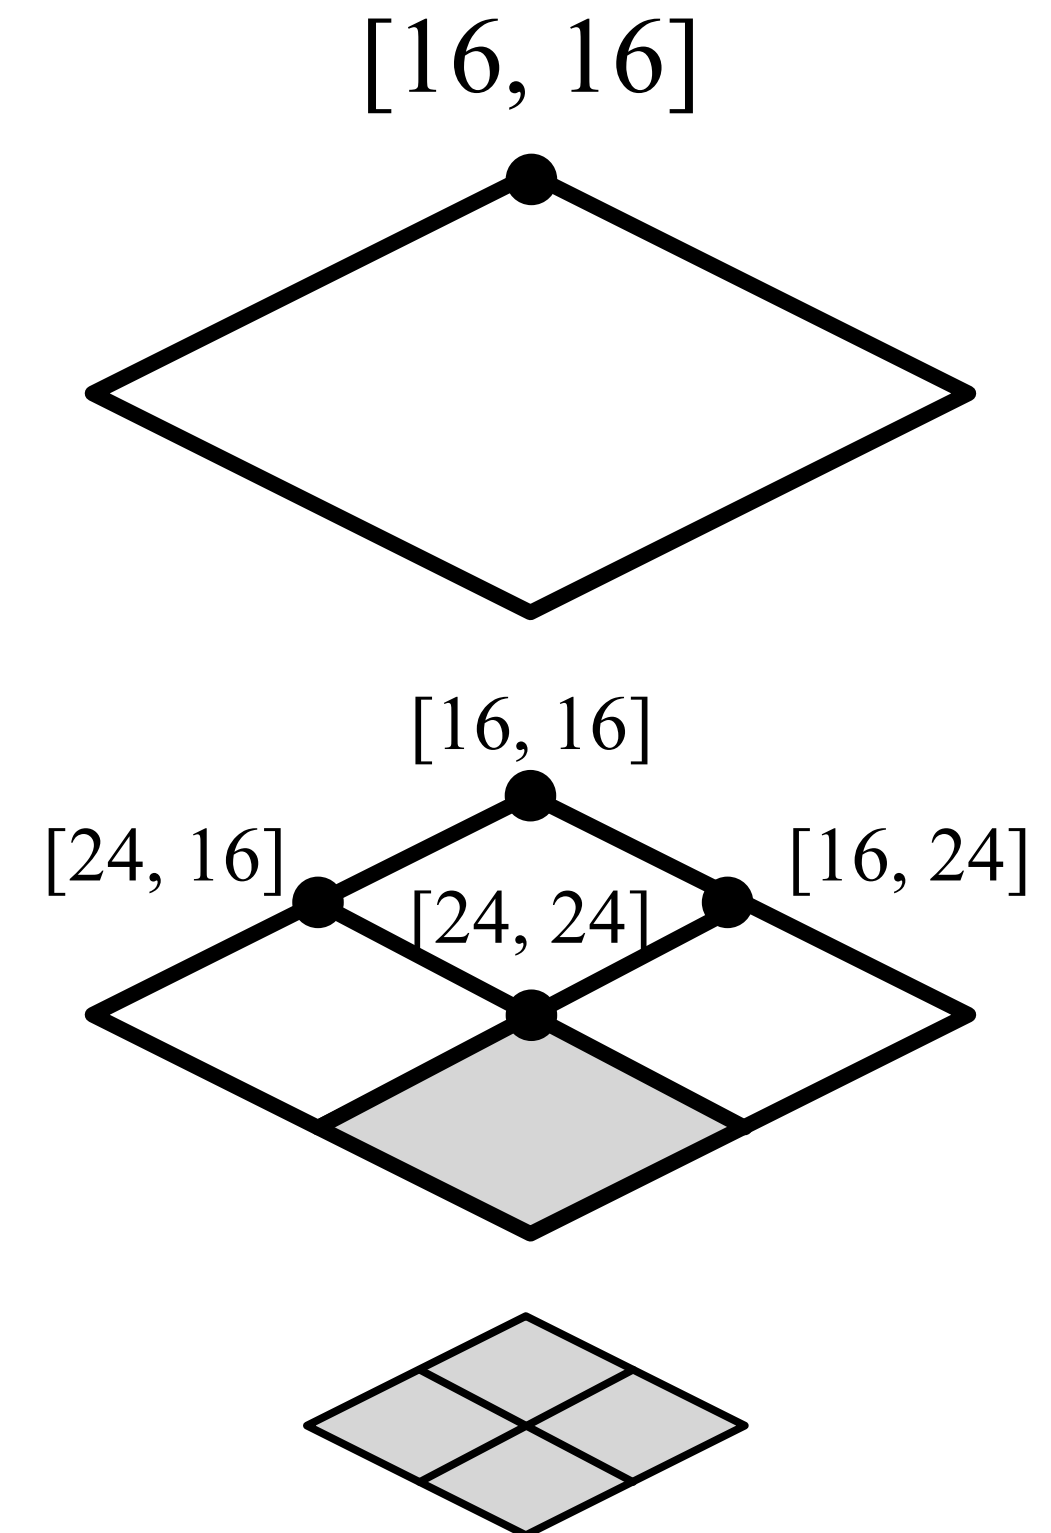


# HHH revisited

- key idea: redefine  $child(i)$  to allow space partitioning
  - $child(i)$ : from bin-tree to quadtree



**bottom-up aggregation**

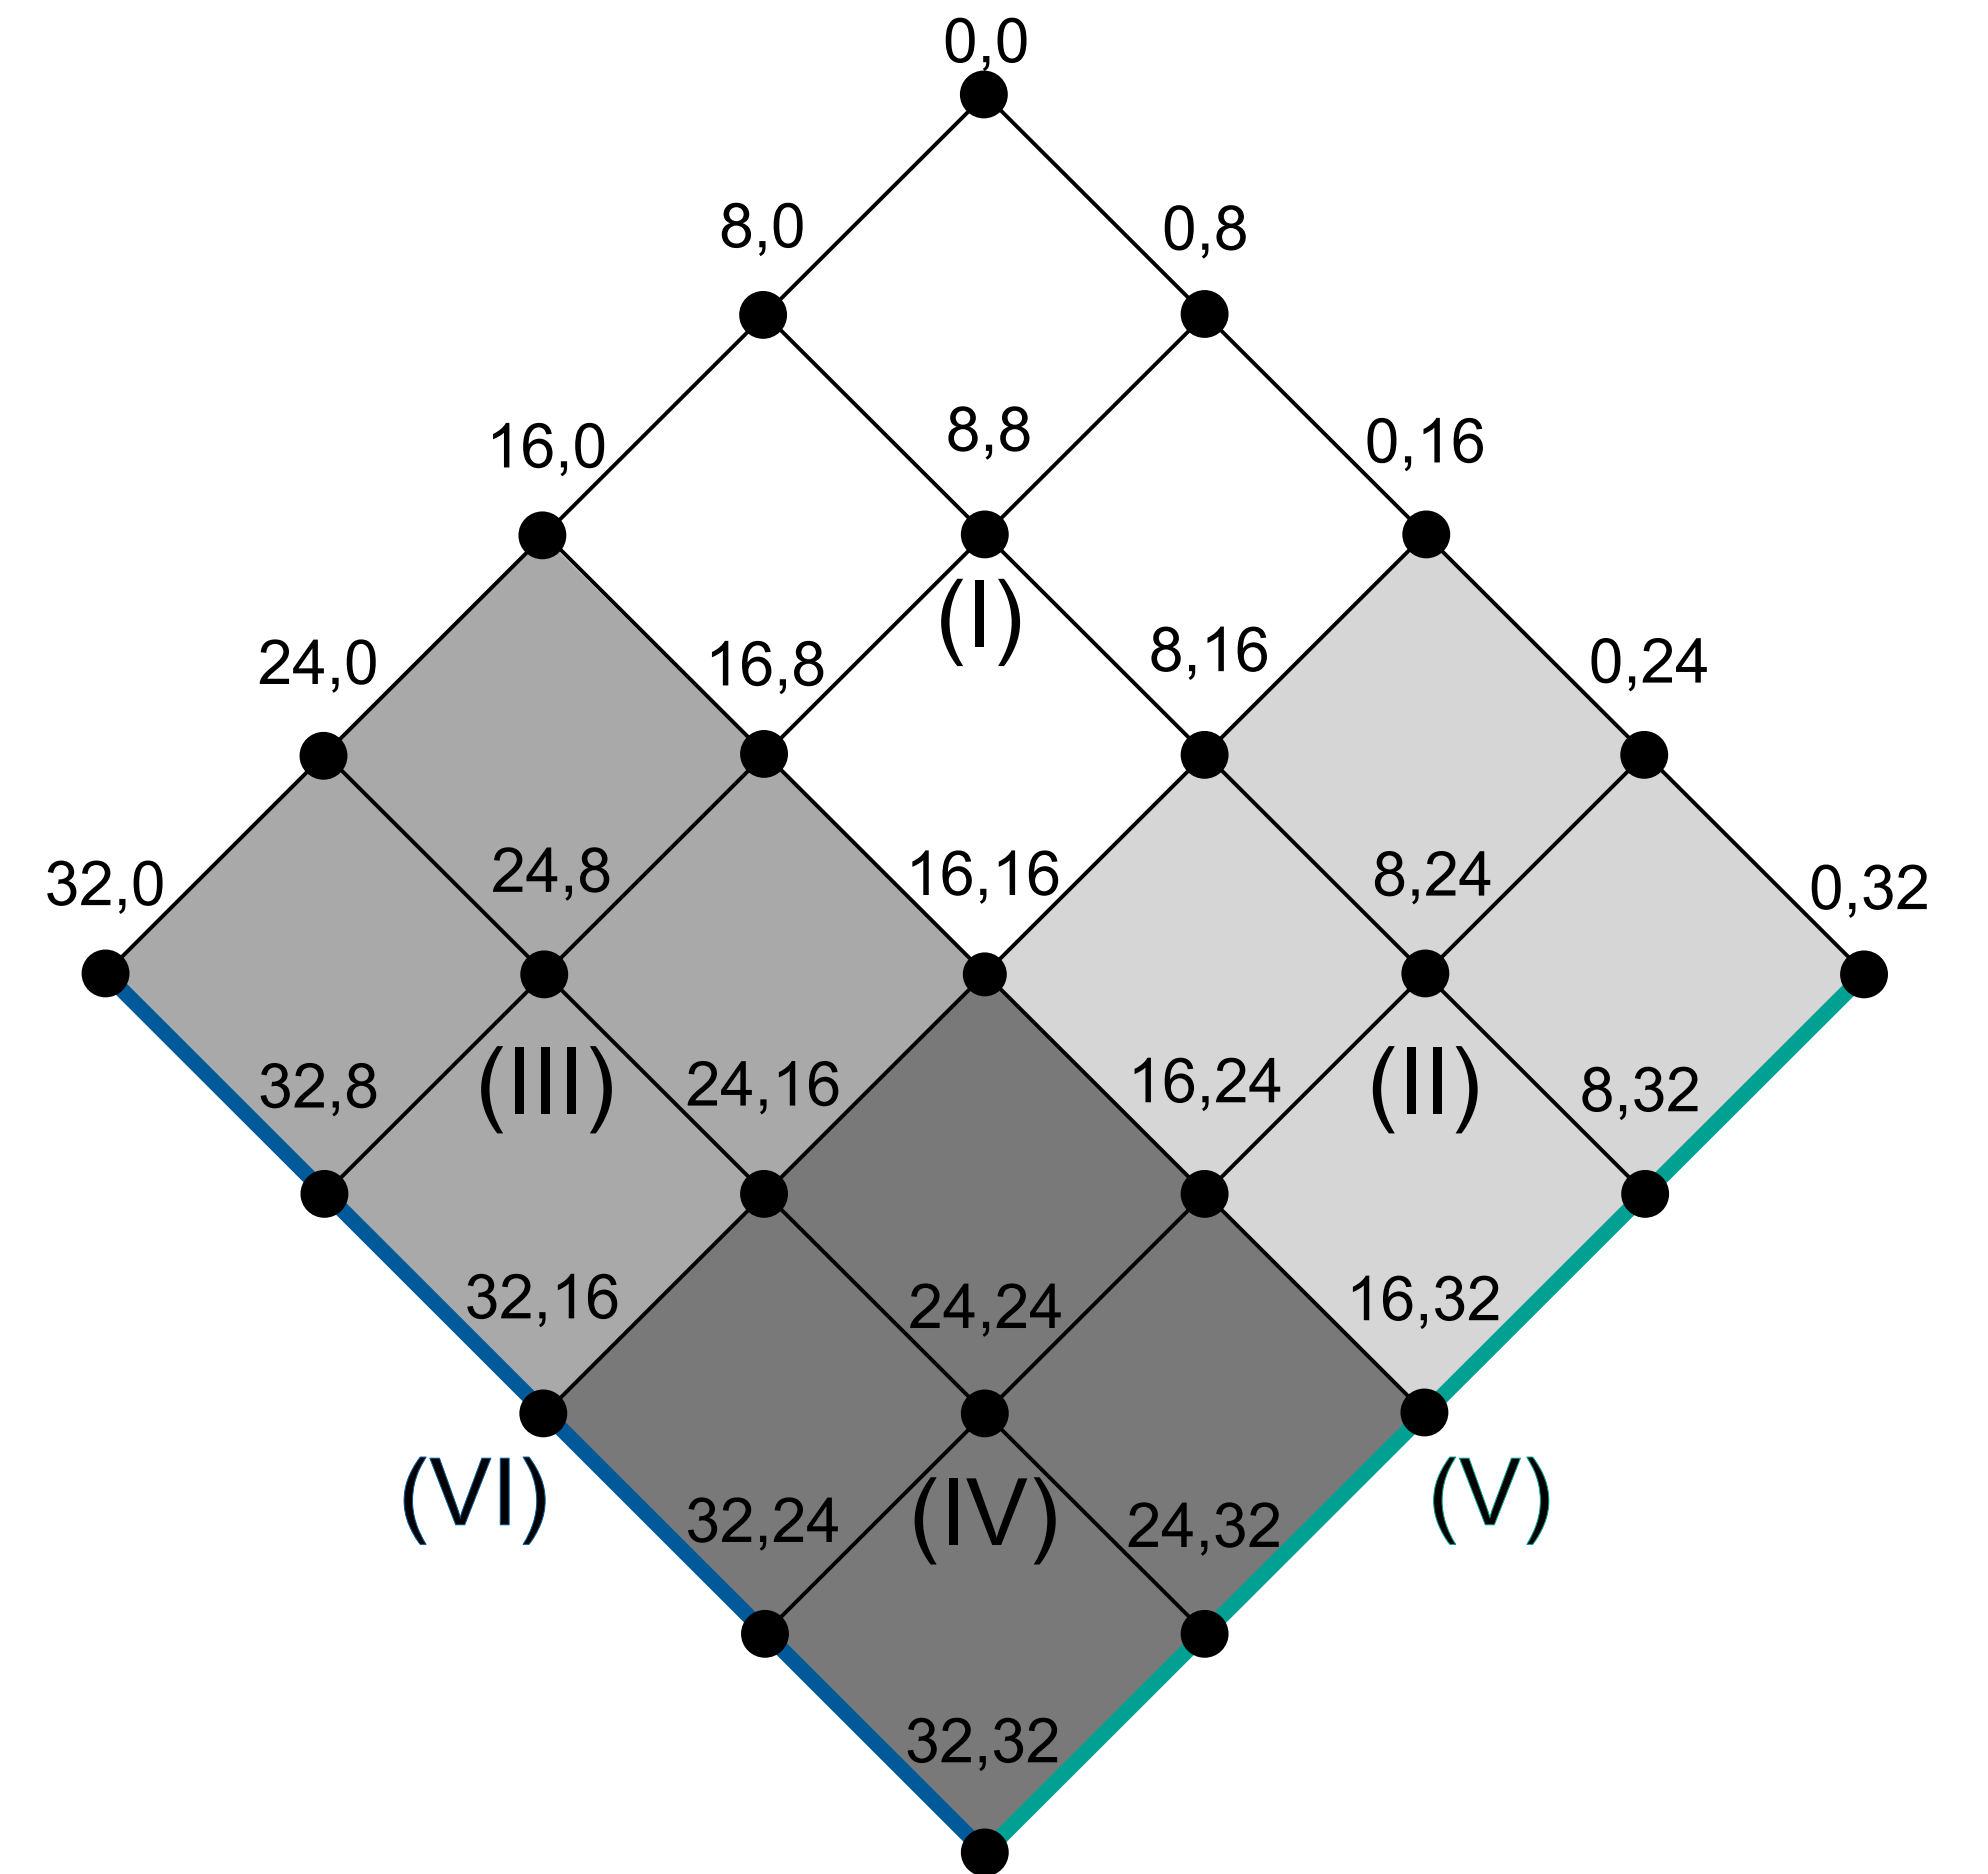


**top-down space partitioning**



# Recursive spatial partitioning

- visit regions from (VI) to (I) recursively
  - 2 bottom edges
    - (VI) left-bottom edge
    - (V) right-bottom edge
  - 4 quadrants
    - (IV) lower quadrant
    - (III) left quadrant
    - (II) right quadrant
    - (I) upper quadrant

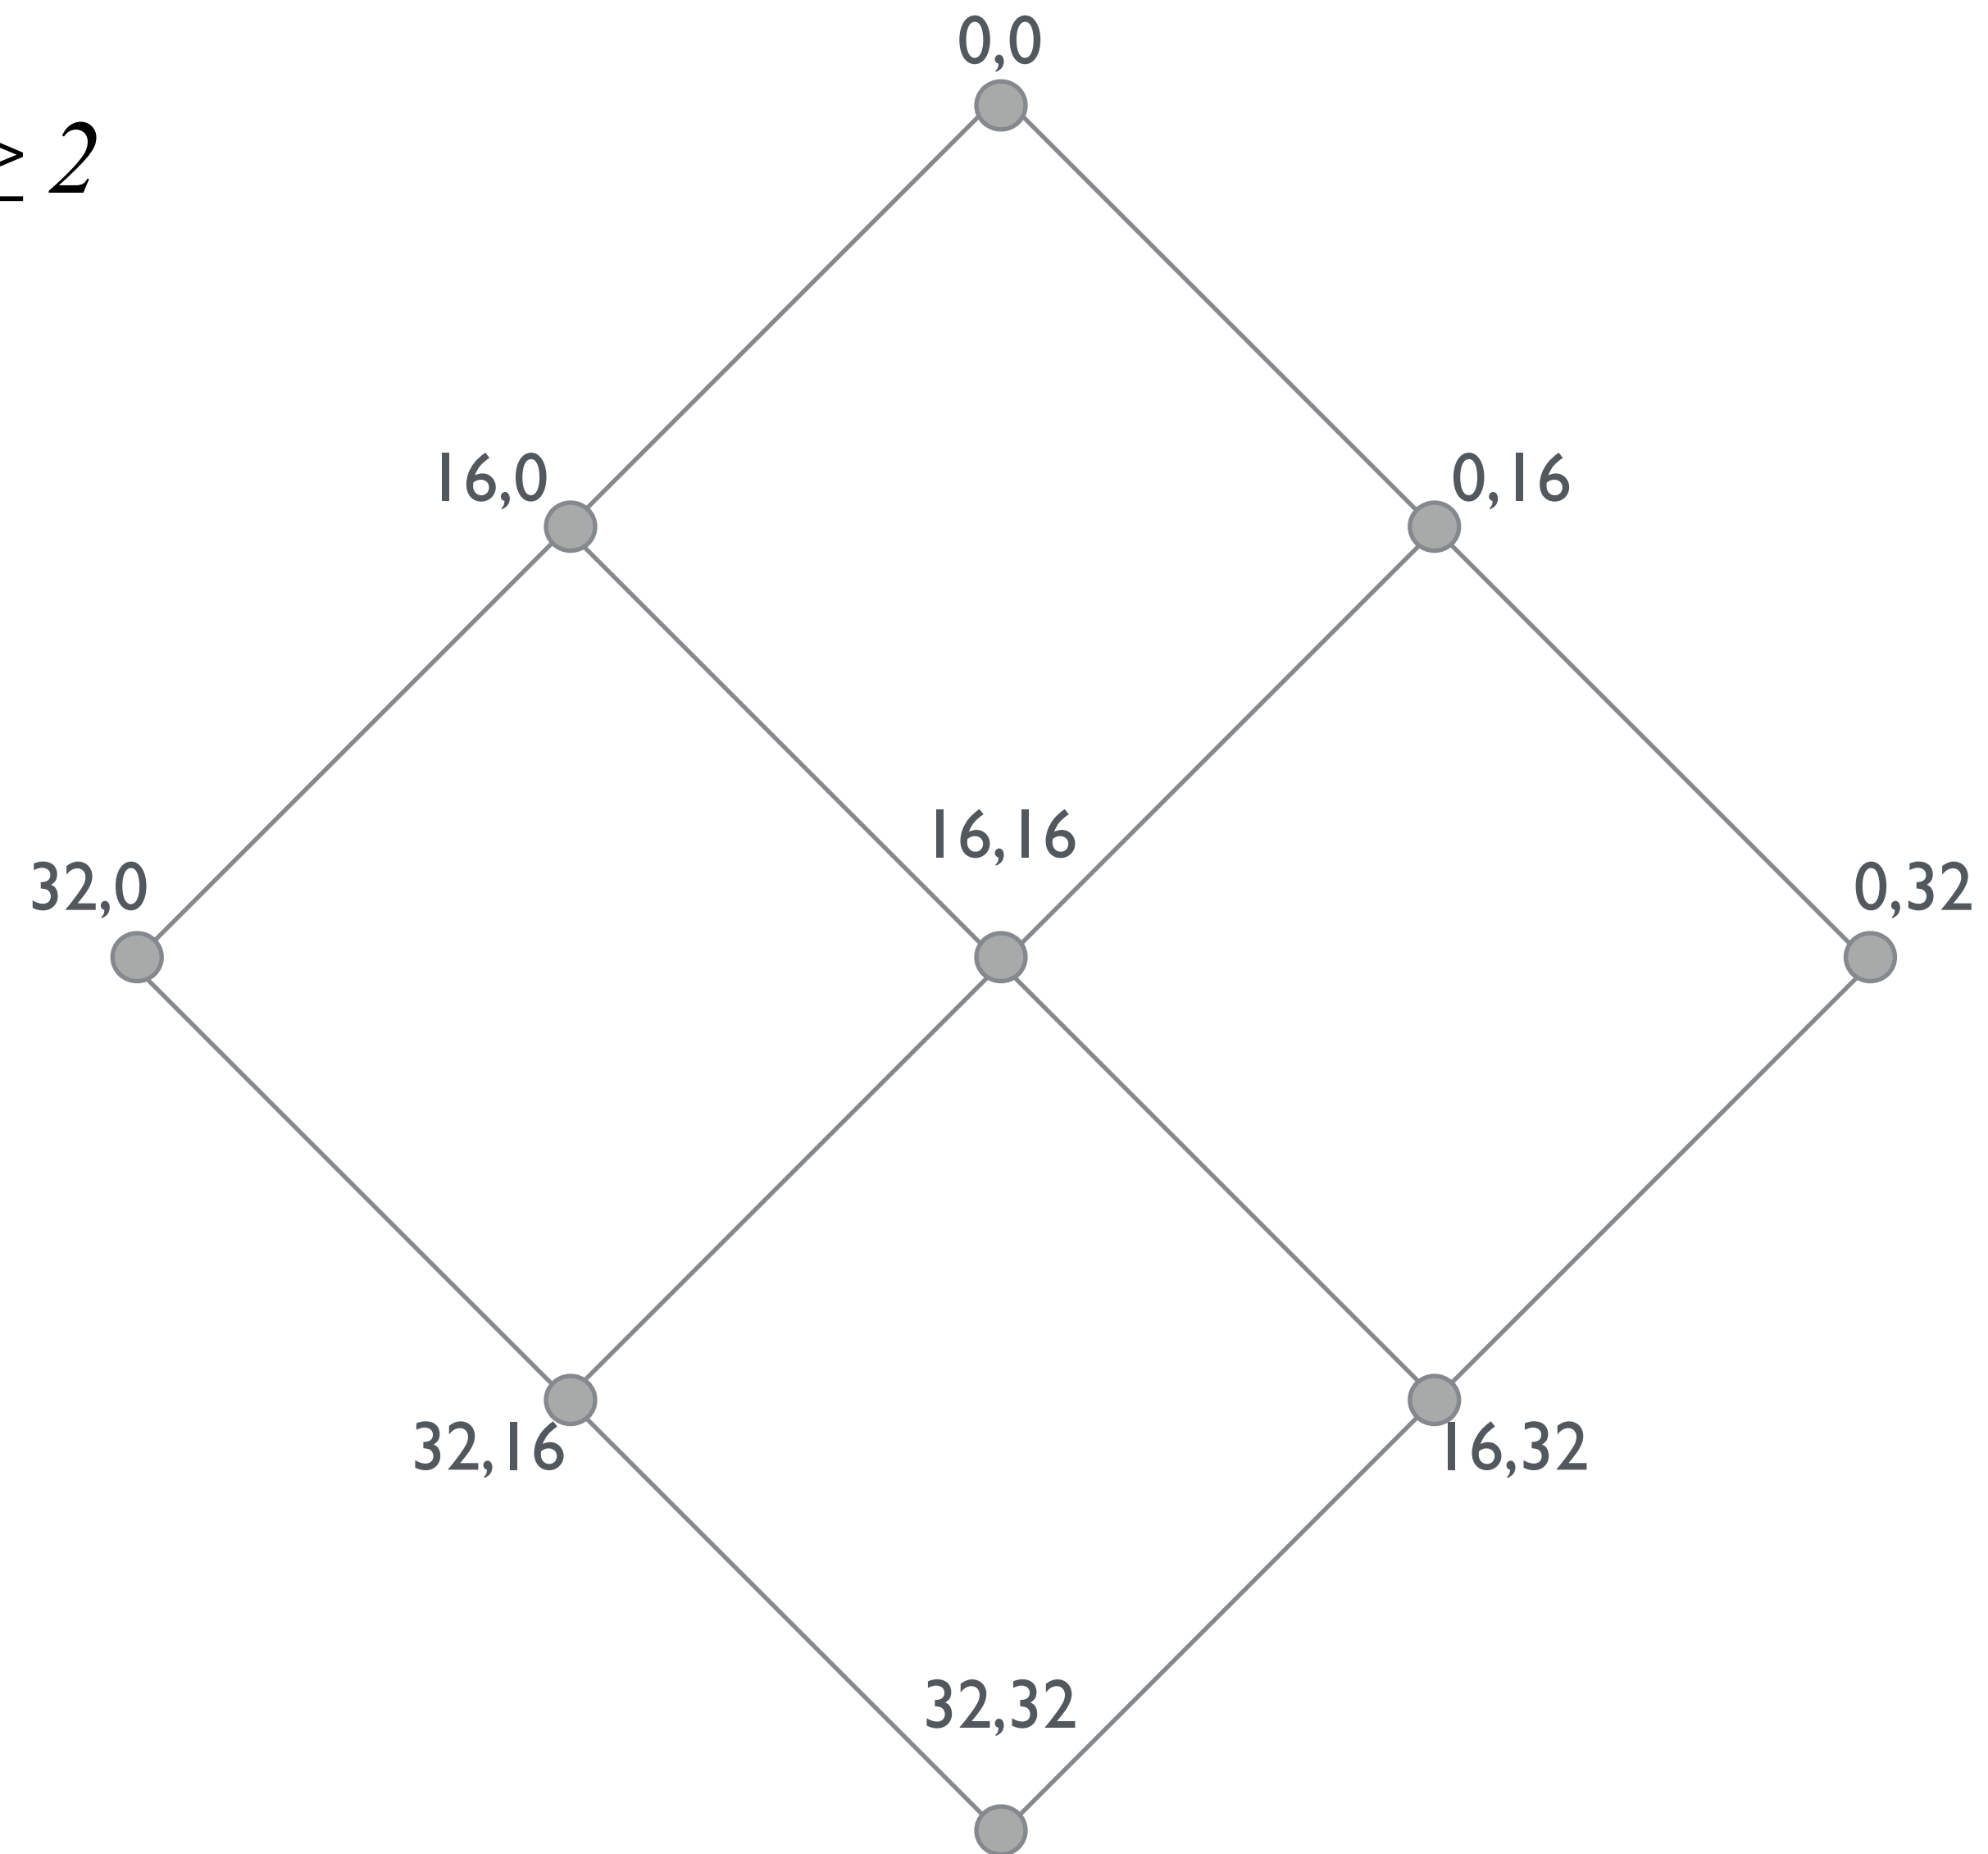


# Recursive Lattice Search (RLS)

- idea: recursively subdivide aggregates by Z-order
- pros
  - recurse only for flows  $\geq$  thresh
  - sub-division needs only parent's sub-flows
  - /32 becomes higher in the order
- cons
  - bias for the first dimension

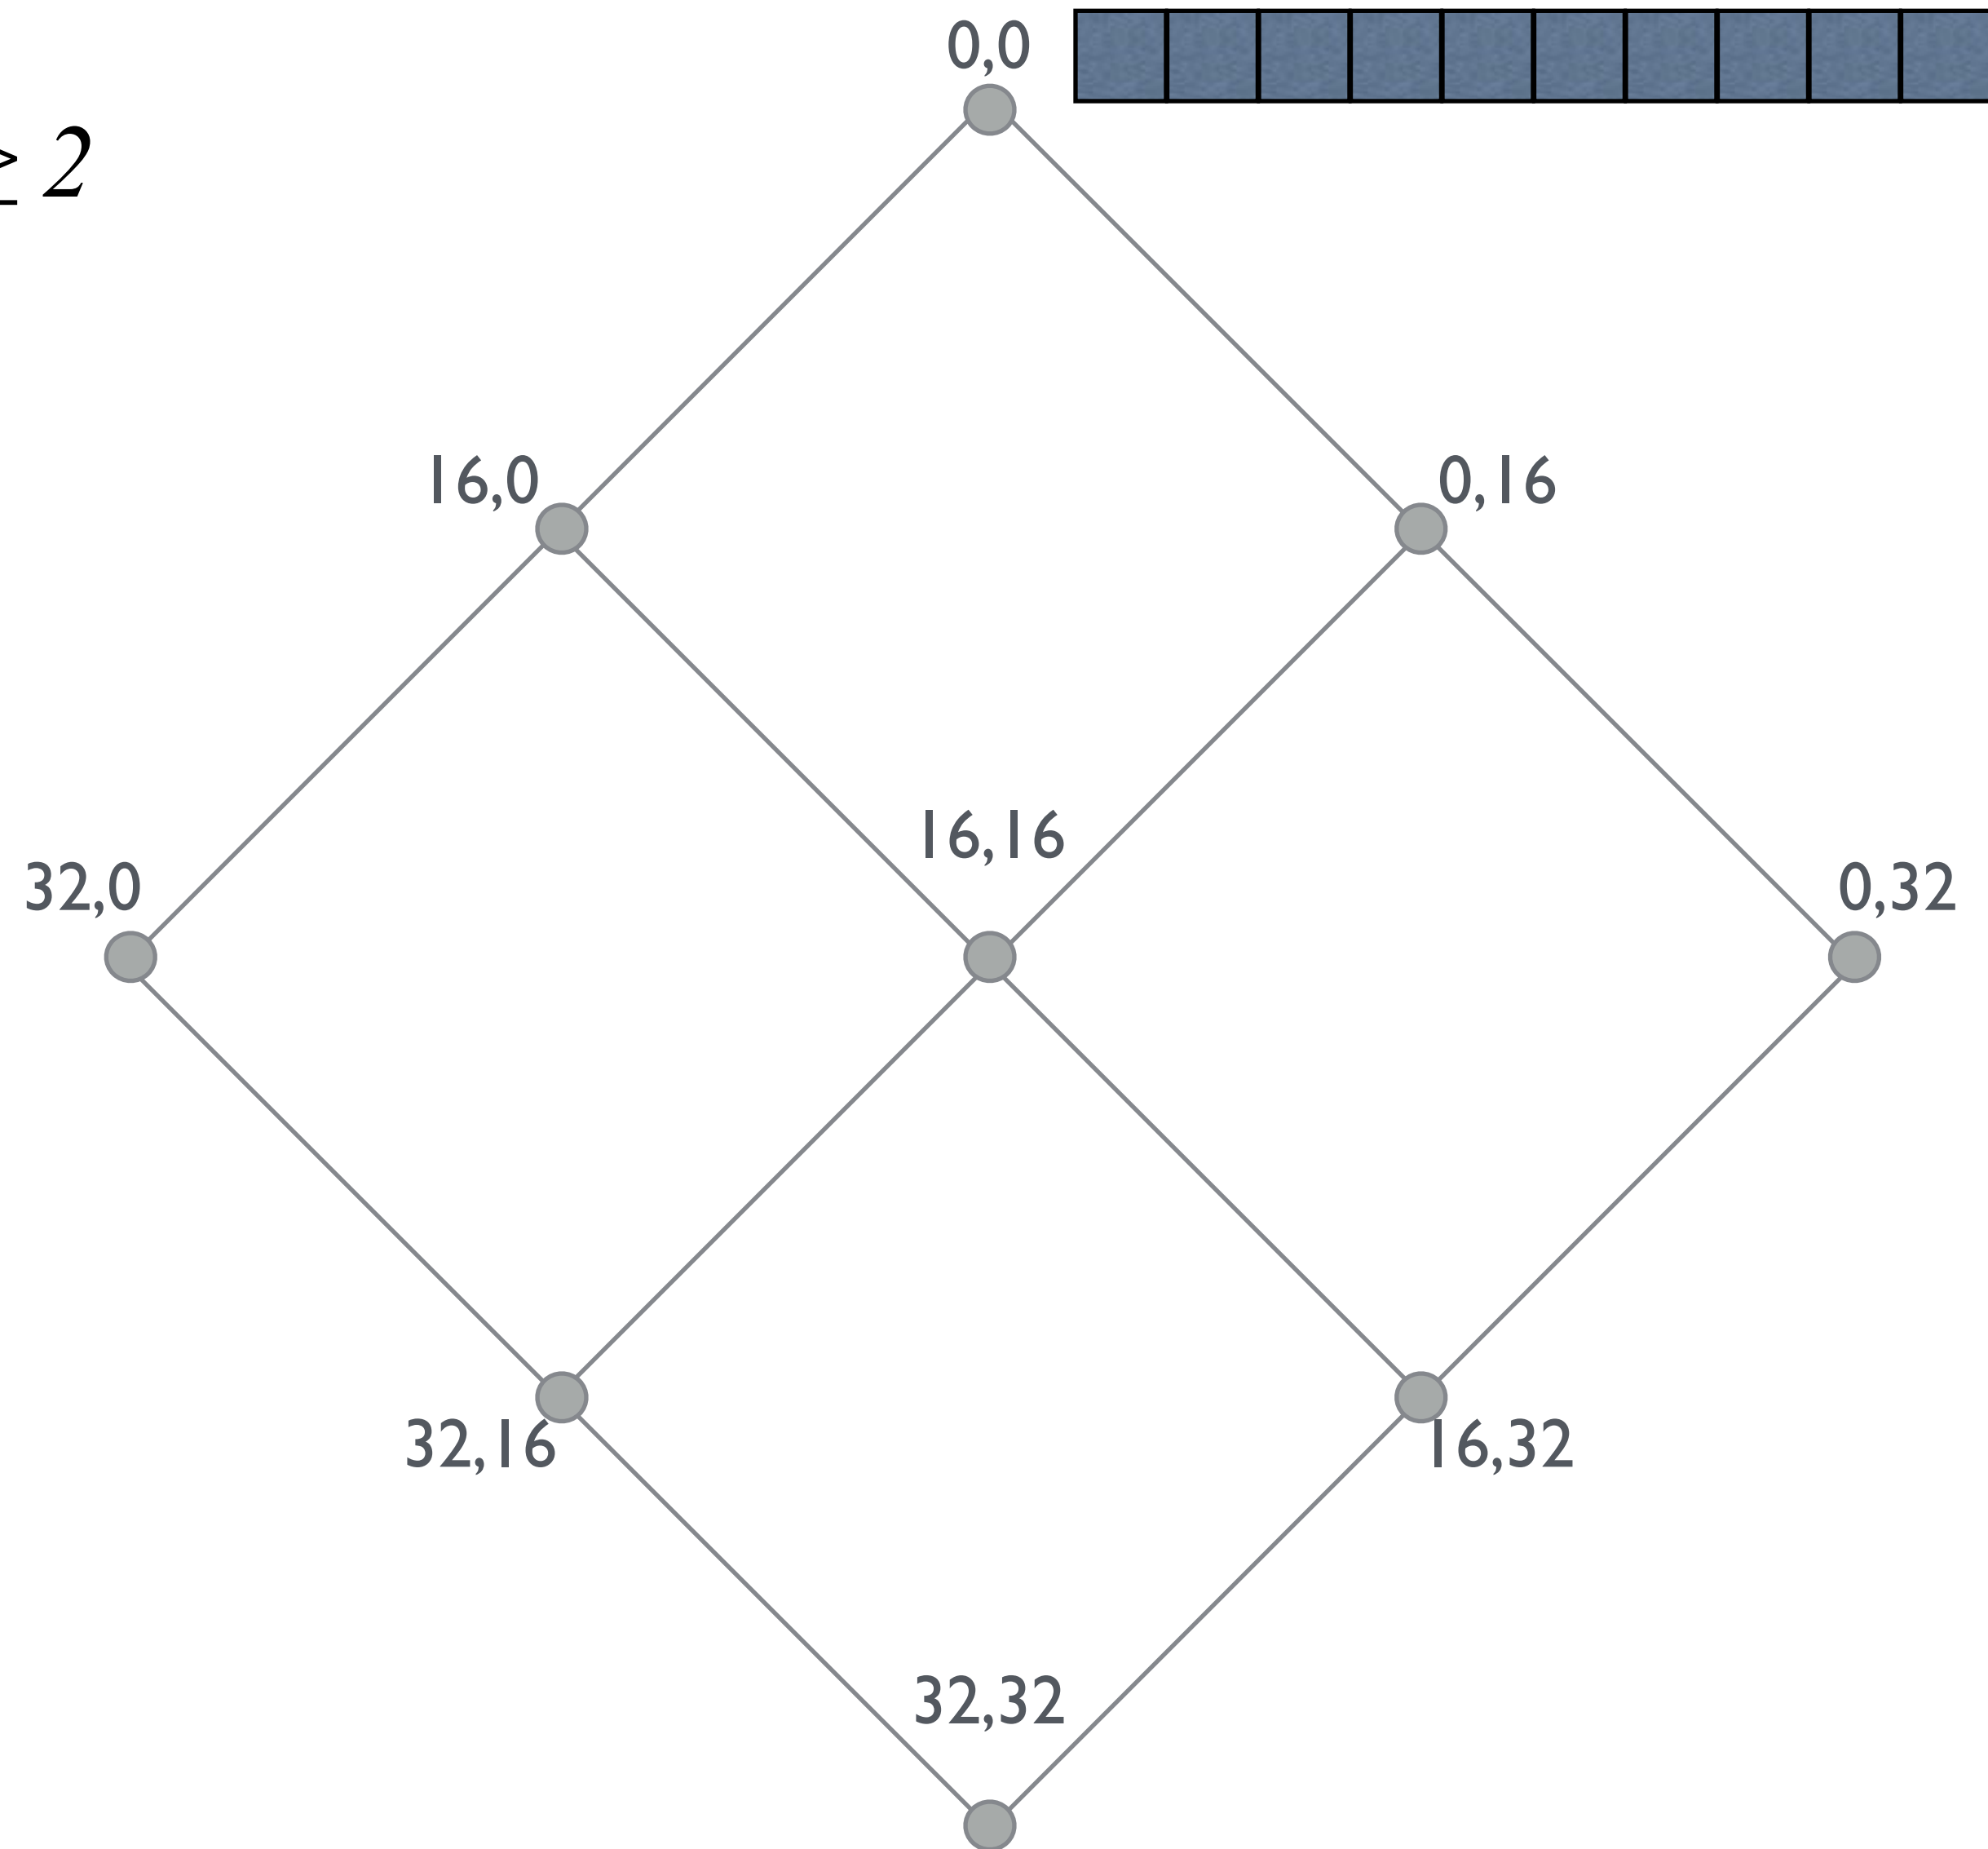
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



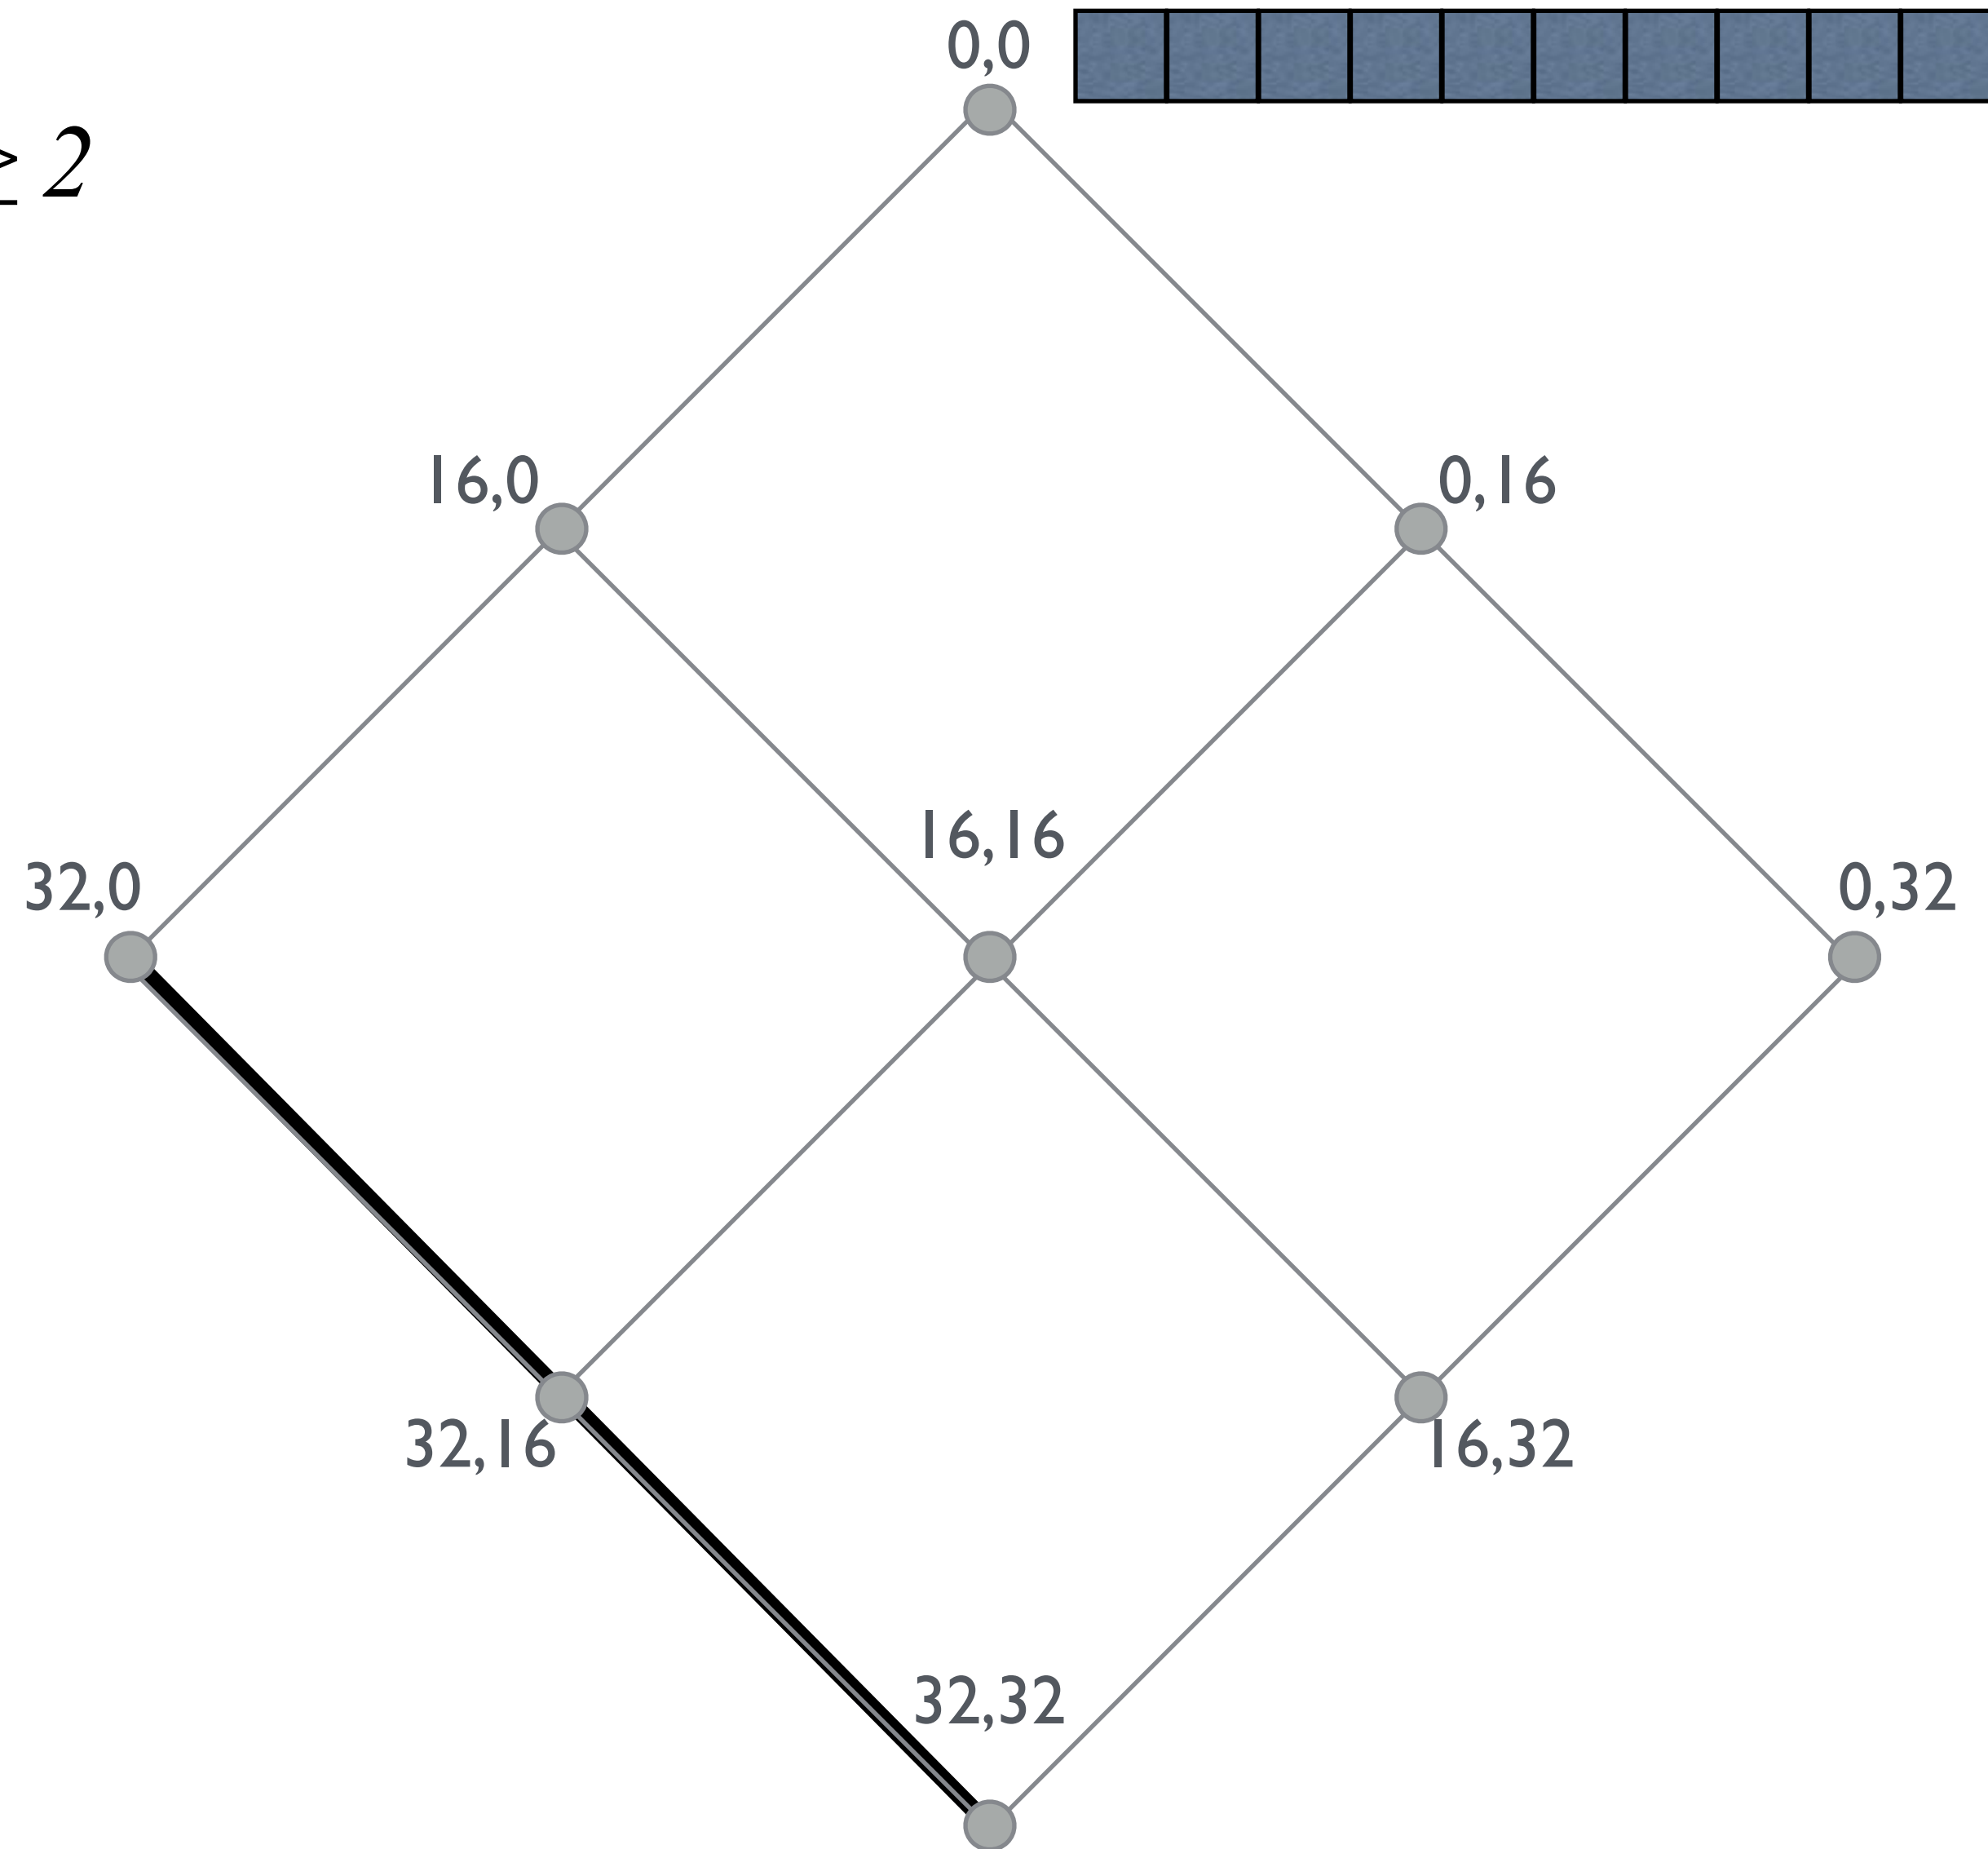
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



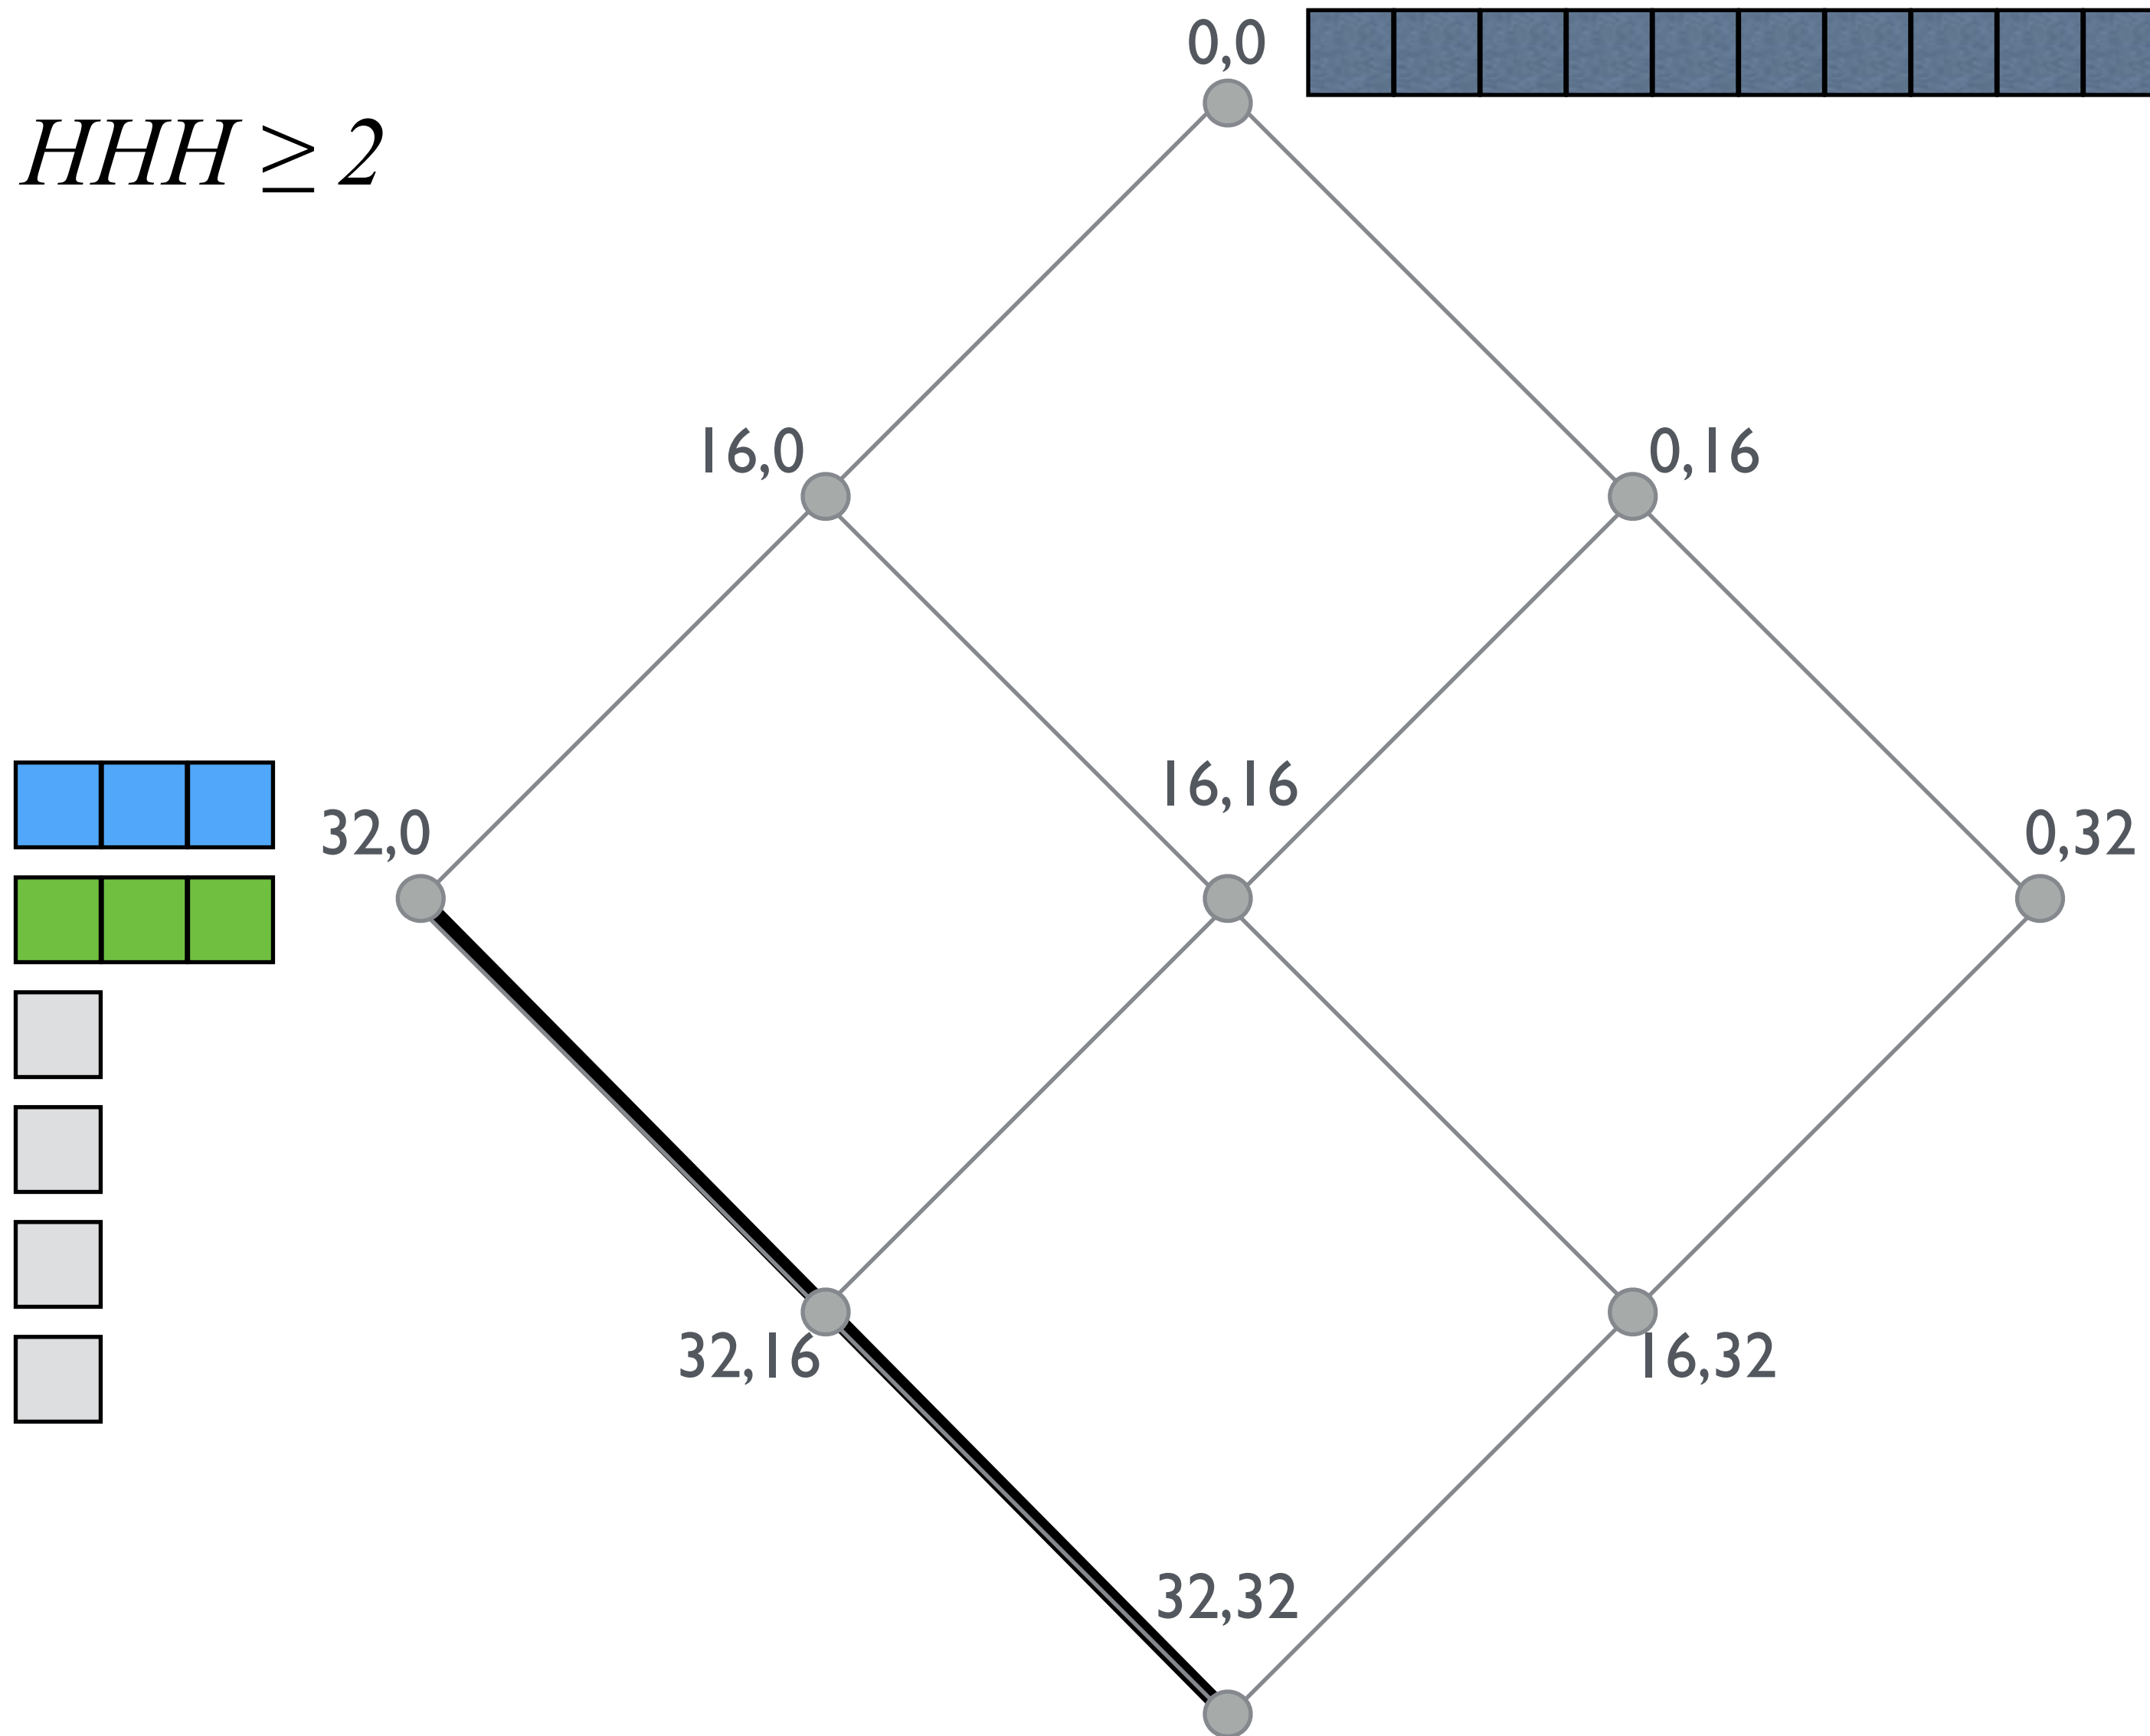
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



# RLS Illustrated

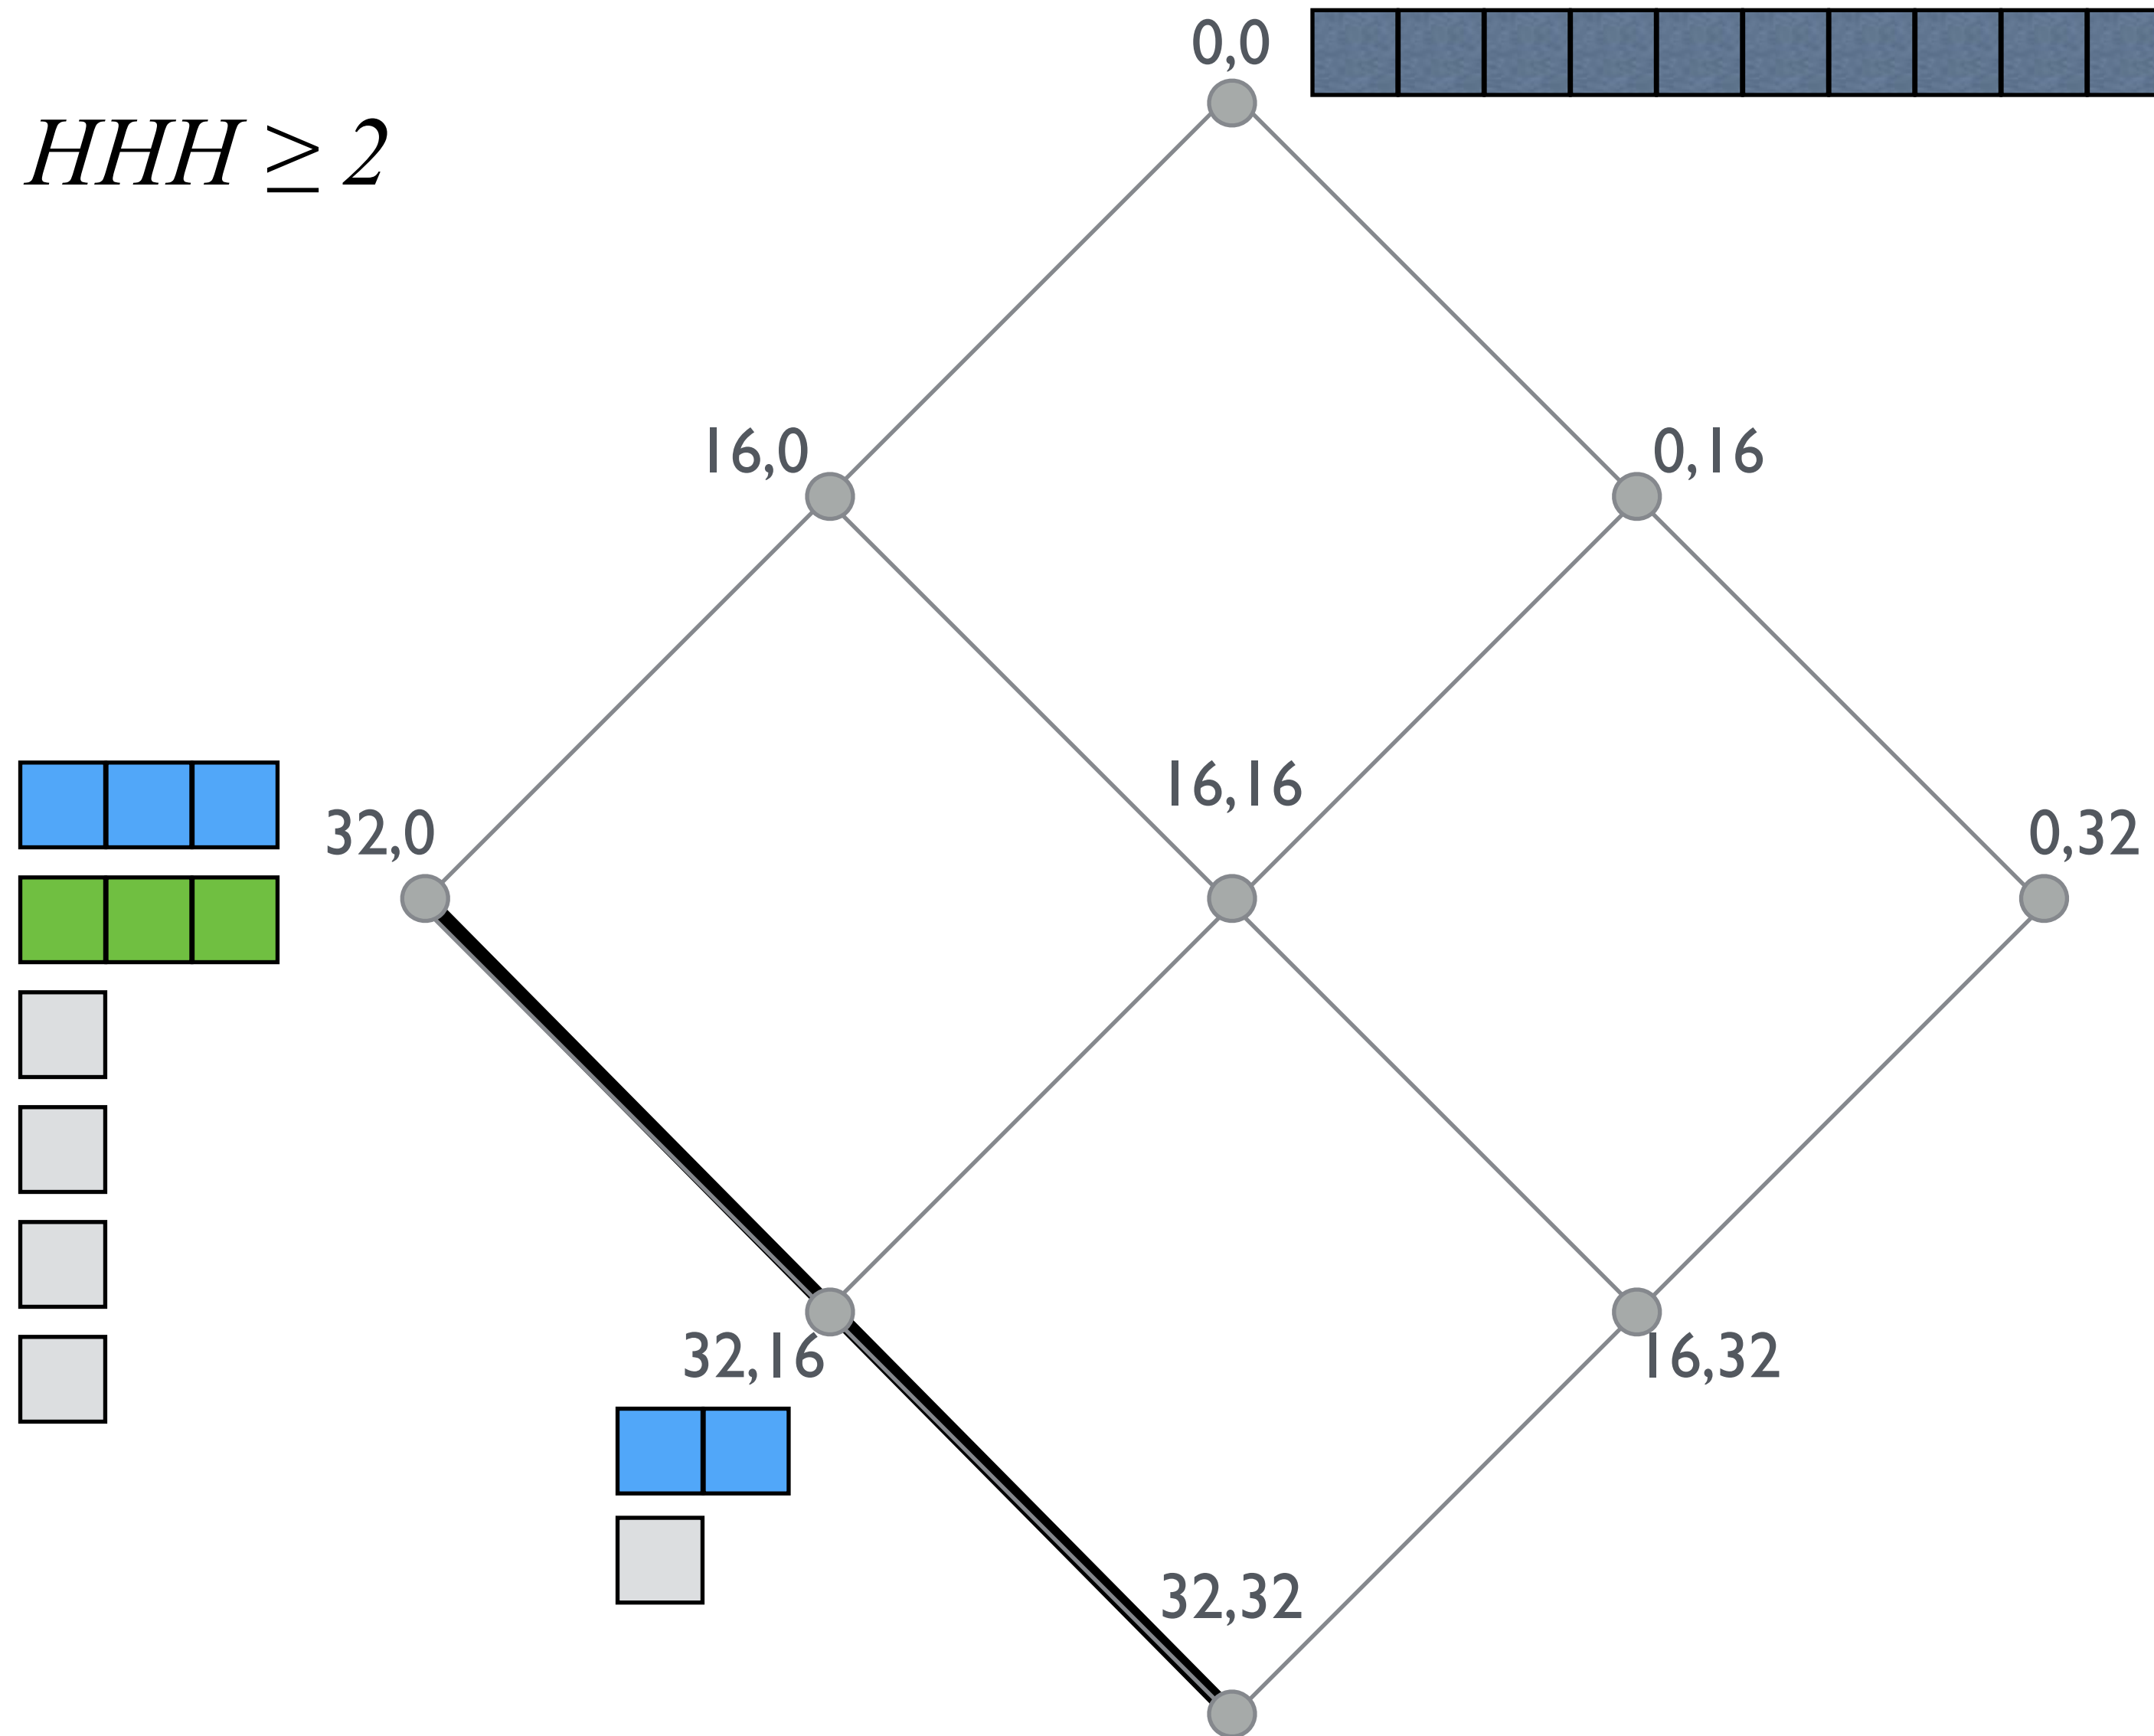
*10 inputs,  $HHH \geq 2$*





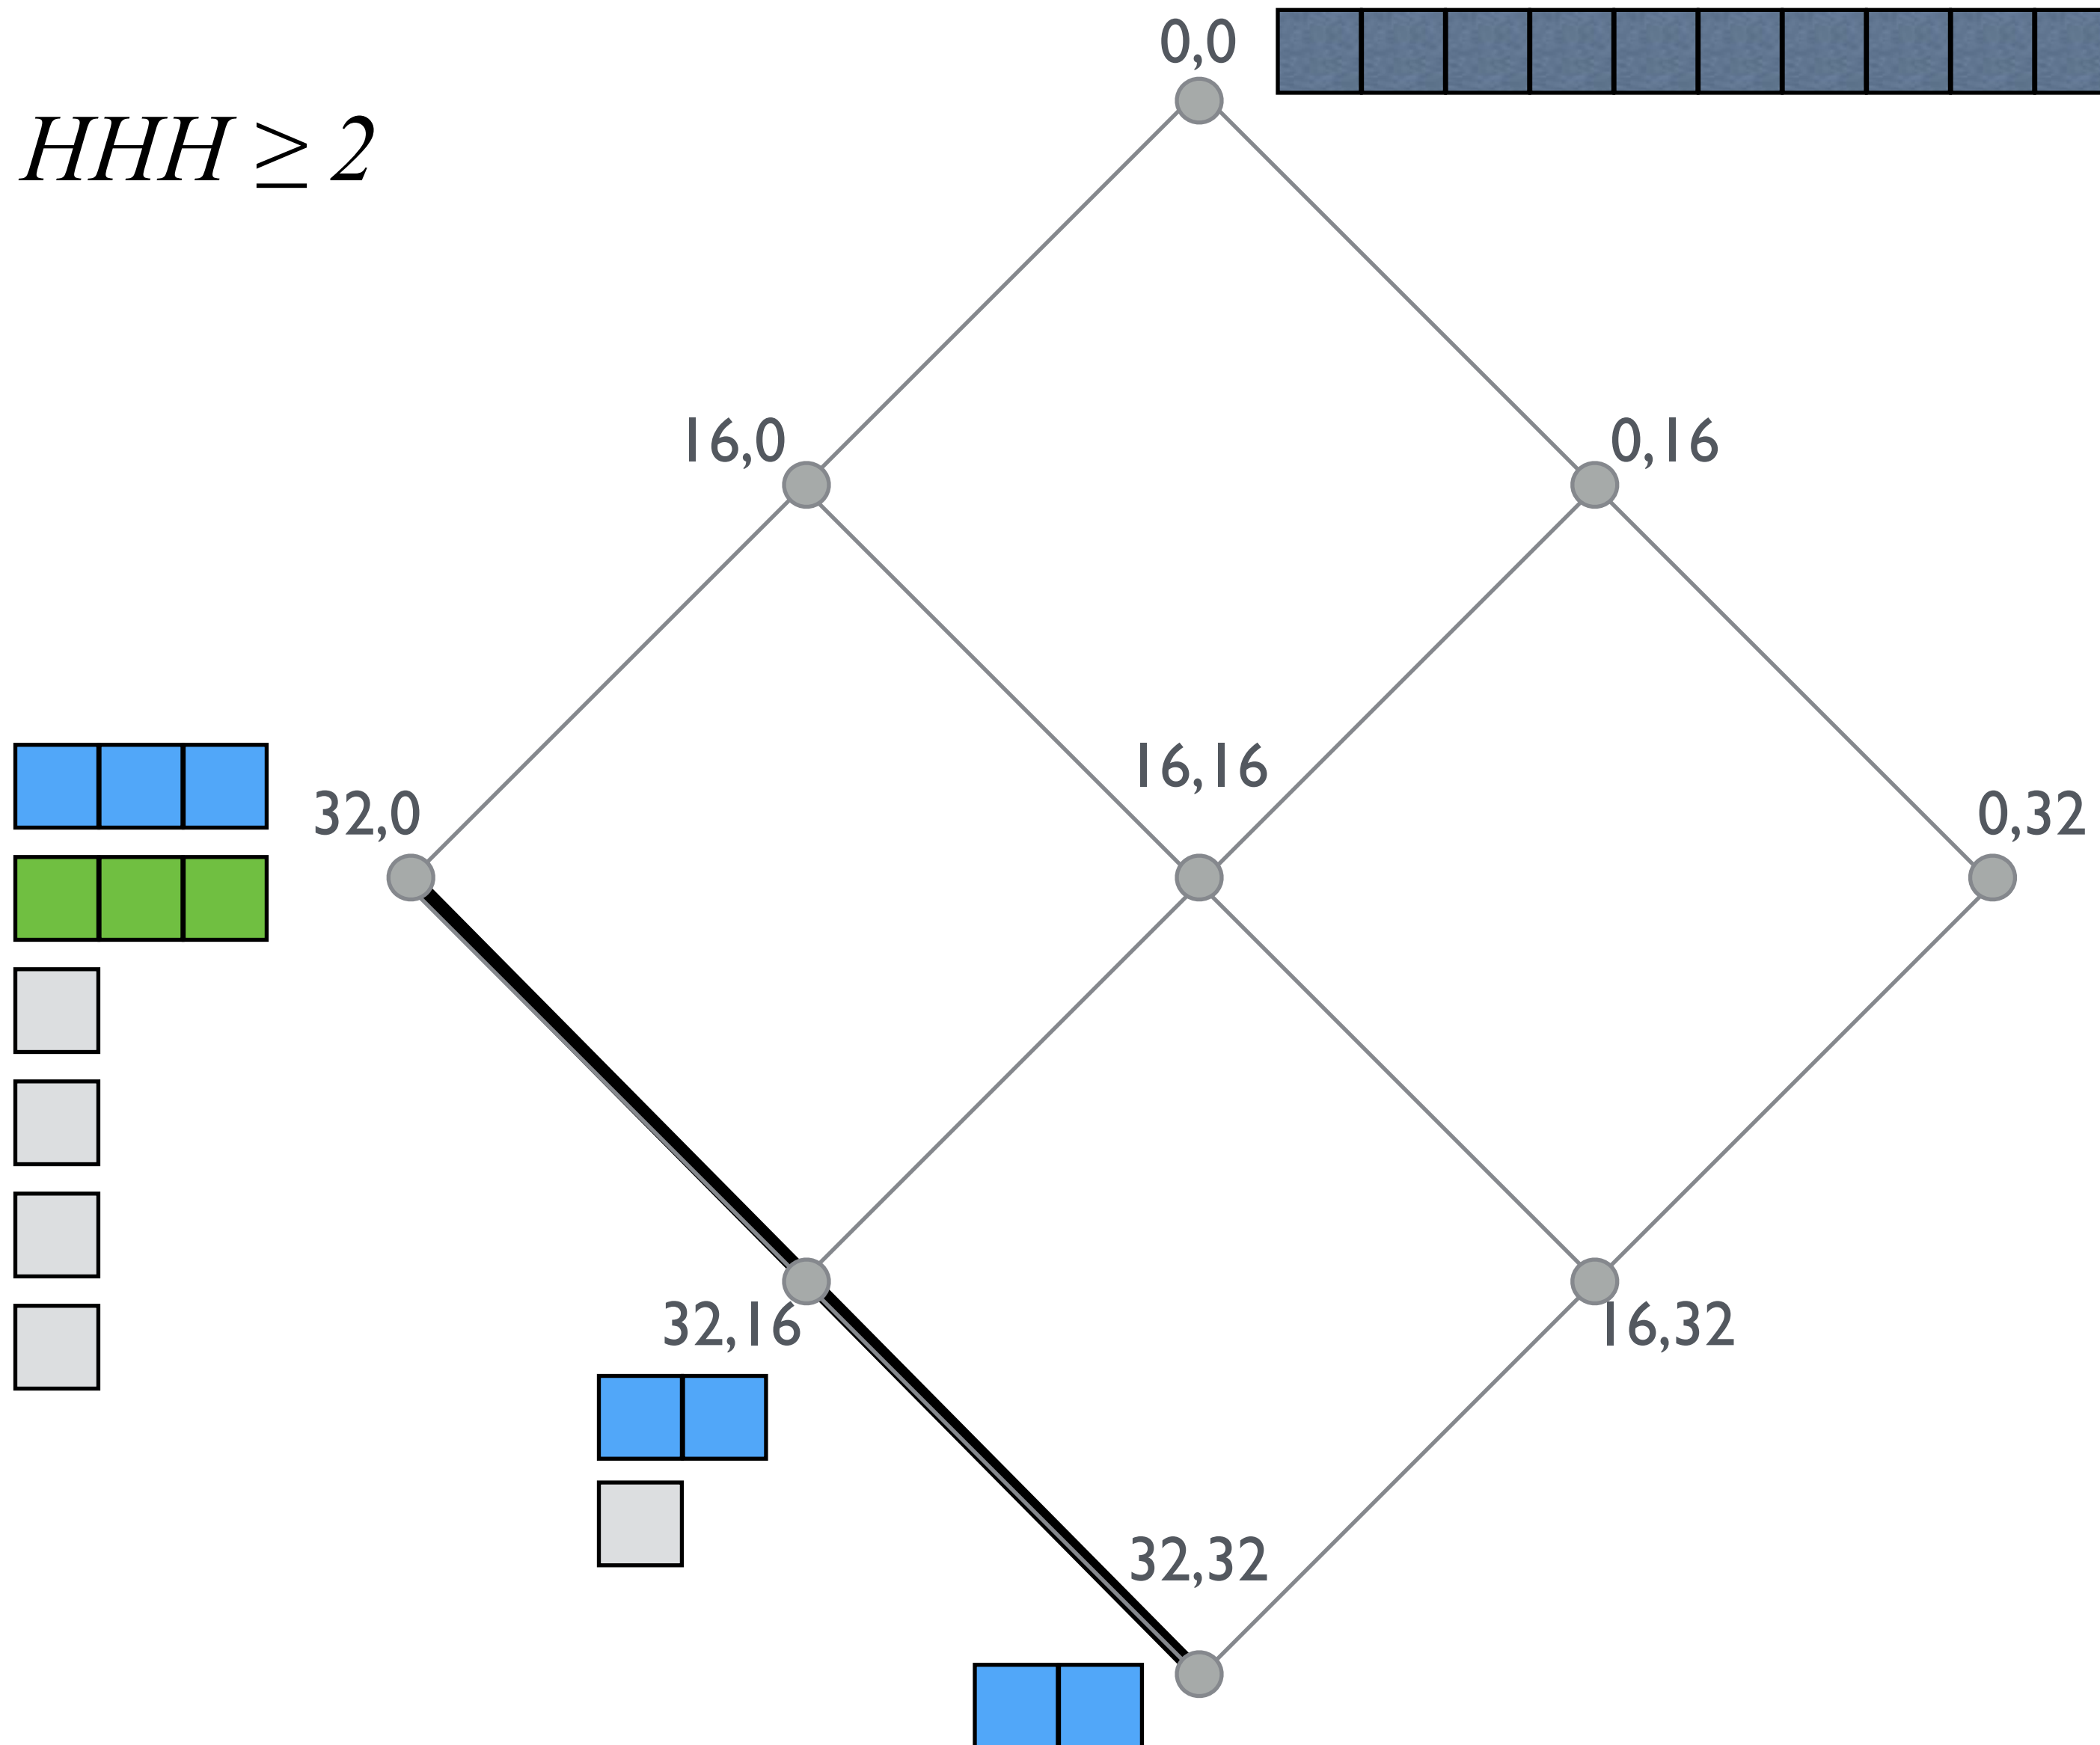
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



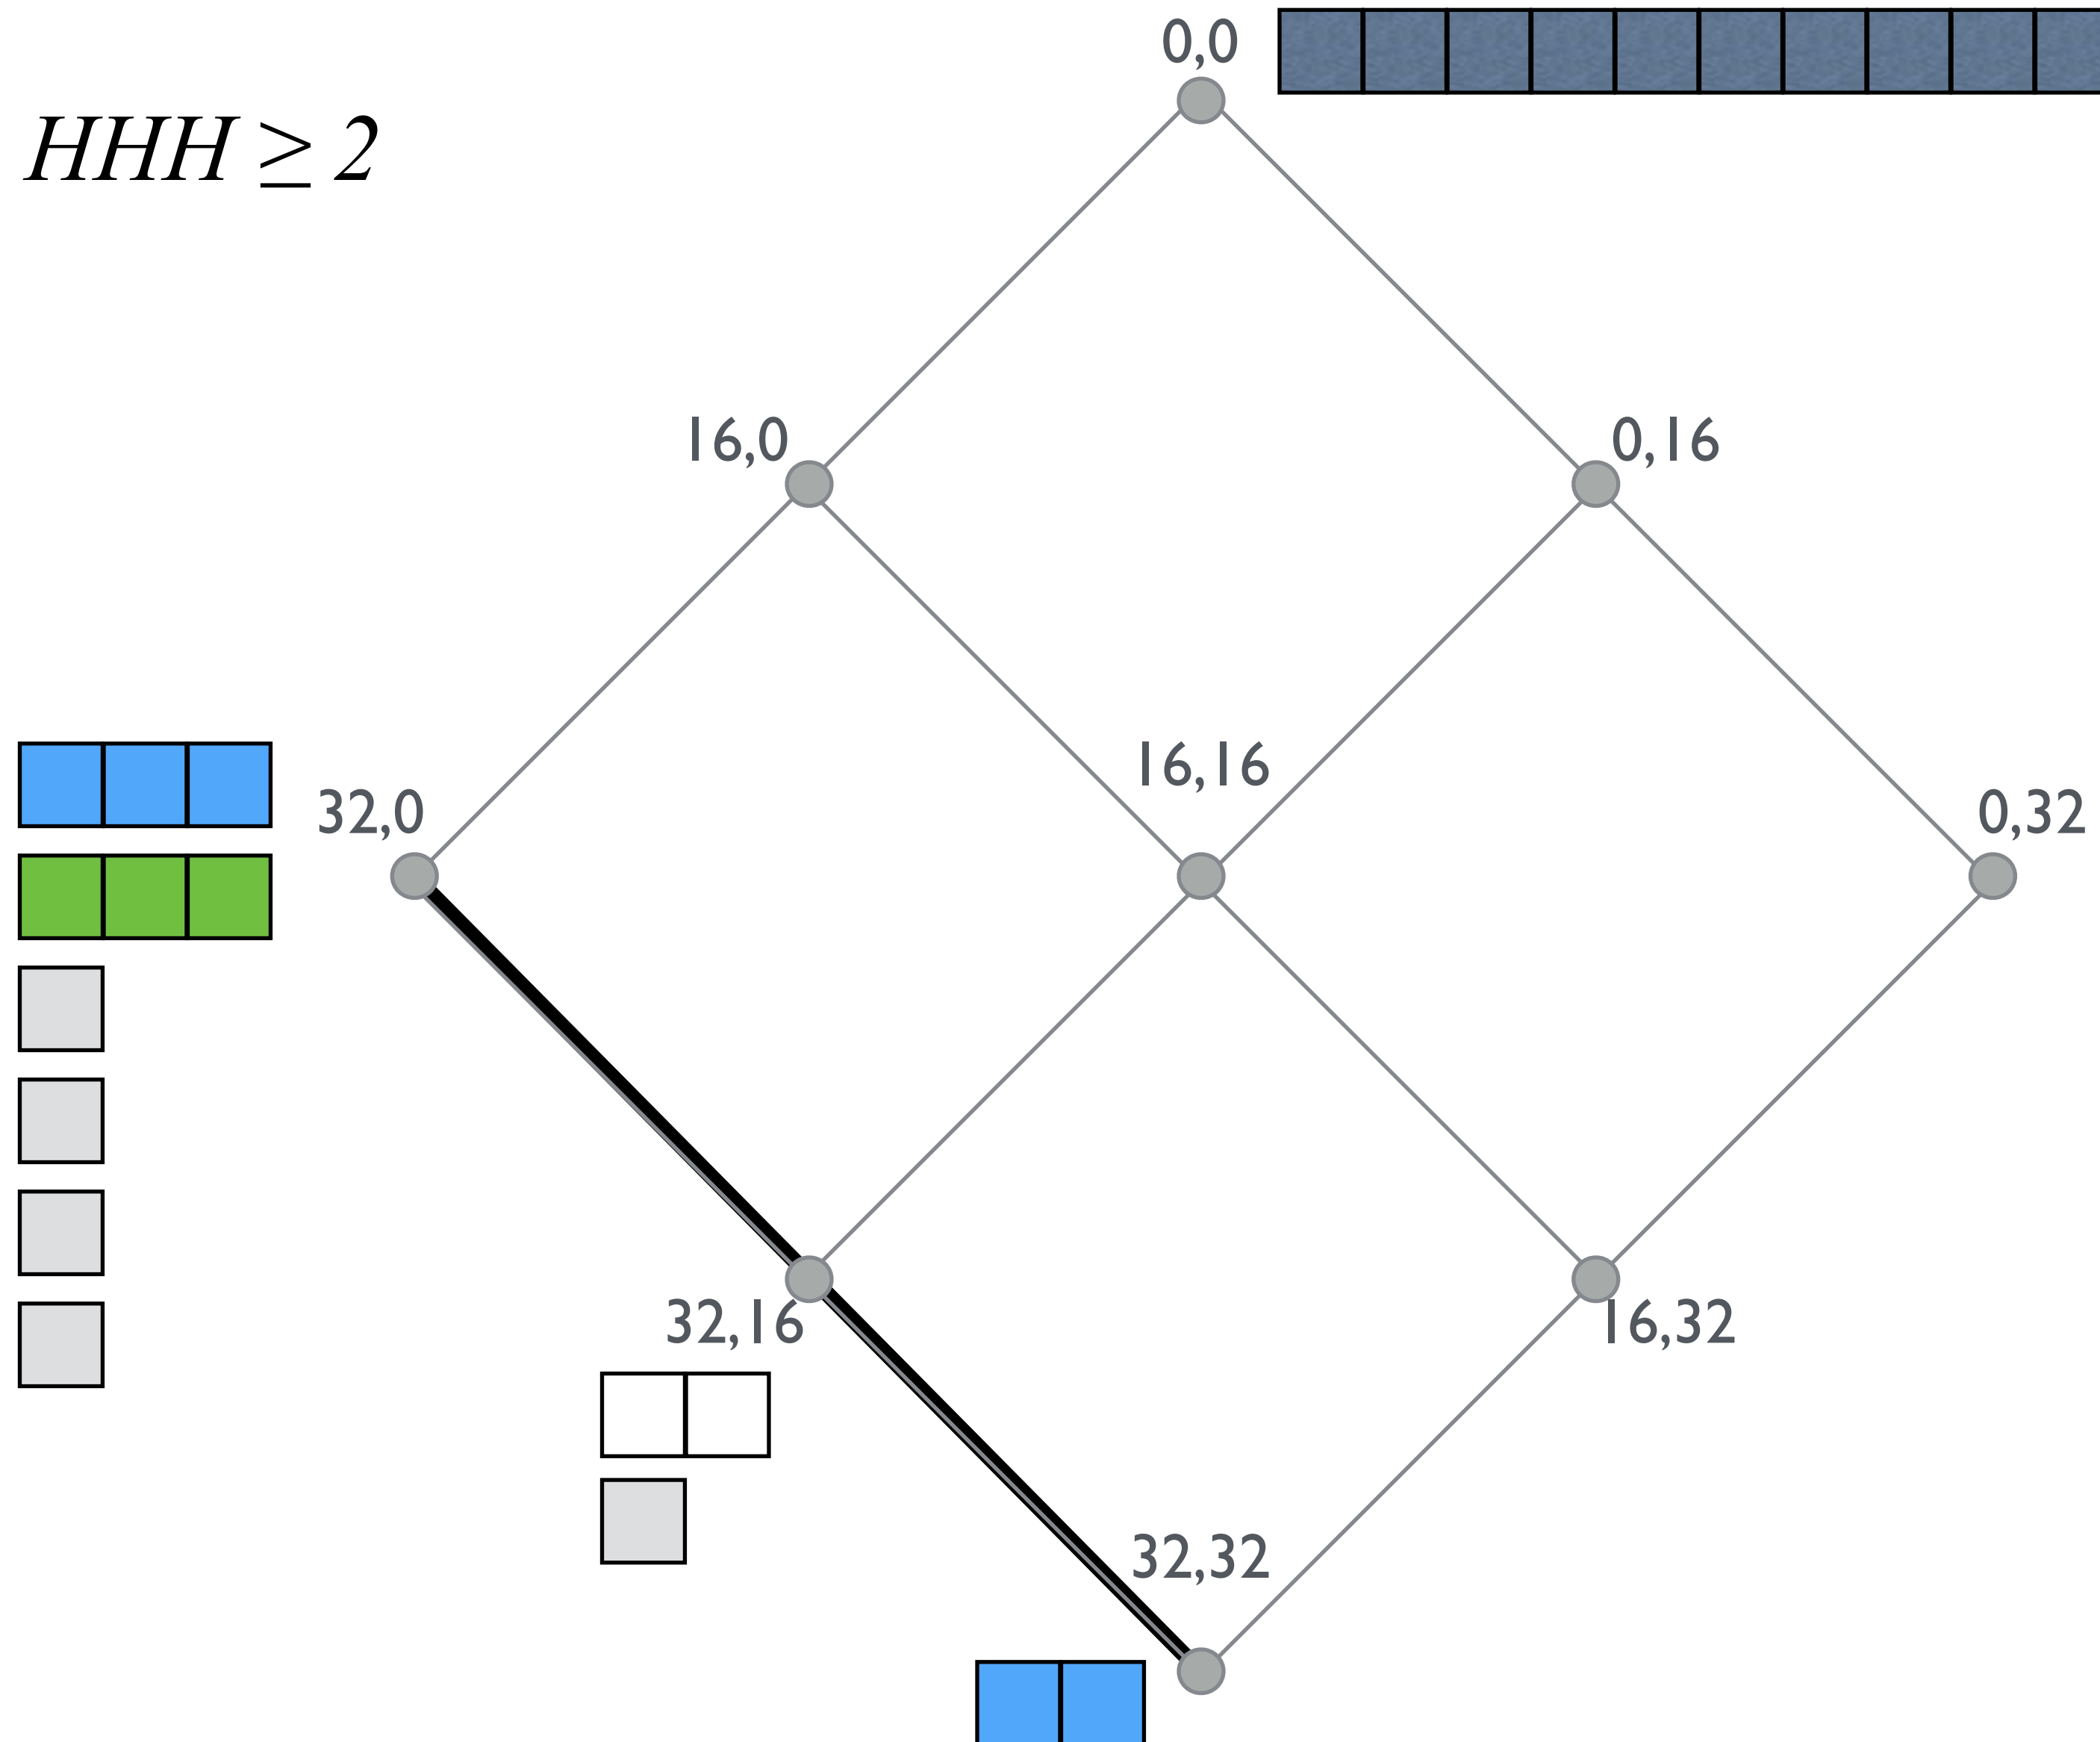
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



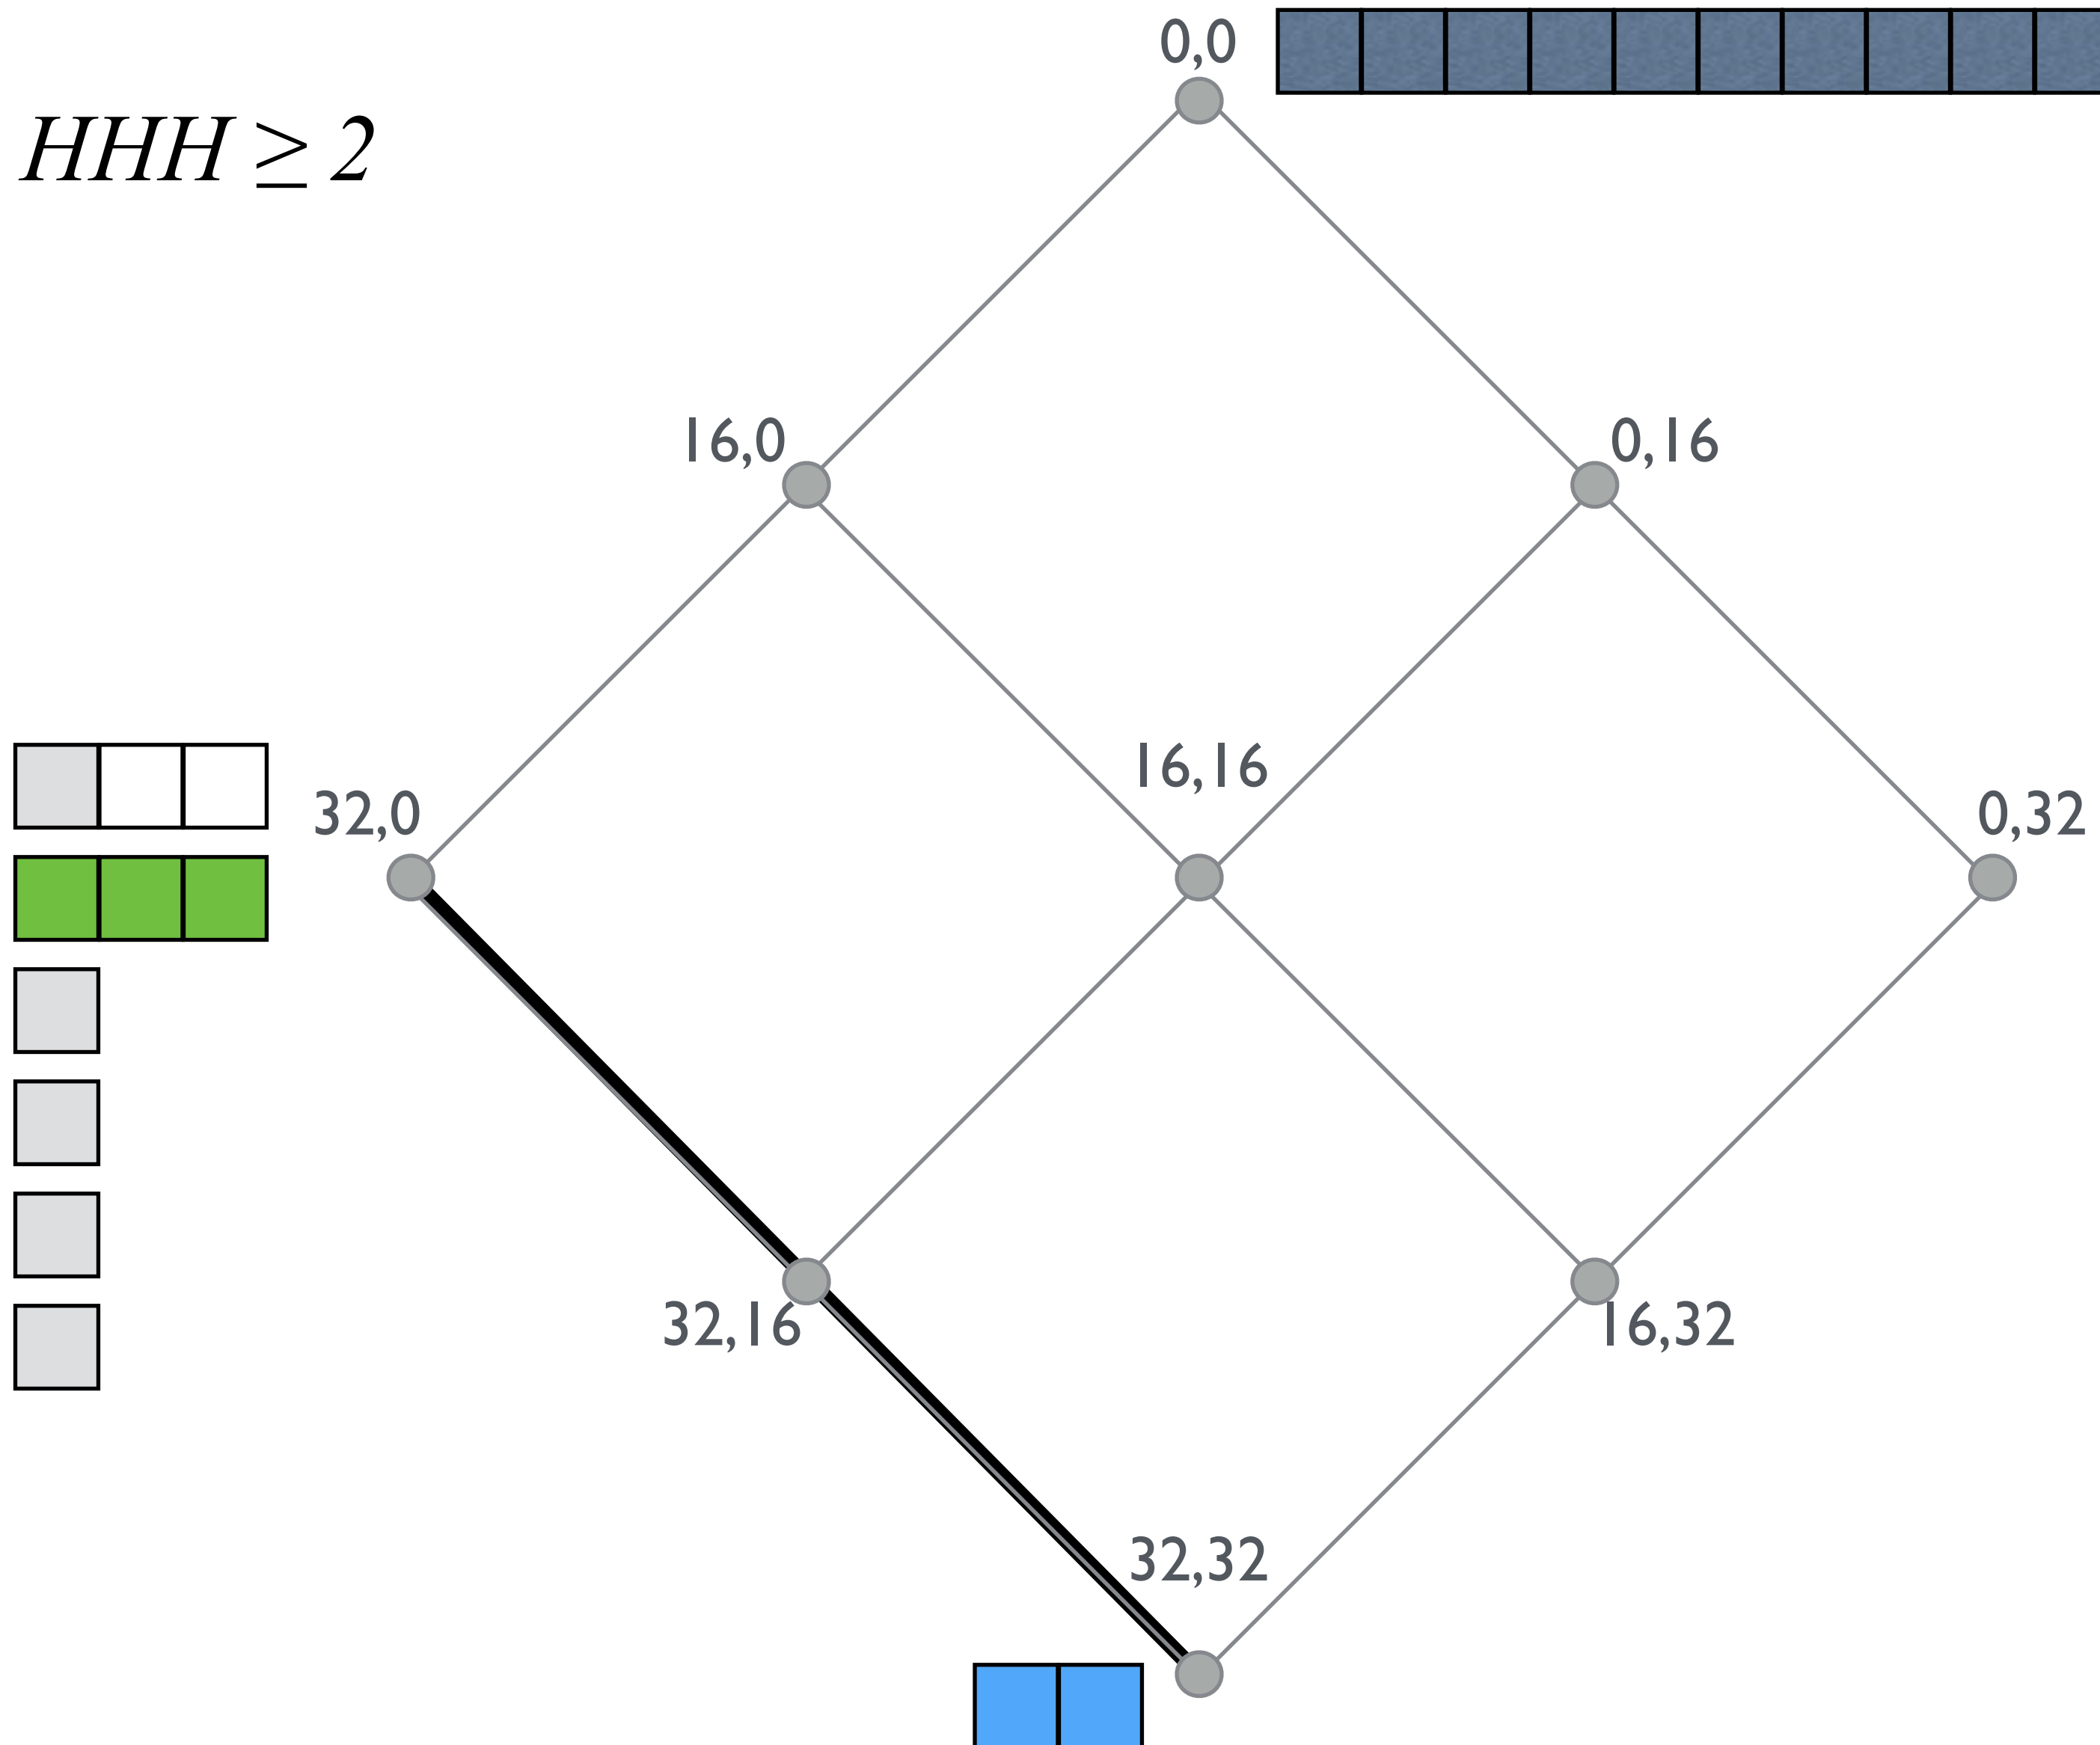
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



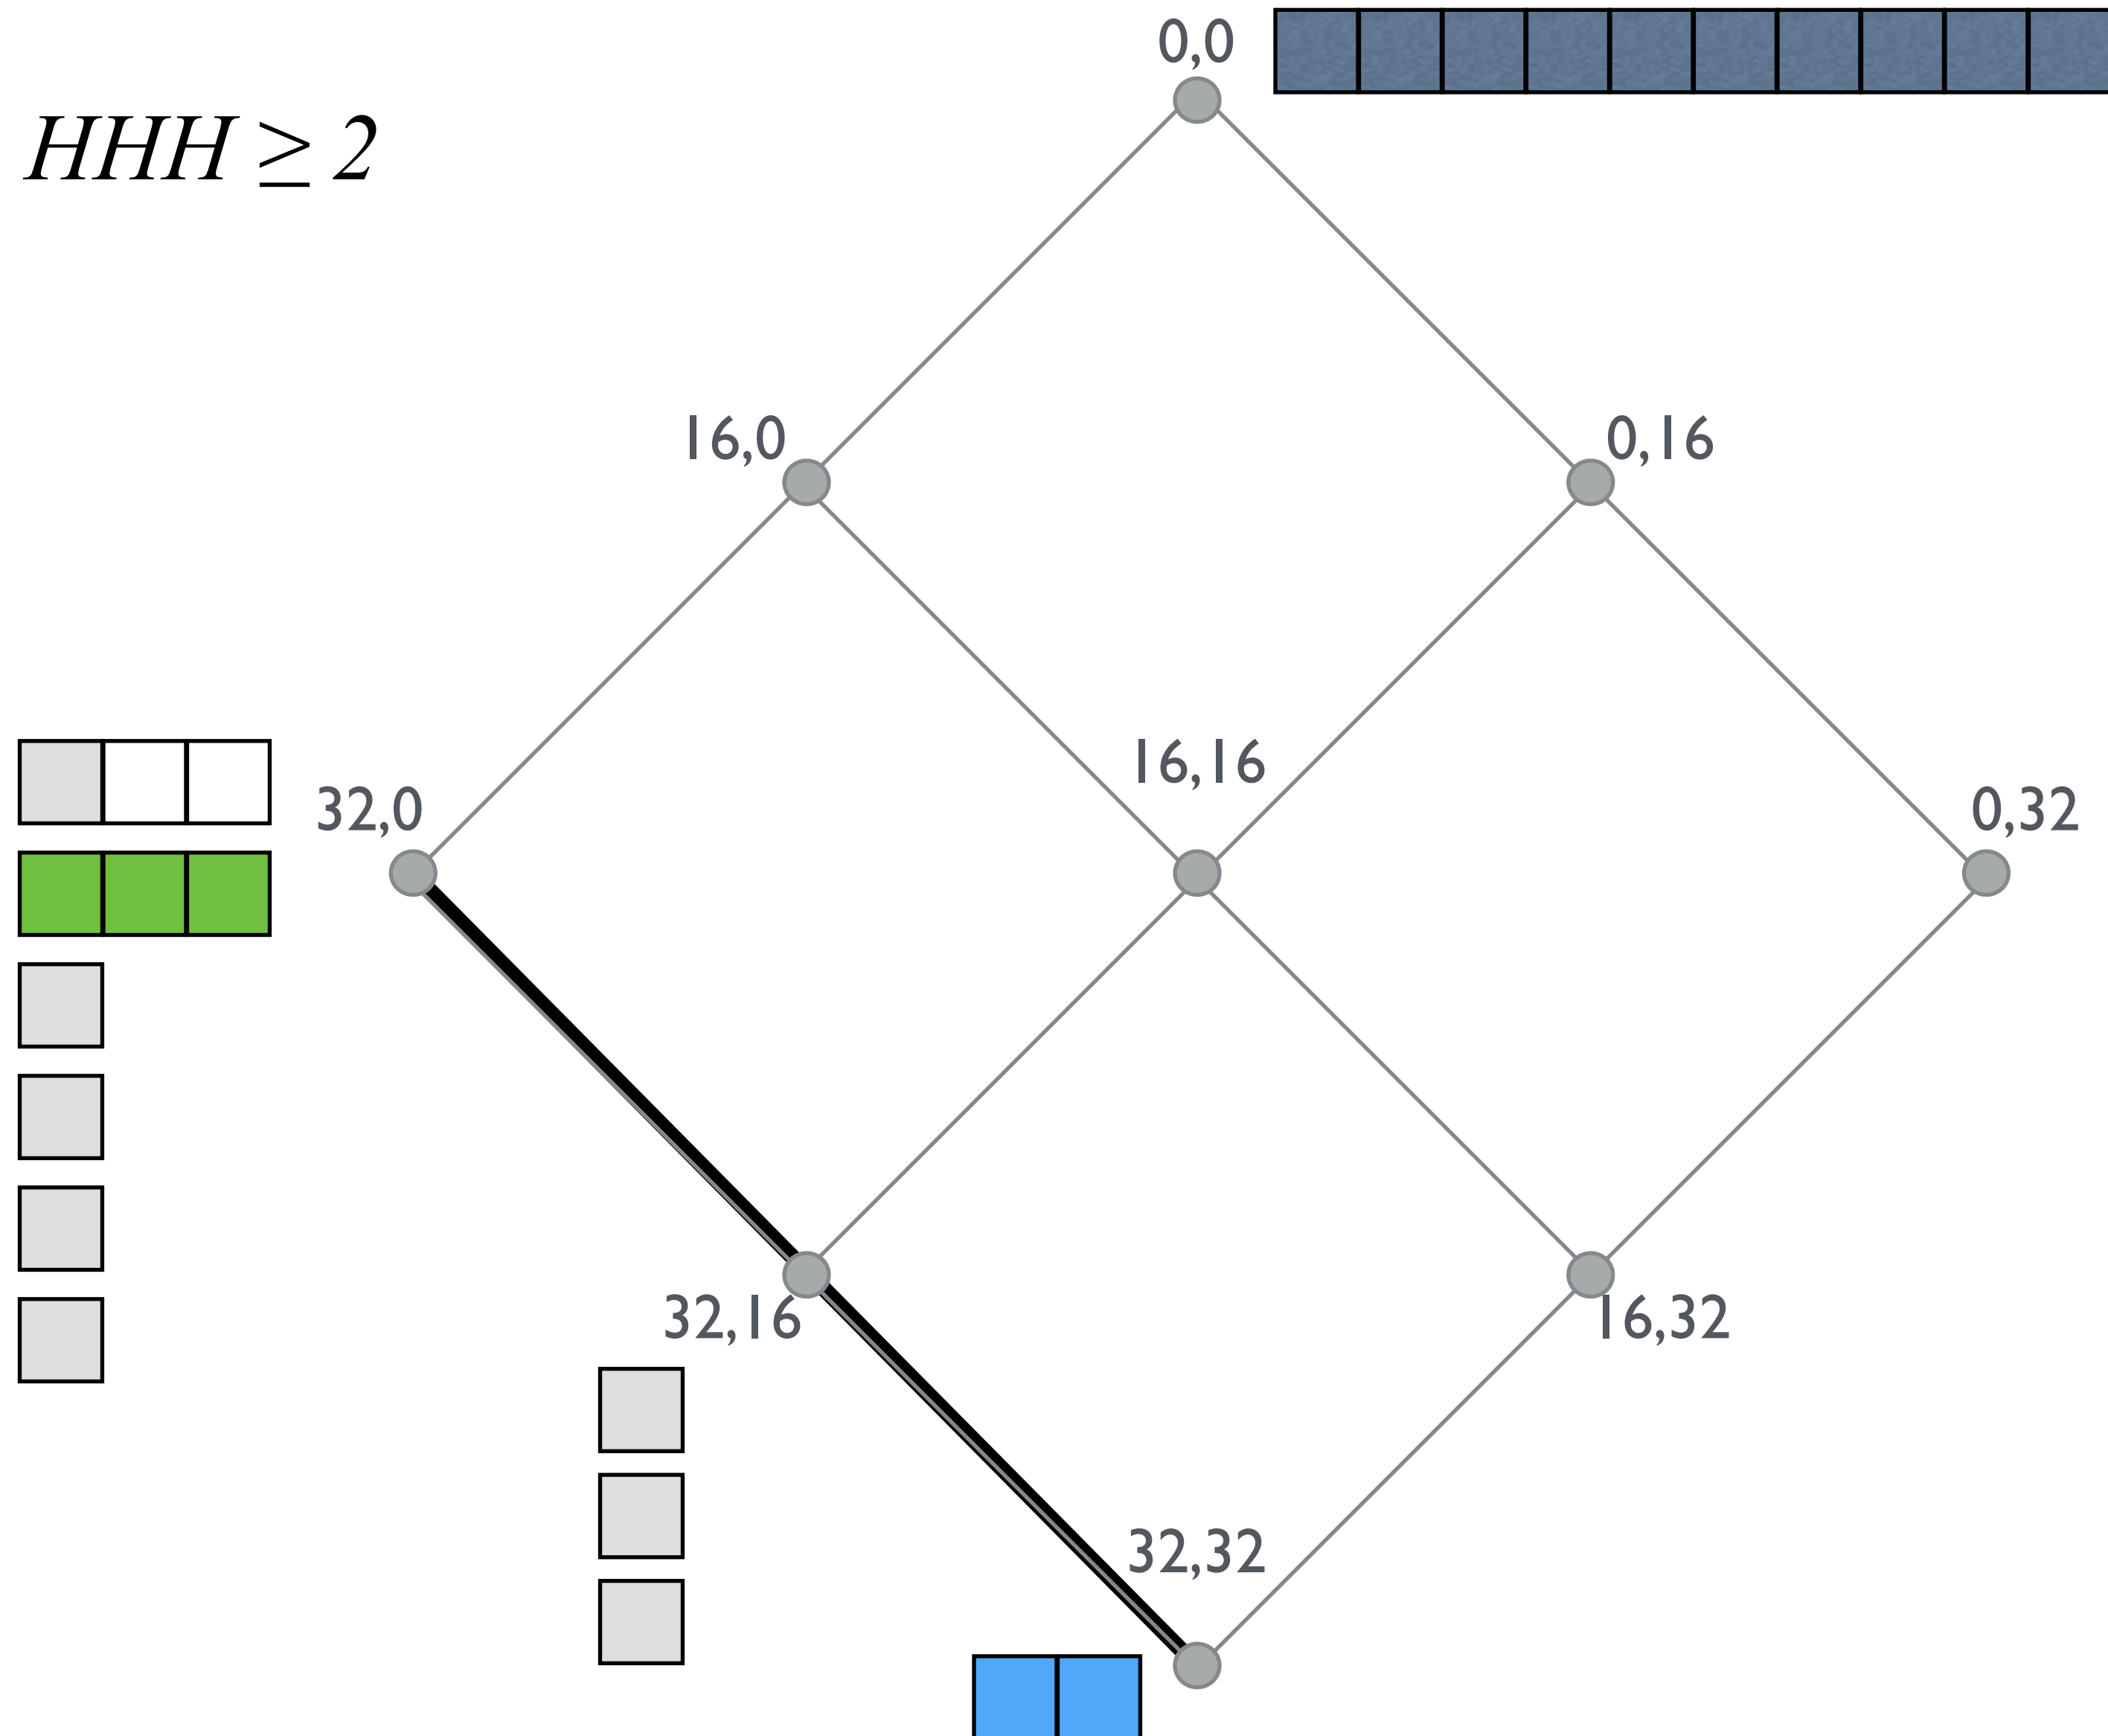
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



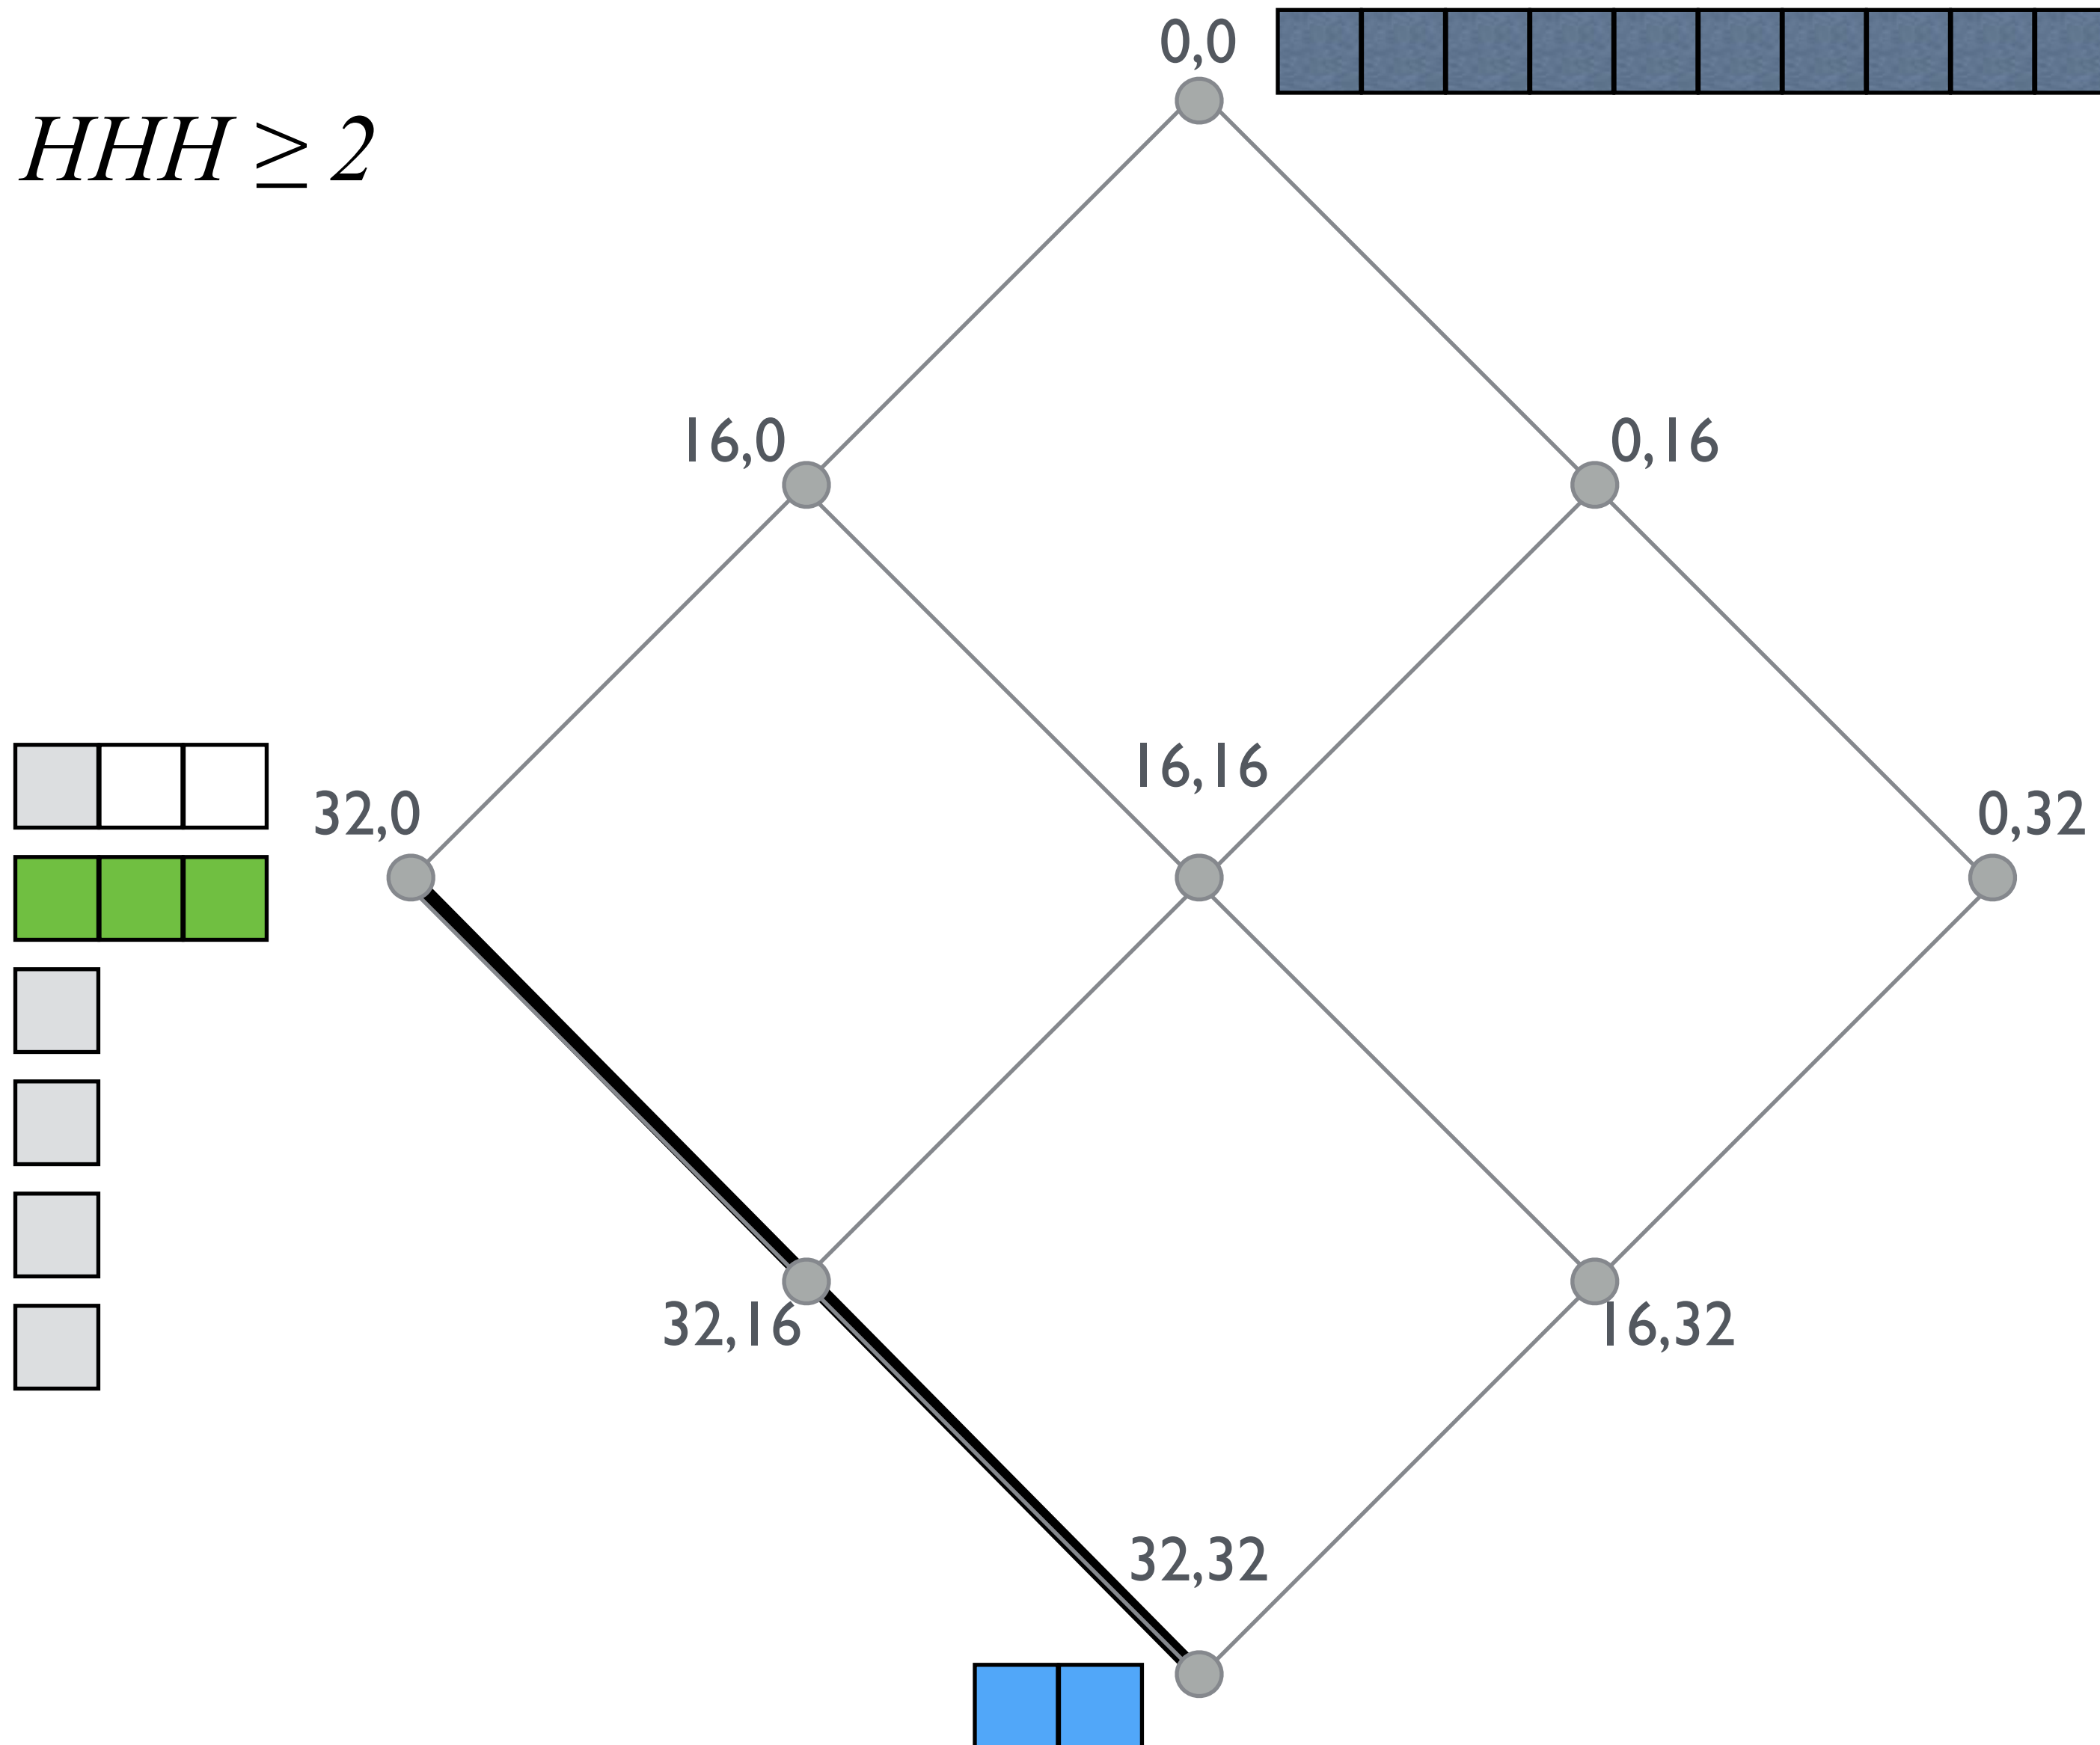
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



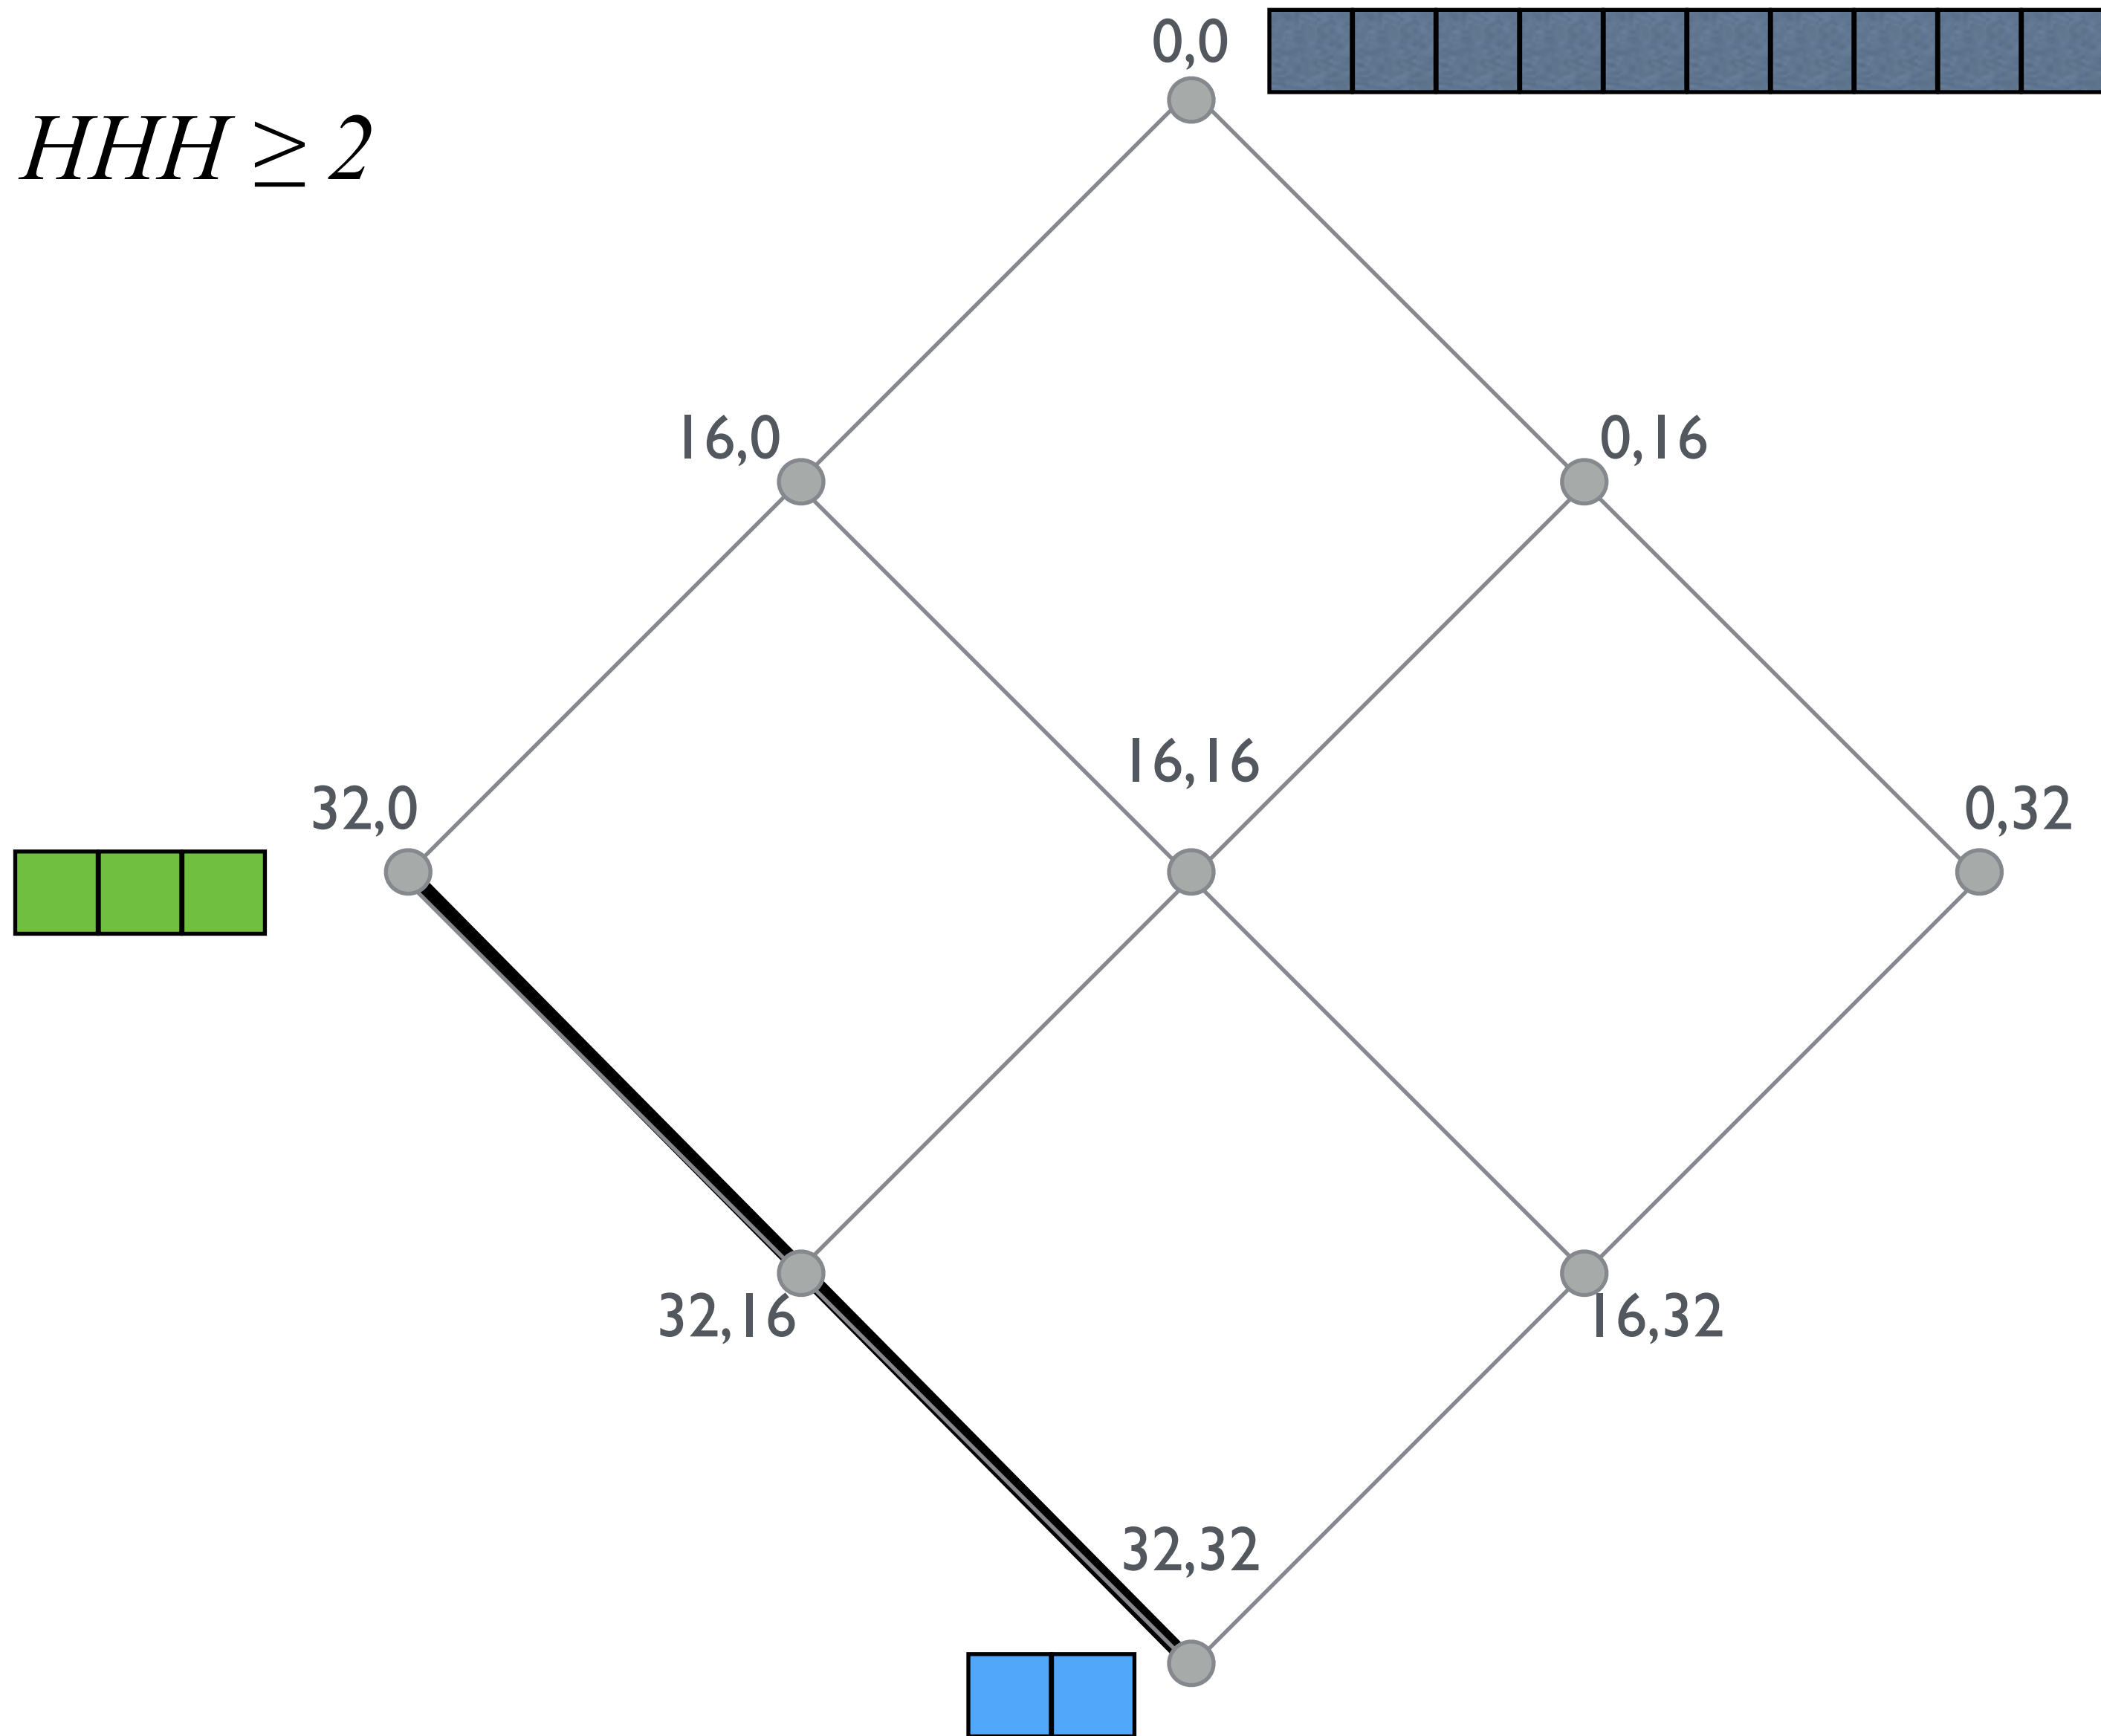
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



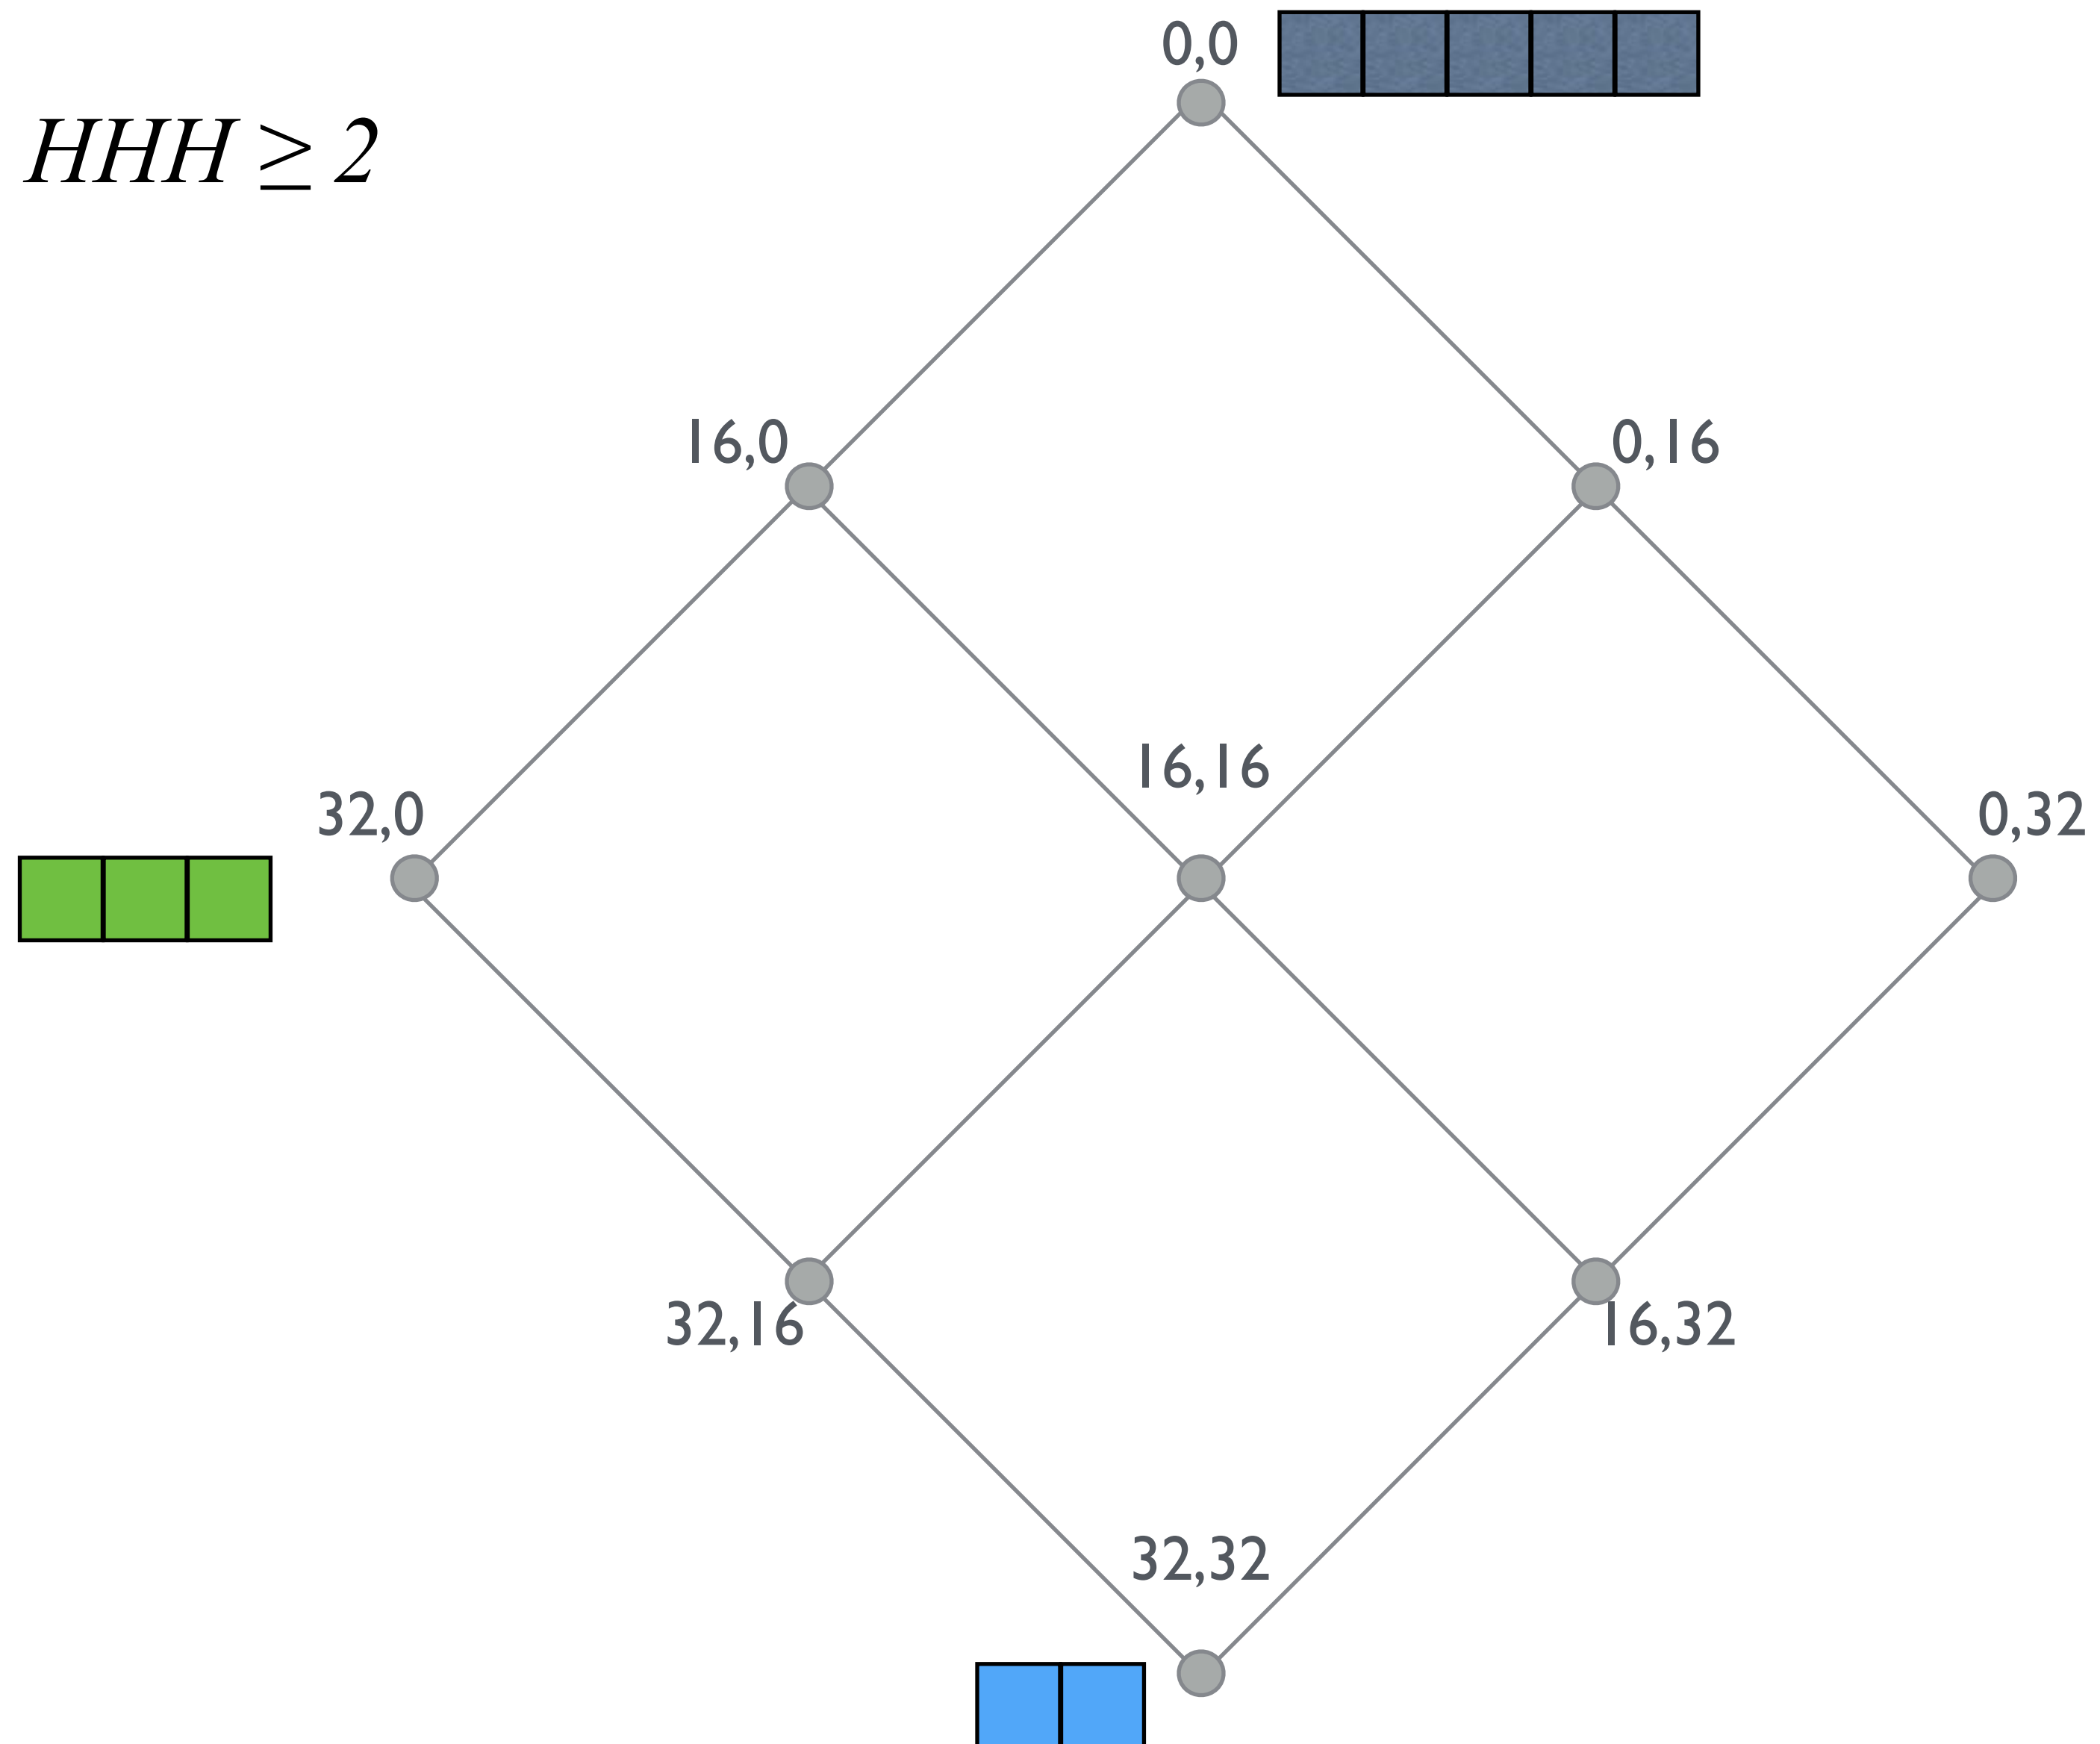
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



# RLS Illustrated

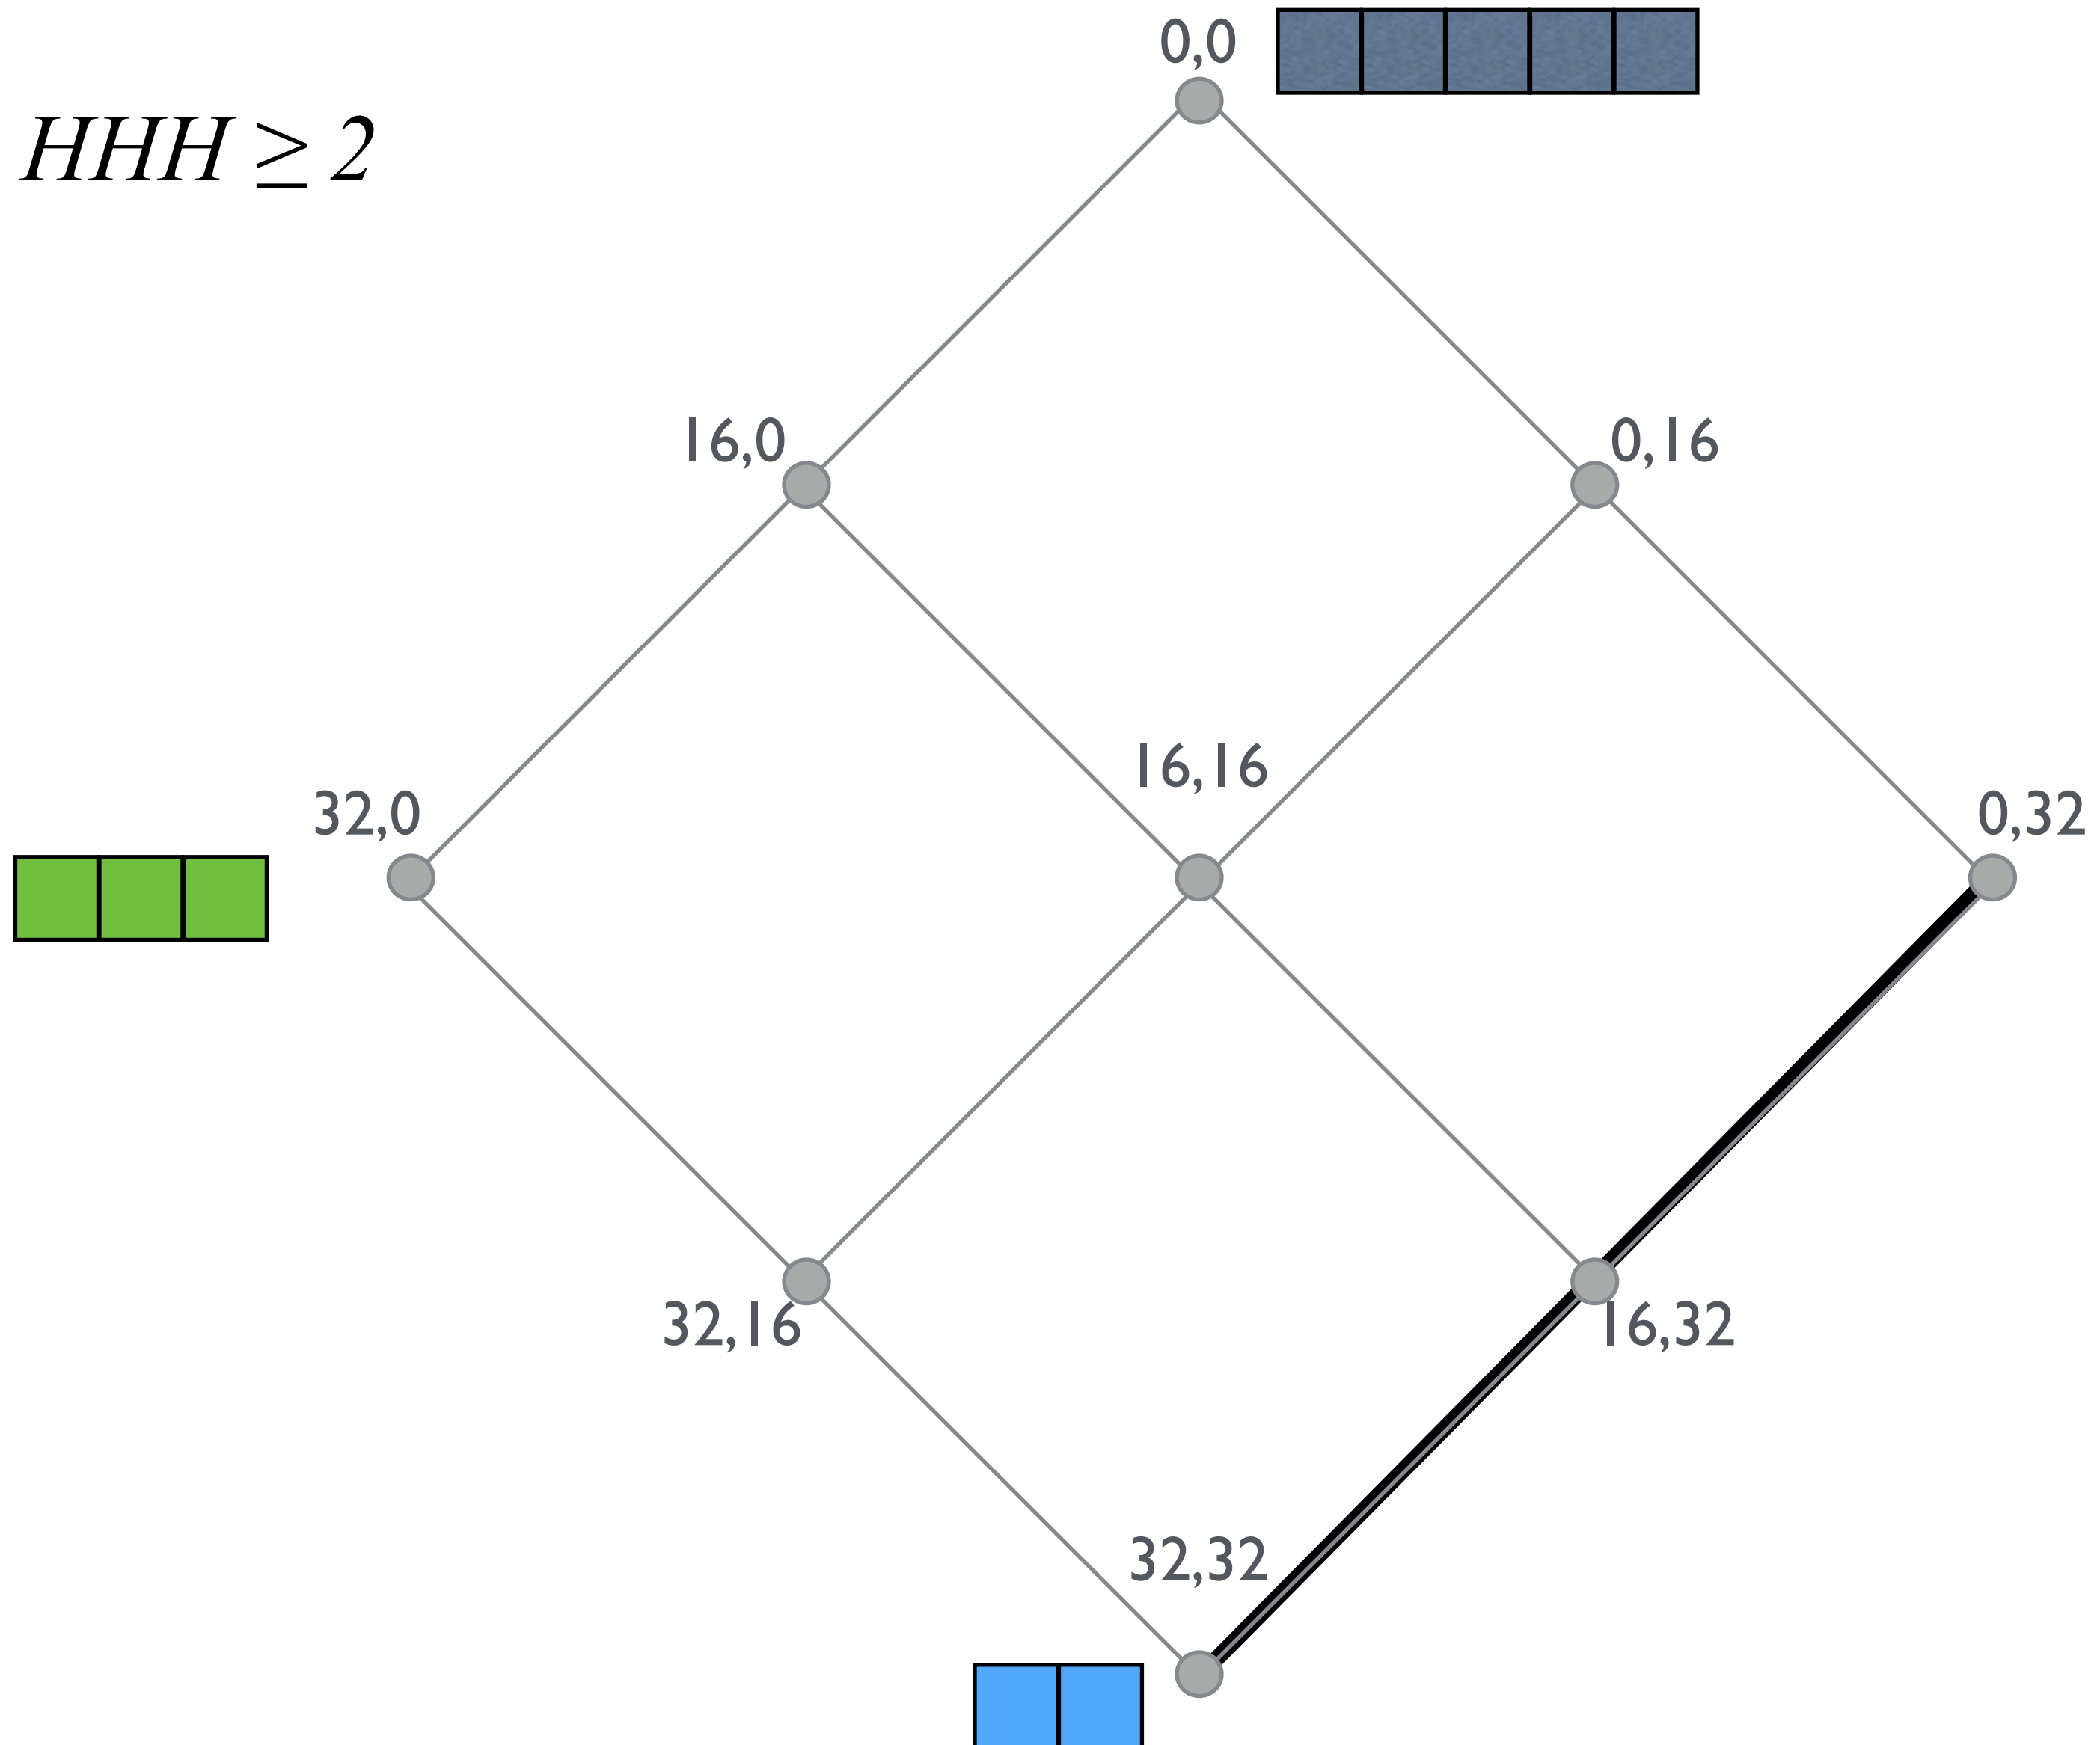
*10 inputs,  $HHH \geq 2$*





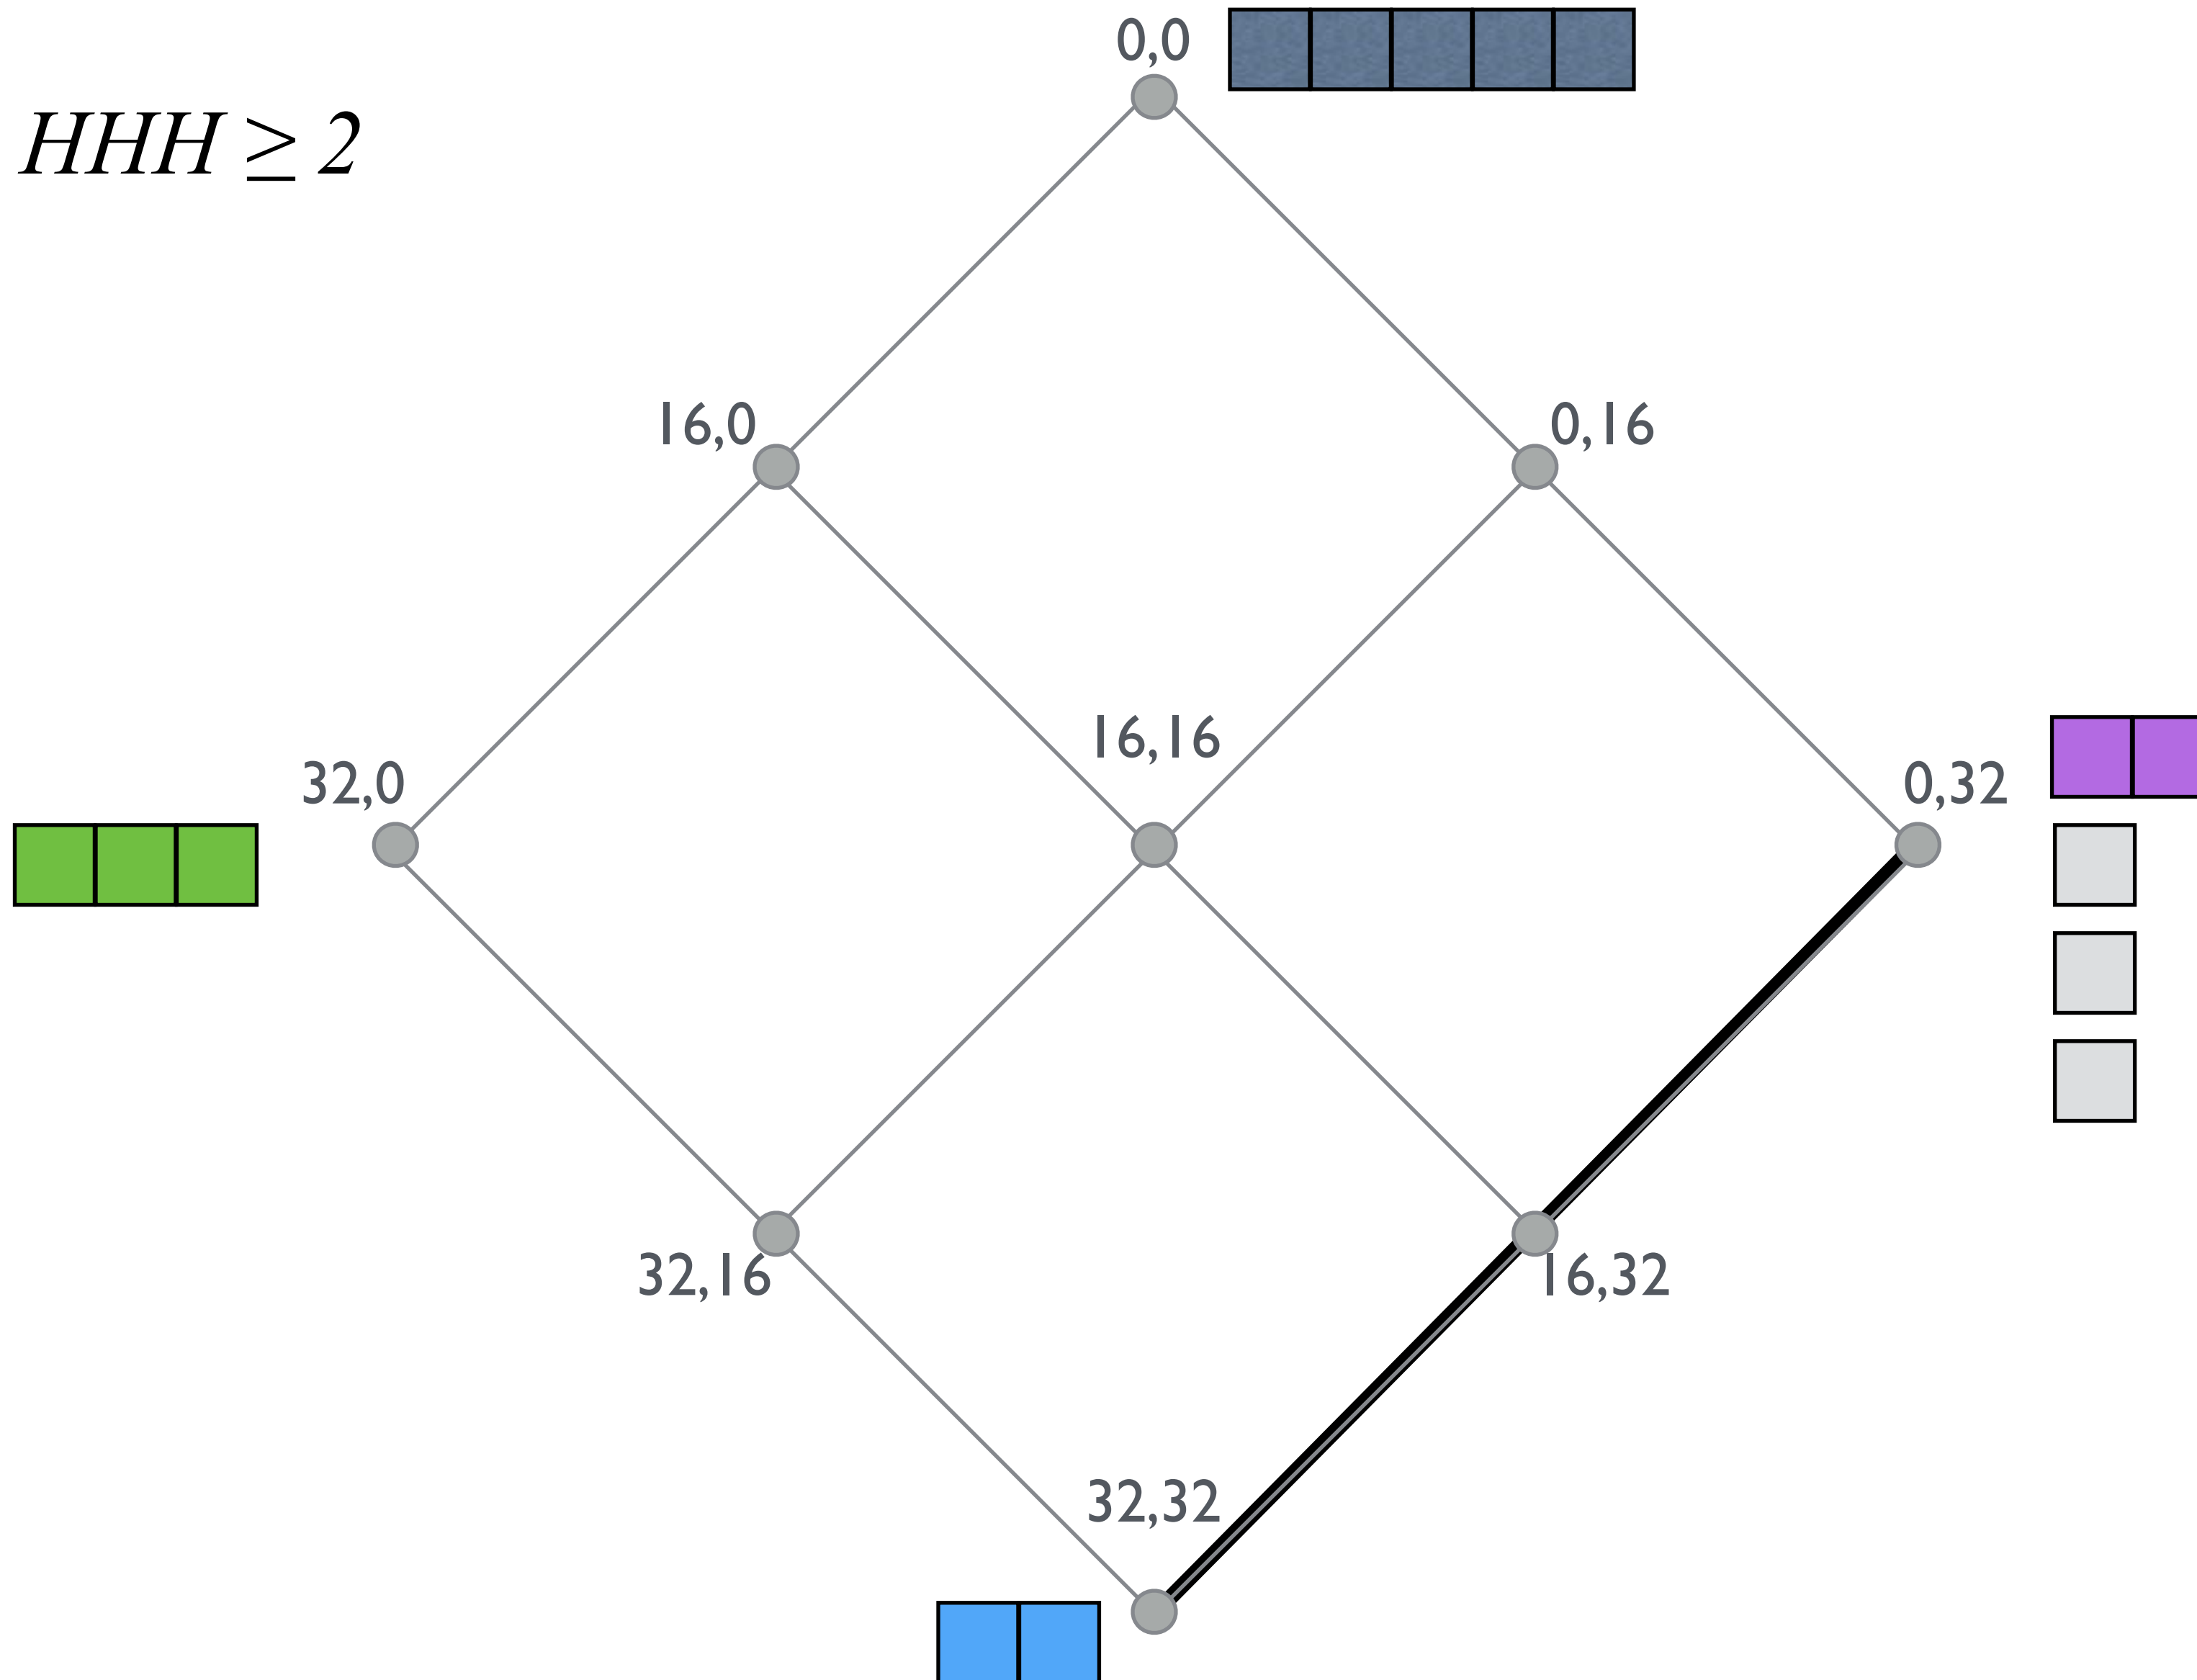
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



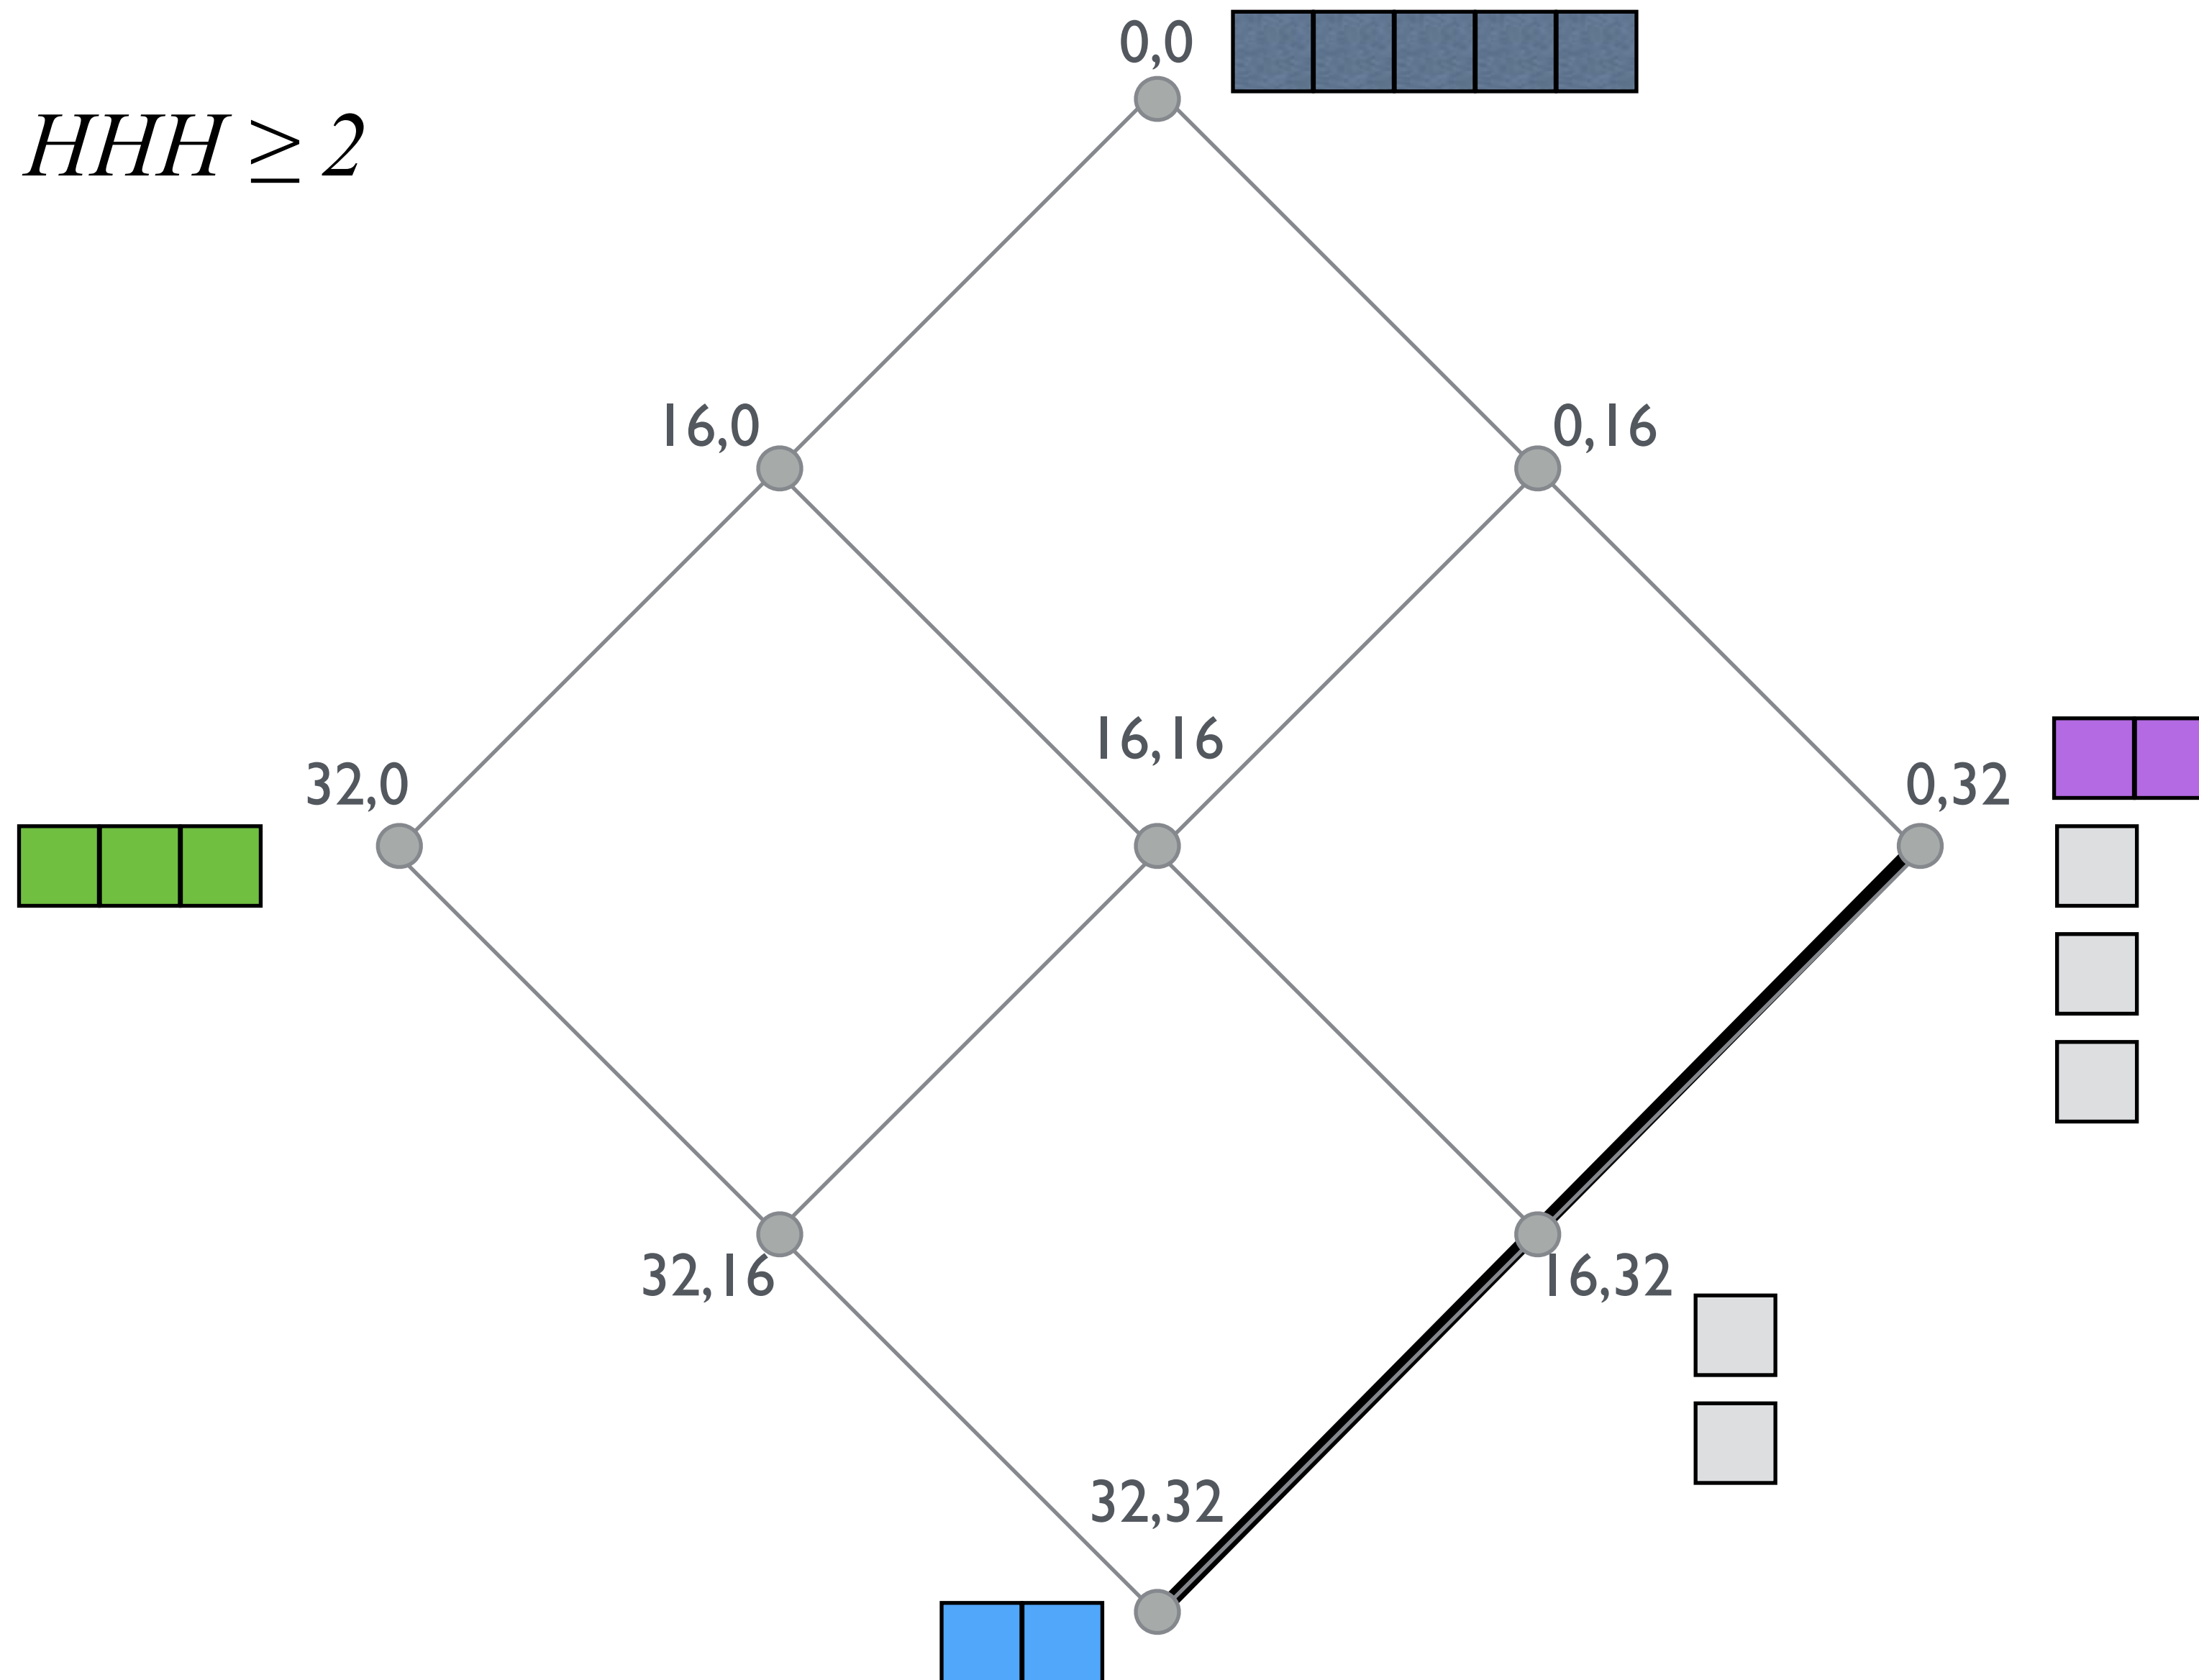
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



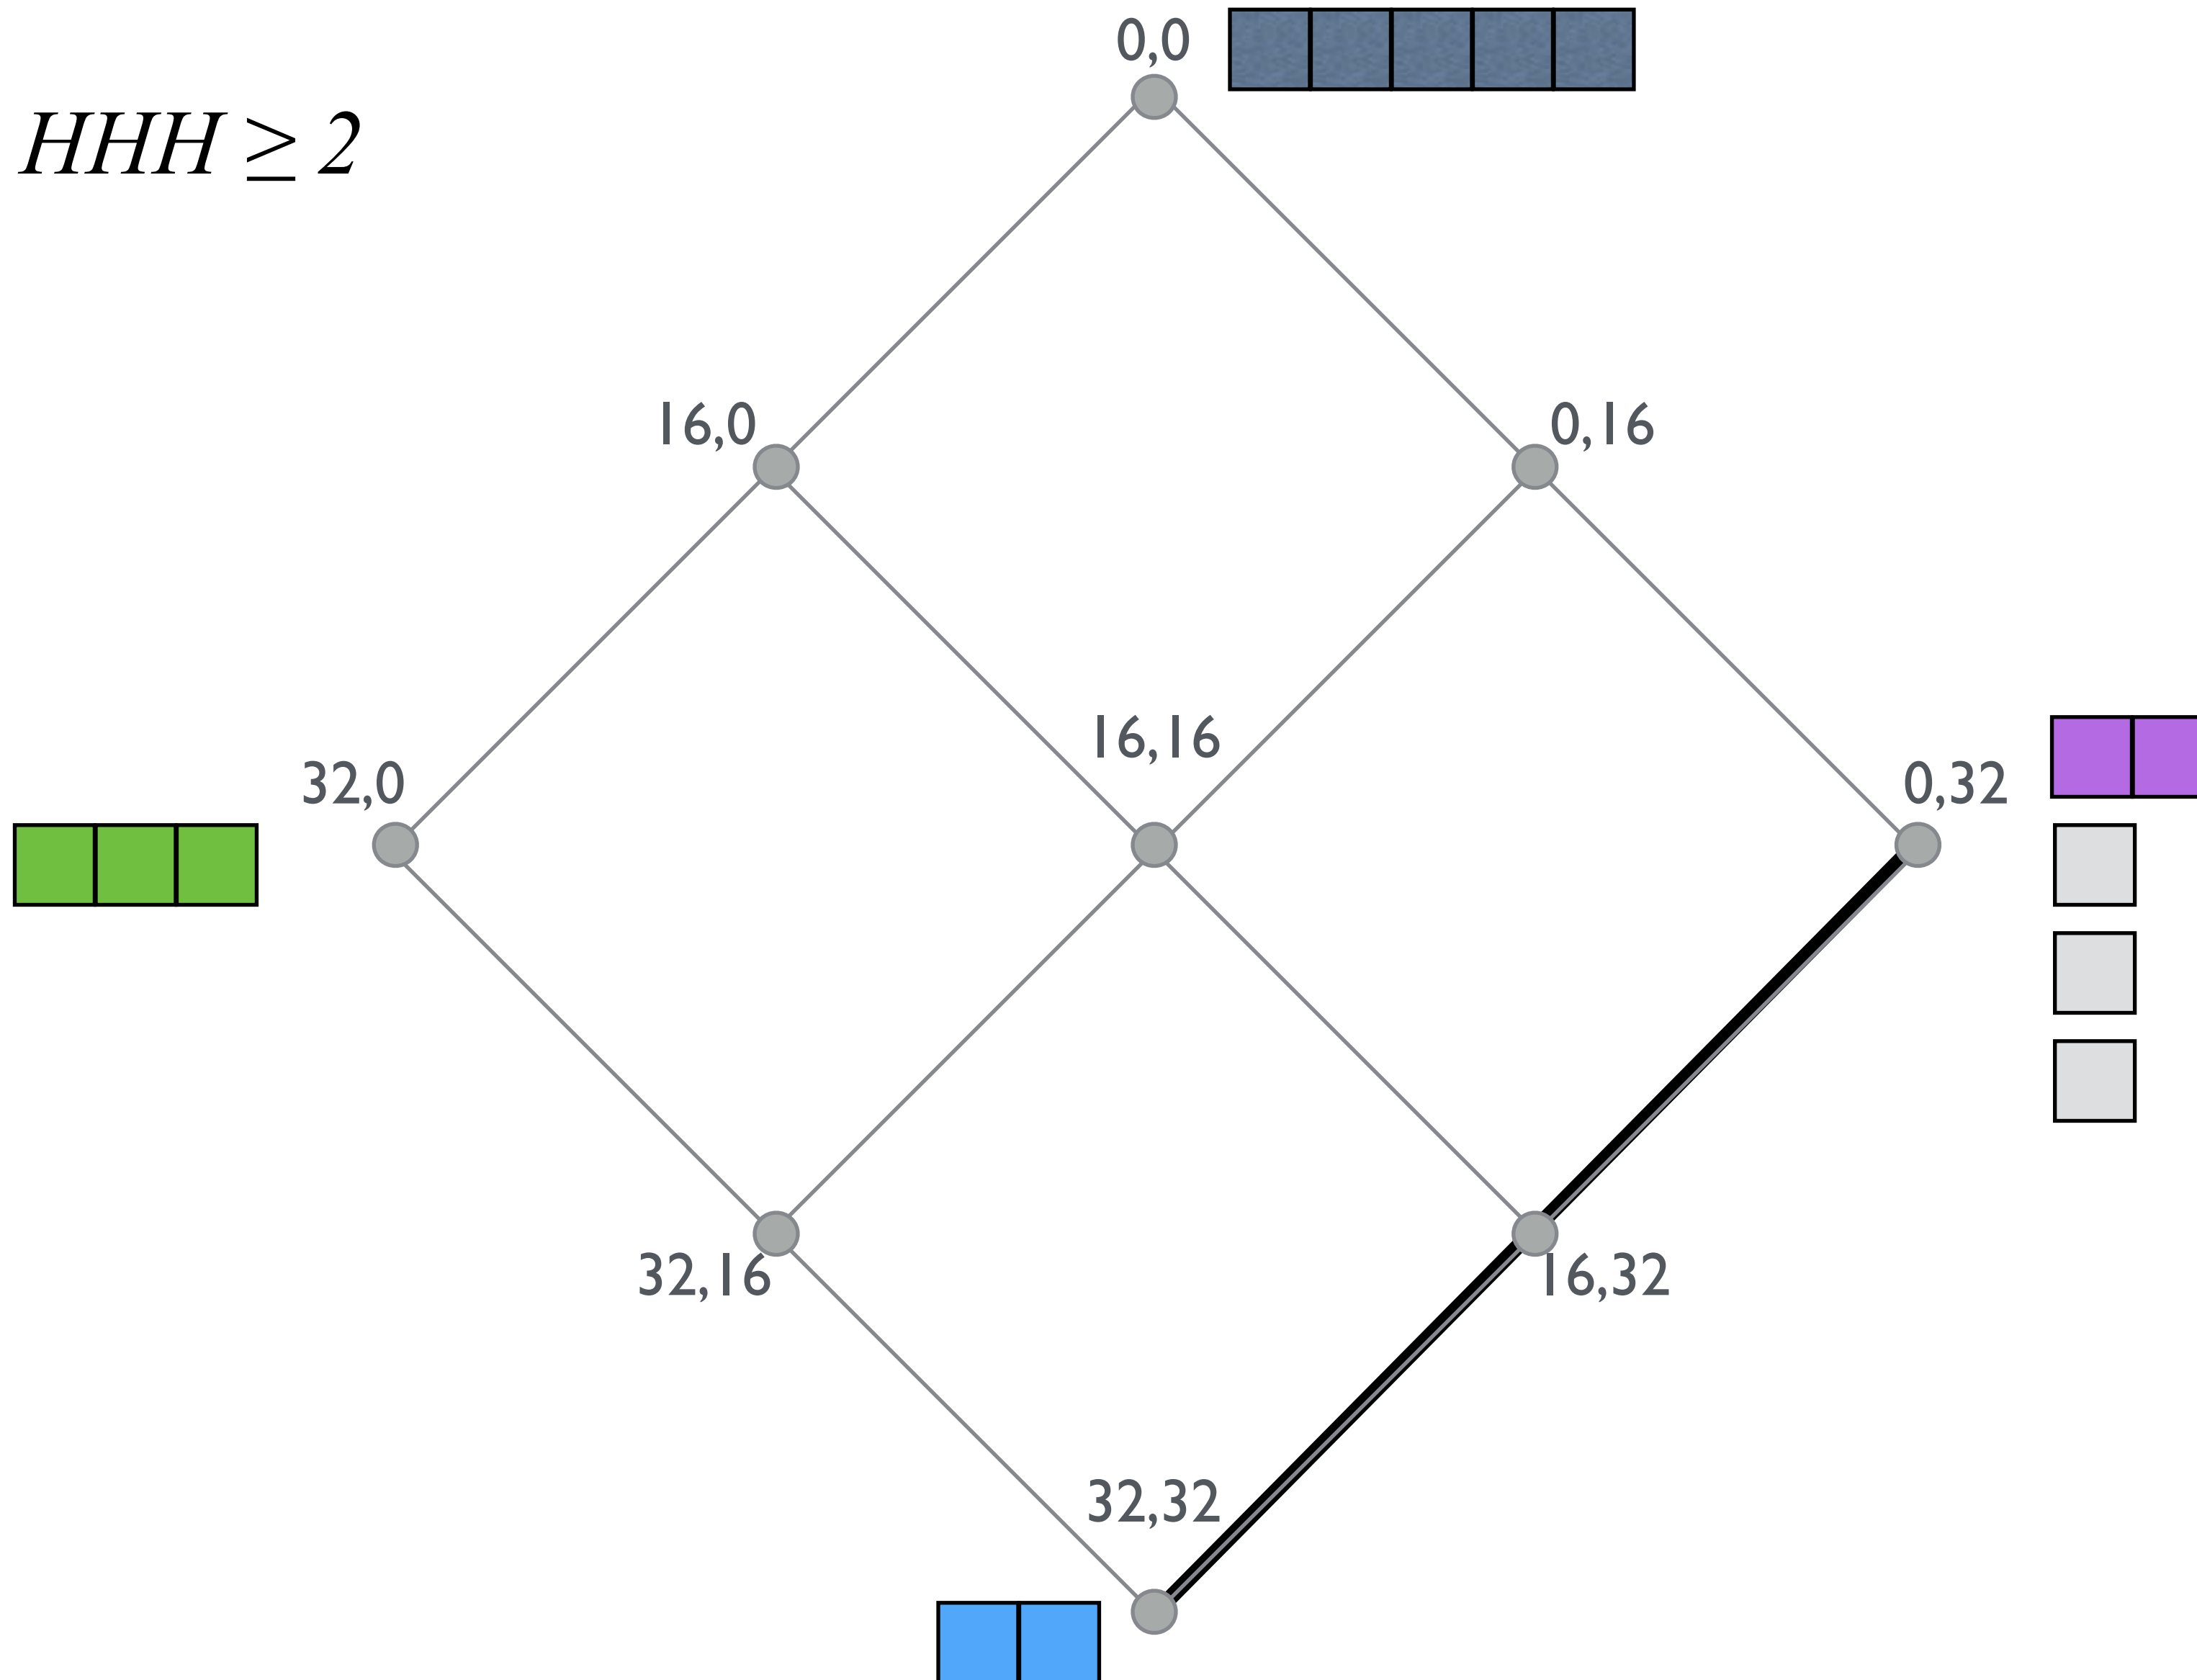
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



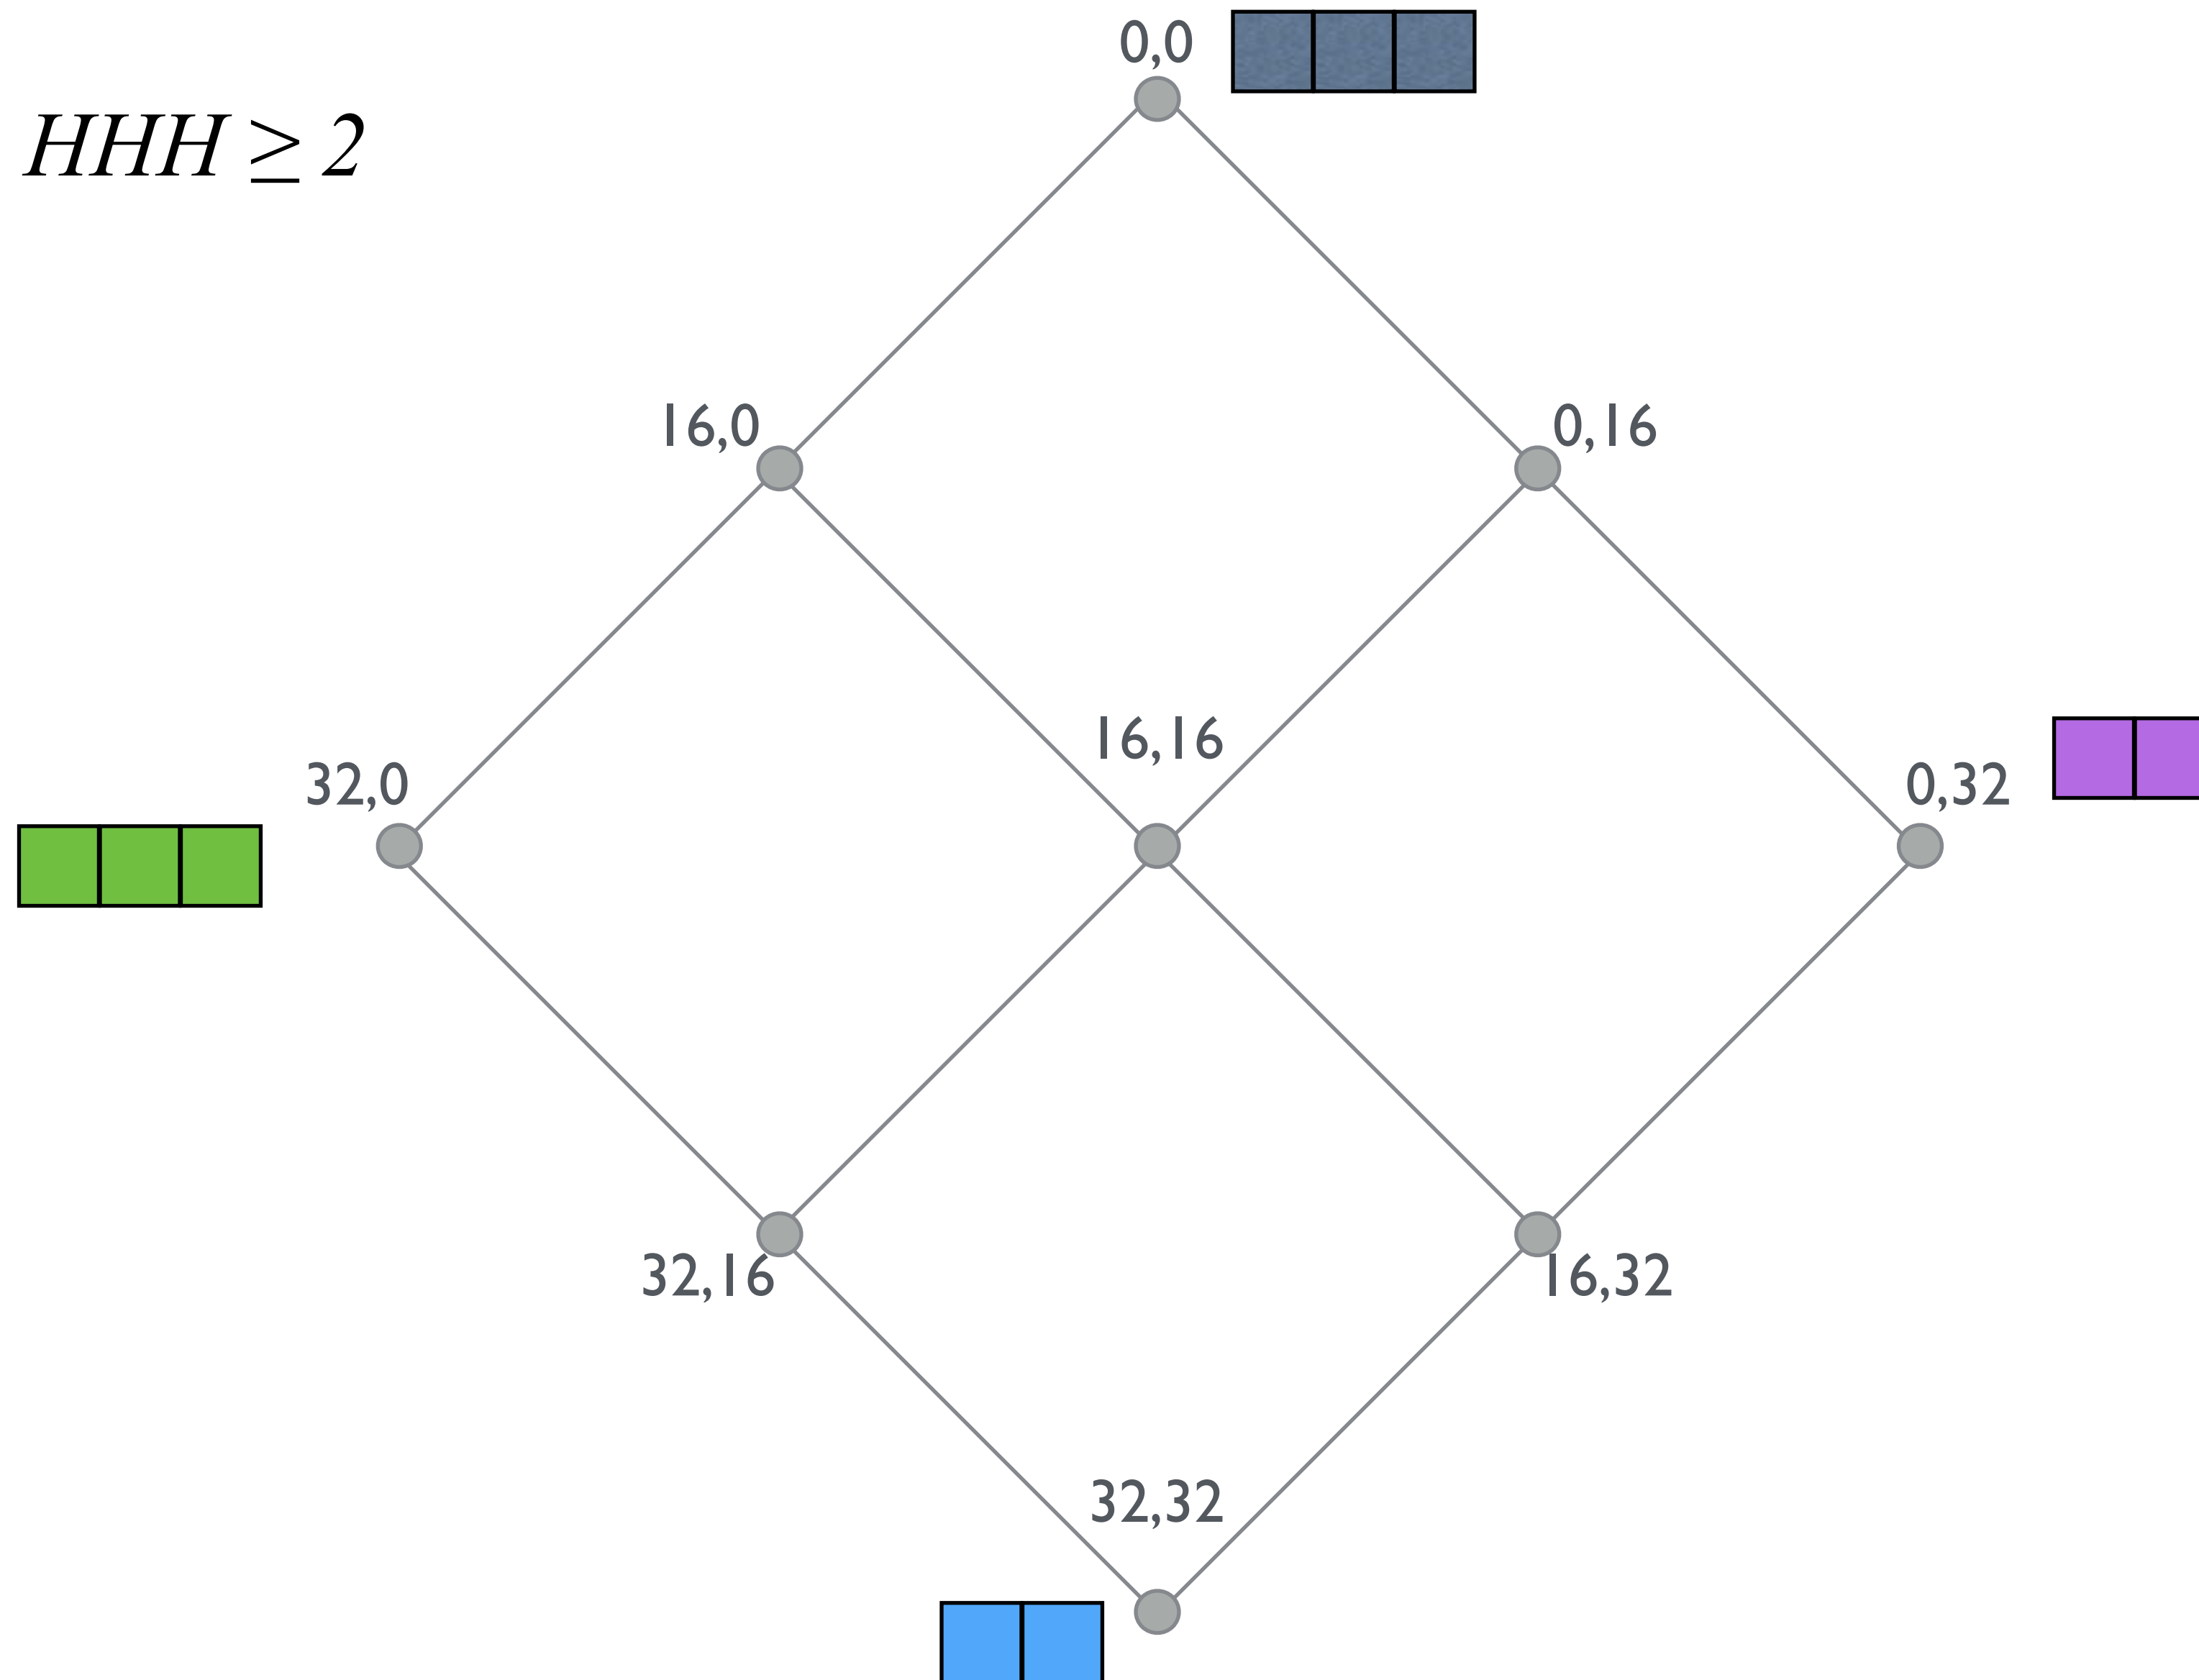
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



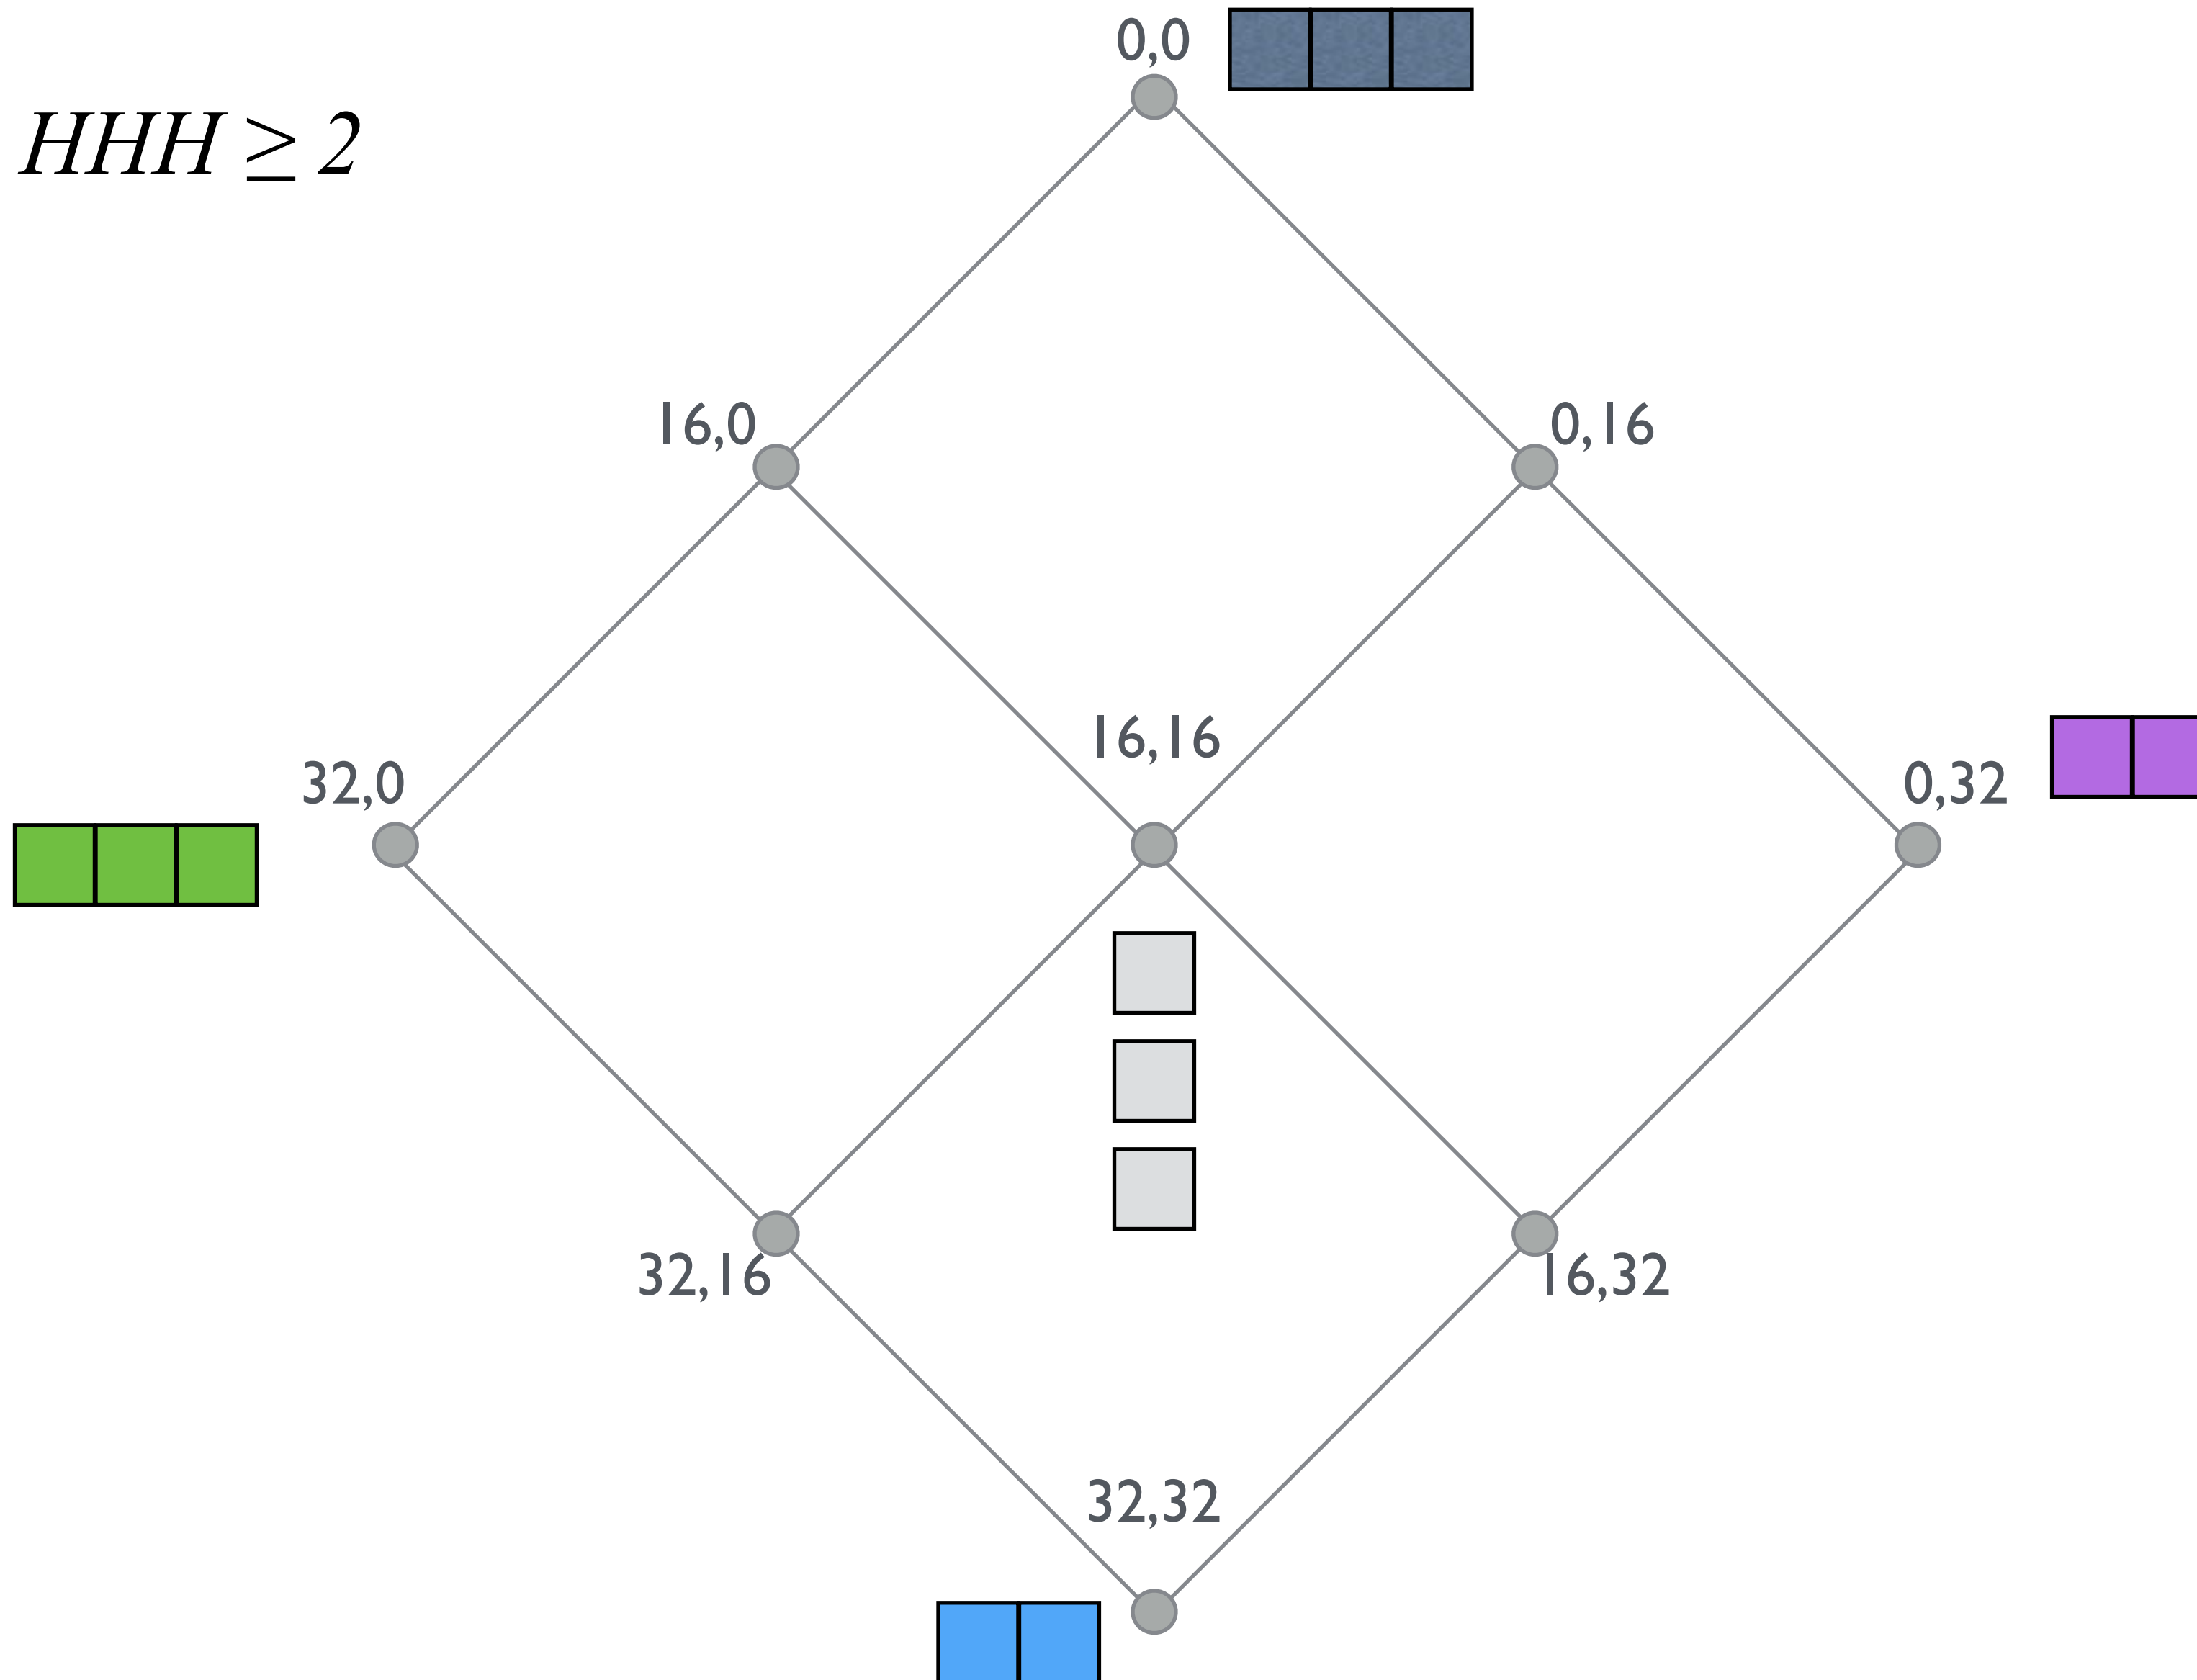
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



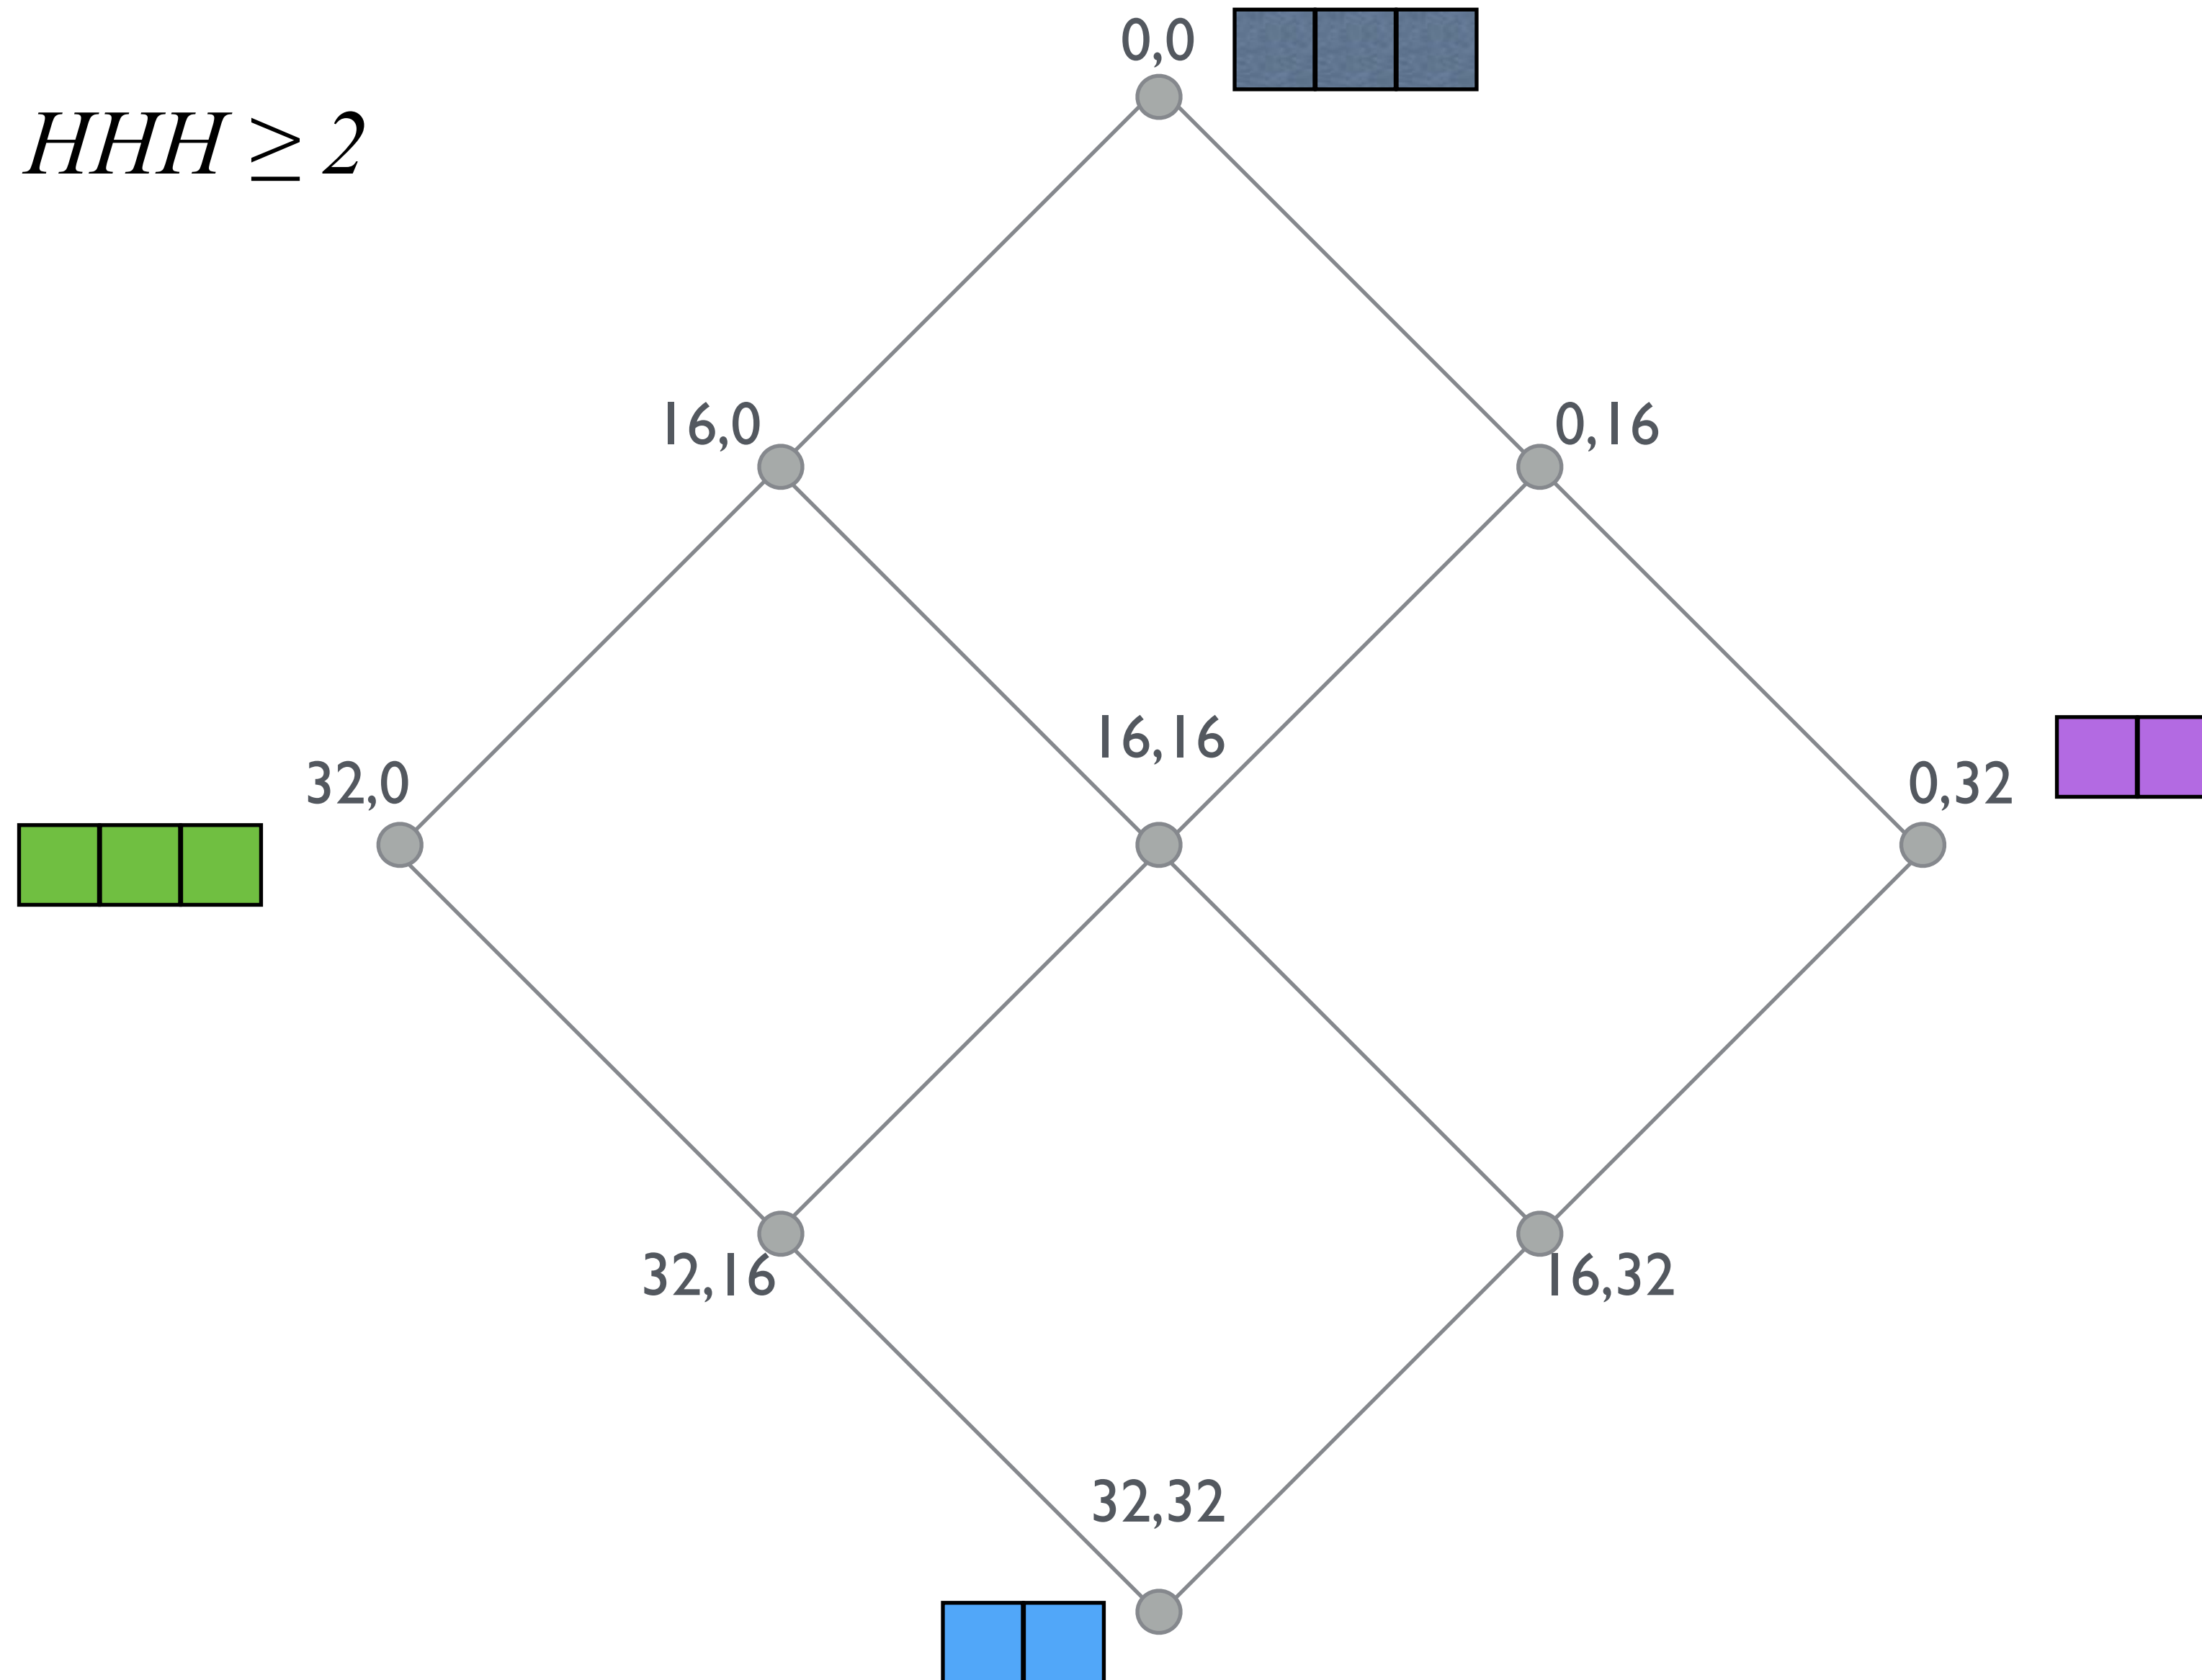
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



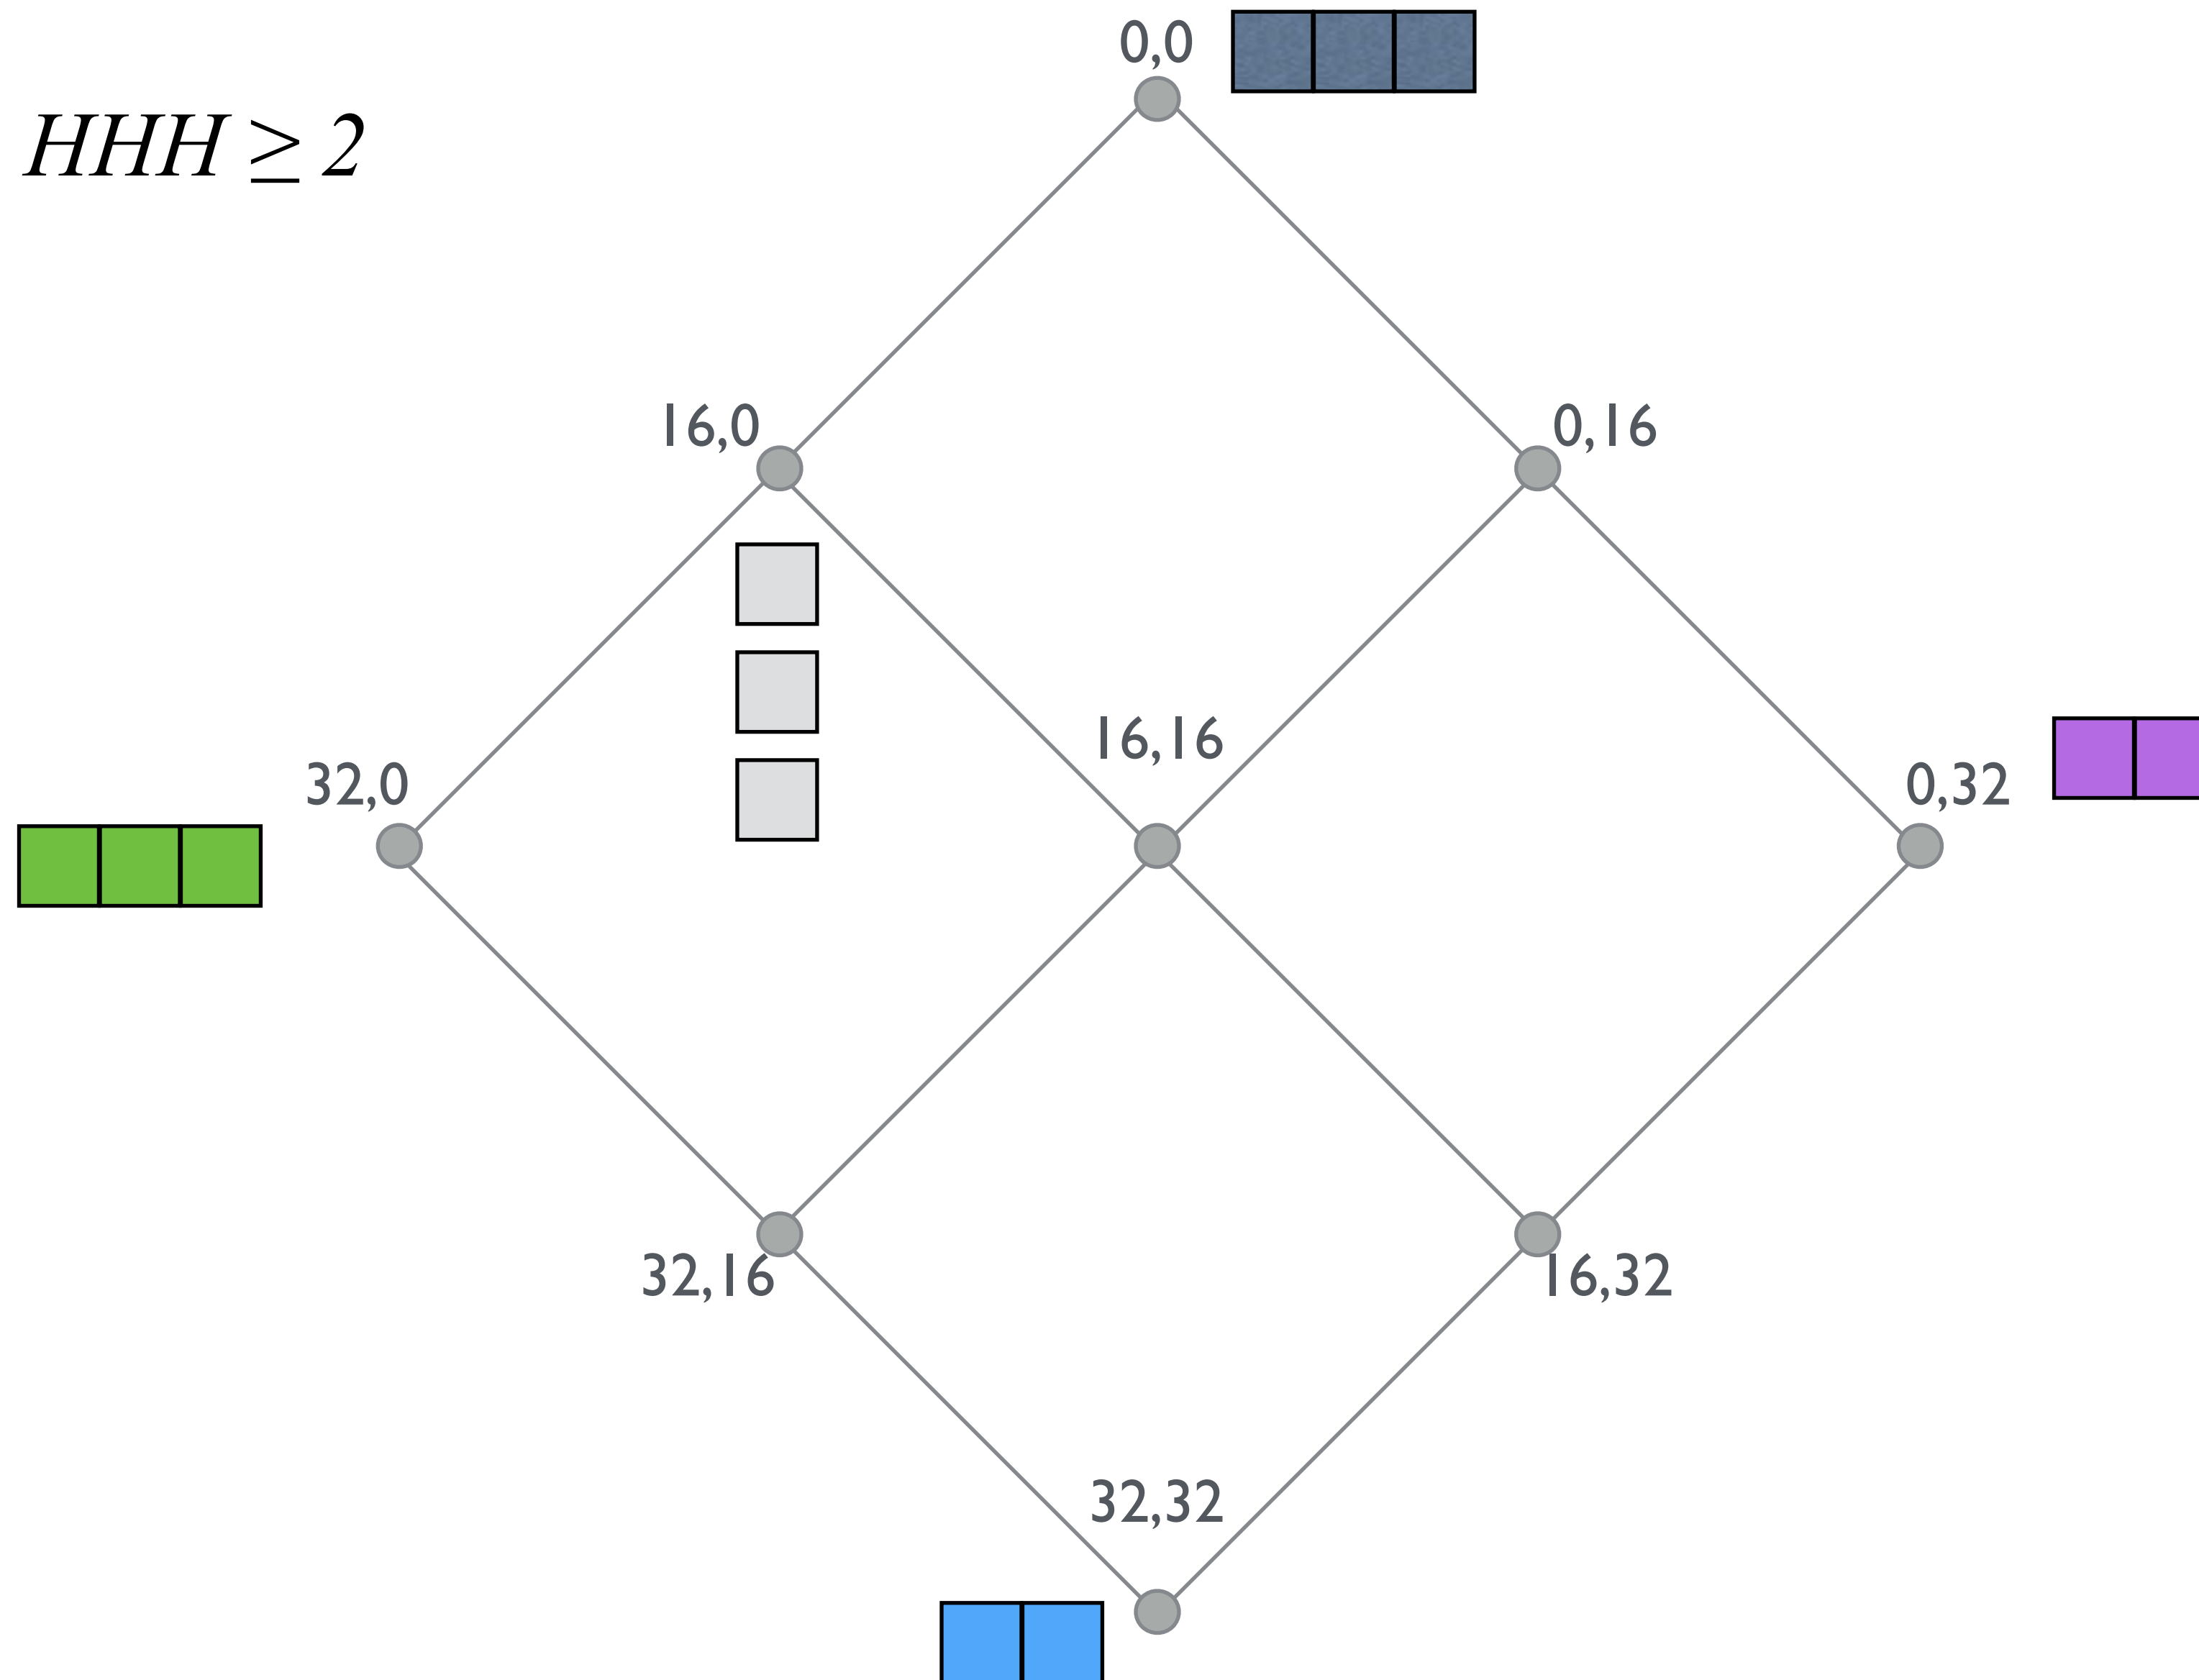
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



# RLS Illustrated

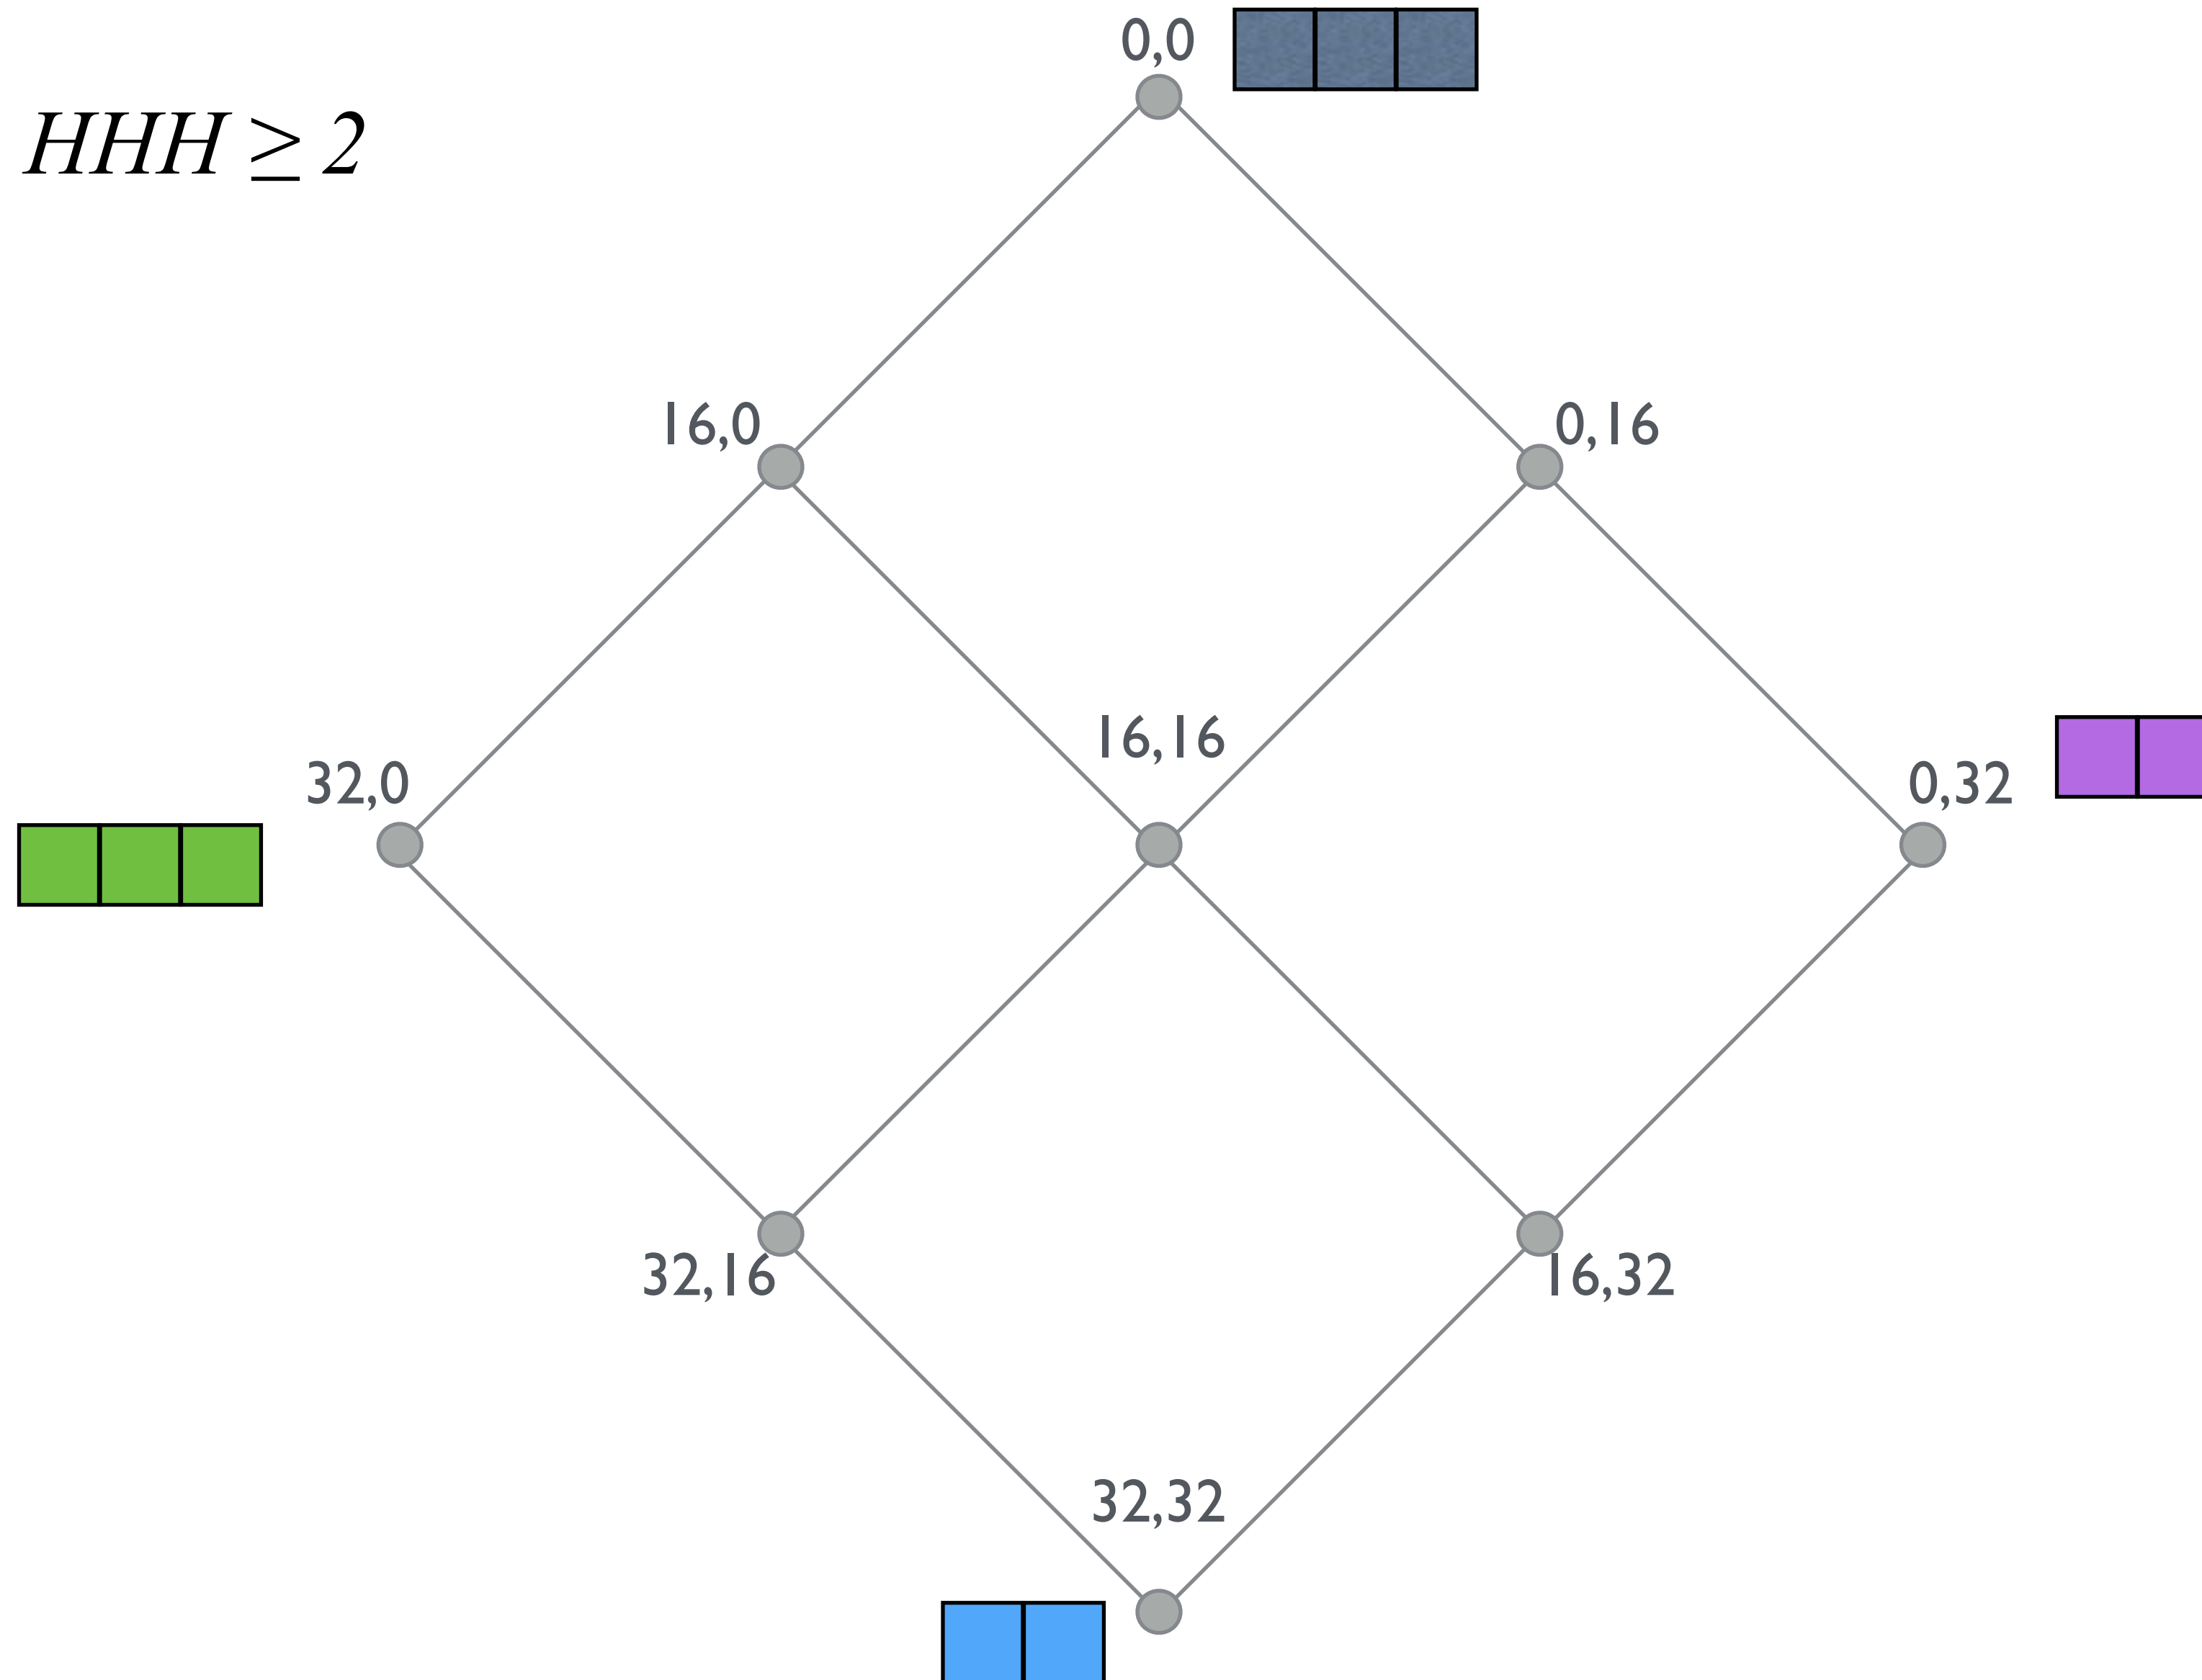
*10 inputs,  $HHH \geq 2$*





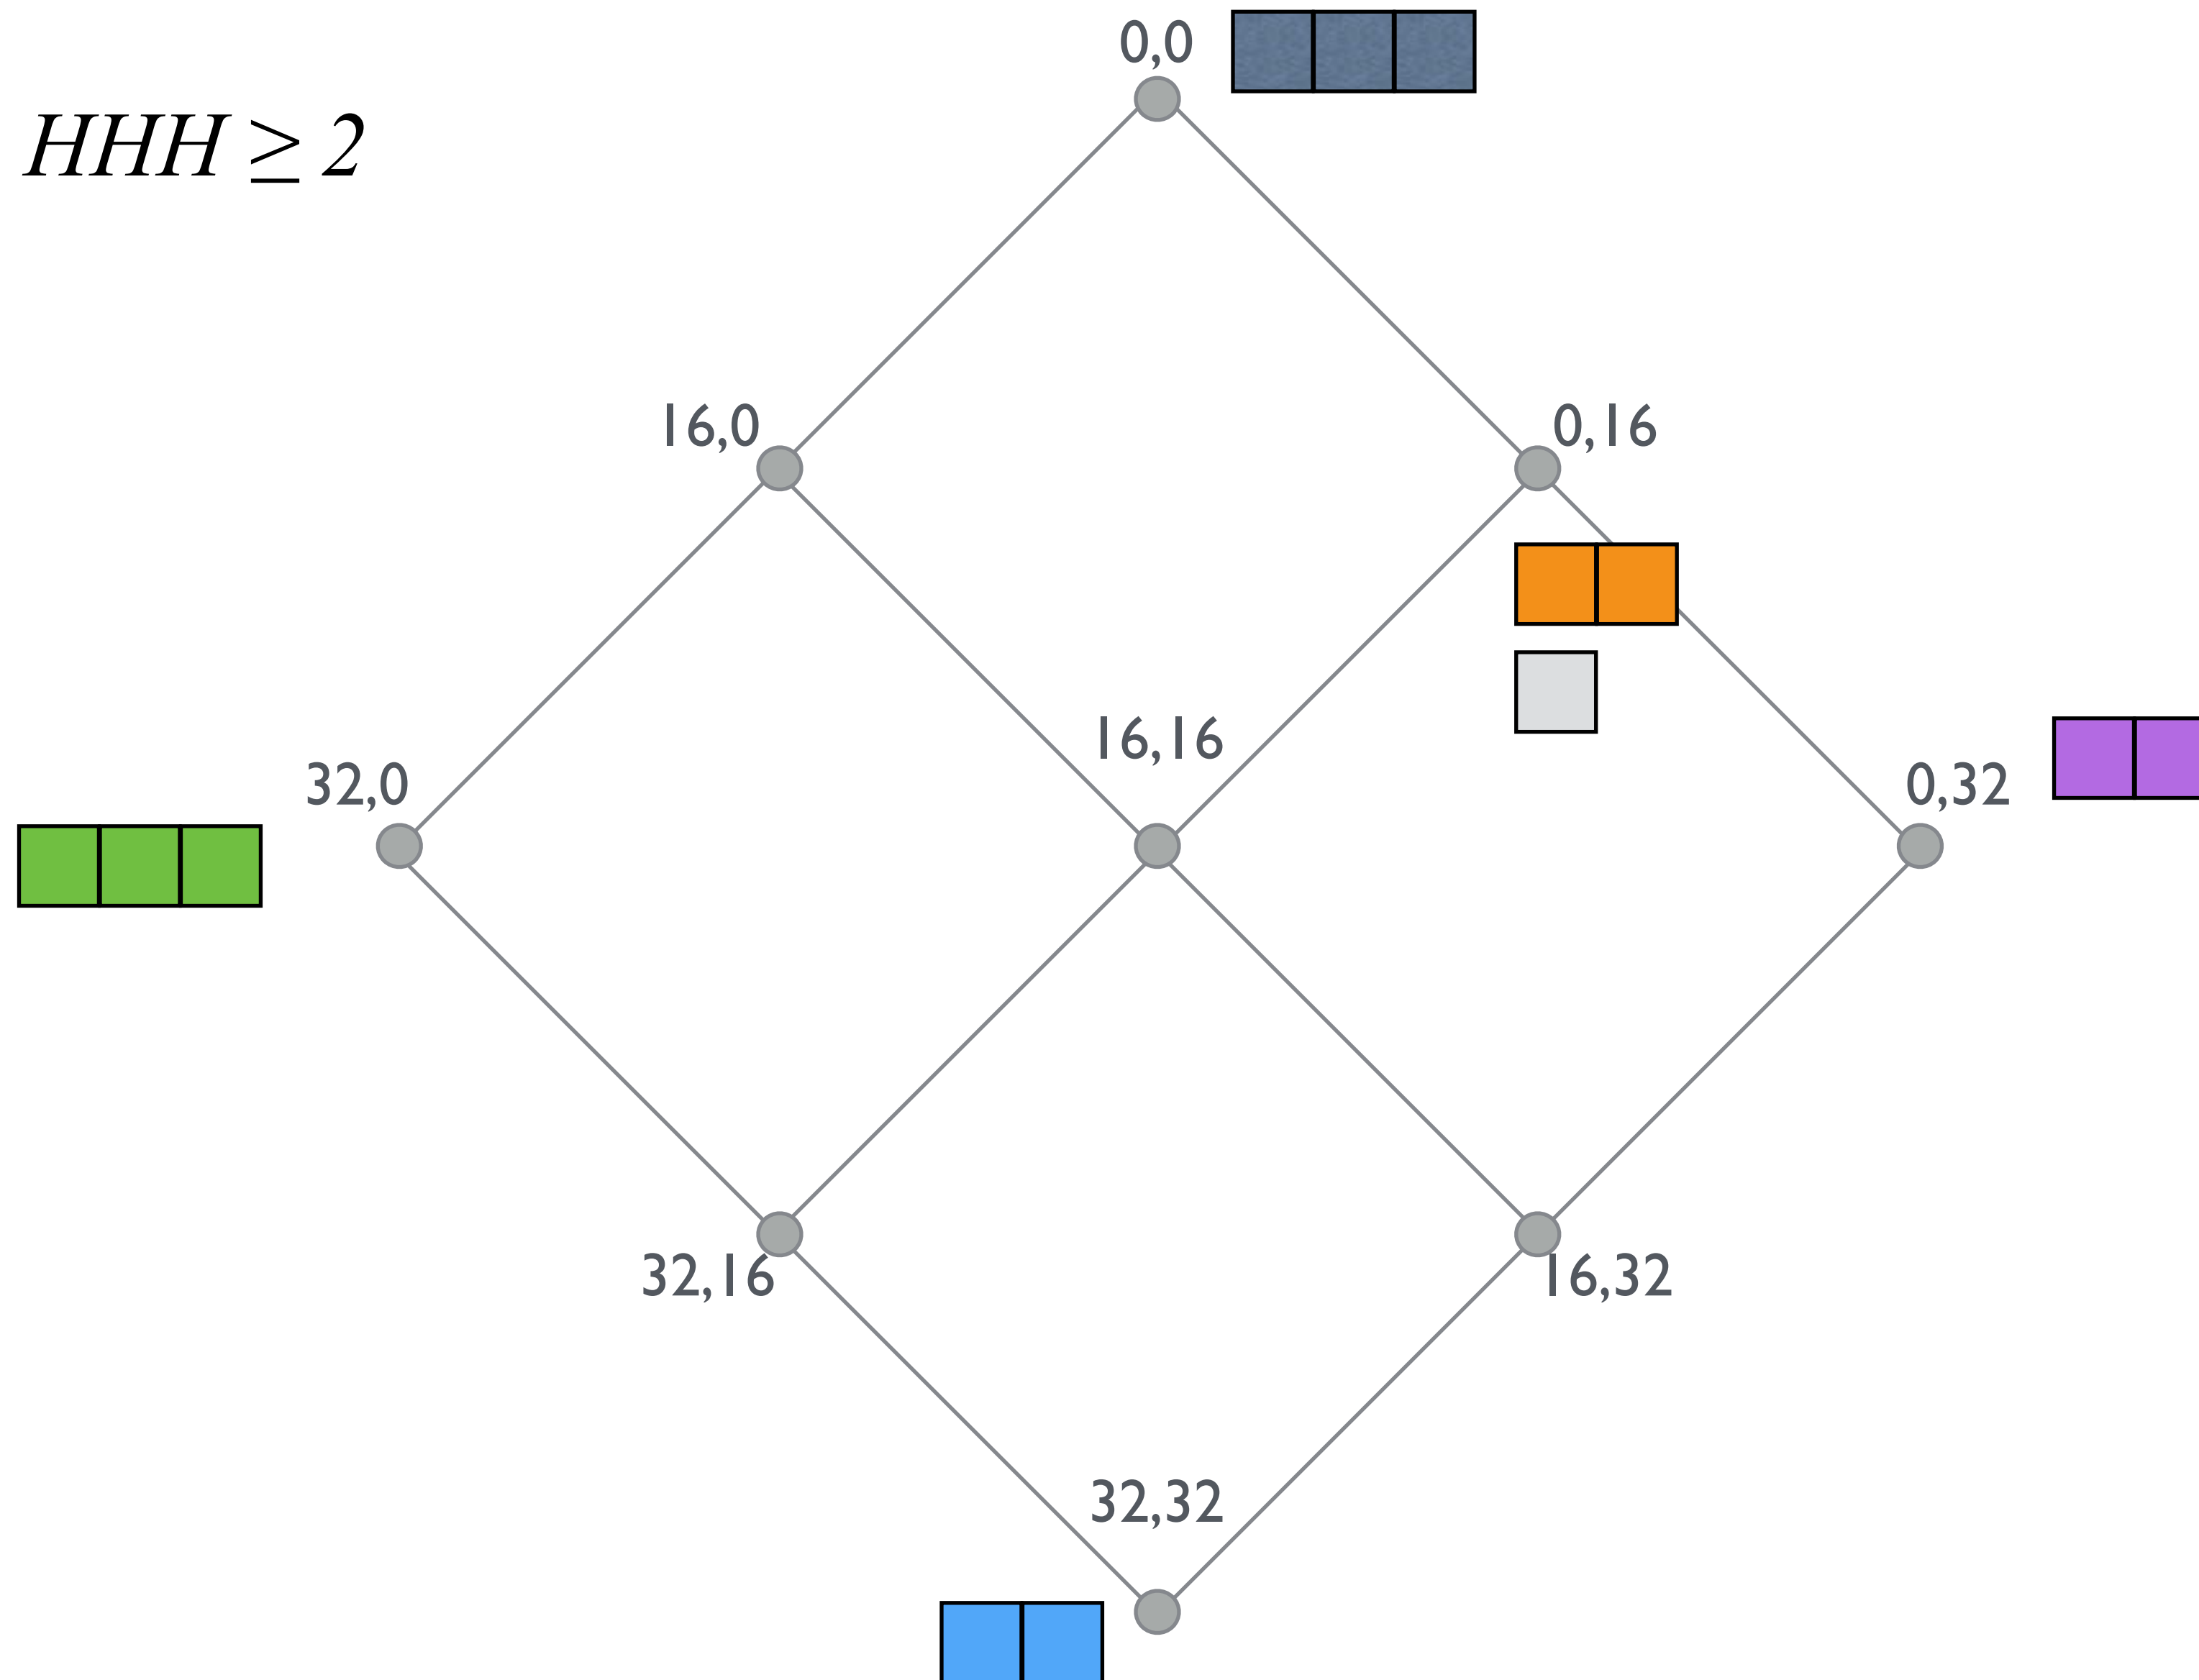
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



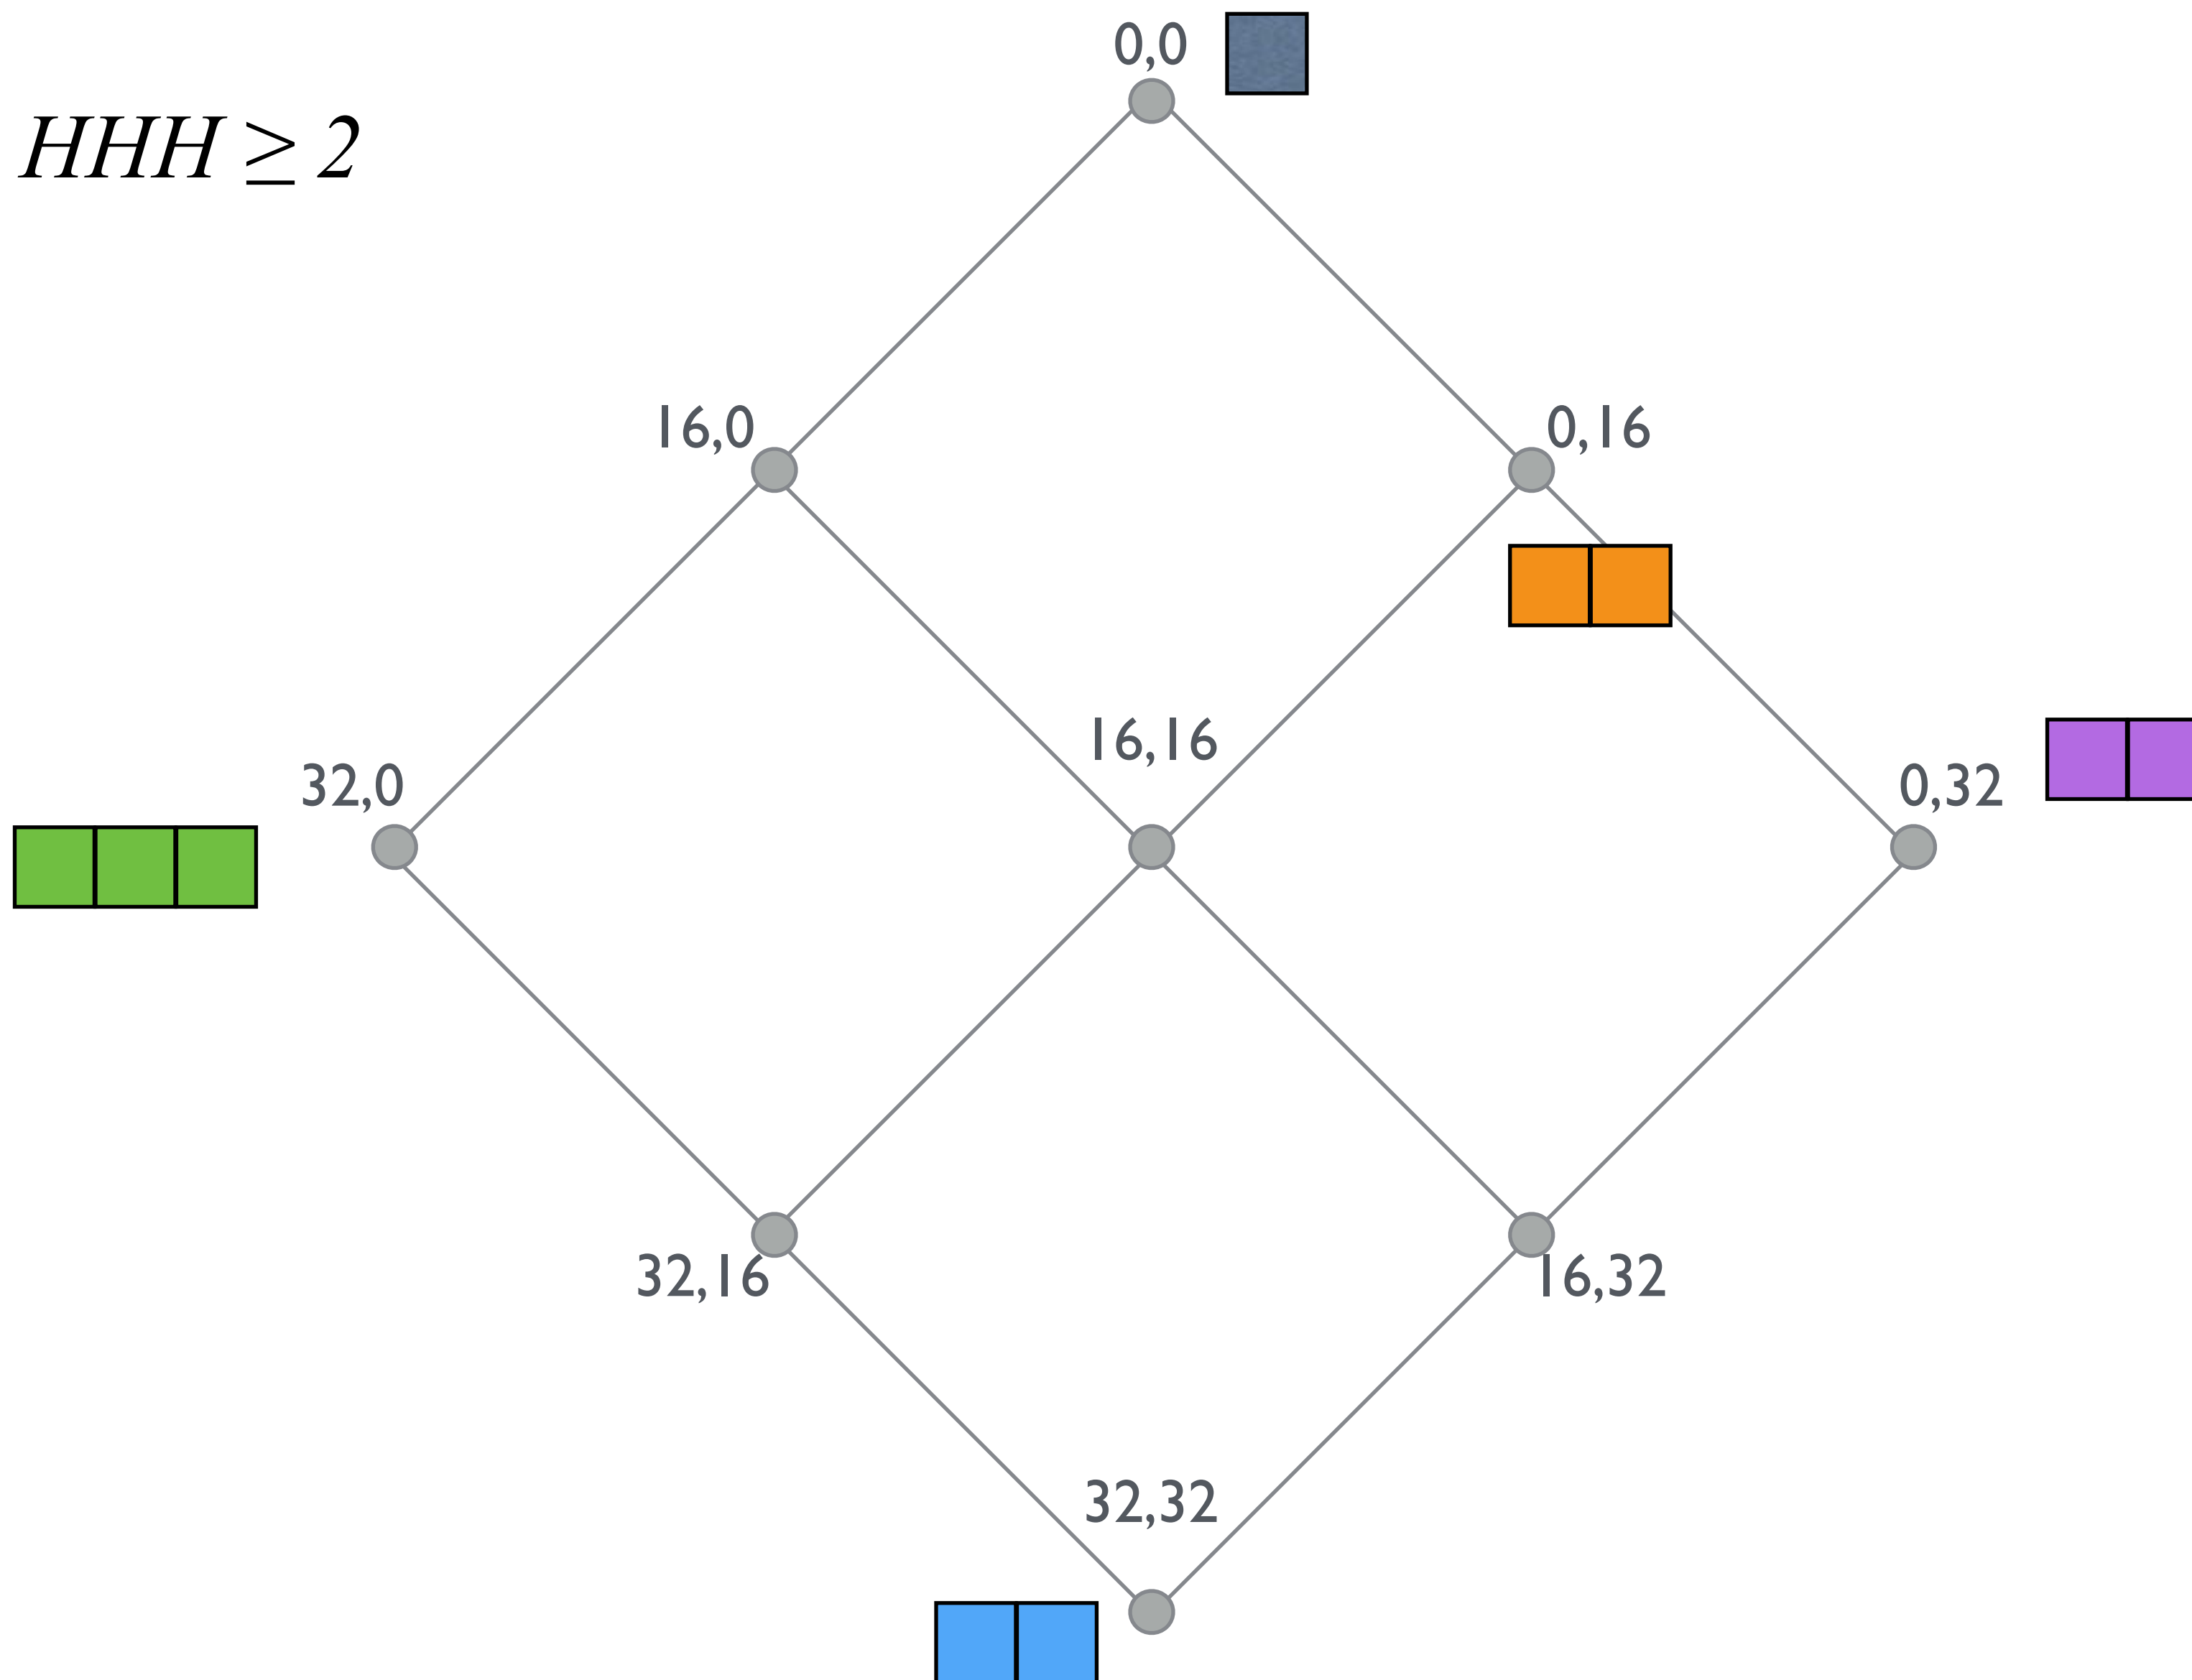
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



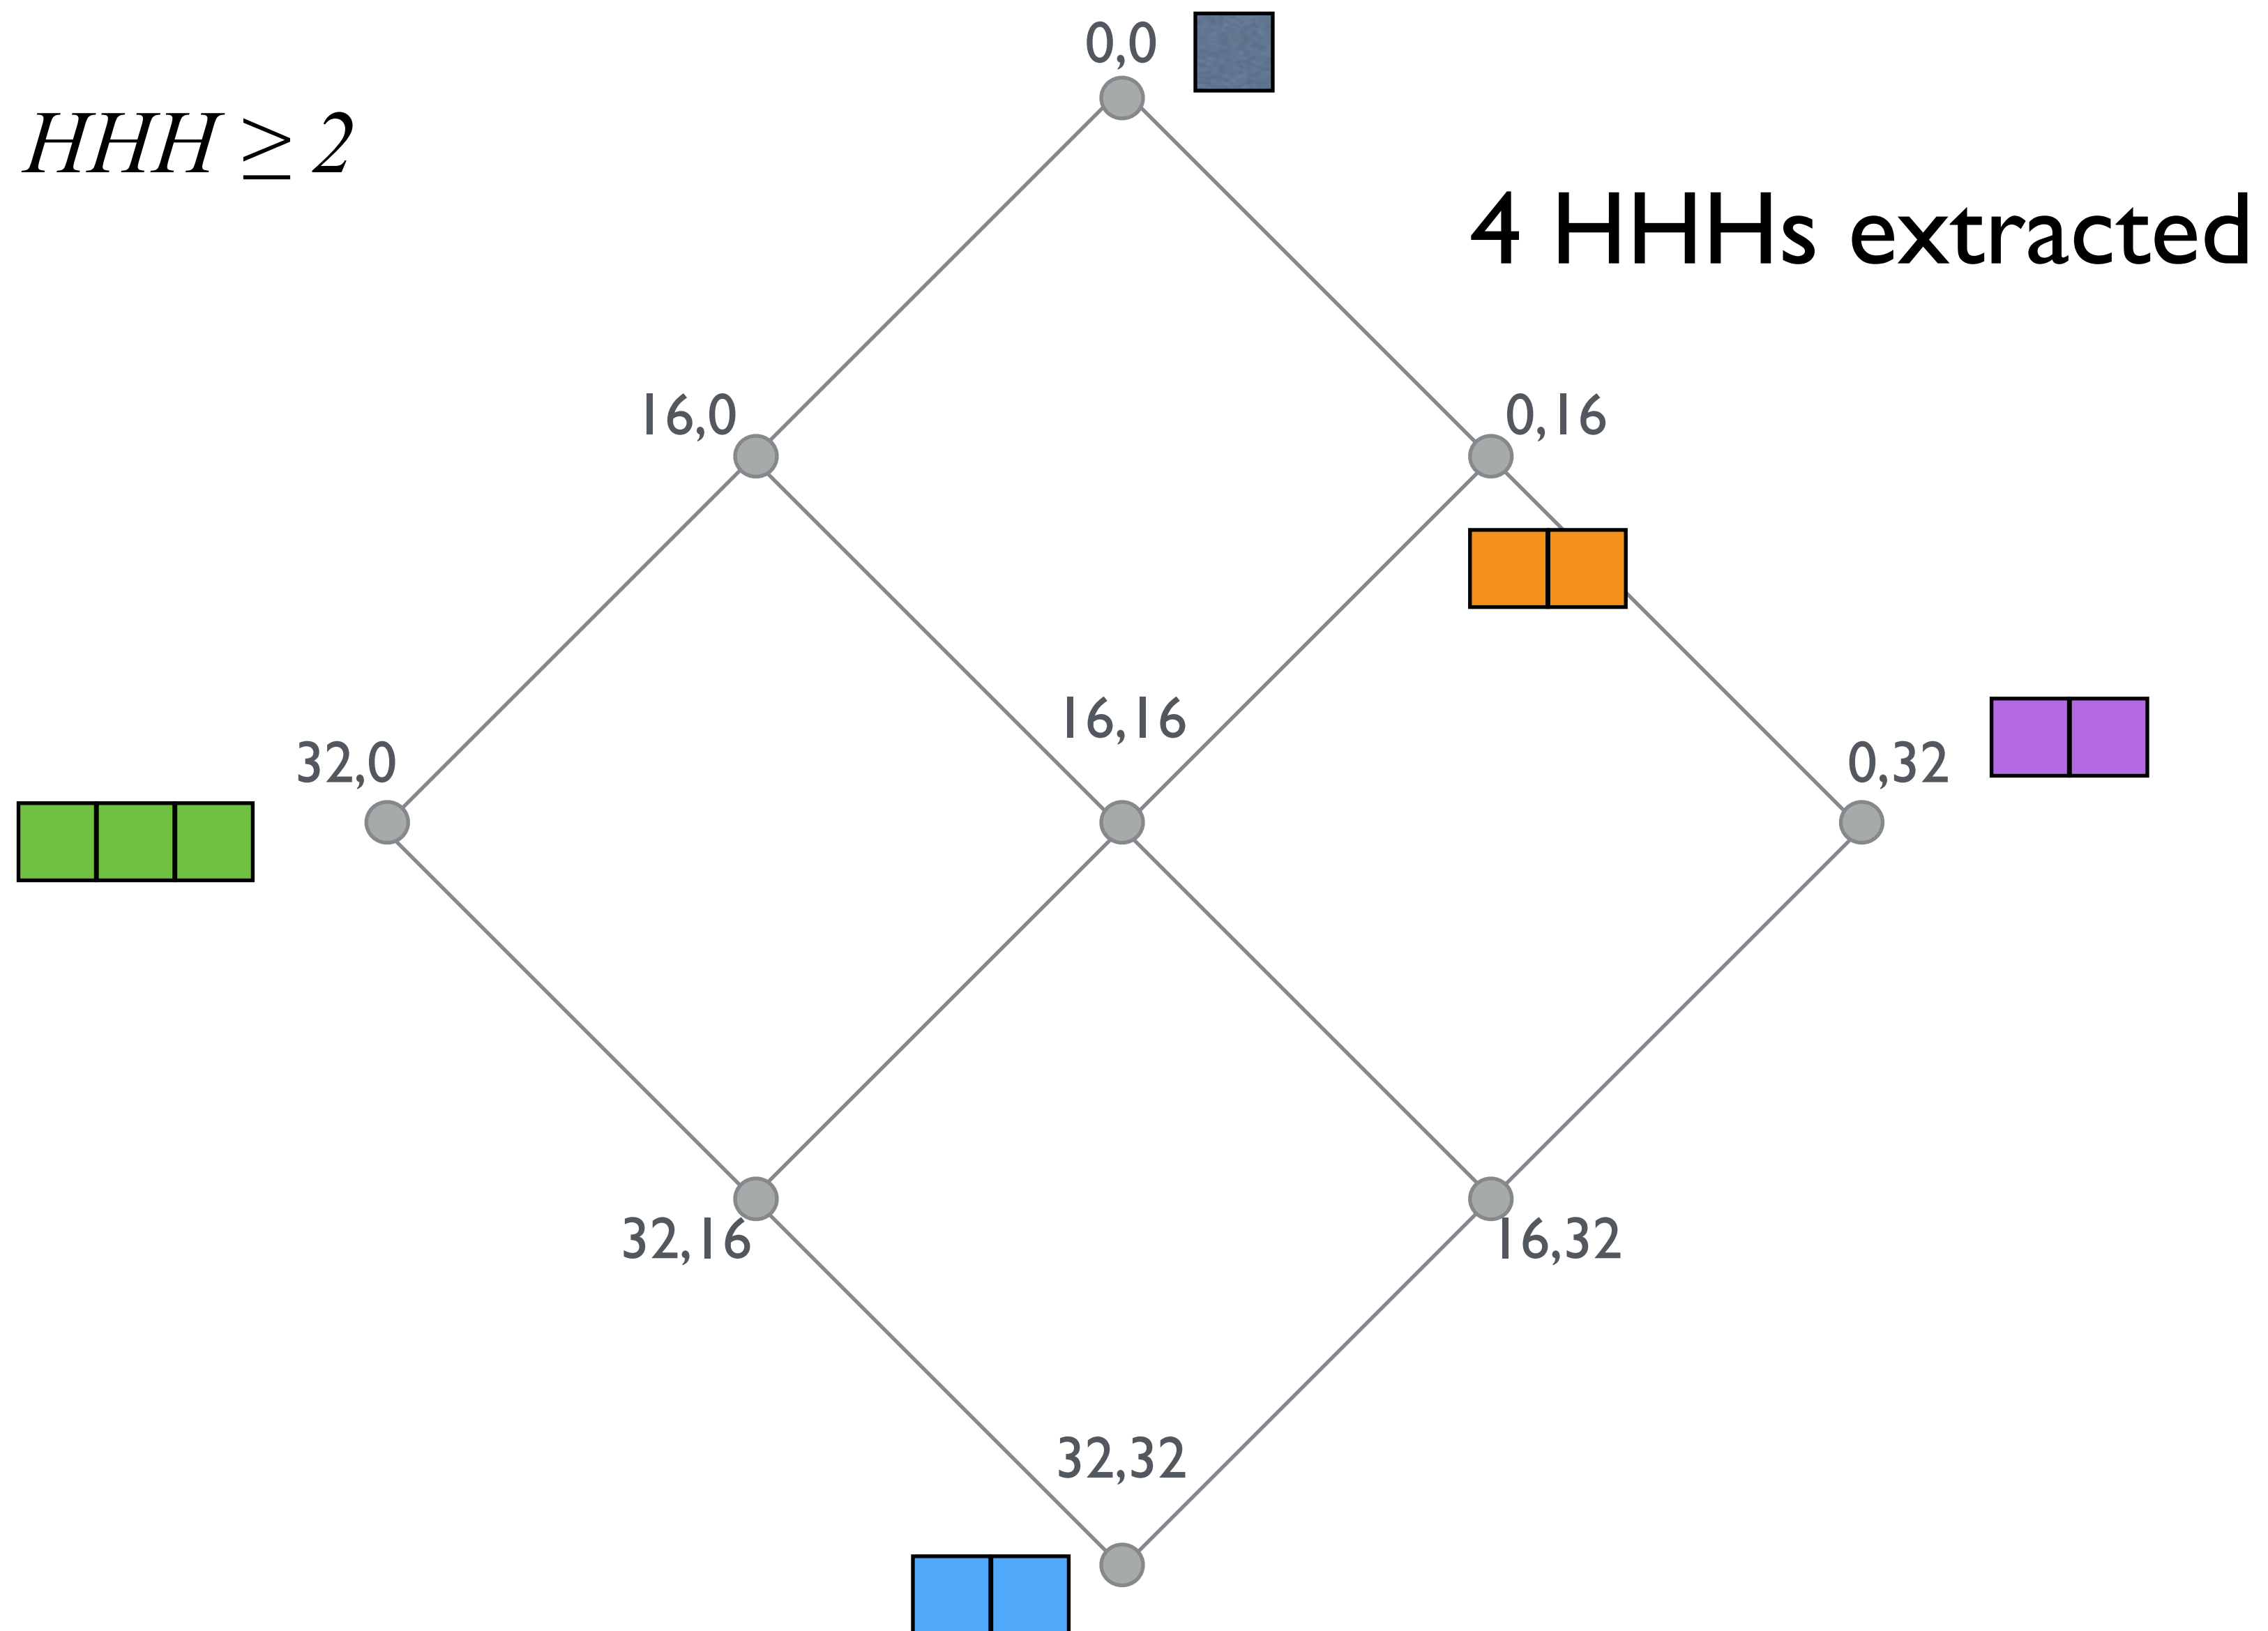
# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



# RLS Illustrated

*10 inputs,  $HHH \geq 2$*



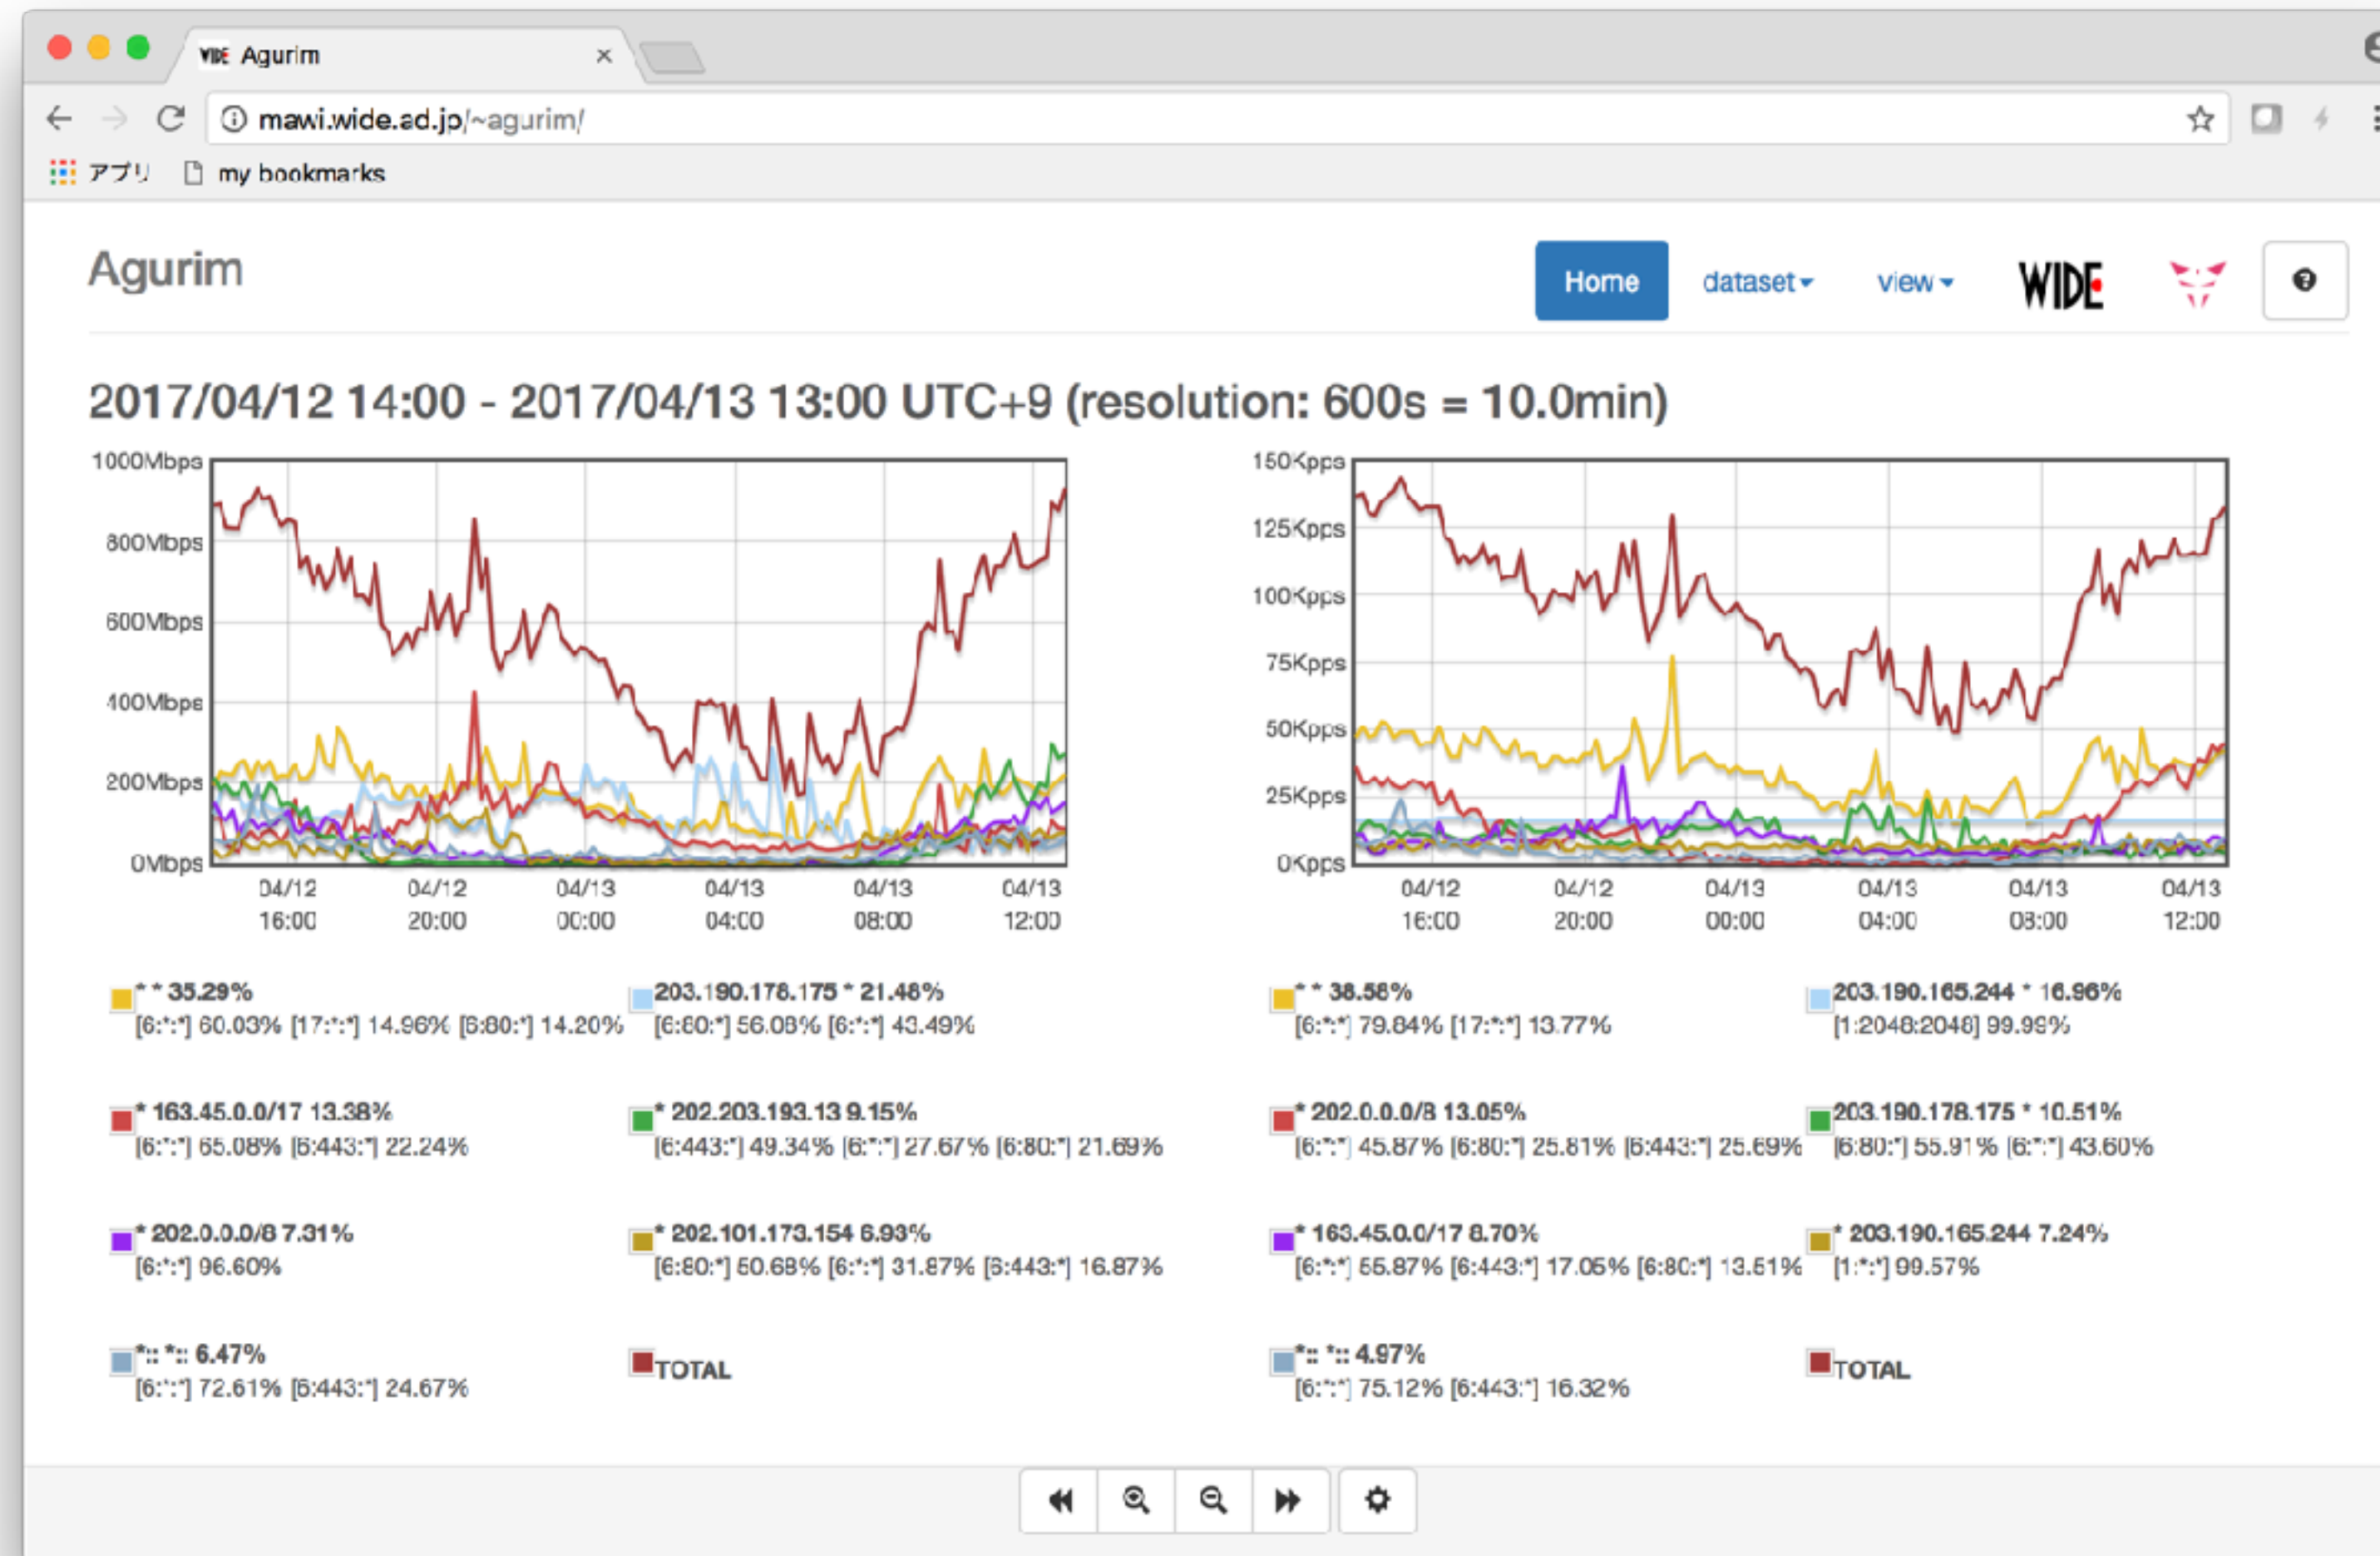
# Evaluation (in the paper)

- ordering bias: (src, dst) vs (dst, src) ➡ negligible
- comparison with Space-Saving: to illustrate differences
  - outputs ➡ much more compact
    - differences due to different definitions
  - speed ➡ 100 times faster for bitwise aggregation
    - but requires more memory (as a non-streaming algo)

# Implementations: RLS in agurim

- agurim: open-source tool
- 2-level HHH
  - main-attribute (src-dst adds), sub-attribute (ports)
- protocol specific heuristics
  - change depth of recursions by protocol knowledge to meet operational needs
- online processing by exploiting multi-core CPU

# agurim Web UI



<http://mawi.wide.ad.jp/~agurim/>

# Summary

- Recursive Lattice Search algorithm for HHH
  - revisit the definition of HHH, apply Z-ordering
  - propose an efficient HHH algorithm
- open-source tool and open datasets from 2013

<http://mawi.wide.ad.jp/~agurim/about.html>



# evaluation in detail

- simulation: code from SpaceSaving [Mitzenmacher2012]
  - quick hack to port agurim's RLS
  - input: a mawi packet trace from 2016-10-20
- order sensitivity: (src,dst) vs. (dst,src)
  - very similar outputs: not sensitive to the order
- comparing with SS (streaming algorithm, overlap rollup)
  - different definitions: just to illustrate major differences
  - outputs: comparable, except nodes in upper lattice
  - performance: 100x faster for bit-wise aggregation!

# order sensitivity

## (src,dst) vs. (dst,src)

region	no	aggregated by (src,dst)		$c'/N(\%)$
		src	dst	
VI	(1)	112.31.100.1/32	163.229.97.230/32	16.5
V	(2)	64.0.0.0/2	202.203.3.13/32	5.2
	(3)	128.0.0.0/1	202.203.3.13/32	5.8
	(4)	*	202.26.162.46/32	6.0
III	(5)	163.229.96.0/23	*	5.0
	(6)	203.179.128.0/20	*	6.8
II	(7)	*	202.203.3.0/24	5.9
	(8)	*	203.179.140.0/23	5.7
	(9)	*	163.229.128.0/17	5.1
I	(10)	0.0.0.0/1	202.192.0.0/12	5.3
	(11)	202.192.0.0/12	*	6.7
	(12)	*	202.0.0.0/7	7.6
	(13)	<b>128.0.0.0/4</b>	*	<b>5.0</b>
	(14)	<b>128.0.0.0/2</b>	*	<b>6.0</b>
	(15)	*	<b>128.0.0.0/2</b>	<b>5.4</b>
-		*	*	<b>2.0</b>
100.0				
aggregated by (dst,src)				
	(1)-(12)	identical to (src,dst)		
I	(13)	<b>128.0.0.0/2</b>	<b>0.0.0.0/2</b>	<b>5.7</b>
	(14)	*	<b>128.0.0.0/3</b>	<b>5.3</b>
	(15)	<b>128.0.0.0/1</b>	*	<b>6.4</b>
-		*	*	<b>1.0</b>

- (1)-(12): identical
- (13)-(15): minor difference
- not sensitive to src-dst order

# HHHs reported by RLS vs. SS

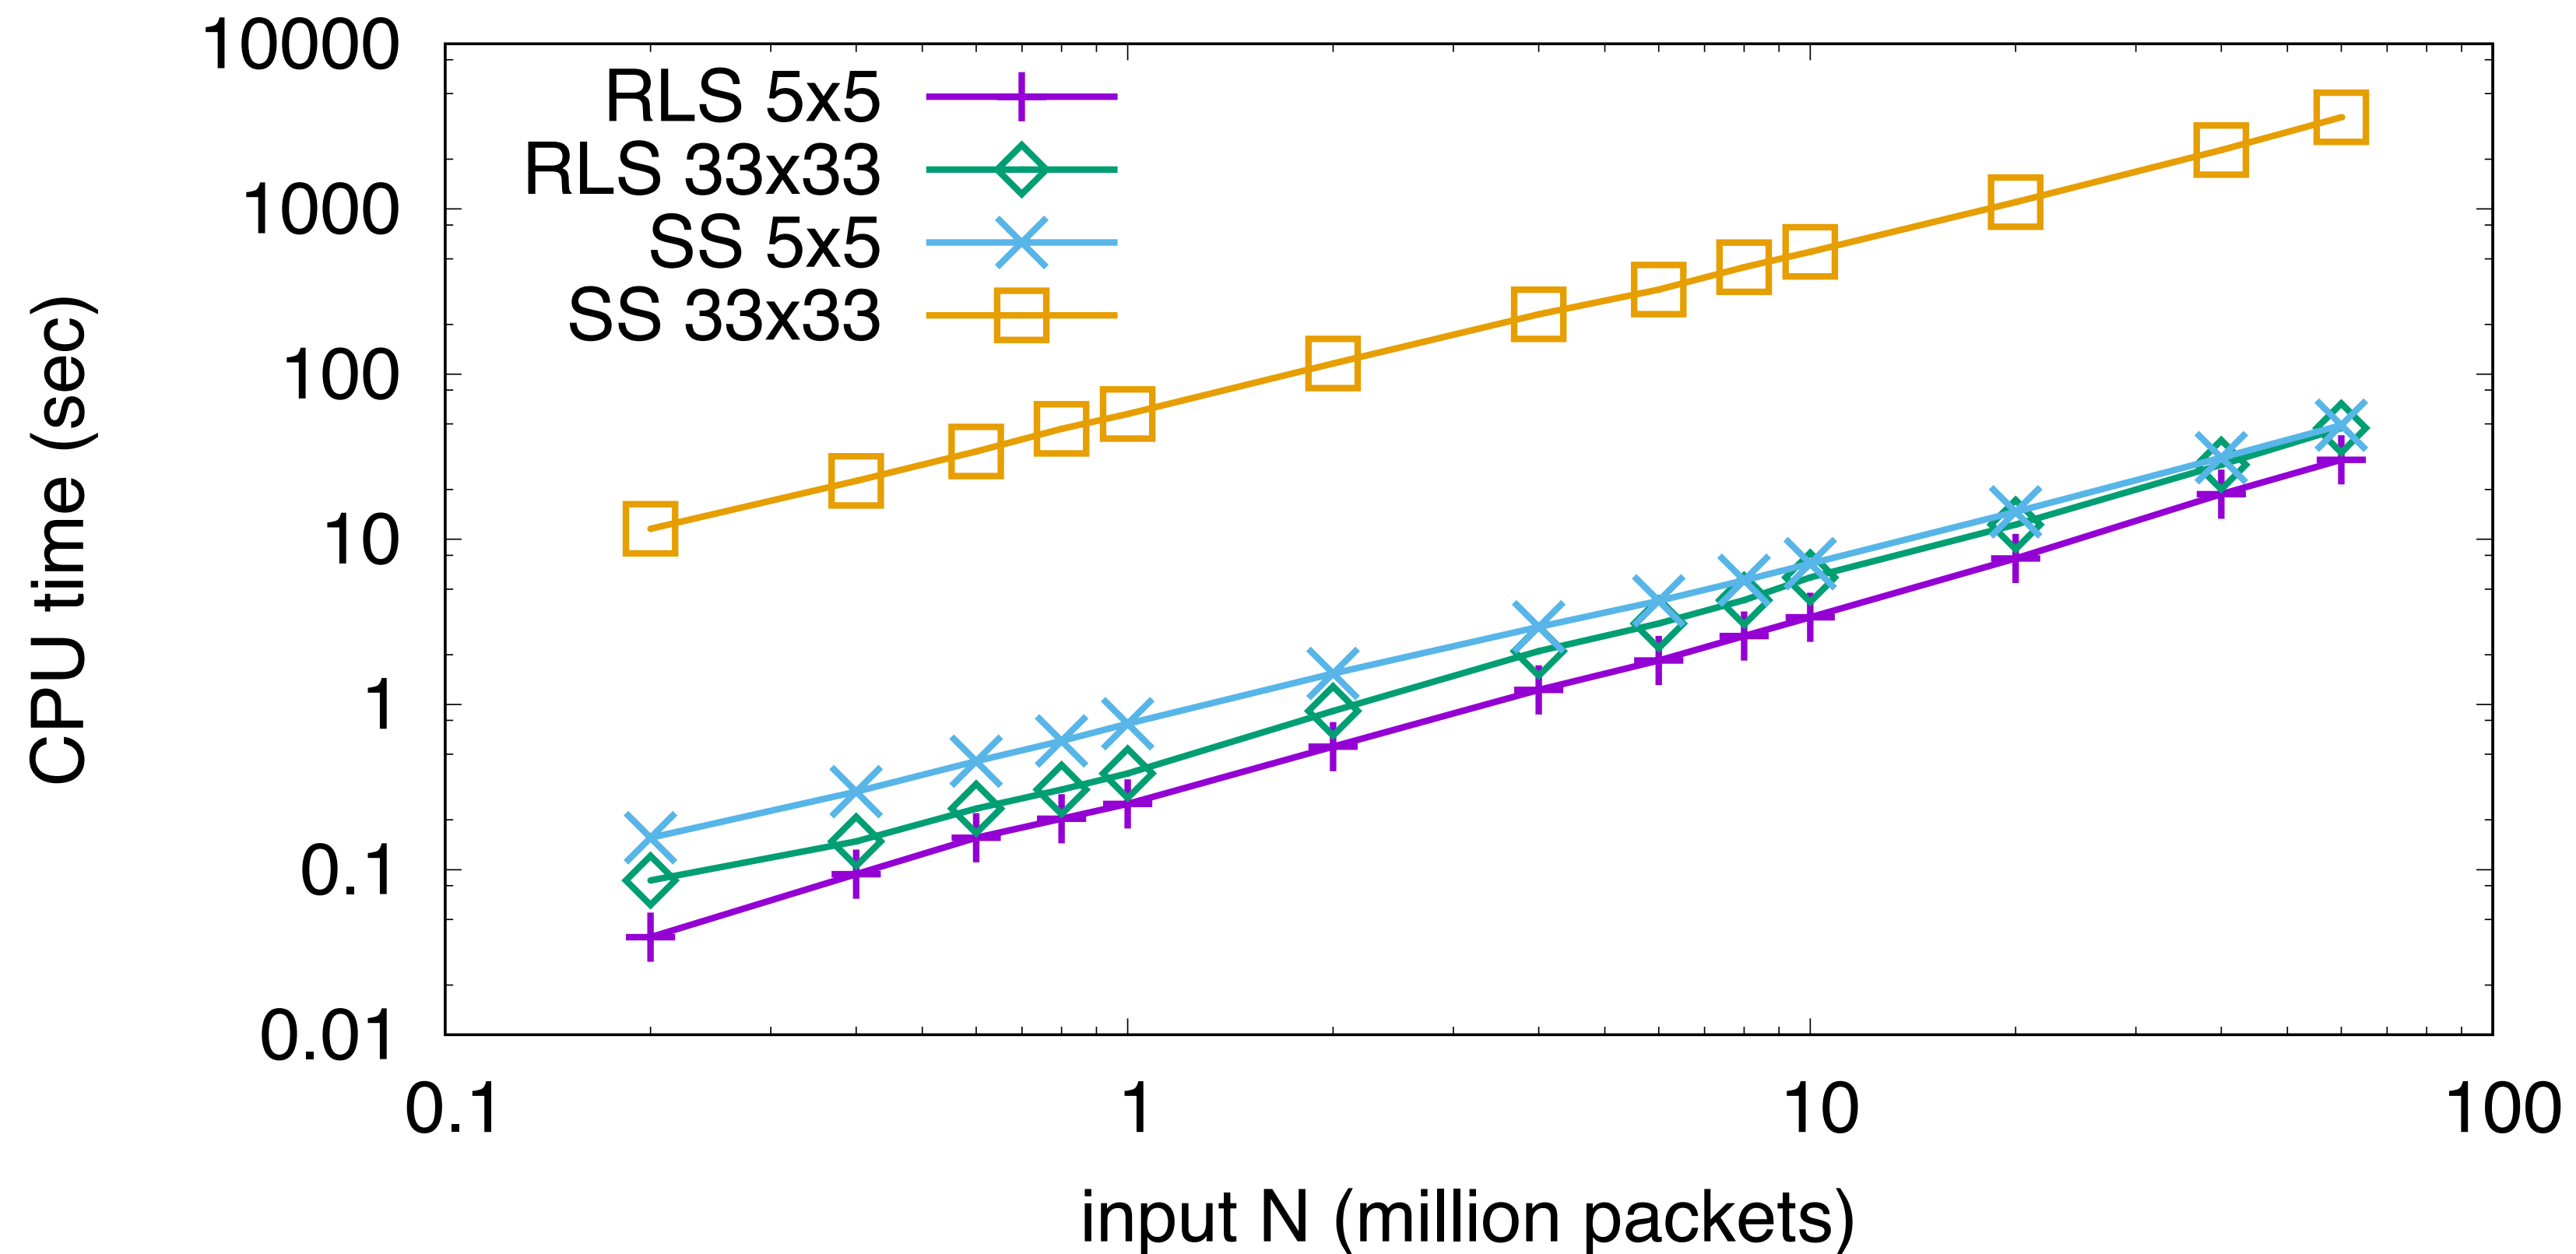
- # of HHHs: RLS:15, SS:52
- missing HHHs: not informative
  - double-counting / short prefix lengths
  - 40 missing HHHs, 35 in (I), 4 in (II), 1 in (III)
- RLS: concise and compact summary

region	no	aggregated by (src,dst)		$c'/N(\%)$
		src	dst	
VI	(1)	112.31.100.1/32	163.229.97.230/32	16.5
	(2)	64.0.0.0/2	202.203.3.13/32	5.2
V	(3)	128.0.0.0/1	202.203.3.13/32	5.8
	(4)	*	202.26.162.46/32	6.0
III	(5)	163.229.96.0/23	*	5.0
	(6)	203.179.128.0/20	*	6.8
II	(7)	*	202.203.3.0/24	5.9
	(8)	*	203.179.140.0/23	5.7
	(9)	*	163.229.128.0/17	5.1
I	(10)	0.0.0.0/1	202.192.0.0/12	5.3
	(11)	202.192.0.0/12	*	6.7
	(12)	*	202.0.0.0/7	7.6
	(13)	<b>128.0.0.0/4</b>	*	<b>5.0</b>
	(14)	<b>128.0.0.0/2</b>	*	<b>6.0</b>
-	*	*	<b>2.0</b>	
			100.0	

no	RLS(%)	SS(%)	missing SS HHHs with their $c'/N(\%)$
(1)	16.5	16.5	-
(2)	5.2	5.2	-
(3)	5.8	5.8	-
(4)	6.0	6.0	-
(5)	5.0	5.0	-
(6)	6.8	6.8	-
(7)	<b>5.9</b>	<b>16.9</b>	-
(8)	5.7	5.7	-
(9)	5.1	5.1	-
(10)	<b>5.3</b>	-	(96/ <b>3</b> ,202.203/ <b>16</b> ):5.4 (0/ <b>2</b> ,202.203/ <b>16</b> ):5.6 (112/ <b>4</b> ,202.192/ <b>12</b> ):5.2 (64/ <b>2</b> ,202.192/ <b>12</b> ):9.0
(11)	6.7	6.7	-
(12)	<b>7.6</b>	-	(0/ <b>1</b> ,203.179.128/ <b>20</b> ):6.0 (128/ <b>2</b> ,202.203/ <b>16</b> ):5.5 (192/ <b>4</b> ,202/ <b>8</b> ):5.1 (*,202.192/ <b>12</b> ):25.5 (16/ <b>4</b> ,202/ <b>7</b> ):5.4 (128/ <b>1</b> ,202.128/ <b>9</b> ):10.6 (64/ <b>2</b> ,202/ <b>7</b> ):15.5 (128/ <b>1</b> ,202/ <b>7</b> ):17.7
(13)	<b>5.0</b>	<b>5.2</b>	-
(14)	<b>6.0</b>	-	(163.229/ <b>16</b> ,0/ <b>1</b> ):6.0 (144/ <b>4</b> ,128/ <b>1</b> ):5.3 (128/ <b>2</b> ,96/ <b>3</b> ):5.0 (128/ <b>3</b> ,0/ <b>1</b> ):5.3 (160/ <b>3</b> ,128/ <b>1</b> ):7.0 (128/ <b>2</b> ,0/ <b>2</b> ):5.7 (128/ <b>2</b> ,0/ <b>1</b> ):11.4
(15)	<b>5.4</b>	<b>33.1</b>	(128/ <b>1</b> ,160/ <b>6</b> ):5.0 (192/ <b>4</b> ,128/ <b>2</b> ):5.2 (0/ <b>1</b> ,128/ <b>2</b> ):22.7 (* ,128/ <b>3</b> ):7.1
-	<b>2.0</b>	-	(202/ <b>7</b> ,0/ <b>2</b> ):5.4 (192/ <b>8</b> ,128/ <b>1</b> ):5.6 (202/ <b>8</b> ,0/ <b>1</b> ):5.7 (202/ <b>7</b> ,128/ <b>1</b> ):6.0 (192/ <b>3</b> ,200/ <b>5</b> ):10.5 (128/ <b>1</b> ,112/ <b>6</b> ):5.1 (112/ <b>5</b> ,128/ <b>1</b> ):21.8 (200/ <b>5</b> ,*):17.0 (192/ <b>4</b> ,128/ <b>1</b> ):13.6 (128/ <b>1</b> ,16/ <b>4</b> ):6.2 (*,200/ <b>5</b> ):42.4 (64/ <b>3</b> ,128/ <b>1</b> ):6.0 (96/ <b>3</b> ,128/ <b>1</b> ):29.7 (128/ <b>1</b> ,64/ <b>2</b> ):10.4 (0/ <b>1</b> ,128/ <b>1</b> ):46.7 (128/ <b>1</b> ,*):53.3 (*,128/ <b>1</b> ):78.3

# CPU time: RLS vs. SS

- RLS: lower cost for finer granularity
  - 100+ times faster for bit-wise aggregation!



# memory usage: RLS vs. SS

- RLS: proportional to inputs (ok for modern PCs)
- SS: fixed memory usage

