Wireless Connectivity:
An Intuitive and Fundamental Guide

Chapter 5: Packets Under the Looking Glass: Symbols and Noise

Petar Popovski
Connectivity Section
Department of Electronic Systems
petarp@es.aau.dk

Contributions to the slides:
Israel Leyva-Mayorga
Radoslaw Kotaba
Abolfazl Amiri
Alexandru-Sabin Bana
Robin J. WIlliams

AALBORG UNIVERSITY
DENMARK

Modules

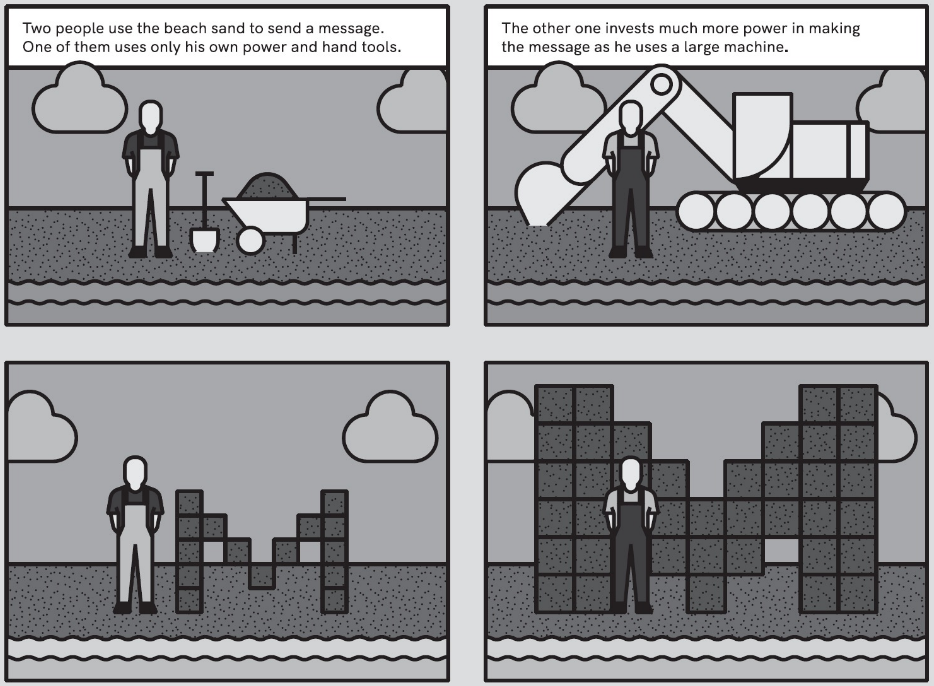
Two people use the beach sand to send a message. One of them uses only his own power and hand tools.
The other one invests much more power in making the message as he uses a large machine.

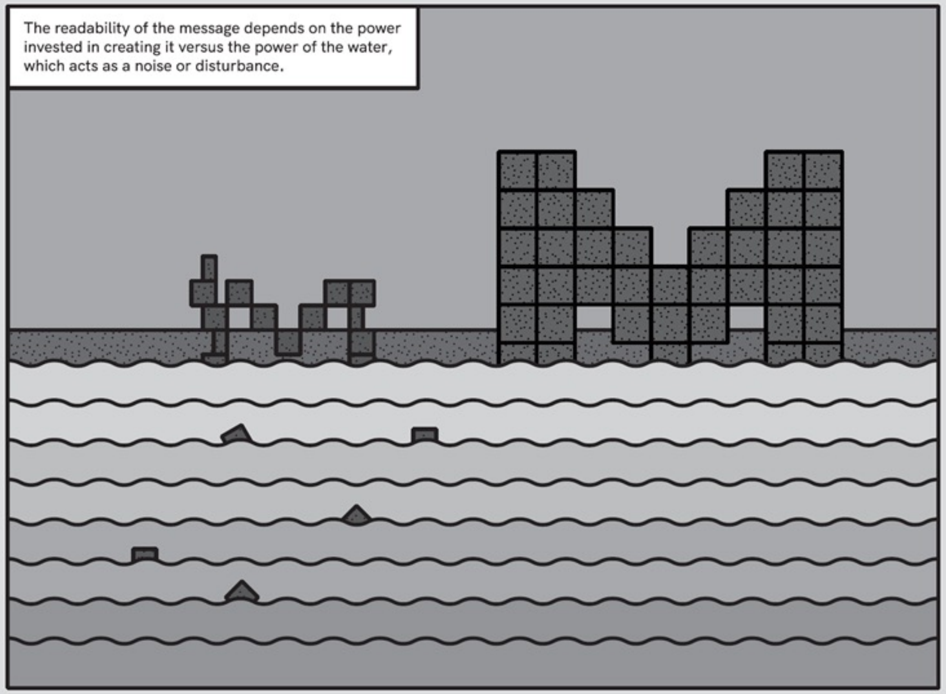
The readability of the message depends on the power invested in creating it versus the power of the water, which acts as a noise or disturbance.

Relation between the data rate and Signal-to-Noise Ratio (SNR)

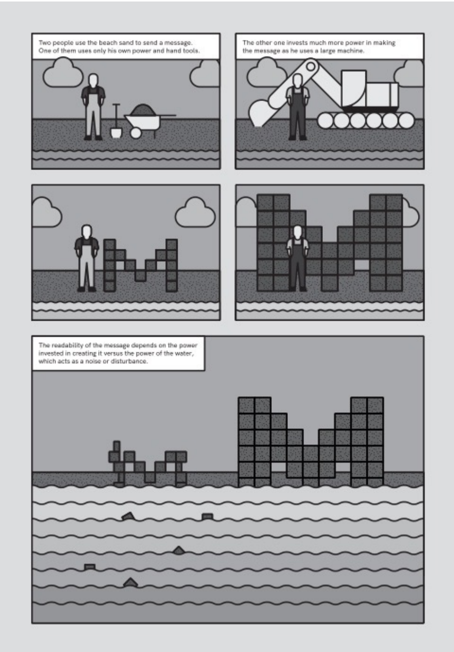

- Only discrete signal points may be perfectly reconstructed at the receiver
- Rather than absolute power, the achievable data rate is determined by the power relative to the noise or interference

What will be learned in this chapter

- Compression and why can we assume that 0 and 1 are equiprobable in the digital messages that should be transmitted
- Baseband representation of information and the notion of modulation constellations
- Signal-to-Noise ratio and expressing transmission cost through power
- Binary and non-binary constellations for modulation
- Collision under a looking glass: Symbol-level interference modeling
- Selection of data rate and adaptive modulation
- The basic idea of superposition coding

Initial digitalization and source coding

How does the system obtain data that should be packetized?

Efficient representation of information: source coding

- Intuitively: Spend less bits on the states that occur more frequently

Example:

- Assume Zoya transmits a sequence of readings to Yoshi
- Each reading can take one of the four values S_1, S_2, S_3, S_4
- Clearly, **2 bits are enough** to represent the readings, for example,

$$S_1 \rightarrow 00,\ S_2 \rightarrow 01,\ S_3 \rightarrow 10,\ S_4 \rightarrow 11$$

But, what if certain values appear **more frequently** than others?
For example $p(S_1) = 0.5$, $p(S_2) = 0.25$, $p(S_3) = 0.125$, $p(S_4) = 0.125$

Compression and entropy

Given a distribution $p_1, p_2, \dots, p_S$, the minimal average number of bits $\bar{B}$ required to describe the system output **cannot be lower than the entropy** of that distribution, defined as:

$$H(p_1, p_2, \dots, p_S) = -\sum_{s=1}^{S} p_s \log_2 p_s$$

In our case $H(0.5, 0.25, 0.125, 0.125) = 1.75$

Turns out, the optimal way to encode the messages above is actually

$$S_1 \rightarrow 0,\ S_2 \rightarrow 10,\ S_3 \rightarrow 110,\ S_4 \rightarrow 111$$

This way, on average $\bar{B} = 0.5\times1 + 0.25\times2 + 0.125\times3\times2 = 1.75$ are transmitted instead of 2, thereby attaining **optimal** compression

A bit of information

The optimal compression has another important feature:

What is the probability that a bit at a random position is either 0 or 1?

In total there are $B = 1.75L$ bits in a sequence of L readings (transmitted values).

S_1	0
S_2	10
S_3	110
S_4	111

- $0.25L$ of values correspond to S_2 which has one 1
- $0.125L$ of values correspond to S_3 which has two $1s$
- $0.125L$ of values correspond to S_4 which has three $1s$

Then, $\frac{0.25L\times1+0.125L\times2+0.125L\times3}{1.75L} = \frac{0.875L}{1.75L} = \frac{1}{2}$ of bits in the stream are $1s$

After optimal compression, 0s and 1s are **equiprobable** and occur **independently**

Many models and algorithms rely on these assumptions

How is information represented at different layers

While its data content can be understood as a sequence of 0's and 1's, this representation needs to be transformed and **modulated** to a different information carrier at different part of the systems

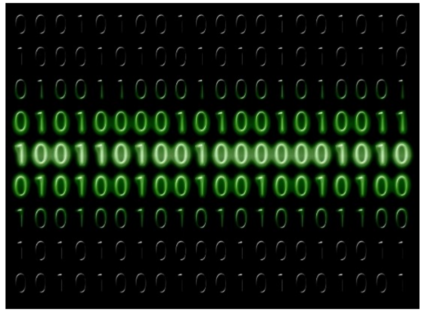

0 36000 29145 2

A step towards physical reality of packets:
baseband representation/signals

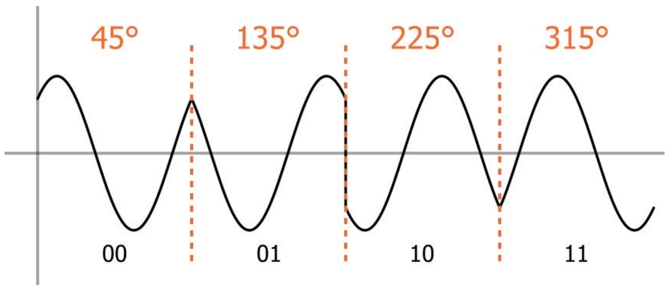

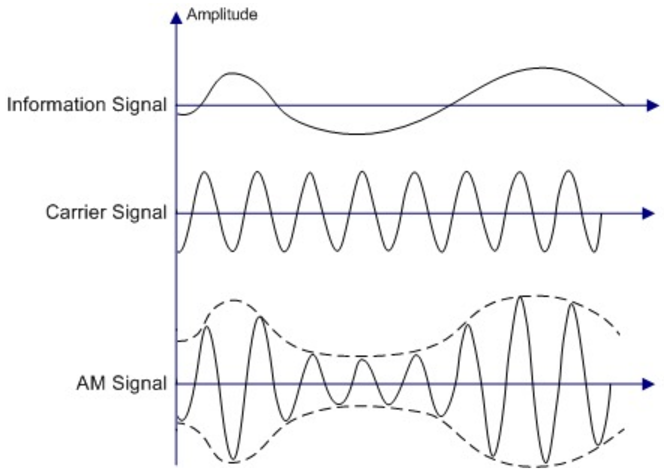

Baseband modules and data representation

A simple model

1. The baseband module takes bits and converts them into complex numbers called **baseband symbols**
2. Every T_s seconds, a symbol is sent by the TxRFModule and, after passing through the wireless medium, received by the RxRFModule
3. Model of the symbol received at RxBaseband
$$y = hz + n$$
where h is the **channel coefficient** and n is the **noise**

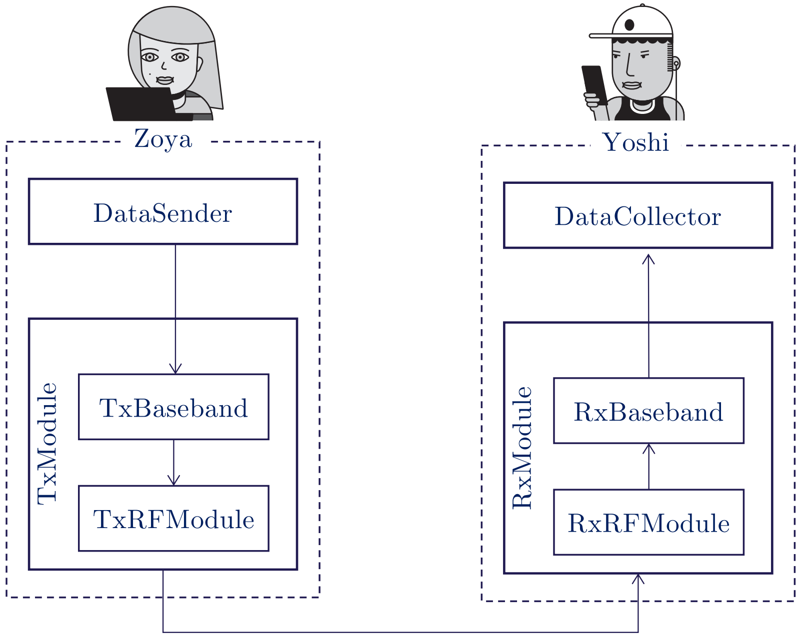

Baseband modulation: Mapping bits to symbols

Take a fictitious case with no noise ($n = 0$) and neutral channel $h = 1$, leading to a perfect reception:

$$y = z$$

Depending on the association, different data rates are possible

- If z represents a single bit, rate is $\frac{1}{T_s}$ [bps]

$z = -1$	$\rightarrow$	0
$z = +1$	$\rightarrow$	1

- If z represents two bits, rate is $\frac{2}{T_s}$ [bps]

$z = -1 - j$	$\rightarrow$	00
$z = +1 - j$	$\rightarrow$	01
$z = -1 + j$	$\rightarrow$	10
$z = +1 + j$	$\rightarrow$	11

One could define M levels, corresponding to $\log_2 M$ bits, obtaining rate $\frac{\log_2 M}{T_s}$ [bps].

However, before getting too excited, we need to look at the assumptions

Challenging the simplified assumptions about the baseband

Parties must know in advance that communication will take place

In line with discussions from previous lectures, but deeper, **think in terms of bits!**

What about **synchronization**?

- It also needs to be done using bits, but how are they different from data?

The **noise** is unknown

- The received signal is always altered in some way and not all processes that affect it can be reversed

The channel and its changes are unknown

Not as chaotic as noise and **can be learned***:* periodic investment of resources

- The *Invite* packet from Chapter 1: help in synchronization and, potentially, channel estimation

The noise cloud

We now focus on the role of the **noise** $(h = 1)$

Yoshi receives a sequence of symbols:

$$y_i = hz_i + n_i = z_i + n_i \text{ for } i = 1, \dots, L$$

It is assumed that:

- The noise samples are **independent** from symbols.
- They are drawn from a **stationary** process.
- They are **zero-mean**.
- They are **circularly symmetric**.
- Their distribution is **Gaussian**.

This leads to an intuitive representation of the noise as the **noise cloud.**

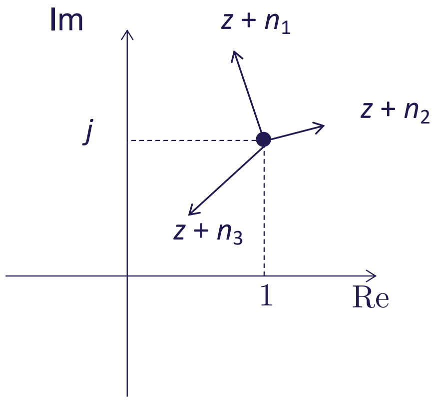

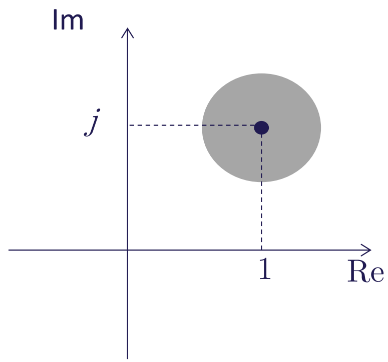

Constellation points

The key issue is designing a suitable **constellation:** a set of discrete points that represent M possible transmitted signals

- *How to chose the best points for binary modulation $M = 2$?*

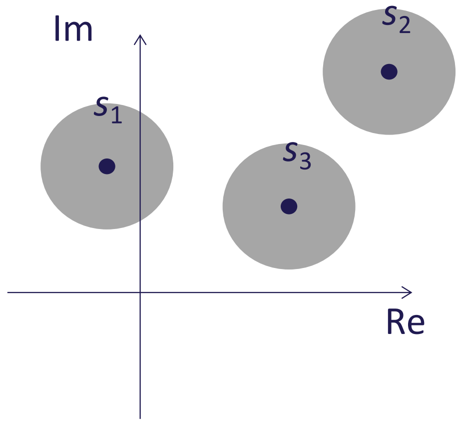

Some observations:

- The more spaced apart, the lower chance of error
- More points = more bits into a single symbol
- *Can we do this indefinitely? What are the consequences?*

Increasing the modulation order (number of points) without overlapping the clouds is possible but:

- Fix the noise clouds but move the points away from the origin
- Decrease the noise clouds, while restricting the area of the signal

Transmission power as a cost

The **magnitude** of the symbol $|z|^2$ directly related to **transmission cost** in terms of **power**

In practice, systems put a limit on the

- **Average transmit power** $P_T = \sum_{i=1}^{S} p(s_i)\,|s_i|^2$
- **Maximum transmit power** $P_{max} = \max_{i=1,\ldots,S} |s_i|^2$

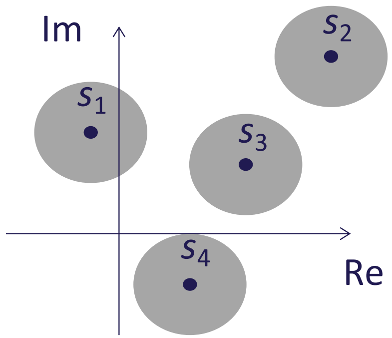

Why is the constellation on the figure not a good one, especially, when $p(s_1) = p(s_2) = p(s_3) = p(s_4)$?

Signal constellations and noise

When the noise is Gaussian,
the ‘radius’ of the cloud is related to its power P_N
and allows error detection

- This might not be necessary (recall CRC)
- Instead divide the plane into complementary **decision regions**

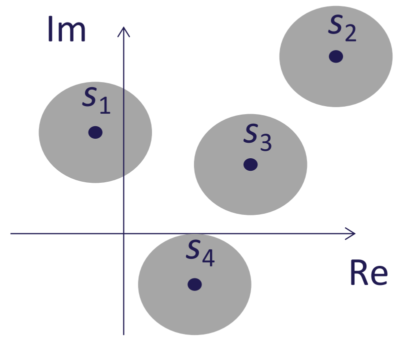

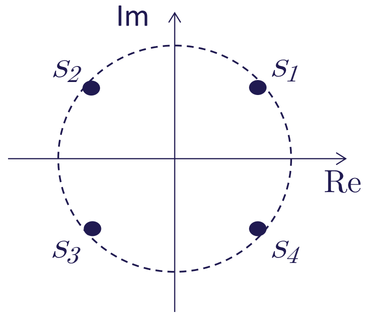

Note on source coding and modulation

Source coding maps input states/symbols to bits,
modulation maps bits to symbols, but is not an inverse of source coding

The source-coded pure random bits are then divided into M-tuples (for M-ary constellation), each M-tuple mapped to a symbol and eventually transmitted

- Symbol does not necessarily contain full state, as can be seen in the example
- If the bitstream were not optimally compressed,
  symmetric modulation would not be optimal

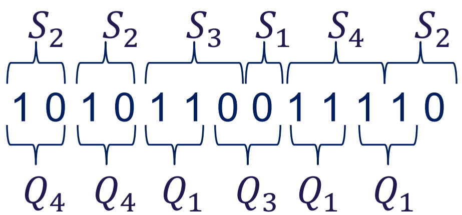

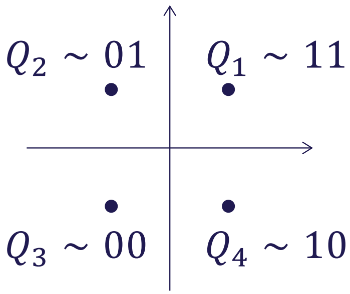

S_1	0
S_2	10
S_3	110
S_4	111

Defining SNR (Signal-to-Noise Ratio)

Recall the discussion on the
increase of the magnitude of the constellation points (P_T) vs.
reducing the radius of clouds (P_N)

- This implies that what matters is their **ratio**, quantity known as ***SNR:***

$$\gamma = P_T / P_N$$

For a signal given by $y = hz_i + n_i$, the SNR is equal to:

$$\gamma = \frac{(hz_i)(hz_i)^*}{n_i n_i^*} = \frac{|h|^2 |z_i|^2}{|n_i|^2} = \frac{|h|^2 P_T}{P_N}$$

Where $|z_i|^2 = P_T$ is due to all constellation points having equal power.

- SNR does not change if we operate with $h^* y$ or $\frac{h^* y}{|h|^2} = z_i + \frac{h^* n_i}{|h|^2}$

Modulation with a binary constellation

A simple modulation is **binary-phase-shift keying (BPSK)**

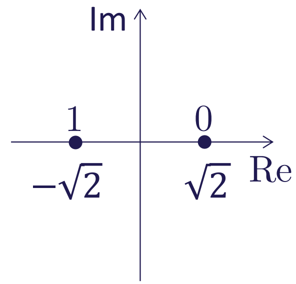

- The symbols are $\mathcal{S}_B = \{-\sqrt{2}, +\sqrt{2}\}$ with average power $P_T = 2$
- Error probability error is a function of the SNR

$$P_B(\gamma) = Q\left(\sqrt{2\gamma}\right) \text{ where } Q(x) = \frac{1}{\sqrt{2\pi}}\int_x^{\infty} e^{-\frac{x^2}{2}} dx$$

A specific noise property in BPSK:

- Since $n = n_{Re} + n_{Im}$, each with power $\frac{P_N}{2}$, the SNR is actually $\gamma = \frac{2|h|^2}{\frac{P_N}{2}} = \frac{4|h|^2}{P_N}$, because the imaginary component has no impact

Modulation with a quaternary constellation

By multiplexing real and imaginary BPSK we obtain **QPSK**

- The symbols are $\mathcal{S}_Q = \{+1+j, -1+j, -1-j, +1-j\}$ with $P_T = 2$
- The bit-error probability is now different from symbol-error probability $P_Q = 1 - \left(1 - P_B(\gamma)\right)^2$
- There are many possible mappings between bit sequences and $\mathcal{S}_Q$
- A smaller Euclidean distance should correspond to a lower Hamming distance

$d_H(00{,}01) = 1$ and $d_H(00{,}11) = 2$, more generally $d_H(\boldsymbol{b}_1, \boldsymbol{b}_2) = \sum_{i=1}^{N} b_{1,i} \oplus\ b_{2,1}$

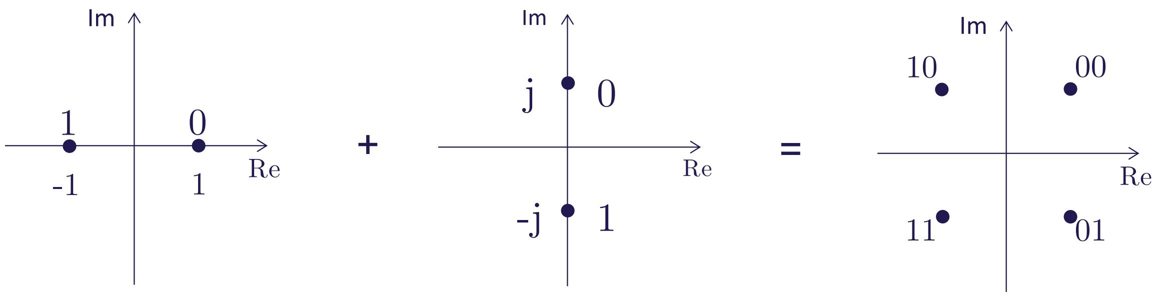

Modulation with constellations of a higher order

A single dimension with M-ary constellation: **Pulse Amplitude Modulation** (**PAM**)

- In each dimension noise maps onto an interval around the symbol
- As long as the intervals are separated well enough, the error probability is low
  - Recall the power constraint
  - Look at the example below, where $P_T = \frac{1}{4}(3^2 + 1^2 + 1^2 + 3^2) = 5$
- Complex constellations with PAM along each dimension

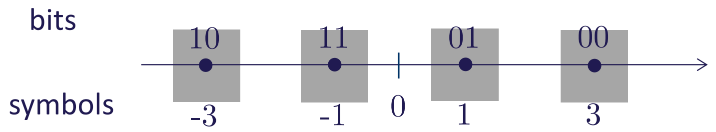

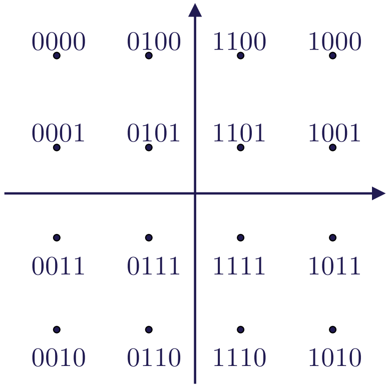

Some thoughts on generalized constellations

The approach explained can be generalized to more than two dimensions

- **Example:** Treat two consecutive received values (u, v) as a single symbol and it is described by two real and two imaginary values

  The average power of a constellation symbol is $P(u, v) = 2P_T$

  Take the specific case of a constellation $\mathcal{S}$ that carry 8 bits (256 values)

  One could achieve that by transmitting two independent 16-QAM points

  What if that is not optimal?
  Could u and v be drawn from different constellations instead?

Symbol-level interference models

Recall lecture 3 on staring to look inside a collision

Enriched model by symbol-level packet representation

Assume: synchronicity, L complex symbols per packet of duration T each symbol duration $\frac{T}{L}$

$$y_B = h_{ZB} \cdot z + h_{XB} \cdot x + n_B$$

Assume: h_{ZB} and h_{XB} are **constant** and **known** by **Basil**

No transmission from a user is equivalent to knowing the data sent by the user

Lorem ipsum
Pantarerei

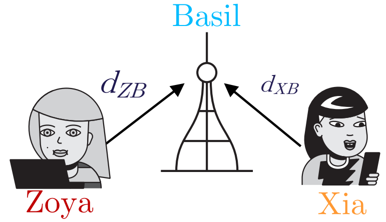

Advanced collision treatment at baseband

$|h_{ZB}| > |h_{XB}|$

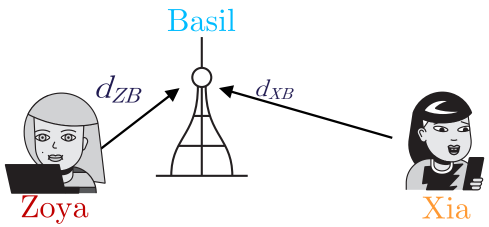

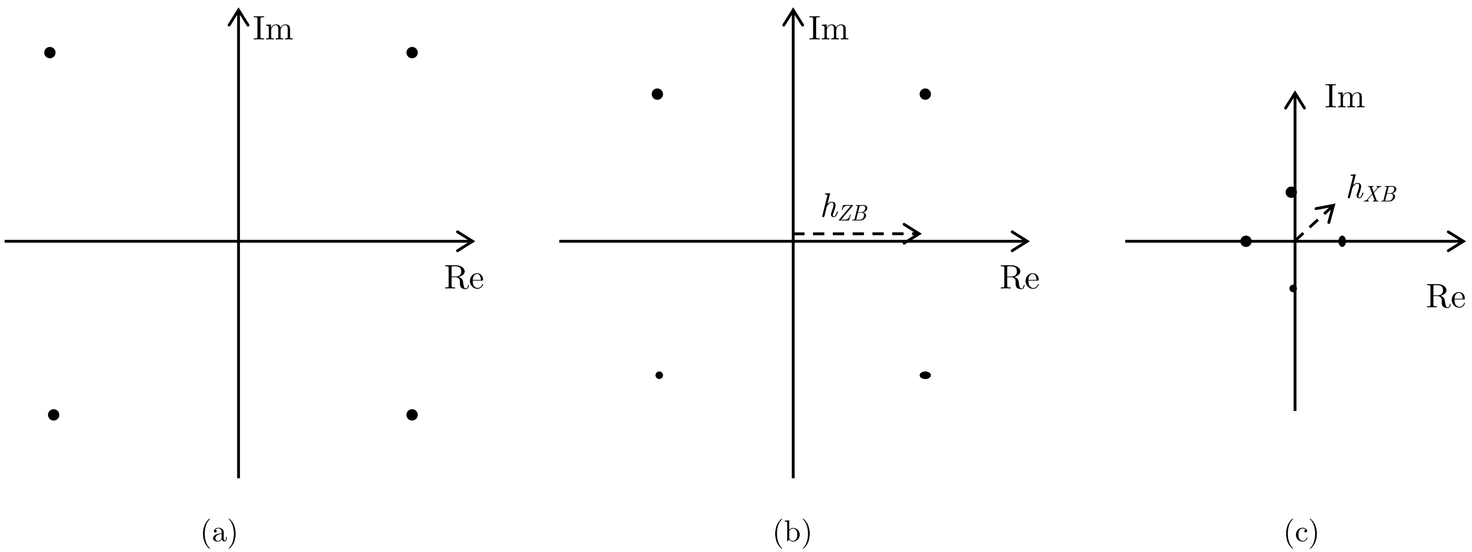

Advanced collision treatment at baseband

Received constellation
(superposition)

$$h_{ZB} \cdot z + h_{XB} \cdot x$$

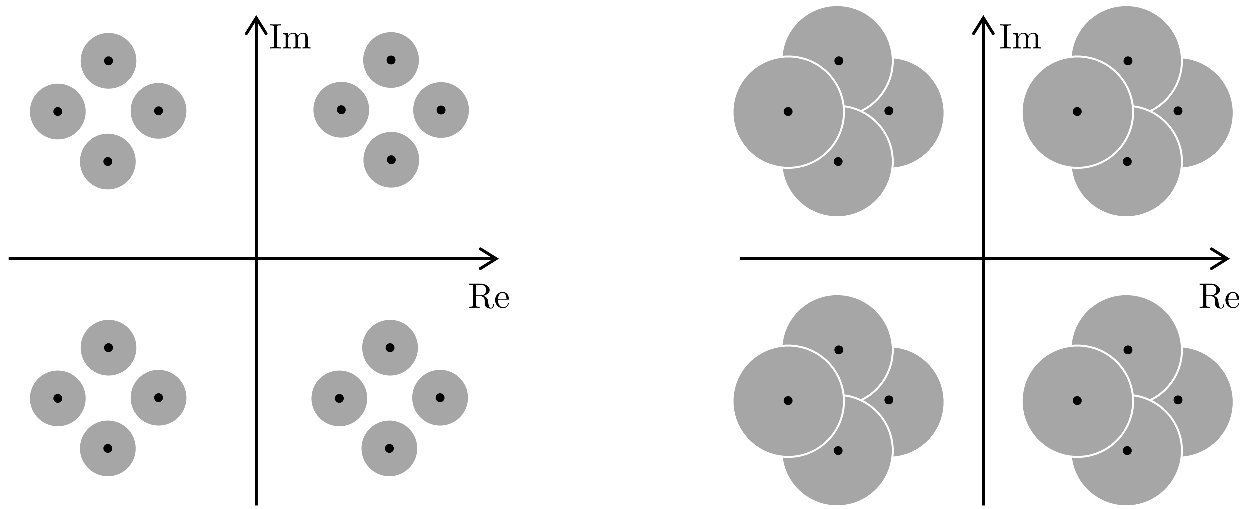

16 points

Unique pair of Tx symbols

Basil performs **joint decoding**

SIC and SINR

Basil observes $h_{ZB} \cdot z$ with
added **noise** plus **interference**

The total noise is $h_{XB} \cdot x + n_B$

Not precisely Gaussian,
but approximated as Gaussian

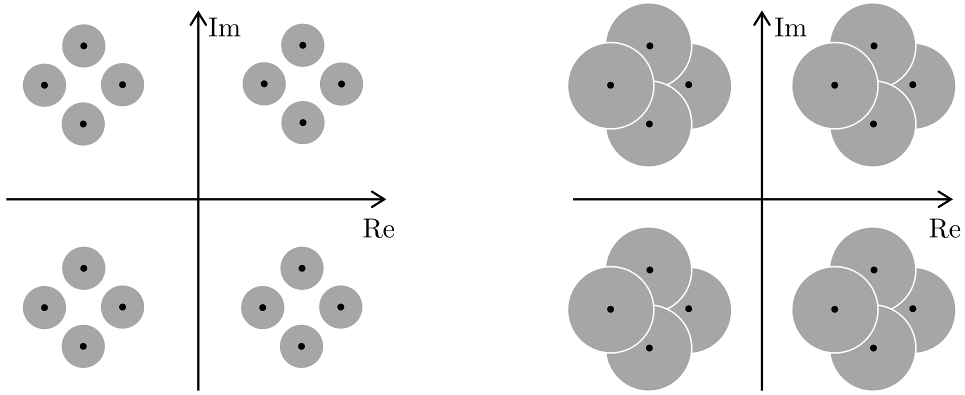

Resulting noise variance is $|h_{XB}|^2 P_T + P_N$

then $SINR = \frac{|h_{ZB}|^2 P_T}{|h_{XB}|^2 P_T + P_N} = \frac{\gamma_{ZB}}{1+\gamma_{XB}}$

SIC and SINR

$$SINR = \frac{|h_{ZB}|^2 P_T}{|h_{XB}|^2 P_T + P_N} = \frac{\gamma_{ZB}}{1 + \gamma_{XB}}$$

Joint decoding and intra-collision SIC
are **not** identical!

Important: SIC is rarely perfect in practice

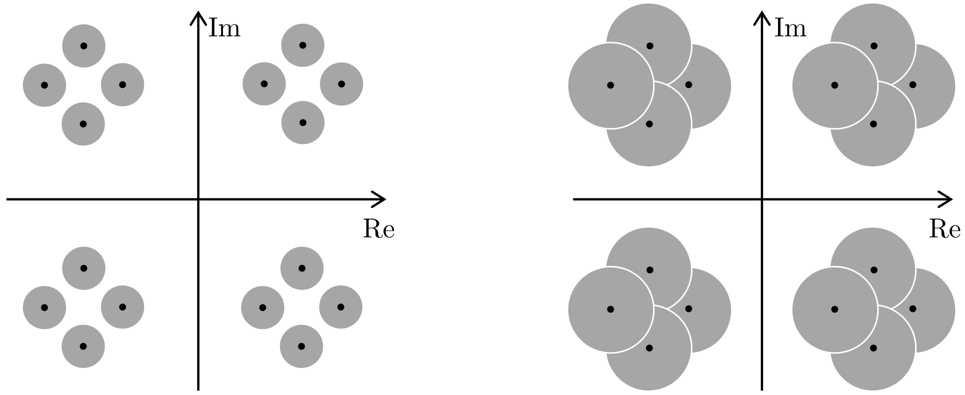

Capture effect at the symbol level

Weak and strong signals

$$SINR_Z = \frac{\gamma_{ZB}}{1 + \gamma_{XB}}$$

$$SINR_X = \frac{\gamma_{XB}}{1 + \gamma_{ZB}}$$

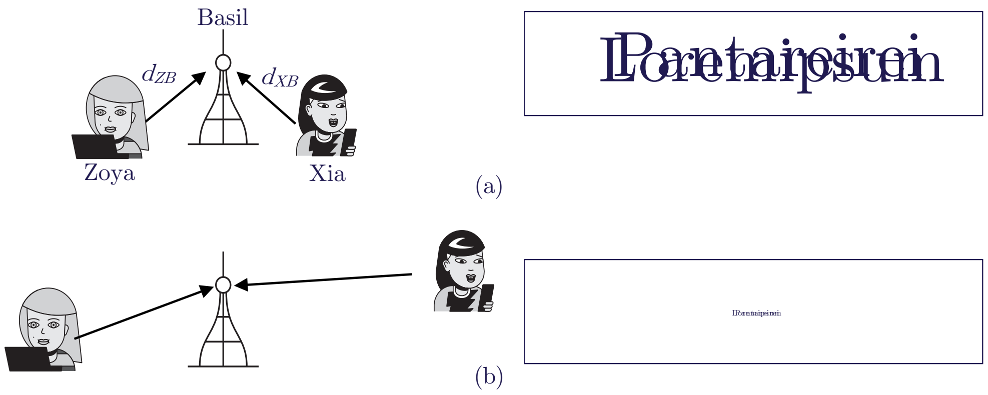

Consider $\gamma_{ZB} = \gamma_{XB}$ very large

Then $SINR_Z = SINR_X \to 1 = 0$ [dB] →

→ **Interference limited** rather than **noise-limited**

Extreme example: $h_{ZB} = h_{XB} = 1$ and no noise, then $y_B = 0$ generates ambiguity among:

$$(s_Z, s_X) \in \{(1 + j, -1 - j), (-1 + j, 1 - j), (-1 - j, 1 + j), (1 - j, -1 + j)\}$$

It follows that it is impossible to have an error-free detection of overlapping points in the superimposed constellation

Randomization of power

Symbol-level perspective

basis to generalize ALOHA:

Fine-tuned power control (not just either *P* or *0)*

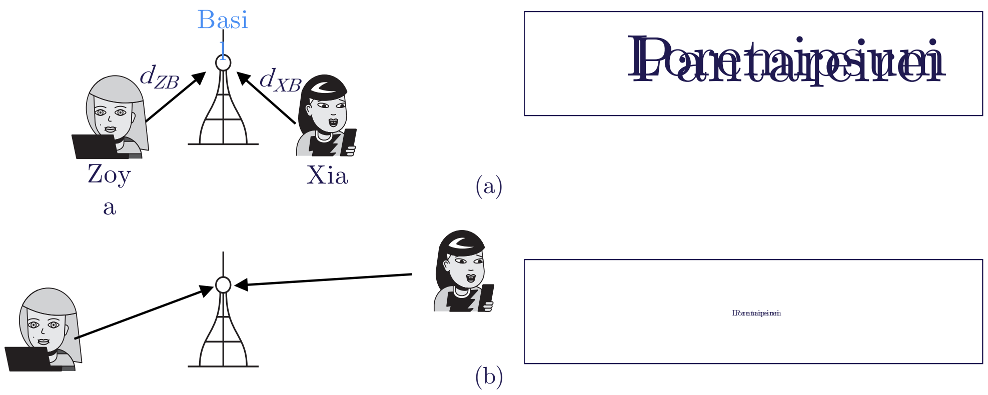

Illustration:

Assume **Zoya** keeps constant power and **Xia** **decreases** to $\gamma_{XB,1} < \gamma_{XB}$

$\gamma_{XB,1}$ should be **sufficiently low** s.t. $SINR_Z$ is **sufficiently high** that **B** decodes **Z**

After **Zoya**'s packet is **decoded** and **canceled**,
Xia's packet can be decoded if $\gamma_{XB,1}$ is **sufficiently high**

Other baseband model goodies

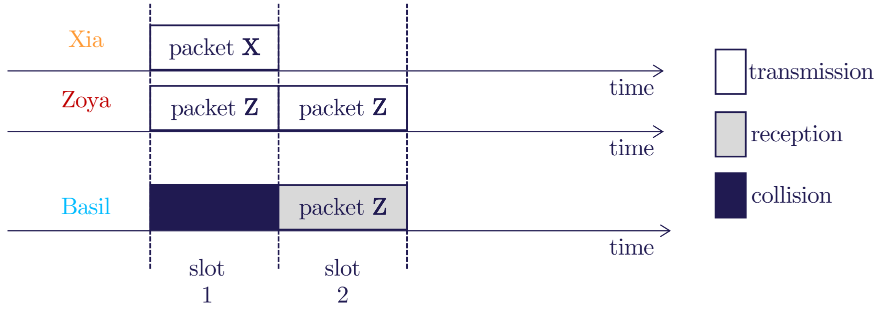

$$y_{B,i,1} = h_{ZB} \cdot z_i + h_{XB} \cdot x_i + n_{B,i,1}$$

and $y_{B,i,2} = h_{ZB} \cdot z_i + n_{B,i,2}$

From slot 2, **Basil** decodes $\boldsymbol{Z}$, reconstructs its baseband symbols z_i, subtracts them from slot 1 signal

Try now to decode X from

$$y'_{B,i,1} = h_{XB} \cdot x_i + n_{B,i,1}$$

Important feature: **scales** to more than 2 **interfering users**

$$y_{B,i} = \sum_{k=1}^{K} h_{kB} \cdot x_{k,i} + n_{B,i}$$

Generalization when packets do not fully overlap

How to select the data rate?

L complex symbols per **packet** of duration **T**

nominal rate $\mathrm{R} = \frac{D}{T}$

BPSK: $D = L$ bits $\rightarrow R_{BPSK} = \frac{L}{T}$

16-QAM: $D = 4L$ bits $\rightarrow R_{16-QAM} = \frac{4L}{T} = 4R_{BPSK}$

Actual data rate depends on the packet error probability (*PEP*)

$$PEP = 1 - (1 - P_S)^L$$

Note: for a **fixed modulation,** $\boldsymbol{P_S}$ **increases** as the **SNR** **decreases**

Adaptive modulation

Assume Zoya transmits to Yoshi with fixed TX power and an SNR γ

Zoya can use **BPSK**, **QPSK**, **16-QAM** achieving symbol error probabilities

$P_B(\gamma) < P_Q(\gamma) < P_{16}(\gamma)$, inversely proportional to the Euclidean distance between symbols

Same goes for the ***PEP*** if the number of symbols of a packet is *L*

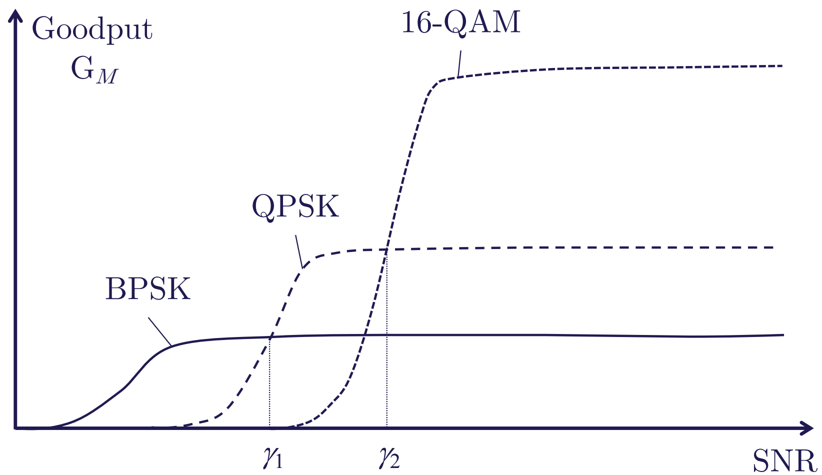

Goodput: $G = \frac{B_F}{T_F} \rightarrow G_M(\gamma) = \frac{B_M}{T}\left(1 - PEP_M(\gamma)\right)$

One can apply **adaptive modulation** if one has **channel state information** (CSI)

All this assumes T is fixed and $L = \frac{T}{T_S}$

T_S cannot be described within the baseband model, reaching its limitations.

Superposition of baseband symbols

Extending from $z = \sqrt{\alpha} z_{B1} + (1 - \sqrt{\alpha}) z_{B2}$ to

$$z_i = \sum_{s=1}^{S} \sqrt{\alpha_s} z_{s,i}$$

with the power constraint $\sum_{s=1}^{S} \alpha_s = 1$

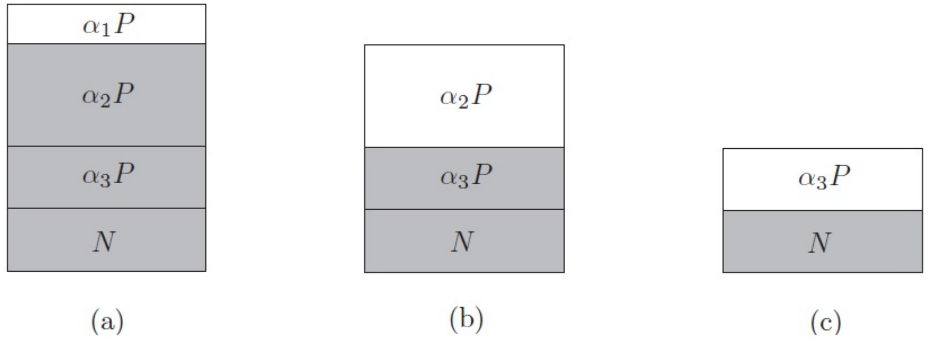

$$SINR_1 = \frac{\alpha_1 P}{\alpha_2 P + \alpha_3 P + N} \qquad SINR_2 = \frac{\alpha_2 P}{\alpha_3 P + N} \qquad SINR_3 = SNR_3 = \frac{\alpha_3 P}{N}$$

If modulations can be chosen independently, the data rate of each can be different
if only the first $K - 1$ packets are decoded correctly, then the achieved data rate is:

$$R' = \sum_{s=1}^{K-1} R_s$$

Even though superposition creates **self-interference** and seems **suboptimal**,
let us look at some elegant solutions it offers

Broadcast and non-orthogonal access

Broadcast = multicast traffic + unicast traffic

If there is no multicast,
we have non-orthogonal multiple access (NOMA)

Signal transmitted by Basil $b = \alpha_Z z + \alpha_X x$

Nominal rates $R_Z = R_X = 2$ bits/symbol

α_Z, R_Z and α_X, R_X are the **parameters** at play

Maximize overall throughput?

- Set the weaker receiver's $\alpha = 0$
- However, this may raise **fairness** issues
- Maximize overall throughput and guaranteeing minimal throughput for each user can work
- Potential **security** issues

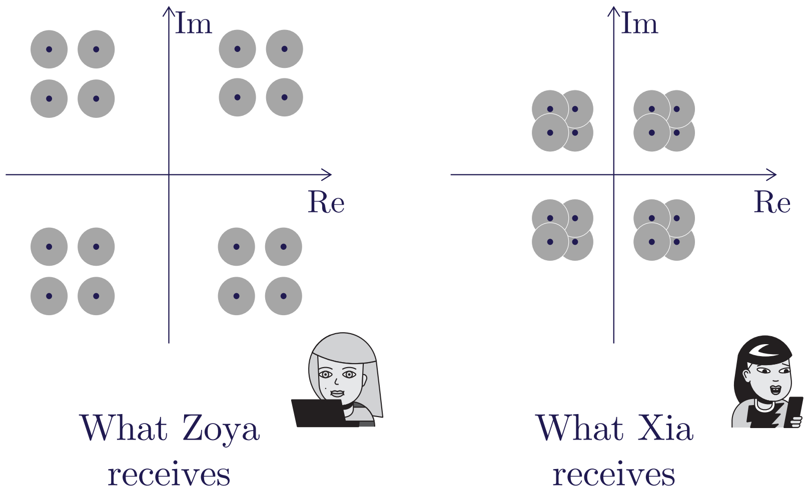

Unequal Error Protection (UEP)

Layers in multimedia content have various importance (video for instance)

Basic data D_B for lower quality, and enhancement data D_E for higher resolution

Goal: maximizing **utility function** $\boldsymbol{u}(D_B, D_E)$

$$E[u] = q_B \cdot u(D_B, 0) + q_{BE} \cdot u(D_B, D_E)$$

Superposition coding to modulate z_B with α_B and z_E with α_E

Particularly suitable when multicasting to multiple receivers

Ensures **graceful degradation** of the multimedia quality across receivers

Outlook and takeaways

- Connection between networking and digital communication: packets vs symbols-bits-modulation-noise
- Baseband complex symbols and noisy versions
- SNR and SINR
- Evaluating symbol error performance and packet error
- New protocol designs: changing transmit power, using superposition coding