Fundamental Bounds on the Interconnect Complexity of Decoder Implementations

. . . towards a theory of “green” communication

Pulkit Grover

with Anant Sahai, Jan Rabaey

Wireless Foundations and Berkeley Wireless Research Center, UC Berkeley

CISS, Johns Hopkins University, 24 March 2011

Motivation:
Where is power consumed in a communication system?

Motivation:
Where is power consumed in a communication system?

Information-theorist's abstraction

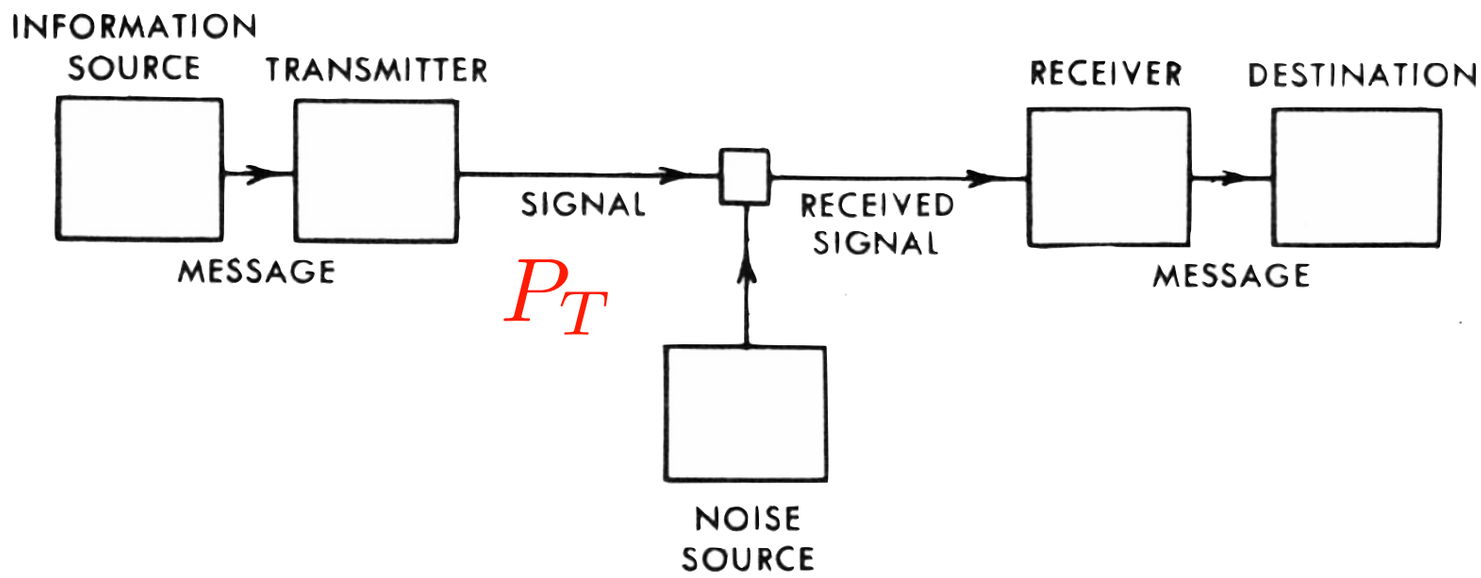

Motivation:
Where is power consumed in a communication system?

Information-theorist's abstraction

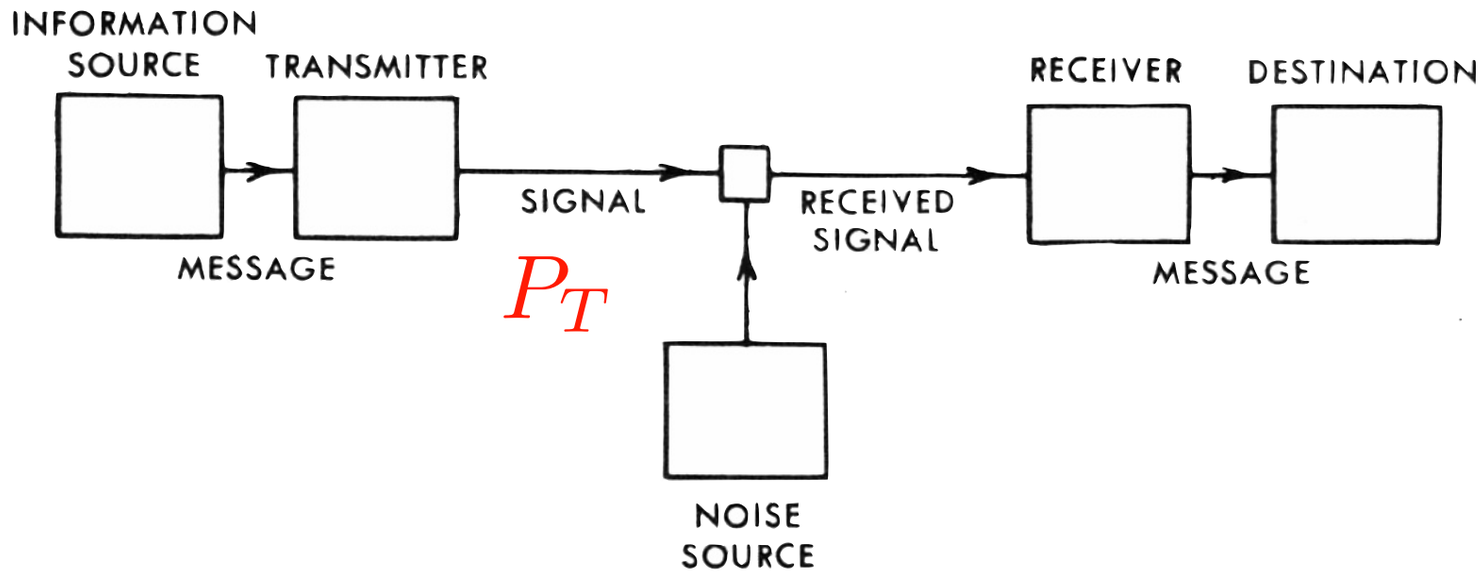

System designer's abstraction:

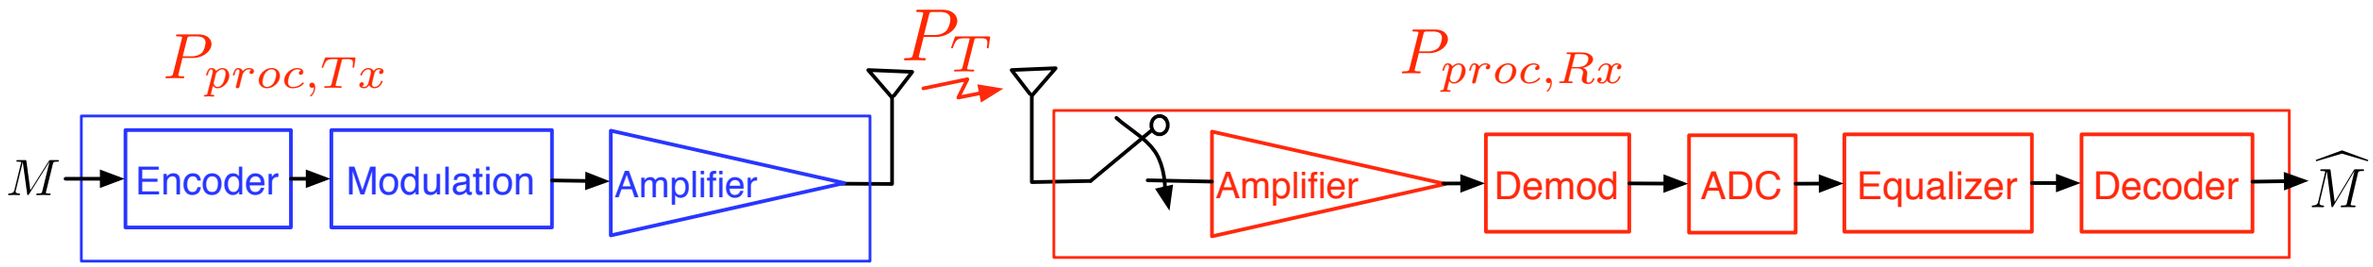

Motivation:
Where is power consumed in a communication system?

Information-theorist's abstraction

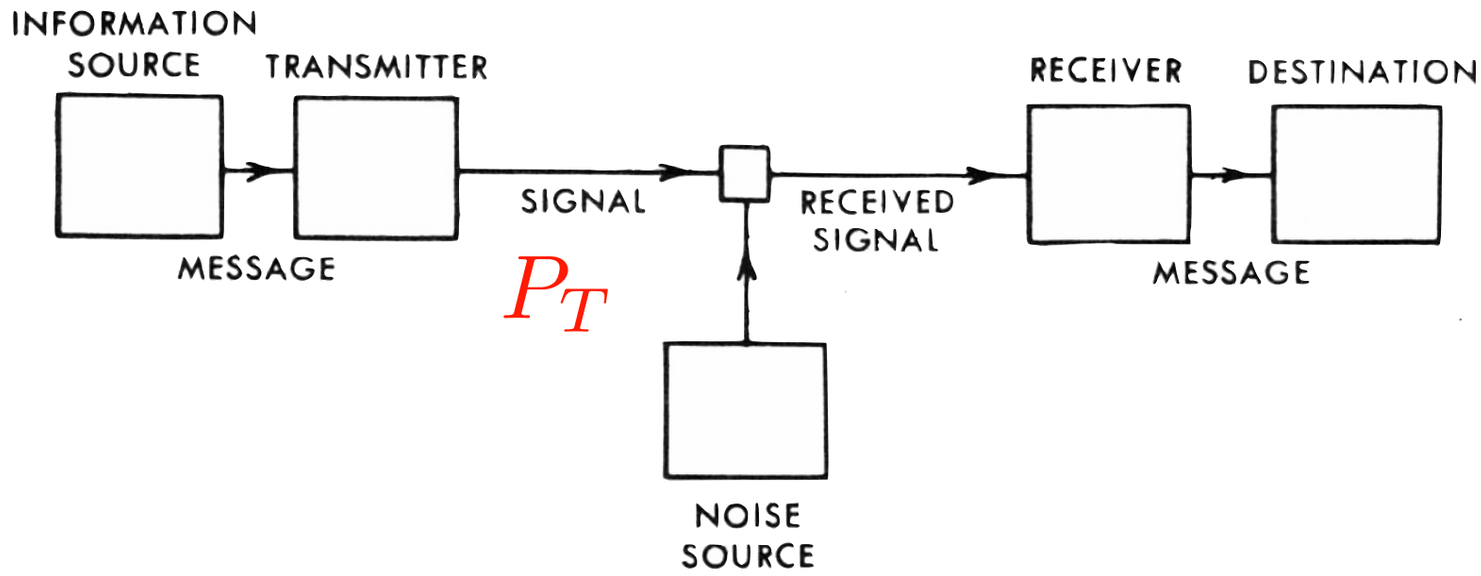

System designer's abstraction:

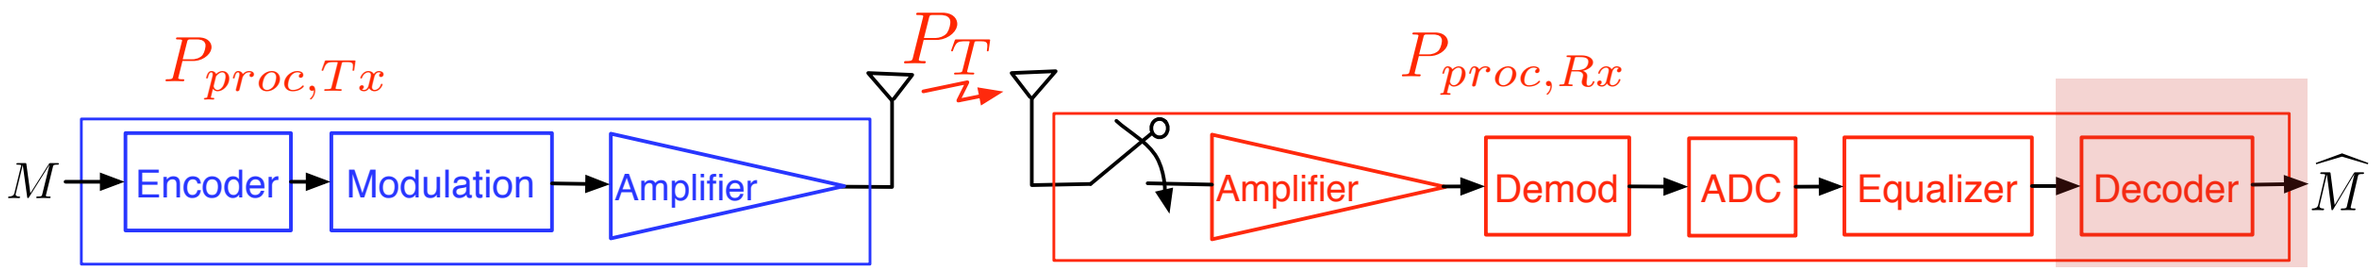

Decoding power is substantial

Decoding power is substantial

Empirical observation: decoding power > transmit power at short distances!

Decoding power is substantial

Empirical observation: decoding power > transmit power at short distances!

"theoretical" verification:

Decoding power is substantial

Empirical observation: decoding power > transmit power at short distances!

"theoretical" verification:

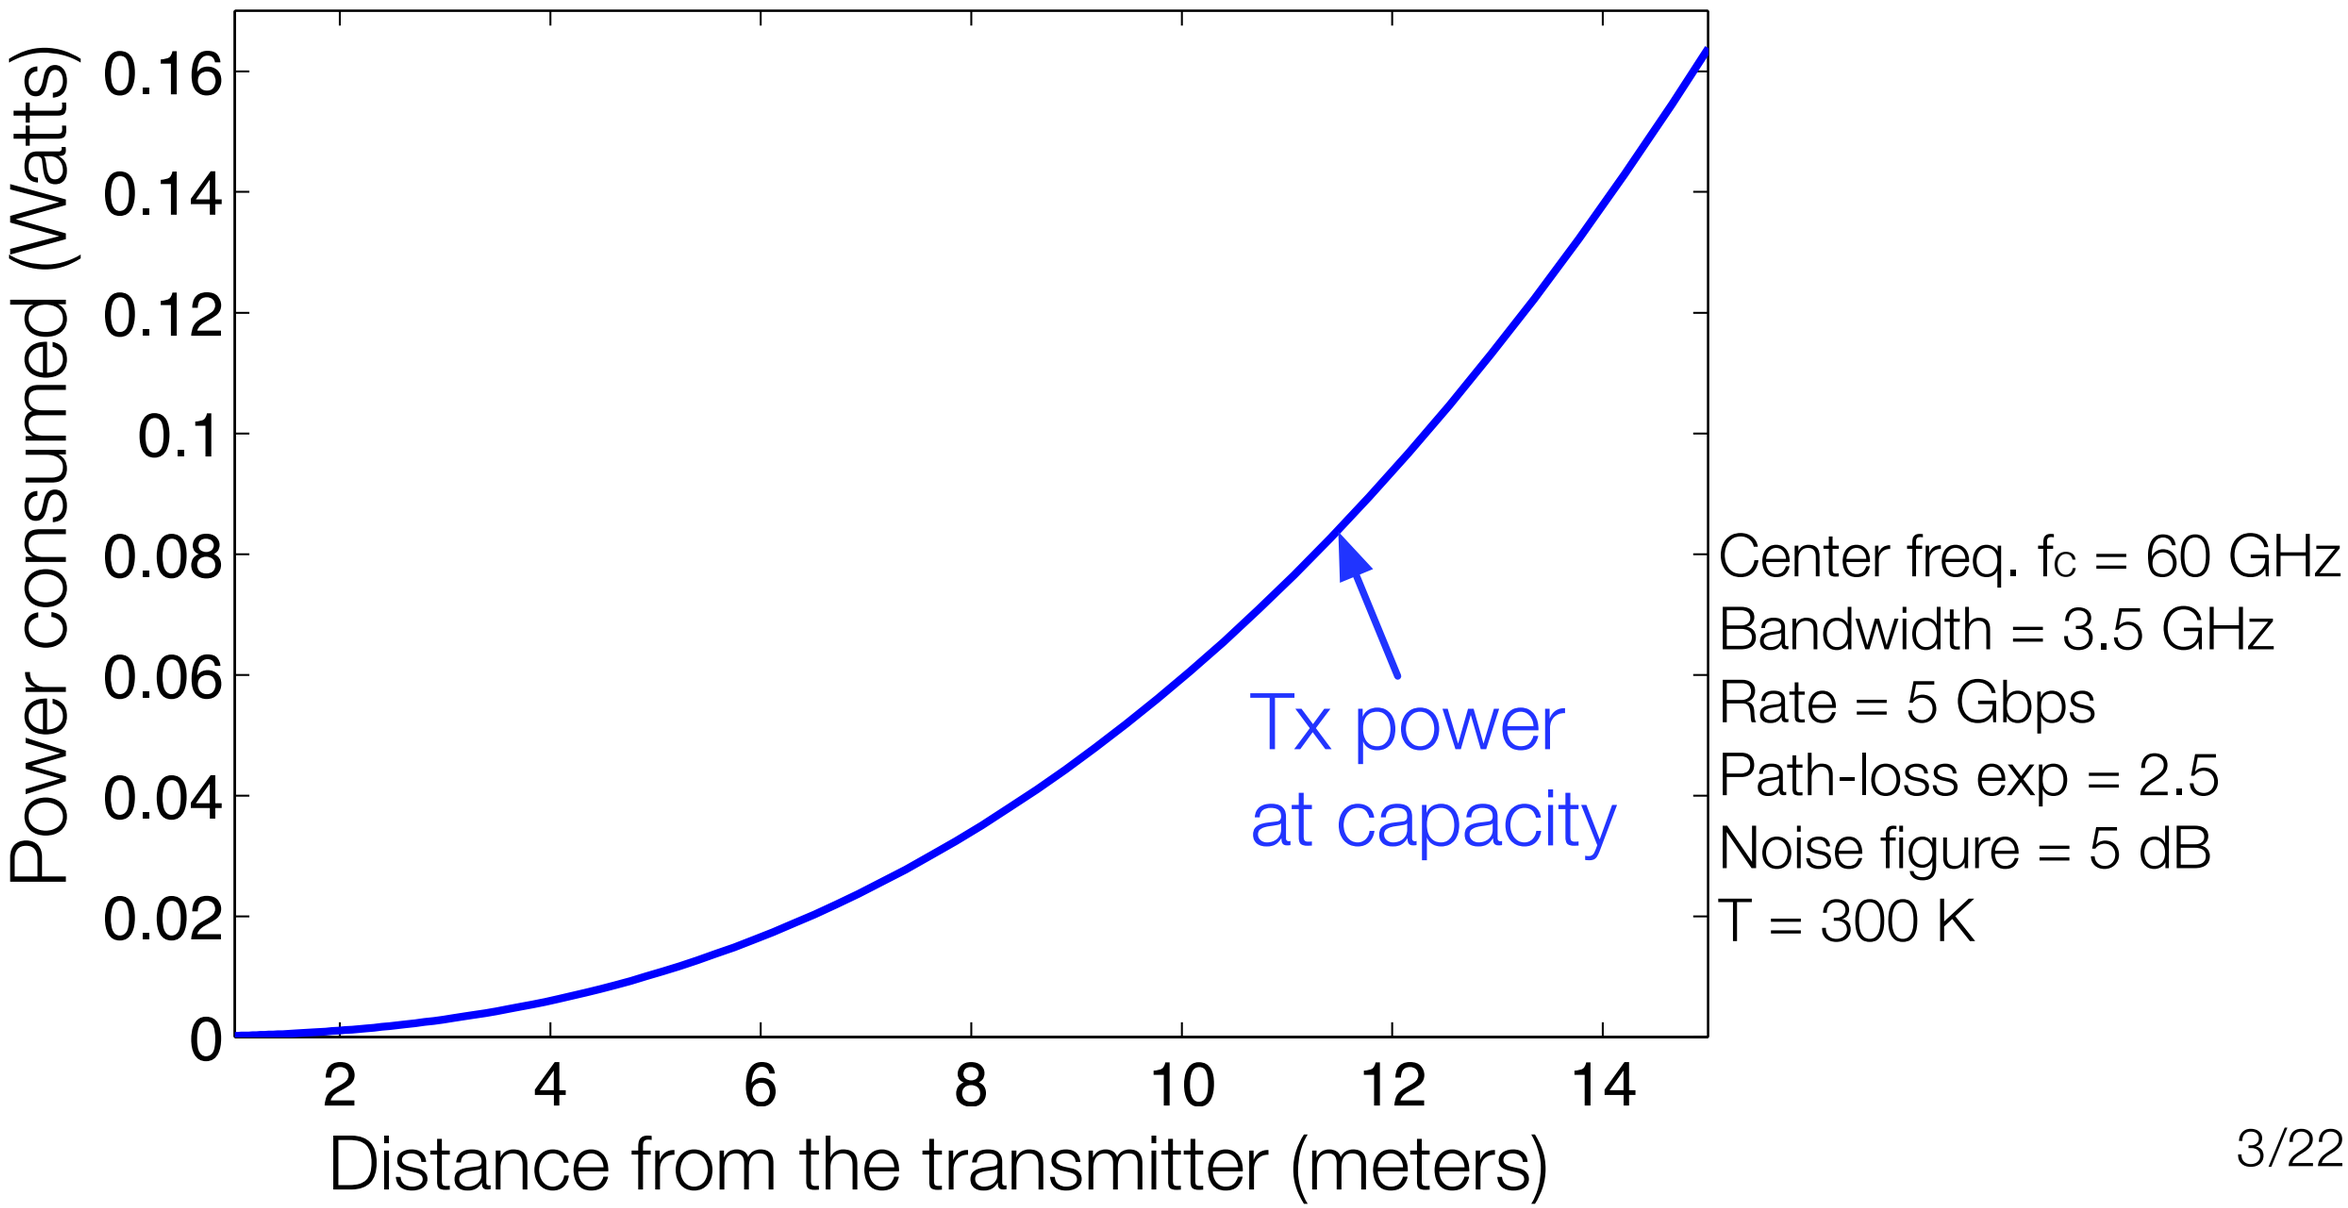

Decoding power is substantial

Empirical observation: decoding power > transmit power at short distances!

"theoretical" verification:

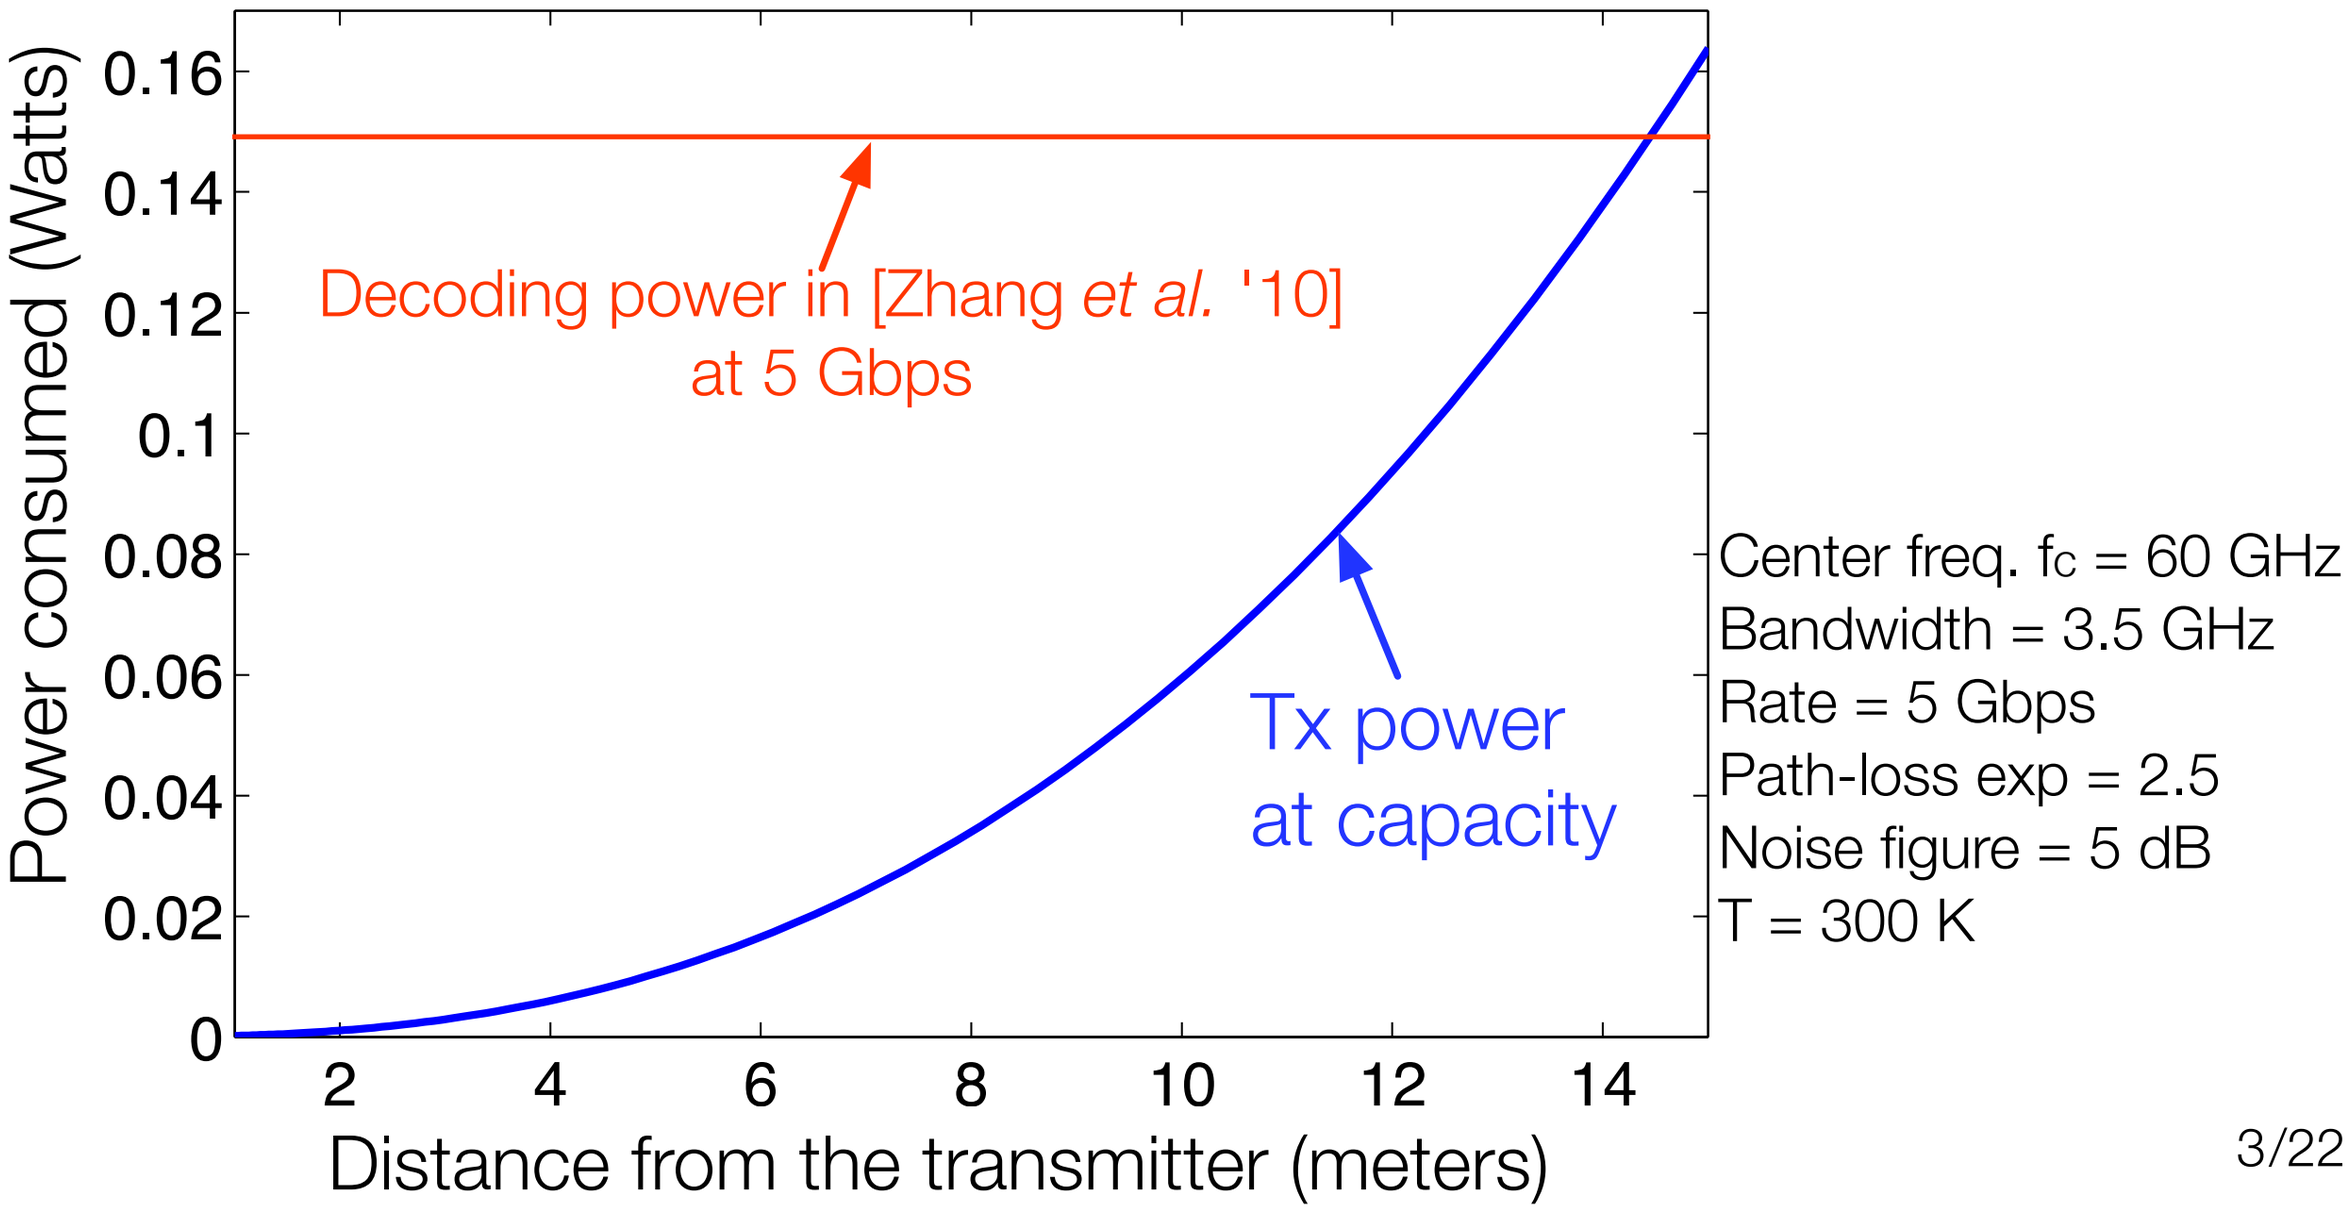

Decoding power is substantial

Empirical observation: decoding power > transmit power at short distances!

"theoretical" verification:

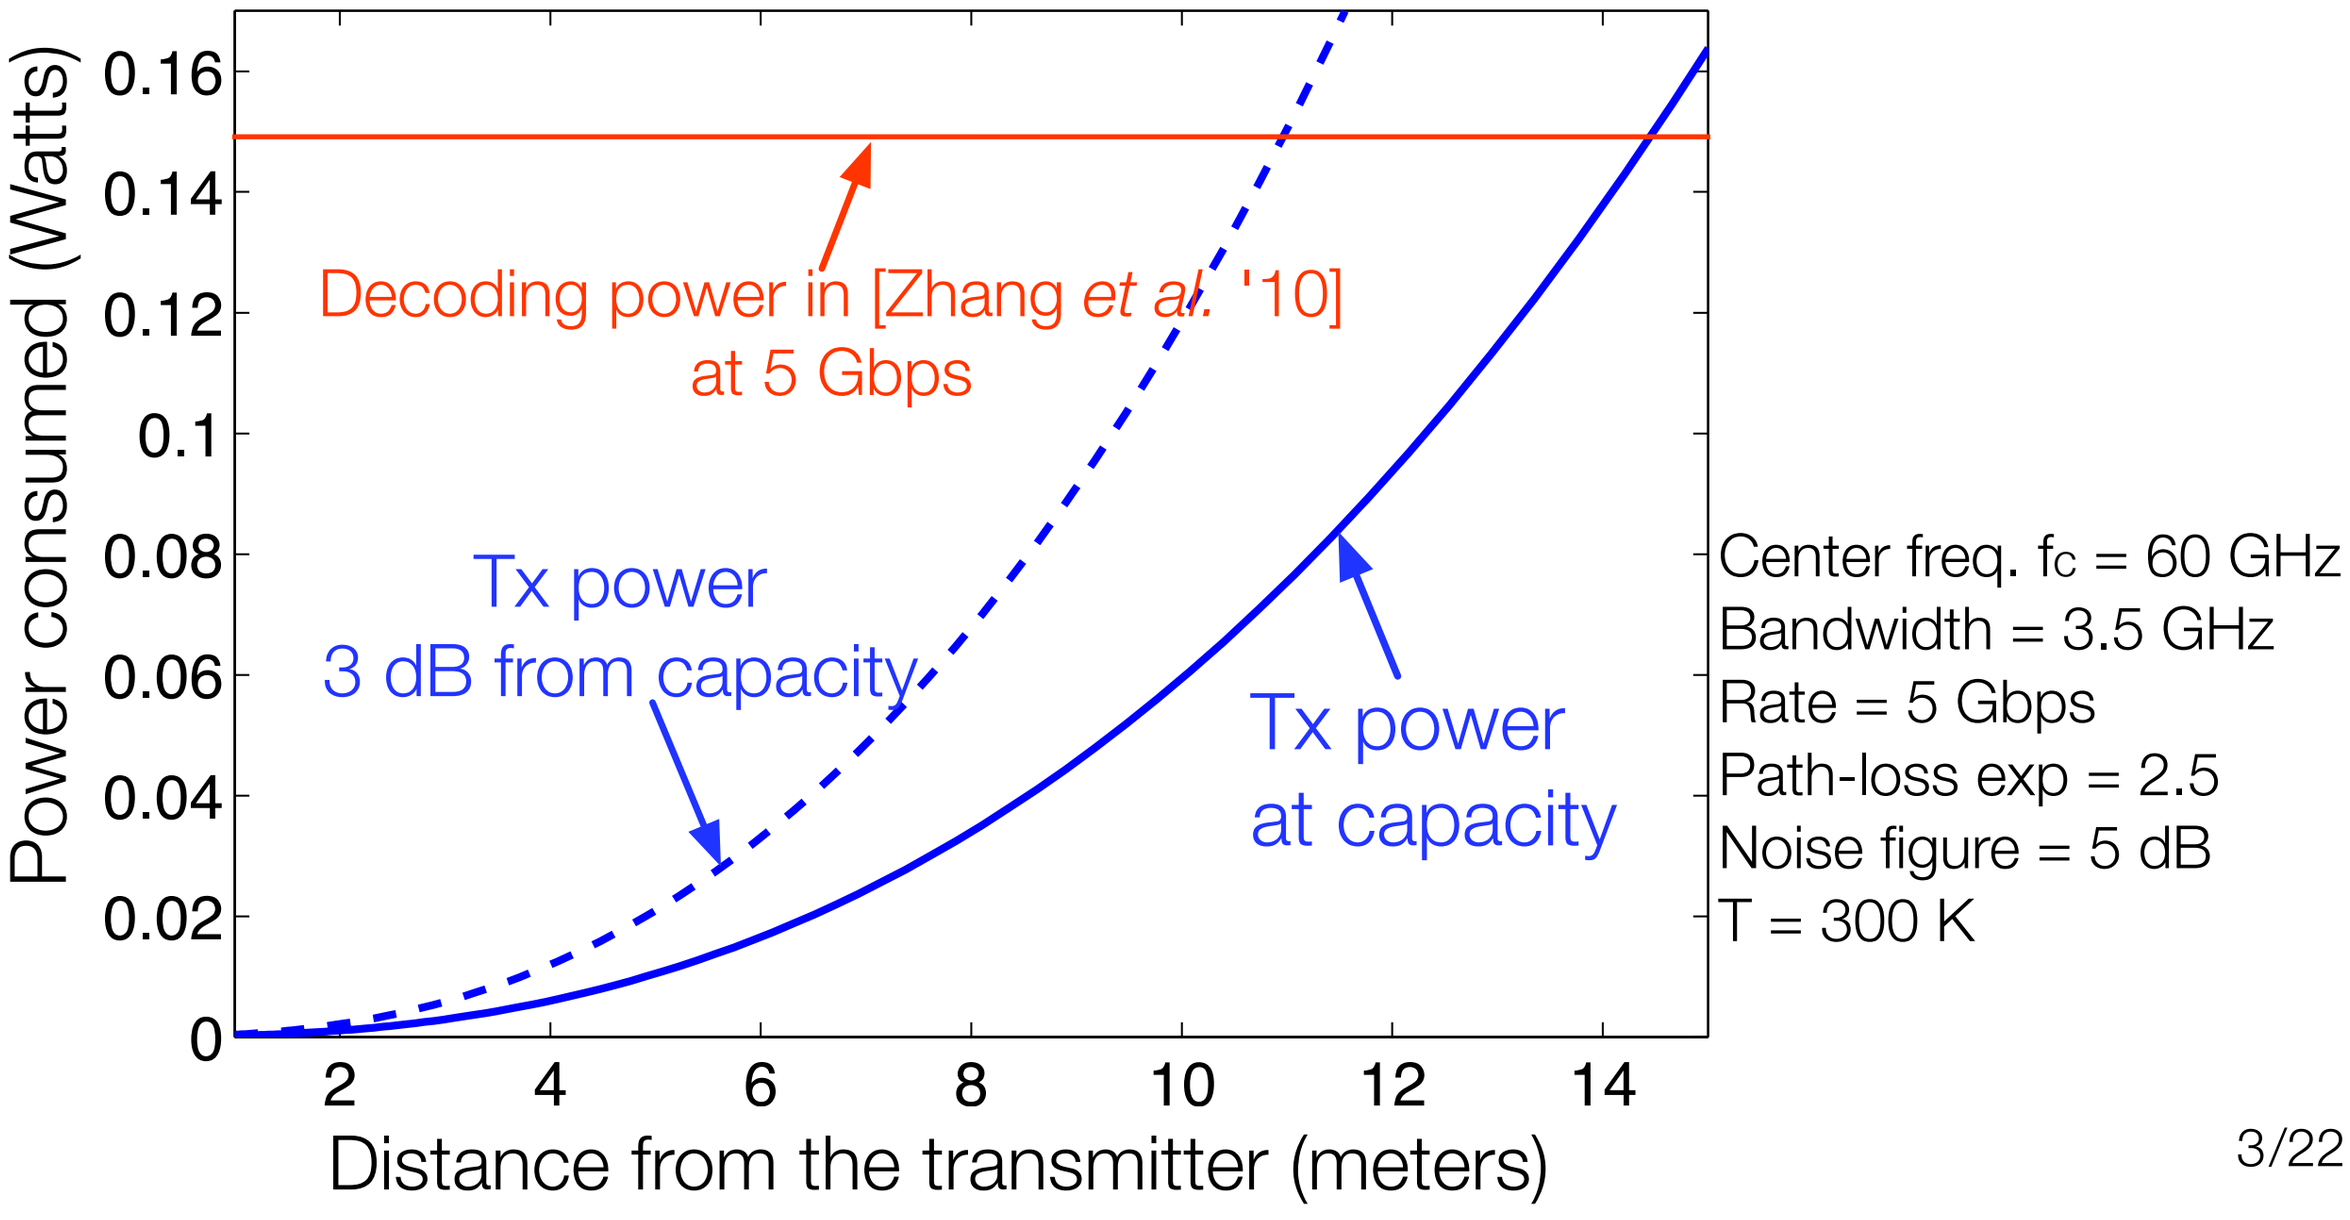

Talk outline

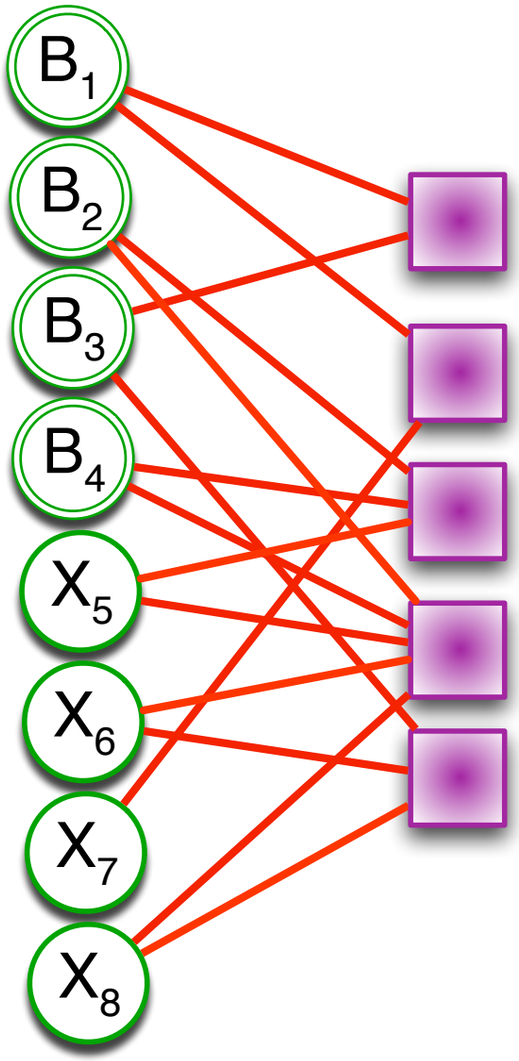

Talk outline

1. How do we build power-efficient communication systems?

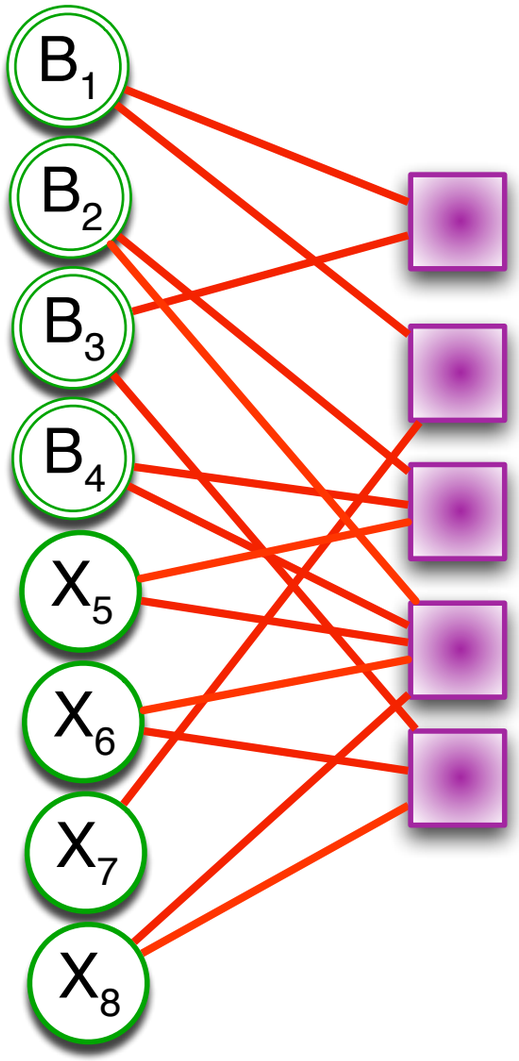

Talk outline

1. How do we build power-efficient communication systems?

   - understanding processing/decoding power

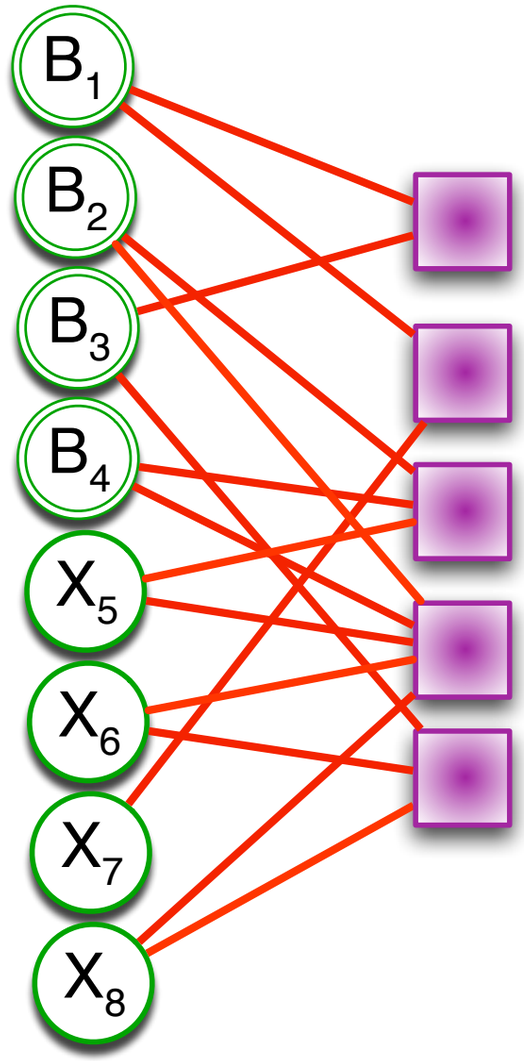

Talk outline

1. How do we build power-efficient communication systems?

   - understanding processing/decoding power

2. “VLSI model” of decoding computation

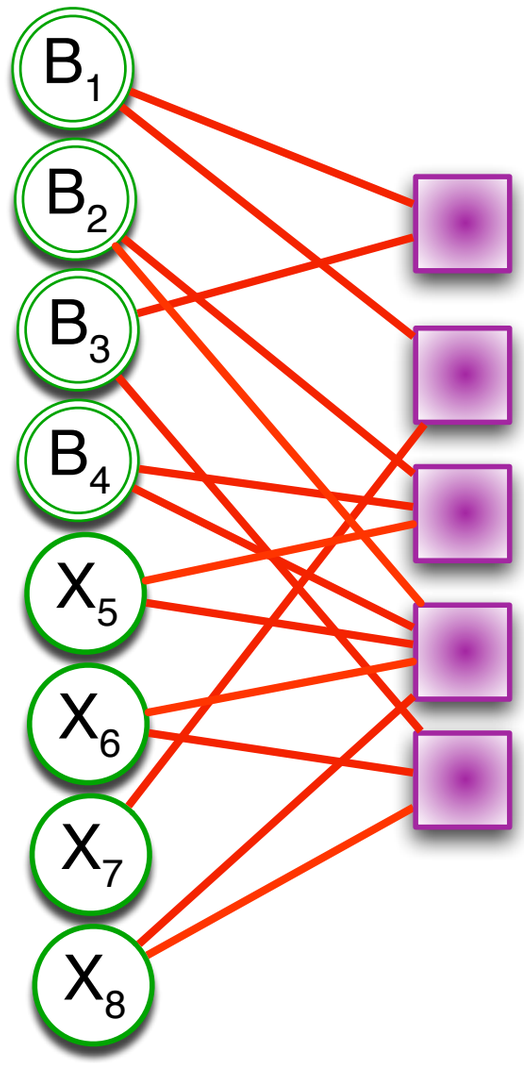

Talk outline

1. How do we build power-efficient communication systems?
   - understanding processing/decoding power
2. “VLSI model” of decoding computation
3. Decoding power model: “Wire model”

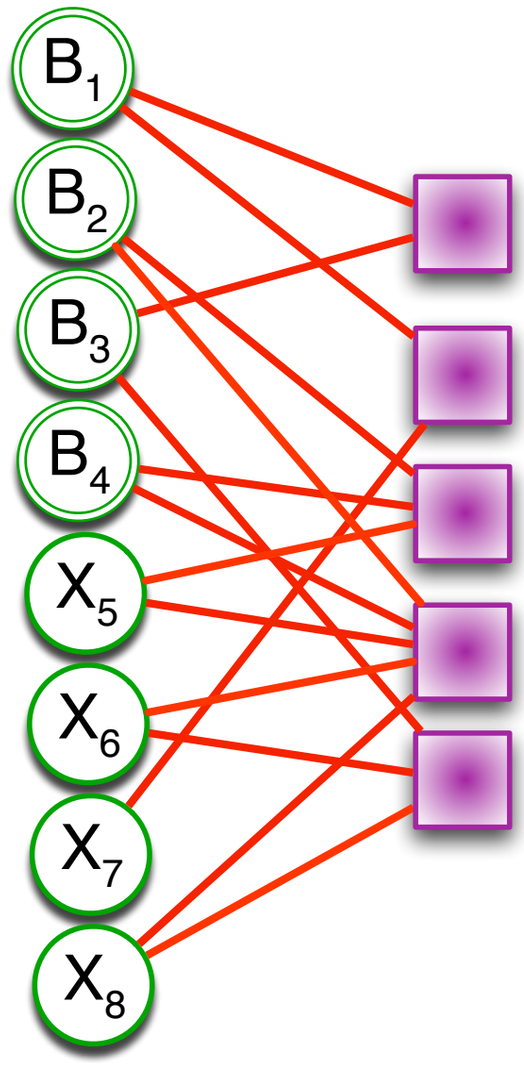

Talk outline

1. How do we build power-efficient communication systems?

   - understanding processing/decoding power

2. “VLSI model” of decoding computation

3. Decoding power model: “Wire model”

   - do not operate too close to capacity!

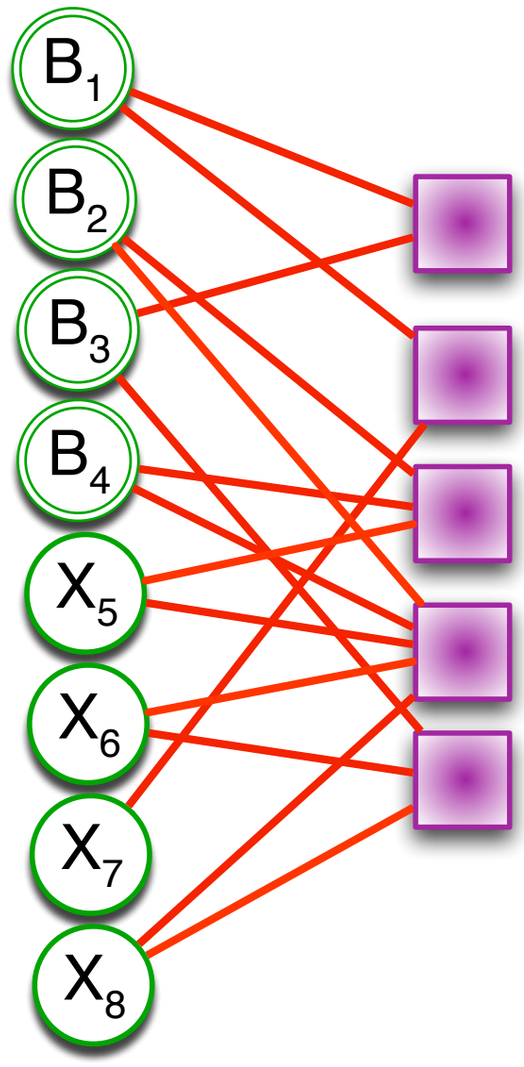

Talk outline

1. How do we build power-efficient communication systems?

   - understanding processing/decoding power

2. “VLSI model” of decoding computation

3. Decoding power model: “Wire model”

   - do not operate too close to capacity!

   - design codes for just the required performance (no better!)

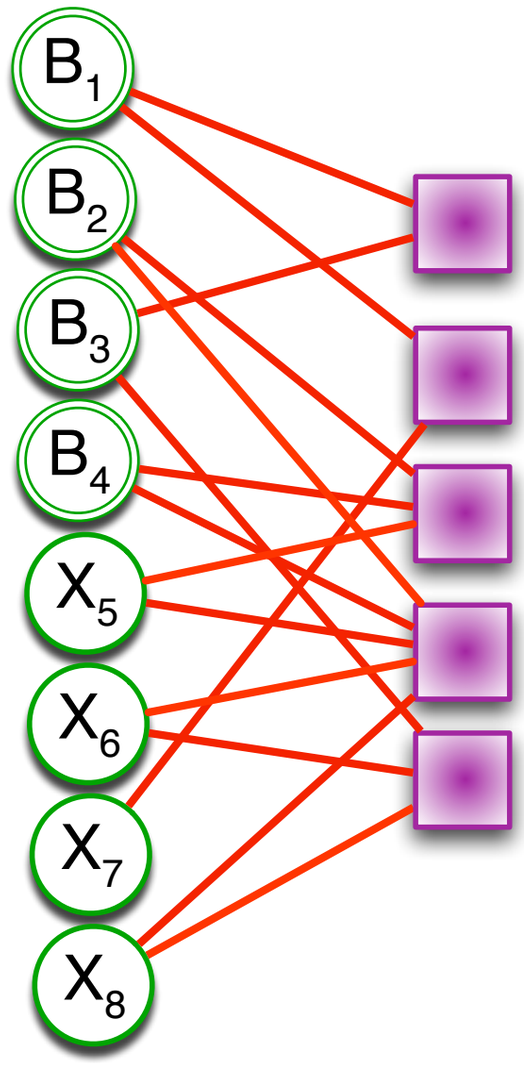

Talk outline

1. How do we build power-efficient communication systems?

   - understanding processing/decoding power

2. “VLSI model” of decoding computation

3. Decoding power model: “Wire model”

   - do not operate too close to capacity!

   - design codes for just the required performance (no better!)

4. Can more complex nodes help?

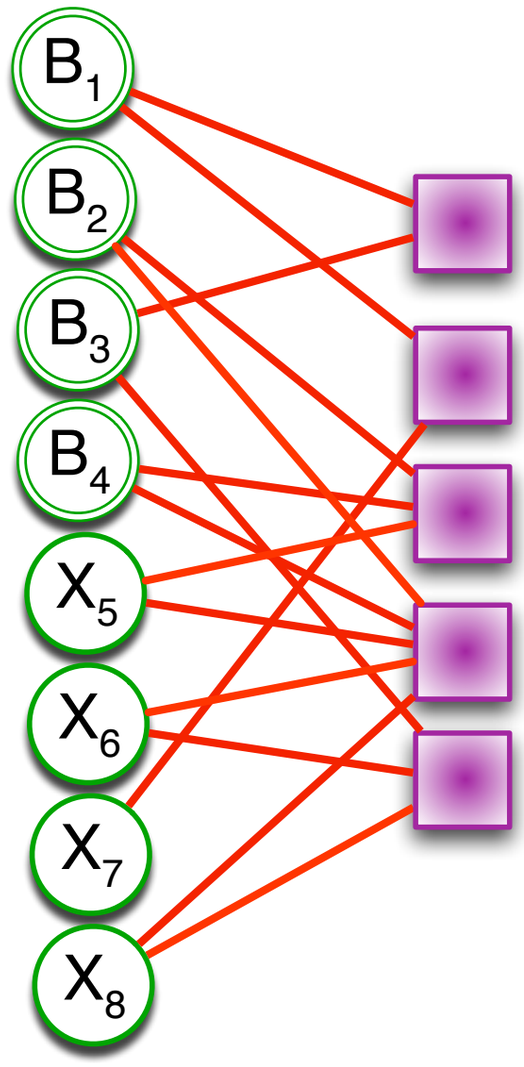

Goal: fundamental limits on transmit + computation power

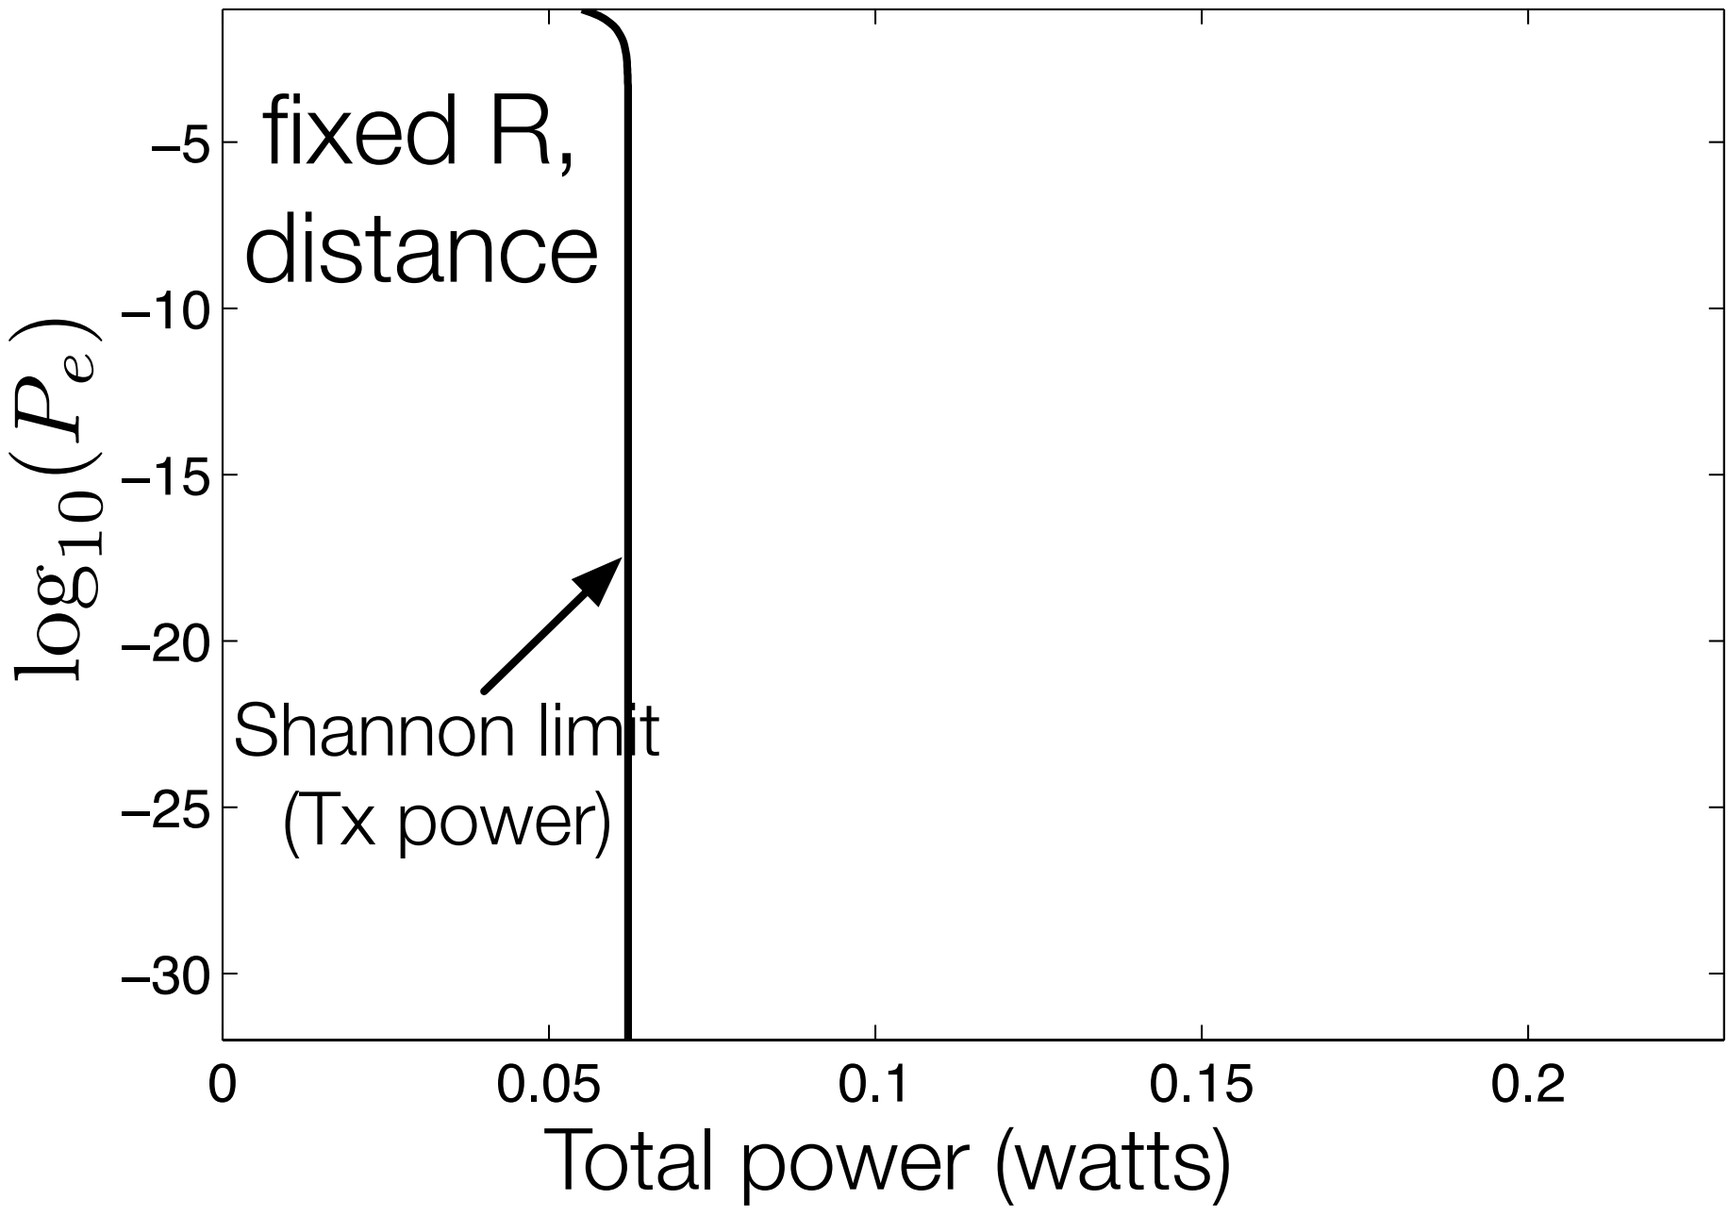

Goal: fundamental limits on transmit + computation power

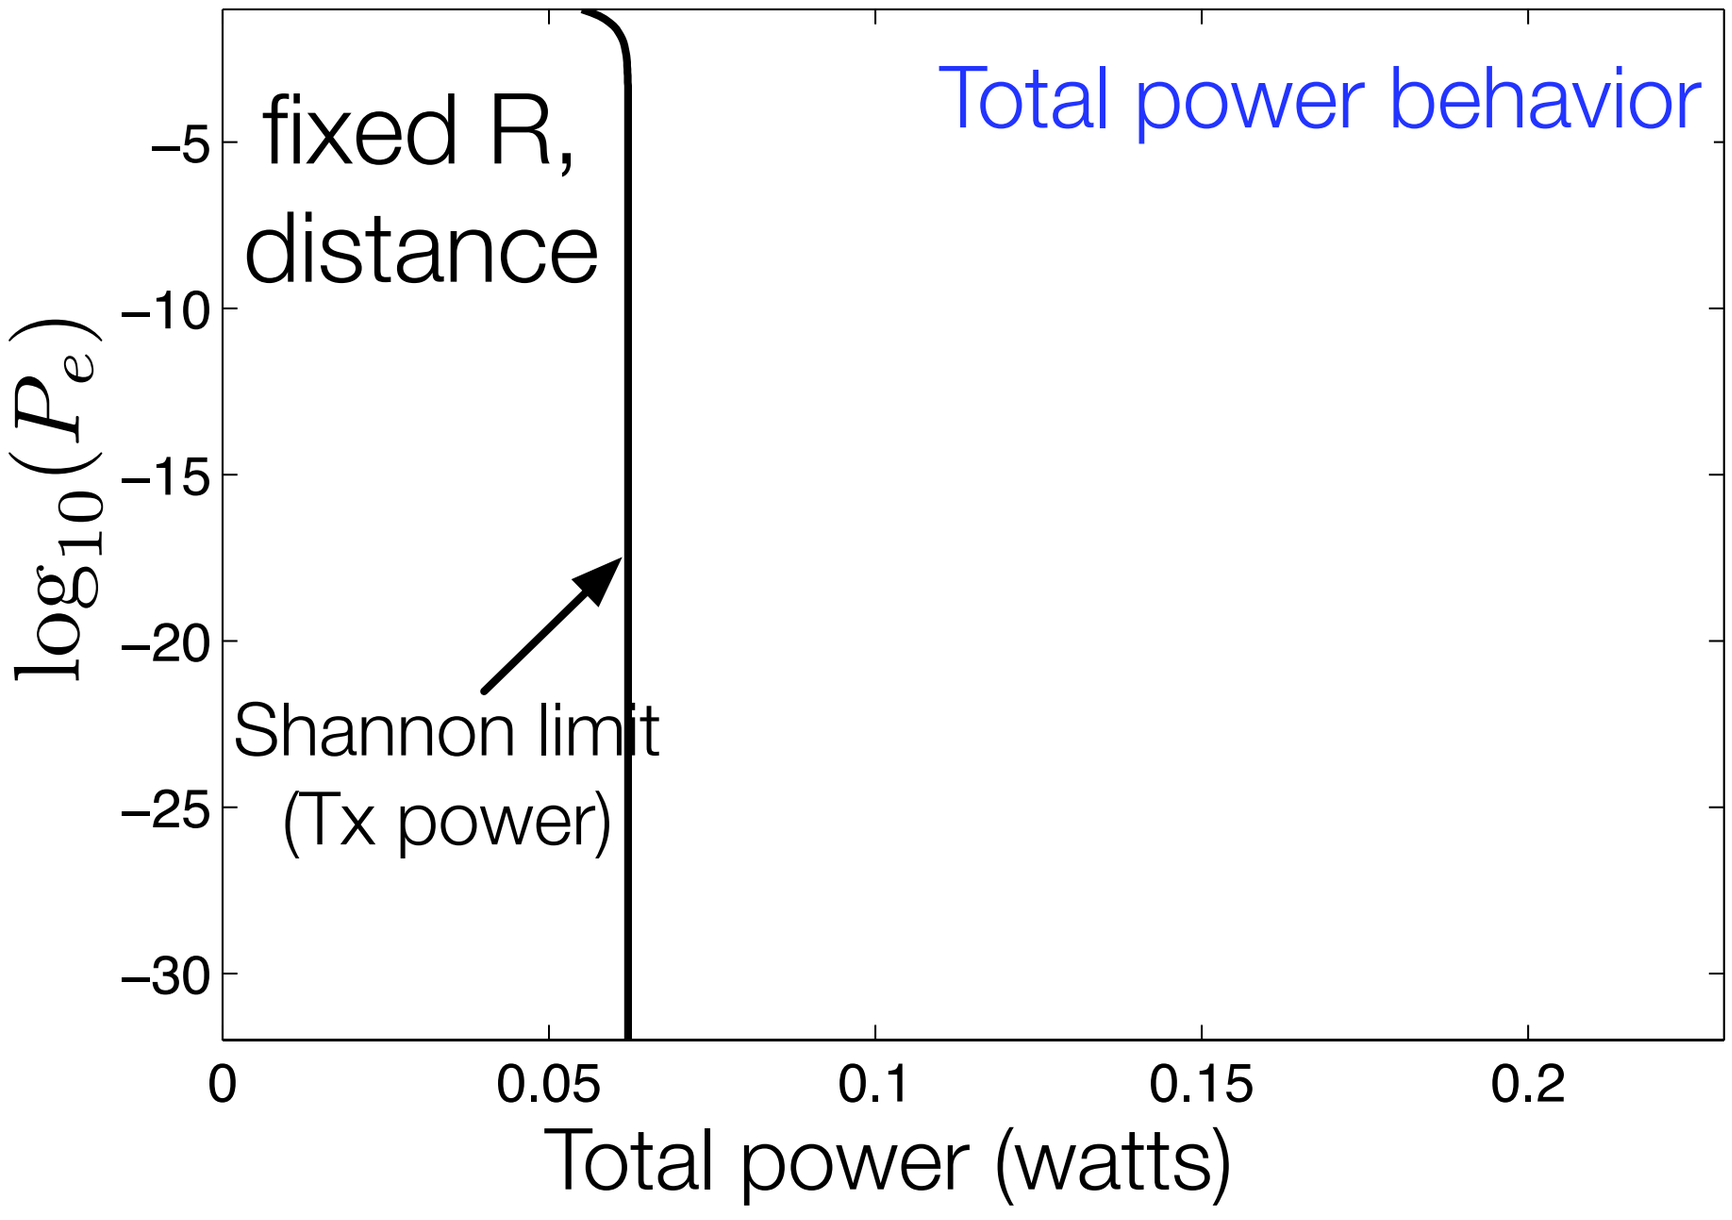

Goal: fundamental limits on transmit + computation power

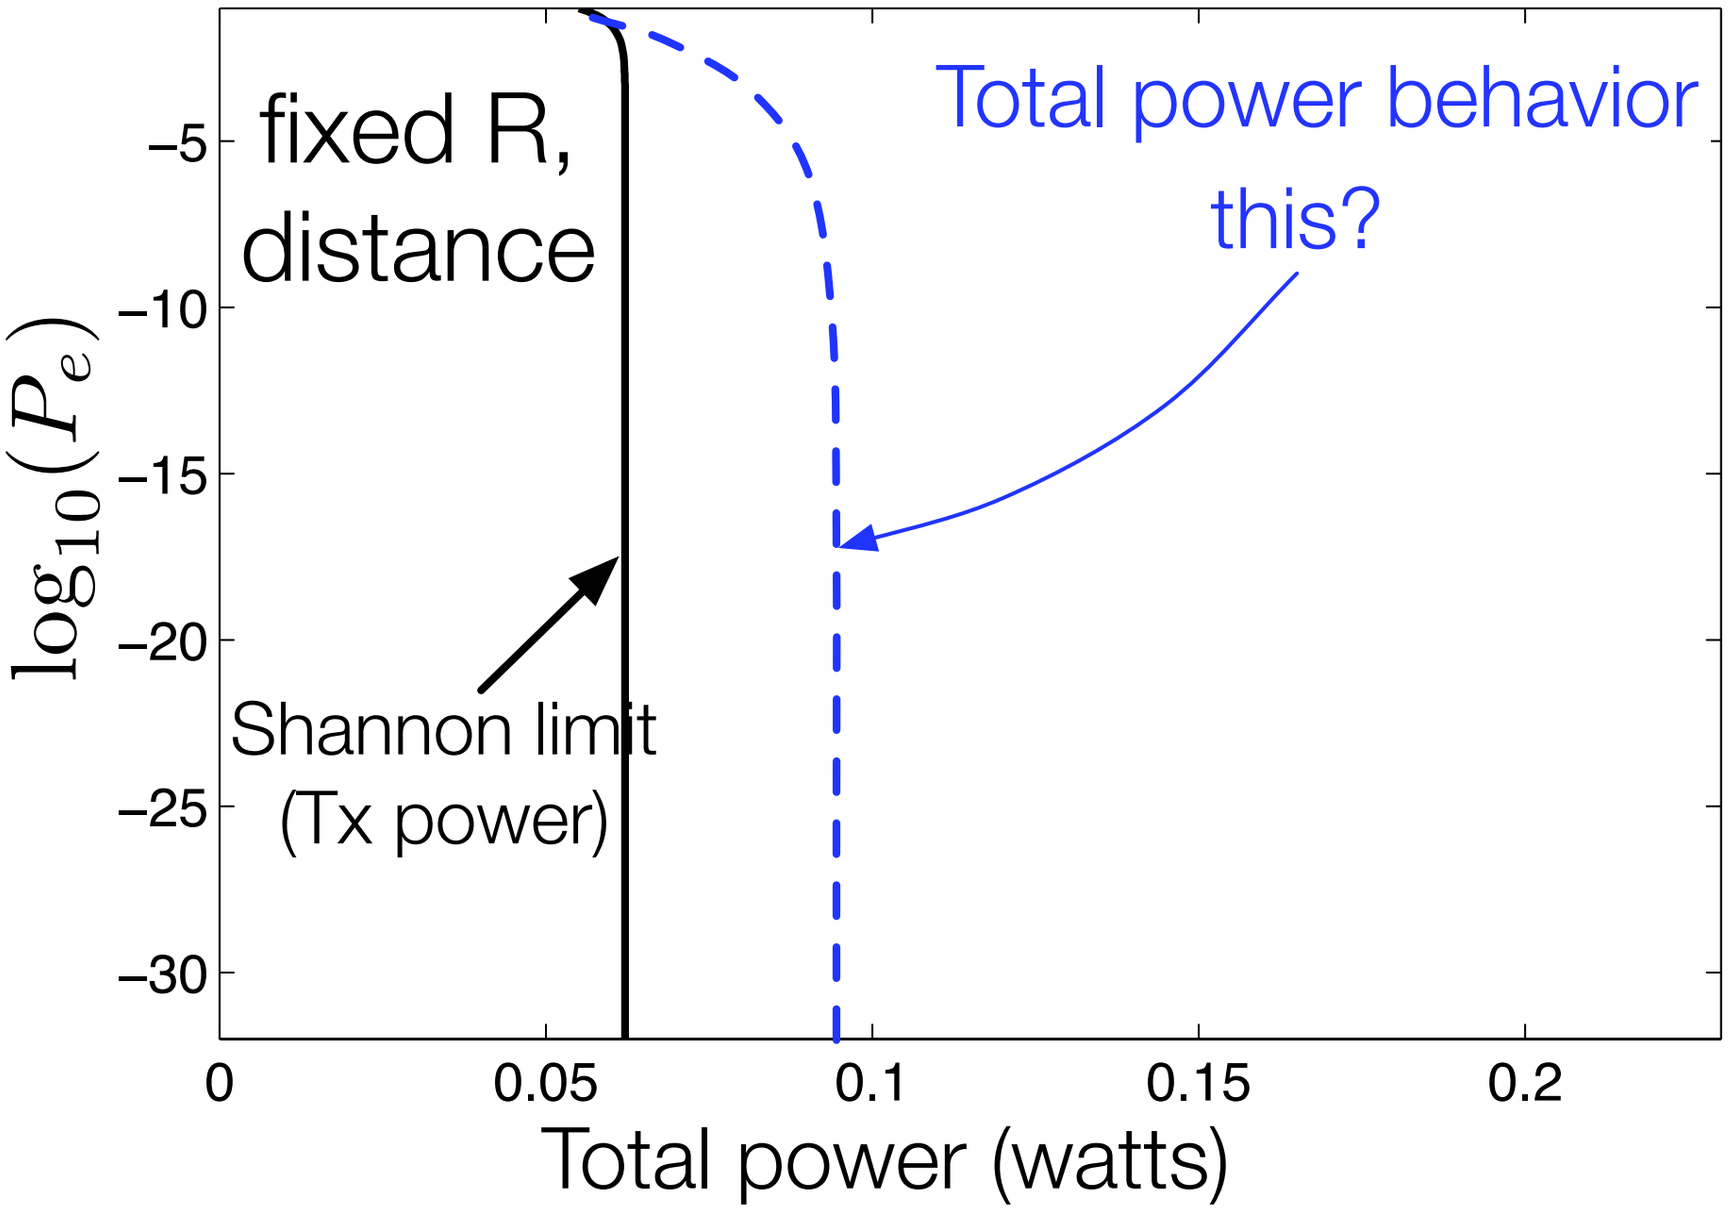

Goal: fundamental limits on transmit + computation power

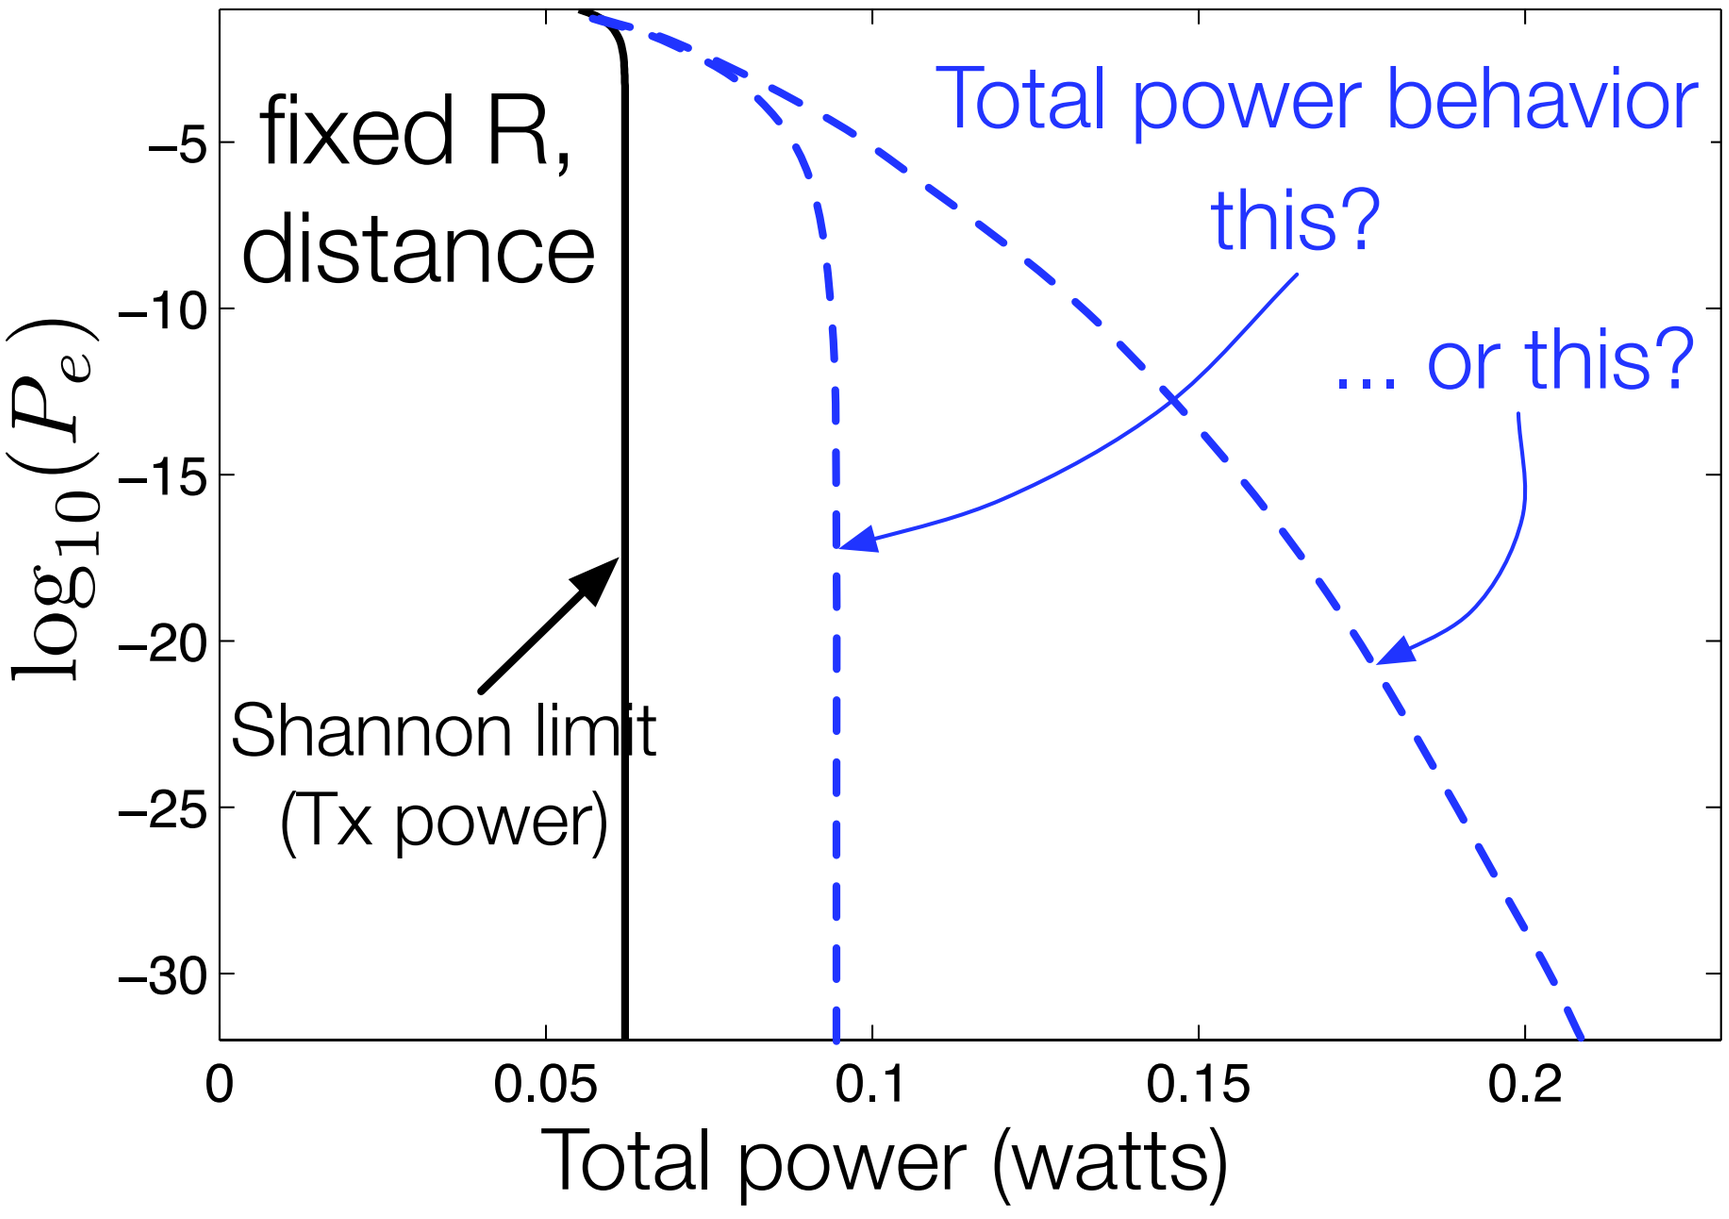

Goal: fundamental limits on transmit + computation power

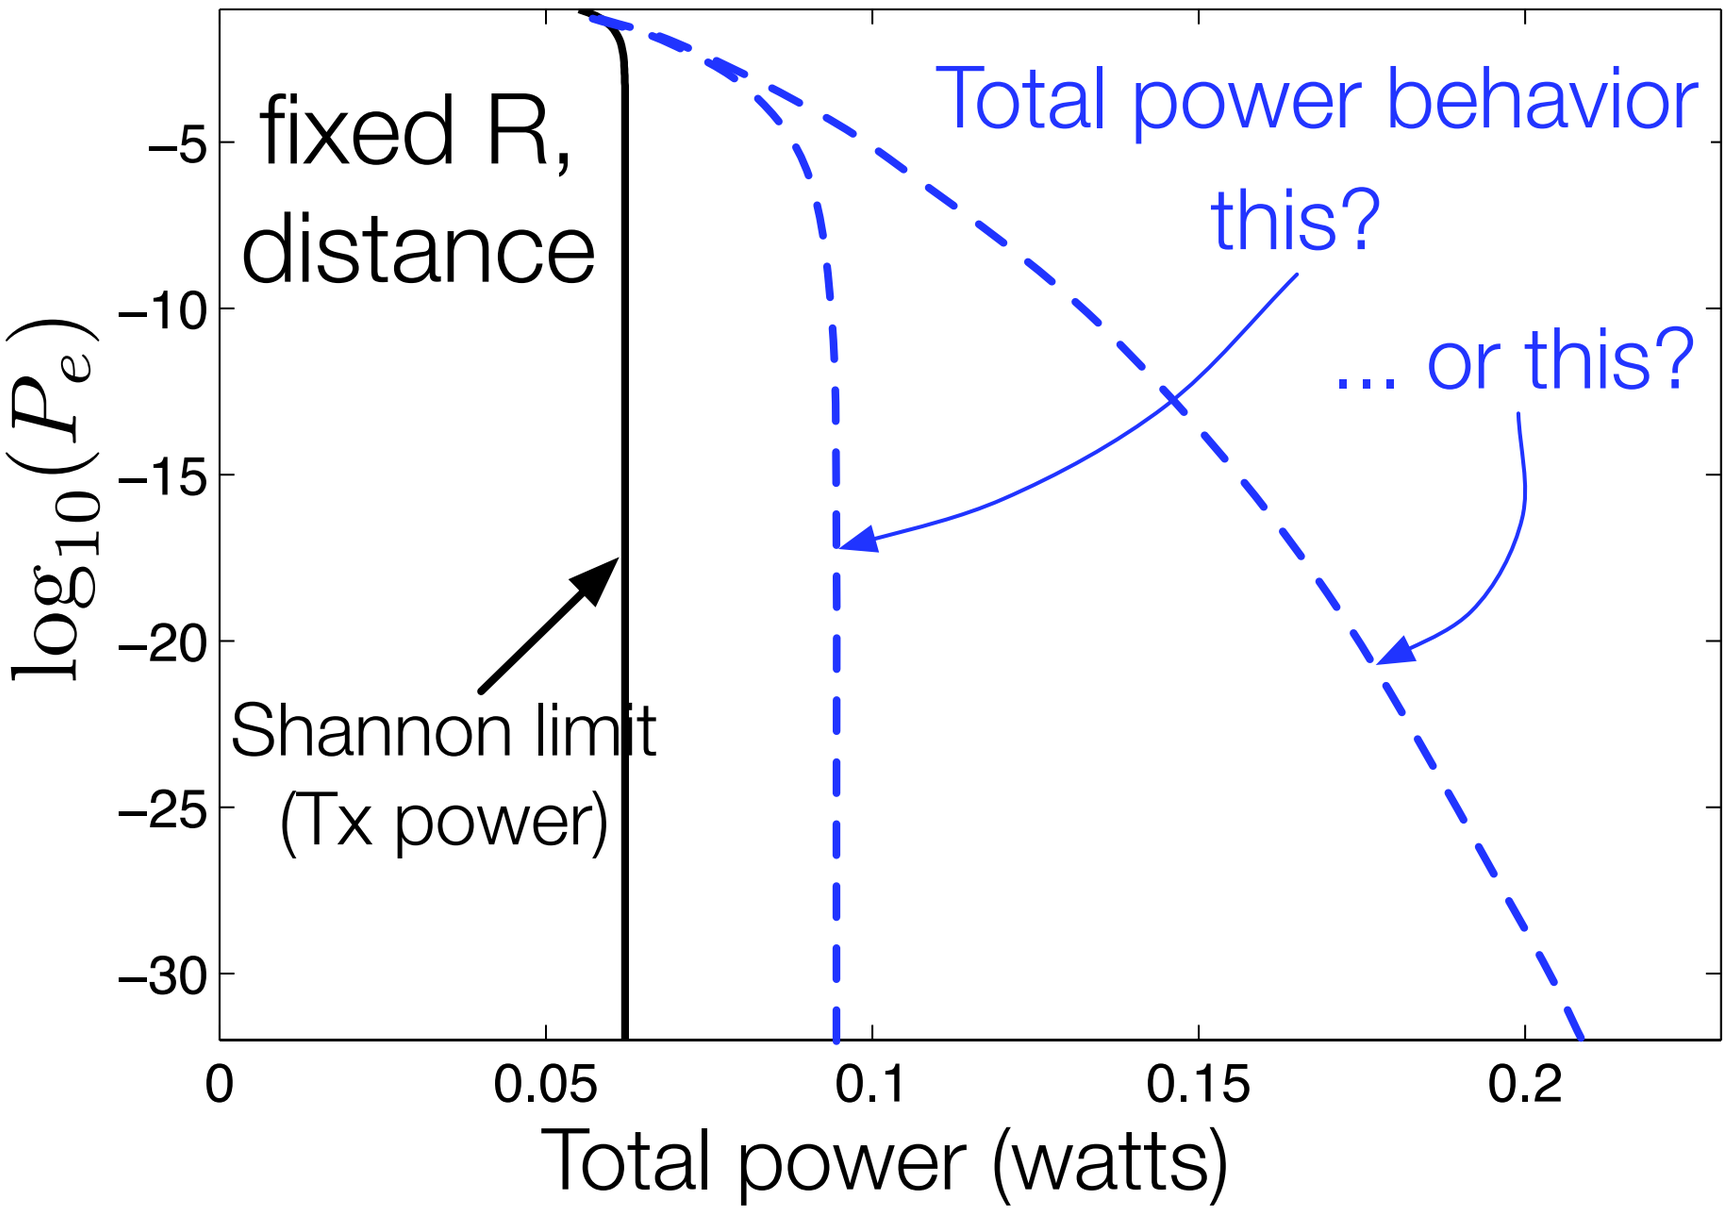

Is attaining Shannon capacity still the goal?

Known fundamental limits

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2} \log \left(1 + \frac{\eta P_T}{N} \right)$$

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2} \log \left(1 + \frac{\eta P_T}{N} \right)$$

Path loss

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2} \log \left(1 + \frac{\eta P_T}{N} \right)$$

Path loss

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2} \log \left(1 + \frac{\eta P_T}{N} \right)$$

Path loss

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Decoding circuits

Since 1947:
Minimize computation power

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2} \log \left(1 + \frac{\eta P_T}{N} \right)$$

Path loss

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Decoding circuits

Since 1947:
Minimize computation power

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2} \log \left(1 + \frac{\eta P_T}{N} \right)$$

Path loss

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Decoding circuits

Since 1947:
Minimize computation power

??

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2}\log\left(1 + \frac{\eta P_T}{N}\right)$$

$$C = 1 - h_b(p)$$

Path loss

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Decoding circuits

Since 1947:
Minimize computation power

??

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Known fundamental limits

Comm

Since 1948:
Minimize transmit power

Fundamental limit:

$$C = \frac{W}{2}\log\left(1 + \frac{\eta P_T}{N}\right)$$

$$C = 1 - h_b(p)$$

$$C = 1 - \epsilon$$

Path loss

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Decoding circuits

Since 1947:
Minimize computation power

??

Absolute fundamental limit: kT ln(2)
[von Neumann, Shannon, Landauer]

Main difficulty in finding fundamental limits to decoding power

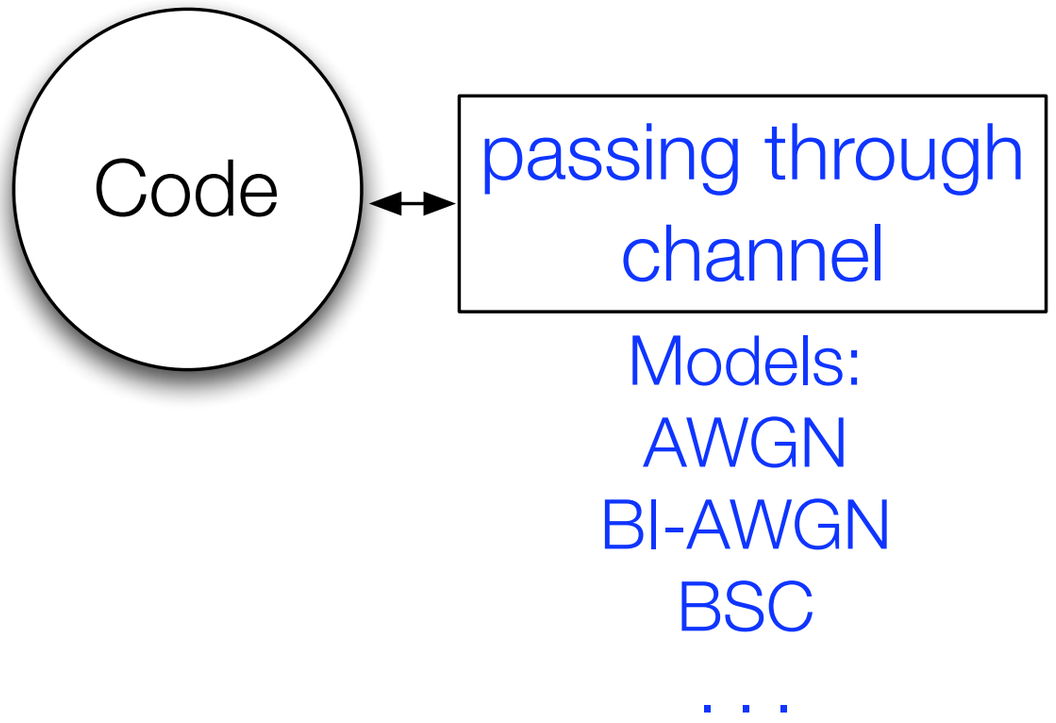

Main difficulty in finding fundamental limits to decoding power

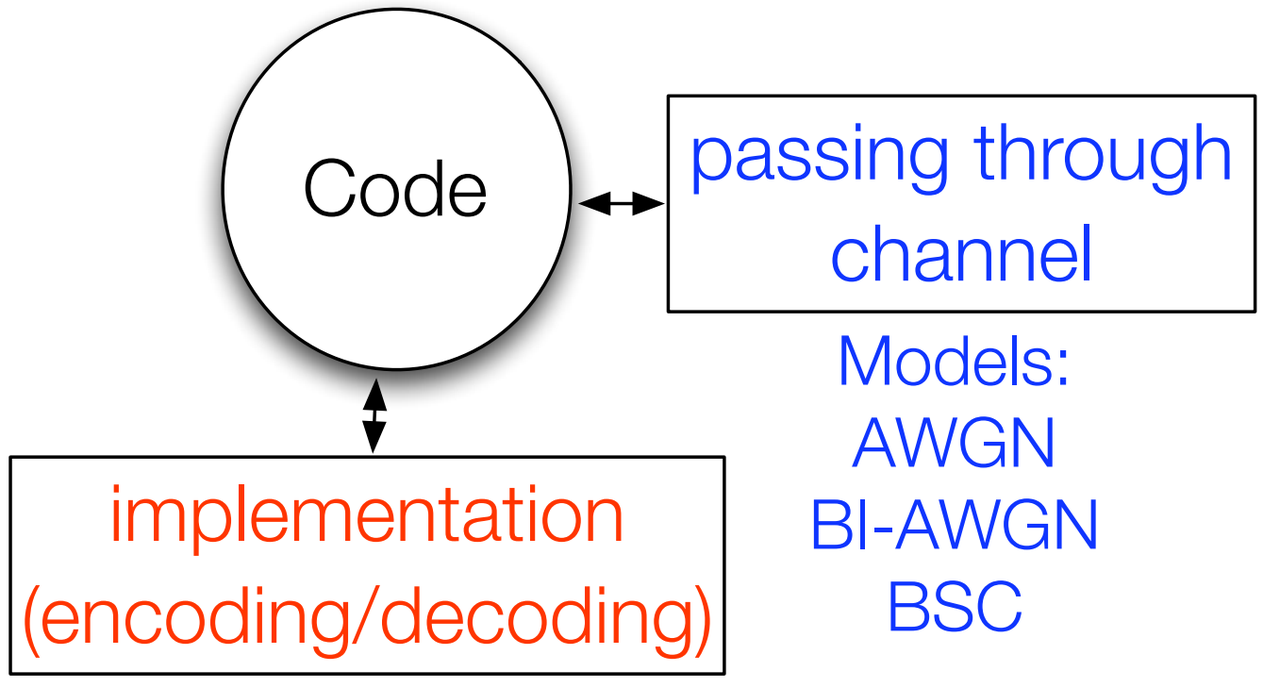

Main difficulty in finding fundamental limits to decoding power

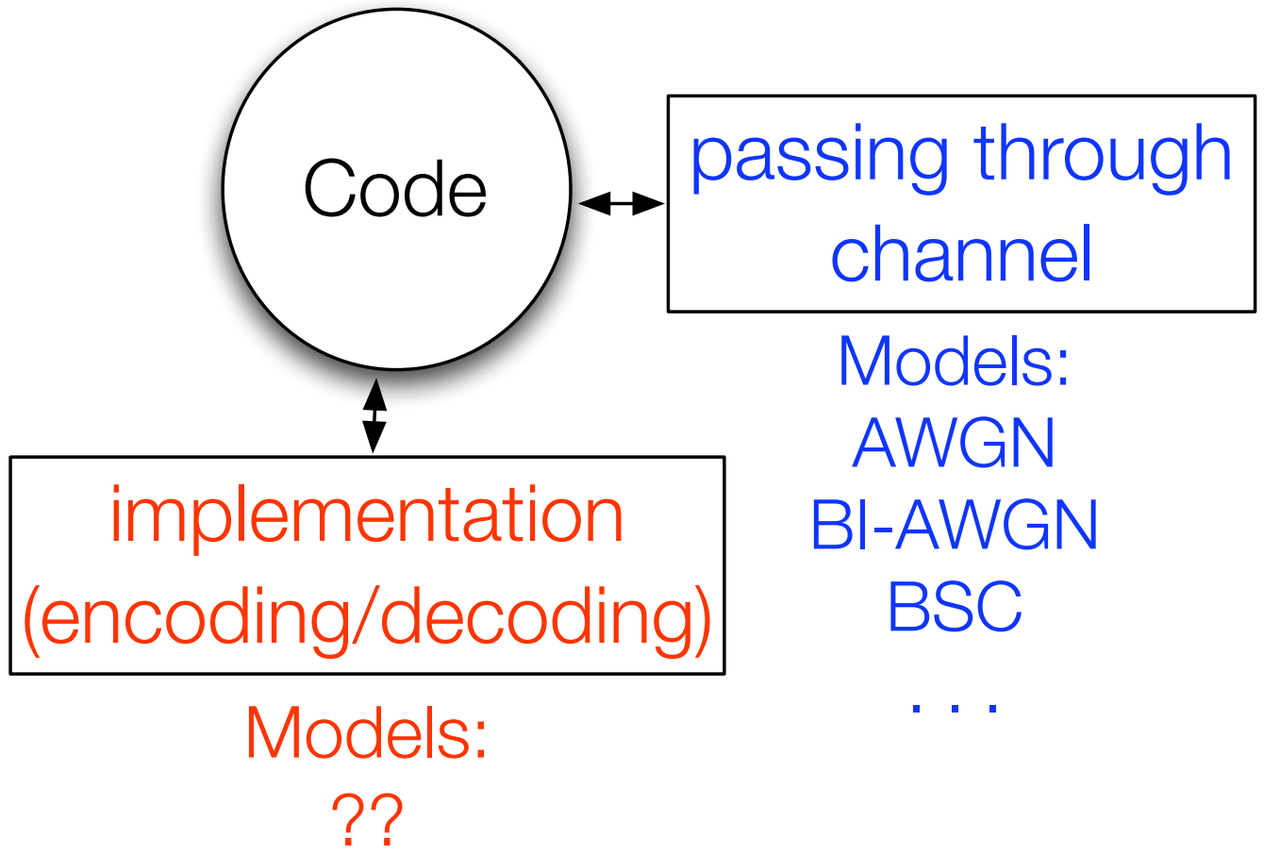

Main difficulty in finding fundamental limits to decoding power

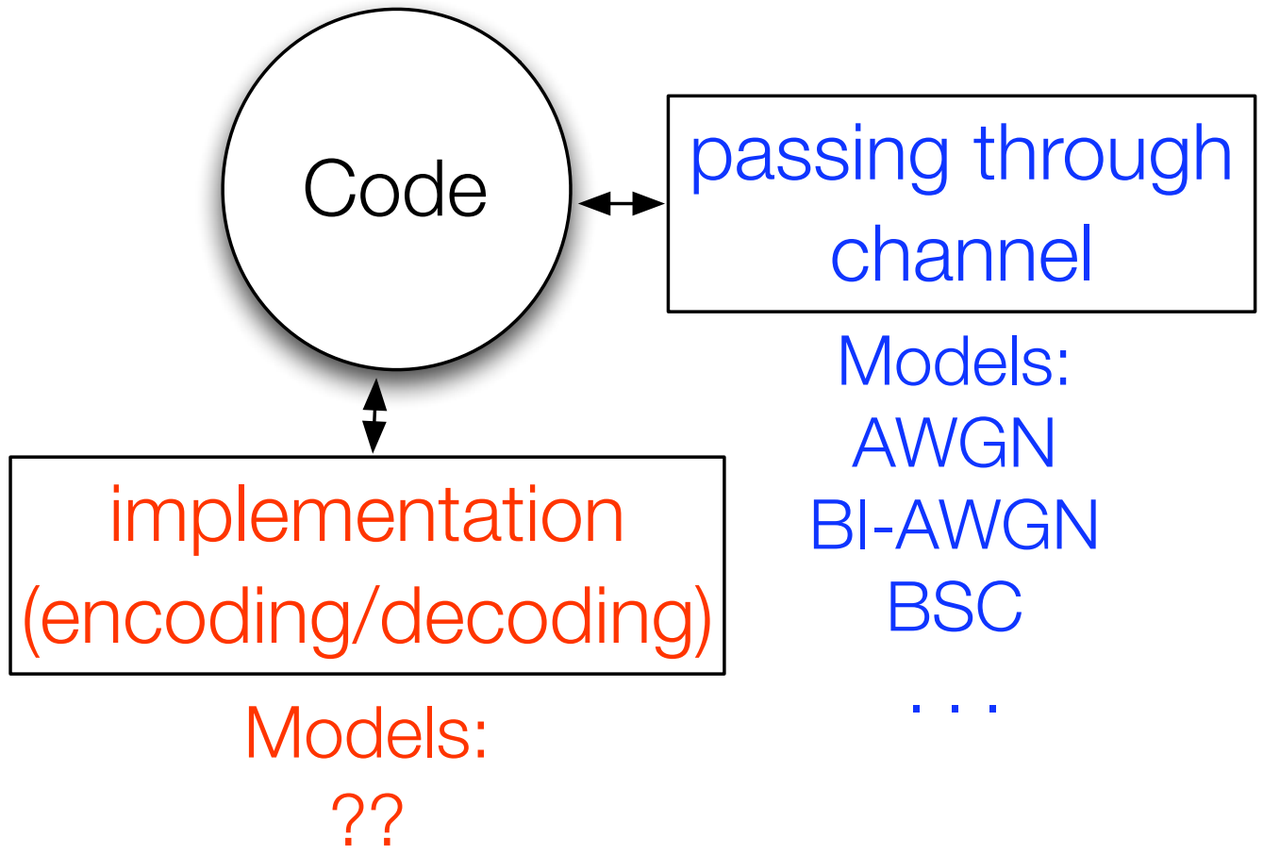

Just as we have channel models to understand transmit power, we need models of decoding implementation to understand decoding power

Talk outline

1. How do we build power-efficient communication systems?

   - understanding processing/decoding power

2. “VLSI model” of decoding computation

3. Decoding power model: “Wire model”

   - design codes for just the required performance (no better!)

4. Can more complex nodes help?

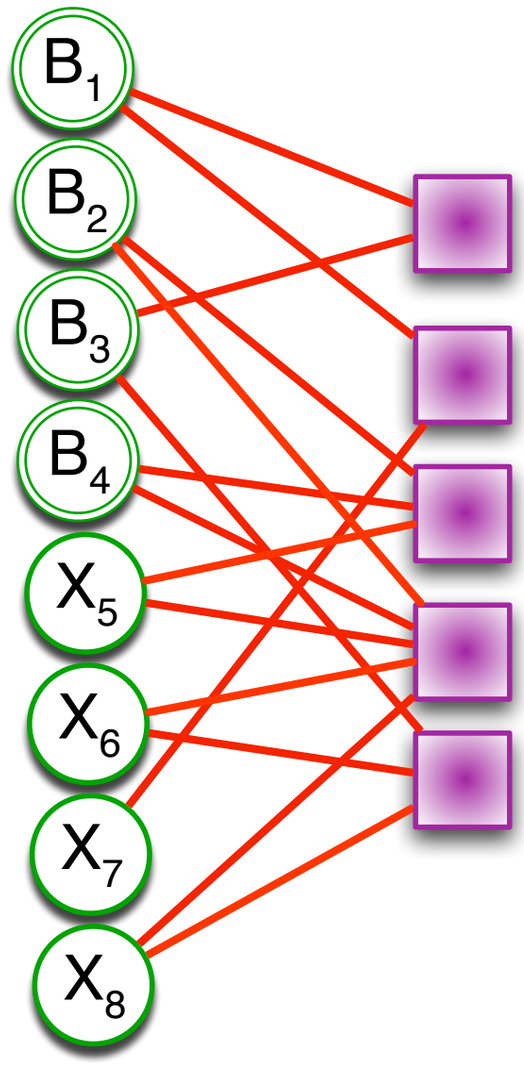

VLSI model of decoding implementation

A message-passing decoder

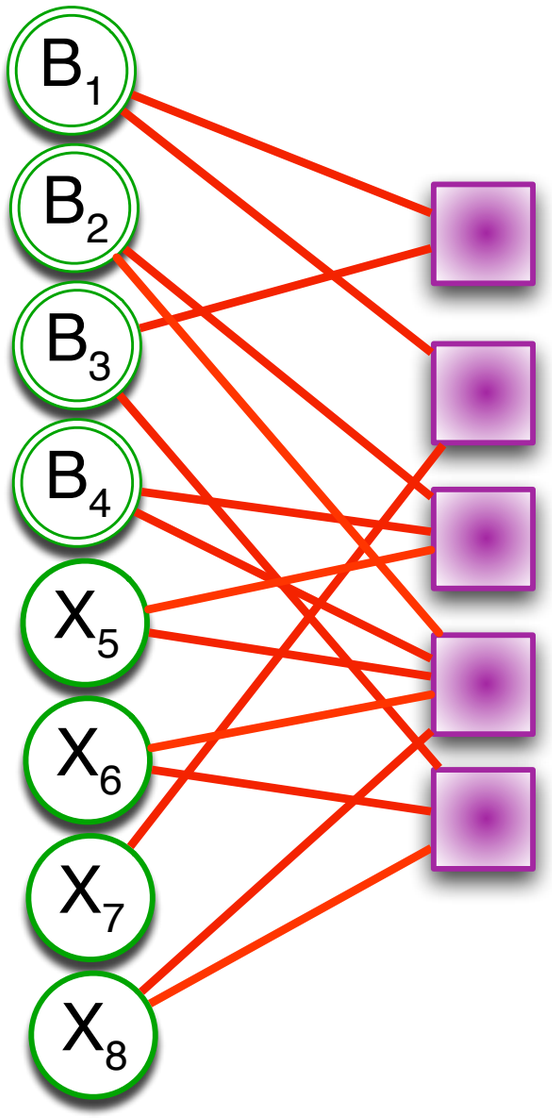

VLSI model of decoding implementation

A message-passing decoder

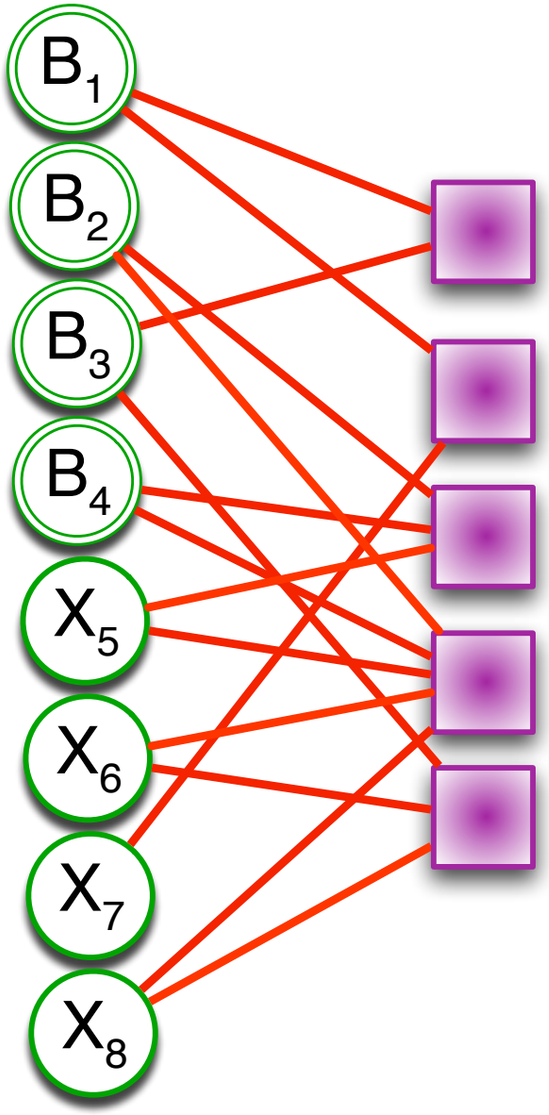

VLSI abstraction of a fully-parallel decoder implementation

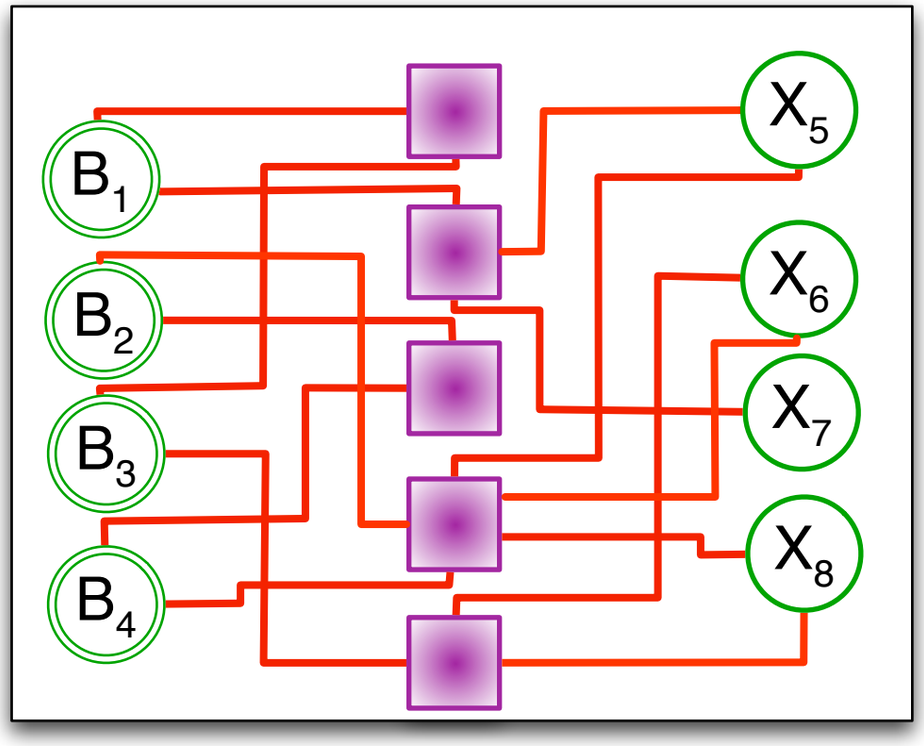

VLSI model of decoding implementation

A message-passing decoder

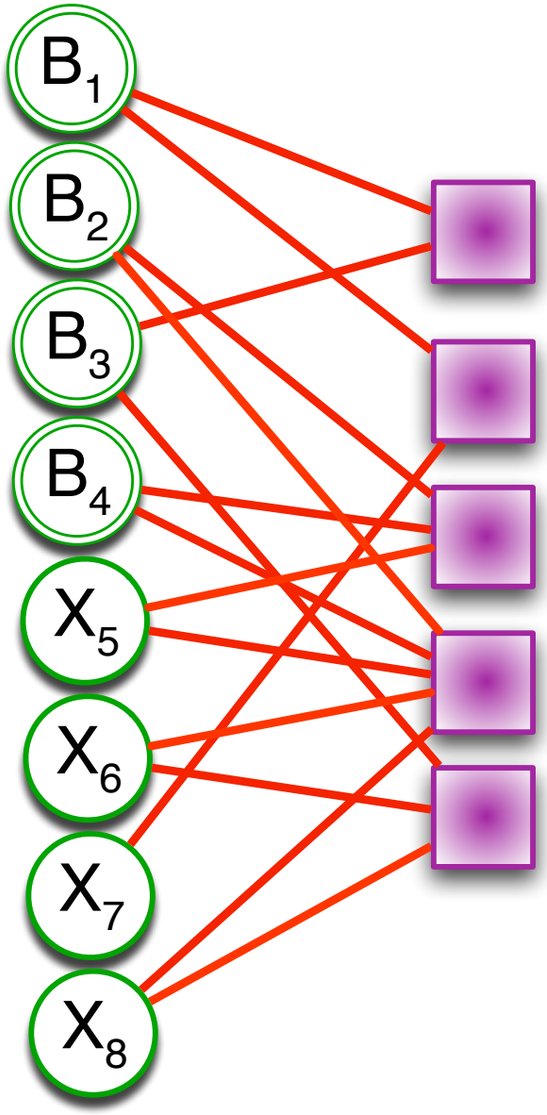

VLSI abstraction of a fully-parallel decoder implementation

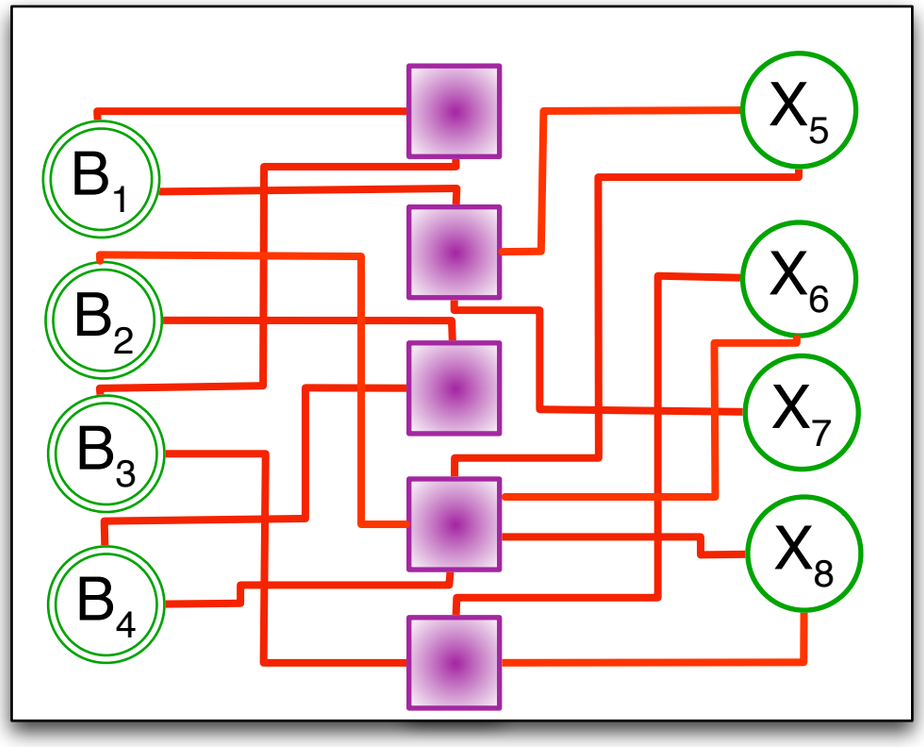

At each “iteration,” each node exchanges messages with all neighbors

Talk outline

1. How do we build power-efficient communication systems?

    - understanding processing/decoding power

2. A “VLSI model” of decoding computation

3. Decoding power model: “Wire model”

    - do not operate too close to capacity!

    - design codes for just the required performance (no better!)

4. Can more complex nodes help?

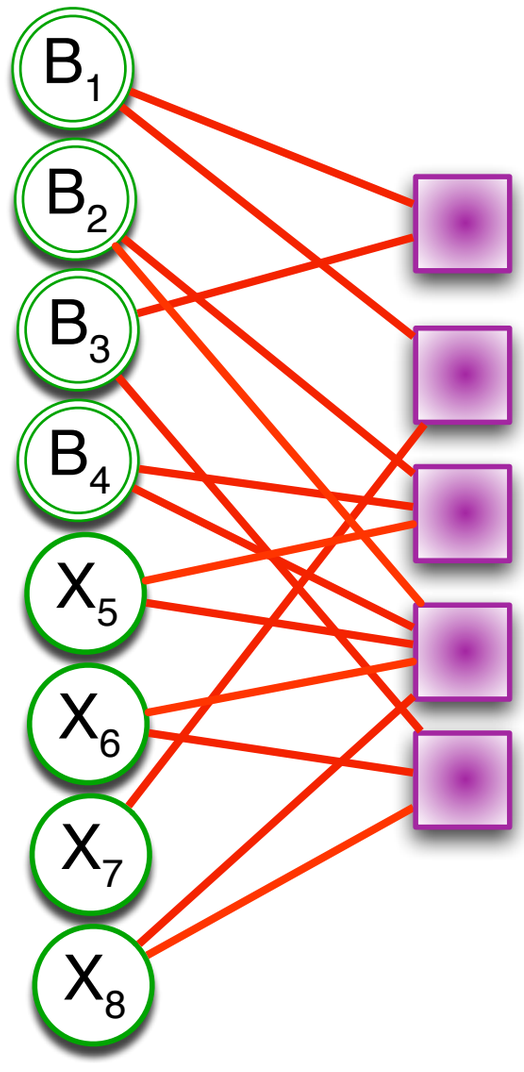

Models of power consumption

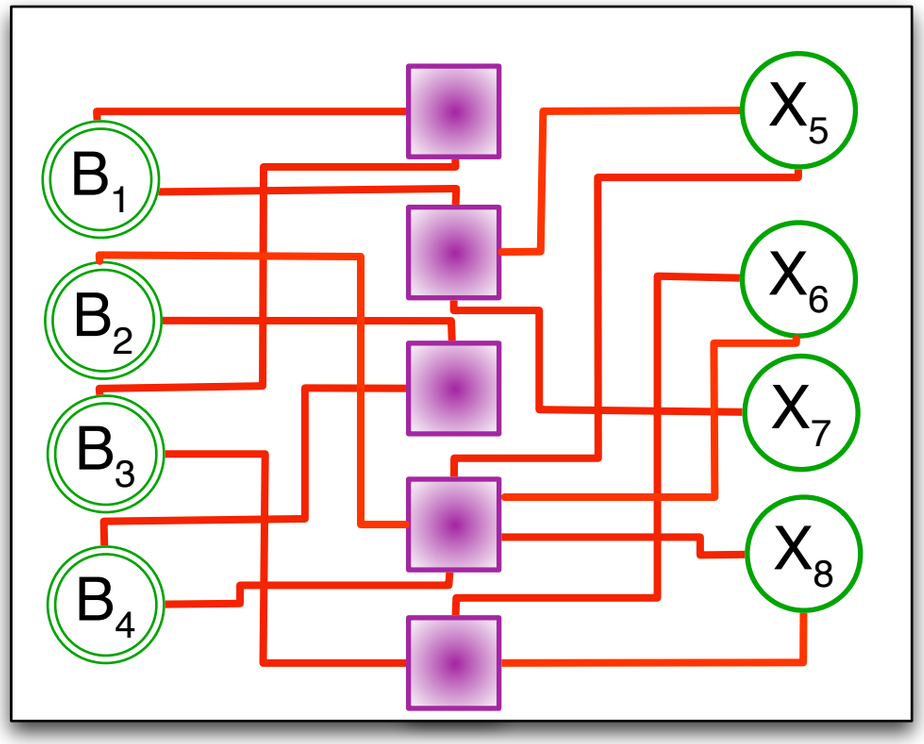

Models of power consumption

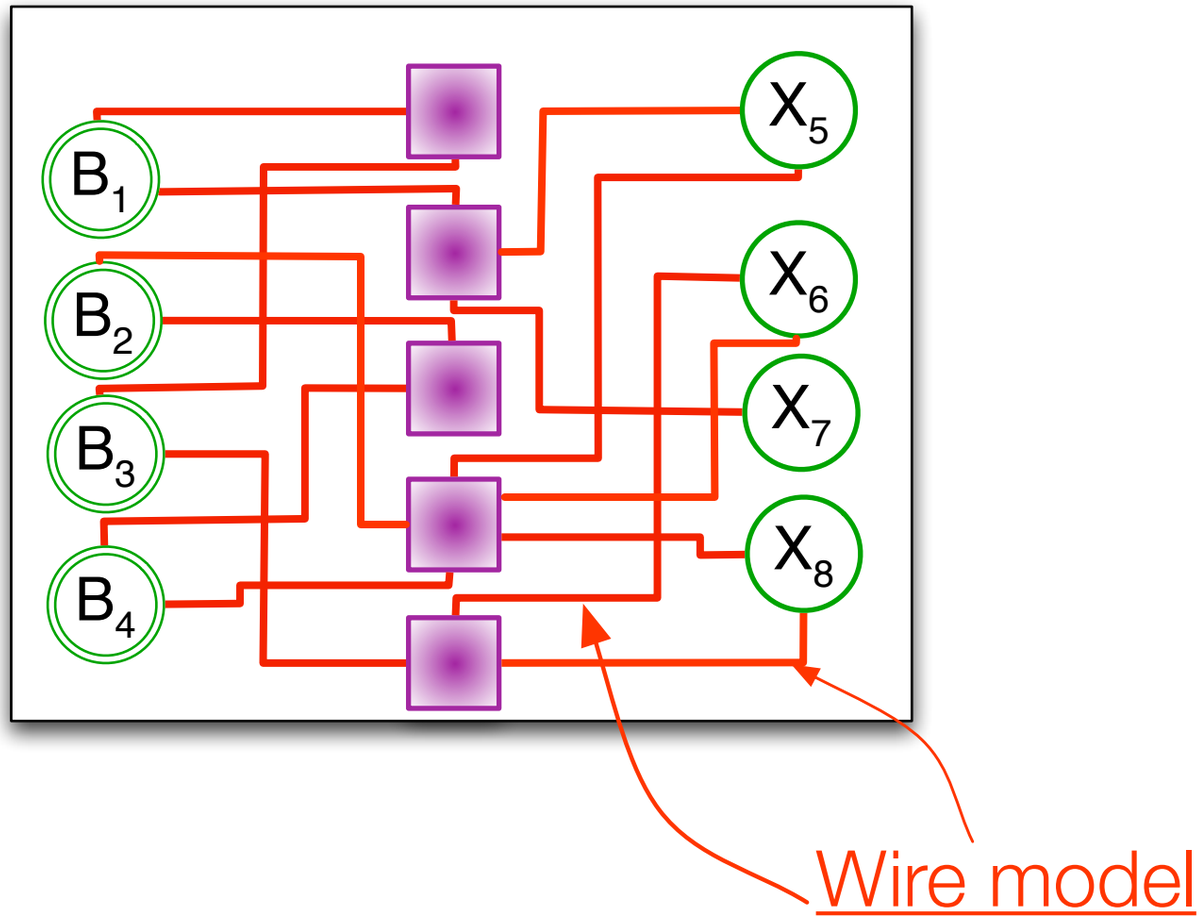

Power consumption increases linearly with the wire-length

Models of power consumption

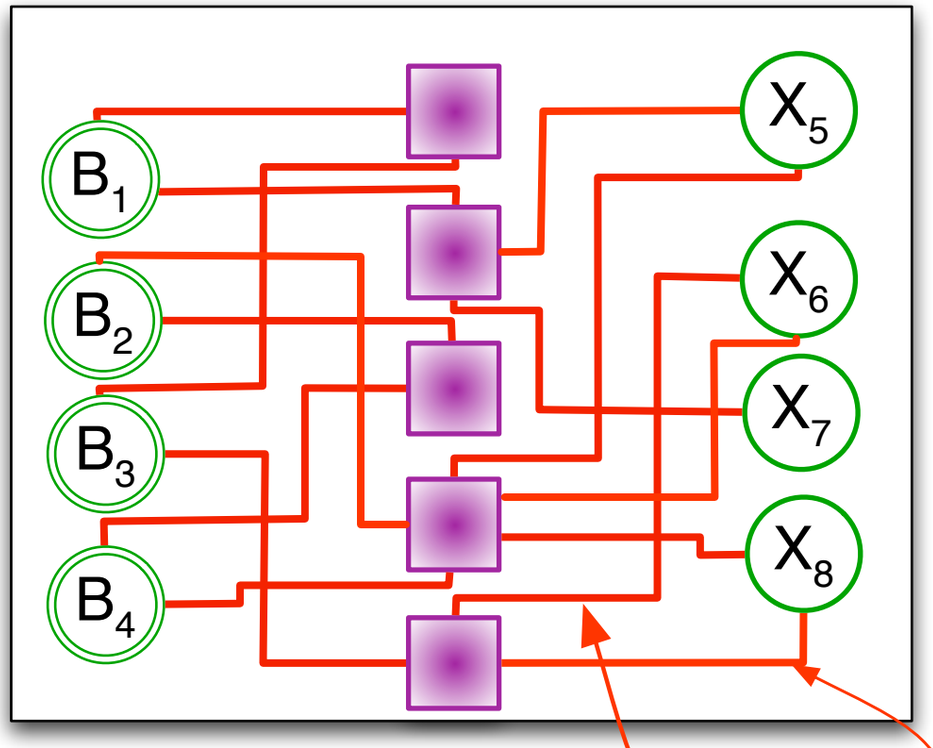

Wire model

Interconnect (wire) power :
dominant sink of dynamic power
microprocessors, FPGAs, ASICs etc.

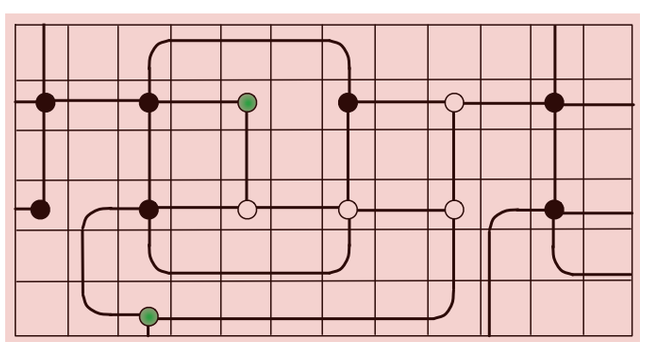

Power consumption increases linearly with the wire-length

Empirical observations

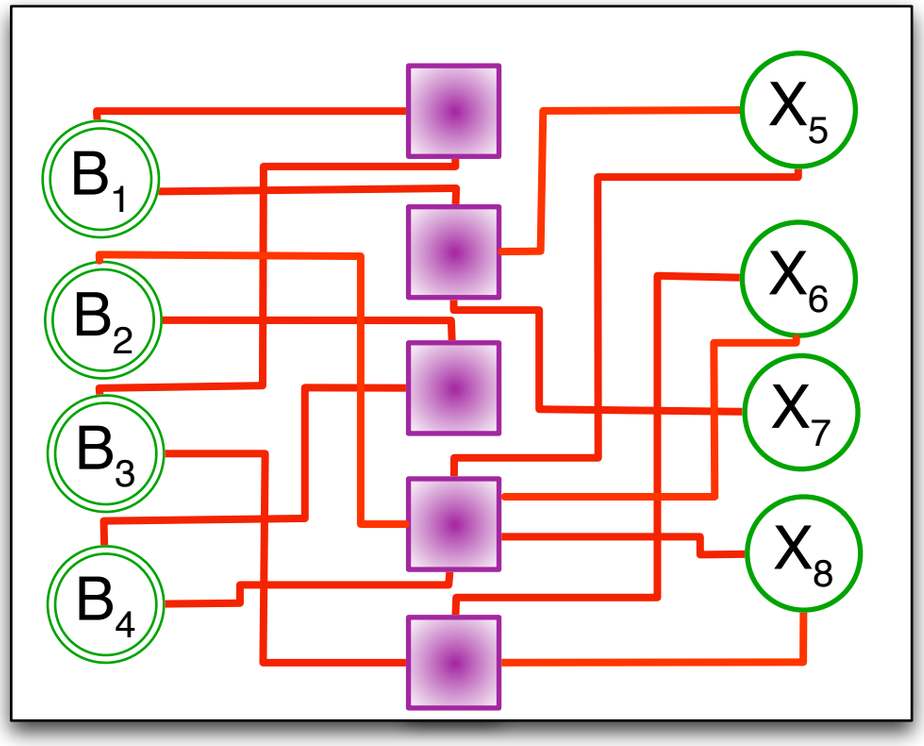

Empirical observations

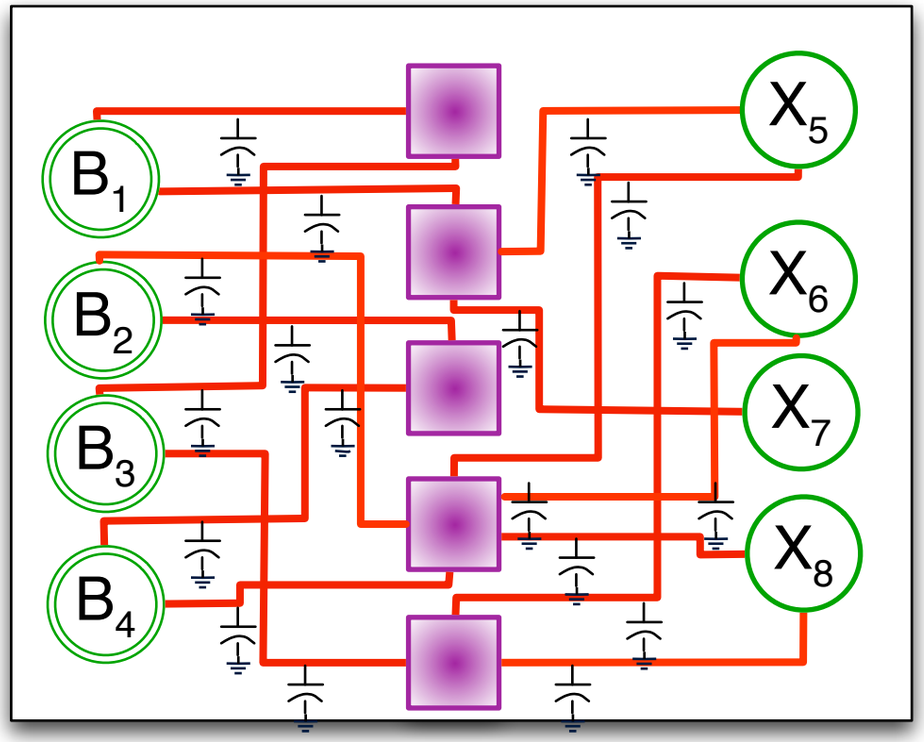

Empirical observations

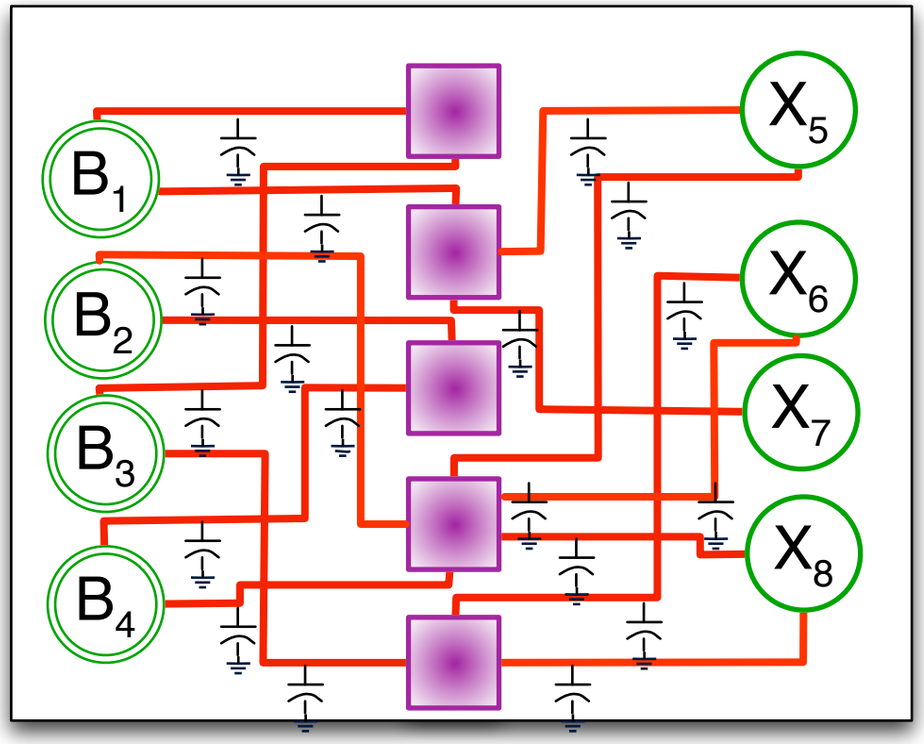

Power consumption increases with wire-length

Empirical observations

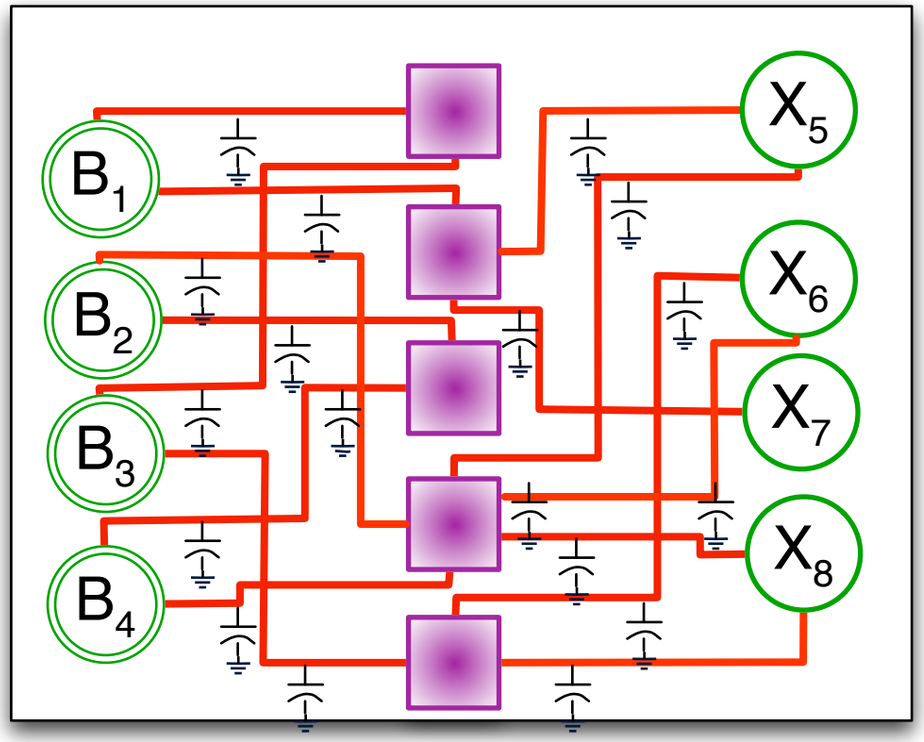

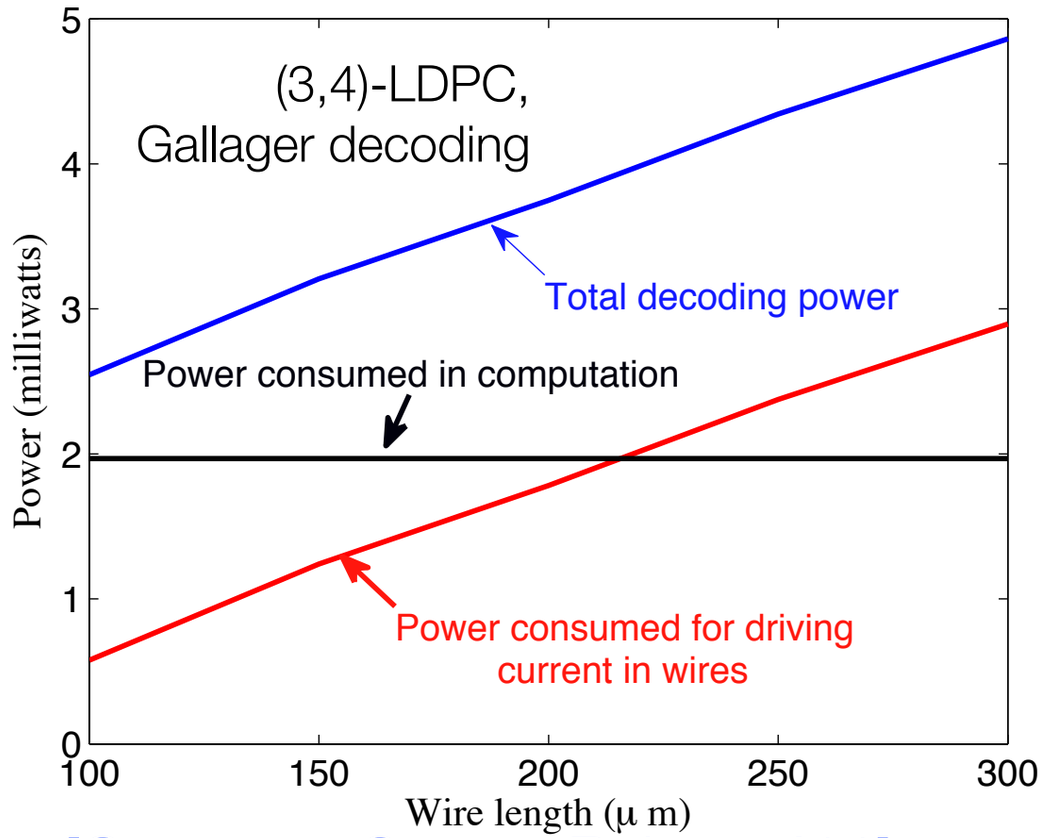

[Ganesan, Grover, Rabaey '11]

Power consumption increases with wire-length

Empirical observations

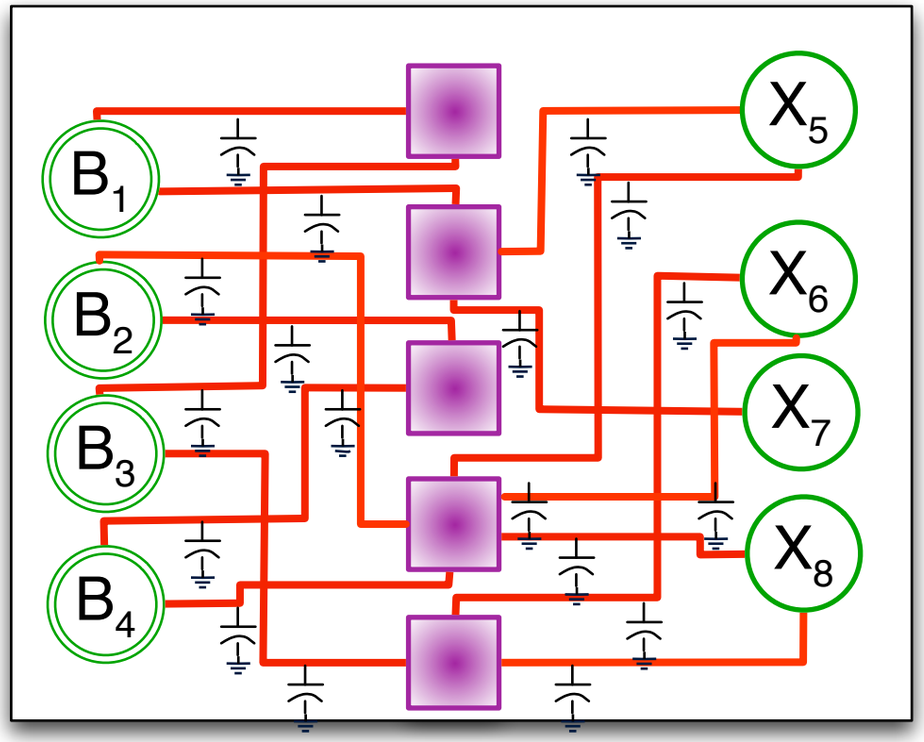

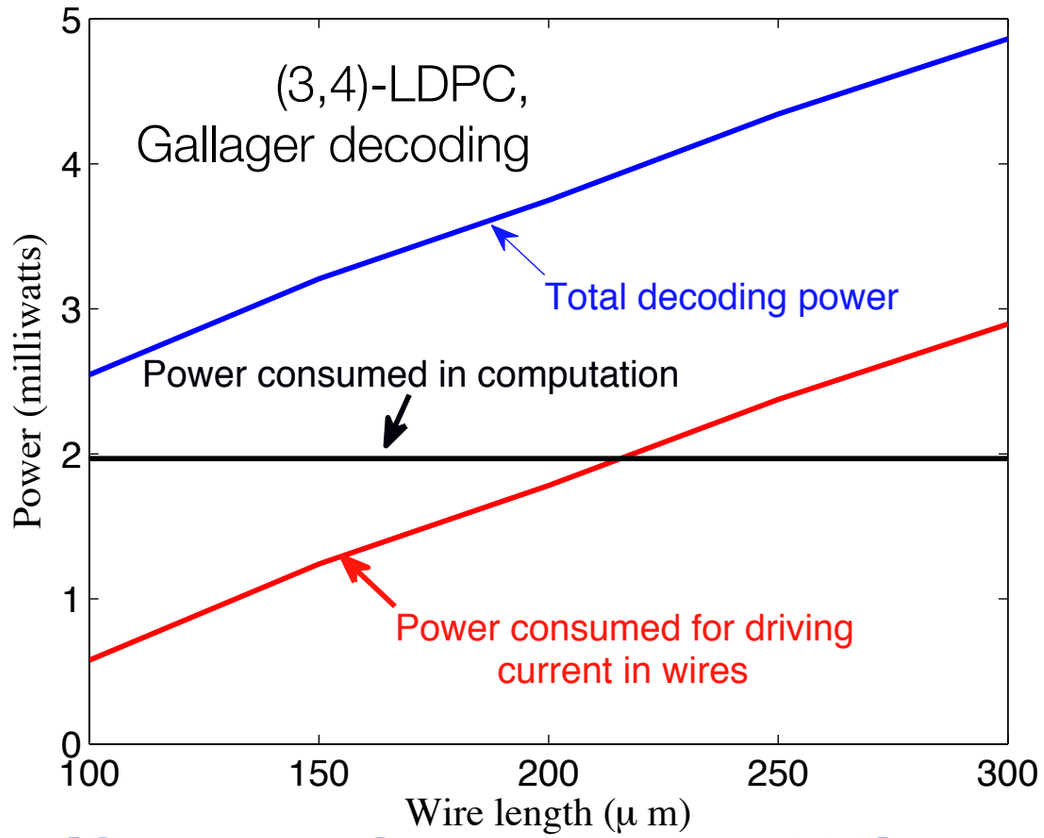

[Ganesan, Grover, Rabaey '11]

Power consumption increases with wire-length

Wire-power > computational node power for simple node computations

Empirical observations

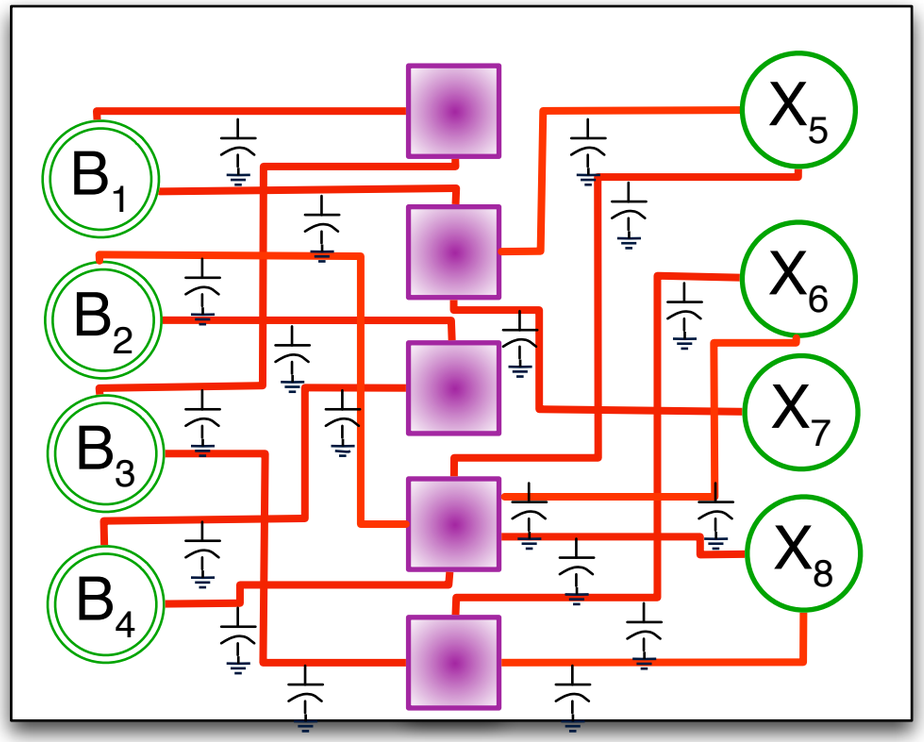

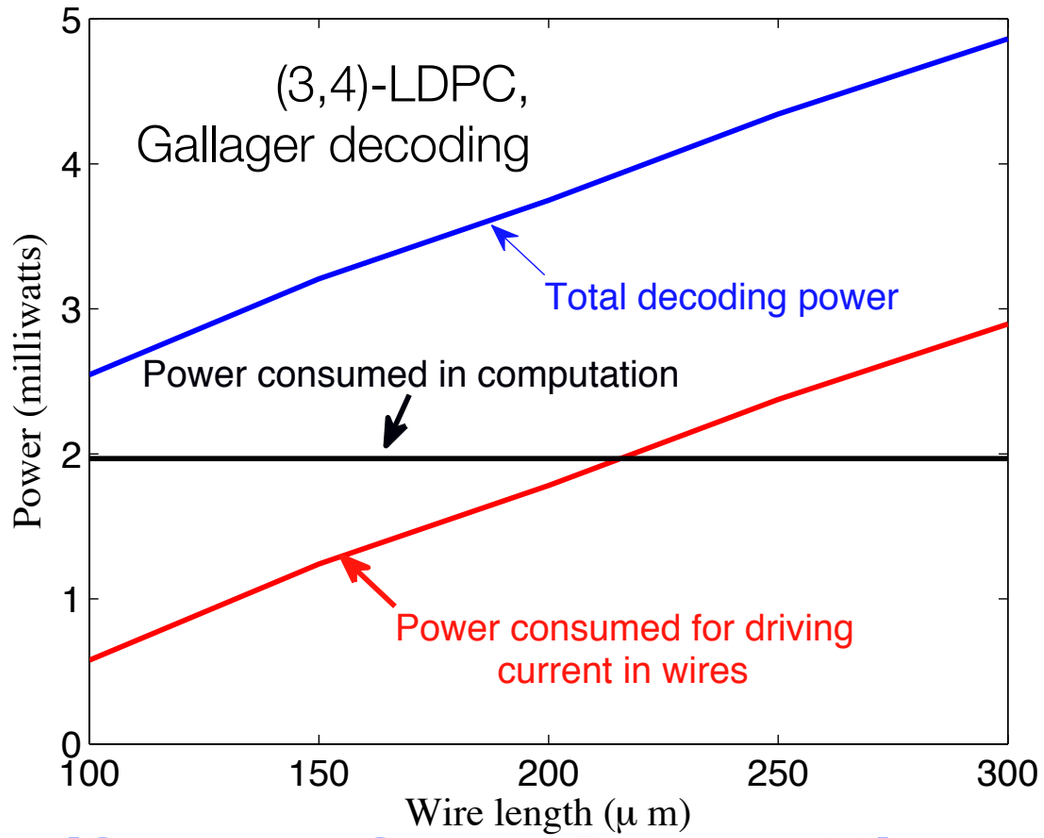

[Ganesan, Grover, Rabaey '11]

Power consumption increases with wire-length

Wire-power > computational node power for simple node computations

Need decoders with short wires!

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

Decoders with short wires: a purely implementation issue?

$\mathcal{B}$

Example:
(3,4) LDPC

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

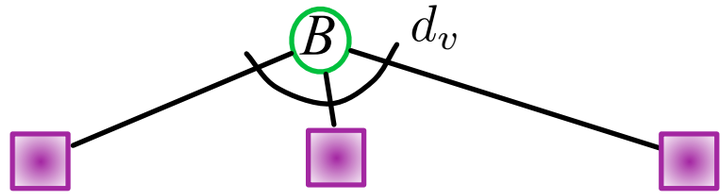

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

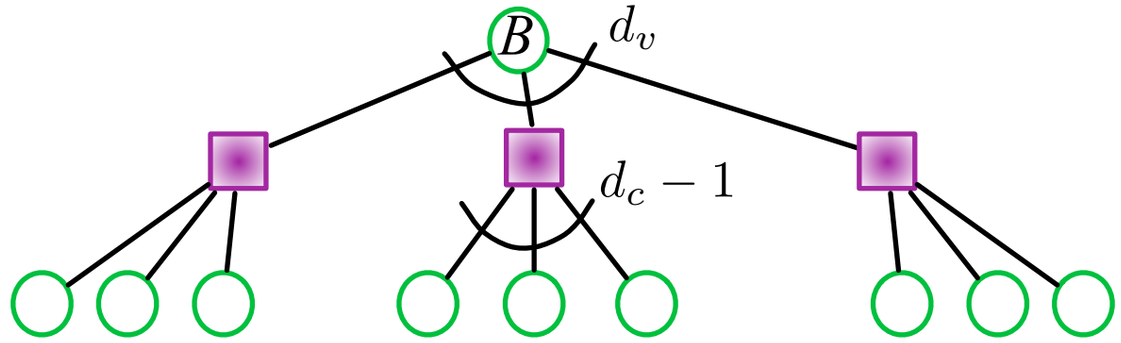

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

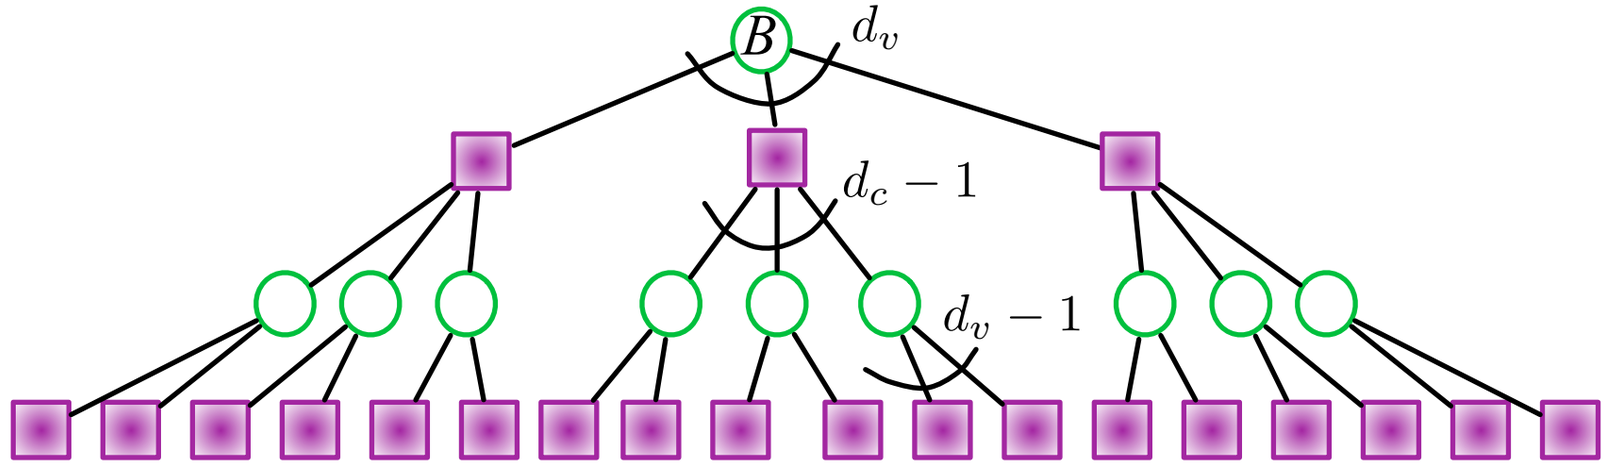

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

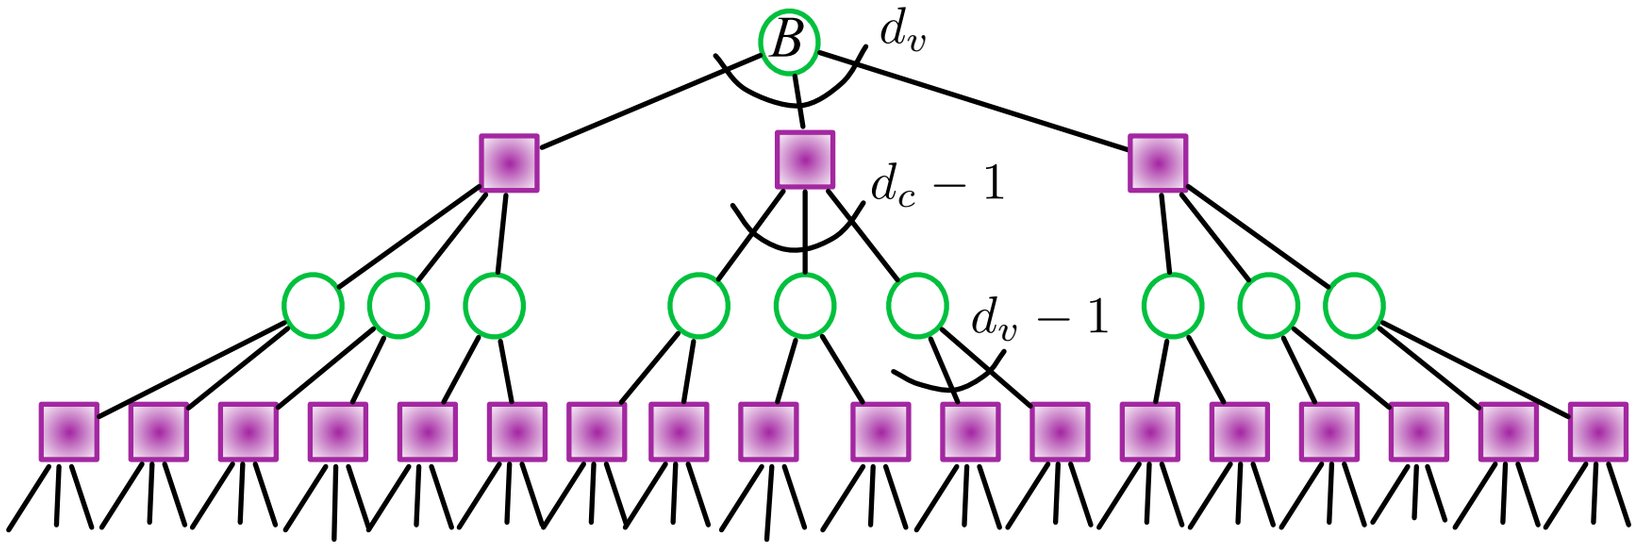

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

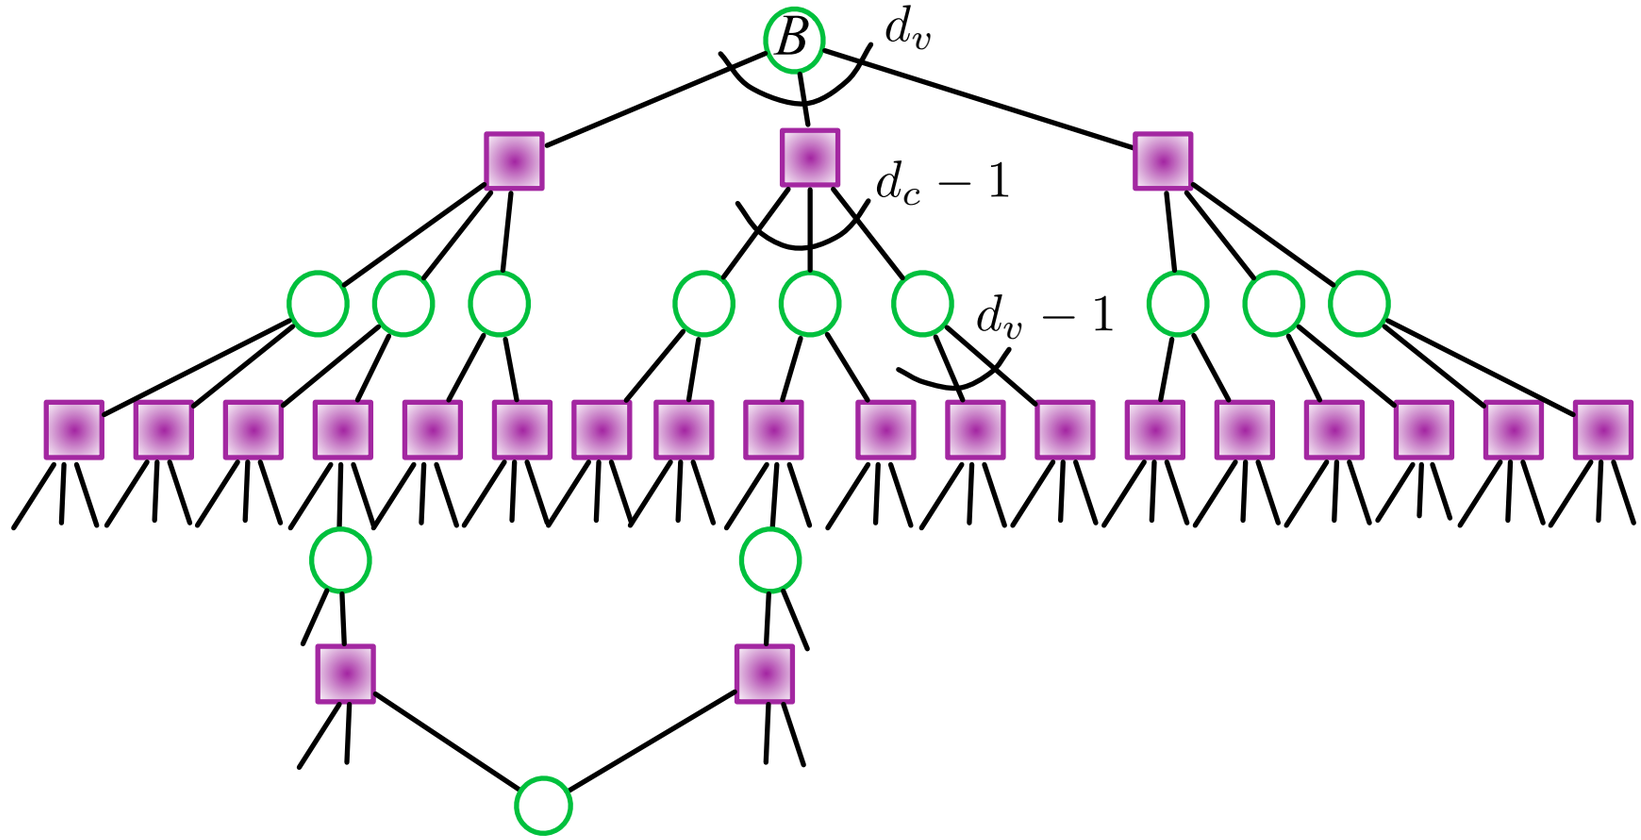

Decoders with short wires: a purely implementation issue?

Example:
(3,4) LDPC

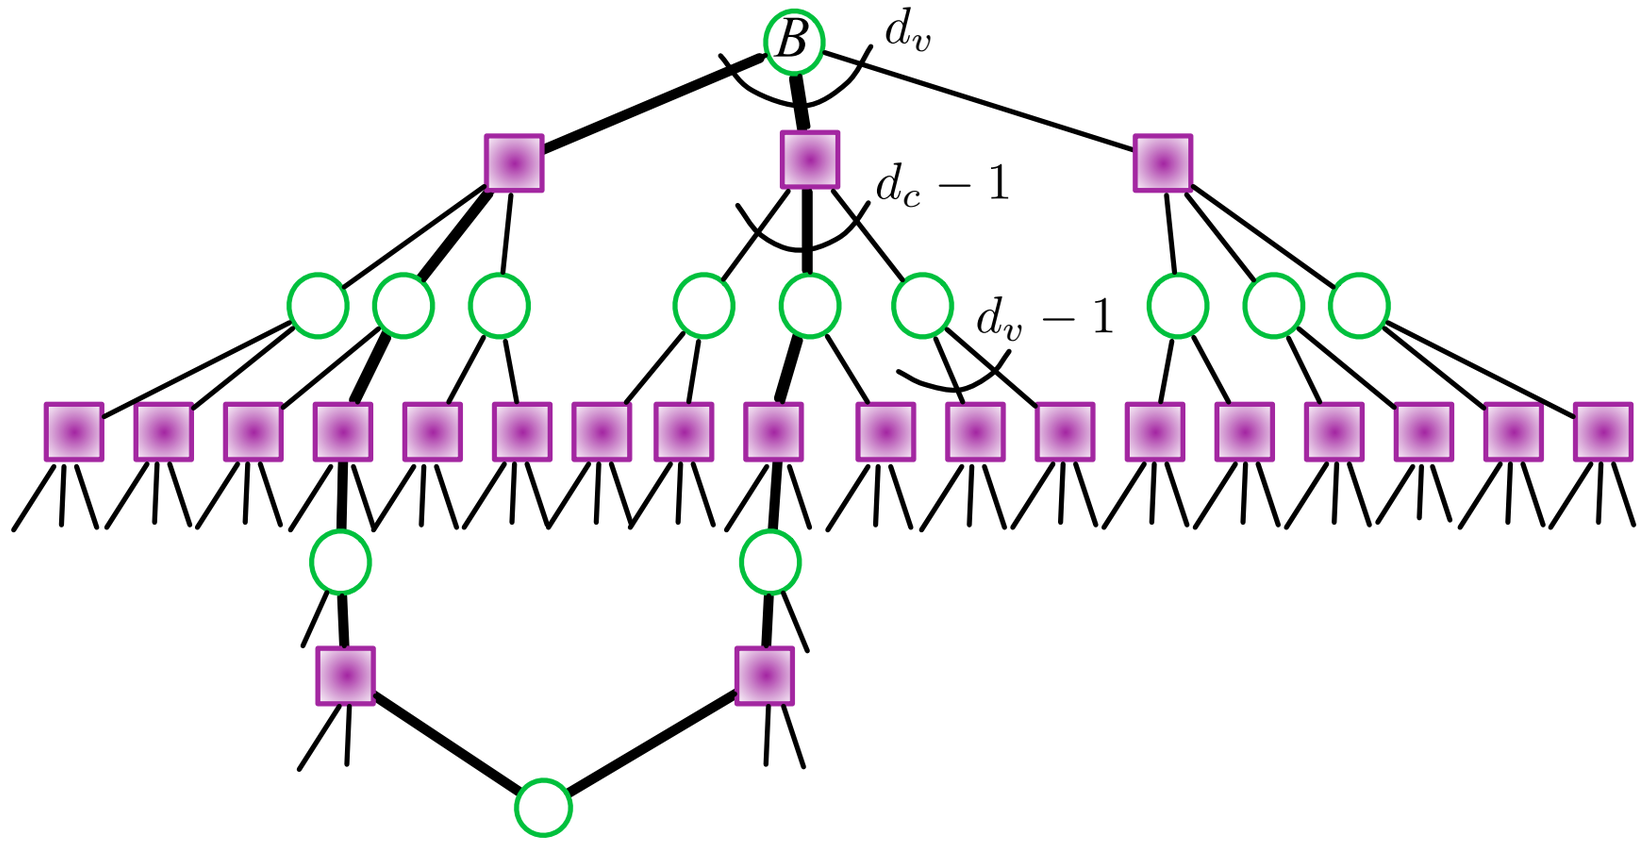

Decoders with short wires: a purely implementation issue?

Example: (3,4) LDPC

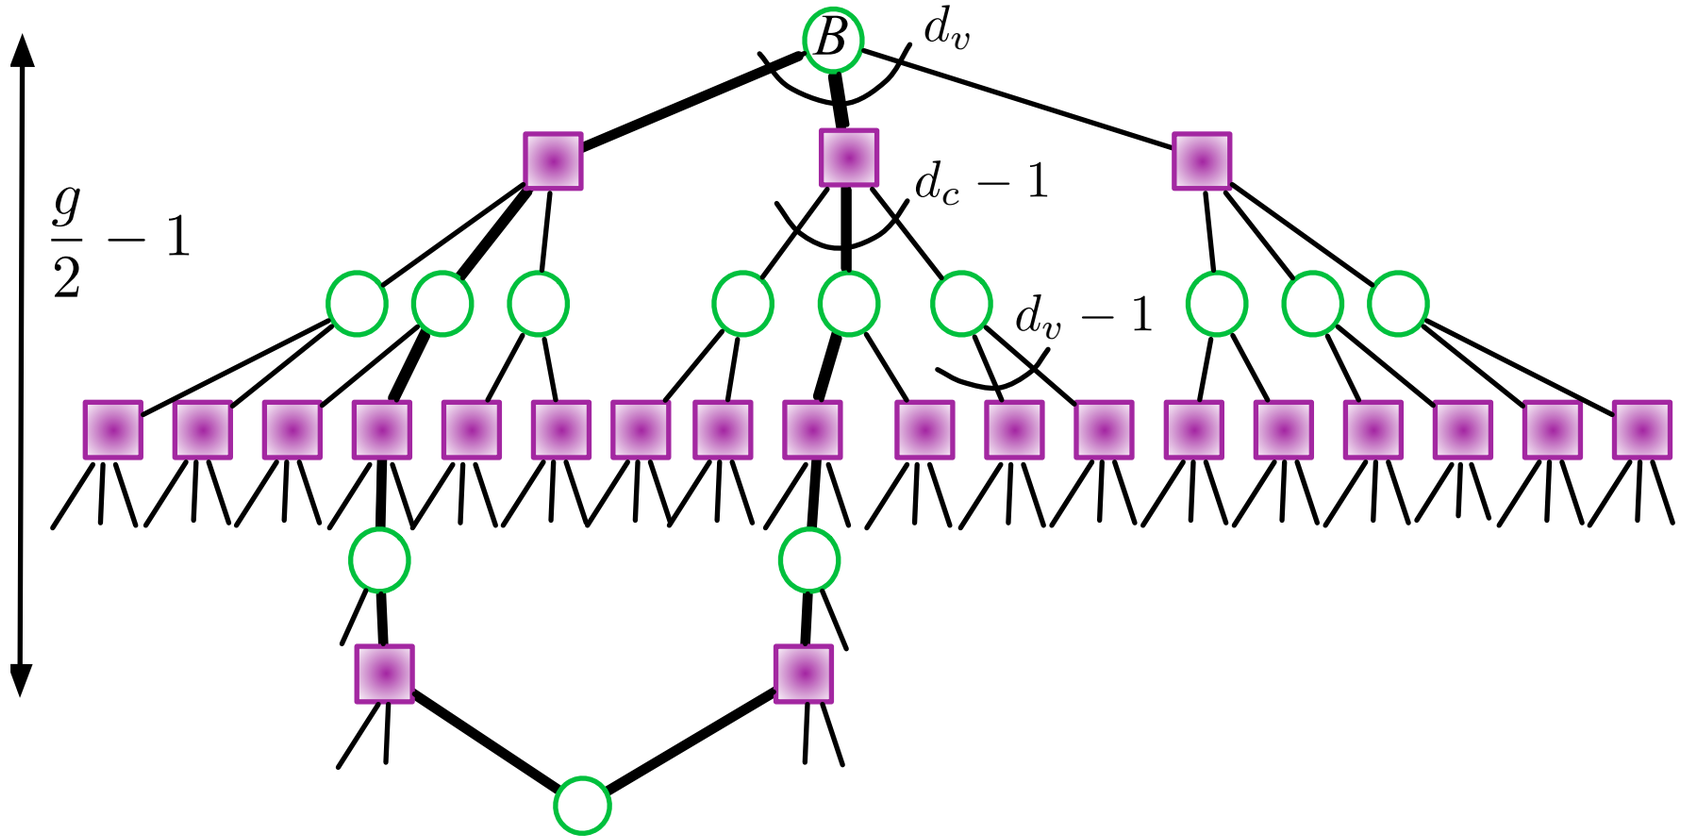

Decoders with short wires: a purely implementation issue?

Example: (3,4) LDPC

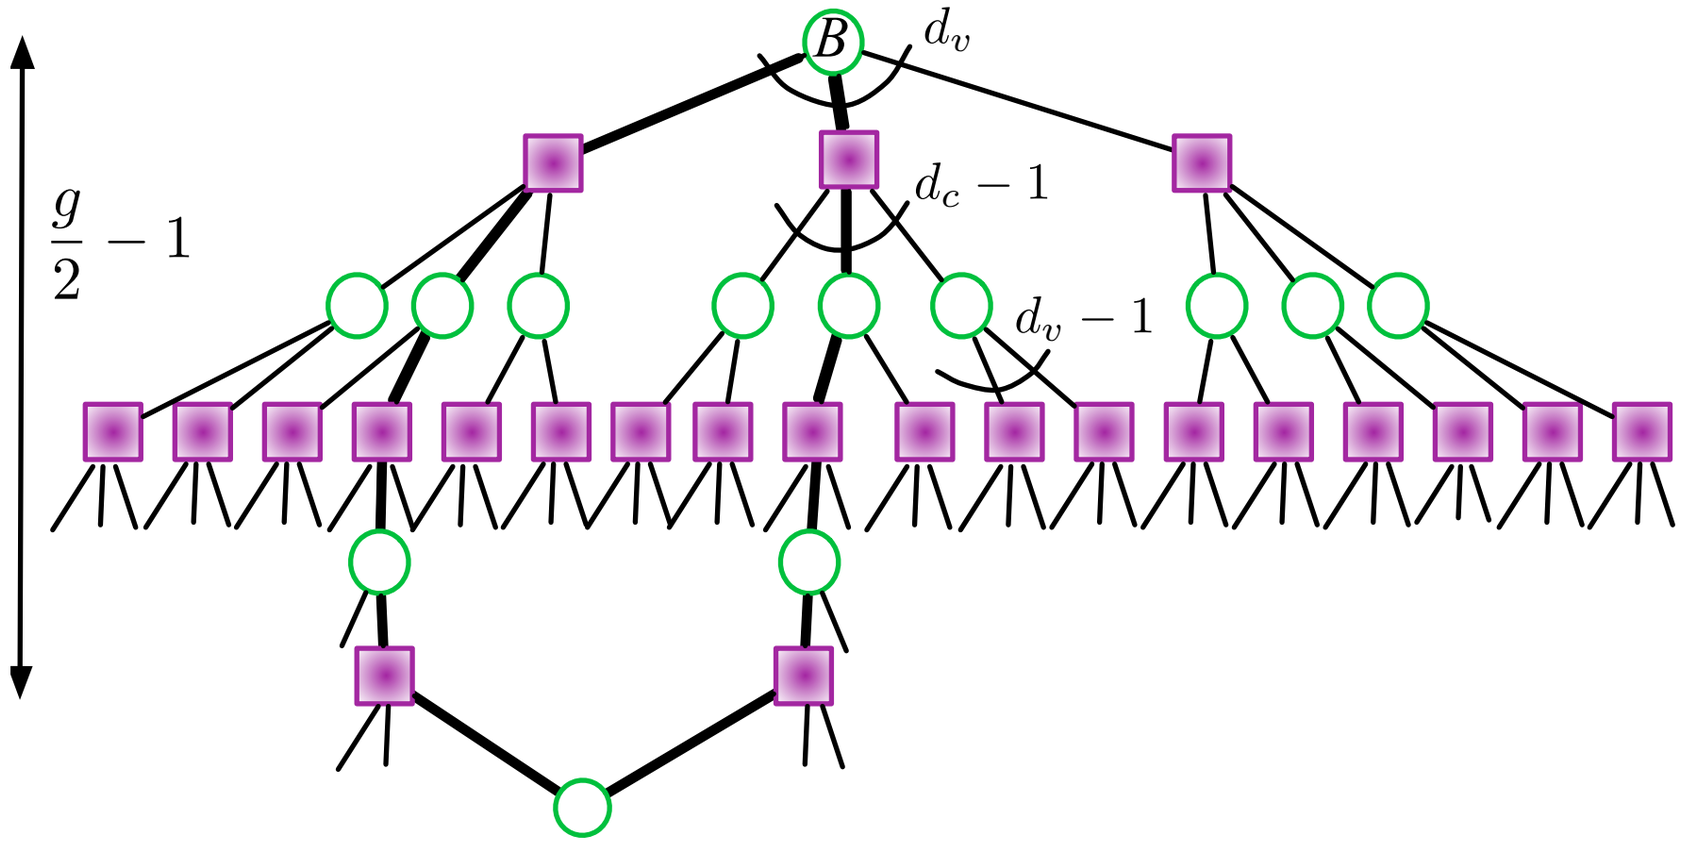

Neighborhood size: exponential in girth
Area: square of (longest wire x girth)

Decoders with short wires: a purely implementation issue?

Example: (3,4) LDPC

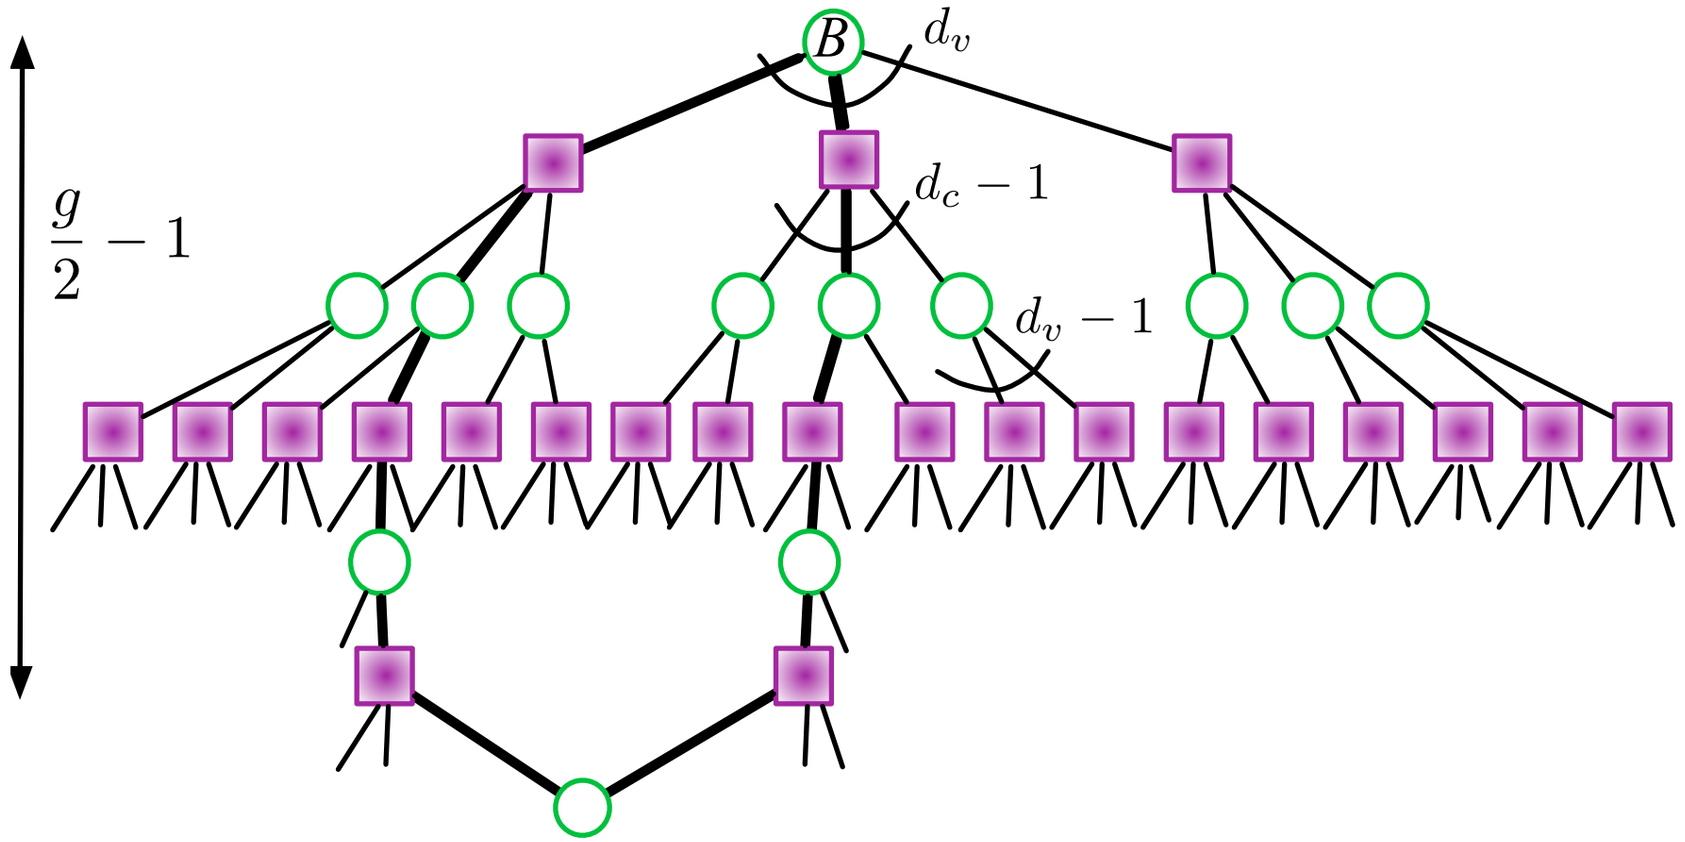

Neighborhood size: exponential in girth
Area: square of (longest wire x girth)

} Longest wire exponential in girth

Decoders with short wires: a purely implementation issue?

Example: (3,4) LDPC

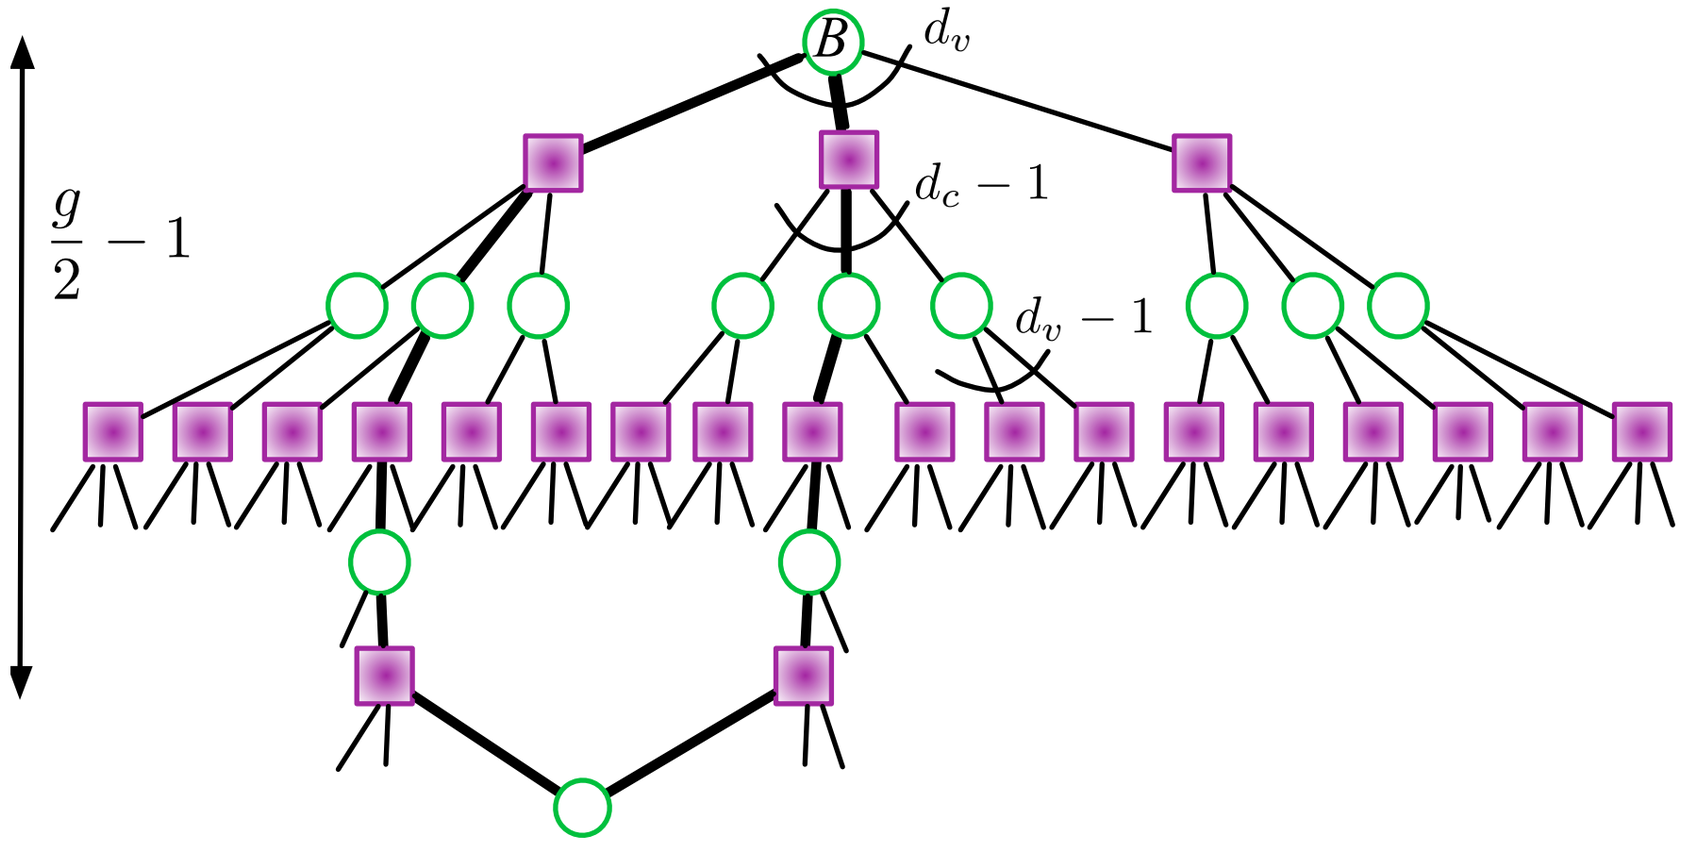

Neighborhood size: exponential in girth
Area: square of (longest wire x girth)

} Longest wire exponential in girth

Theorem [Grover, Sahai '11]

$$W_{\max} \geq \frac{\sqrt{A_{node}}}{\sqrt{\pi}\left(\frac{g}{2}-1\right)} \left((d_v-1)^{\lfloor \frac{g}{4}-\frac{1}{2} \rfloor}(d_c-1)^{\lfloor \frac{g}{4}-\frac{3}{4} \rfloor}\right)$$

Fundamental bounds on interconnect lengths

Fundamental bounds on interconnect lengths

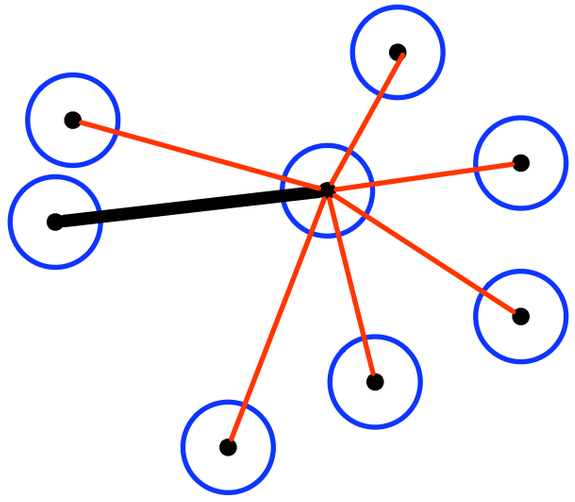

Fundamental bounds on interconnect lengths

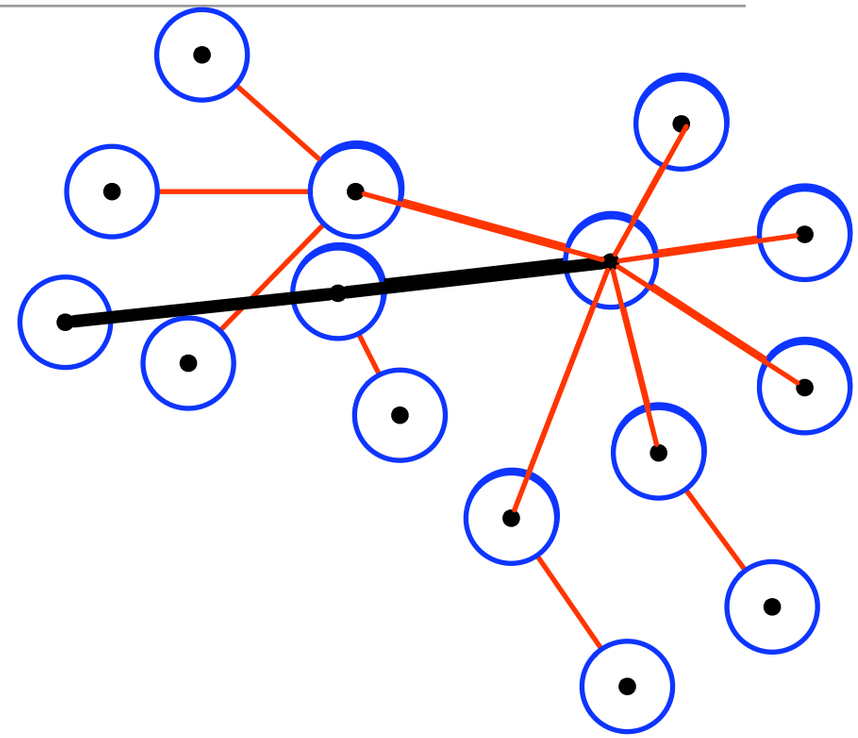

Fundamental bounds on interconnect lengths

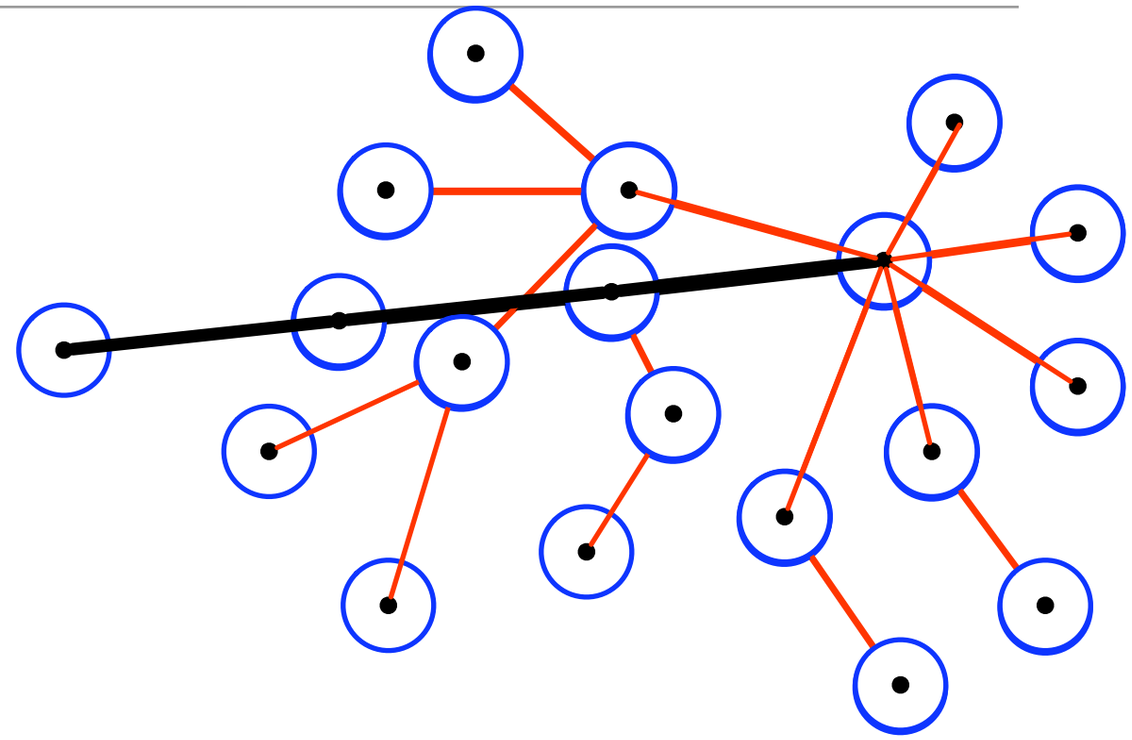

Fundamental bounds on interconnect lengths

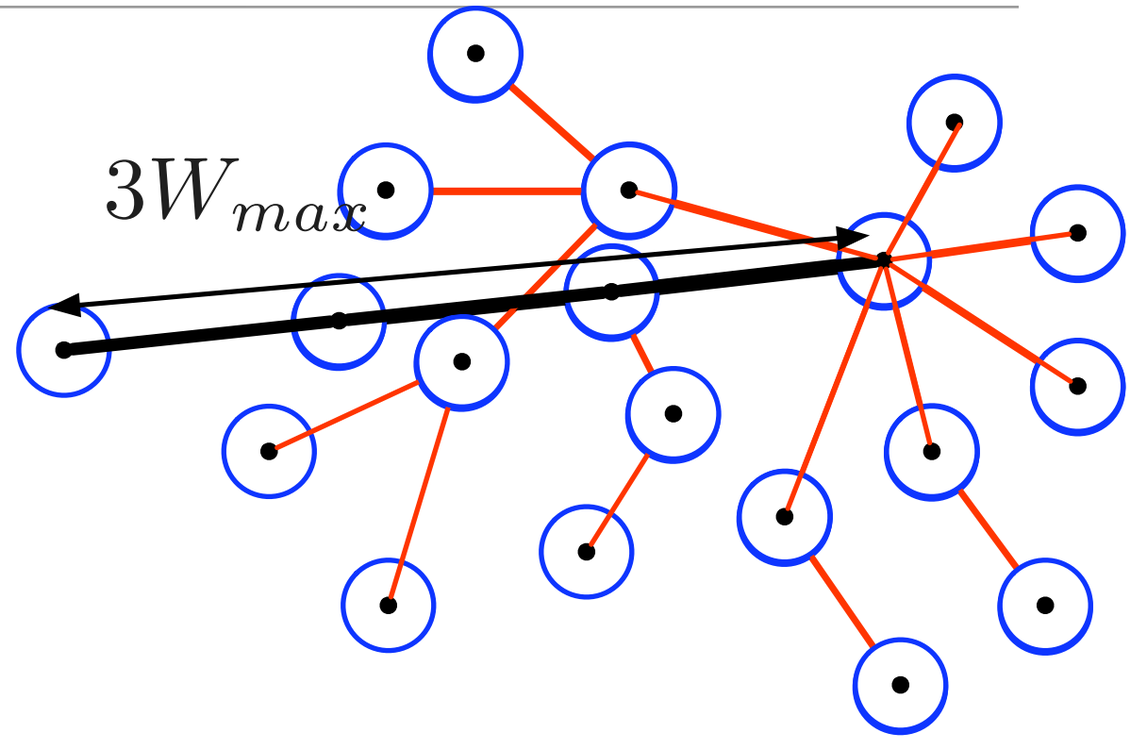

Fundamental bounds on interconnect lengths

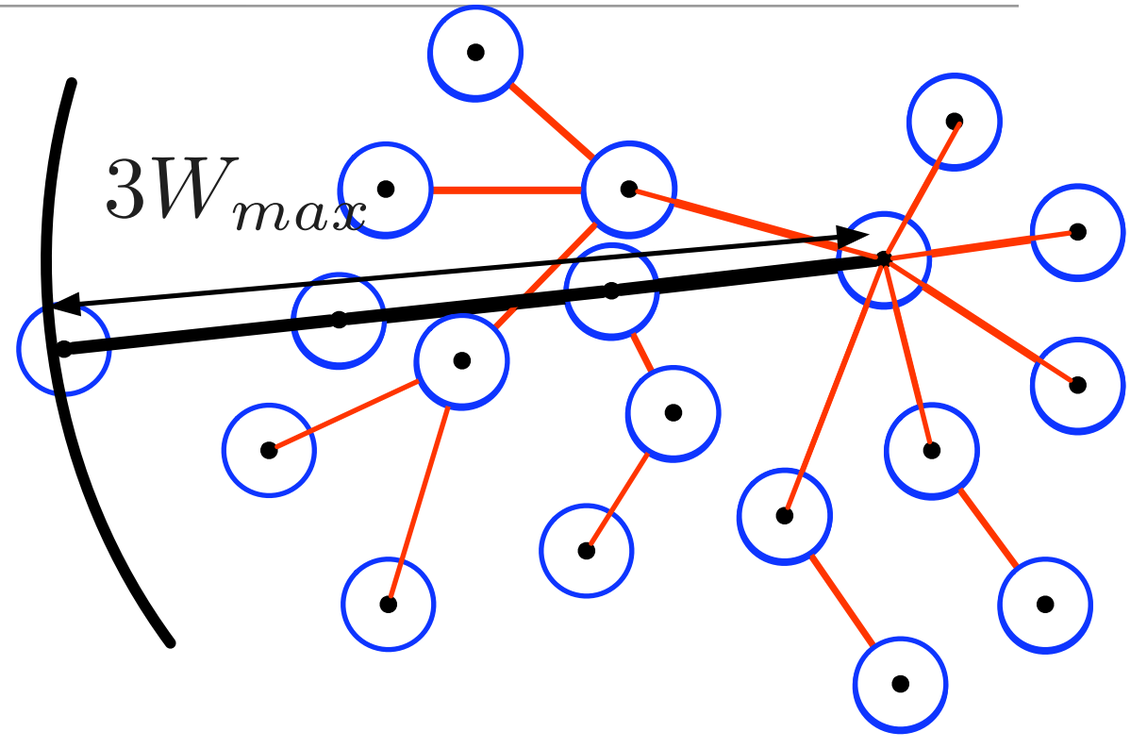

Fundamental bounds on interconnect lengths

"Visible universe" of a node:

$$\pi \left(iter \times W_{max} \right)^2$$

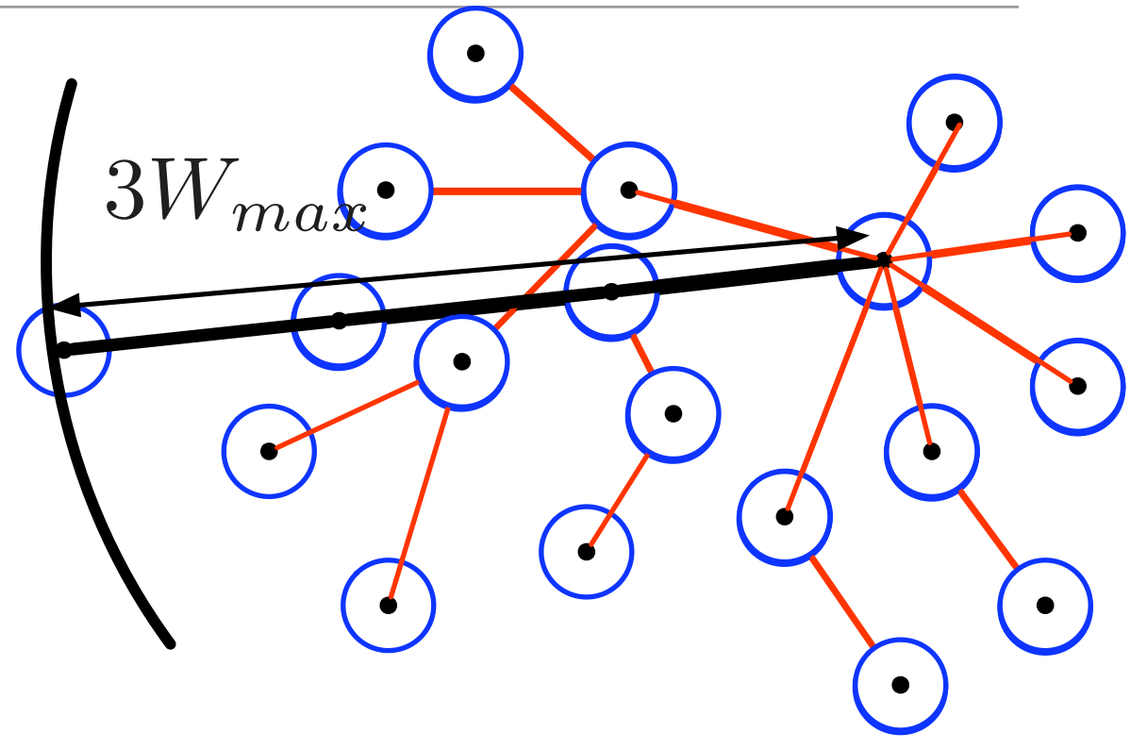

Fundamental bounds on interconnect lengths

"Visible universe" of a node:

$$\pi \left(iter \times W_{max} \right)^2$$

Decentralized decision-making

[Tsitsiklis '84, Boyd et al. '06]

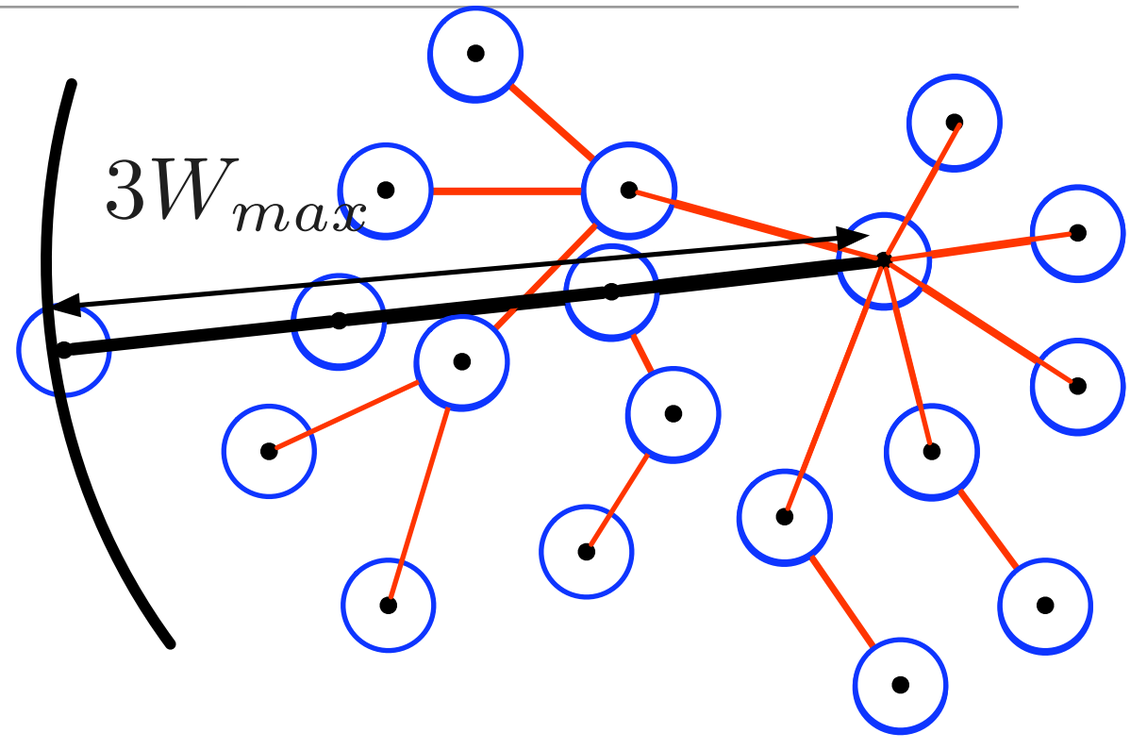

Fundamental bounds on interconnect lengths

"Visible universe" of a node:

$$\pi \left(iter \times W_{max} \right)^2$$

Decentralized decision-making
[Tsitsiklis '84, Boyd et al. '06]

Decoding as decentralized decision-making
[Sahai, Grover '07]

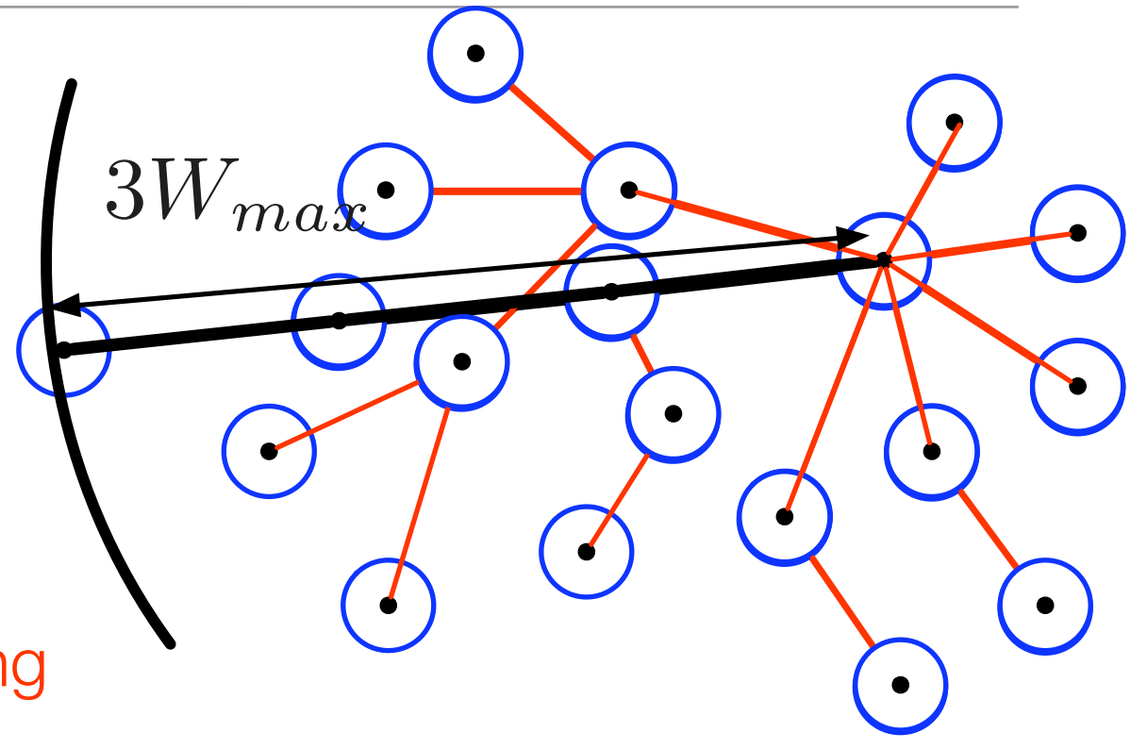

Fundamental bounds on interconnect lengths

"Visible universe" of a node:

$$\pi \left(iter \times W_{max} \right)^2$$

Decentralized decision-making
[Tsitsiklis '84, Boyd et al. '06]

Decoding as decentralized decision-making
[Sahai, Grover '07]

$$P_e \approx e^{-n\beta(C-R)^2}$$

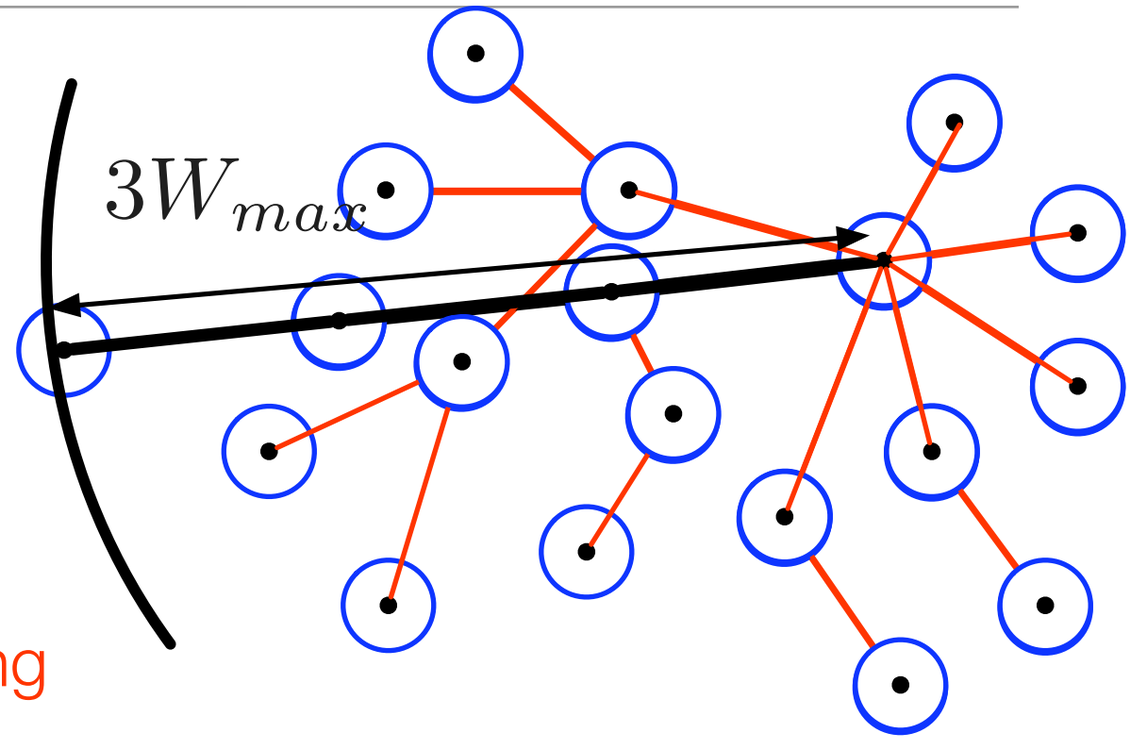

Fundamental bounds on interconnect lengths

"Visible universe" of a node:

$$\pi \left(iter \times W_{max} \right)^2$$

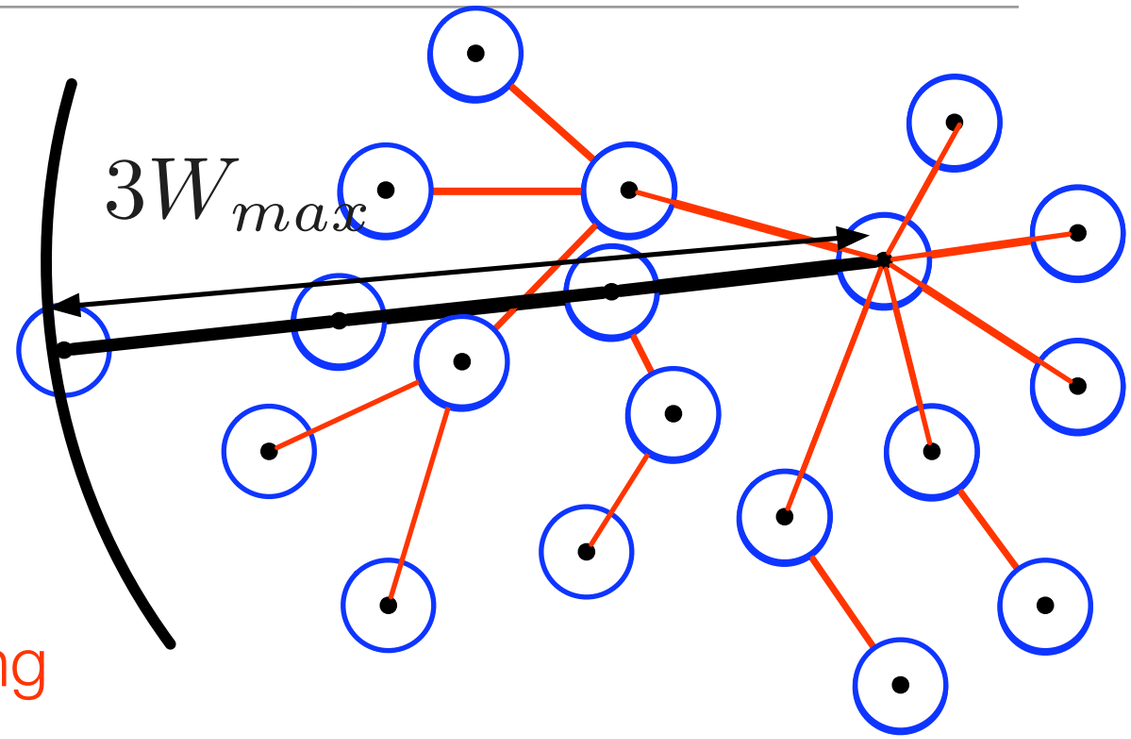

Decentralized decision-making
[Tsitsiklis '84, Boyd et al. '06]

Decoding as decentralized decision-making
[Sahai, Grover '07]

$$P_e \approx e^{-n\beta(C-R)^2}$$

Theorem [Grover, Sahai CISS'11]

For *any* code/decoder in t iterations,

$$iter \times W_{max} \gtrsim \Theta \left(\sqrt{A_{node}} \frac{\sqrt{\log \frac{1}{P_e}}}{C(P_T) - R} \right)$$

Fundamental bounds on interconnect lengths

"Visible universe" of a node:

$$\pi \left(iter \times W_{max} \right)^2$$

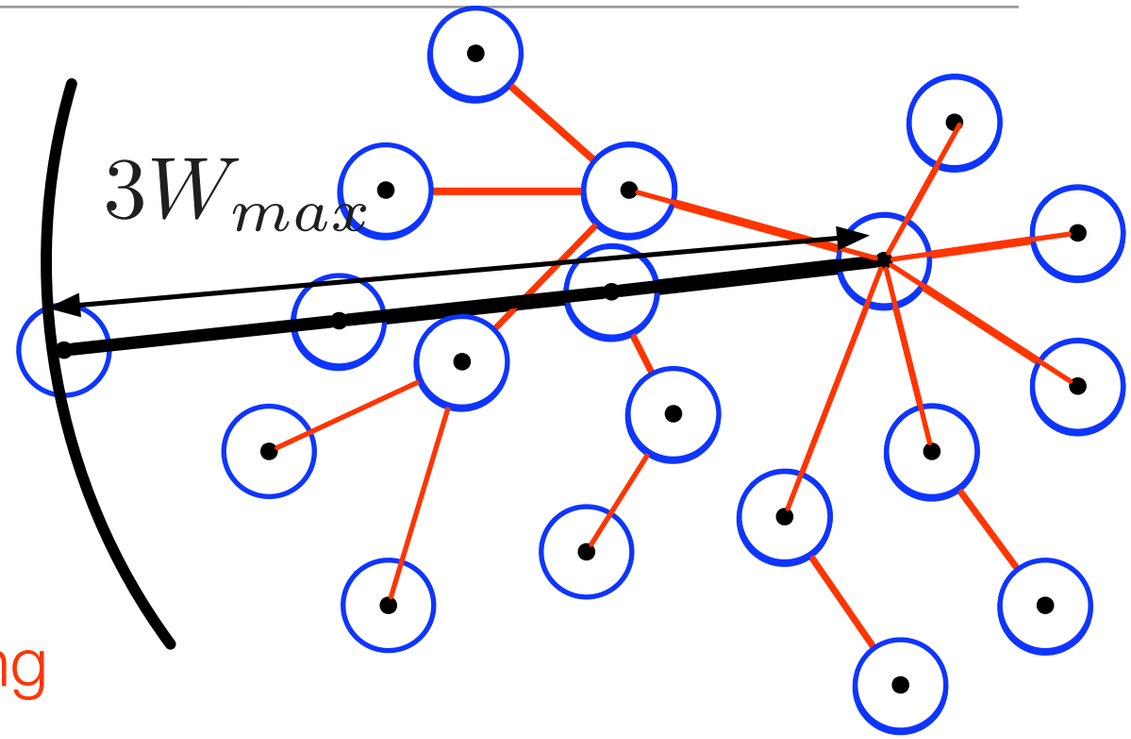

Decentralized decision-making
[Tsitsiklis '84, Boyd et al. '06]

Decoding as decentralized decision-making
[Sahai, Grover '07]

$$P_e \approx e^{-n\beta(C-R)^2}$$

Theorem [Grover, Sahai CISS'11]

For *any* code/decoder in t iterations,

$$iter \times W_{max} \gtrsim \Theta \left(\sqrt{A_{node}} \frac{\sqrt{\log \frac{1}{P_e}}}{C(P_T) - R} \right)$$

Wire model: $P_{decoding} \propto iter \times W_{max}$

“Wire model”
Lower bounds on power consumption

“Wire model”
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

“Wire model”
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

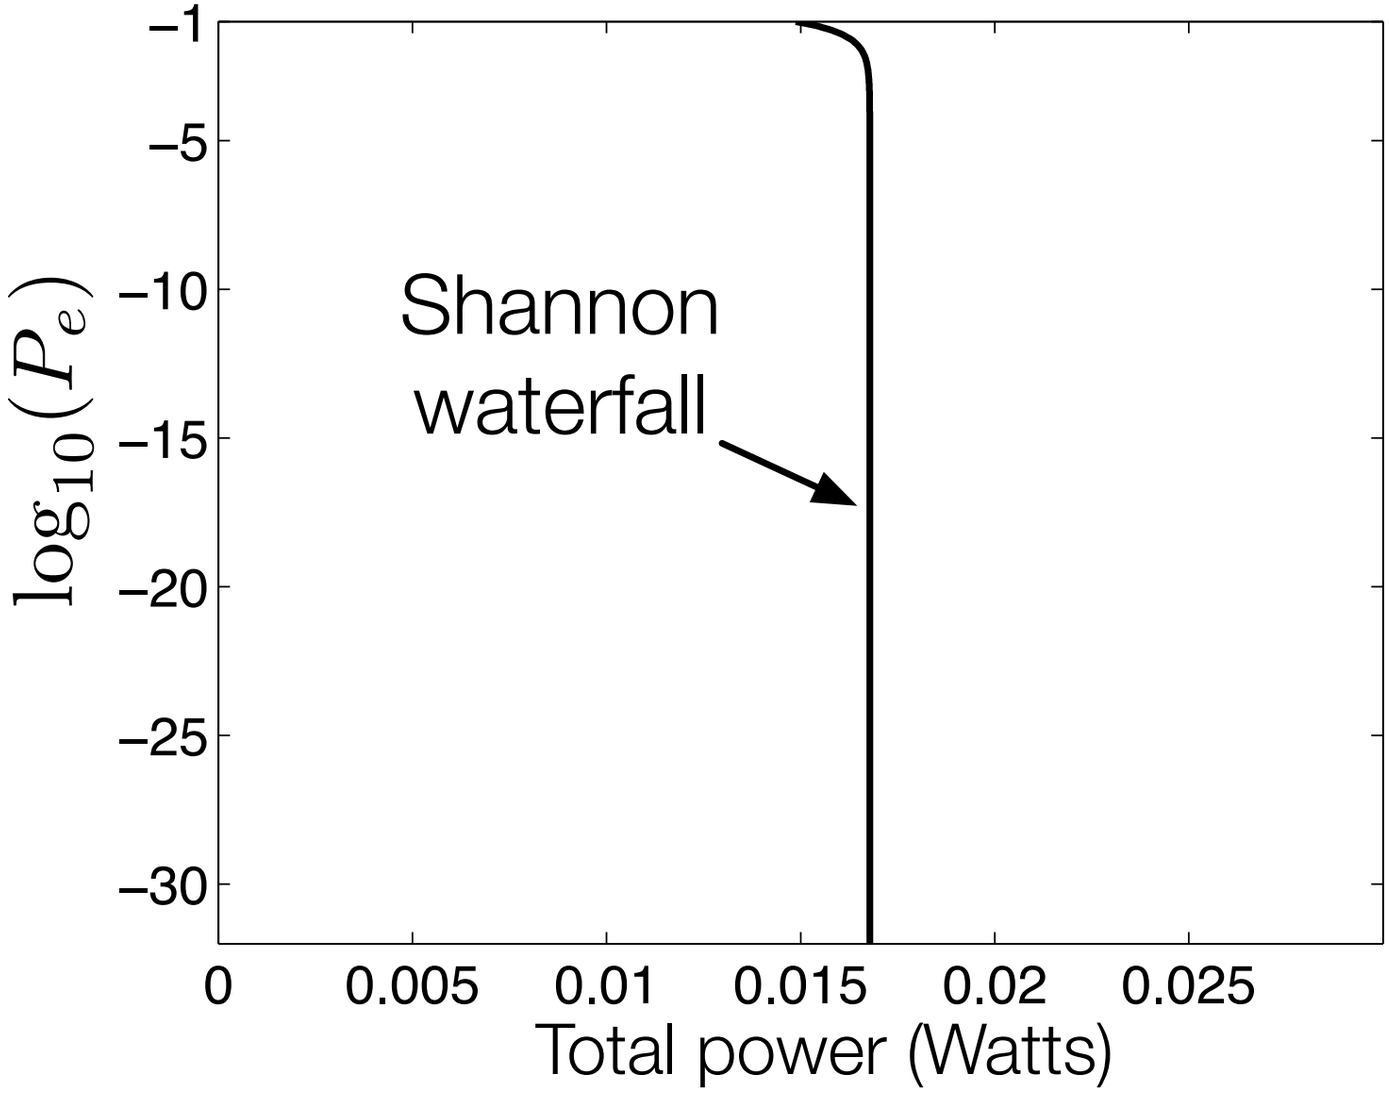

“Wire model”
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

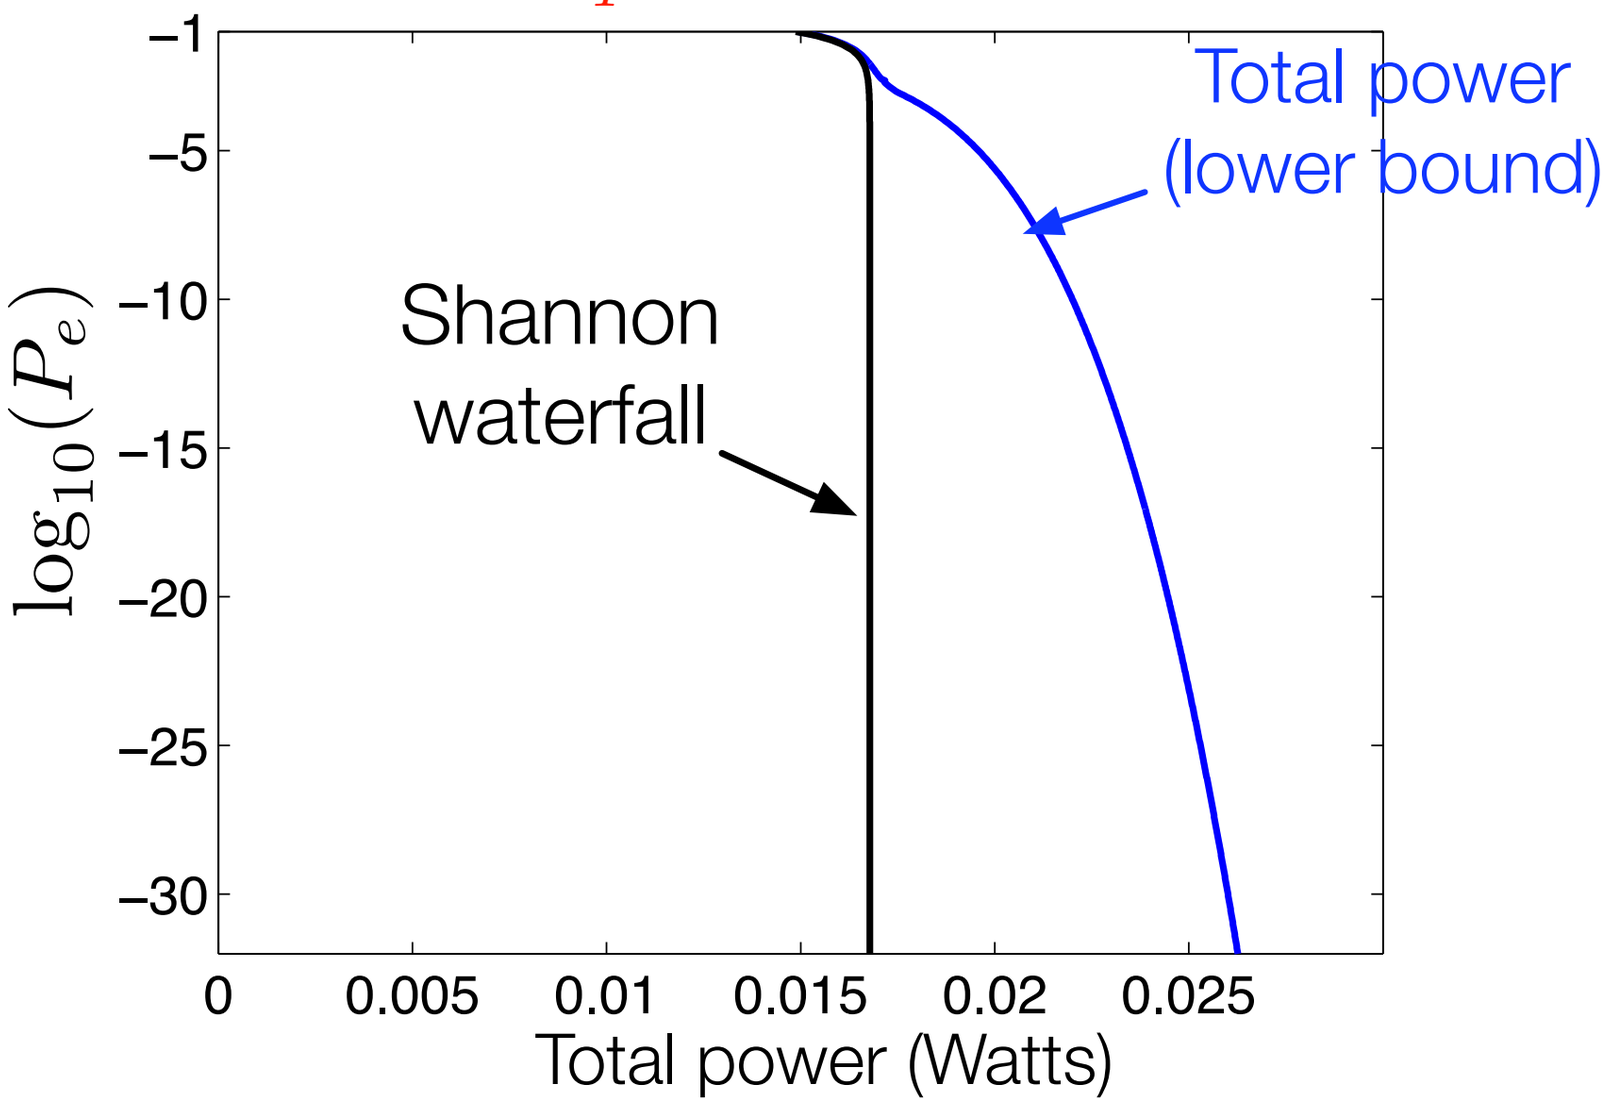

“Wire model”
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

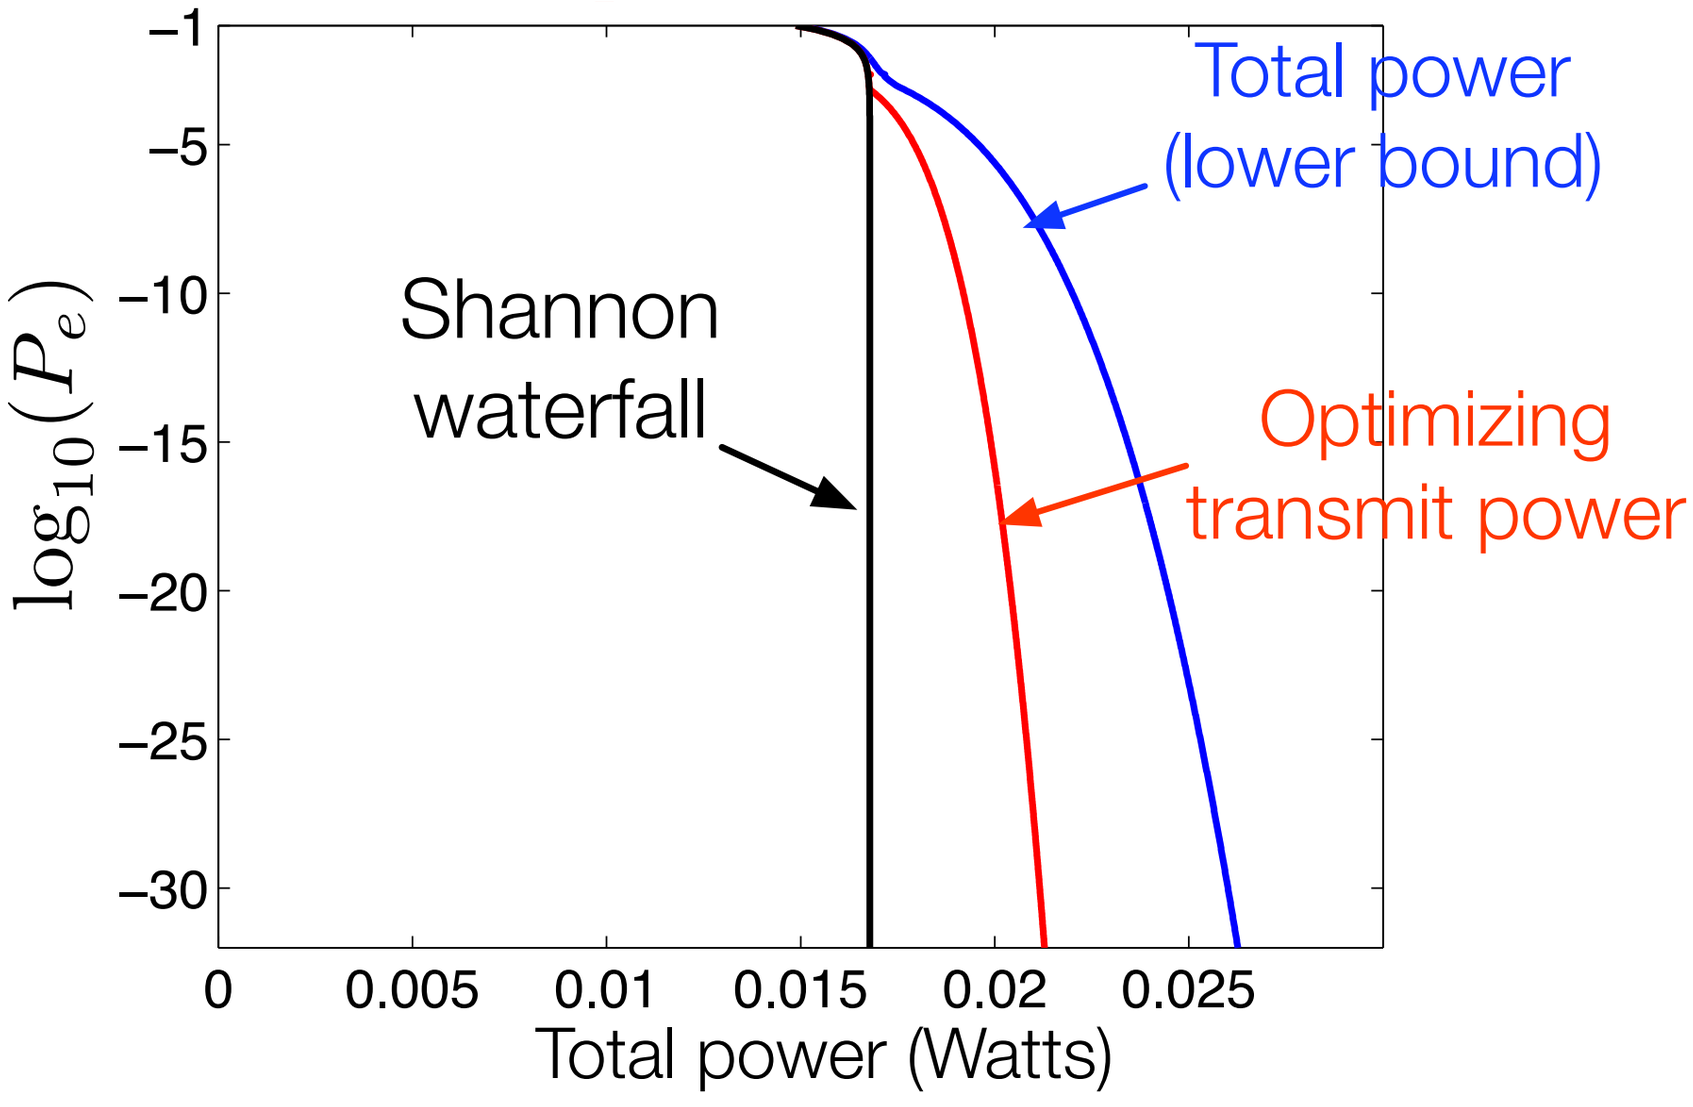

"Wire model"
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

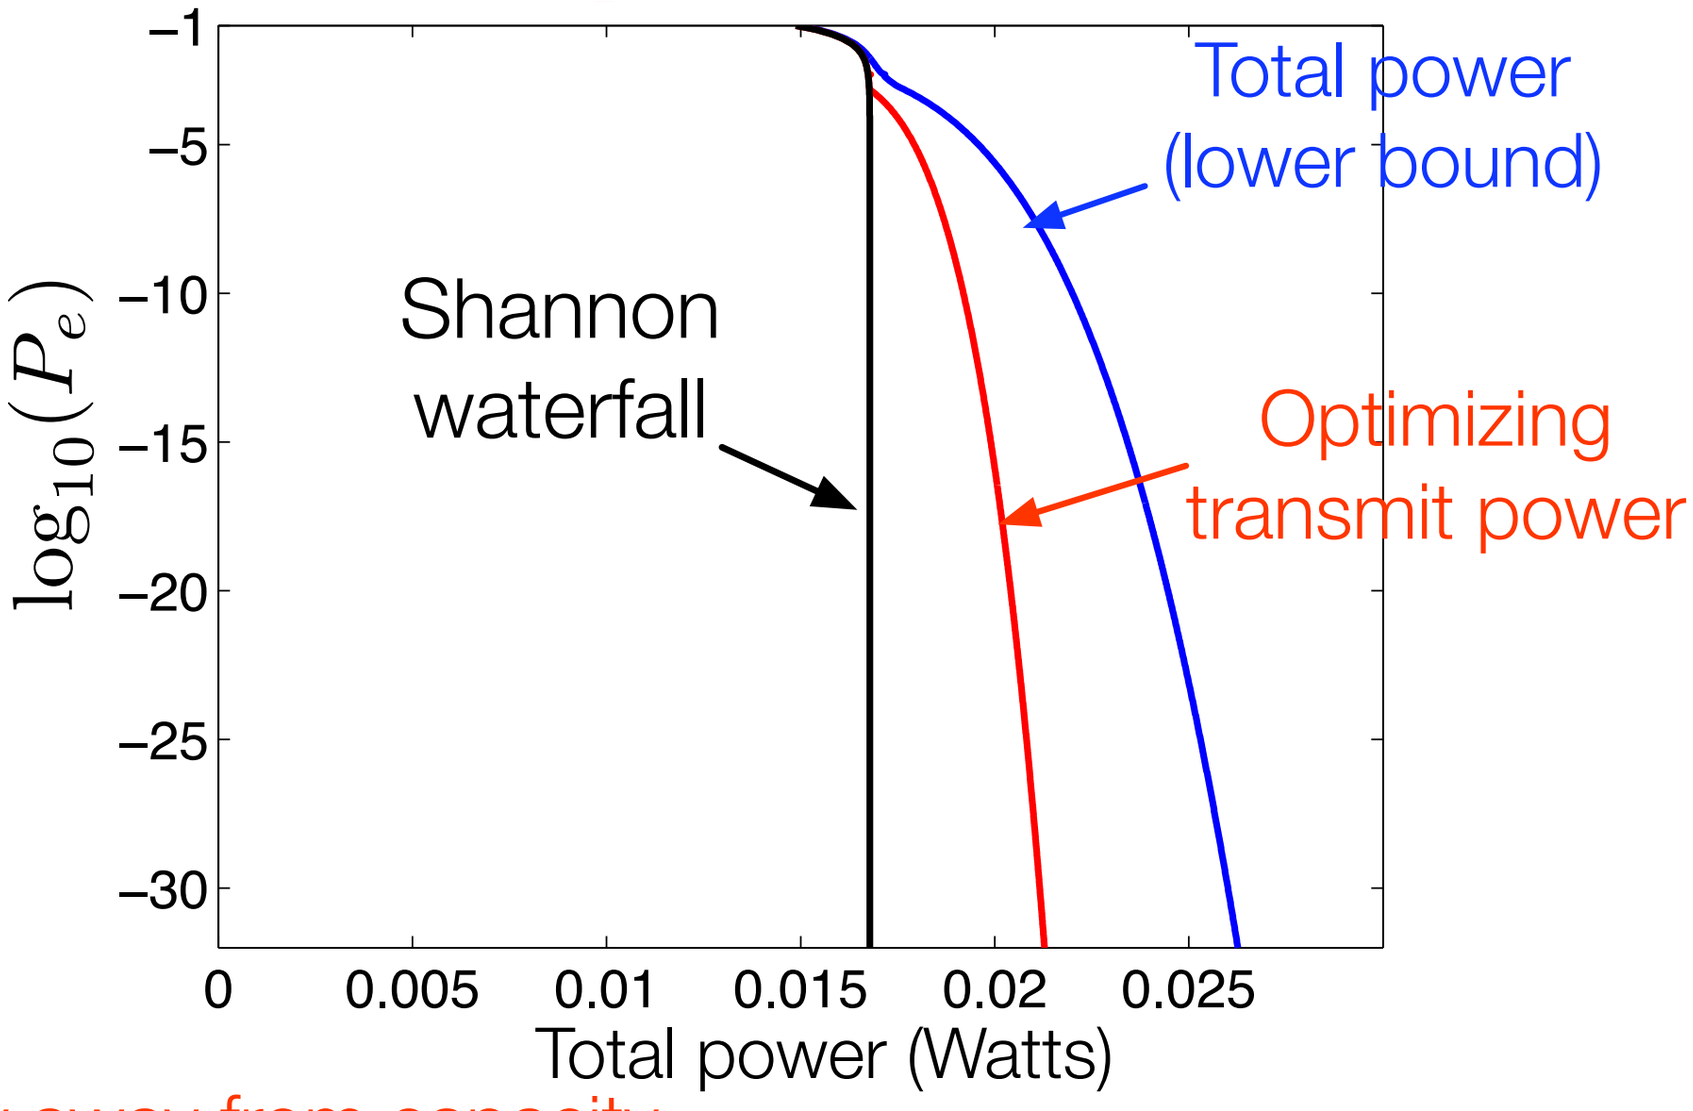

Moral: stay away from capacity

“Wire model”
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

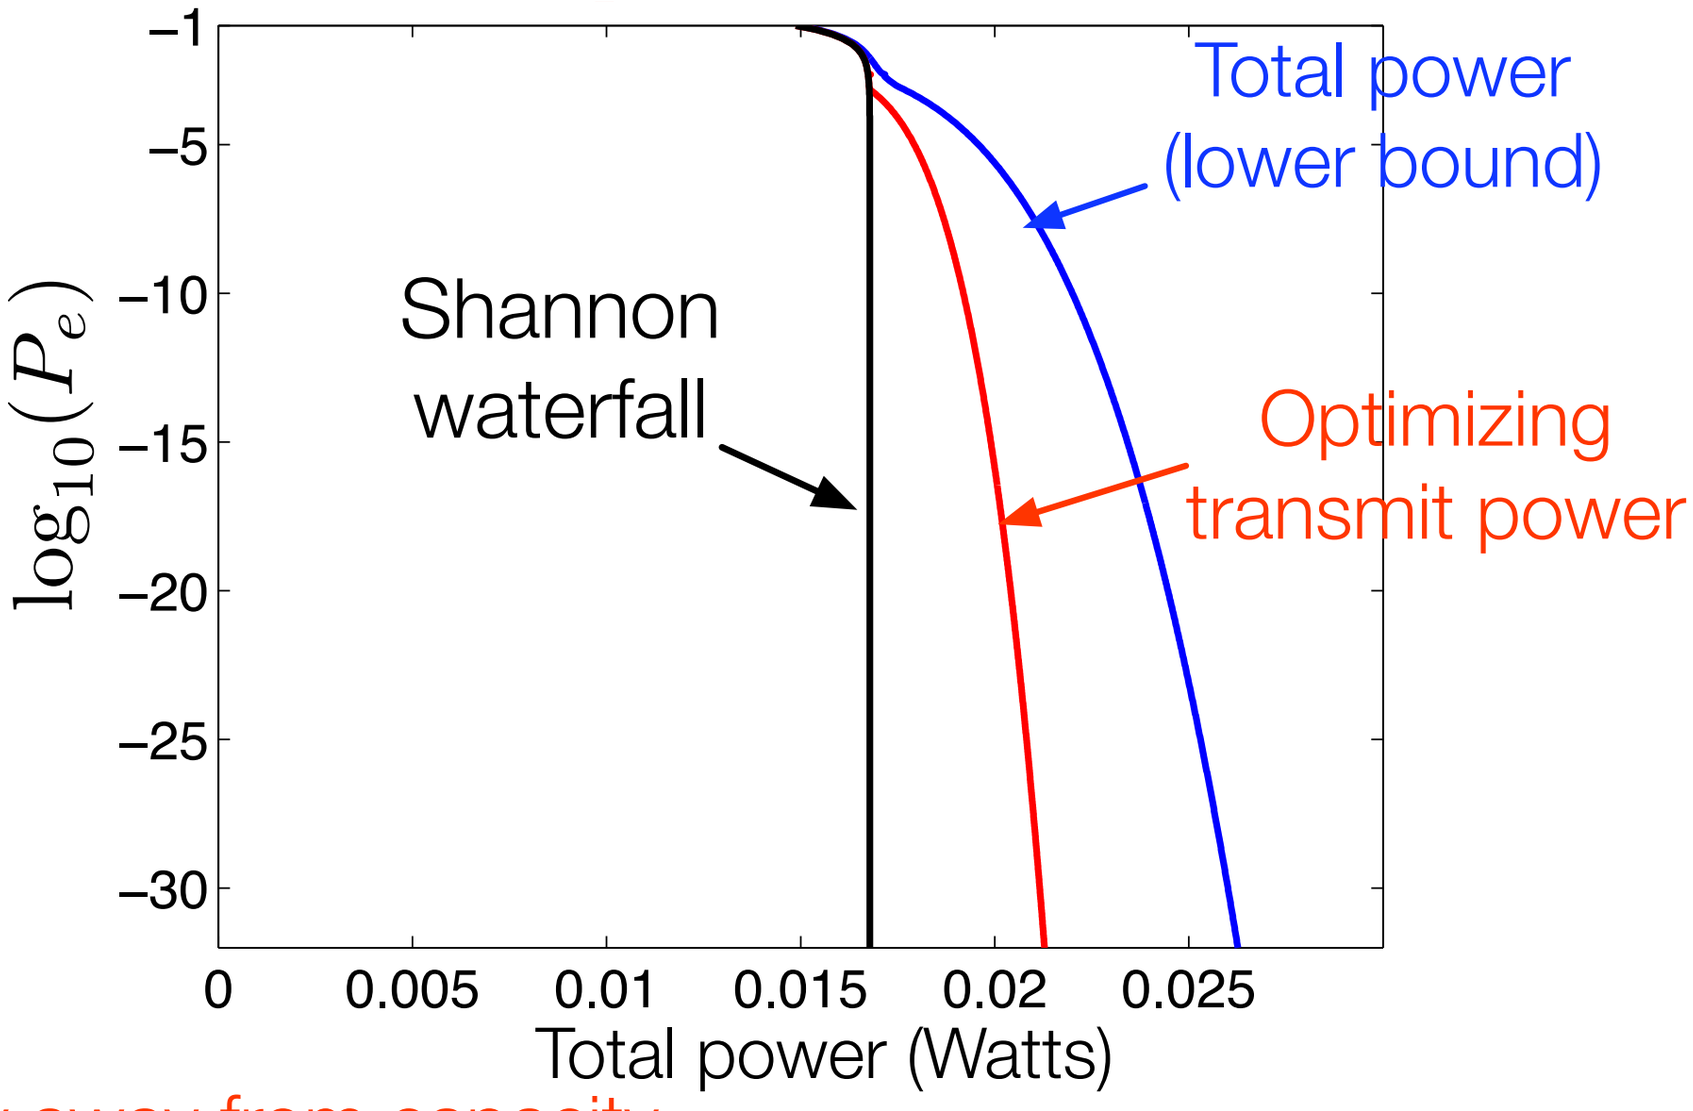

Moral: stay away from capacity

design codes for just the required performance

"Wire model"
Lower bounds on power consumption

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

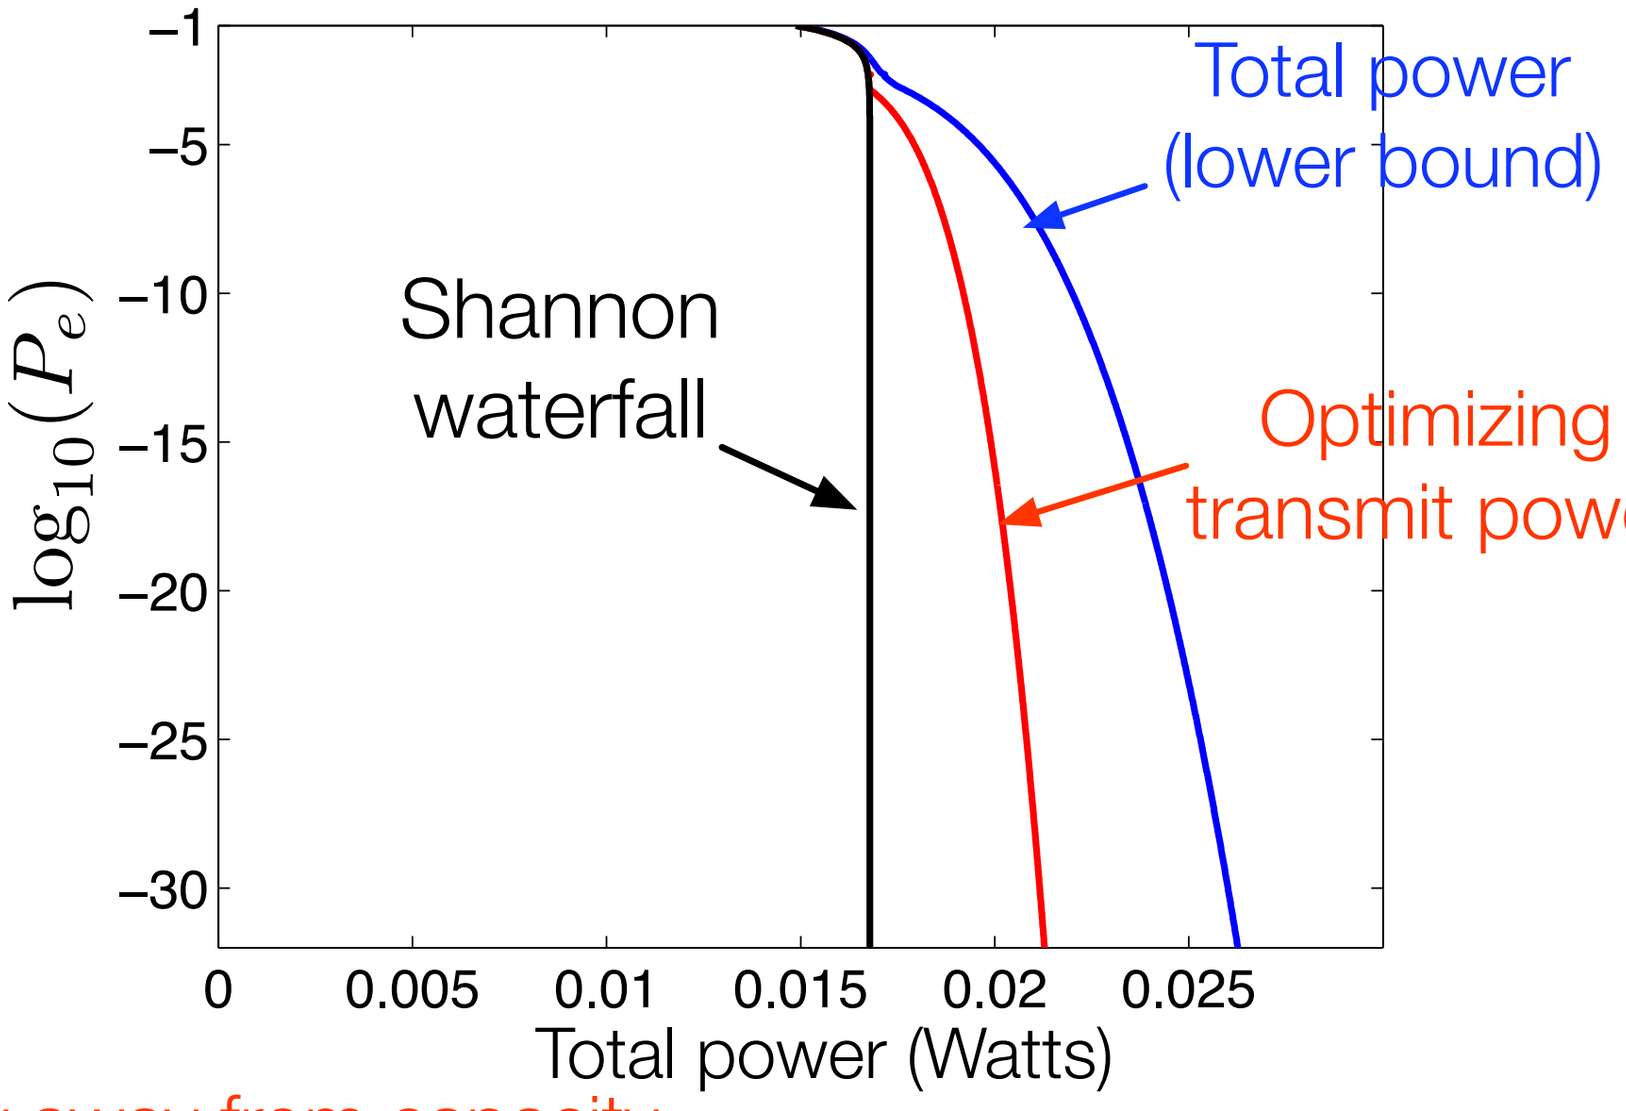

Moral: stay away from capacity

design codes for just the required performance

. . . and no better

Constructing codes with short wires, good performance

Constructing codes with short wires, good performance

---

short wires: by small block-lengths?

Constructing codes with short wires, good performance

---

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager ’63][Furedi *et al.* ’95]

Constructing codes with short wires, good performance

---

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager '63][Furedi *et al.* '95]
[Grover, Sahai '11]

Constructing codes with short wires, good performance

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager '63][Furedi *et al.* '95]

[Grover, Sahai '11]

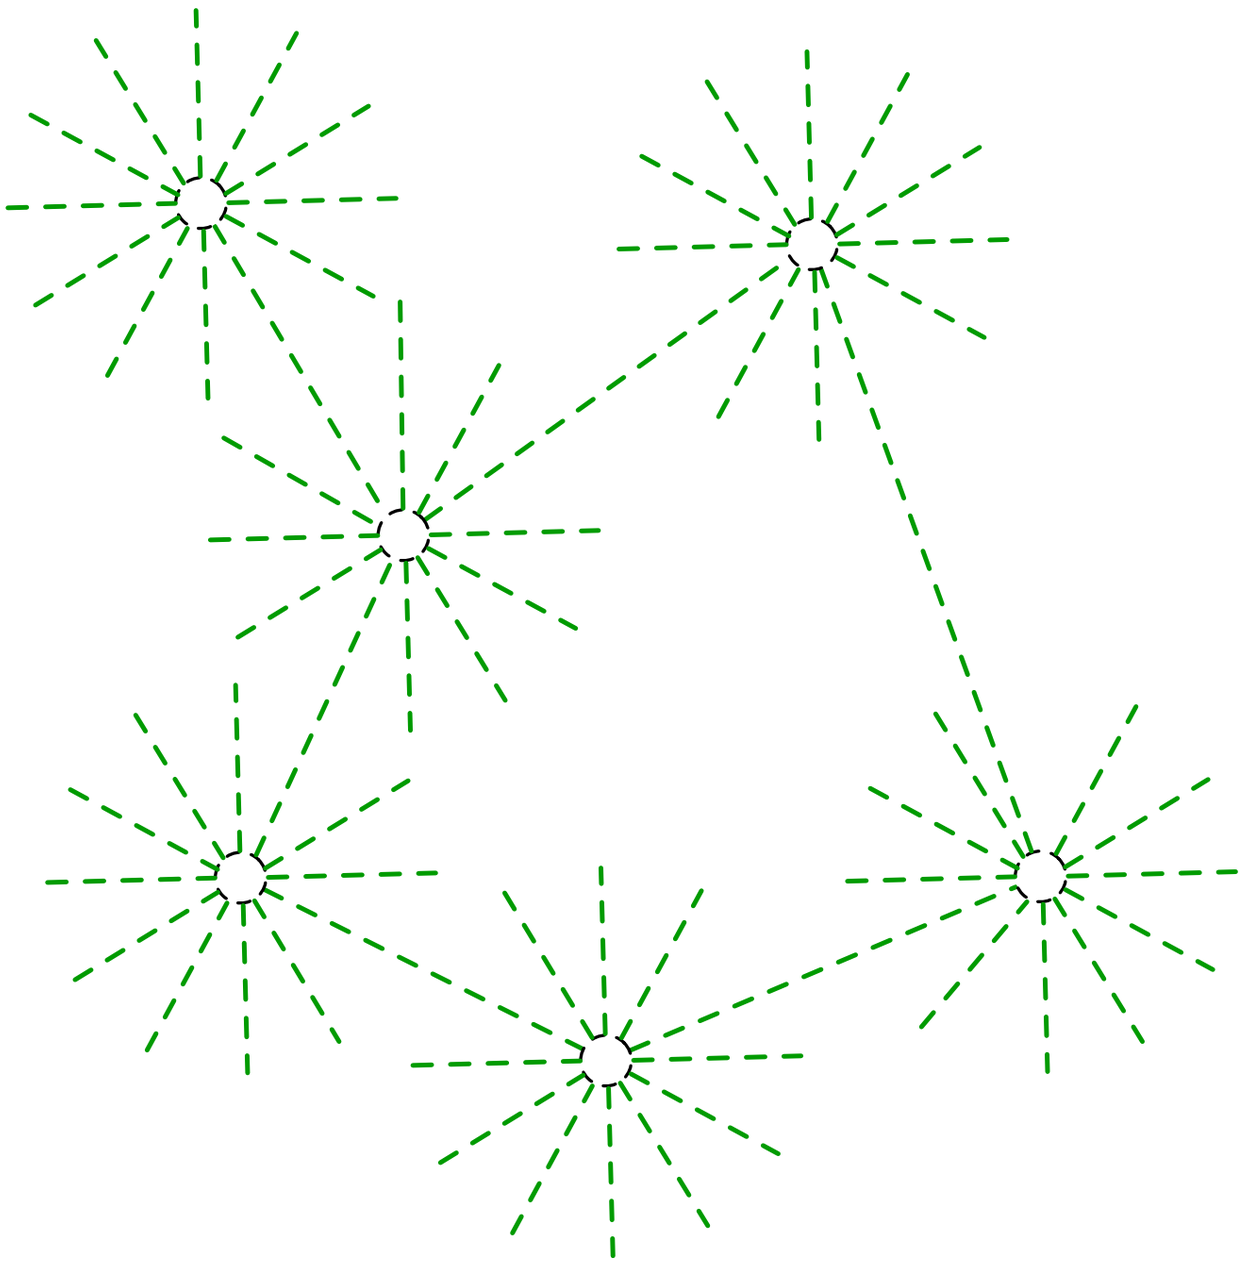

Constructing codes with short wires, good performance

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager '63][Furedi *et al.* '95]

[Grover, Sahai '11]

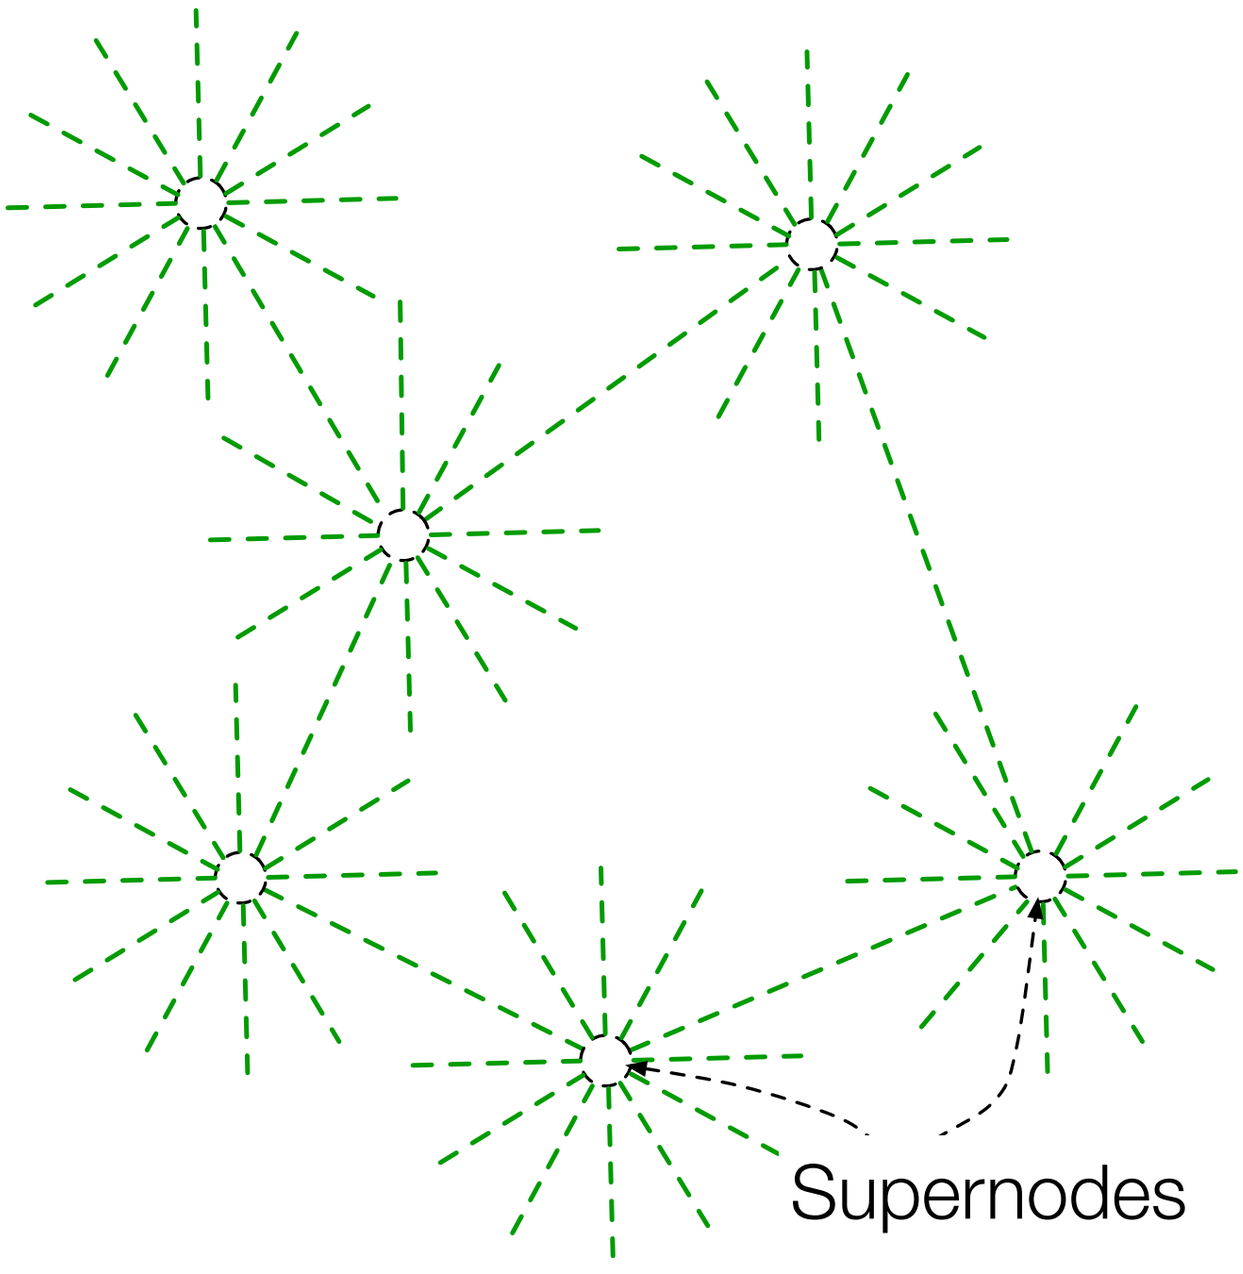

Constructing codes with short wires, good performance

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager '63][Furedi *et al.* '95]

[Grover, Sahai '11]

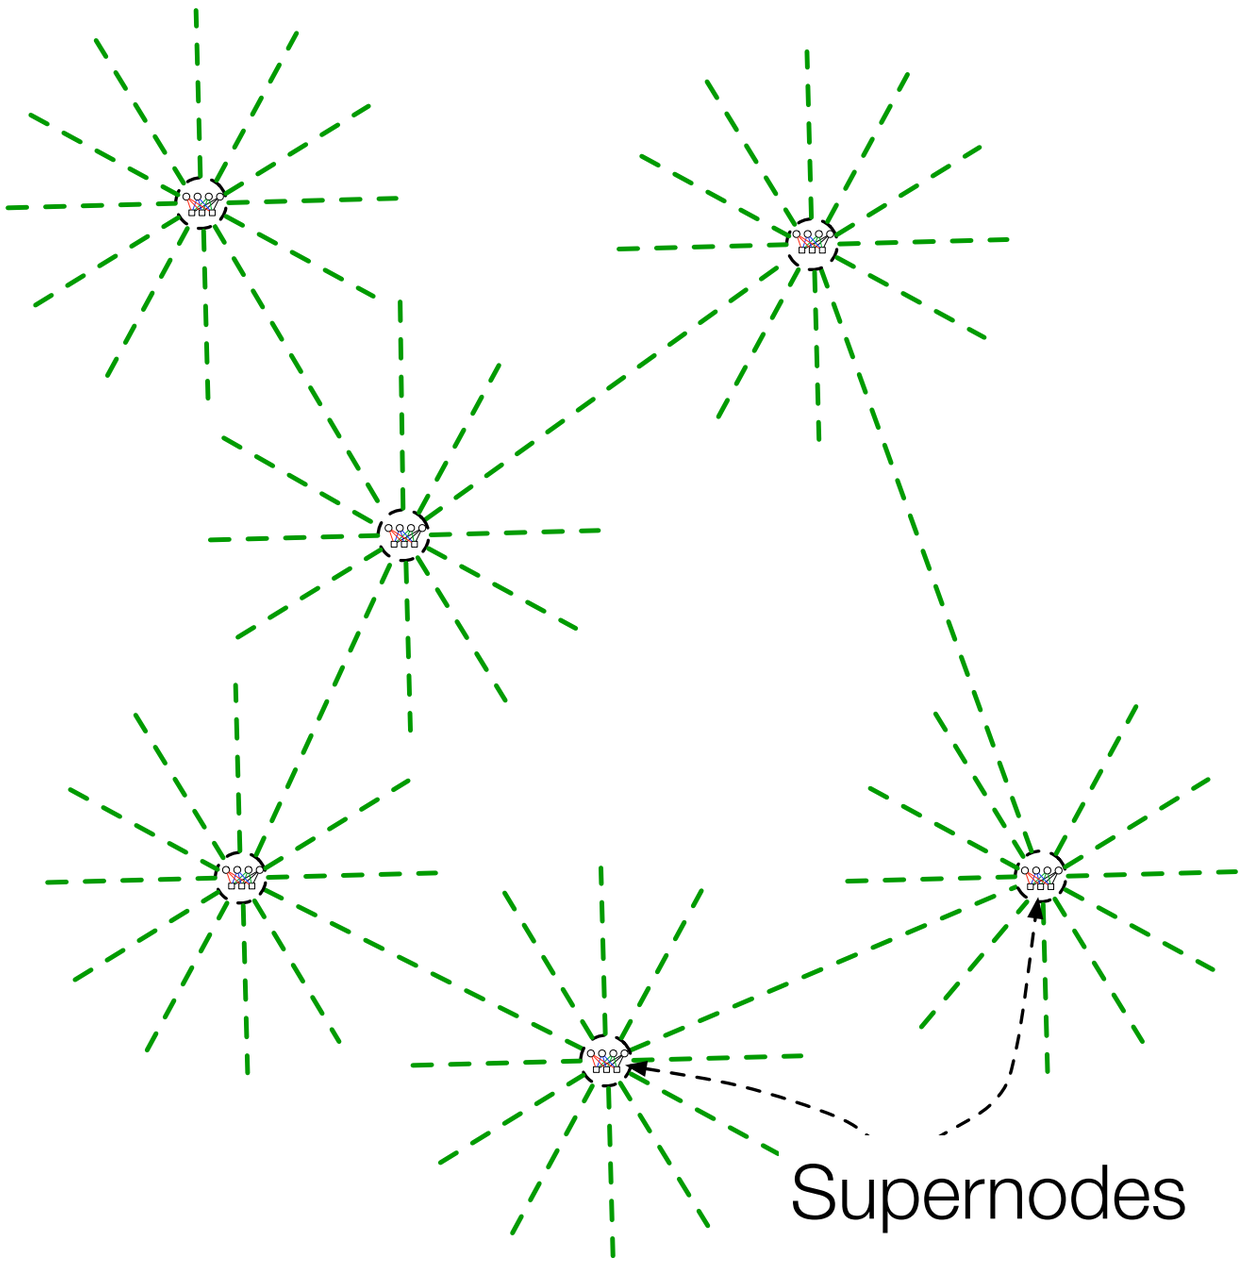

Constructing codes with short wires, good performance

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager ’63][Furedi *et al.* ’95]

[Grover, Sahai ’11]

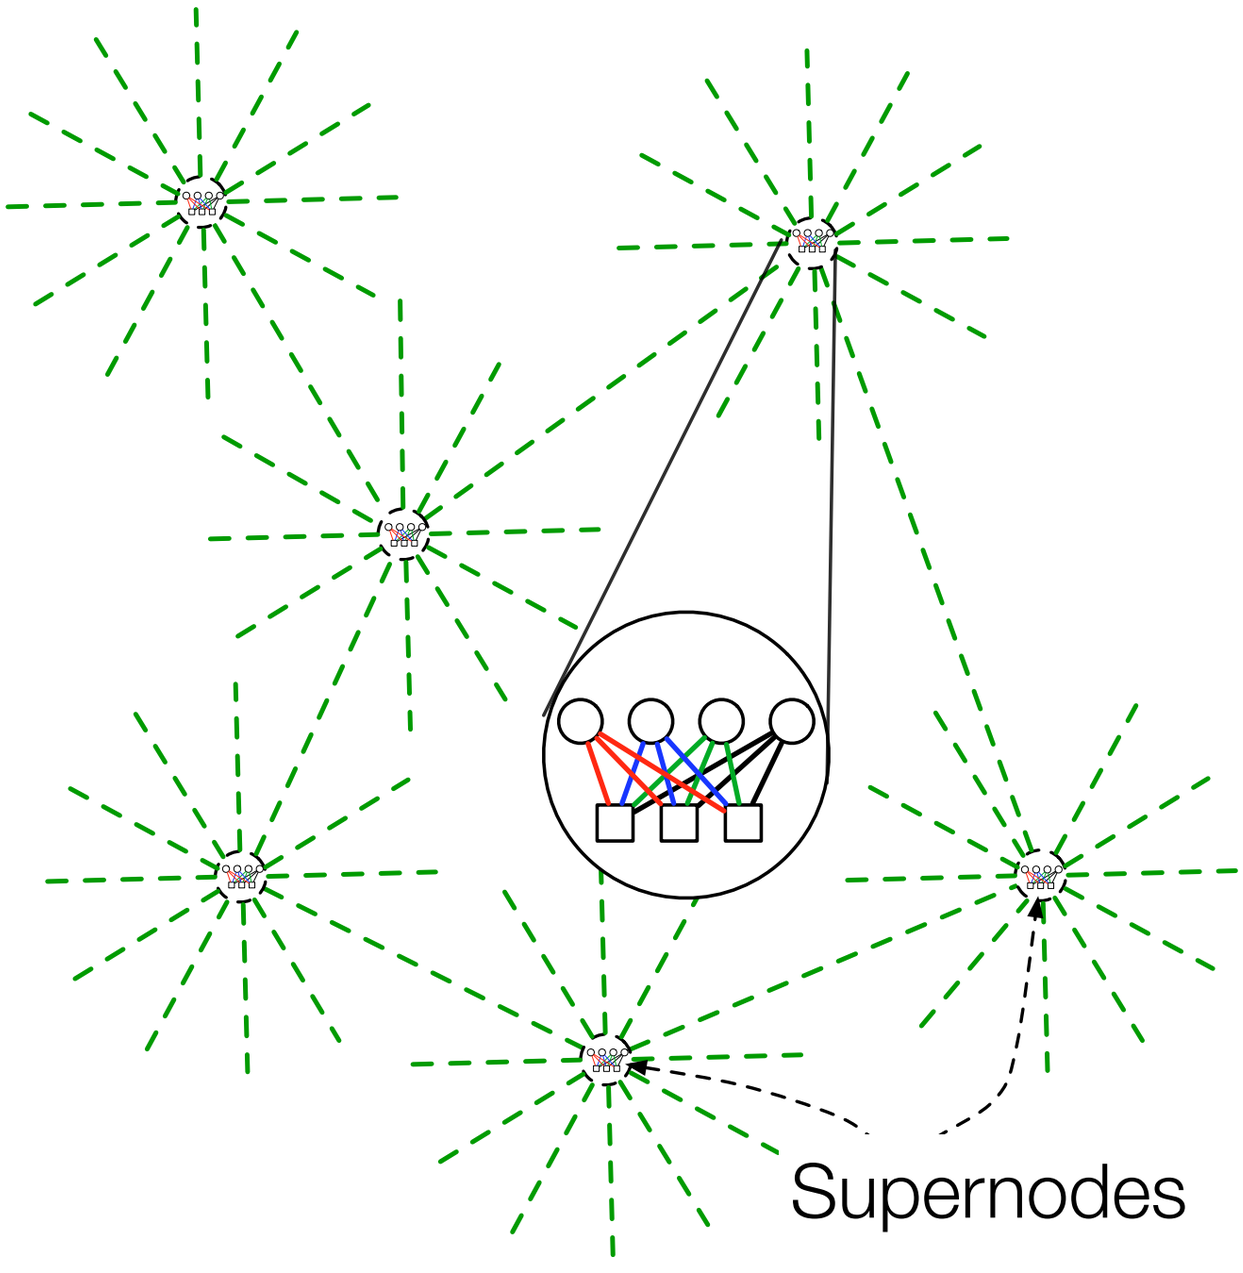

Constructing codes with short wires, good performance

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager '63][Furedi *et al.* '95]

[Grover, Sahai '11]

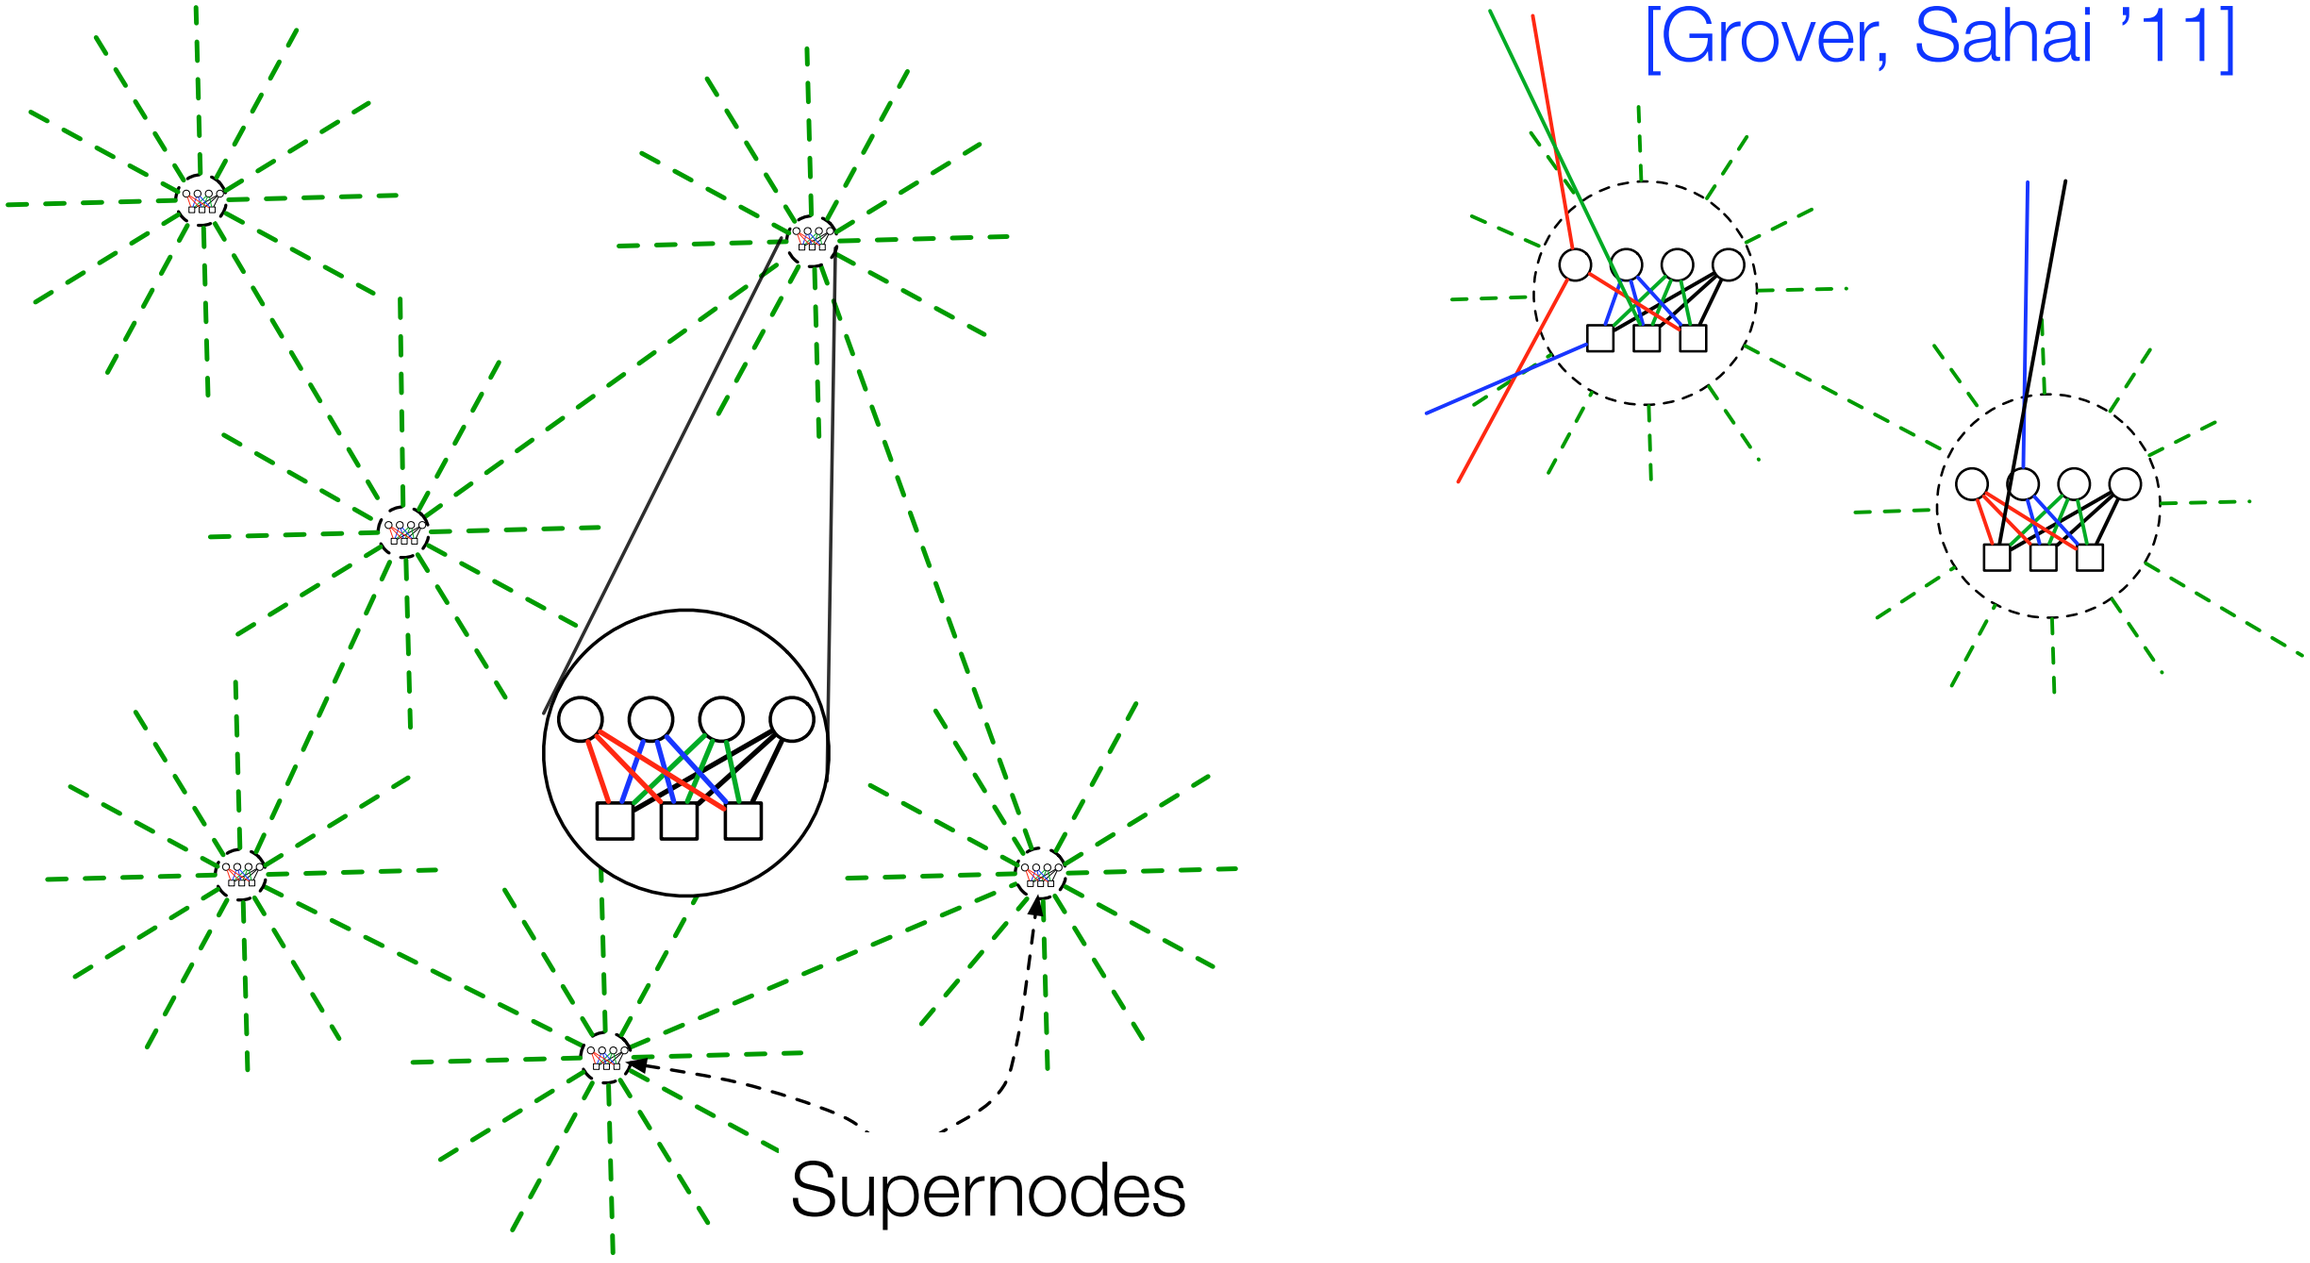

Constructing codes with short wires, good performance

short wires: by small block-lengths?

LDPCs with small block-lengths, large girth [Gallager '63][Furedi *et al.* '95]

[Grover, Sahai '11]

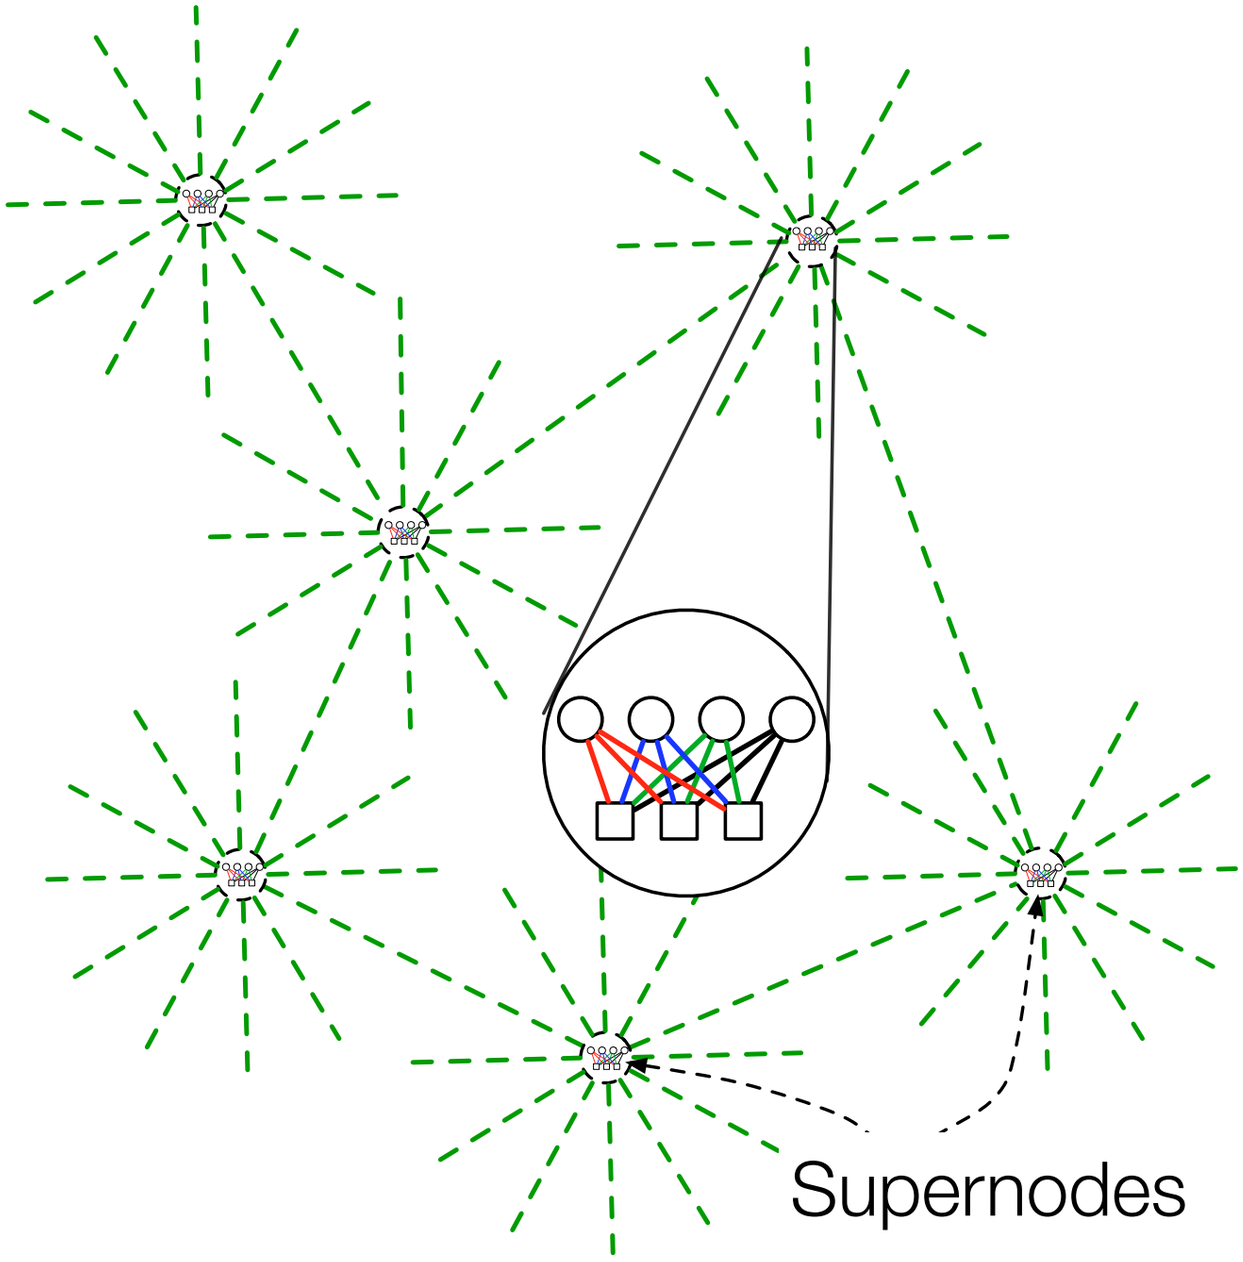

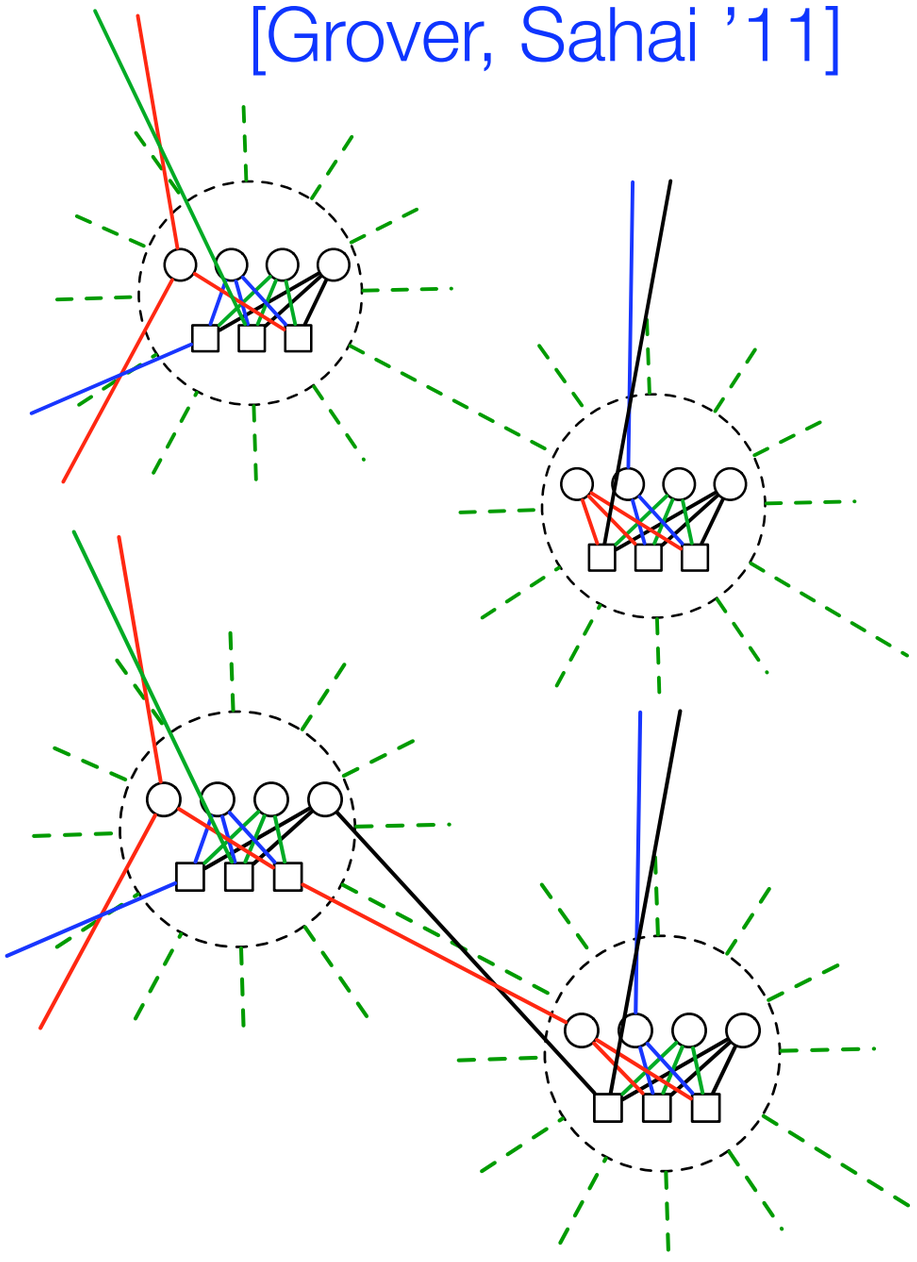

Achievability: code constructions for an “upper bound”

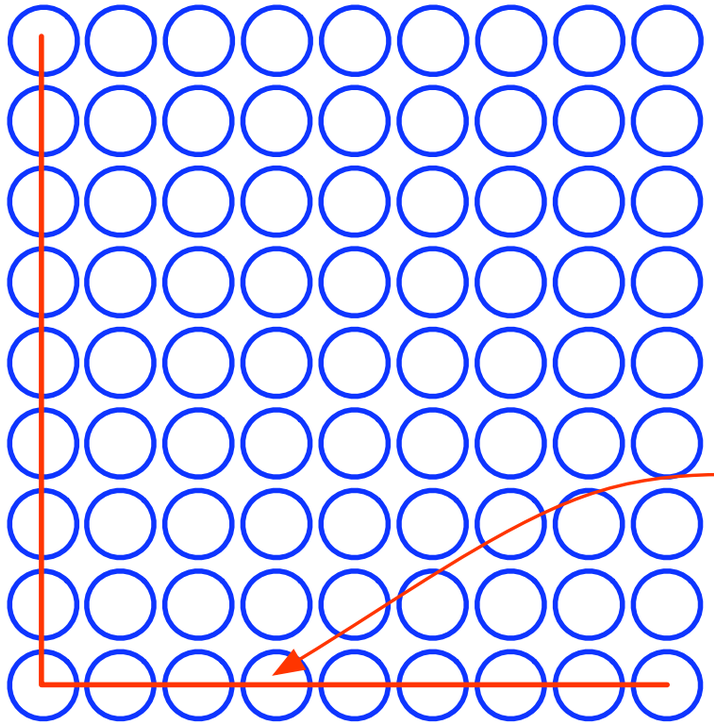

Longest wire in our model
(with Mahattan wiring)

Achievability: code constructions for an “upper bound”

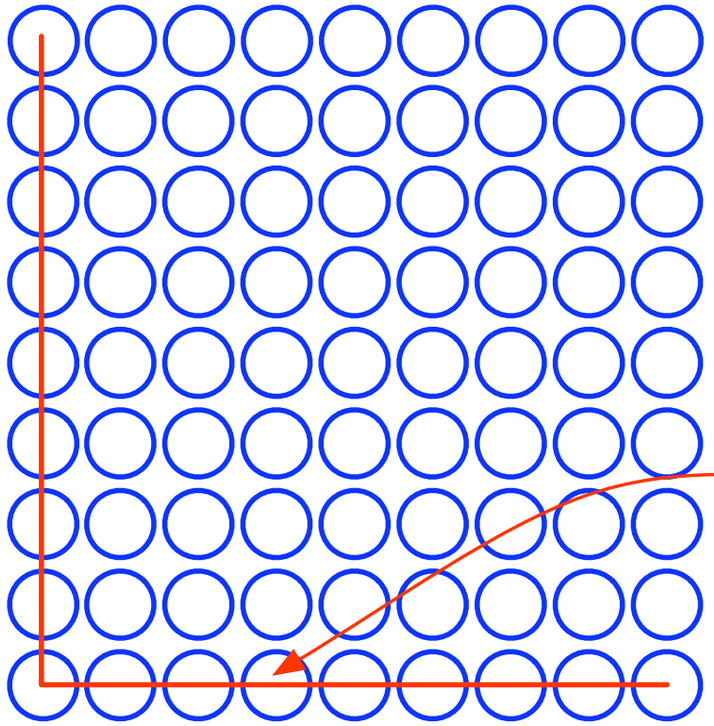

Theorem [Grover, Sahai CISS‘11]

There exist (d_v, d_c)-regular LDPC codes such that

$$W_{max} \leq 4\left(\sqrt{2(d_v+d_c)d_v d_c q^{\frac{3}{4}g_{code}-a}}+1\right)\sqrt{\frac{A_{node}}{\pi}}$$

$$t \times W_{max} = O\left(\left(\log\frac{1}{P_e}\right)^{\zeta}\right)$$

Talk outline

1. How do we build power-efficient communication systems?

    - understanding processing/decoding power

2. A “VLSI model” of decoding computation

3. Decoding power model: “Wire model”

    - do not operate too close to capacity!

    - design codes for just the required performance (no better!)

4. Can more complex nodes help?

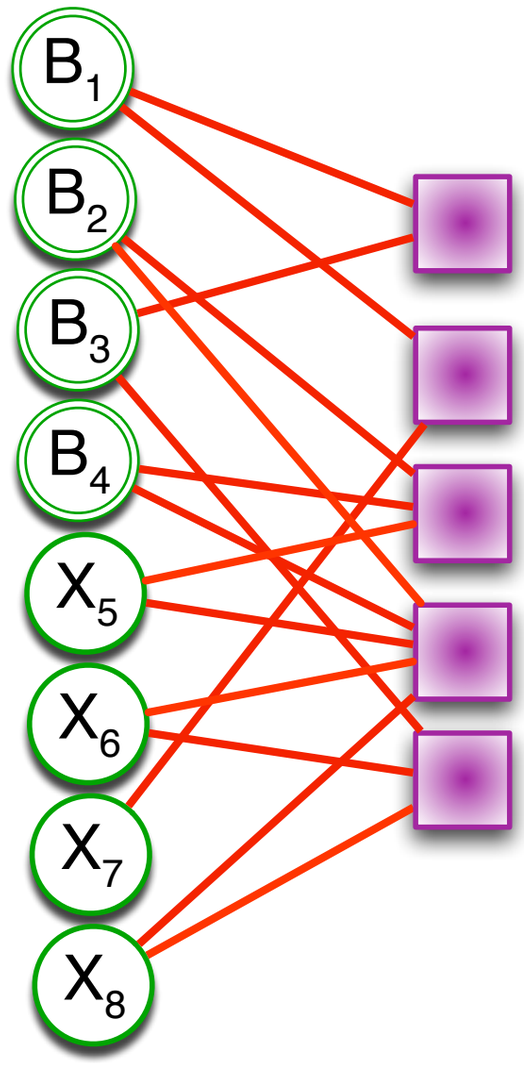

Models of power consumption

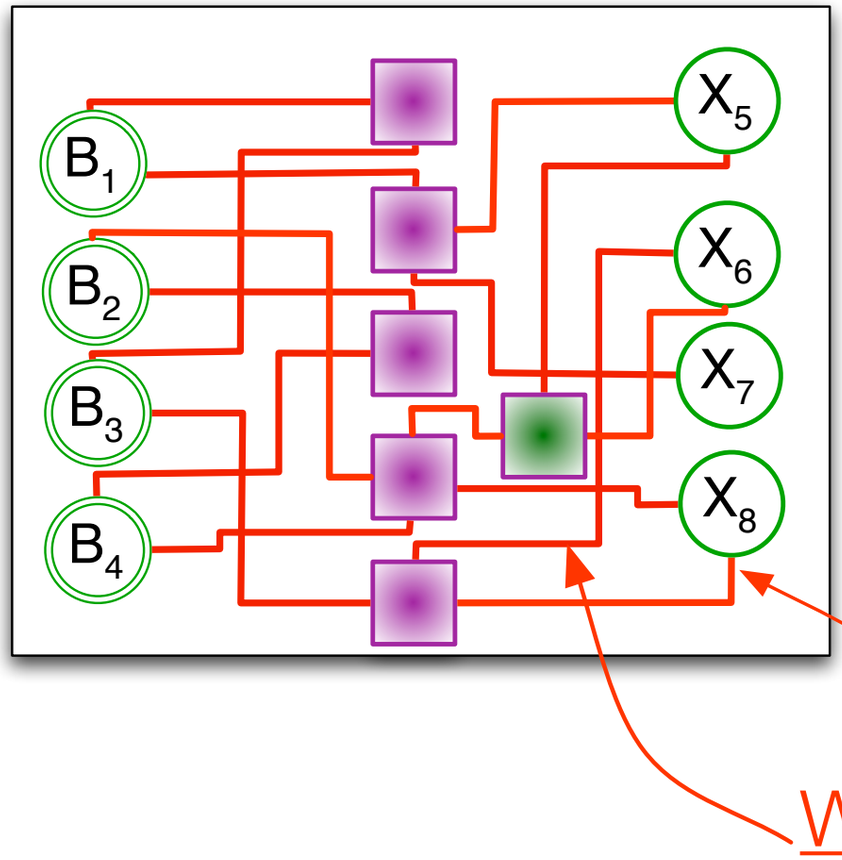

Power consumption increases linearly with the wire-length

Models of power consumption

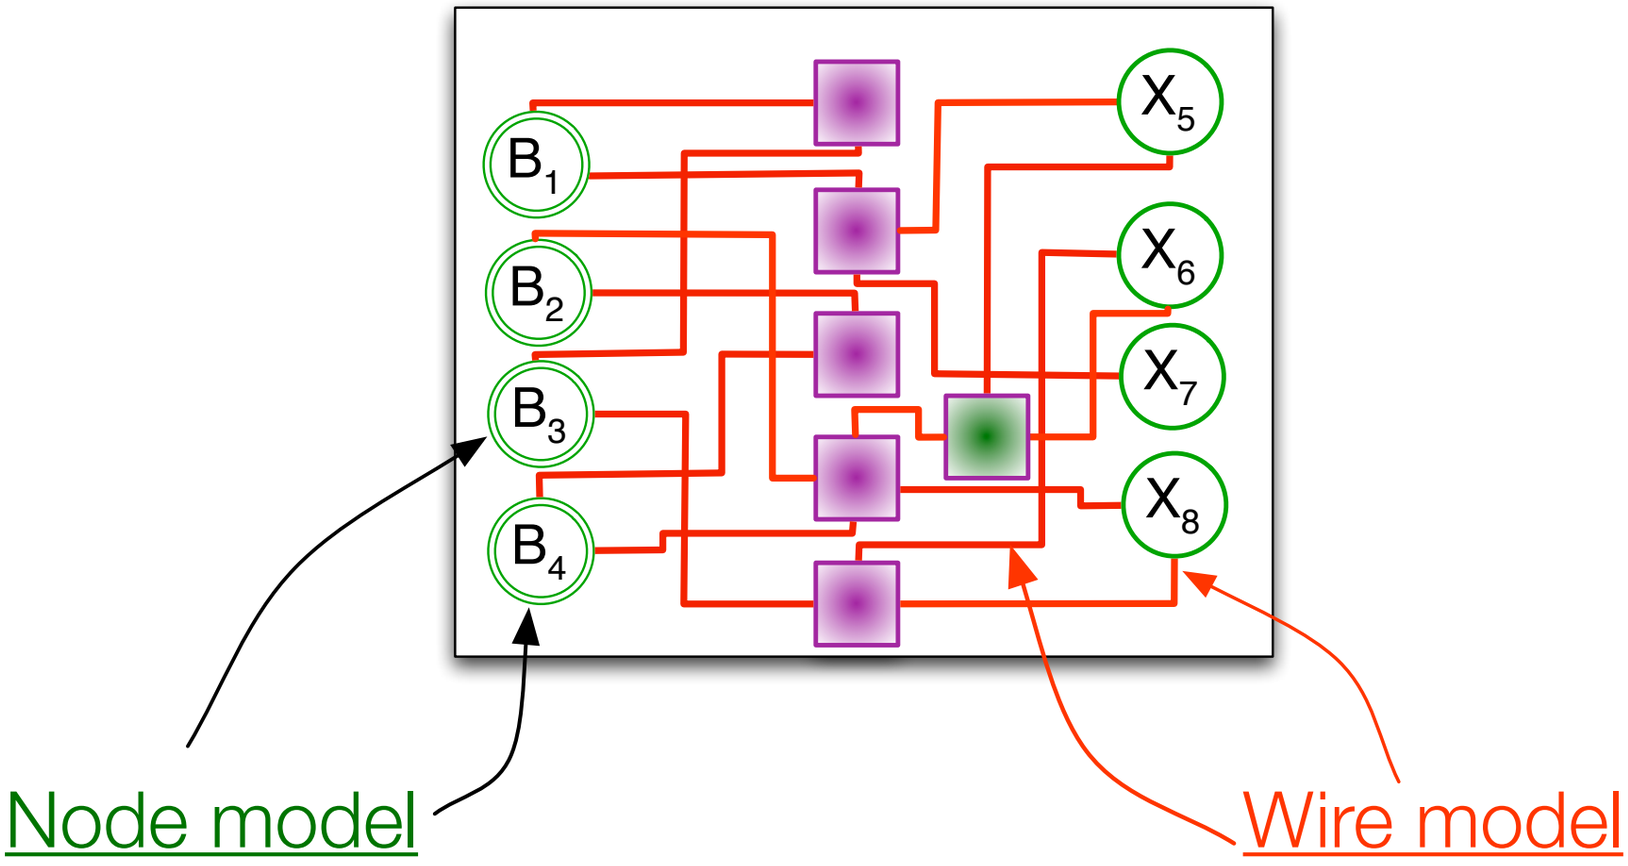

Each node consumes fixed power per iteration

Power consumption increases linearly with the wire-length

Can more complex nodes help?

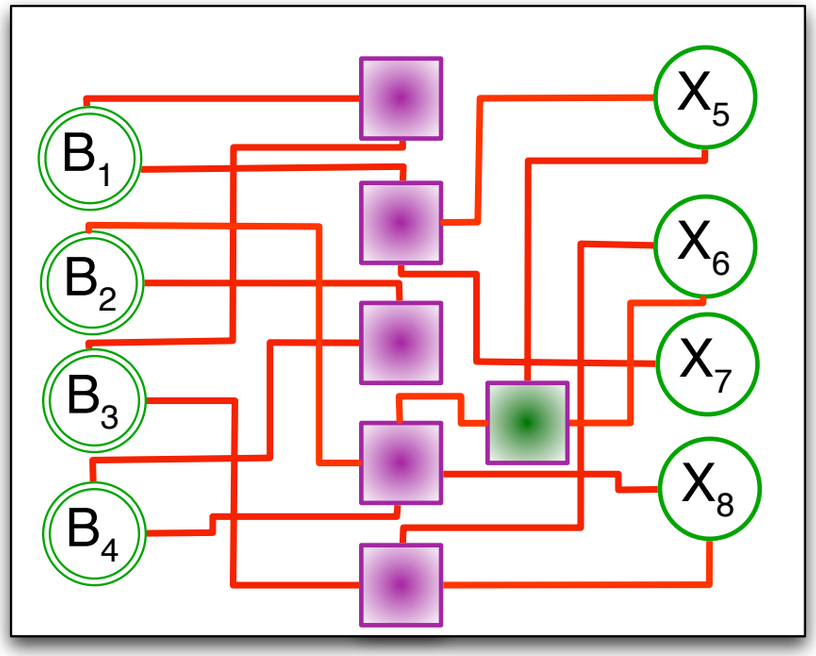

Can more complex nodes help?

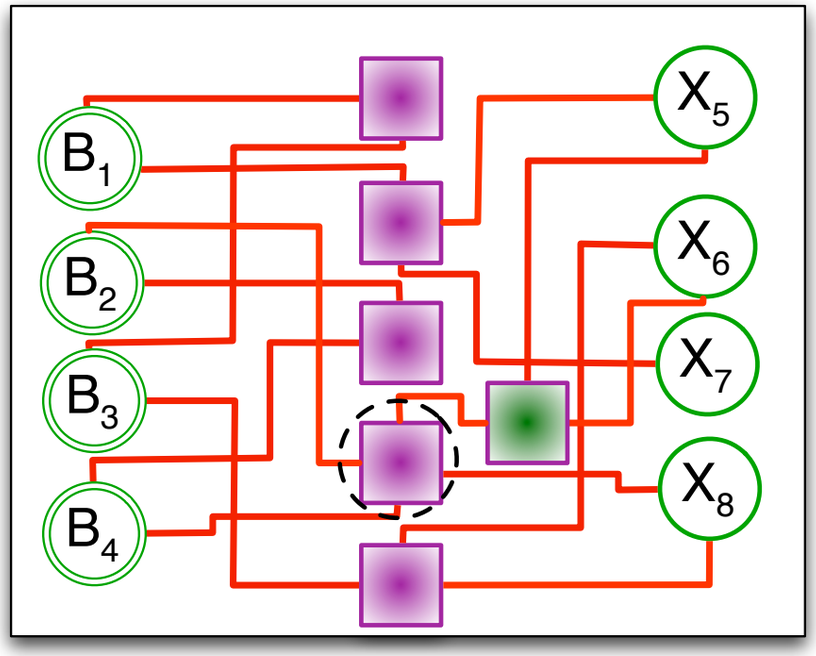

$$\alpha = 4$$

Can more complex nodes help?

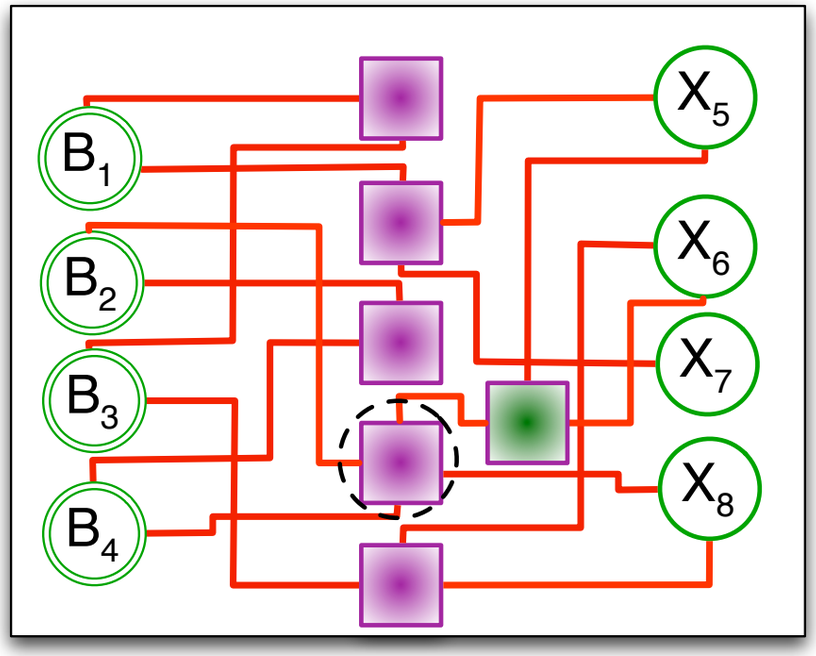

$$\alpha = 4$$

"Visible universe" of a bit is limited by the connectivity

Can more complex nodes help?

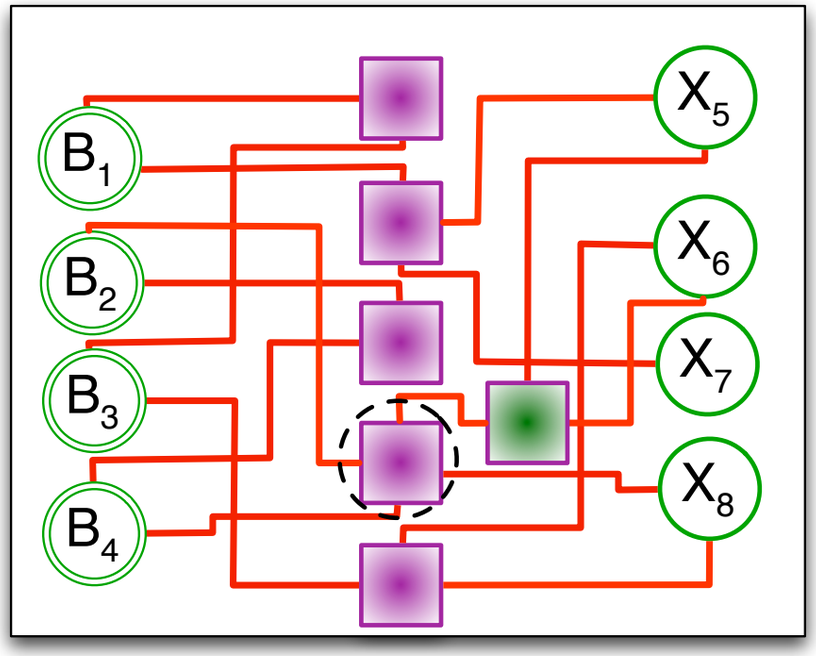

$$\alpha = 4$$

"Visible universe" of a bit is limited by the connectivity

Can more complex nodes help?

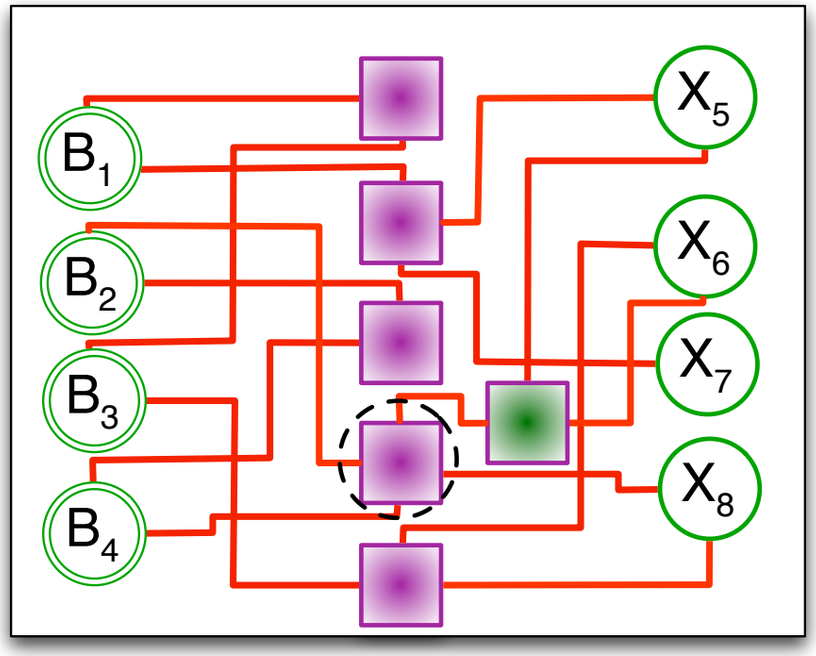

$$\alpha = 4$$

"Visible universe" of a bit is limited by the connectivity

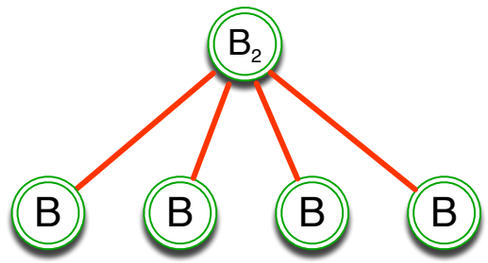

Can more complex nodes help?

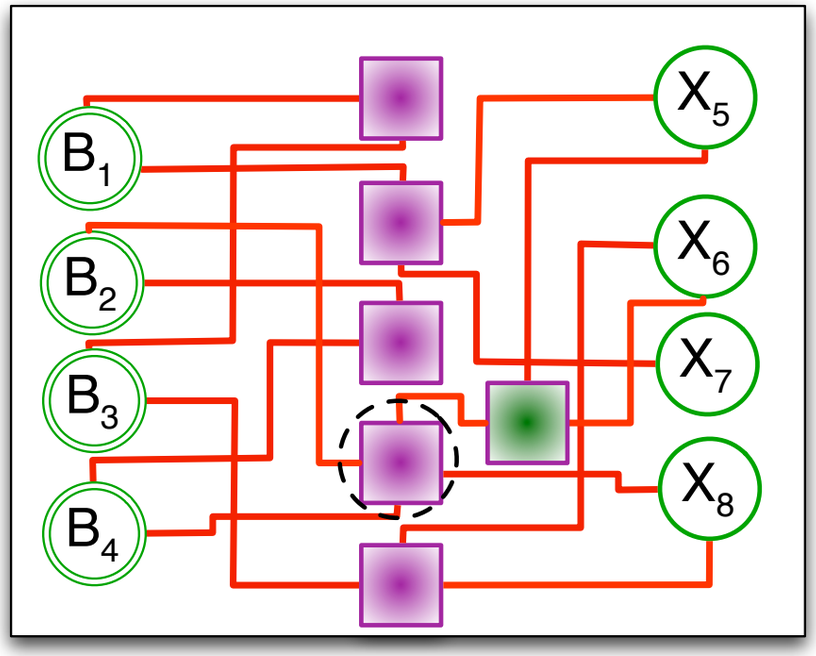

$\alpha = 4$

"Visible universe" of a bit is limited by the connectivity

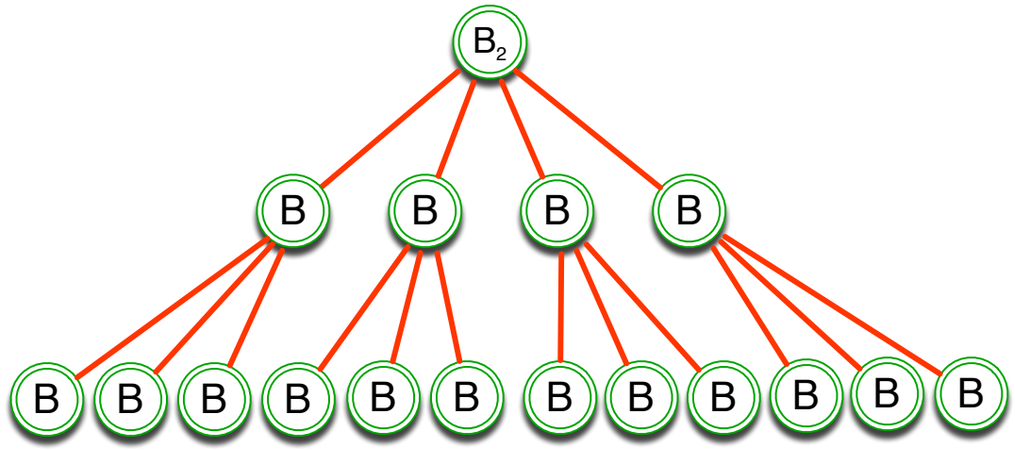

Can more complex nodes help?

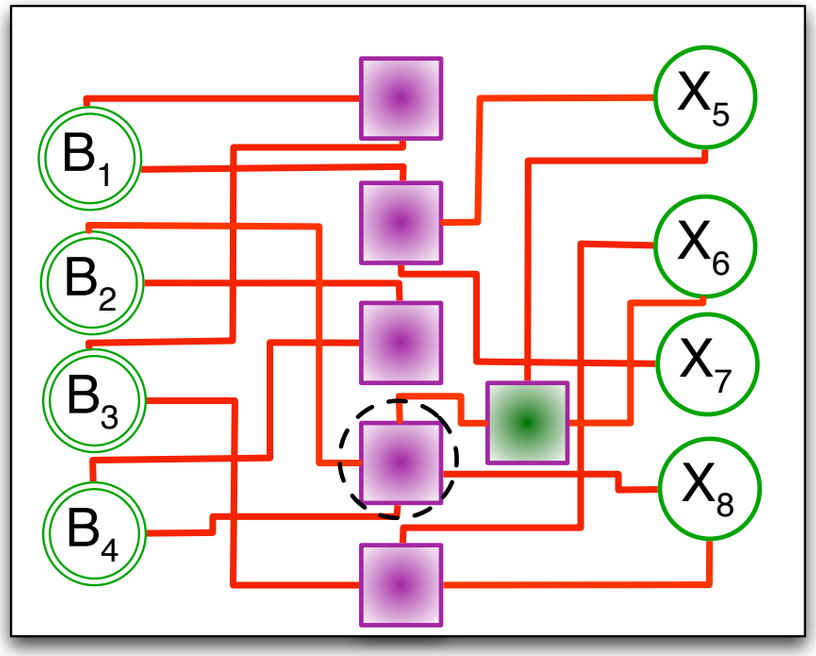

$\alpha = 4$

"Visible universe" of a bit is limited by the connectivity

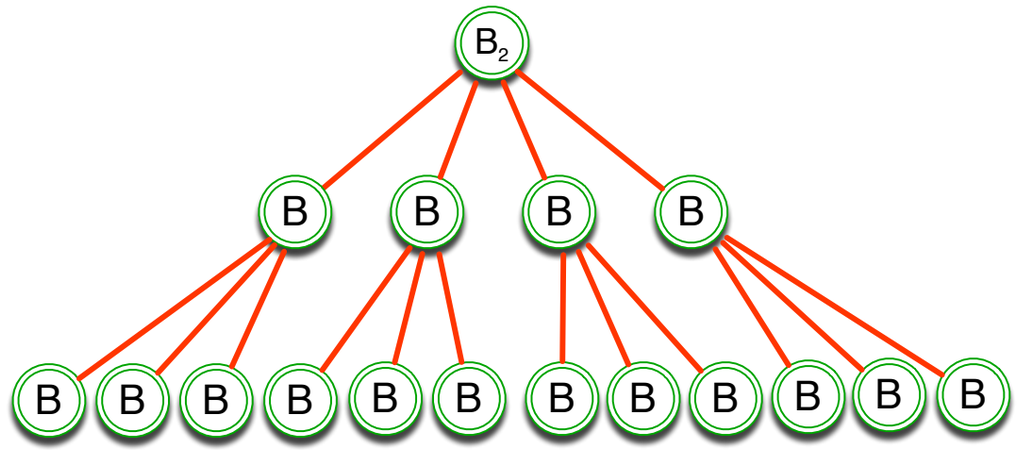

Theorem* [Grover, Sahai ISIT '08]

$$\text{iterations} \gtrsim \frac{1}{\log(\alpha - 1)} \log \left(\frac{\log \frac{1}{P_e}}{(C(P_T) - R)^2} \right)$$

* precise result for any P_e and any gap appears in [Grover, Woyach, Sahai '11]

Can more complex nodes help?

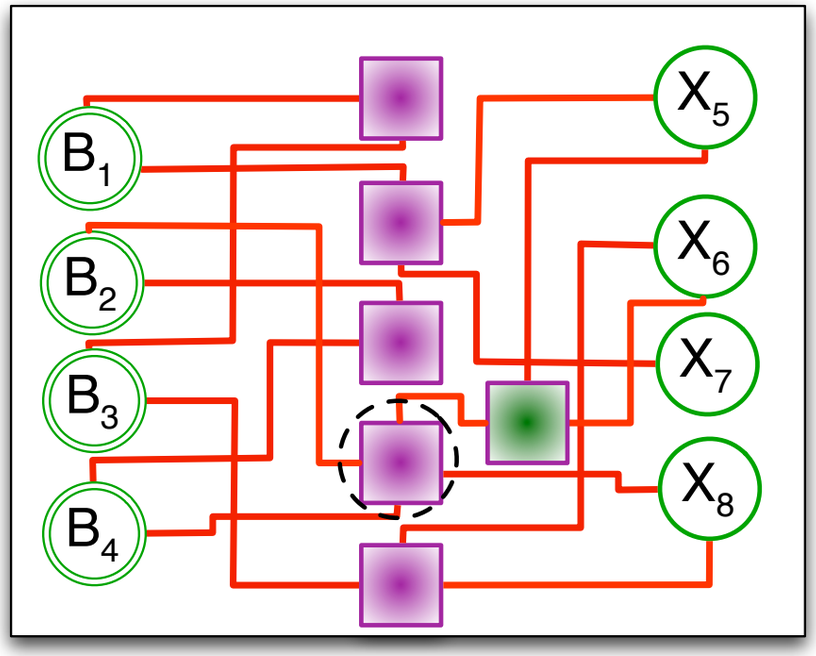

$\alpha = 4$

"Visible universe" of a bit is limited by the connectivity

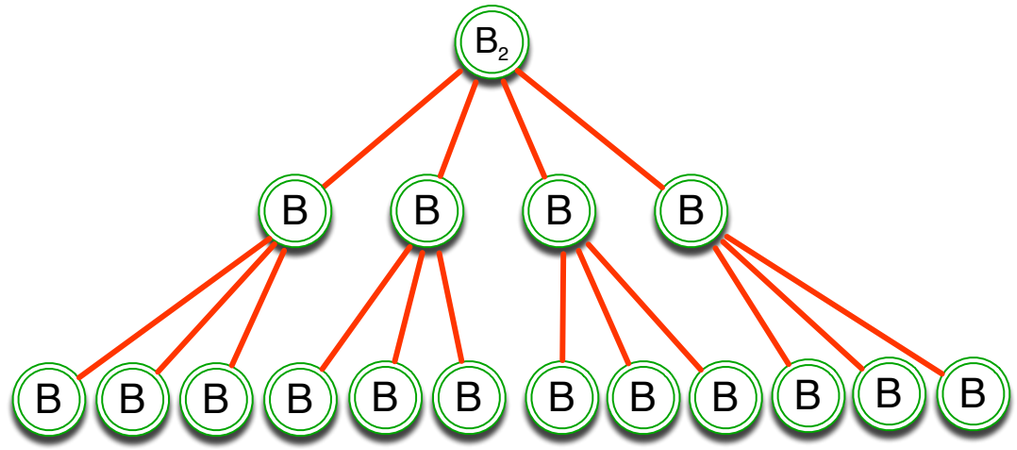

Theorem* [Grover, Sahai ISIT '08]

$$\text{iterations} \gtrsim \frac{1}{\log(\alpha - 1)} \log \left(\frac{\log \frac{1}{P_e}}{(C(P_T) - R)^2} \right)$$

. . . for any code, and any message-passing algorithm

* precise result for any P_e and any gap appears in [Grover, Woyach, Sahai '11]

Node model:
Fundamental bounds on total power

"Node model": Every node consumes constant energy per iteration

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

Node model:
Fundamental bounds on total power

"Node model": Every node consumes constant energy per iteration

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

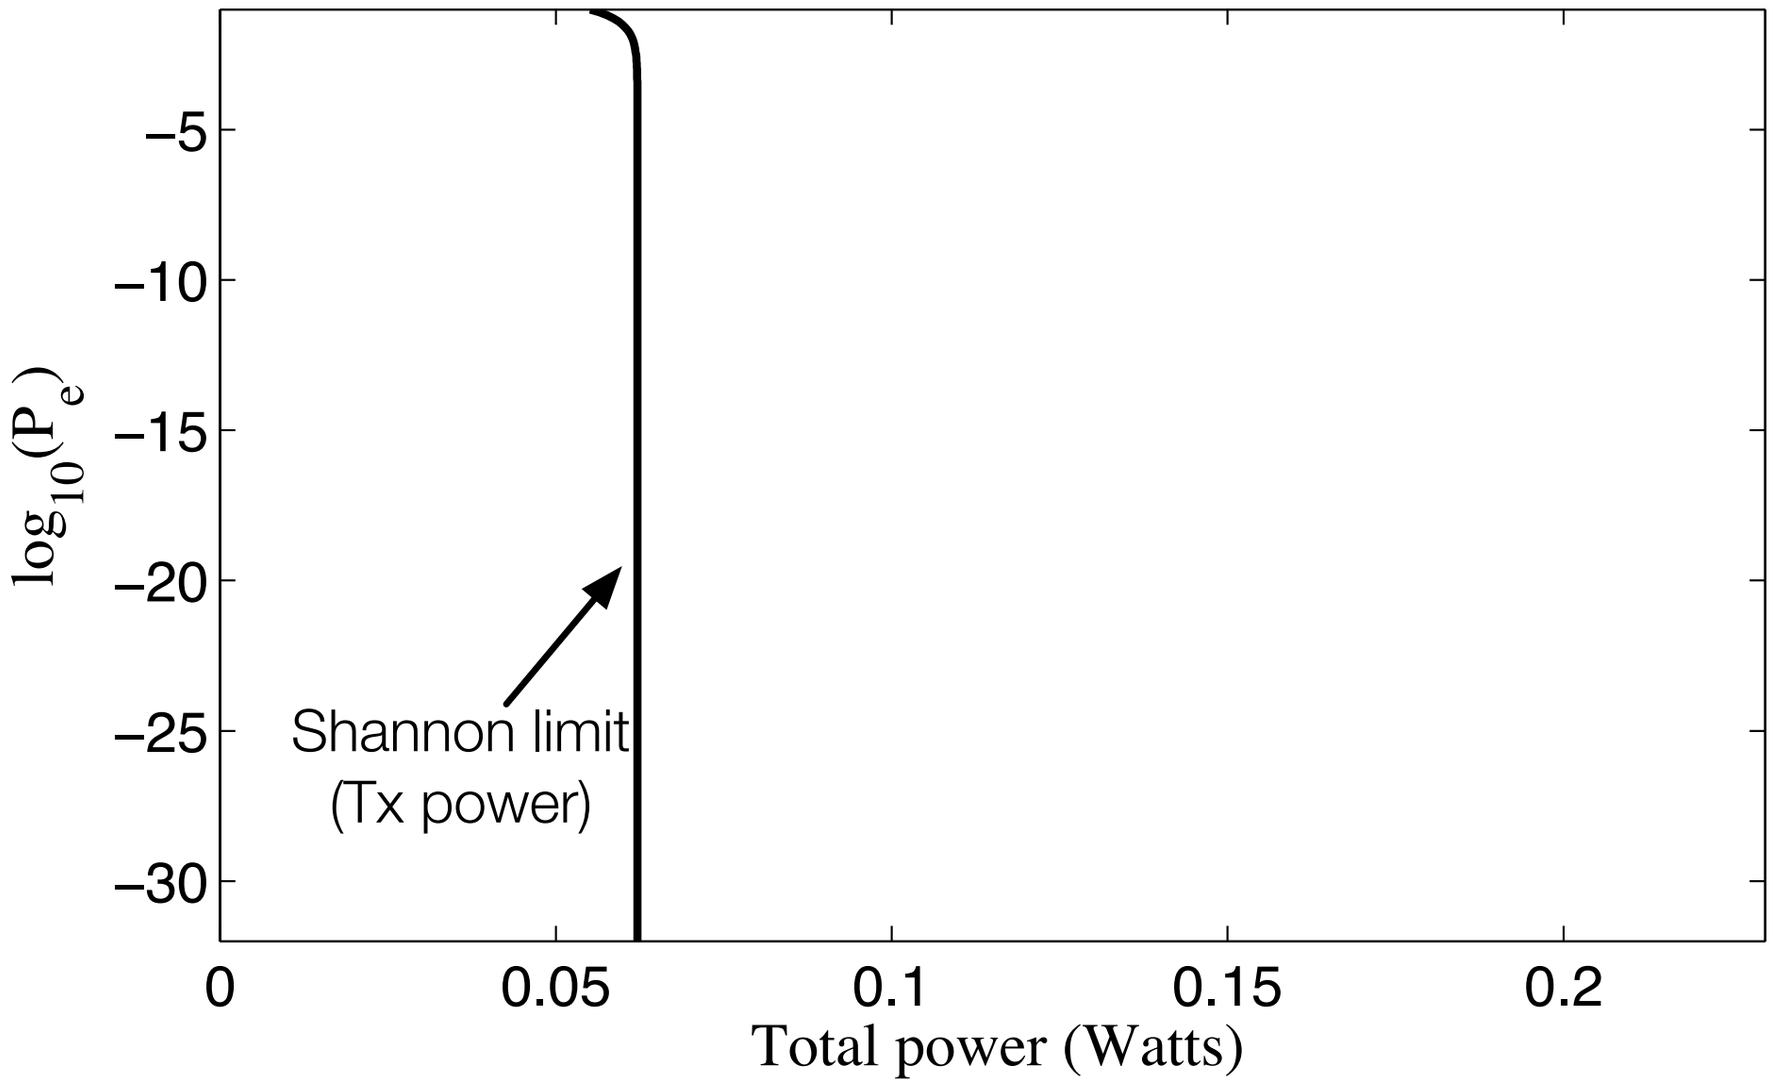

Node model:
Fundamental bounds on total power

"Node model": Every node consumes constant energy per iteration

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

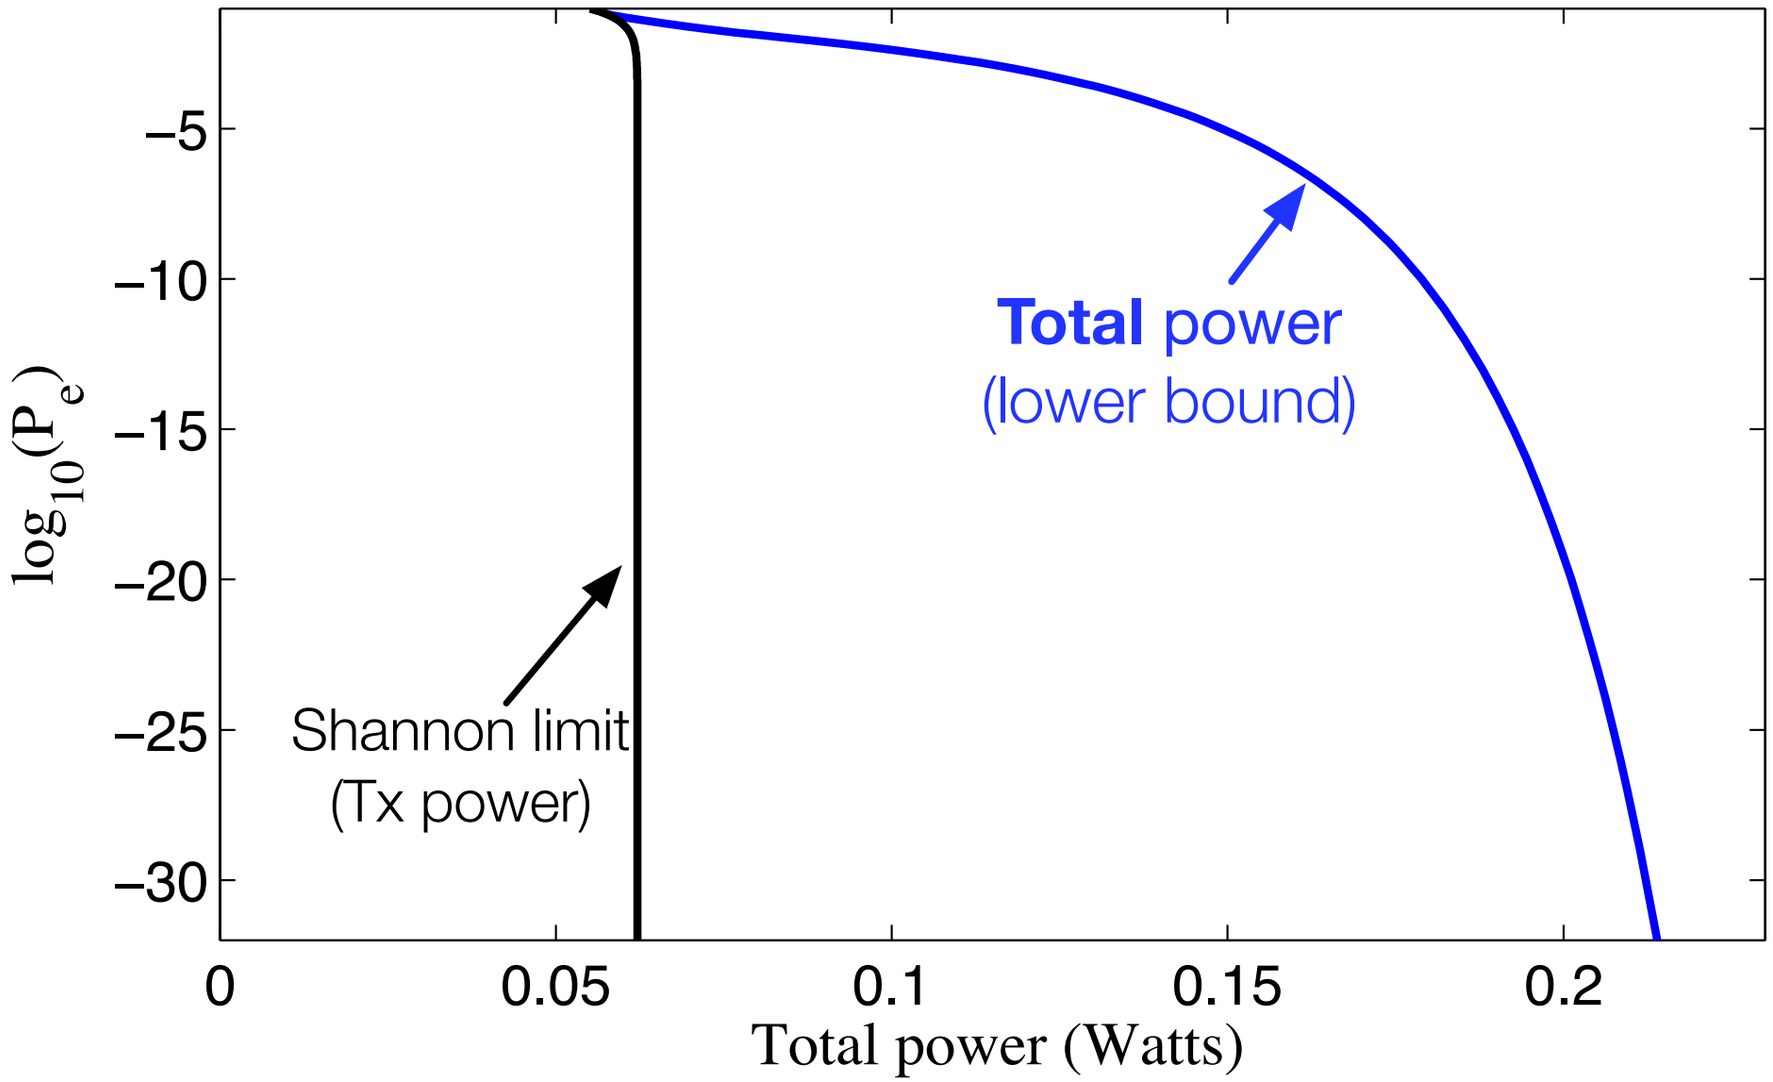

Node model:
Fundamental bounds on total power

"Node model": Every node consumes constant energy per iteration

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

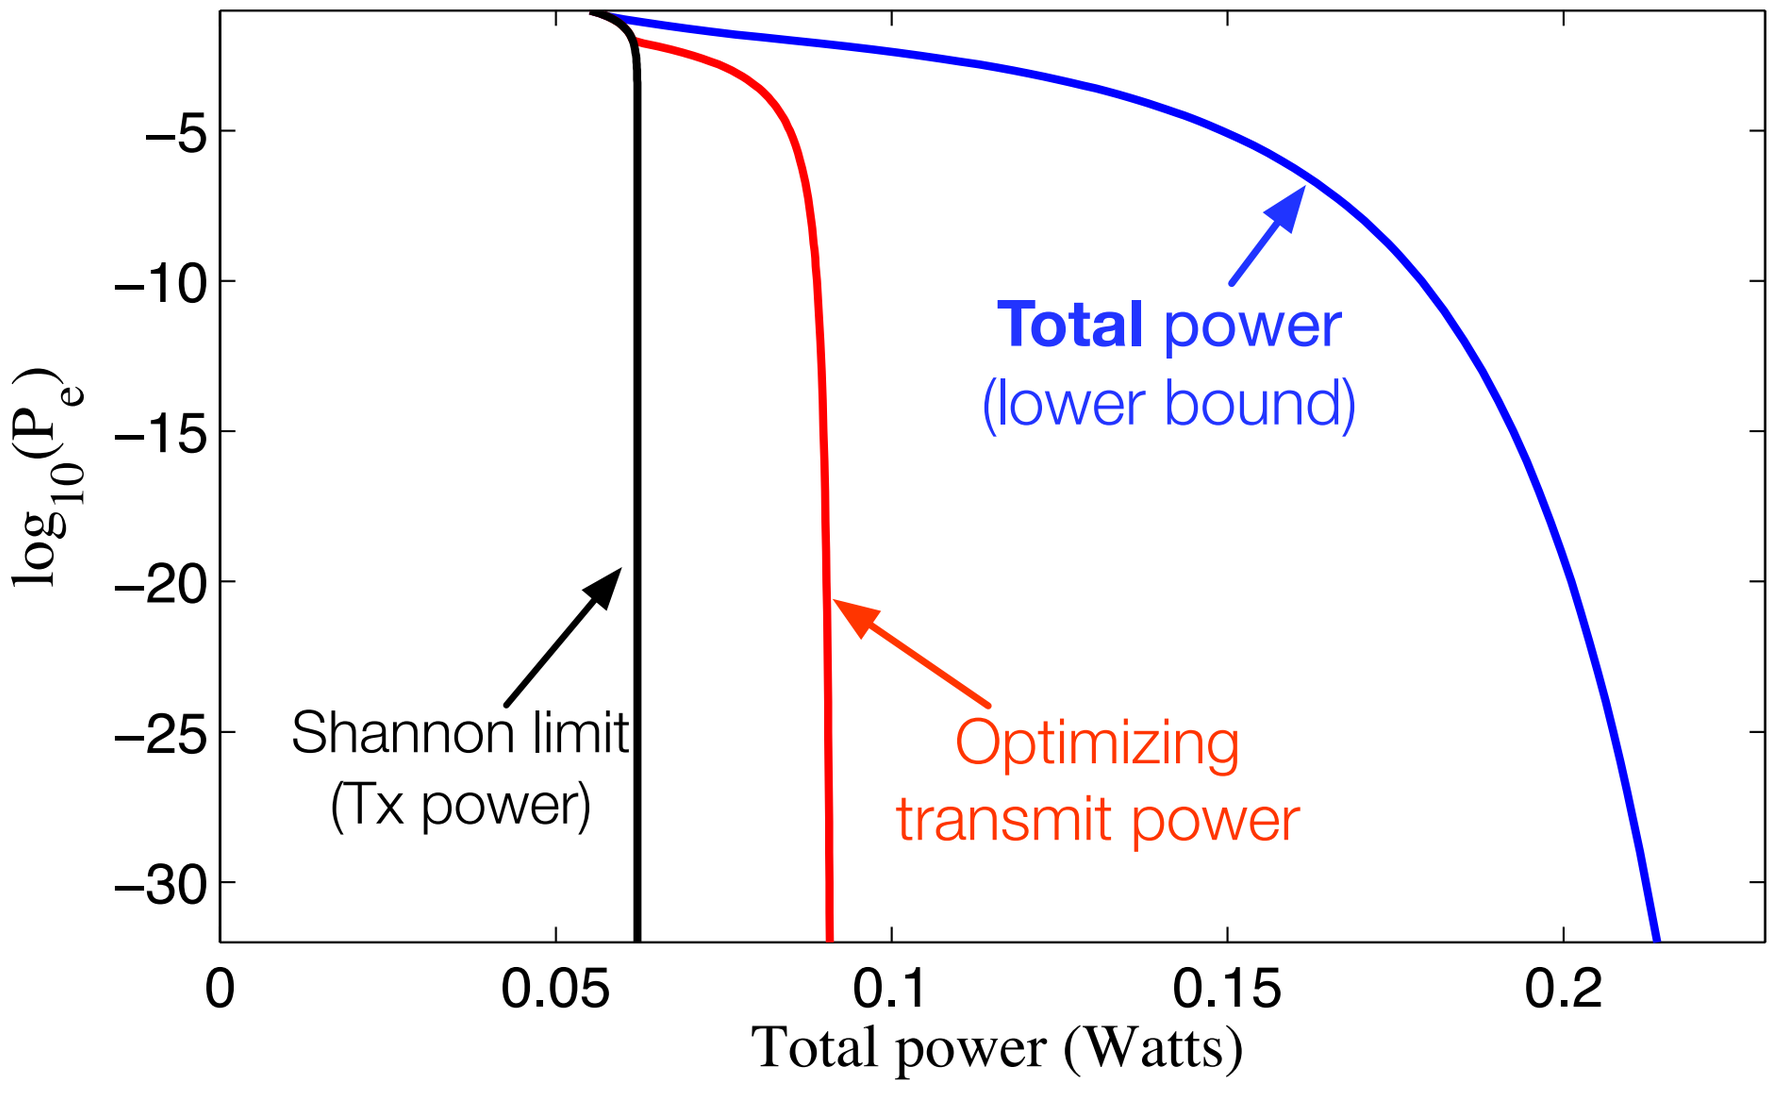

Node model:
Fundamental bounds on total power

"Node model": Every node consumes constant energy per iteration

$$P_{total} = \min_{P_T} P_T + P_{decoding}$$

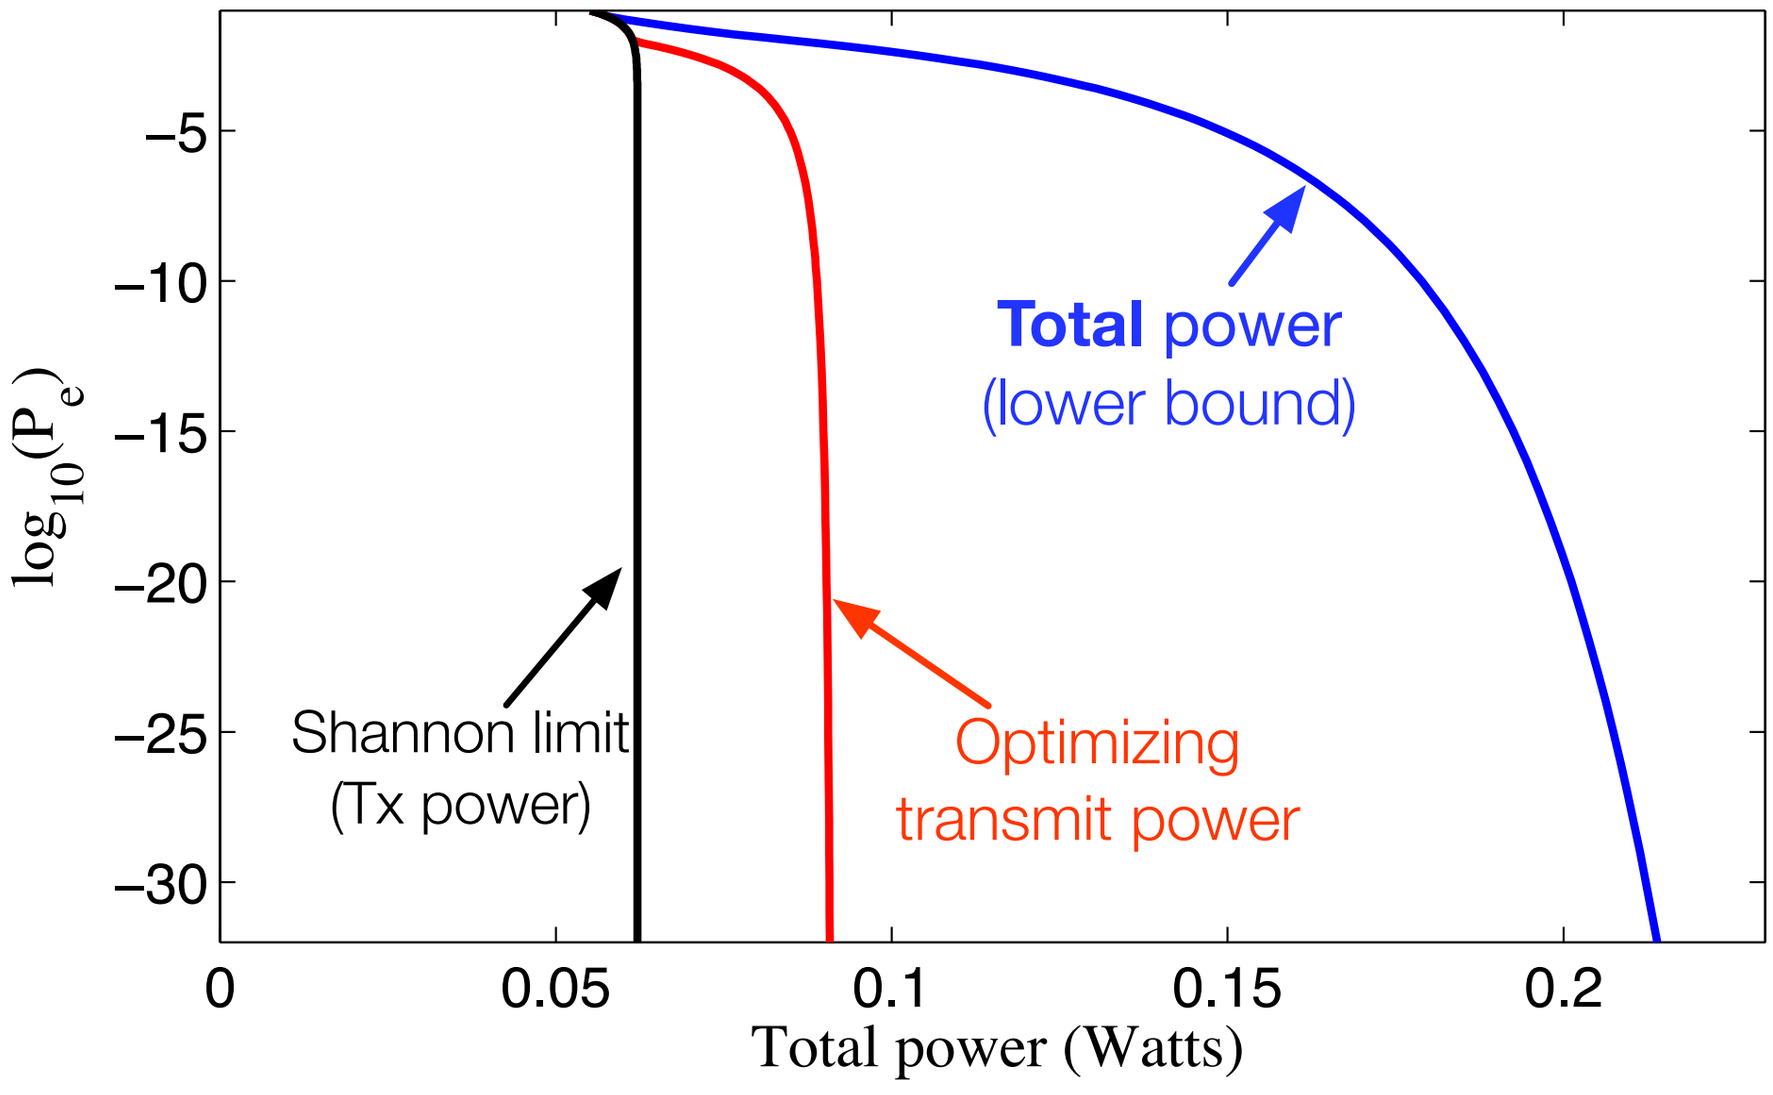

Moral: do not operate too close to capacity

Conclusions

Conclusions

- Decoding power can be large

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation

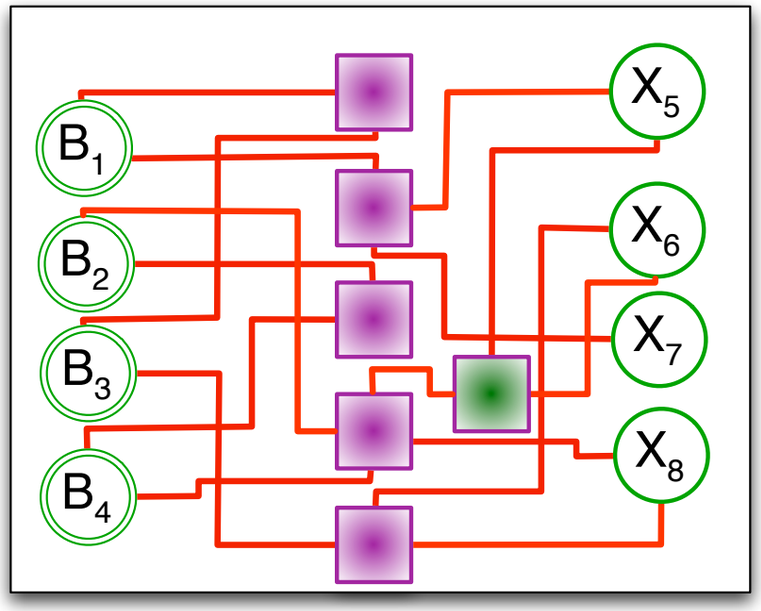

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation
- “Wire/Node model”

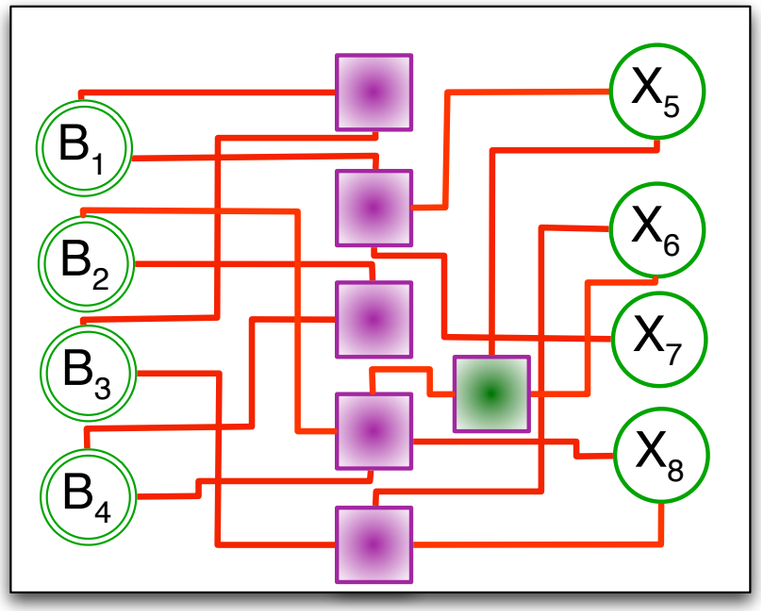

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation
- “Wire/Node model”
  - stay away from capacity!

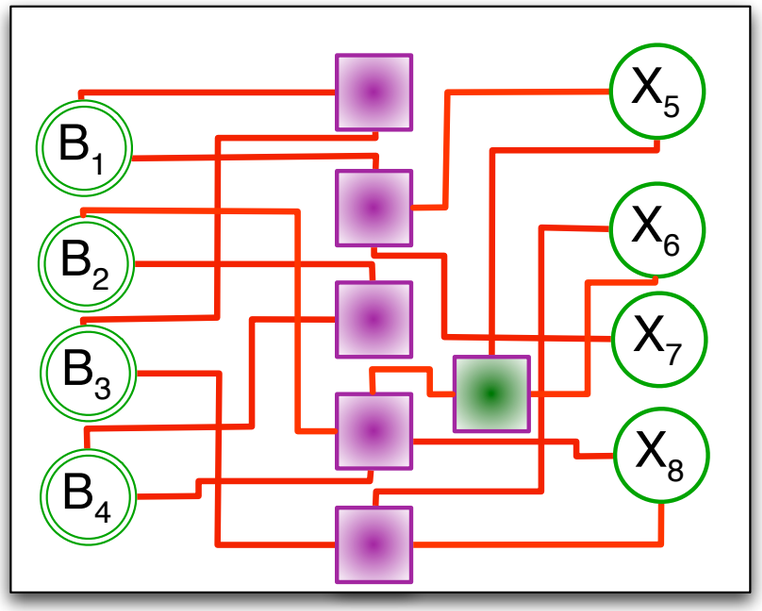

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation
- "Wire/Node model"
  - stay away from capacity!
  - design codes/decoders for just the required performance (no better!)

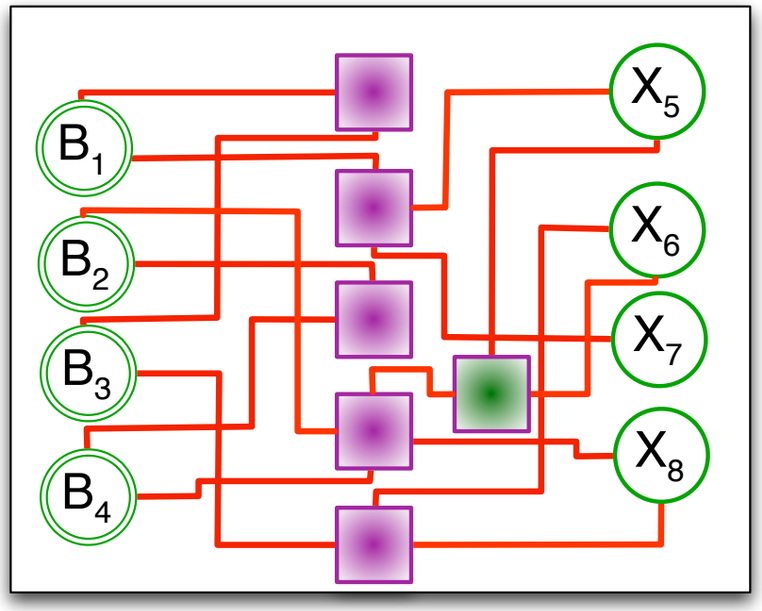

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation
- "Wire/Node model"
  - stay away from capacity!
  - design codes/decoders for just the required performance (no better!)
- Other architectures? (e.g. not fully parallel, memory access, sleeping nodes?)

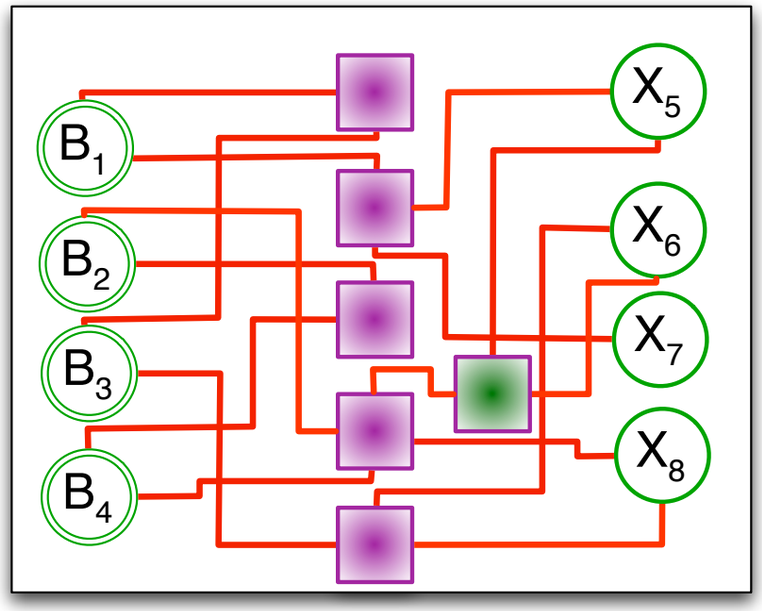

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation
- "Wire/Node model"
  - stay away from capacity!
  - design codes/decoders for just the required performance (no better!)
- Other architectures? (e.g. not fully parallel, memory access, sleeping nodes?)
- Multiple users, equalization, ADC?

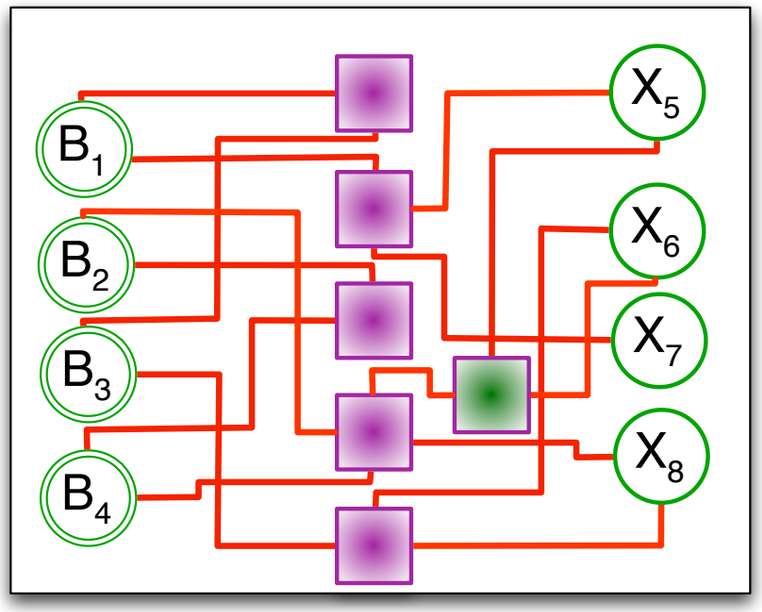

Conclusions

- Decoding power can be large
- VLSI model of decoding implementation
- "Wire/Node model"
  - stay away from capacity!
  - design codes/decoders for just the required performance (no better!)
- Other architectures? (e.g. not fully parallel, memory access, sleeping nodes?)
- Multiple users, equalization, ADC?

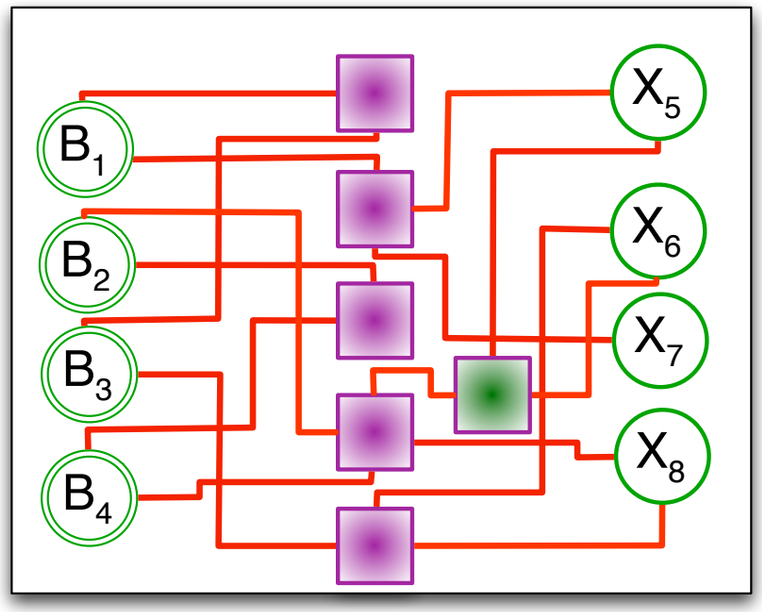

Handouts/papers/summary: http://www.eecs.berkeley.edu/~pulkit

Capacity-approaching LDPCs or regular LDPCs?

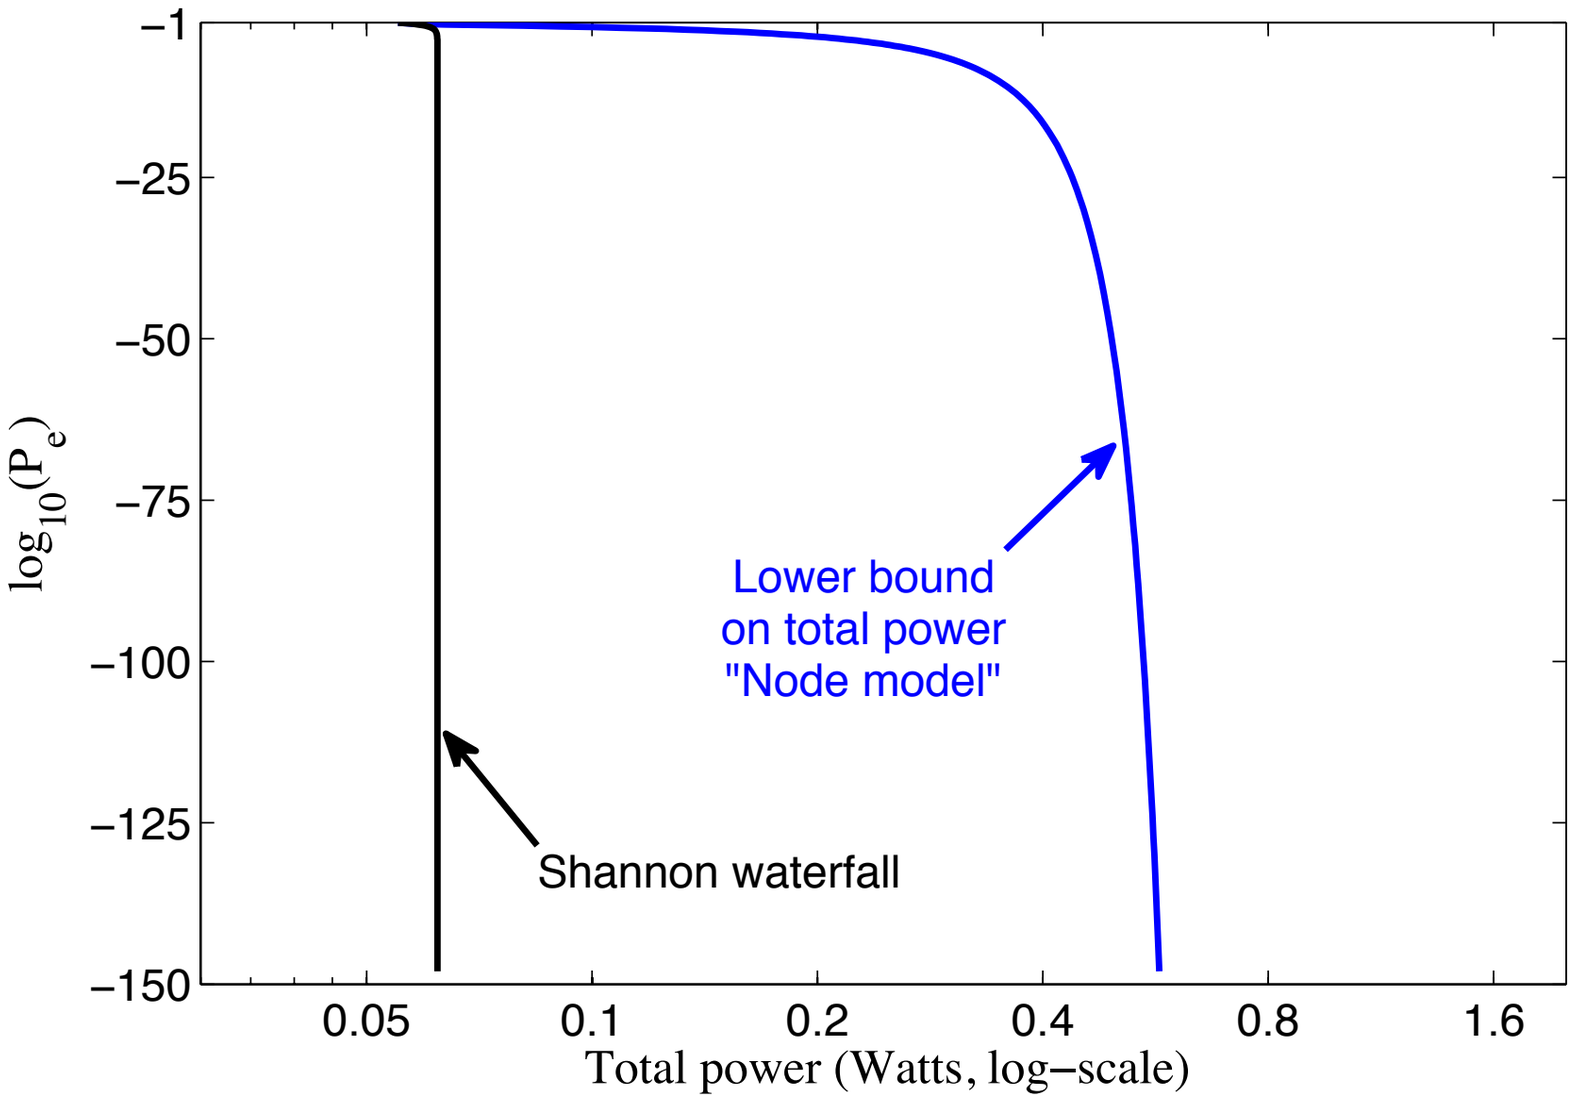

energy per operation E = 3 pJ,
distance = 10 m,
f = 3 GHz
path-loss exponent = 3,
maximum node connectivity= 4
T = 300 K
Rate = 1/4 bits/ch-use

[Grover, Sahai '09]

Capacity-approaching LDPCs or regular LDPCs?

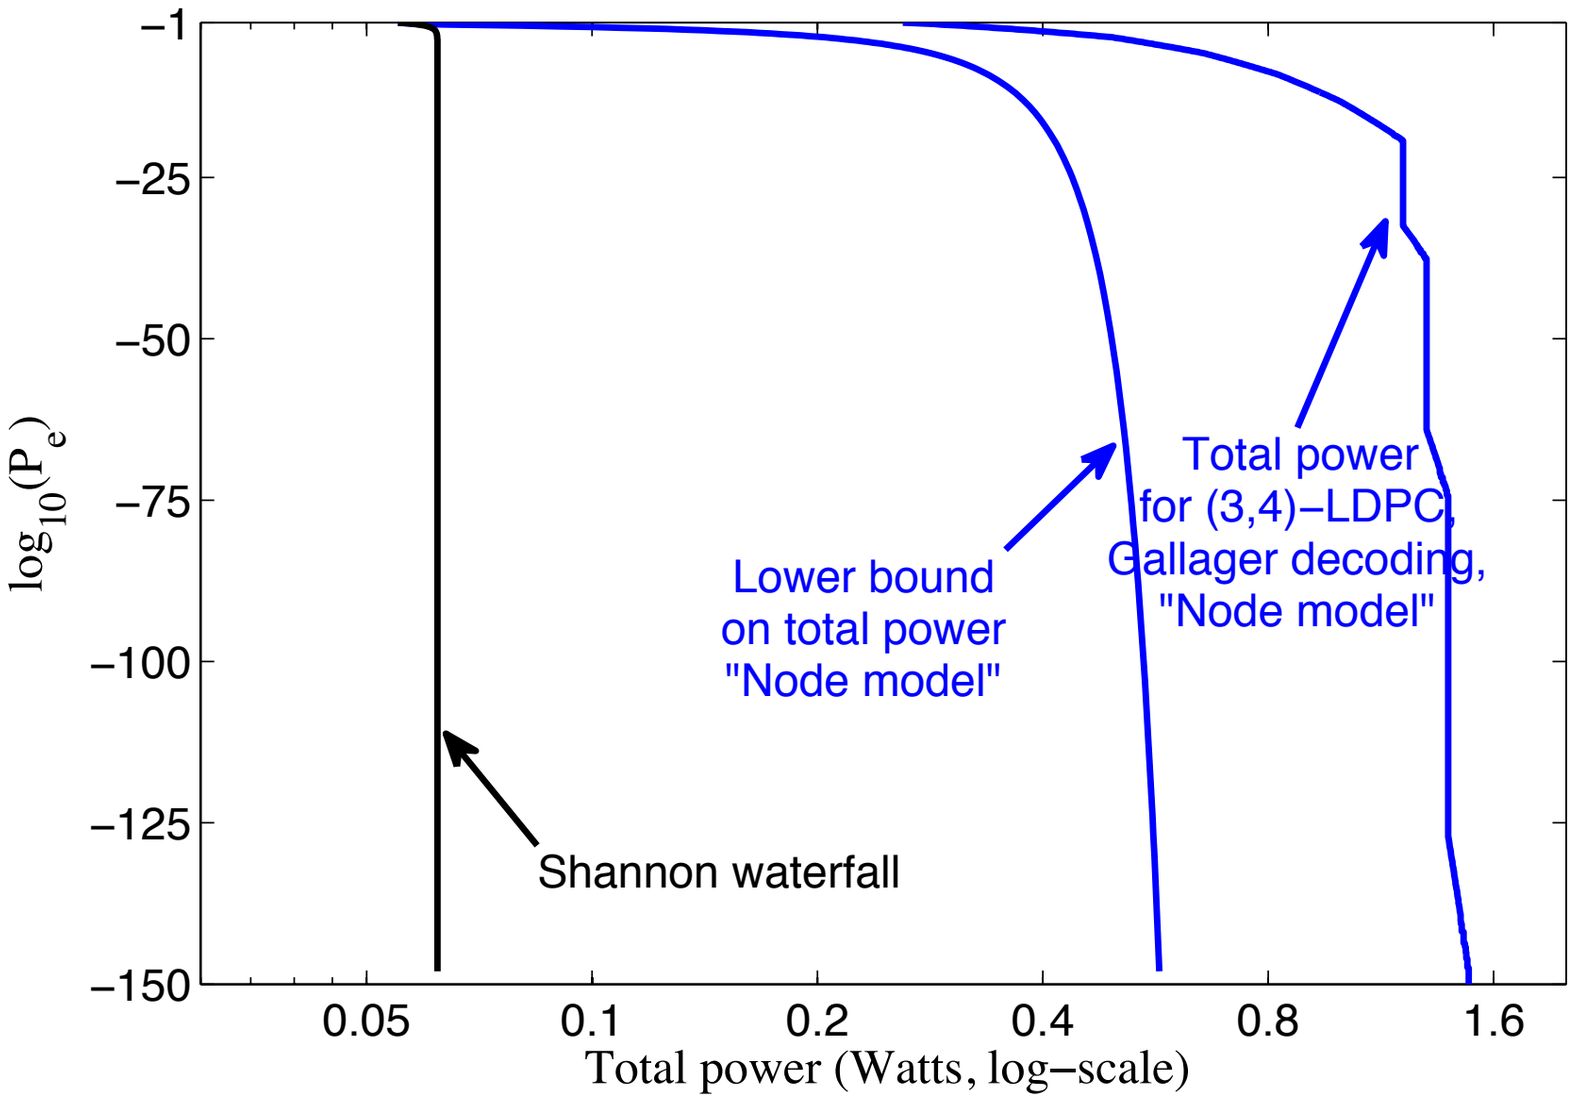

energy per operation E = 3 pJ,
distance = 10 m,
f = 3 GHz
path-loss exponent = 3,
maximum node connectivity= 4
T = 300 K
Rate = 1/4 bits/ch-use

[Grover, Sahai '09]

Capacity-approaching LDPCs or regular LDPCs?

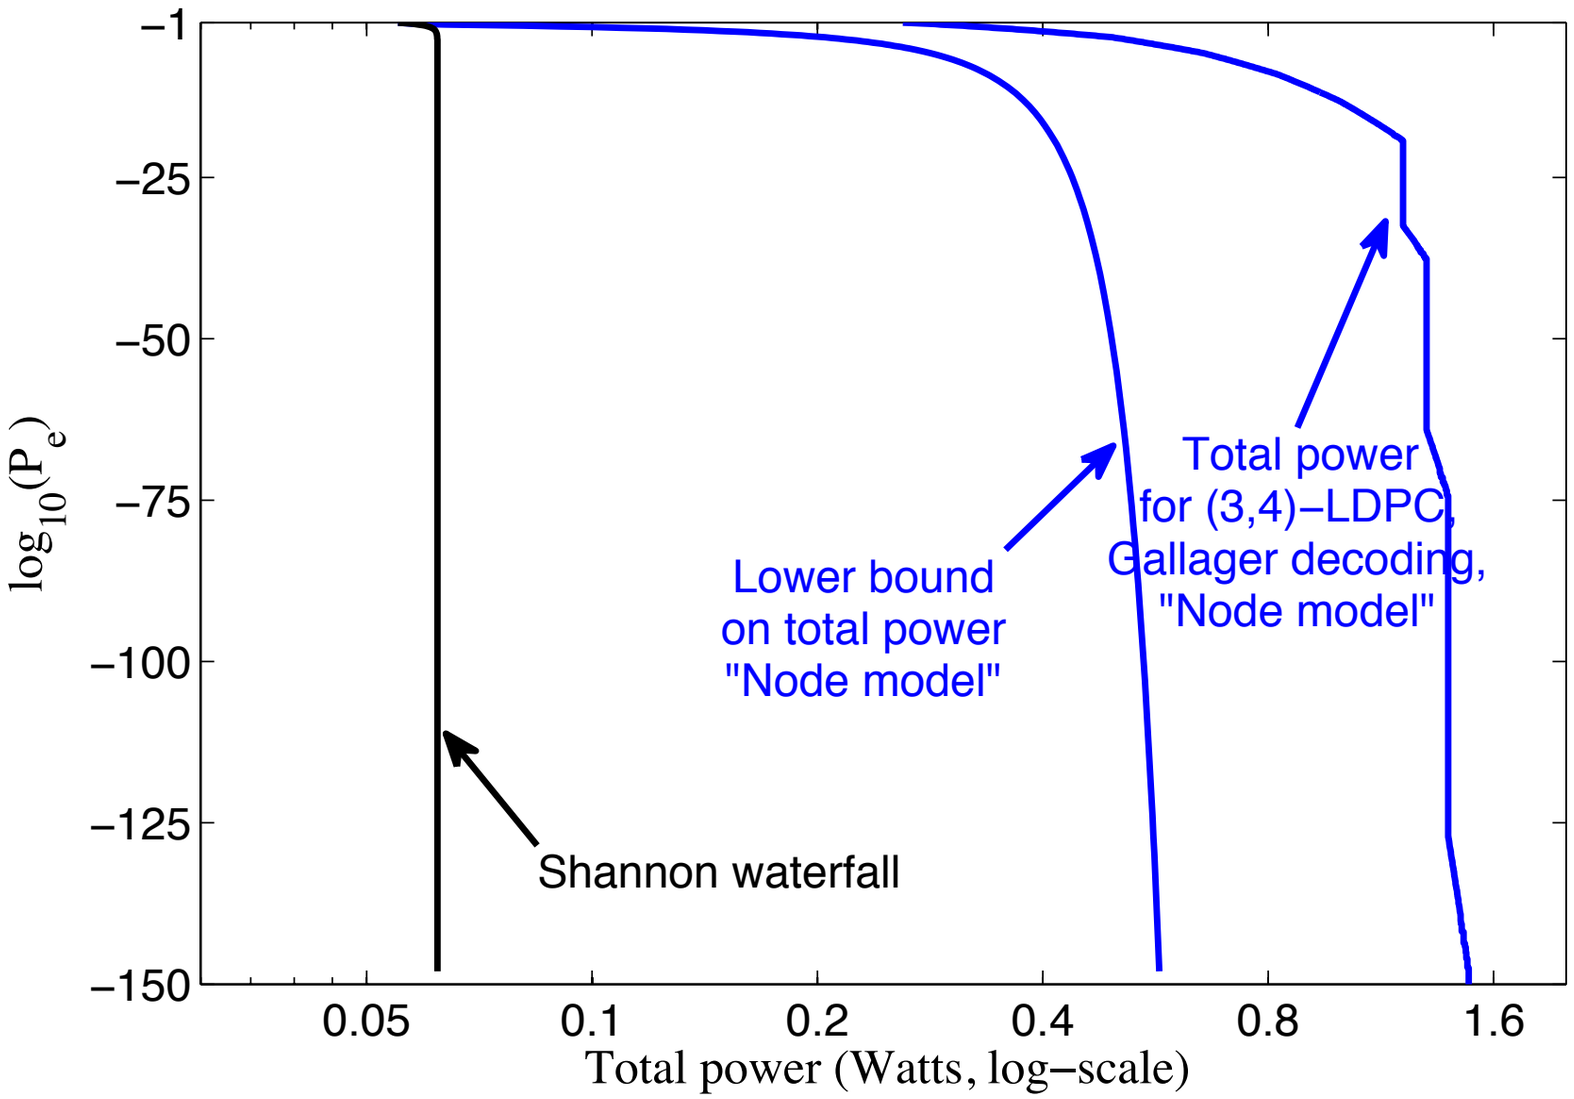

energy per operation E = 3 pJ,
distance = 10 m,
f = 3 GHz
path-loss exponent = 3,
maximum node connectivity= 4
T = 300 K
Rate = 1/4 bits/ch-use

[Grover, Sahai '09]

Morals: 1. Regular LDPCs (not capacity-approaching LDPCs) are order optimal!

Capacity-approaching LDPCs or regular LDPCs?

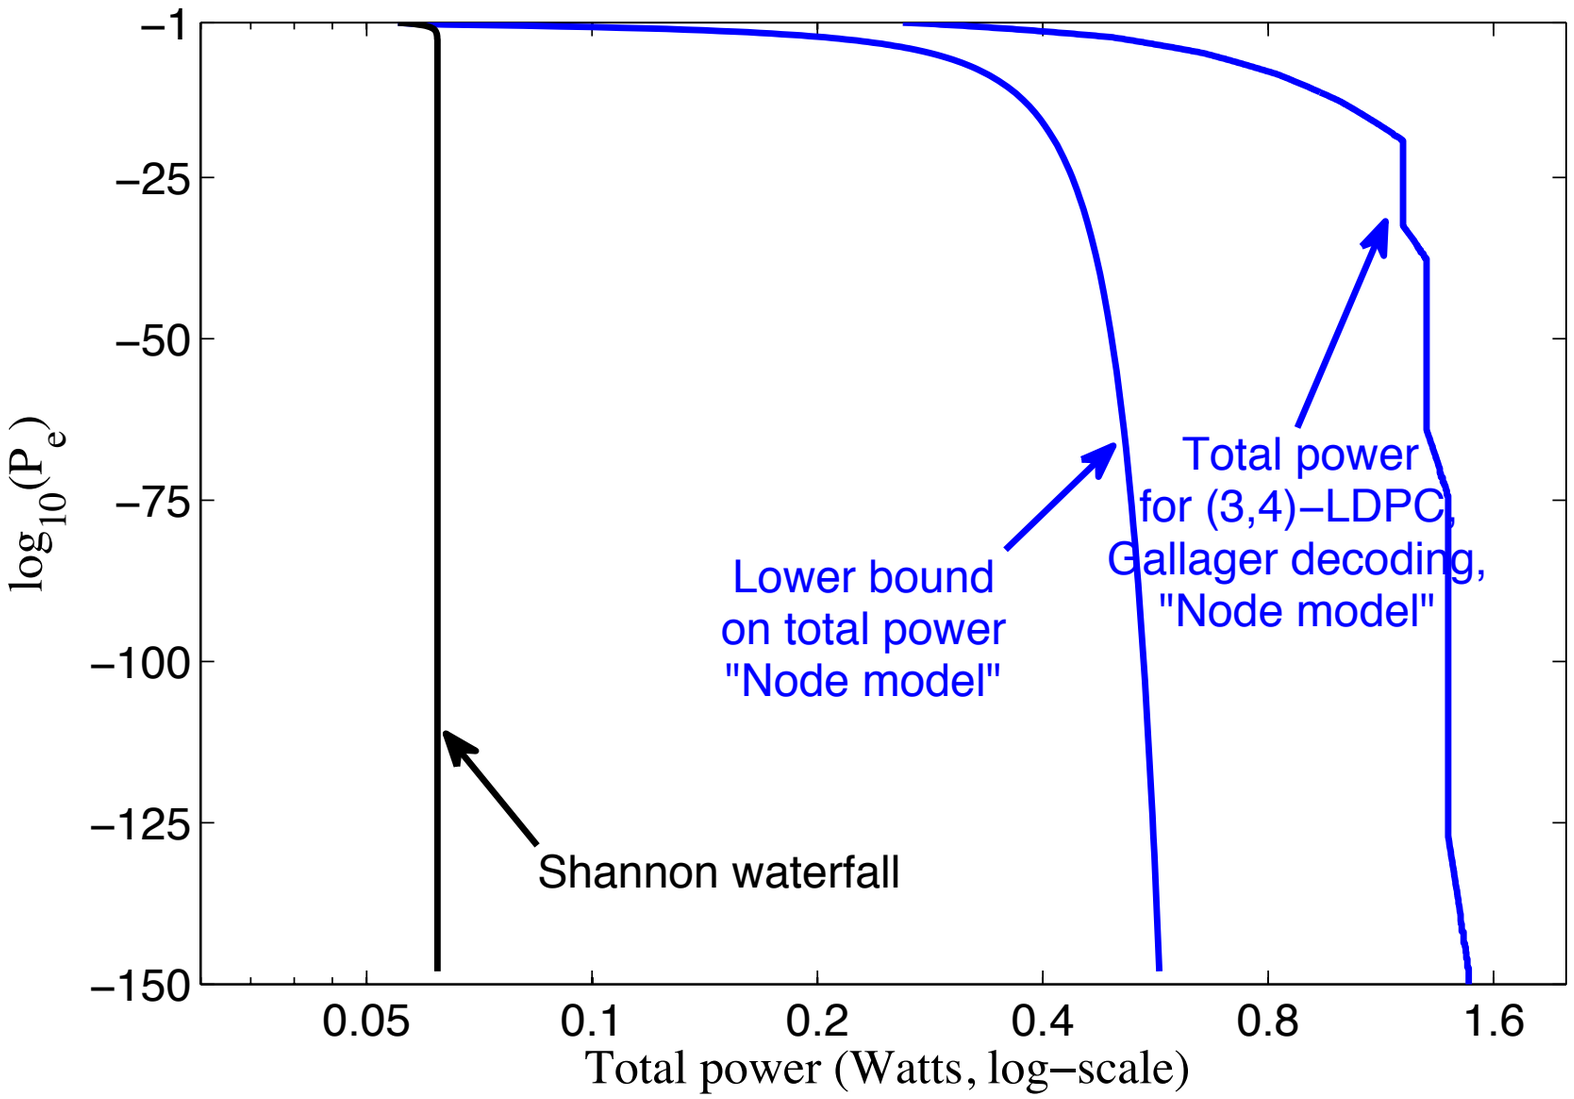

energy per operation E = 3 pJ,
distance = 10 m,
f = 3 GHz
path-loss exponent = 3,
maximum node connectivity= 4
T = 300 K
Rate = 1/4 bits/ch-use

[Grover, Sahai '09]

Morals: 1. Regular LDPCs (not capacity-approaching LDPCs) are order optimal!

2. need new joint code/decoder designs

How much *density* of links can you attain?

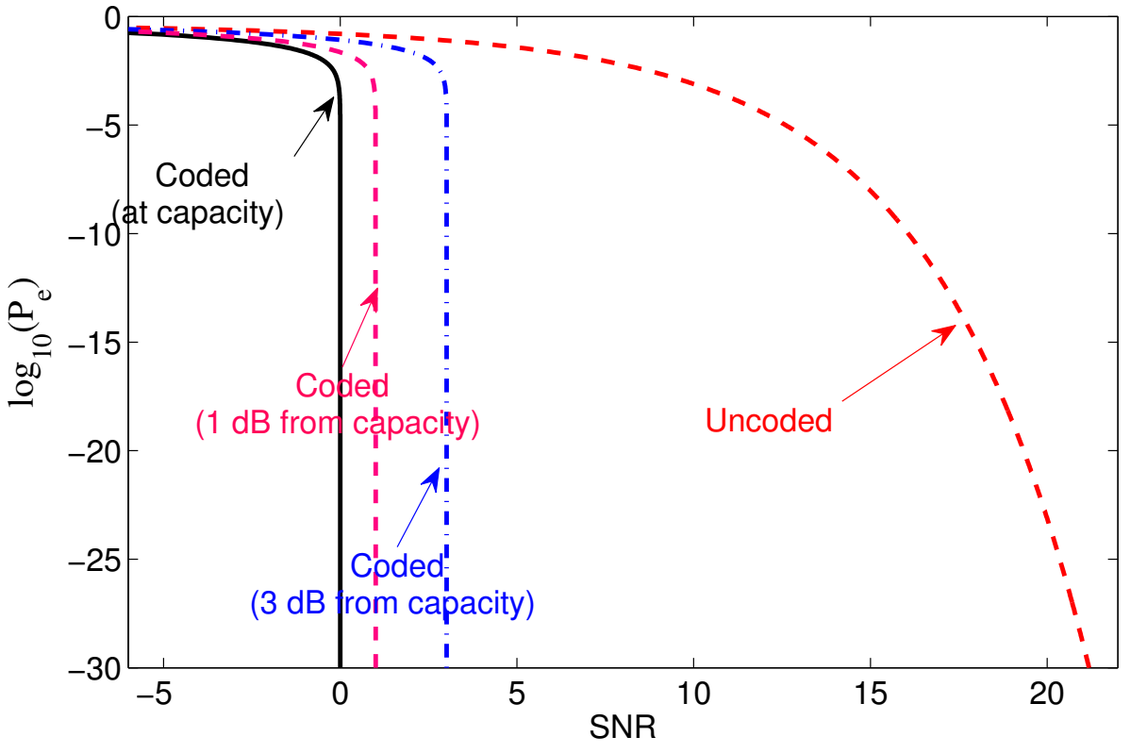

Traditional picture: Waterfall

How much *density* of links can you attain?

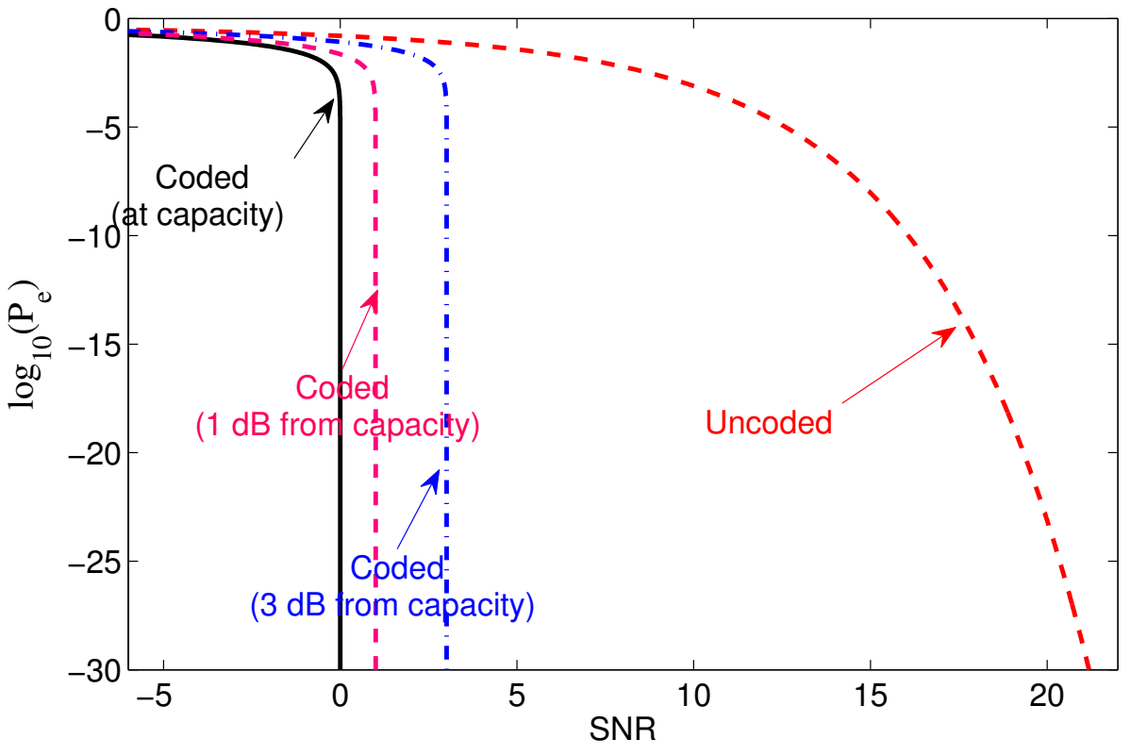

Traditional picture: Waterfall

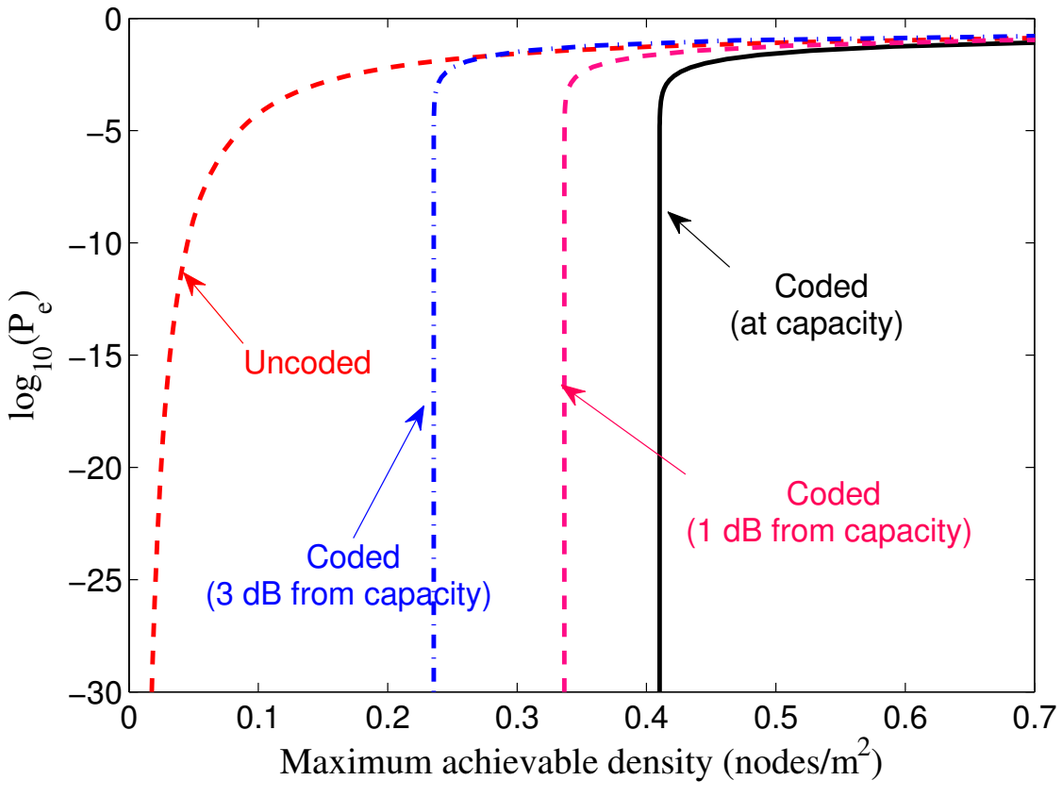

Density picture
[Grover, Woyach, Palaiyanur, Sahai ’10]

“Wire model” of power consumption

$$P \propto \frac{1}{2}\text{Capacitance} \times V^2 \times t$$

“Wire model” of power consumption

$$P \propto \frac{1}{2}\text{Capacitance} \times V^2 \times t$$

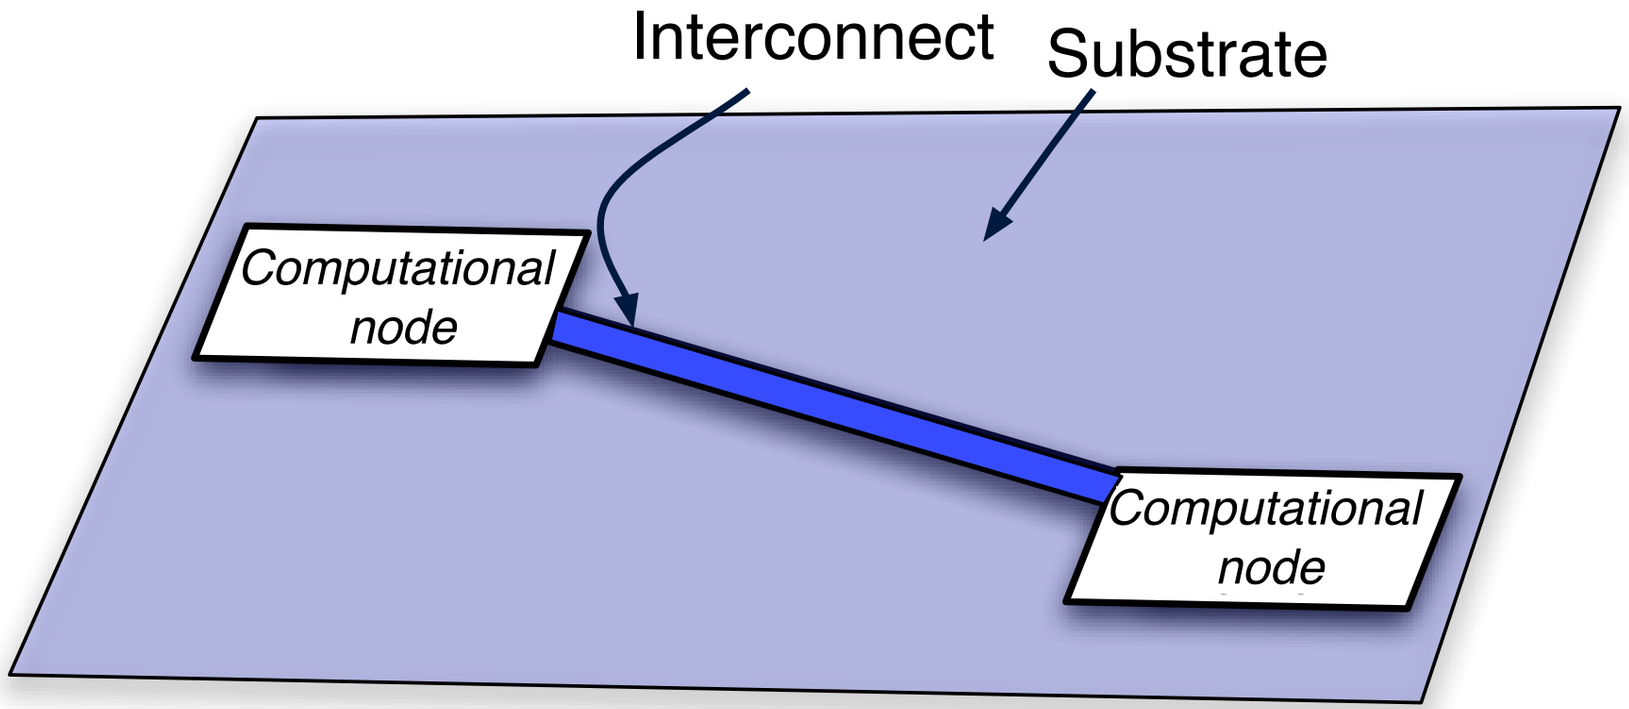

“Wire model” of power consumption

$$P \propto \frac{1}{2}\text{Capacitance} \times V^2 \times t$$

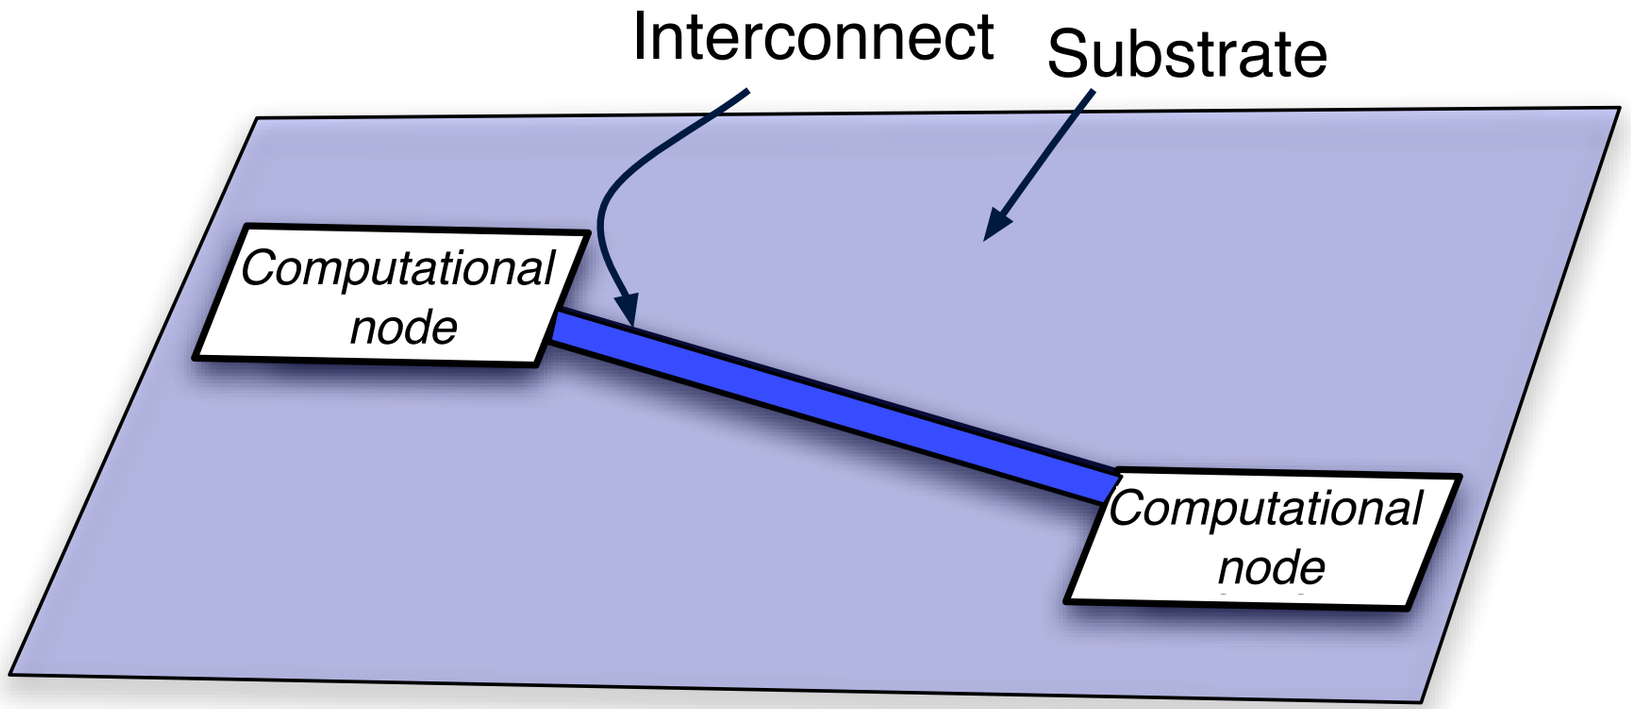

Capacitance $\propto$ Wire length

“Wire model” of power consumption

$$P \propto \frac{1}{2}\text{Capacitance} \times V^2 \times t$$

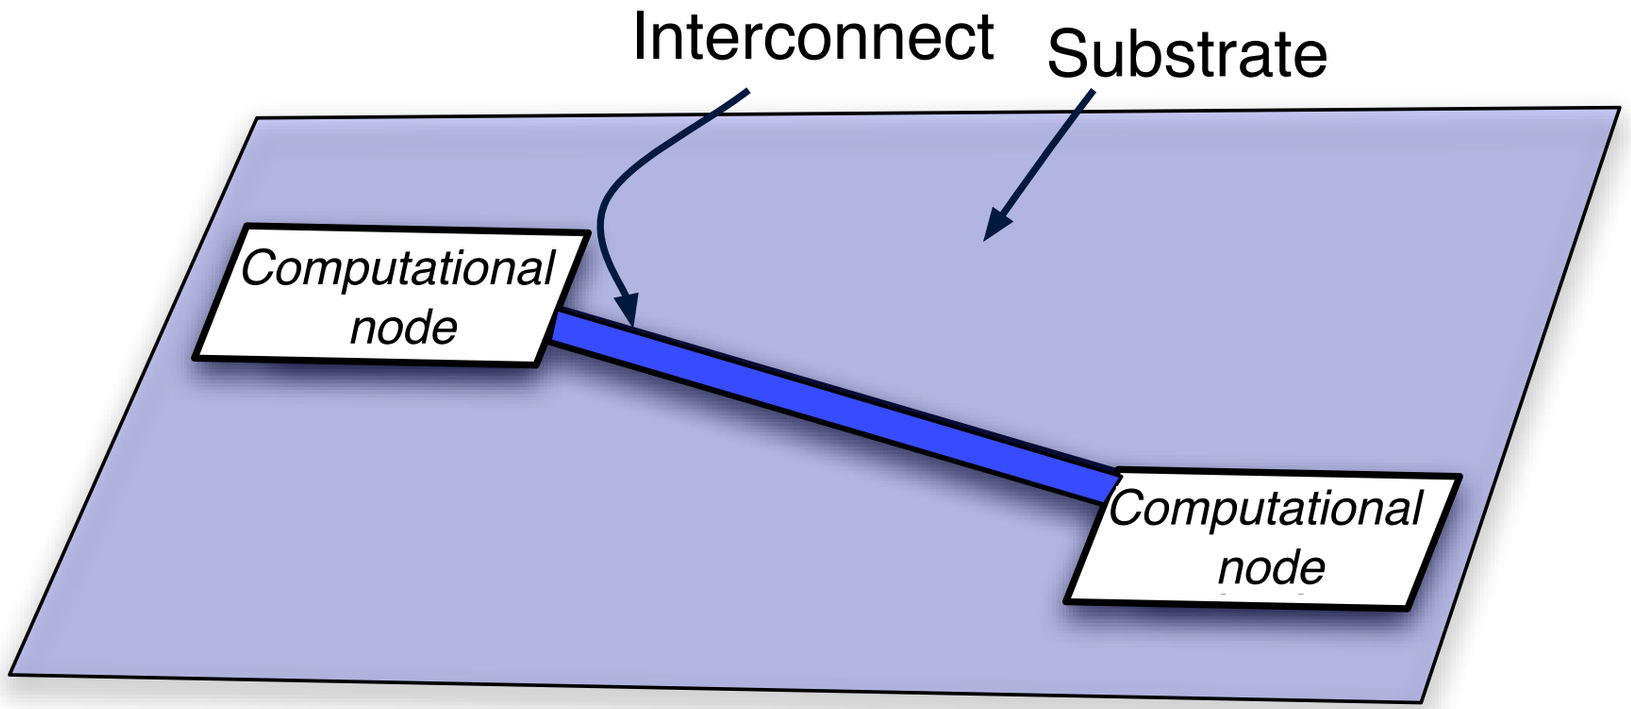

$$\text{Capacitance} \propto \text{Wire length}$$

$$\text{Thus, } P \propto t \times \text{Wire length}$$

Ongoing work

- Multiuser setups:
  - broadcast [Grover, Sahai ISIT '09]
  - collection of point-to-point links [Grover, Woyach, Sahai JSAC '11]
- Accounting for receiver power beyond just the decoding power
  - ADC, equalizer, etc. Understanding various tradeoffs. [ongoing with Nikolic, Park]

Reducing power consumption:
Option 2: “black-box” abstractions for processing power

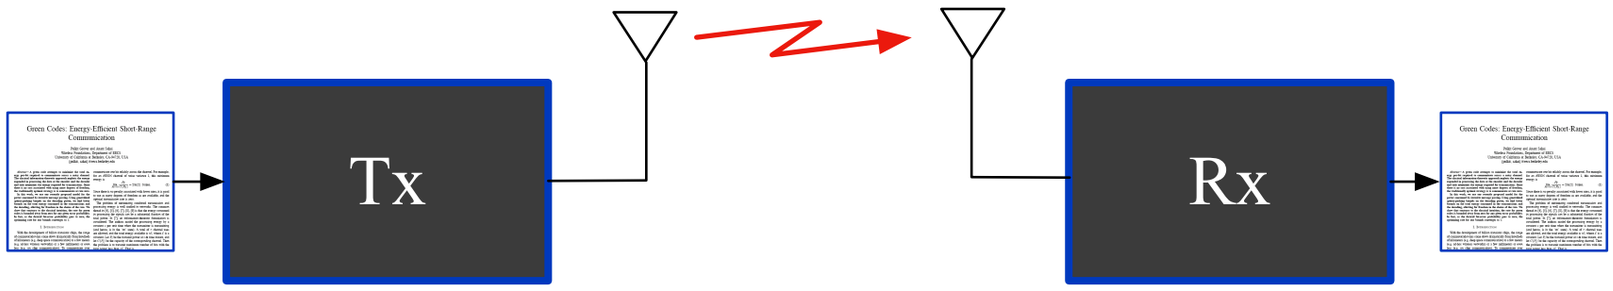

The Tx and Rx consume a fixed amount of power when “on”

[Cui, Goldsmith Bahai]
[Massaad, Medard, Zheng]
[Prabhakaran]

Reducing power consumption:
Option 2: “black-box” abstractions for processing power

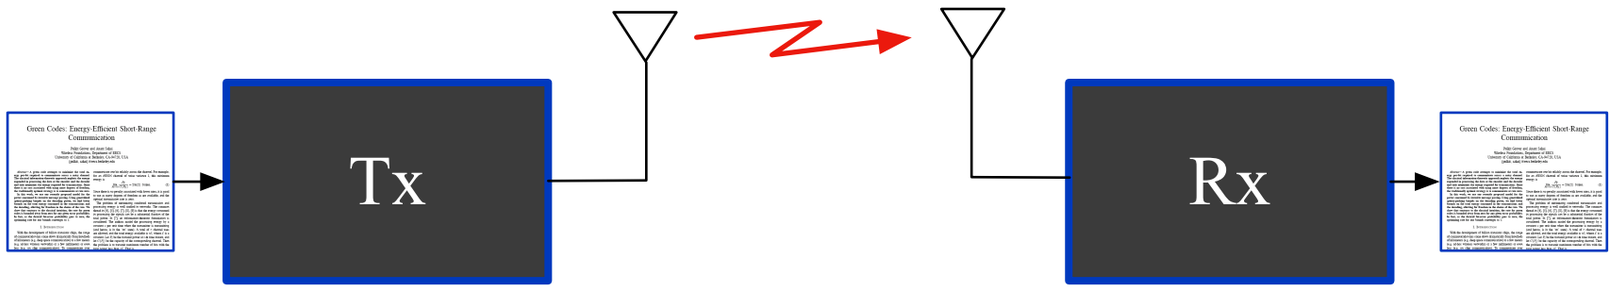

The Tx and Rx consume a fixed amount of power when “on”

[Cui, Goldsmith Bahai]
[Massaad, Medard, Zheng]
[Prabhakaran]

upshot: bursty transmissions (“go to sleep”)

Decoders with short wires: a purely implementation issue?

Decoders with short wires: a purely implementation issue?

Comm

Optimize performance

Decoders with short wires: a purely implementation issue?

Comm

Optimize performance

Circuits

Minimize wire lengths

Decoders with short wires: a purely implementation issue?

Comm

Optimize performance

Circuits

Minimize wire lengths

Theorem [Grover, Sahai '11]

For a regular (c,d)-LDPC code

$$W_{max} \geq \frac{\sqrt{A_{node}}}{\sqrt{\pi}\left(\frac{g}{2}-1\right)}\left((d_v-1)^{\lfloor \frac{g}{4}-\frac{1}{2} \rfloor}(d_c-1)^{\lfloor \frac{g}{4}-\frac{3}{4} \rfloor}\right)$$

Decoders with short wires: a purely implementation issue?

Comm

Optimize performance

Circuits

Minimize wire lengths

Theorem [Grover, Sahai '11]

For a regular (c,d)-LDPC code

$$W_{max} \geq \frac{\sqrt{A_{node}}}{\sqrt{\pi}\left(\frac{g}{2}-1\right)}\left((d_v-1)^{\lfloor \frac{g}{4}-\frac{1}{2} \rfloor}(d_c-1)^{\lfloor \frac{g}{4}-\frac{3}{4} \rfloor}\right)$$

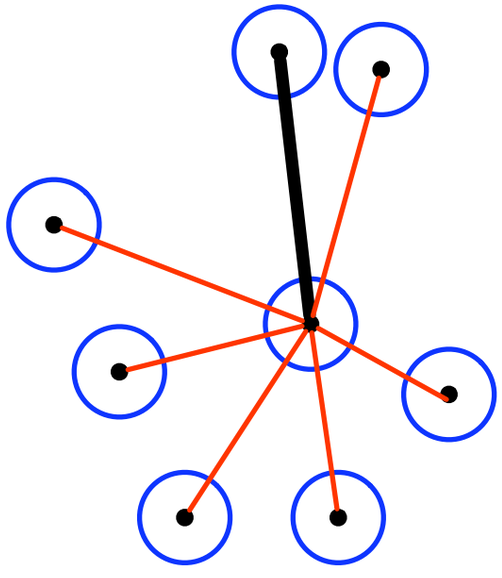

Decoders with short wires: a purely implementation issue?

Comm

Optimize performance

Circuits

Minimize wire lengths

Theorem [Grover, Sahai '11]

For a regular (c,d)-LDPC code

$$W_{max} \geq \frac{\sqrt{A_{node}}}{\sqrt{\pi}\left(\frac{g}{2}-1\right)}\left((d_v-1)^{\lfloor \frac{g}{4}-\frac{1}{2} \rfloor}(d_c-1)^{\lfloor \frac{g}{4}-\frac{3}{4} \rfloor}\right)$$

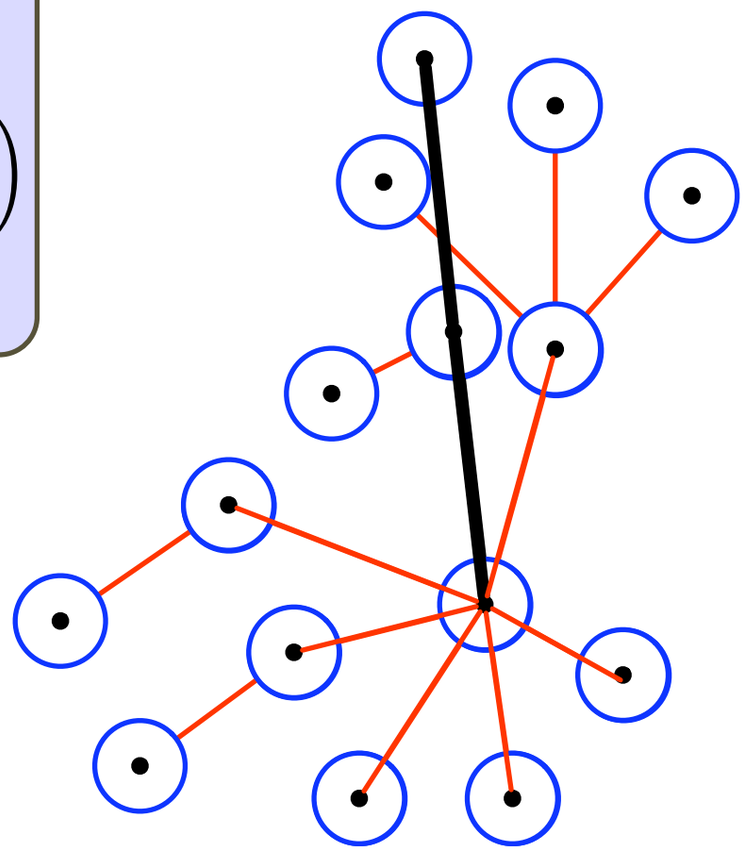

Decoders with short wires: a purely implementation issue?

Comm

Optimize performance

Circuits

Minimize wire lengths

Theorem [Grover, Sahai '11]

For a regular (c,d)-LDPC code
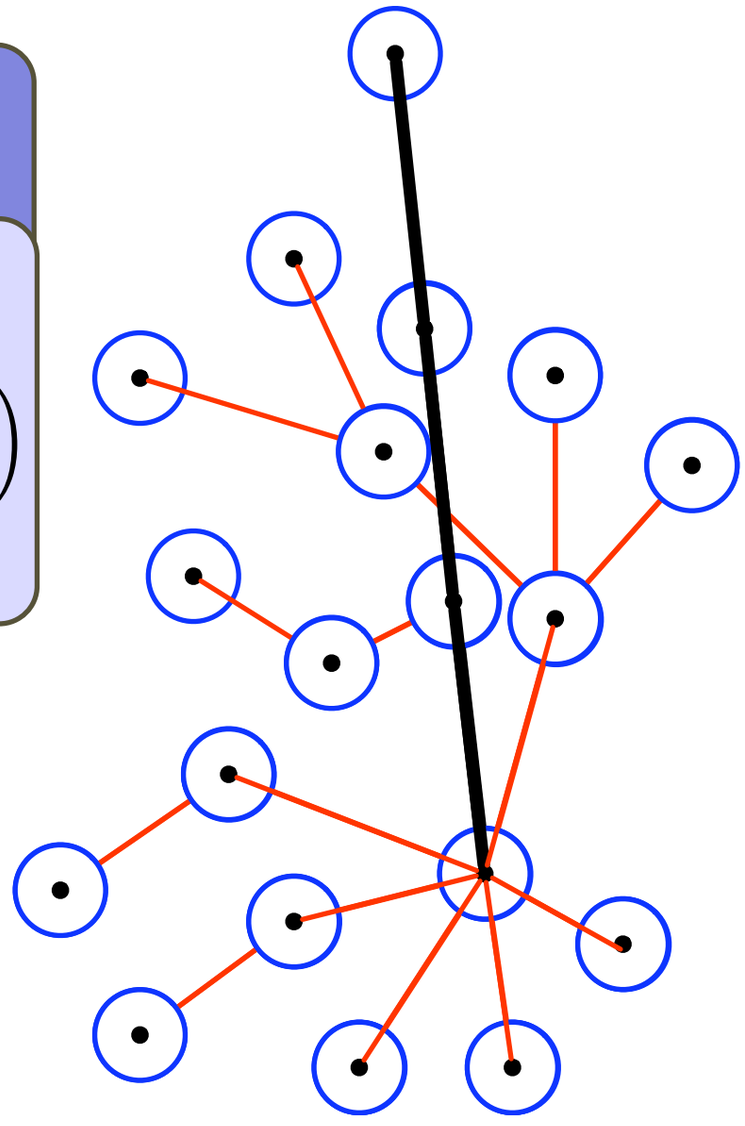
$$W_{max} \geq \frac{\sqrt{A_{node}}}{\sqrt{\pi}\left(\frac{g}{2}-1\right)}\left((d_v-1)^{\lfloor \frac{g}{4}-\frac{1}{2} \rfloor}(d_c-1)^{\lfloor \frac{g}{4}-\frac{3}{4} \rfloor}\right)$$