Conectando os Pontos - A Contribuição de Robert Metcalfe Para a Computação

Claudio Miceli de Farias -Universidade Federal do Rio de Janeiro
cmicelifarias@cos.ufrj.br

UFRJ faz 100 ANOS 1920 | 2020
UNIVERSIDADE FEDERAL DO RIO DE JANEIRO

Engenharia de Sistemas e Computação PESC

NCE UFRJ Instituto Tércio Pacitti de Aplicações e Pesquisas Computacionais

PR-2 / UFRJ

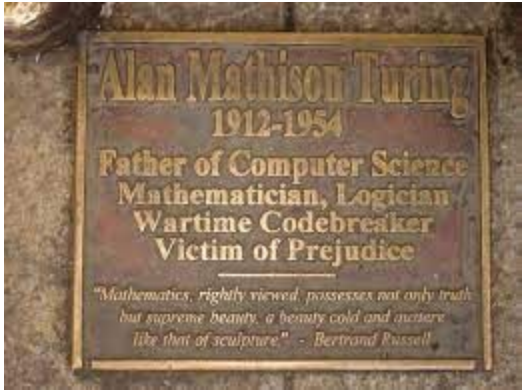

acm TURING AWARD


Yoshua Bengio
Geoffrey Hinton
Yann LeCun


Edwin Catmull
Patrick Hanrahan

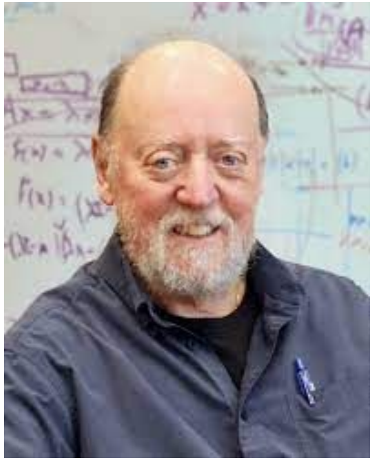


Alfred Aho
Jeffrey Ullman

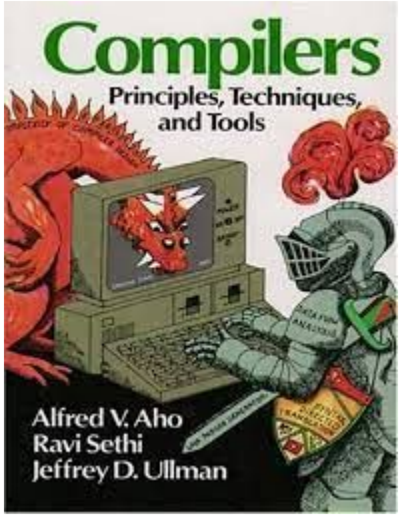

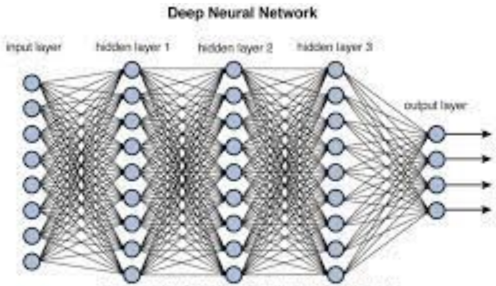

Figure 12.2 Deep network architecture with multiple layers.

Bob MetCalfe - vencedor em 2023 pela ethernet

Não parece tão legal … =(

SERÁ??

First things first: Quem é Bob MetCalfe?

- Aluno do MIT
- Capitão do time de Tênis do MIT - quase foi jogador
- PhD em Harvard - Applied Mathematics - mas não de primeira
- Trabalhou na Xerox Palo Alto Research center (PARC) em 1973
- Em 79 sai da Xerox e funda a Tricom
- Em 2022, virou professor de empreendedorismo no MIT
- Muito multifacetado
- Trabalhou na Infoworld e teve uma coluna semanal

https://cacm.acm.org/magazines/2023/6/273232-making-connections/fulltext

Ótimo, mas qual o problema?

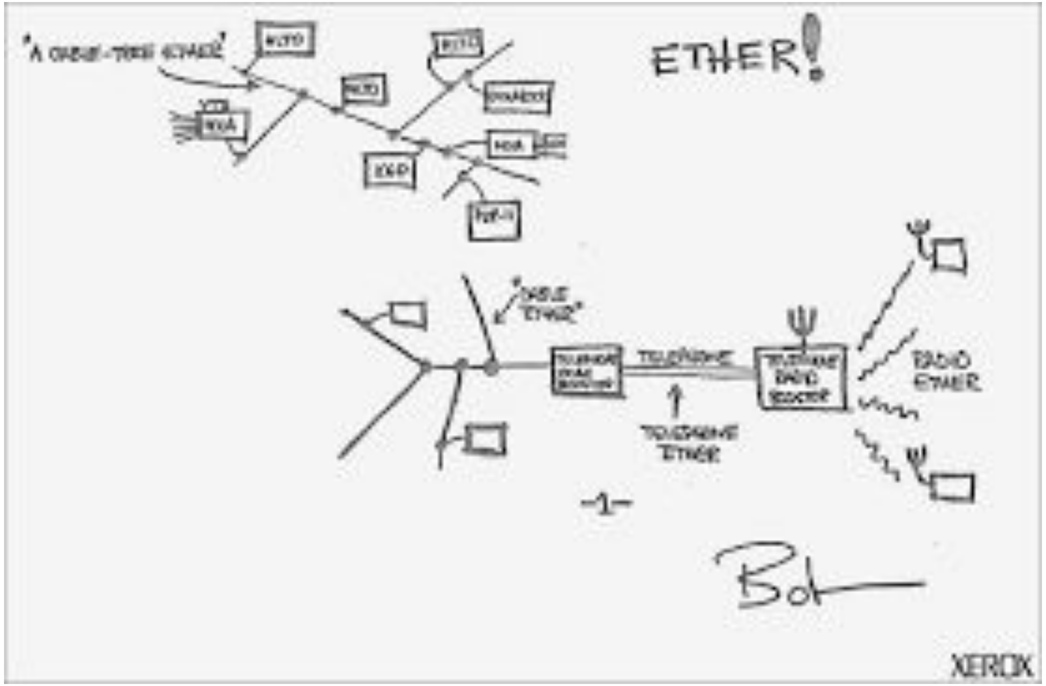

Ok, mas e essa tal de ethernet?

IEEE 802 Working Group
&
Executive Committee Study Group
Home Pages

http://www.ieee802.org/dots.html
Outubro 2009

Active Working Groups and Study Groups

- 802.1 Higher Layer LAN Protocols Working Group
  - Link Security Executive Committee Study Group is now part of 802.1
- 802.3 Ethernet Working Group
- 802.11 Wireless LAN Working Group
- 802.15 Wireless Personal Area Network (WPAN) Working Group
- 802.16 Broadband Wireless Access Working Group
- 802.17 Resilient Packet Ring Working Group
- 802.18 Radio Regulatory TAG
- 802.19 Coexistence TAG
- 802.20 Mobile Broadband Wireless Access (MBWA) Working Group
- 802.21 Media Independent Handoff Working Group
- 802.22 Wireless Regional Area Networks

Inactive Working Groups and Study Groups

- 802.2 Logical Link Control Working Group
- 802.5 Token Ring Working Group

Disbanded Working Groups and Study Groups

- 802.4 Token Bus Working Group (material no longer available on this web site)
- 802.6 Metropolitan Area Network Working Group (material no longer available on this web site)
- 802.7 Broadband TAG (material no longer available on this web site)
- 802.8 Fiber Optic TAG (material no longer available on this web site)
- 802.9 Integrated Services LAN Working Group (material no longer available on this web site)
- 802.10 Security Working Group (material no longer available on this web site)
- 802.12 Demand Priority Working Group
- 802.14 Cable Modem Working Group (material no longer available on this web site)
- QOS/FC Executive Committee Study Group (material no longer available on this web site)

Computer Systems

G. Bell, S. Fuller and D. Siewiorek, Editors

Ethernet: Distributed Packet Switching for Local Computer Networks

Robert M. Metcalfe and David R. Boggs
Xerox Palo Alto Research Center

Ethernet is a branching broadcast communication system for carrying digital data packets among locally distributed computing stations. The packet transport mechanism provided by Ethernet has been used to build systems which can be viewed as either local computer networks or loosely coupled multiprocessors. An Ethernet's shared communication facility, its Ether, is a passive broadcast medium with no central control. Coordination of access to the Ether for packet broadcasts is distributed among the contending transmitting stations using controlled statistical arbitration. Switching of packets to their destinations on the Ether is distributed

One can characterize distributed computing as a spectrum of activities varying in their degree of decentralization, with one extreme being remote computer networking and the other extreme being multiprocessing. Remote computer networking is the loose interconnection of previously isolated, widely separated, and rather large computing systems. Multiprocessing is the construction of previously monolithic and serial computing systems from increasingly numerous and smaller pieces computing in parallel. Near the middle of this spectrum is local networking, the interconnection of computers to gain the resource sharing of computer networking and the parallelism of multiprocessing.

The separation between computers and the associated bit rate of their communication can be used to divide the distributed computing spectrum into broad activities. The product of separation and bit rate, now about 1 gigabit-meter per second (1 Gbmps), is an indication of the limit of current communication technology and can be expected to increase with time:

Activity	Separation	Bit rate
Remote networks	> 10 km	< .1 Mbps
Local networks	10–.1 km	.1–10 Mbps
Multiprocessors	< .1 km	> 10 Mbps

1.1 Remote Computer Networking

Computer networking evolved from telecommunications *terminal-computer* communication, where the object was to connect remote terminals to a central computing facility. As the need for *computer-computer* interconnection grew, computers themselves were used

Calma aí, então ethernet não é cabo?

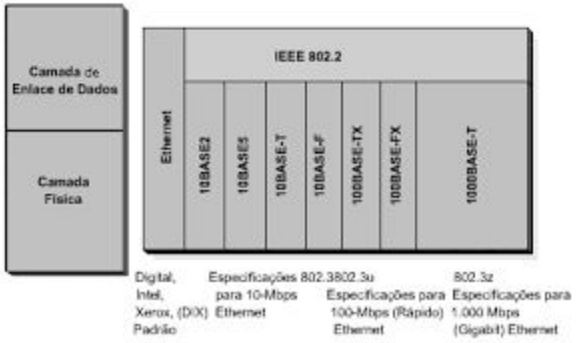

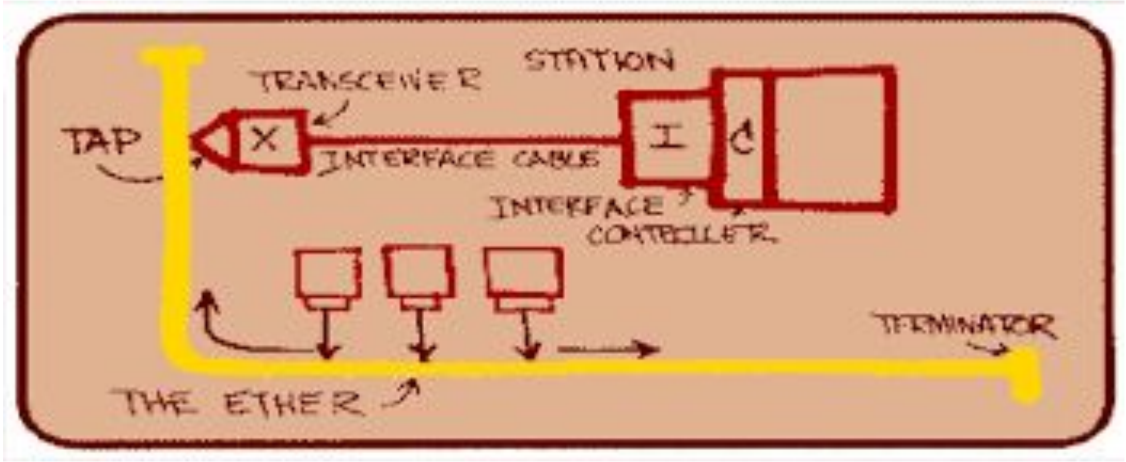

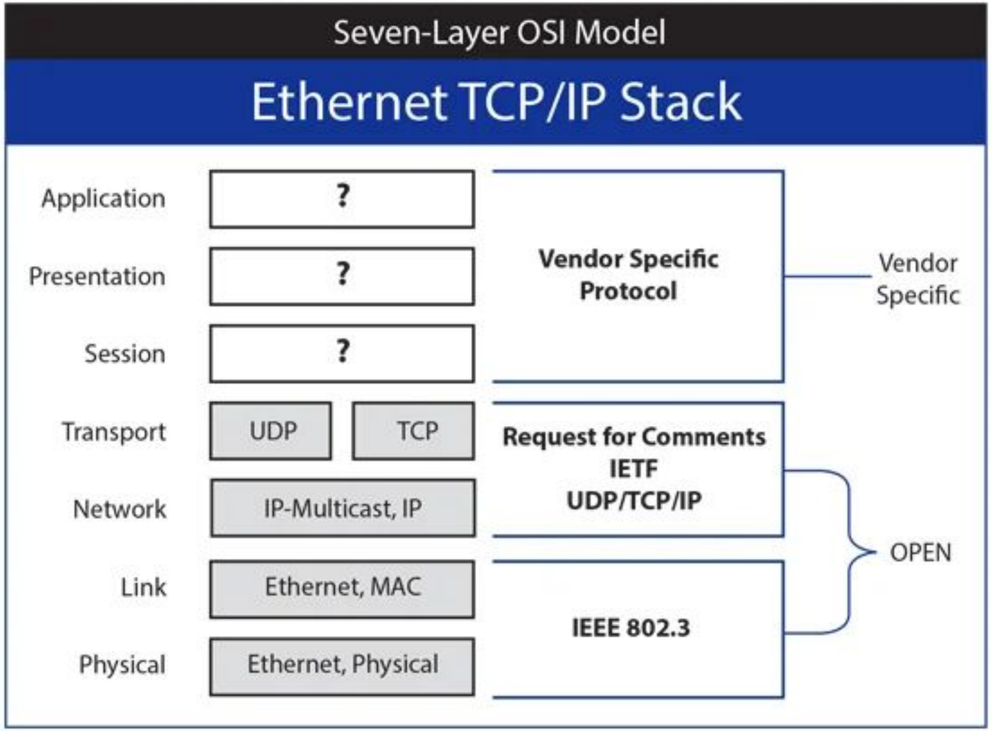
Seven-Layer OSI Model
Ethernet TCP/IP Stack
Application
Presentation
Session
Transport
Network
Link
Physical
?
?
?
UDP
TCP
IP-Multicast, IP
Ethernet, MAC
Ethernet, Physical
Vendor Specific Protocol
Request for Comments
IETF
UDP/TCP/IP
IEEE 802.3
Vendor Specific
OPEN

Por quê Ethernet?

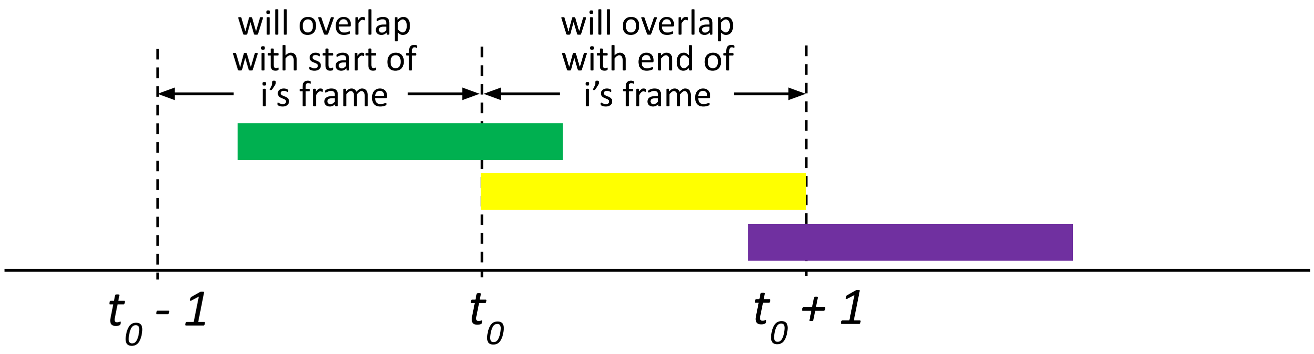

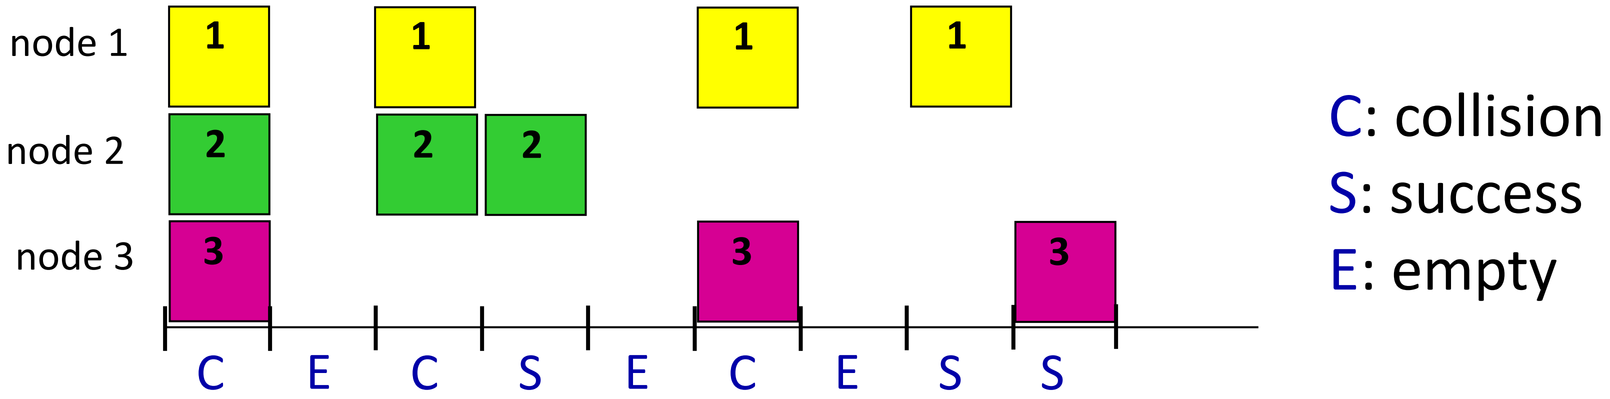
node 1
node 2
node 3
1
1
1
1
2
2
2
3
3
3
C E C S E C E S S
C: collision
S: success
E: empty

eficiência: long-run fraction of successful slots (many nodes, all with many frames to send)

- *suponha: N* nós com vários quadros a transmitir, cada um com probabilidade p
  - prob de uma dado nó ter sucesso na transmissão no slot = $p(1-p)^{N-1}$
  - prob que algum nó tenha sucesso = $Np(1-p)^{N-1}$
  - a eficiência máxima: encontre p^* que maximize $Np(1-p)^{N-1}$
  - para muitos nós, o limite de $Np^*(1-p^*)^{N-1}$ quando *N* tende ao infinito é:

1/e = .37

- *No melhor dos casos:* o canal está ocupado 37% do tempo!

CSMA (carrier sense multiple access)

CSMA simples: ouça antes de transmitir:

- se o canal estiver vazio: transmitir o quadro inteiro
- se o canal estiver ocupado: abortar transmissão

- analogia humana: não interrompa os outros!

CSMA/CD: CSMA com *detecção de colisão*

- colisões detectadas em um intervalo curto
- transmissões que colidem são abortadas, reduzindo o desperdício do canal de comunicação
- a detecção de colisão é fácil no meio cabeado e complicado no sem fio

- analogia humana: conversa educada com alternância entre os falantes

CSMA: colisões

- colisões ainda podem ocorrer:
  - O atraso de propagação faz com que dois nós possam não ouvir que houve um início de transmissão
- colisão: O tempo total de transmissão do quadro é perdido
  - A distância e a propagação tem um papel vital na determinação da probabilidade de colisão

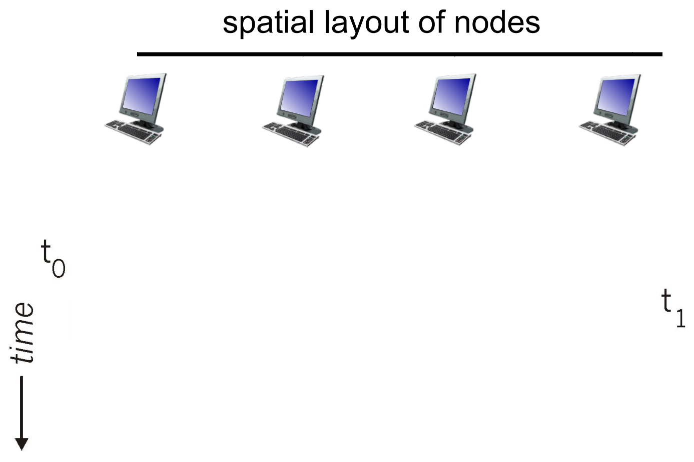

CSMA/CD:

- CSMA/CD reduz o tempo gasto com colisões
  - Uma vez que eu detectei a colisão aborta a comunicação

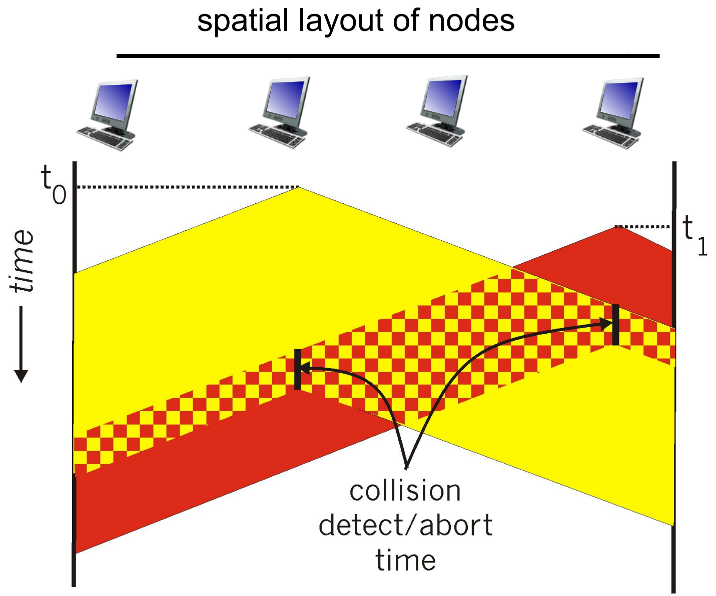

Ethernet CSMA/CD

1. Placa de rede recebe um datagrama da camada de rede e cria um quadro
2. Se o canal de comunicação estiver:
   Não utilizado: começa a transmissão do quadro.
   Se ocupado: Espera até que o canal esteja livre e transmite
3. Se a placa transmitiu o quadro sem colisão, acabou!
4. Se detectar outra transmissão: aborta e envia um sinal de jam
5. Depois de abortar, entra *binary (exponential) backoff:*
   - depois da colisão *m*-ésima, a placa de rede escolhe um K aleatório de $\{0,1,2, \ldots, 2^m-1\}$. Espera $K \cdot 512$ bit times, volta ao passo 2
   - Quanto mais colisões maior o backoff

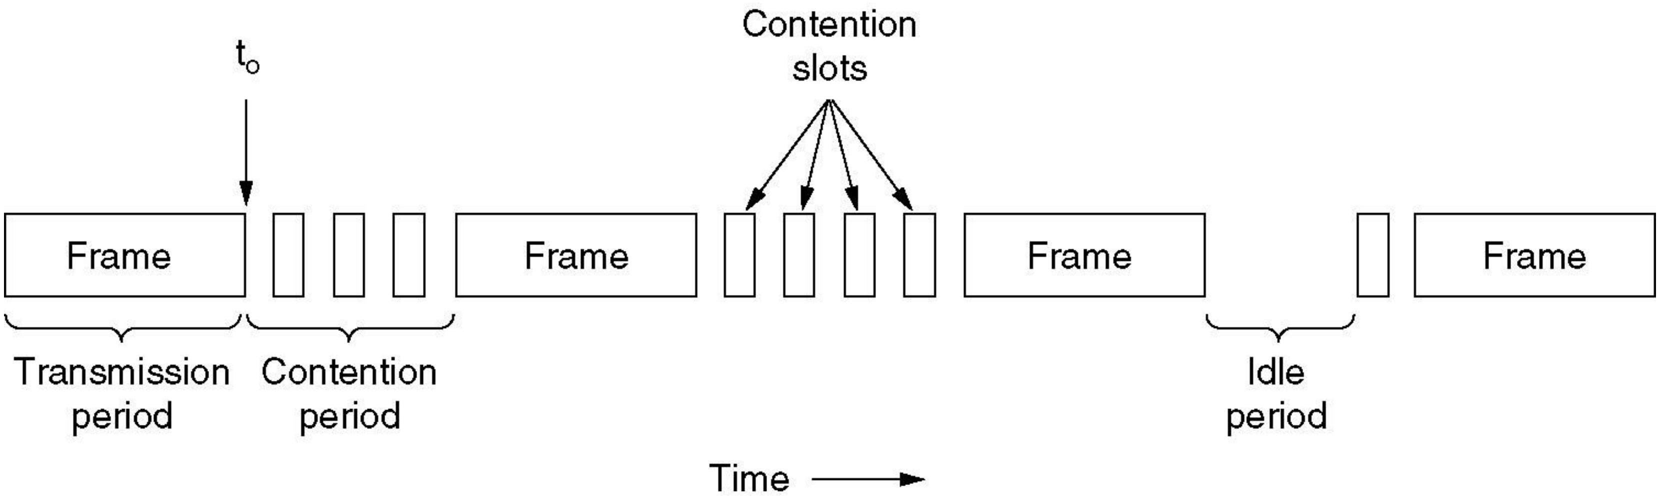

Contention slots
t₀
Frame
Frame
Frame
Frame
Transmission period
Contention period
Idle period
Time

CSMA/CD - eficiência

- T_{prop} = max prop delay entre 2 nós em uma LAN
- t_{trans} = tempo para transmitir um quadro no tamanho máximo

$$efficiency = \frac{1}{1 + 5t_{prop}/t_{trans}}$$

- eficiência tende a 1
  - se t_{prop} tende a 0
  - se t_{trans} tende ao infinito
- desempenho muito superior ao ALOHA: simples, barato e descentralizado!

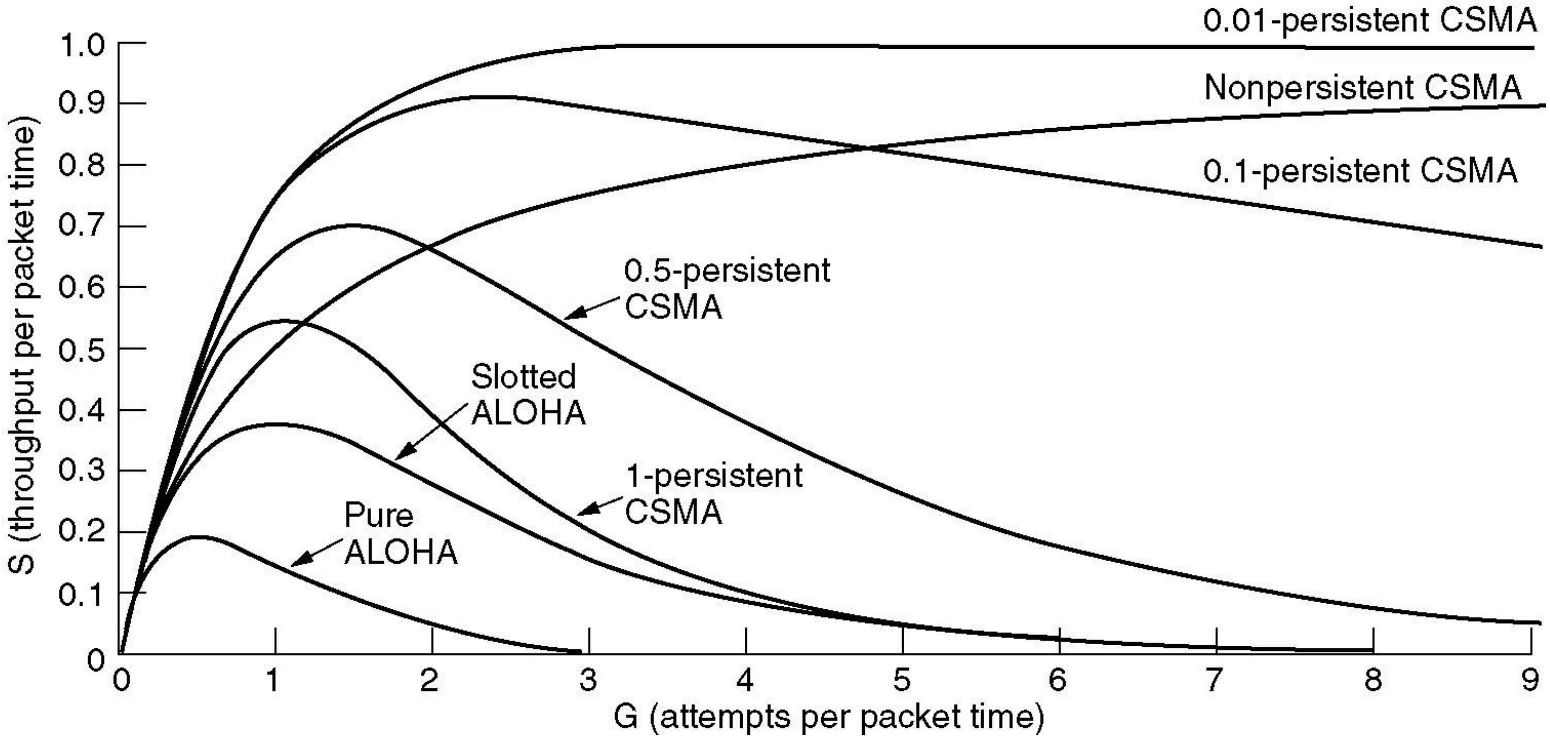
0.01-persistent CSMA
Nonpersistent CSMA
0.1-persistent CSMA
0.5-persistent CSMA
Slotted ALOHA
1-persistent CSMA
Pure ALOHA
S (throughput per packet time)
G (attempts per packet time)
1.0
0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
0 1 2 3 4 5 6 7 8 9

IEEE | IEEE 802

IEEE 802 Working Group
&
Executive Committee Study Group
Home Pages

http://www.ieee802.org/dots.html
Outubro 2009

Active Working Groups and Study Groups

- 802.1 Higher Layer LAN Protocols Working Group
    - Link Security Executive Committee Study Group is now part of 802.1
- 802.3 Ethernet Working Group
- 802.11 Wireless LAN Working Group
- 802.15 Wireless Personal Area Network (WPAN) Working Group
- 802.16 Broadband Wireless Access Working Group
- 802.17 Resilient Packet Ring Working Group
- 802.18 Radio Regulatory TAG
- 802.19 Coexistence TAG
- 802.20 Mobile Broadband Wireless Access (MBWA) Working Group
- 802.21 Media Independent Handoff Working Group
- 802.22 Wireless Regional Area Networks

Inactive Working Groups and Study Groups

- 802.2 Logical Link Control Working Group
- 802.5 Token Ring Working Group

Disbanded Working Groups and Study Groups

- 802.4 Token Bus Working Group (material no longer available on this web site)
- 802.6 Metropolitan Area Network Working Group (material no longer available on this web site)
- 802.7 Broadband TAG (material no longer available on this web site)
- 802.8 Fiber Optic TAG (material no longer available on this web site)
- 802.9 Integrated Services LAN Working Group (material no longer available on this web site)
- 802.10 Security Working Group (material no longer available on this web site)
- 802.12 Demand Priority Working Group
- 802.14 Cable Modem Working Group (material no longer available on this web site)
- QOS/FC Executive Committee Study Group (material no longer available on this web site)

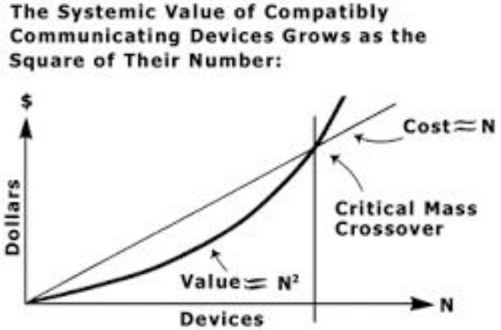
The Systemic Value of Compatibly Communicating Devices Grows as the Square of Their Number:
$
Dollars
Cost ≃ N
Critical Mass Crossover
Value ≃ N²
Devices
N

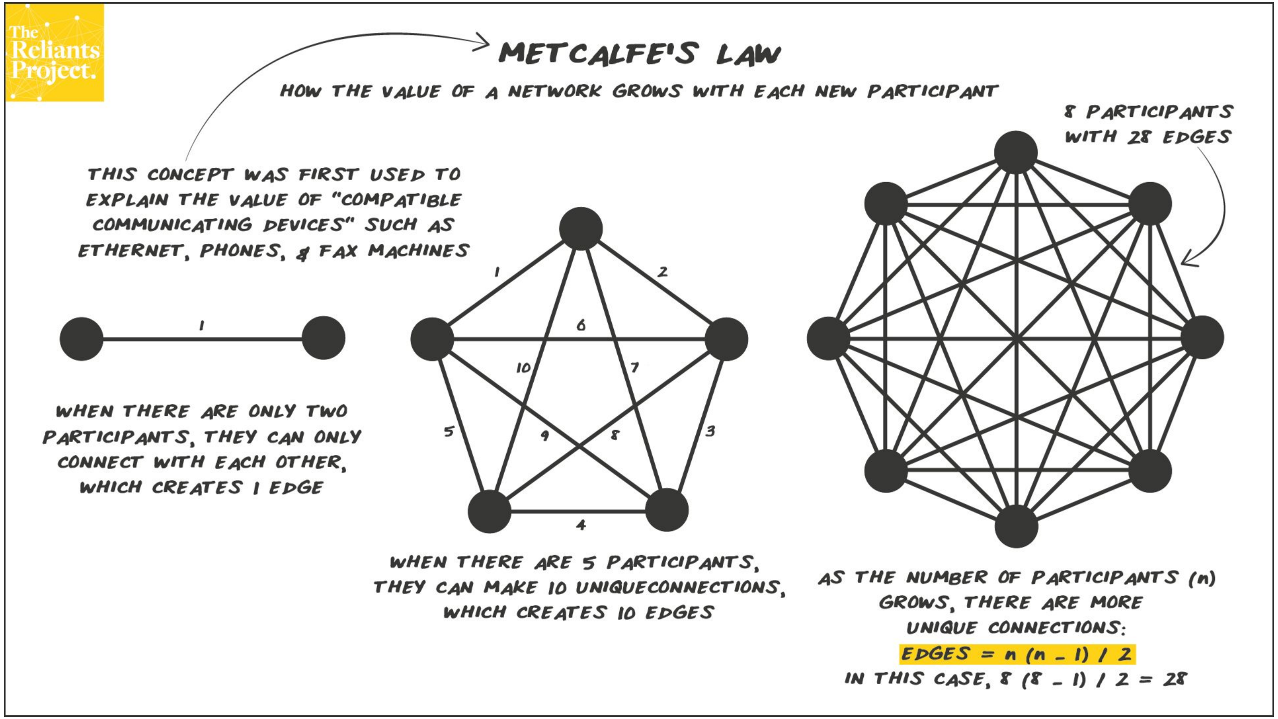
The Reliants Project.
METCALFE'S LAW
HOW THE VALUE OF A NETWORK GROWS WITH EACH NEW PARTICIPANT
THIS CONCEPT WAS FIRST USED TO EXPLAIN THE VALUE OF "COMPATIBLE COMMUNICATING DEVICES" SUCH AS ETHERNET, PHONES, & FAX MACHINES
1
WHEN THERE ARE ONLY TWO PARTICIPANTS, THEY CAN ONLY CONNECT WITH EACH OTHER, WHICH CREATES 1 EDGE
1
2
3
4
5
6
7
8
9
10
WHEN THERE ARE 5 PARTICIPANTS, THEY CAN MAKE 10 UNIQUECONNECTIONS, WHICH CREATES 10 EDGES
8 PARTICIPANTS WITH 28 EDGES
AS THE NUMBER OF PARTICIPANTS (n) GROWS, THERE ARE MORE UNIQUE CONNECTIONS:
EDGES = n (n - 1) / 2
IN THIS CASE, 8 (8 - 1) / 2 = 28

Onde isso está sendo usado hoje?

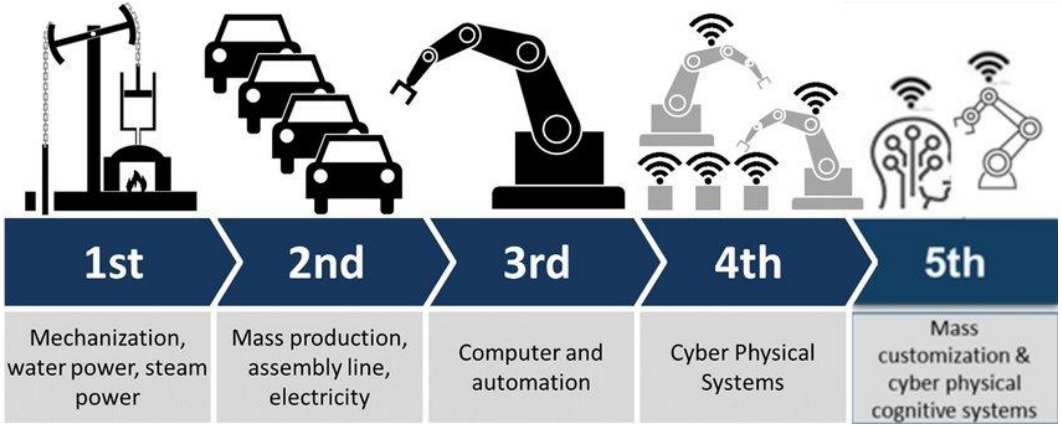

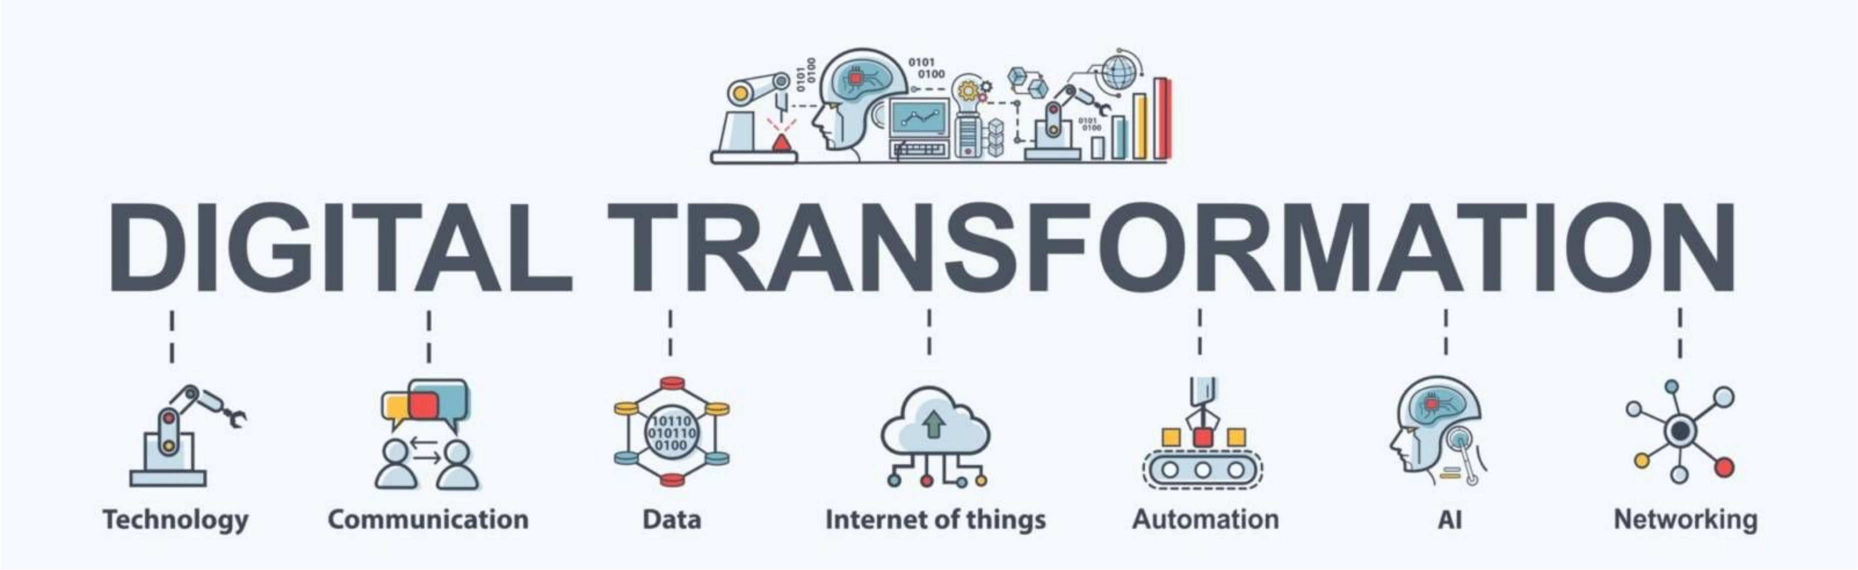
DIGITAL TRANSFORMATION
Technology
Communication
Data
Internet of things
Automation
AI
Networking

"I lucked out and got a problem nobody had ever had before -- a building full of PCs that needed high-speed connection to a laser printer and the Internet (known as ARPAnet then)."

– Robert Metcalfe


LIÇÕES
APRENDIDAS
gerenciamento
de projetos

Então quem é o Uncle Bob mais legal?

Bob Metcalfe

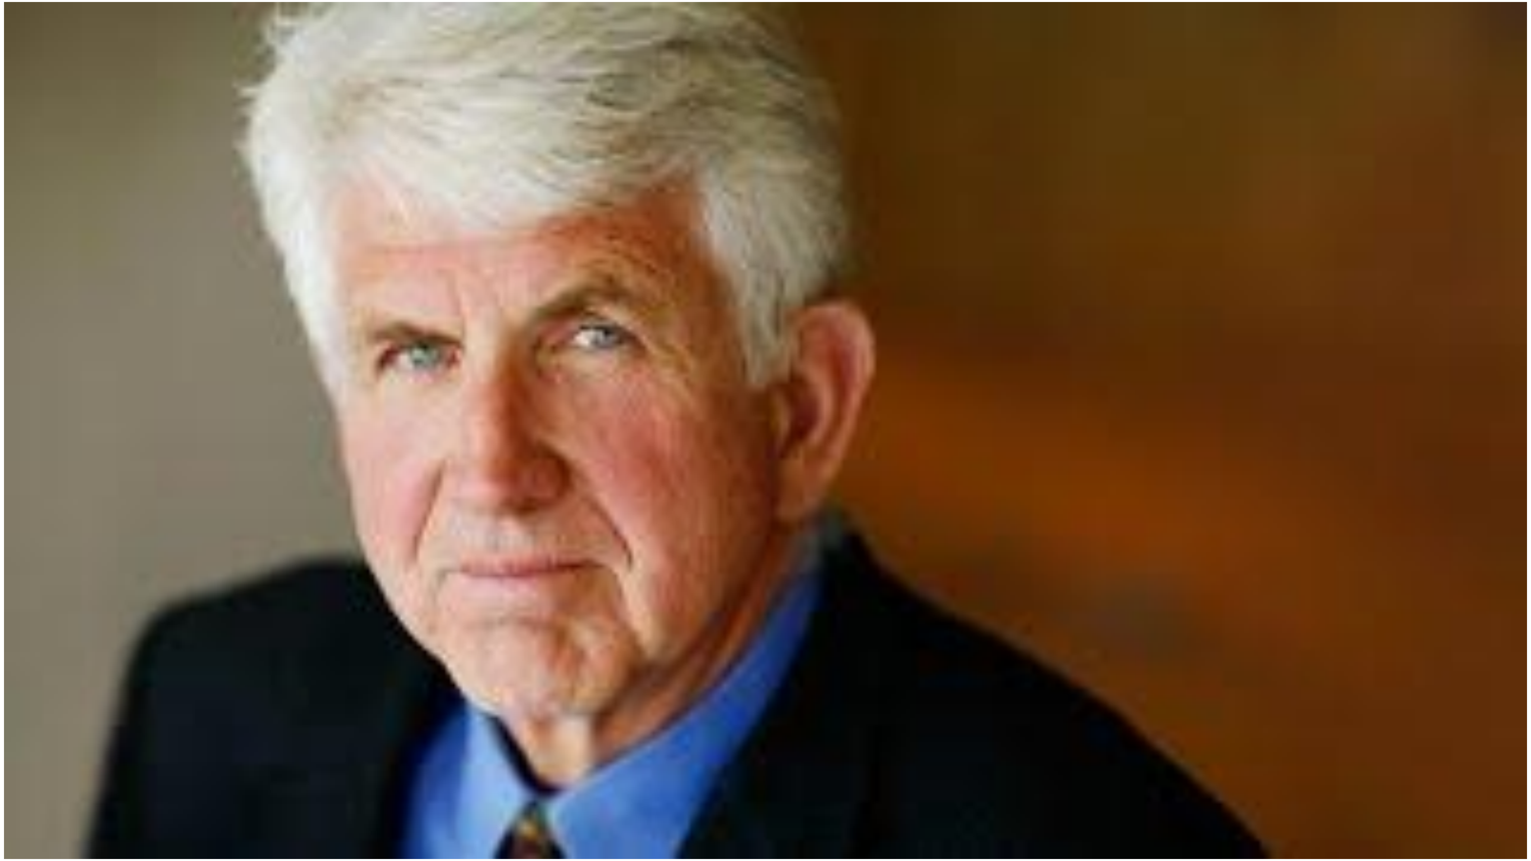

Conectando os Pontos - A Contribuição de Robert Metcalfe Para a Computação

Claudio Miceli de Farias -Universidade Federal do Rio de Janeiro
cmicelifarias@cos.ufrj.br

Instituto Tércio Pacitti de Aplicações e Pesquisas Computacionais

Engenharia de Sistemas e Computação PESC

UFRJ faz 100 ANOS 1920 | 2020

UNIVERSIDADE FEDERAL DO RIO DE JANEIRO

PR-2 / UFRJ