Internet Protocols

Thomas Fuhrmann

Network Architectures
Computer Science Department
Technical University Munich

Introduction

In this course you will learn …

… to understand the technology in this box.

Bildquelle: Online-Katalog

What is communication?

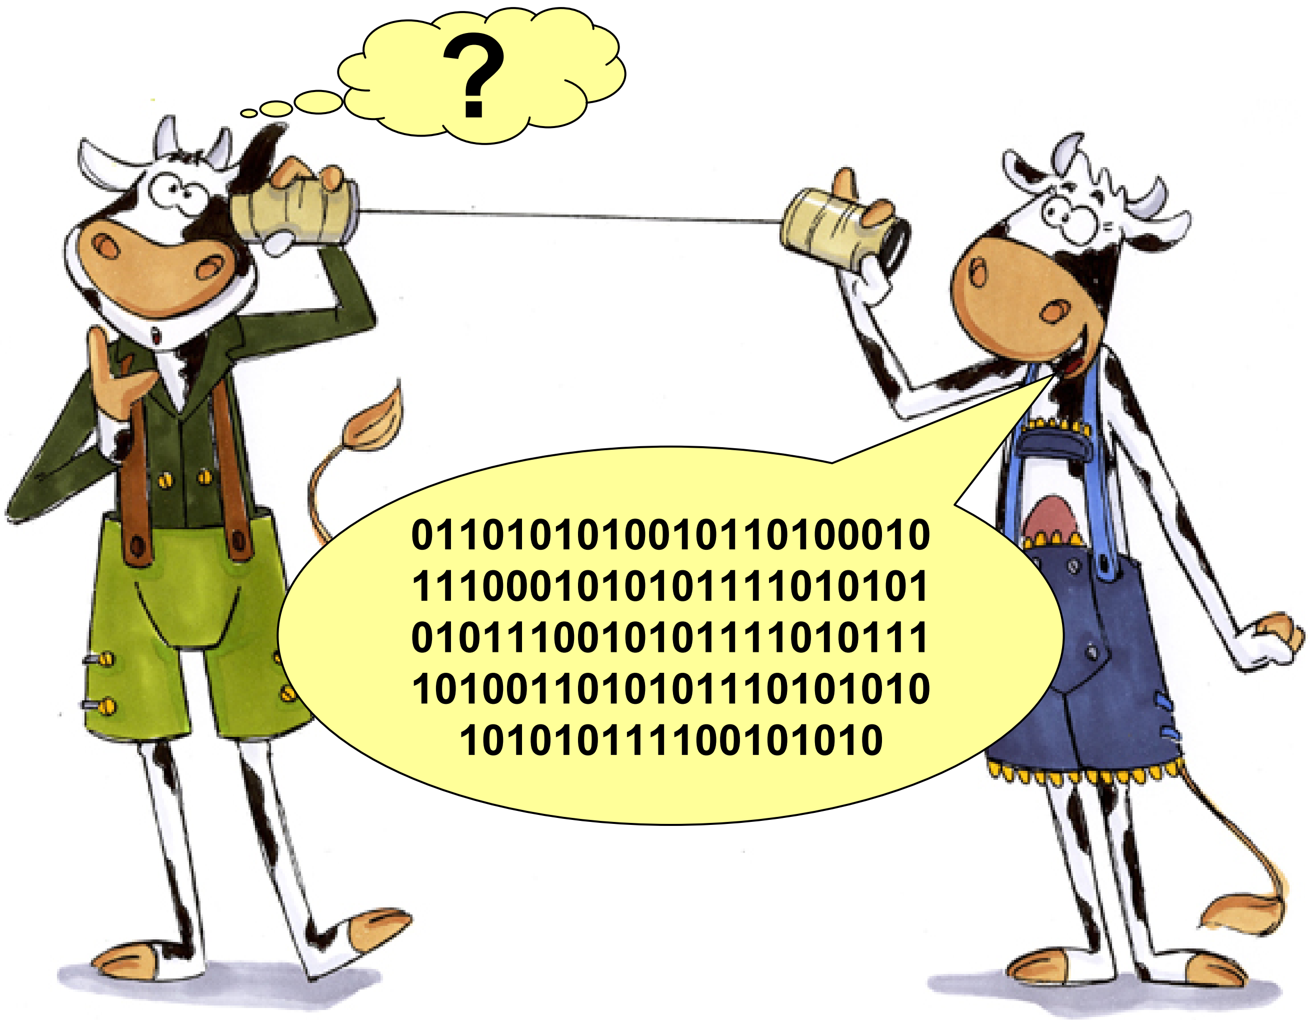

Bild: © Leeni Arts 2007

Social Interactions

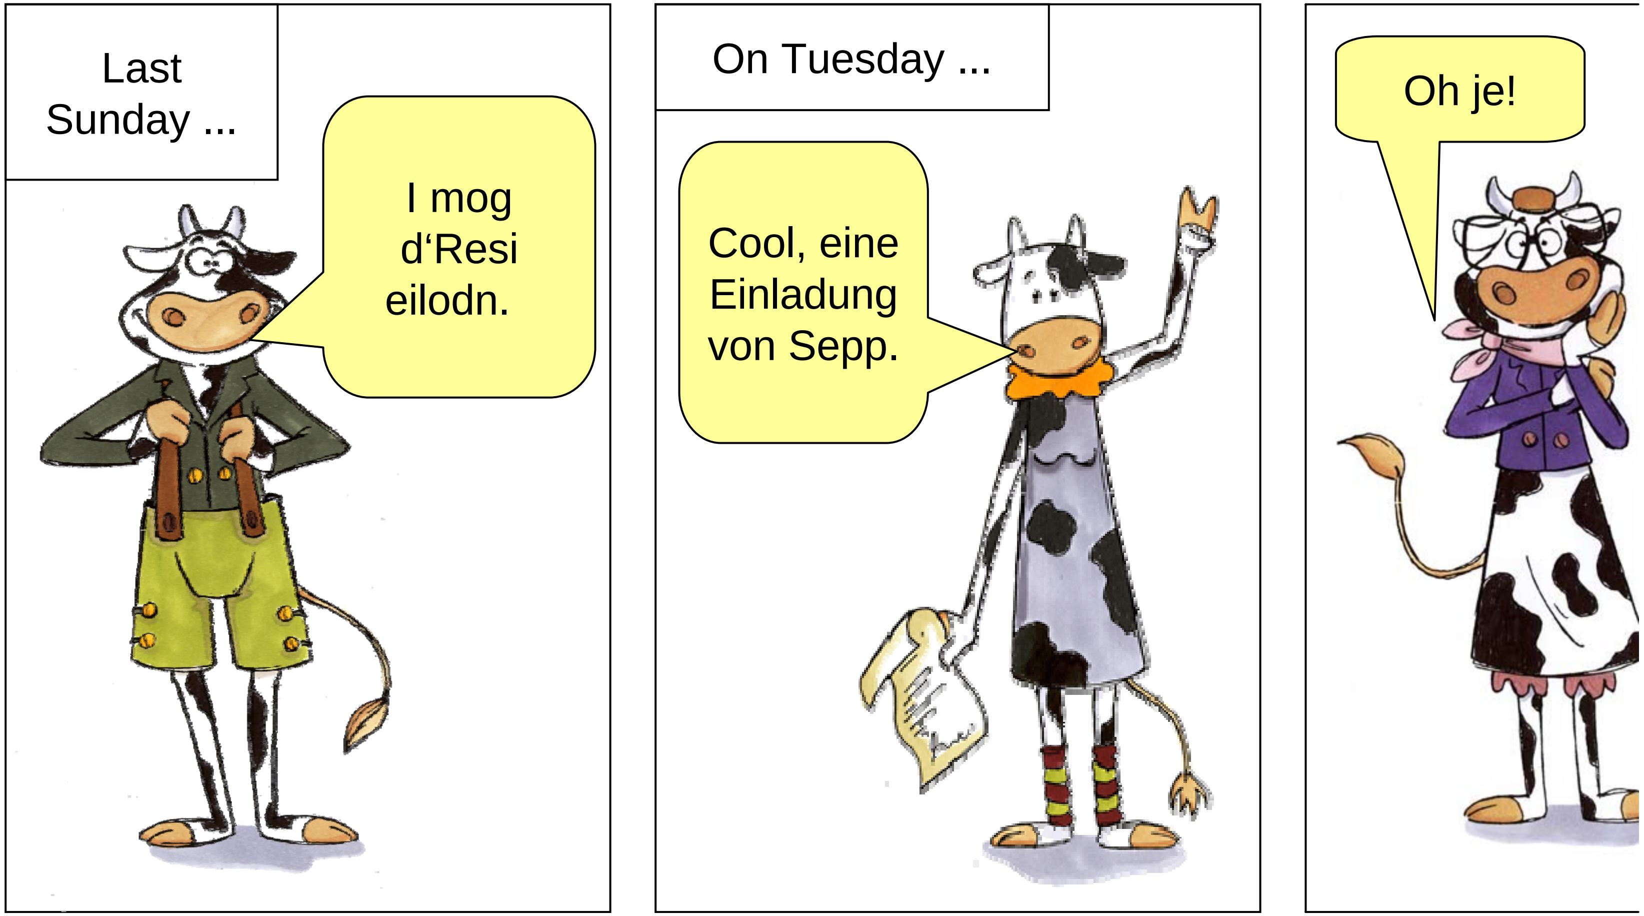

Bilder: © Leeni Arts 2007

Path-Time Diagram

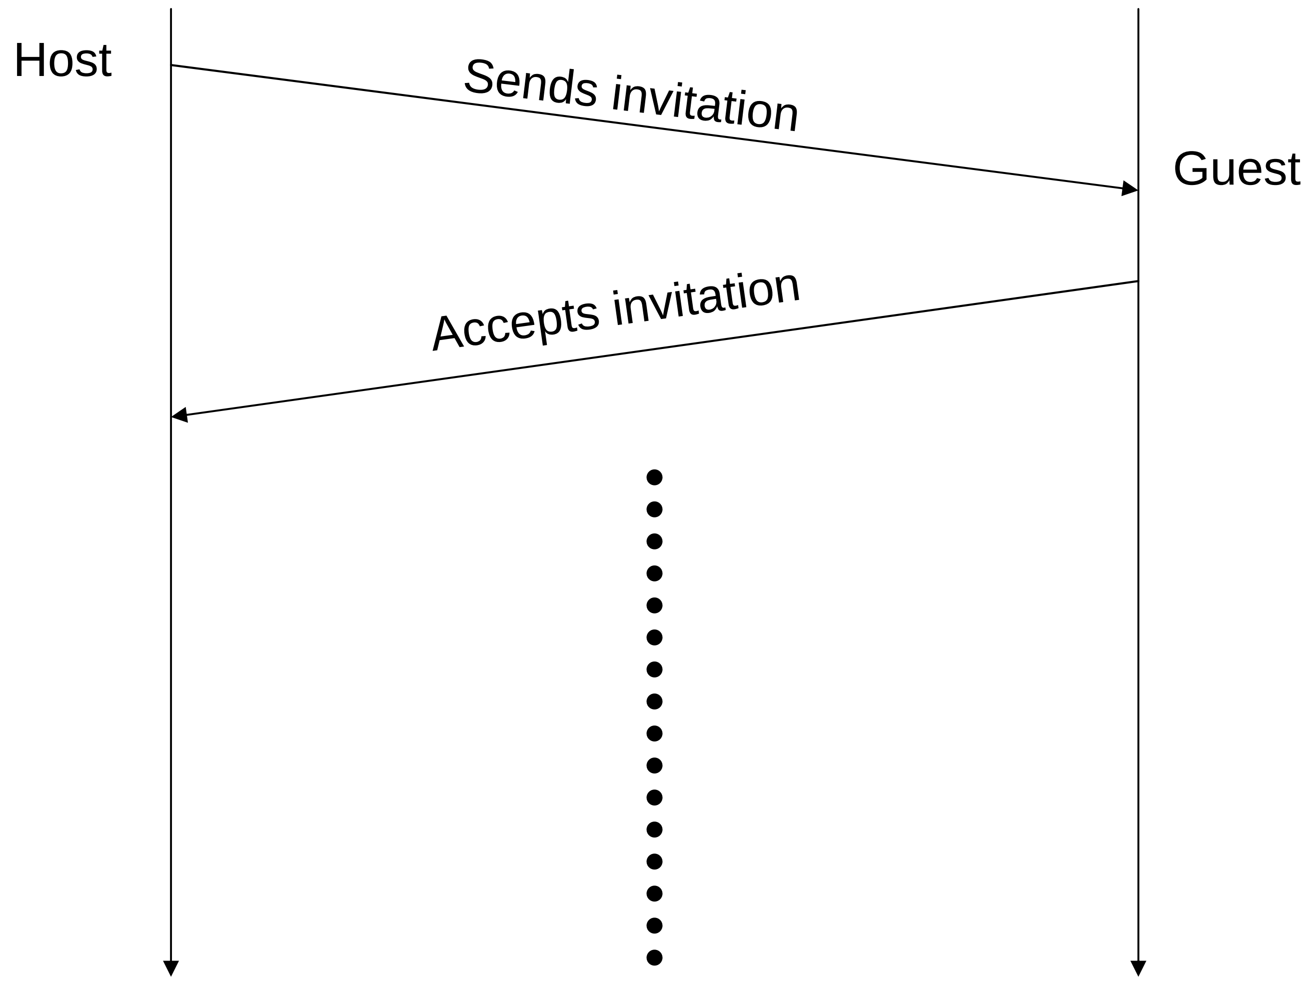

State Diagram

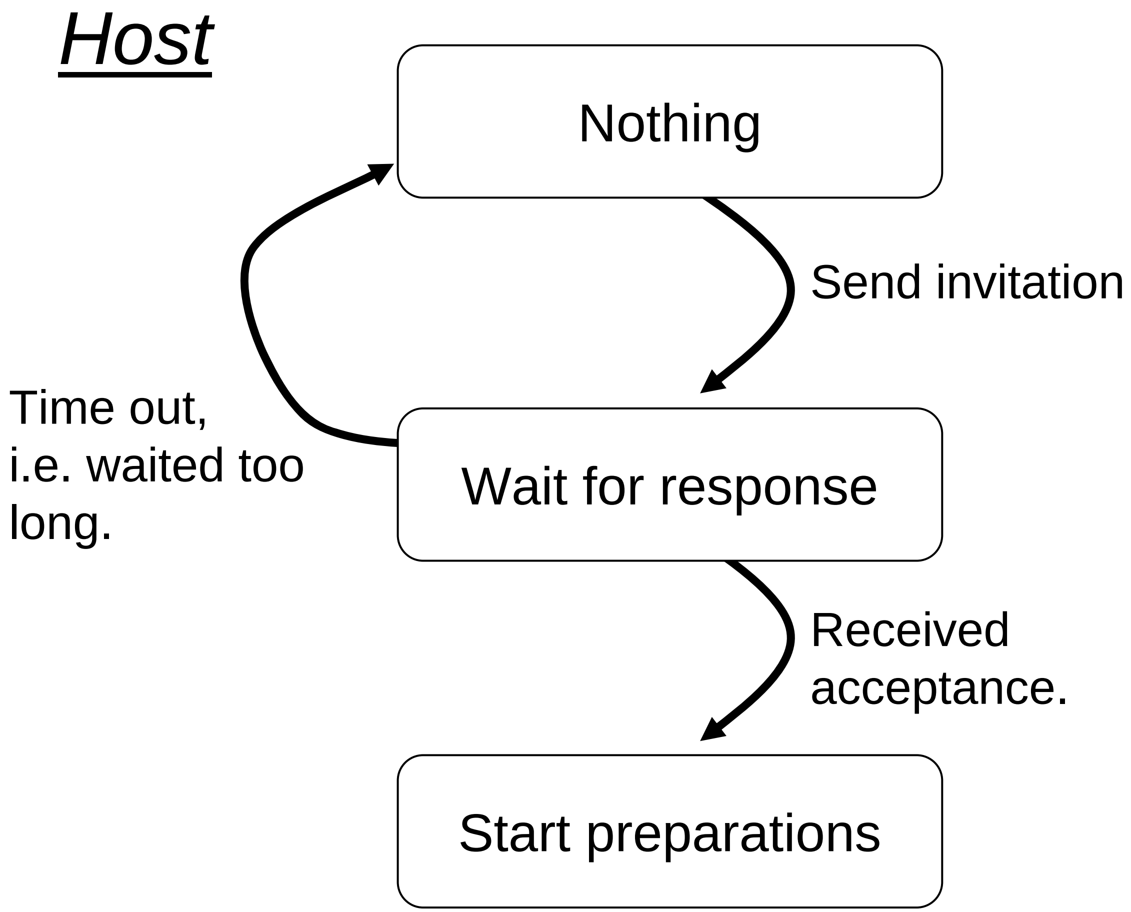

Protokoll – Eine Definition

πρωτόκολλον

aus πρώτος „erster“ und κόλλα „Leim“
„[den amtlichen Papyrusrollen] vorgeleimtes Blatt“

Ein Protokoll hält oder legt fest, zu welchem Zeitpunkt (absolut) oder in welcher Reihenfolge (relativ) welche Entität (wer oder was) welche Operation ausgeführt hat oder ausführen wird. Protokollierung bezeichnet dabei die Operation des Festlegens bzw. Niederschreibens der drei Protokollbestandteile Zeitpunkt, Entität, Operation.

1. Introduction to Protocols

2. Signals, Data, Information

3. Layering

4. Internet Overview

Protokolle verbinden physisch getrennte Entitäten

Daten → Daten‘

Die physische Welt: Daten werden auf physischen Medien gespeichert und übertragen.

Protokoll

Information, Daten, Signale

Information: Gegenstände unseres Denkens

Abstrakte Welt

Konventionen zur Darstellung von *Denkinhalten*

Daten: formalisierte Darstellung

„Baum“

Physikalische Welt

Konventionen zur Darstellung von *Daten*

Signale: reale Darstellung von Daten

Information, Daten, Signale

Information	Daten	Signale
Die Bedeutung, die ein Mensch aufgrund der den Daten zugrundeliegenden Vereinbarungen (Konventionen) diesen beilegen kann. Menschen und Maschinen können Daten handhaben, aber nur der Mensch kann aus Daten Informationen gewinnen.	Darstellung von Sachverhalten (Fakten), Konzepten, Vorstellungen und Anweisungen in formalisierter Weise, die für die Kommunikation, Interpretation und die Verarbeitung durch Menschen und/oder technische Mittel geeignet ist.	Physikalische Darstellung (Repräsentation, Codierung) von Daten durch charakteristische räumliche und/oder zeitliche Veränderungen der Werte physikalischer Größen.

Information, Daten, Signale

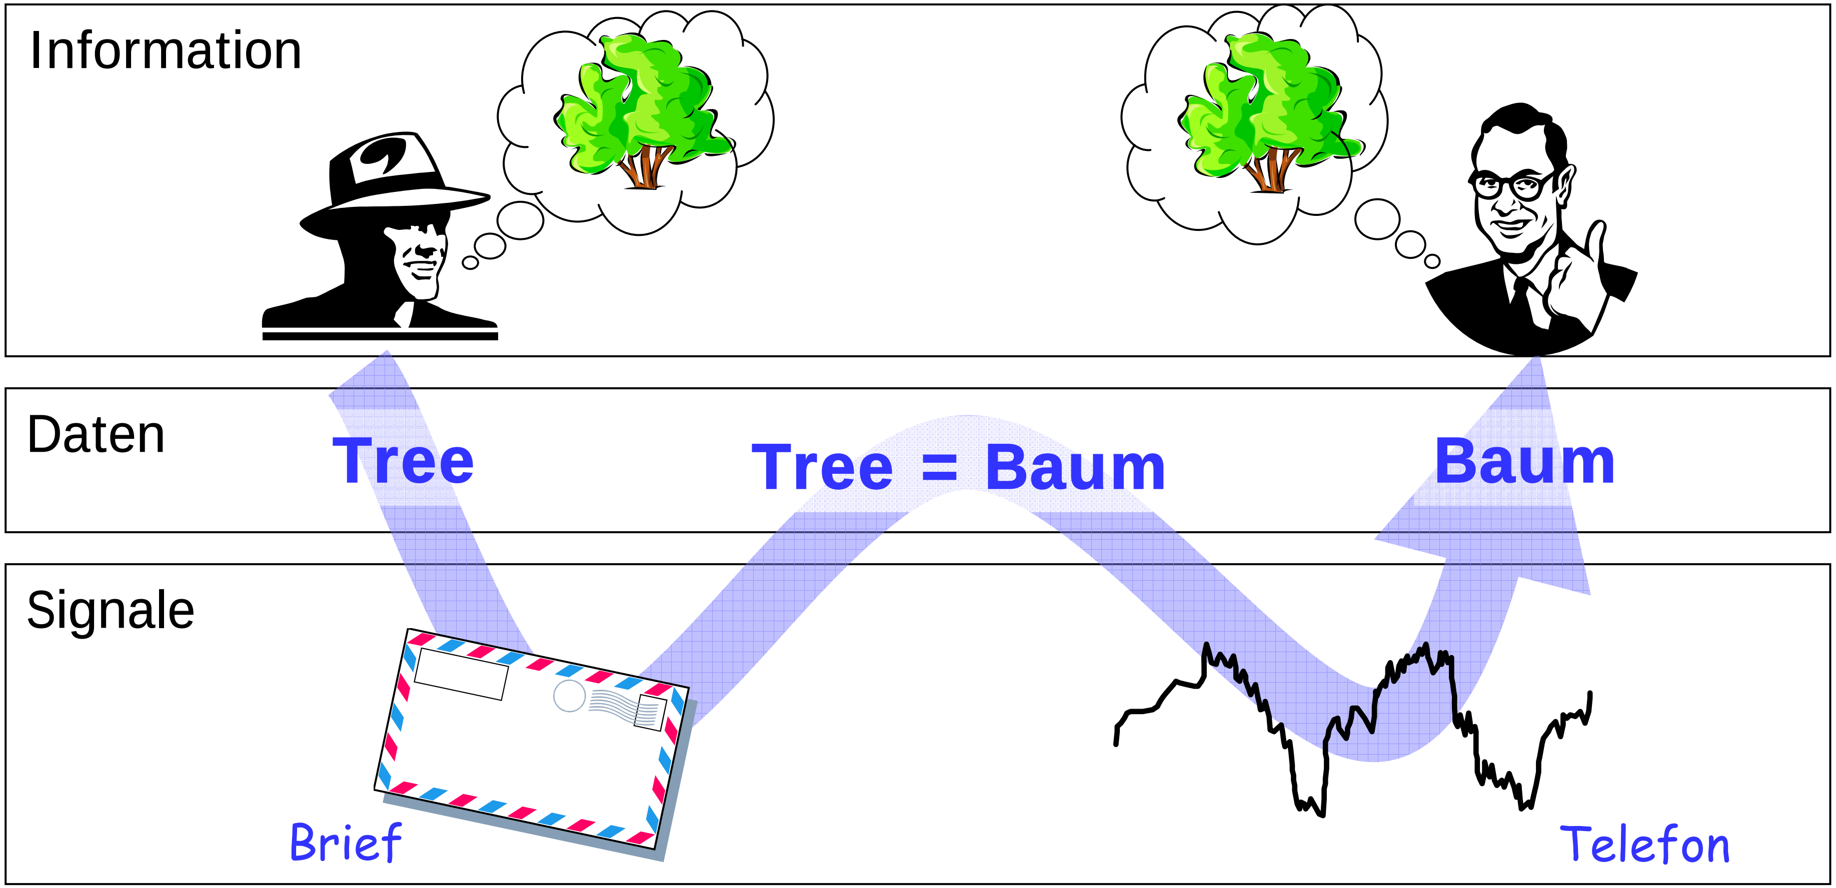

1. Introduction to Protocols

2. Signals, Data, Information

3. Layering

4. Internet Overview

Das ISO/OSI-Referenzmodell

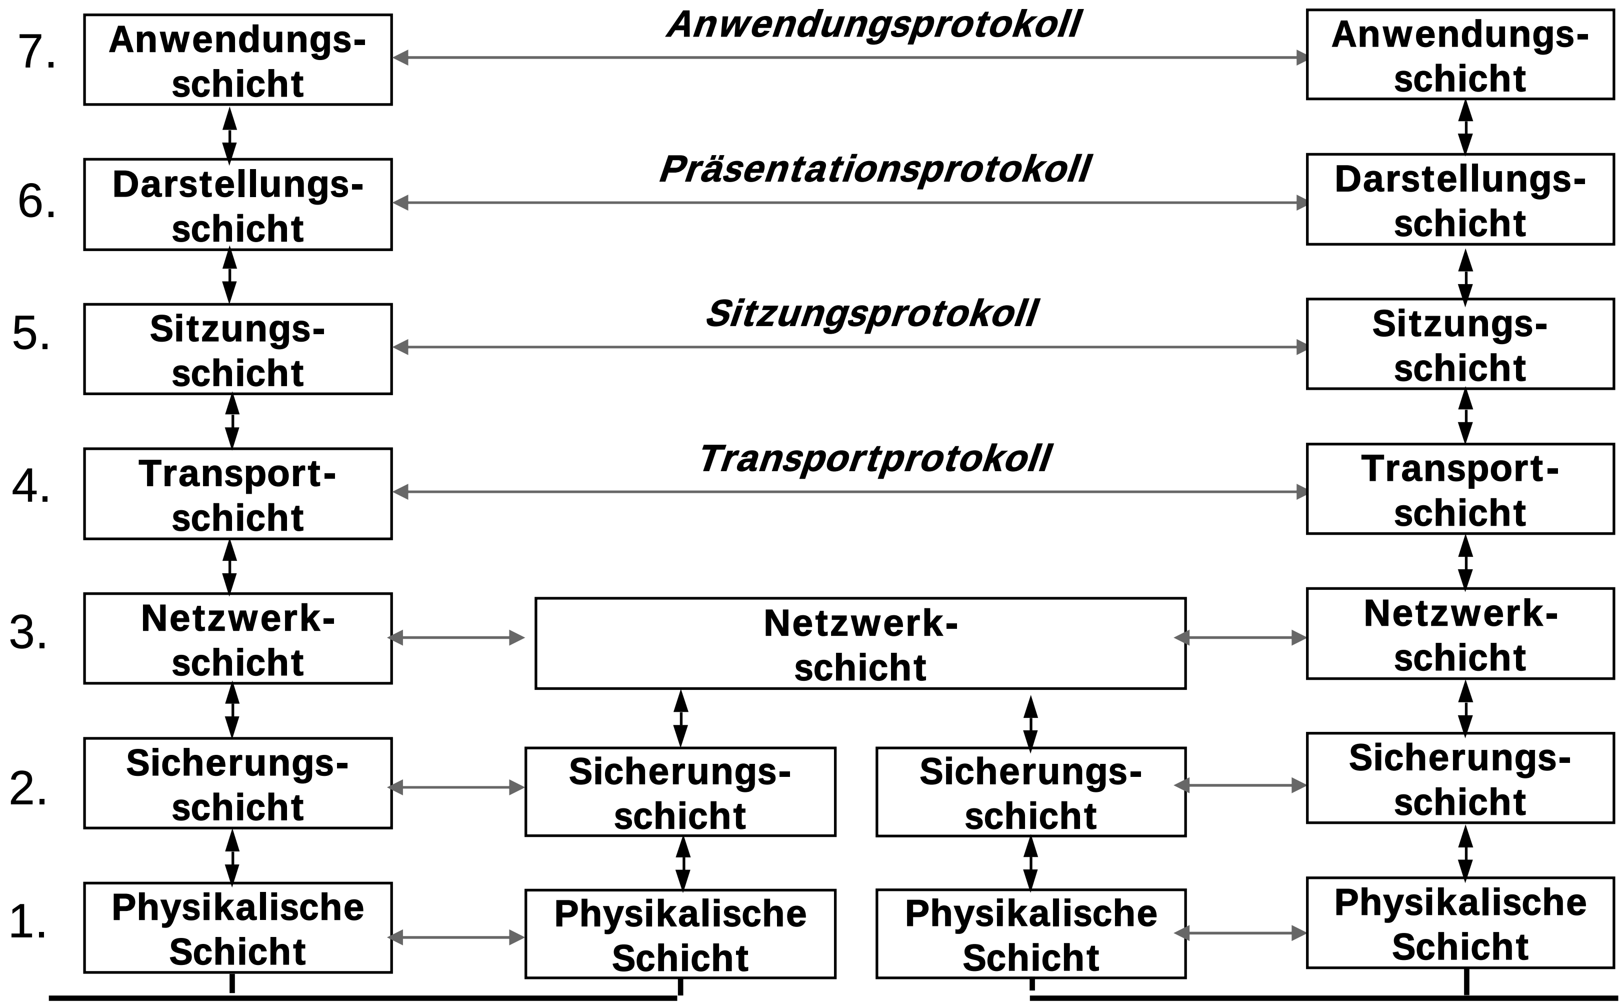

The TCP/IP Model

7.	Application Layer	E-Mail, WWW, Chat, …
4.	Transport Layer	TCP (=Transmission Control Protocol), UDP (=User Datagram Protocol)
3.	Network Layer	IP (=Internet Protocol)
2.	Data Link Layer	Ethernet, Modem, DSL, …
1.	Physical Layer	Copper, Fiber, Radio waves, …

The TCP/IP protocol stack is compared to an hour glass: Lots of different PHYs at the bottom and lots of applications on top, but only IP in the middle.

A Note on Layering

- Layers are not per se uniquely defined.
  - Systems sometimes suggest natural boundaries between different functionalities (→Interface).
  - Generally, layers should simplify matters.
  - A good layer and interface design allows the substitution of protocols in the layers.

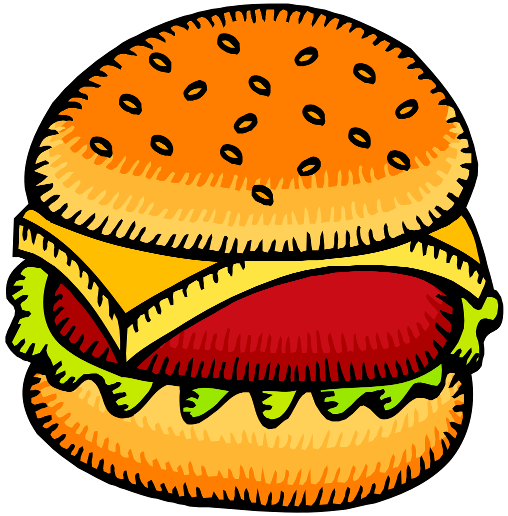

Network layering is not MacDonals layering!!!

- In practice …
  - Systems combine different functionality in one layer;
  - Systems need more and more cross-layering because the original design did not foresee current requirements.

Layers versus Planes

- Data plane
  - Actual data transport between stations
- Control plane
  - Setup and tear-down of circuits, etc.
  - Signaling of quality of service requirements, etc.
- Management plane
  - Status and maintenance info of networking equipment
  - Management of the entire network

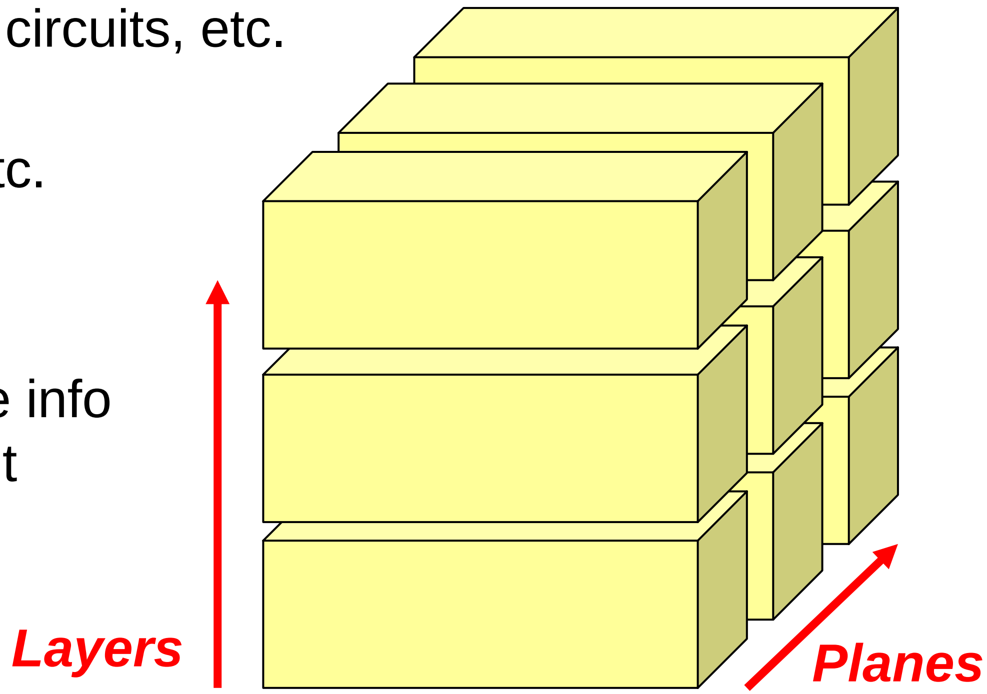

Das Grundmodell der Datenkommunikation

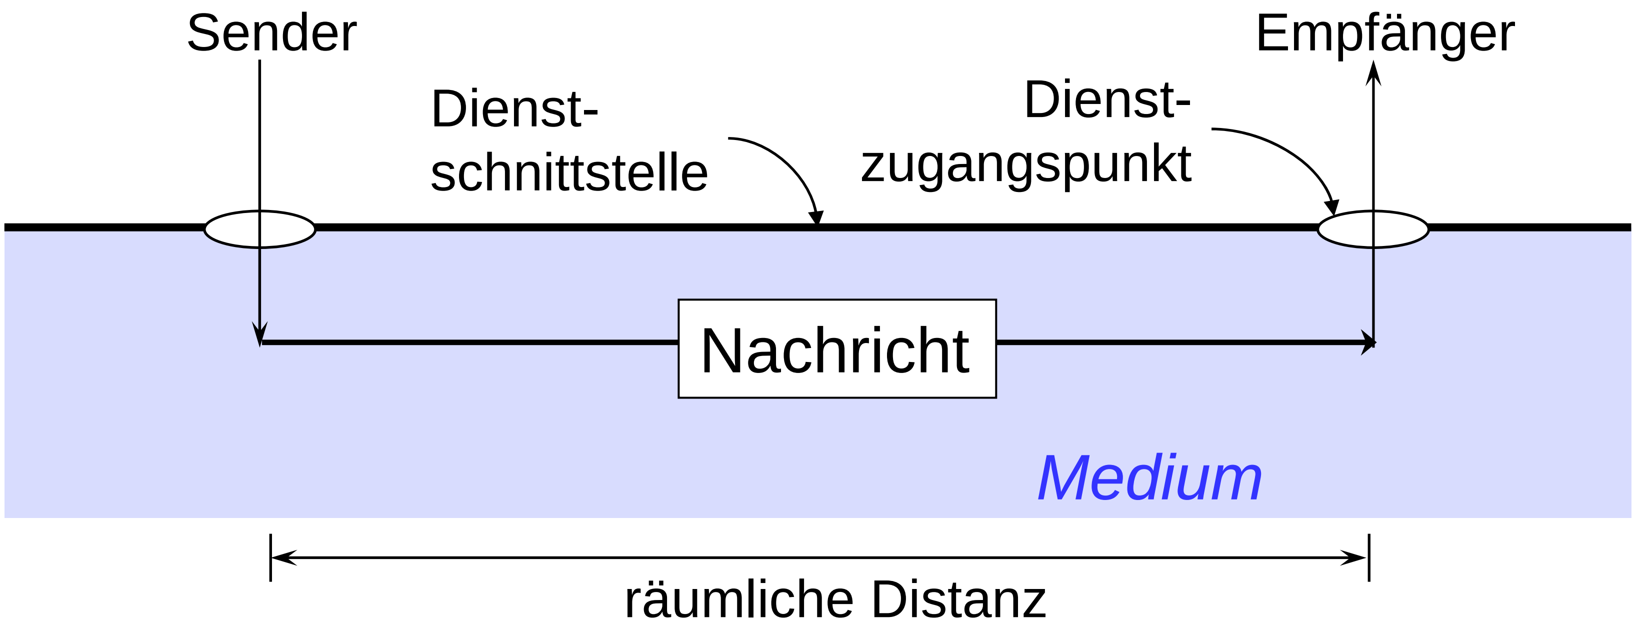

- Teilnehmer agieren als *Sender* oder *Empfänger.*
- Die Dienstnutzung durch Teilnehmer erfolgt an einer speziellen *Dienstschnittstelle* unter Nutzung eines *Dienstzugangspunkts*.
- Durch das *Medium* wird die räumliche Distanz überbrückt.

Kommunikation in Schichten

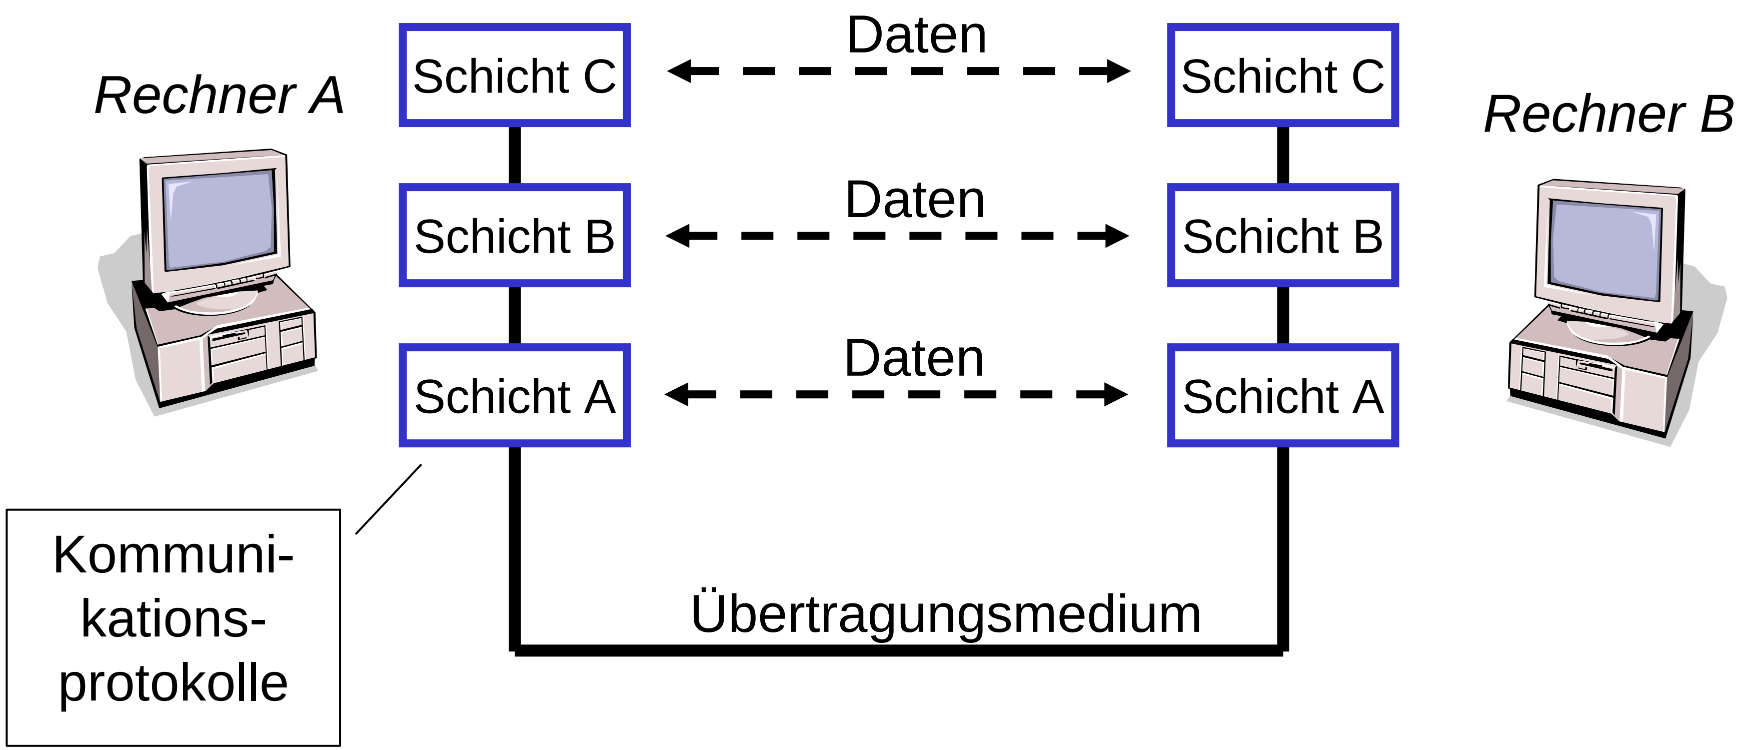

Die Schichtung jeweils aufeinander aufbauender Protokolle heißt Protokollstapel

Protokoll-Schnittstellen

Zwei benachbarte Schichten aus dem Protokollstapel arbeiten als *Dienstgeber* und *Dienstnehmer* zusammen

- Die jeweils untere Schicht erbringt einen *Dienst* für die darüber liegende Schicht
- Diese Dienste müssen in einer *Schnittstellen-definition* festgelegt sein

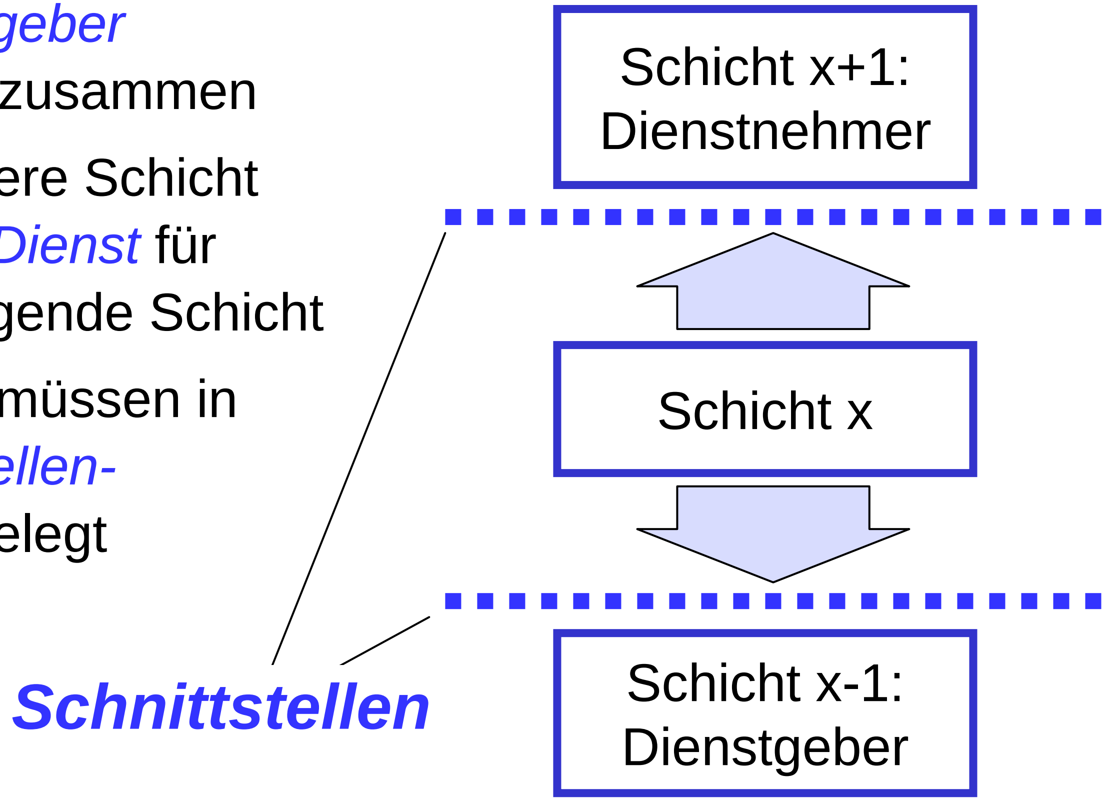

Scheinbare & tatsächliche Kommunikation

- Alle Kommunikationsprotokolle laufen innerhalb ihrer jeweiligen Schicht ab, d.h. ihre Kommunikation findet *scheinbar* direkt zwischen Rechner A und B statt (*horizontal*)
- Tatsächlich erbringen die tiefer liegenden Schichten diesen Kommunikationsdienst, d.h. die *tatsächliche* Kommunikation verläuft innerhalb des Protokollstapels (*vertikal*)

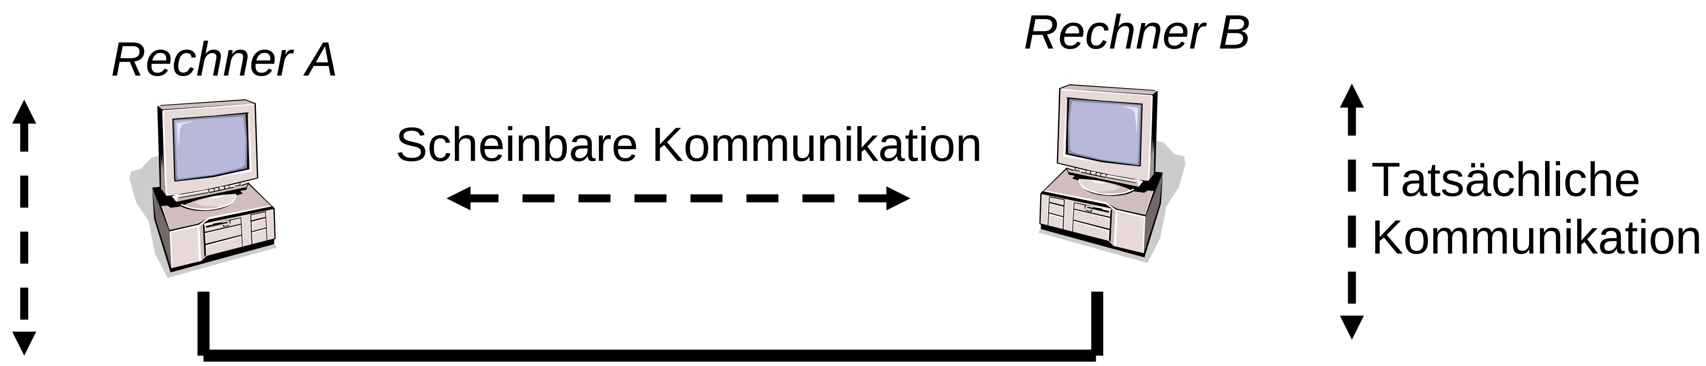

1. Introduction to Protocols

2. Signals, Data, Information

3. Layering

4. Internet Overview

The Internet versus the Telephone Network

- How was the ARPANET different from existing telecommunication networks?
  - Machine-to-machine communication
  - Packet switched network
  - Self-healing properties (wrt routing, data rate adjustment, etc.)
  - A network of autonomous networks.

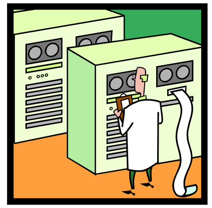

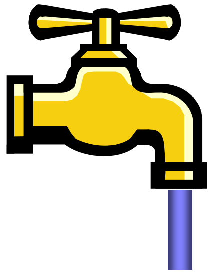

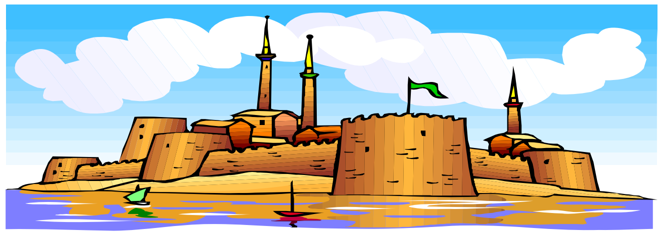

Machines versus People

- Telephone calls had been well understood ($\rightarrow$e.g. Erlang formula).
- People probabilistically start a call with a probabilistic duration.
- Calls are admitted as long as resources suffice.

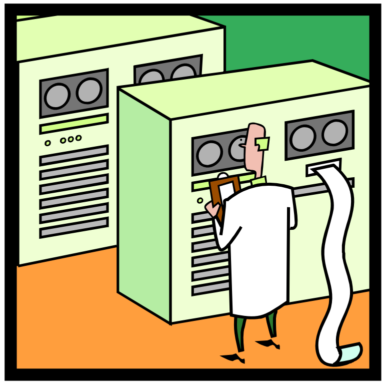

- Remote connections have low but bursty data volume ($\rightarrow$ e.g. lines of text) and need low latency.
- File transfers may be large, but have loose latency requirements.

Packets versus Circuits

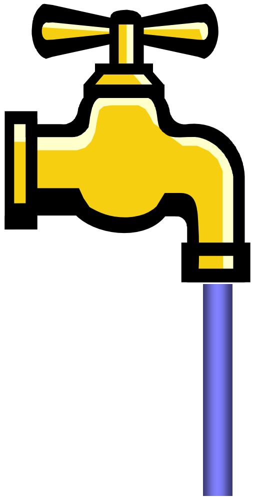

- Circuit switched connections provide a constant data rate which is know in advance.
- Circuit switched connections provide very low latency.
- Connections are set-up and then used for a considerable length of time.
- Circuit switched networks and/or connections require expensive equipment.

- Packets switched networks potentially introduce considerable queuing delay.
- Packets may be lost when resource demand exceeds available resources.
- Packet switching can provide large bandwidth with inexpensive equipment.

Self-Organization versus Administration

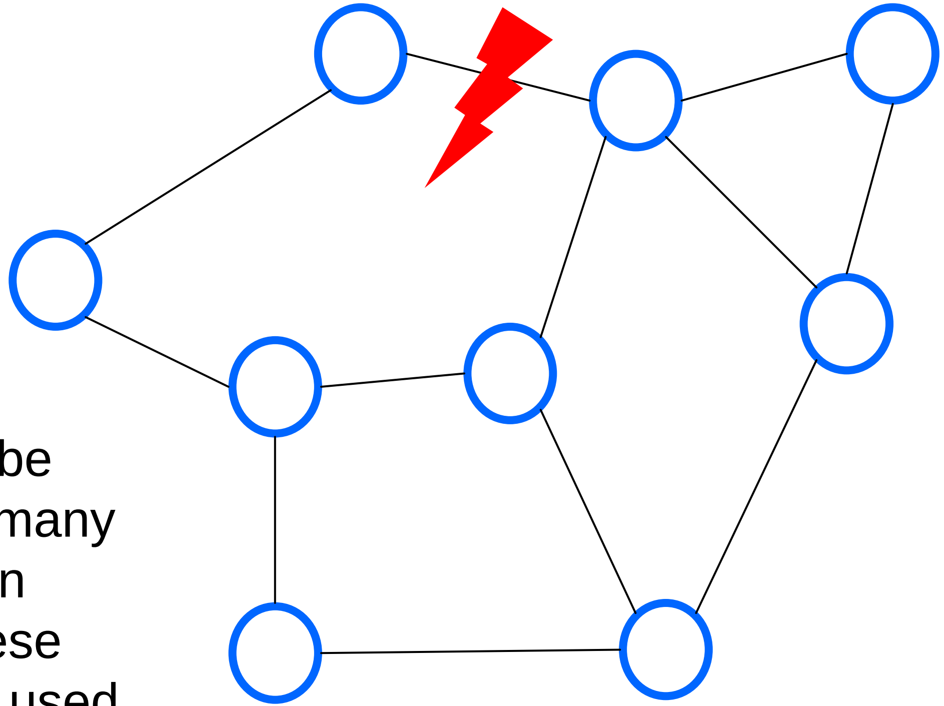

The ARPANET was designed to be self-organizing in the sense that many alternative paths were available in case a component failed; and these alternatives should be found and used without manual set-up.

If necessary, the data rate should automatically adjust to the varying conditions ($\rightarrow$best effort network).

“It was from the RAND study that the false rumor started claiming that the ARPANET was somehow related to building a network resistant to nuclear war. This was never true of the ARPANET.”

Source: Barry M. Leiner, et al.
A Brief History of the Internet.

Autonomy versus Centralization

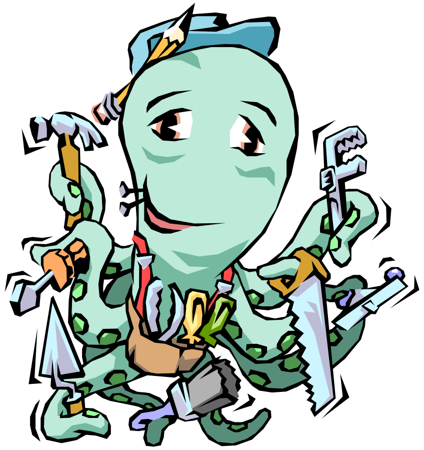

- Telecommunication companies strive to provide full service.
- The telecom provider manages and controls everything, often including the terminal device.

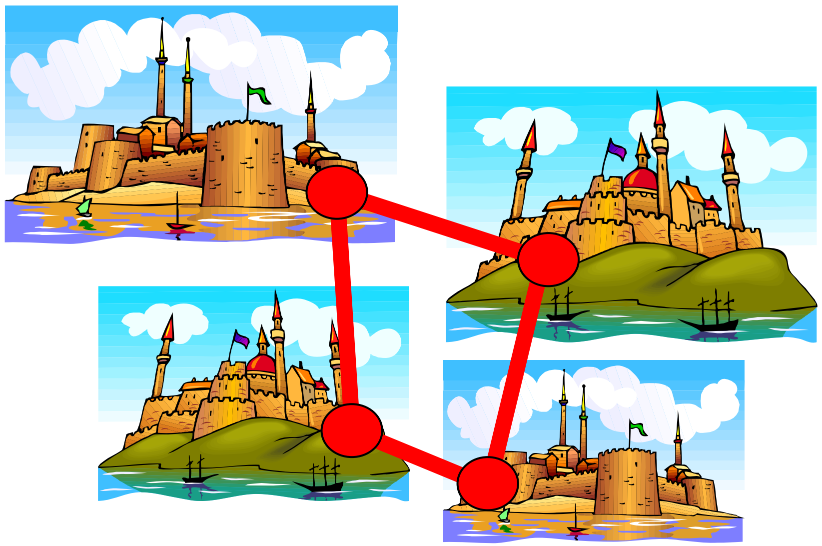

- The ARPANET was designed as federation of autonomous networks.
- Minimal regulation, maximal freedom.
- Consensus by technical discussion and democratic process.

Kahn‘s Four Rules

Four ground rules were critical to Kahn's early thinking:

- Each distinct network would have to stand on its own and no internal changes could be required to any such network to connect it to the Internet.
- Communications would be on a best effort basis. If a packet didn't make it to the final destination, it would shortly be retransmitted from the source.
- Black boxes would be used to connect the networks; these would later be called gateways and routers. There would be no information retained by the gateways about the individual flows of packets passing through them, thereby keeping them simple and avoiding complicated adaptation and recovery from various failure modes.
- There would be no global control at the operations level.

Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff . *A Brief History of the Internet.*

The End-to-End Principle

Reliable systems
need reliability
mechanisms at the
terminal device!

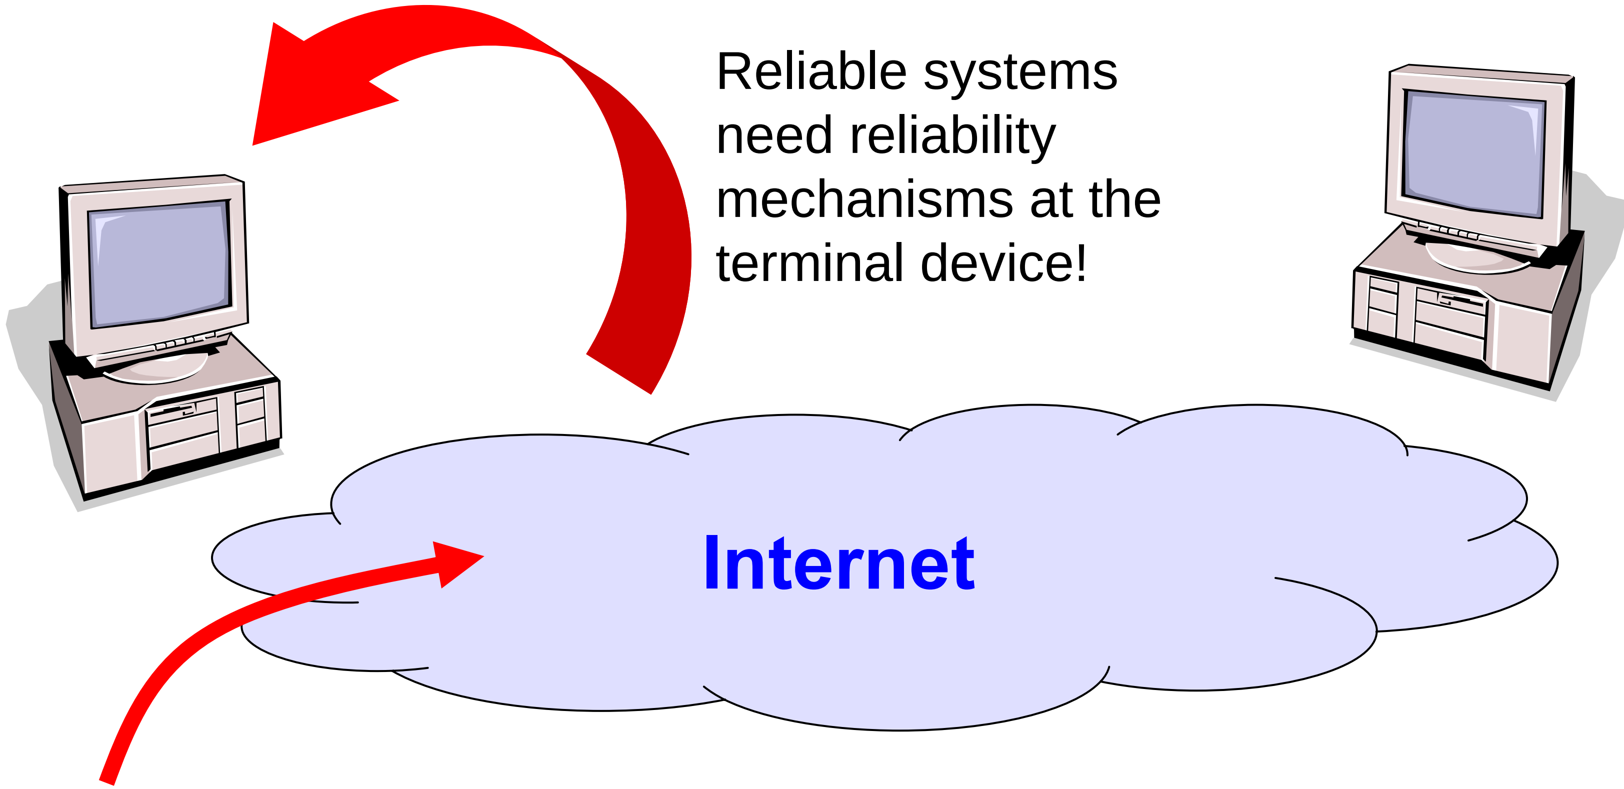

In network devices
may optimize …

Jerome H. Saltzer, David P. Reed, and David D. Clark
End-to-end arguments in system design
ACM Transactions on Computer Systems 2, 4
Nov. 1984, pp. 277-288

Key Issues of the Early Internet

Other key issues that needed to be addressed were:

- Algorithms to prevent lost packets from permanently disabling communications and enabling them to be successfully retransmitted from the source.
- Providing for host to host "pipelining" so that multiple packets could be enroute from source to destination at the discretion of the participating hosts, if the intermediate networks allowed it.
- Gateway functions to allow it to forward packets appropriately. This included interpreting IP headers for routing, handling interfaces, breaking packets into smaller pieces if necessary, etc.
- The need for end-end checksums, reassembly of packets from fragments and detection of duplicates, if any.
- The need for global addressing.
- Techniques for host to host flow control.
- Interfacing with the various operating systems.
- There were also other concerns, such as implementation efficiency, internetwork performance, but these were secondary considerations at first.

Barry M. Leiner, et al. *A Brief History of the Internet.*

ARPANET Key Applications

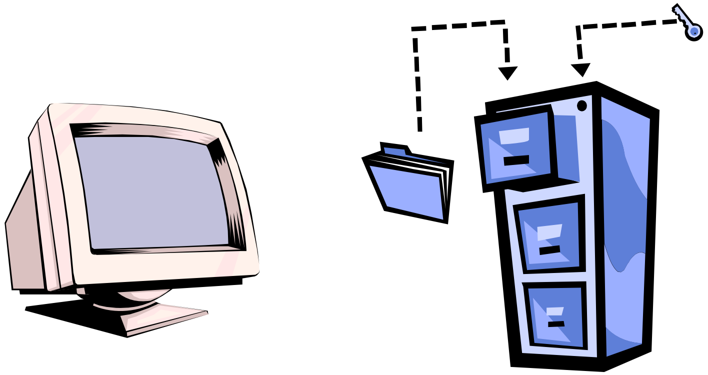

- The ARPANET founders targeted two applications:
  - Login to remote computers (such as SSH today)
  - File transfer (such as FTP & HTTP today)

- But it turned out that electronic mail was the actual killer application!

- Not until two decades later the World Wide Web was invented …

Request for Comments – RFC 1

```
Network Working Group                                     Steve Crocker
Request for Comments: 1                                            UCLA
                                                           7 April 1969

Introduction

     The software for the ARPA Network exists partly in the IMPs and
     partly in the respective HOSTs. BB&N has specified the software of
     the IMPs and it is the responsibility of the HOST groups to agree
     on HOST software.

     During the summer of 1968, representatives from the initial four
     sites met several times to discuss the HOST software and initial
     experiments on the network. There emerged from these meetings a
     working group of three, Steve Carr from Utah, Jeff Rulifson from
     SRI, and Steve Crocker of UCLA, who met during the fall and winter.
     The most recent meeting was in the last week of March in Utah. Also
     present was Bill Duvall of SRI who has recently started working
     with Jeff Rulifson.

     Somewhat independently, Gerard DeLoche of UCLA has been working on
     the HOST-IMP interface.

     I present here some of the tentative agreements reached and some of
     the open questions encountered. Very little of what is here is firm
     and reactions are expected.
```

Request for Comments – RFC 4980

```
Network Working Group                                              C. Ng
Request for Comments: 4980                    Panasonic Singapore Labs
Category: Informational                                         T. Ernst
                                                                   INRIA
                                                                 E. Paik
                                                                      KT
                                                              M. Bagnulo
                                                                    UC3M
                                                            October 2007

          Analysis of Multihoming in Network Mobility Support

Status of This Memo

   This memo provides information for the Internet community.  It does
   not specify an Internet standard of any kind.  Distribution of this
   memo is unlimited.

Abstract

   This document is an analysis of multihoming in the context of network
   mobility (NEMO) in IPv6.  As there are many situations in which
   mobile networks may be multihomed, a taxonomy is proposed to classify
   the possible configurations.  The possible deployment scenarios of
   multihomed mobile networks are described together with the associated
   issues when network mobility is supported by RFC 3963 (NEMO Basic
   Support).  Recommendations are offered on how to address these
   issues.
```

Summary

- The ARPANET started out very informally.
  - People dreamt of accessing mainframe computers across the continent.
  - Telephone networks provided dial-in access, but no internetworking.
- Meanwhile, the Internet has become highly regulated.
  - Packet switching and TCP/IP have become ubiquitous.
  - Telephony is just one application of the Internet.
- The core protocols are frozen due to the commercial interest.
  - IPv4 is ubiquitous albeit IPv6 has been defined and implemented for about a decade.
- But, the Internet is flexible enough to allow new approaches outside the core protocols.
  - Example: Peer-to-peer networking

Overview of This Lecture

- The Internet Protocols
  - Forwarding & Routing (→IP)
  - Interactive access &
    Bulk data transfer (→TCP)
  - Scalability & Hierarchies
  - Naming (DNS & Mail)
  - The World Wide Web
  - Multimedia & QoS
  - Mobility & Wireless Access
  - Commercialization & Security
  - Overlay networks & Peer-to-peer
- What is good with these protocols?
- What went wrong?
- What is going on in Internet research?

Learning the facts

Not in this lecture

Understanding the design rationale

Questions

- Explain the differences between circuit and packet switching?
- What does ‚best effort network‘ mean?
- Why is circuit switching so expensive?
- Why is packet switching so cheap?
- Is best effort packet switching able to carry voice communication?
- What happens when someone cuts an important fiber?
- Who owns the Internet?
- You’ve invented a new protocol. What do you do?
- Is triple play (=voice, Internet, and TV) more Internet or more telco?
- What’s more fair, time or volume tariffs? Or flat rates?
- When will the exponential growth of the Internet end?

Design Goals of this Lecture

- I want you …
  - To listen, think, and ask questions, and
  - To read additional material.
- For the exam …
  - You do not need to learn everything by heart,
  - But even if you did, you would not be excellent.
  - You need to understand!

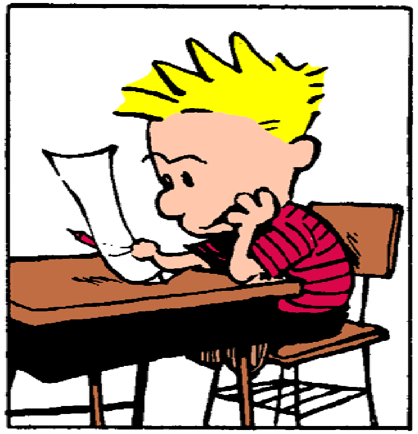

- What to you expect me to do or not to do?

Questions?

Thomas Fuhrmann

CS VIII – Network Architectures
Technical University Munich, Germany

IBDS System Architecture
University of Karlsruhe, Germany

fuhrmann@net.in.tum.de