dr. heap

Netherlands Defence Academy NLDA

Computer Networks

New Technologies and the Armed Forces (NTK)

Computer Networks

Hans-Dieter A. Hiep

9 November, 2026

CC BY 4.0

The content reflects solely the findings of the presenter and should not be interpreted as representing the views of the Ministry of Defence, the Netherlands Defence Academy (NLDA), or any other governmental institution.

Overview

Part I. First hour

1. Theory & Practice
2. The physical layer
3. The data plane

Part II. Second hour

4. The control plane
5. Learning from 9/11
6. Regaining the cyber high ground

Overview

Part I. First hour

1. **Theory & Practice**
2. The physical layer
3. The data plane

Part II. Second hour

4. The control plane
5. Learning from 9/11
6. Regaining the cyber high ground

Overall picture

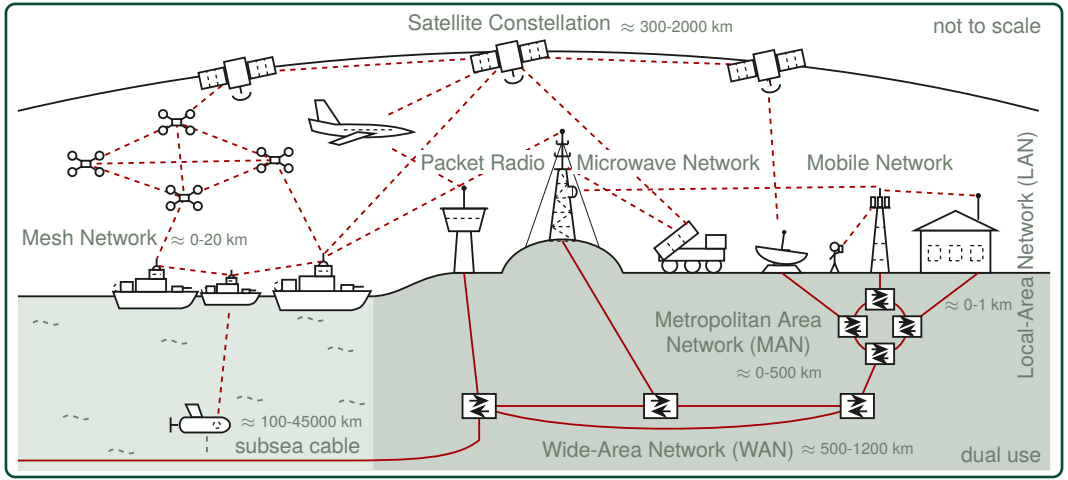

SAGE

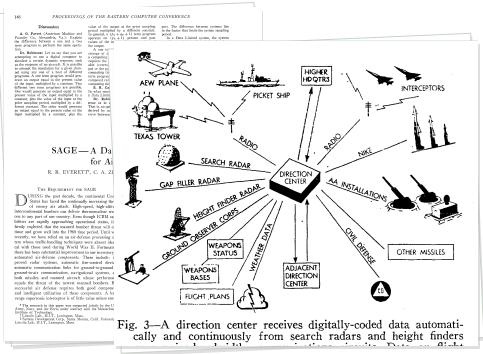

Fig. 3—A direction center receives digitally-coded data automatically and continuously from search radars and height finders

SAGE—A Data-Processing System for Air Defense Everett et al. (1957)

sensors actuators

computers networks

operators

- Semi-Automatic Ground Environment Air Defence System
- Response to Soviet Union's atomic bomb tests (1949)
- Early computer network (1951-1980s)

- Air-defence functions: air surveillance, identification, weapons control, etc.
- Per sector, complete situational picture: 100 personnel behind consoles
- Connected by **telephone**
- Reliability: independent, redundant "stand-by" computers

MIT Lincoln Laboratory

SAGE

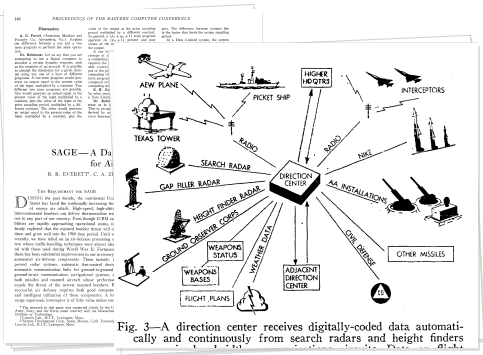

Fig. 3—A direction center receives digitally-coded data automatically and continuously from search radars and height finders

SAGE—A Data-Processing System for Air Defense Everett et al. (1957)

sensors actuators

computers networks

operators

- Semi-Automatic Ground Environment Air Defence System
- Response to Soviet Union's atomic bomb tests (1949)
- Early computer network (1951-1980s)

- Air-defence functions: air surveillance, identification, weapons control, etc.
- Per sector, complete situational picture: 100 personnel behind consoles
- Connected by **telephone**
- Reliability: independent, redundant "stand-by" computers

MIT Lincoln Laboratory

SAGE

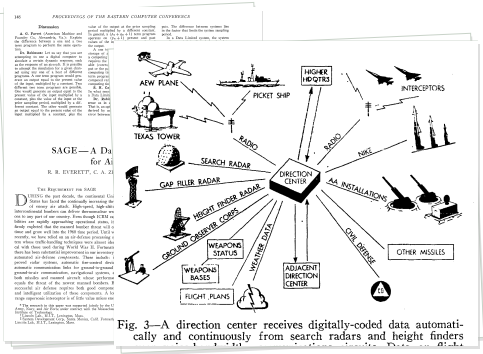

Fig. 3—A direction center receives digitally-coded data automatically and continuously from search radars and height finders

SAGE—A Data-Processing System for Air Defense Everett et al. (1957)

sensors actuators

computers networks

operators

- Semi-Automatic Ground Environment Air Defence System
- Response to Soviet Union's atomic bomb tests (1949)
- Early computer network (1951-1980s)

- Air-defence functions: air surveillance, identification, weapons control, etc.
- Per sector, complete situational picture: 100 personnel behind consoles
- Connected by **telephone**
- Reliability: independent, redundant "stand-by" computers

MIT Lincoln Laboratory

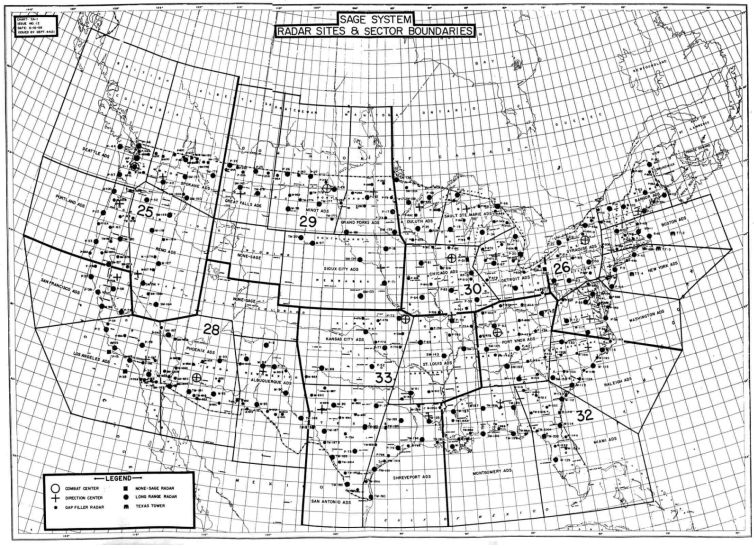

(Source: The Online Air Defense Radar Museum)

Simplicity is prerequisite for reliability.

— Edsger W. Dijkstra

Maxims:

- keep problems small enough for **one mind** to hold
- **trust but verify**: ask for proof, no blind trust
- **angelic/demonic non-determinism**:
  hope for the best, be prepared for the worst

Simplicity is prerequisite for reliability.

— Edsger W. Dijkstra

Maxims:

- keep problems small enough for **one mind** to hold
- **trust but verify**: ask for proof, no blind trust
- **angelic/demonic non-determinism**:
  hope for the best, be prepared for the worst

Simplicity is prerequisite for reliability.

— Edsger W. Dijkstra

Maxims:

- keep problems small enough for **one mind** to hold
- **trust but verify**: ask for proof, no blind trust
- **angelic/demonic non-determinism**:
  hope for the best, be prepared for the worst

Simplicity is prerequisite for reliability.

— Edsger W. Dijkstra

Maxims:

- keep problems small enough for **one mind** to hold
- **trust but verify**: ask for proof, no blind trust
- **angelic/demonic non-determinism**:
  hope for the best, be prepared for the worst

Simple idea: theory of communication

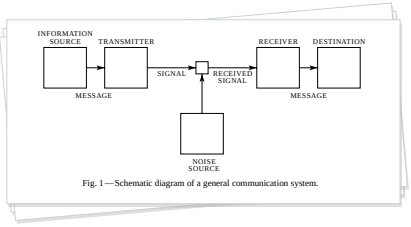

Fig. 1 — Schematic diagram of a general communication system.

A mathematical theory of communication Shannon (1948)

- Channel: source $\longrightarrow$ destination
- Classical bit: binary digit
- 00110101... (sequence of bits)
- Measure of information
- Function: identical input/output
- Capacity: **# bits per second**
- Noise decreases capacity

Simple idea: theory of communication

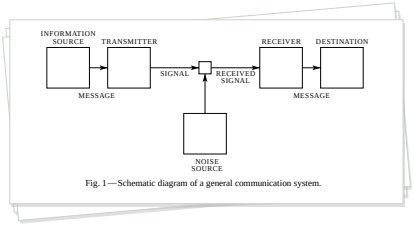

Fig. 1 — Schematic diagram of a general communication system.

A mathematical theory of communication Shannon (1948)

- Channel: source $\longrightarrow$ destination
- Classical bit: binary digit
- 00110101... (sequence of bits)
- Measure of information
- Function: identical input/output
- Capacity: **# bits per second**
- Noise decreases capacity

Simple idea: theory of communication

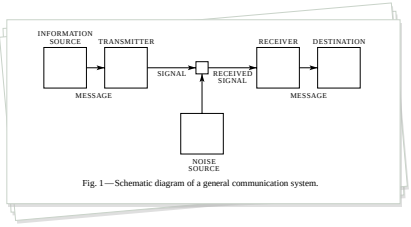

Fig. 1 — Schematic diagram of a general communication system.

A mathematical theory of communication Shannon (1948)

- Channel: source $\longrightarrow$ destination
- Classical bit: binary digit
- 00110101... (sequence of bits)
- Measure of information
- Function: identical input/output
- Capacity: **# bits per second**
- Noise decreases capacity

Simple idea: theory of communication

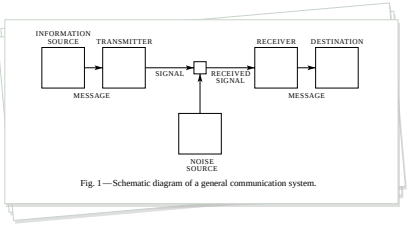

Fig. 1 — Schematic diagram of a general communication system.

A mathematical theory of communication Shannon (1948)

- Channel: source $\longrightarrow$ destination
- Classical bit: binary digit
- 00110101... (sequence of bits)
- Measure of information
- Function: identical input/output
- Capacity: **# bits per second**
- Noise decreases capacity

Simple idea: theory of communication

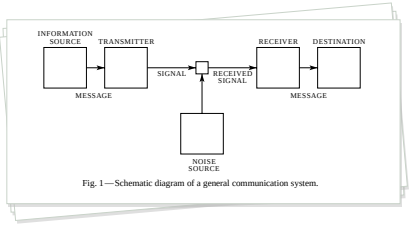

Fig. 1 — Schematic diagram of a general communication system.

A mathematical theory of communication Shannon (1948)

- Channel: source $\longrightarrow$ destination
- Classical bit: binary digit
- 00110101... (sequence of bits)
- Measure of information
- Function: identical input/output
- Capacity: **# bits per second**
- Noise decreases capacity

An intuition of bits

Bit = **b**inary dig**it**

1 bit answers one **yes/no** question:

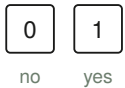

A **path**: each bit branches in two

n bit is a **counter**, one place a bit:

$8^0 + 4^1 + 2^1 + 1^0 = 6$

An intuition of bits

Bit = **b**inary dig**it**

1 bit answers one **yes/no** question:

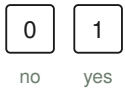

A **path**: each bit branches in two

n bit is a **counter**, one place a bit:

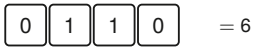

$= 6$

$$8^0 + 4^1 + 2^1 + 1^0 = 6$$

An intuition of bits

Bit = **b**inary dig**it**

1 bit answers one **yes/no** question:

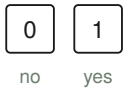

A **path**: each bit branches in two

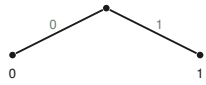

n bit is a **counter**, one place a bit:



$= 6$

$$8^0 + 4^1 + 2^1 + 1^0 = 6$$

An intuition of bits

Bit = **b**inary dig**it**

1 bit answers one **yes/no** question:

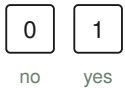

A **path**: each bit branches in two

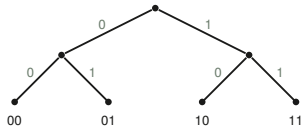

n bit is a **counter**, one place a bit:



$= 6$

$$8^0 + 4^1 + 2^1 + 1^0 = 6$$

An intuition of bits

Bit = **b**inary dig**it**

1 bit answers one **yes/no** question:

0	1
no	yes

n bit is a **counter**, one place a bit:

0	1	1	0	$= 6$

$$8^0 + 4^1 + 2^1 + 1^0 = 6$$

A **path**: each bit branches in two

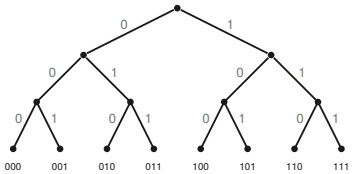

An intuition of bits

Bit = **b**inary dig**it**

1 bit answers one **yes/no** question:

0	1
no	yes

n bit is a **counter**, one place a bit:

0	1	1	0

$= 6$

$$8^0 + 4^1 + 2^1 + 1^0 = 6$$

A **path**: each bit branches in two

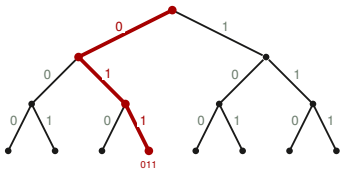

An intuition of trits

Trit = **tr**inary dig**it**

1 trit answers **three-way**:

$-$	0	$+$
negative	neutral	positive
-1	0	$+1$

A **path**: each trit branches in three

- Setun: early electronic Soviet computer (1958)
- **Ternary computers: the Setun and the Setun 70** Brusentsov and Ramil Alvarez (2011)
- Not continued due to ‘weird politics’ (the machine was too cheap?)

An intuition of trits

Trit = **tr**inary dig**it**

1 trit answers **three-way**:

−	0	+
negative	neutral	positive
−1	0	+1

A **path**: each trit branches in three

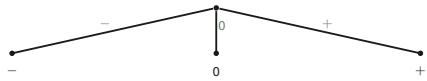

- Setun: early electronic Soviet computer (1958)
- **Ternary computers: the Setun and the Setun 70** Brusentsov and Ramil Alvarez (2011)
- Not continued due to ‘weird politics’ (the machine was too cheap?)

An intuition of trits

Trit = **tr**inary dig**it**

1 trit answers **three-way**:

−	0	+
negative	neutral	positive
−1	0	+1

A **path**: each trit branches in three

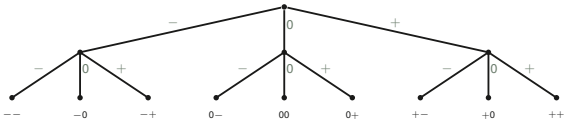

- Setun: early electronic Soviet computer (1958)
- **Ternary computers: the Setun and the Setun 70** Brusentsov and Ramil Alvarez (2011)
- Not continued due to ‘weird politics’ (the machine was too cheap?)

An intuition of trits

Trit = **tr**inary dig**it**

1 trit answers **three-way**:

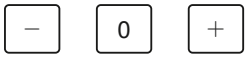

A **path**: each trit branches in three

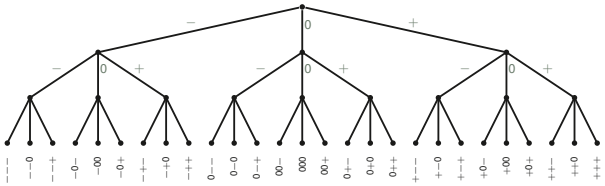

- Setun: early electronic Soviet computer (1958)
- **Ternary computers: the Setun and the Setun 70** Brusentsov and Ramil Alvarez (2011)
- Not continued due to ‘weird politics’ (the machine was too cheap?)

An intuition of trits

Trit = **tr**inary dig**it**

1 trit answers **three-way**:

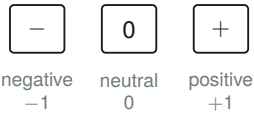

A **path**: each trit branches in three

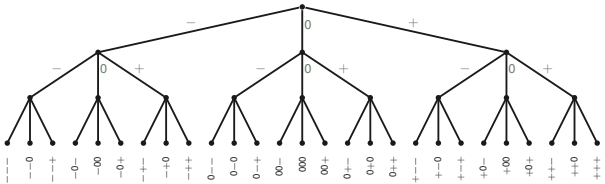

- Setun: early electronic Soviet computer (1958)
- **Ternary computers: the Setun and the Setun 70** Brusentsov and Ramil Alvarez (2011)
- Not continued due to ‘weird politics’ (the machine was too cheap?)

Number of possibilities

Bit sequence of length n:

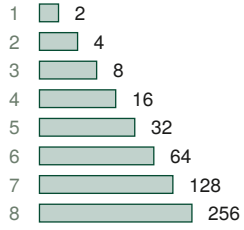

2^n

Trit sequence of length n:

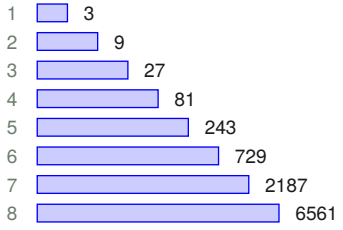

3^n

Exponential growth: each further bit doubles the possibilities, each further trit triples them.

Theory of networks

Network

A **network** is a graph $G = (V, E)$ with **capacity** labels.
A capacity $c(v)$ on a vertex $v \in V$ makes it a **node**.
A capacity $c(e)$ on an edge $e \in E$ makes it a **link**.

How much can flow from s to t at once?

Theory of networks

Network

A **network** is a graph $G = (V, E)$ with **capacity** labels.
A capacity $c(v)$ on a vertex $v \in V$ makes it a **node**.
A capacity $c(e)$ on an edge $e \in E$ makes it a **link**.

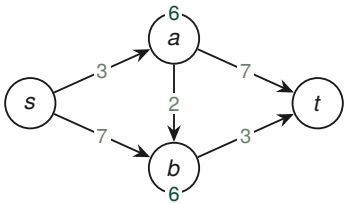

How much can flow from s to t at once?

Theory of networks

Network

A **network** is a graph $G = (V, E)$ with **capacity** labels.
A capacity $c(v)$ on a vertex $v \in V$ makes it a **node**.
A capacity $c(e)$ on an edge $e \in E$ makes it a **link**.

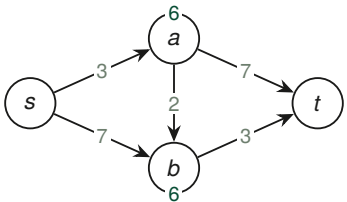

How much can flow from s to t at once?

Where is the cheapest cut?

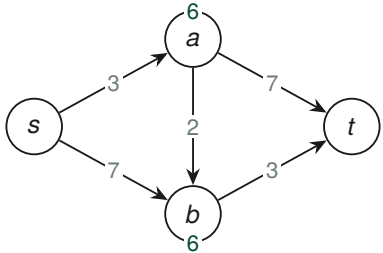

- How to separate s from t?
- Everything out of s: $3 + 7 = 10$.
- Everything into t: $7 + 3 = 10$.
- $s \to a$ **and** $b \to t$**:** $3 + 3 = 6$

Where is the cheapest cut?

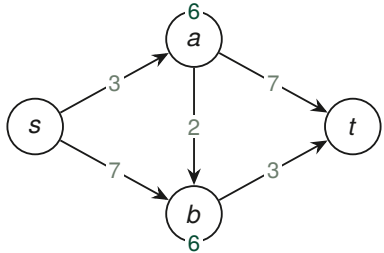

- How to separate s from t?
- Everything out of s: $3 + 7 = 10$.
- Everything into t: $7 + 3 = 10$.
- $s \to a$ **and** $b \to t$**:** $3 + 3 = 6$

Where is the cheapest cut?

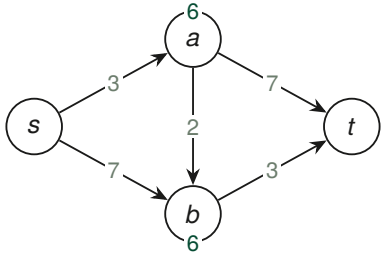

- How to separate s from t?
- Everything out of s: $3 + 7 = 10$.
- Everything into t: $7 + 3 = 10$.
- $s \to a$ **and** $b \to t$**:** $3 + 3 = 6$

Where is the cheapest cut?

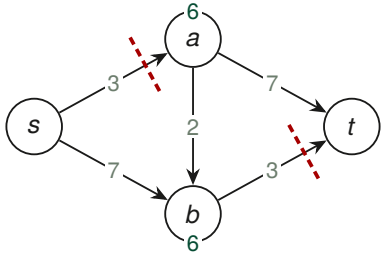

- How to separate s from t?
- Everything out of s: $3 + 7 = 10$.
- Everything into t: $7 + 3 = 10$.
- $s \to a$ **and** $b \to t$**:** $3 + 3 = 6$

Simple idea: distributed communication

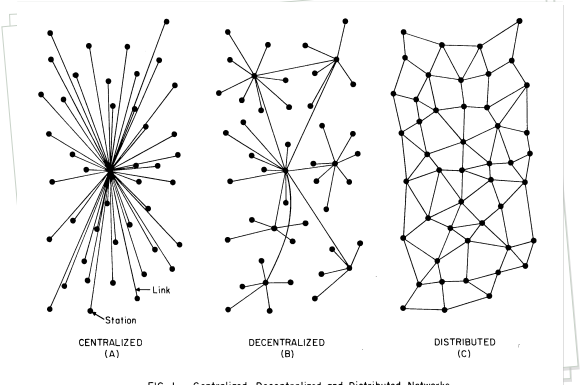

On distributed communications: I. Introduction to distributed communications networks
Baran (1964)

In the news: data centres under attack



Georgia Tech News Center Murphy (2026)



U.S. Department of Justice press release (2021)

- Russian **Yandex** data centres, Ukrainian **ISPs**. Who's next?

In the news: data centres under attack



Georgia Tech News Center Murphy (2026)



U.S. Department of Justice press release (2021)

- Russian **Yandex** data centres, Ukrainian **ISPs**. Who's next?

Design philosophy of Internet

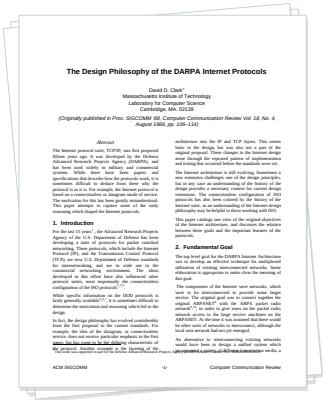

The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA. 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial networks. This paper discusses the philosophy and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has frequently been misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years*, the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[1,2]

While specific information on the DOD protocols is fairly generally available[3,4], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standard. For example, the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but comes to be the dominant characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated pattern of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols[1,2] has also been influenced by this history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The first two networks to be connected were the original ARPANET[5] with the ARPA packet radio network[6,7], to order to give users on the packet radio network access to the large computers on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet erupted.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet

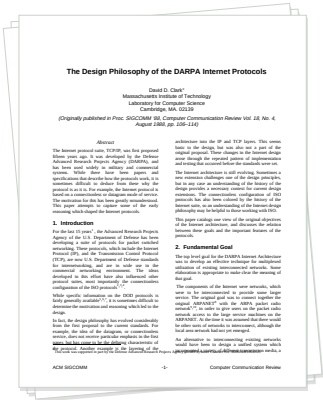

The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial systems. While there have been papers and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has been greatly misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years[1], the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet-switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[2,3,4]

While specific information on the DOD protocols is fairly generally available[5,6,7], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example, the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but has come to be the defining characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the vigorous process of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols has also been colored by the history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The original goal was to be able to use the original ARPANET[8] with the ARPA packet radio network[9,10], in order to give users on the packet radio network access to the large service machines on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet emerged.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet

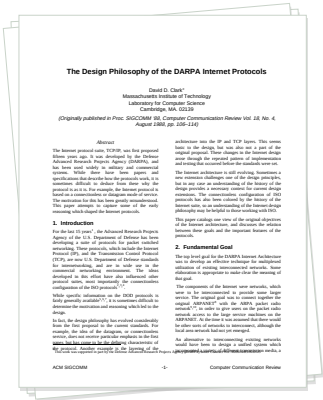

The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA. 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial networks. This paper discusses the philosophy and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has frequently been misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years*, the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[1,2]

While specific information on the DOD protocols is fairly generally available[3,4], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example: the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but comes to be the dominant characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated pattern of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols[1,2] has also been influenced by this history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The first two networks to be connected were the original ARPANET* with the ARPA packet radio network[5,6], in order to give users on the packet radio network access to the large computers on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet erupted.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet

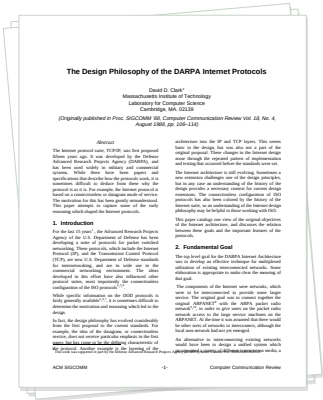

The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial systems. While there have been papers and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has been greatly misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years[1], the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[2,3,4]

While specific information on the DOD protocols is fairly generally available[5,6,7], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example, the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but has come to be the defining characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated process of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols has also been colored by the history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The original goal was to be able to use the single ARPANET[8] with the ARPA packet radio network[9,10], in order to give users on the packet radio network access to the large service machines on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet emerged.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet



The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial systems. While there have been papers and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has been greatly misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years[1], the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols[2,3].

While specific information on the DOD protocols is fairly generally available[4,5,6], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example, the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but has come to be the defining characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated process of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols has also been colored by the history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The original goal was to be able to use the single ARPANET[7] with the ARPA packet radio network[8,9], in order to give users on the packet radio network access to the large service machines on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet emerged.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet



The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial systems. While there have been papers and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has been greatly misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years[1], the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[2,3,4]

While specific information on the DOD protocols is fairly generally available[5,6,7], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example, the idea of the datagram, or connectionless service, does not receive particular emphasis in the first paper, but has come to be the defining characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated process of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols has also been colored by the history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The original goal was to be able to use the single ARPANET[8] with the ARPA packet radio network[9,10], in order to give users on the packet radio network access to the large service machines on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet emerged.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet

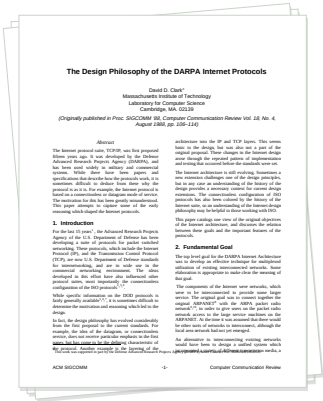

The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA. 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial networks. This paper discusses the philosophy and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has not been documented well. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years*, the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment as well. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[1,2]

While specific information on the DOD protocols is fairly generally available[3,4], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example: the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but comes to be the defining characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated pattern of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols[1,2] has also been influenced by this history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The first two networks to be interconnected were the original ARPANET and the ARPA packet radio network[5,6], in order to give users on the packet radio network access to the large computers on the ARPANET. At the time it was assumed that there would be other sorts of networks to interconnect, although the local area network had not yet emerged.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

Design philosophy of Internet



The Design Philosophy of the DARPA Internet Protocols

David D. Clark*
Massachusetts Institute of Technology
Laboratory for Computer Science
Cambridge, MA. 02139

(Originally published in Proc. SIGCOMM '88, Computer Communication Review Vol. 18, No. 4, August 1988, pp. 106–114)

Abstract

The Internet protocol suite, TCP/IP, was first proposed fifteen years ago. It was developed by the Defense Advanced Research Projects Agency (DARPA), and has been used widely in military and commercial networks. This paper discusses the philosophy and specifications that describe how the protocols work, it is sometimes difficult to deduce from these why the protocol is as it is. For example, the Internet protocol is based on a connectionless or datagram mode of service. The motivation for this has frequently been misunderstood. This paper attempts to capture some of the early reasoning which shaped the Internet protocols.

1. Introduction

For the last 15 years*, the Advanced Research Projects Agency of the U.S. Department of Defense has been developing a suite of protocols for packet switched networking. These protocols, which include the Internet Protocol (IP), and the Transmission Control Protocol (TCP), are now U.S. Department of Defense standards for internetworking, and are in wide use in the commercial networking environment. The ideas developed in this effort have also influenced other protocol suites, most importantly the connectionless configuration of the ISO protocols.[1,2]

While specific information on the DOD protocols is fairly generally available[3,4], it is sometimes difficult to determine the motivation and reasoning which led to the design.

In fact, the design philosophy has evolved considerably from the first proposal to the current standards. For example: the idea of the datagram, or connectionless service, does not receive particular emphasis in the first papers, but comes to be the dominant characteristic of the protocol. Another example is the layering of the architecture into the IP and TCP layers. This seems basic to the design, but was also not a part of the original proposal. These changes in the Internet design arose through the repeated pattern of implementation and testing that occurred before the standards were set.

The Internet architecture is still evolving. Sometimes a new extension challenges one of the design principles, but in any case an understanding of the history of the design provides a necessary context for current design extensions. The connectionless configuration of ISO protocols[5,6,7] has also been influenced by this history of the Internet suite, so an understanding of the Internet design philosophy may be helpful to those working with ISO.

This paper catalogs one view of the original objectives of the Internet architecture, and discusses the relation between these goals and the important features of the protocols.

2. Fundamental Goal

The top level goal for the DARPA Internet Architecture was to develop an effective technique for multiplexed utilization of existing interconnected networks. Some elaboration is appropriate to make clear the meaning of that goal.

The components of the Internet were networks, which were to be interconnected to provide some larger service. The first two networks to be connected were the original ARPANET[8] with the ARPA packet radio network[9,10], in order to give users on the packet radio network access to the large resources on the ARPANET. As the time it was assumed that there would be other sorts of networks to interconnect, although the local network technology had not yet emphasized.

An alternative to interconnecting existing networks would have been to design a unified system which incorporated a variety of different transmission media, a

ACM SIGCOMM -1- Computer Communication Review

The design philosophy of the DARPA internet protocols
Clark (1988)

- Inter-net: ‘Inter network’
- Network: (administrative) boundary of control
- Gateway: border between networks (IMP)

1. Continue despite loss of *networks* or *gateways* (**survivability**)
2. Support multiple communications services
3. Variety of networks (**interoperability**)
4. Distributed management of its resources
5. Cost effective
6. Permit attachment with low effort (**extensibility**)
7. Resources used must be accountable

The five layers of the TCP/IP model

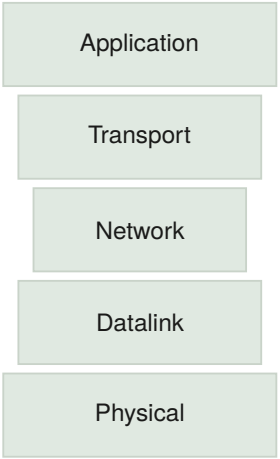

The five layers of the TCP/IP model

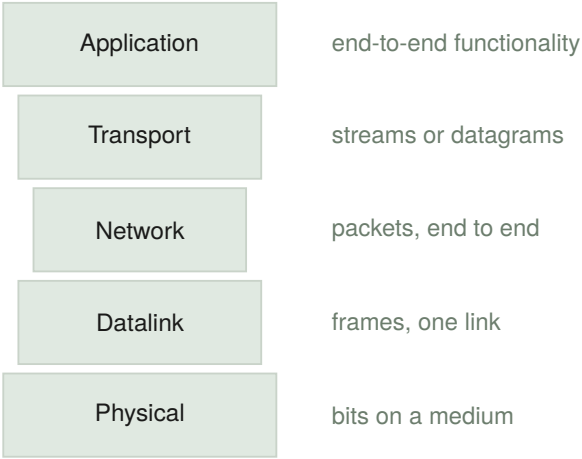

The five layers of the TCP/IP model

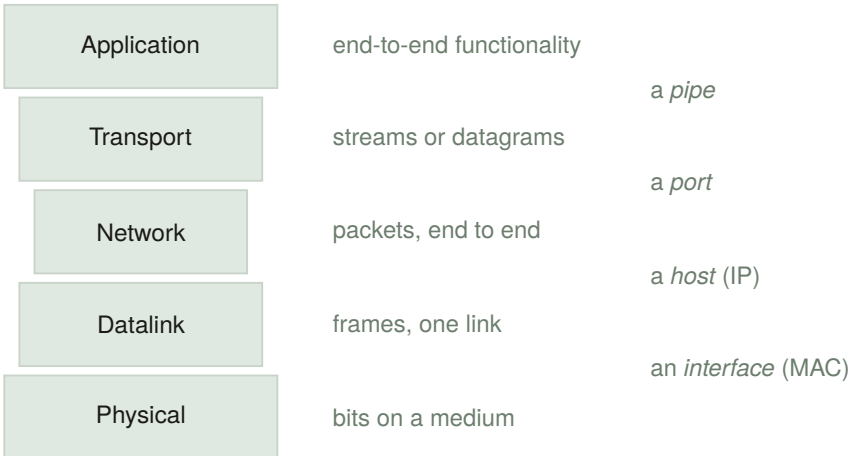

Overview

Part I. First hour

1. Theory & Practice
2. **The physical layer**
3. The data plane

Part II. Second hour

4. The control plane
5. Learning from 9/11
6. Regaining the cyber high ground

Where we are

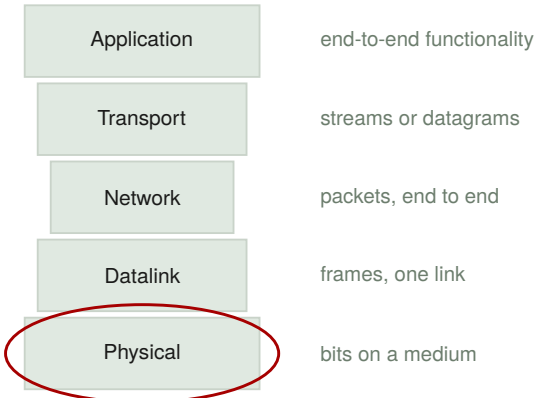
Application
end-to-end functionality
Transport
streams or datagrams
Network
packets, end to end
Datalink
frames, one link
Physical
bits on a medium

The electromagnetic field

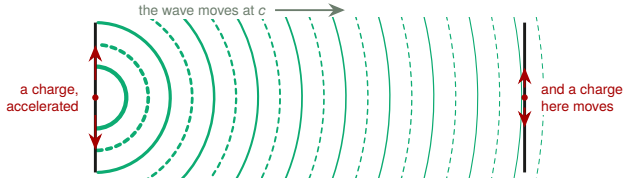

- The field is everywhere, all the time.
- A charge that **accelerates** creates waves in the field.
- Frequency of waves (kHz, MHz, GHz, THz).

- **Disturbances in the field**: in vacuum travel at c, speed of light
- **Larmor formula**: an accelerating charge radiates

$$P = \frac{q^2 a^2}{6\pi\varepsilon_0 c^3}$$

power P in **watt** (W)

3Blue1Brown, *How wiggling charges give rise to light*

The electromagnetic field

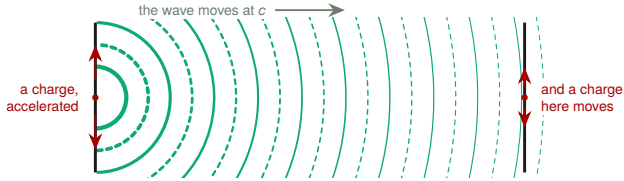

- The field is everywhere, all the time.
- A charge that **accelerates** creates waves in the field.
- Frequency of waves (kHz, MHz, GHz, THz).
- **Disturbances in the field**: in vacuum travel at c, speed of light
- **Larmor formula**: an accelerating charge radiates

$$P = \frac{q^2 a^2}{6\pi\varepsilon_0 c^3}$$

power P in **watt** (W)

3Blue1Brown, *How wiggling charges give rise to light*

The electromagnetic field

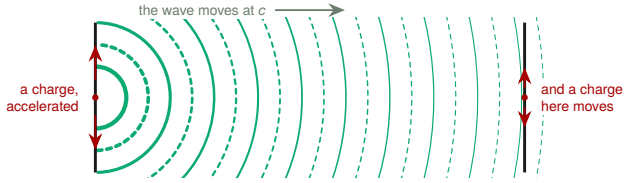

- The field is everywhere, all the time.
- A charge that **accelerates** creates waves in the field.
- Frequency of waves (kHz, MHz, GHz, THz).

- **Disturbances in the field**: in vacuum travel at c, speed of light
- **Larmor formula**: an accelerating charge radiates

$$P = \frac{q^2 a^2}{6\pi\varepsilon_0 c^3}$$

power P in **watt** (W)

3Blue1Brown, *How wiggling charges give rise to light*

The electromagnetic field

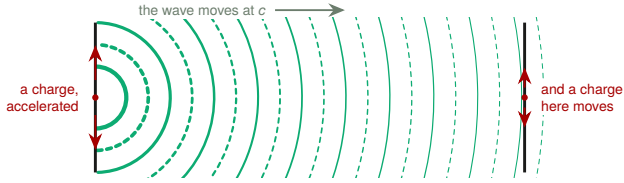

- The field is everywhere, all the time.
- A charge that **accelerates** creates waves in the field.
- Frequency of waves (kHz, MHz, GHz, THz).

- **Disturbances in the field**: in vacuum travel at c, speed of light
- **Larmor formula**: an accelerating charge radiates

$$P = \frac{q^2 a^2}{6\pi\varepsilon_0 c^3}$$

power P in **watt** (W)

3Blue1Brown, *How wiggling charges give rise to light*

The electromagnetic field

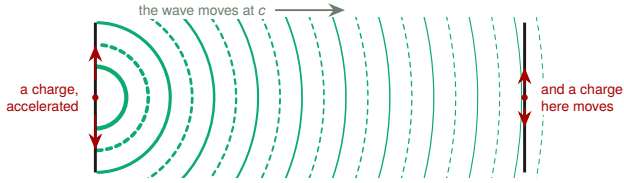

- The field is everywhere, all the time.
- A charge that **accelerates** creates waves in the field.
- Frequency of waves (kHz, MHz, GHz, THz).

- **Disturbances in the field**: in vacuum travel at c, speed of light
- **Larmor formula**: an accelerating charge radiates

$$P = \frac{q^2 a^2}{6\pi\varepsilon_0 c^3}$$

power P in **watt** (W)

3Blue1Brown, *How wiggling charges give rise to light*

One shared medium

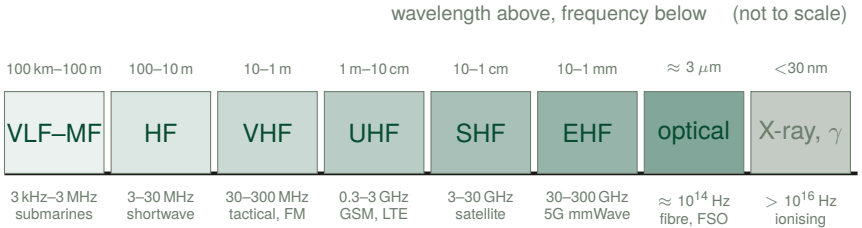

- **Higher frequency**: more bandwidth, shorter reach, stopped by more things.
- **Path loss**: signal power reduction over distance.
- **Band assignment**: by a regulator to coordinate usage.

International Telecommunication Union (ITU)

One shared medium

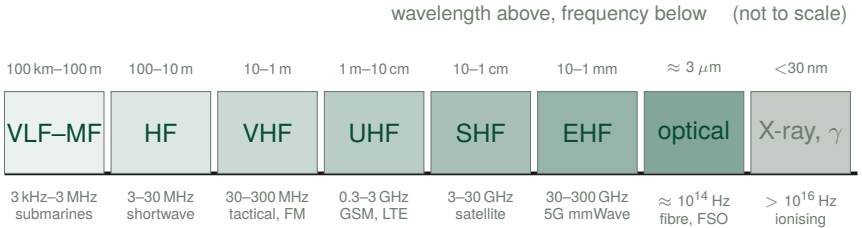

- **Higher frequency**: more bandwidth, shorter reach, stopped by more things.
- **Path loss**: signal power reduction over distance.
- **Band assignment**: by a regulator to coordinate usage.

International Telecommunication Union (ITU)

One shared medium

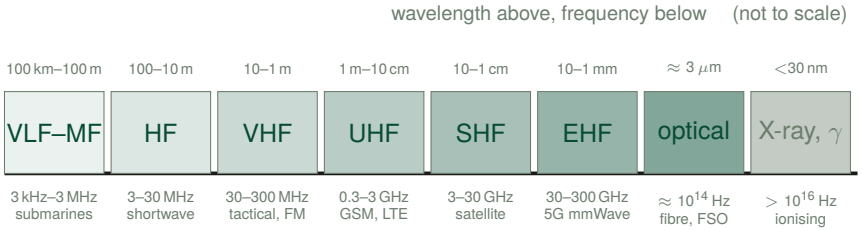

- **Higher frequency**: more bandwidth, shorter reach, stopped by more things.
- **Path loss**: signal power reduction over distance.
- **Band assignment**: by a regulator to coordinate usage.

International Telecommunication Union (ITU)

Bandwidth is not capacity

Bandwidth B: a width of spectrum, in Hz.

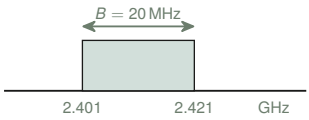

Capacity C: bits per second the channel carries *without error*.

Shannon–Hartley theorem

$$C = B \log_2\left(1 + \frac{S}{N}\right)$$

- One Wi-Fi channel: $B = 20$ MHz, up to about 140 Mbit/s.
- Doubling S/N gives *one* bit per hertz extra.
- Jamming only has to raise N.

Bandwidth is not capacity

Bandwidth B: a width of spectrum, in Hz.

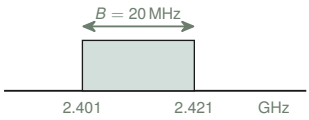

Capacity C: bits per second the channel carries *without error*.

Shannon–Hartley theorem

$$C = B \log_2\left(1 + \frac{S}{N}\right)$$

- One Wi-Fi channel: $B = 20$ MHz, up to about 140 Mbit/s.
- Doubling S/N gives *one* bit per hertz extra.
- Jamming only has to raise N.

Bandwidth is not capacity

Bandwidth B: a width of spectrum, in Hz.

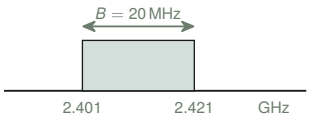

Capacity C: bits per second the channel carries *without error*.

Shannon–Hartley theorem

$$C = B \log_2\left(1 + \frac{S}{N}\right)$$

- One Wi-Fi channel: $B = 20$ MHz, up to about 140 Mbit/s.
- Doubling S/N gives *one* bit per hertz extra.
- Jamming only has to raise N.

Bandwidth is not capacity

Bandwidth B: a width of spectrum, in Hz.

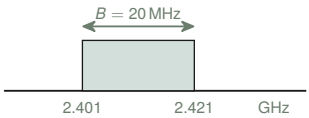

Capacity C: bits per second the channel carries *without error*.

Shannon–Hartley theorem

$$C = B \log_2\left(1 + \frac{S}{N}\right)$$

- One Wi-Fi channel: $B = 20$ MHz, up to about 140 Mbit/s.
- Doubling S/N gives *one* bit per hertz extra.
- Jamming only has to raise N.

Bandwidth is not capacity

Bandwidth B: a width of spectrum, in Hz.

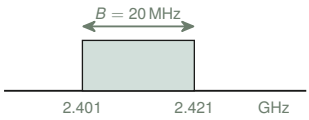

Capacity C: bits per second the channel carries *without error*.

Shannon–Hartley theorem

$$C = B \log_2\left(1 + \frac{S}{N}\right)$$

- One Wi-Fi channel: $B = 20$ MHz, up to about 140 Mbit/s.
- Doubling S/N gives *one* bit per hertz extra.
- Jamming only has to raise N.

Bandwidth is not capacity

Bandwidth B: a width of spectrum, in Hz.

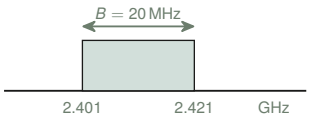

Capacity C: bits per second the channel carries *without error*.

Shannon–Hartley theorem

$$C = B \log_2\left(1 + \frac{S}{N}\right)$$

- One Wi-Fi channel: $B = 20$ MHz, up to about 140 Mbit/s.
- Doubling S/N gives *one* bit per hertz extra.
- Jamming only has to raise N.

An intuition for S/N

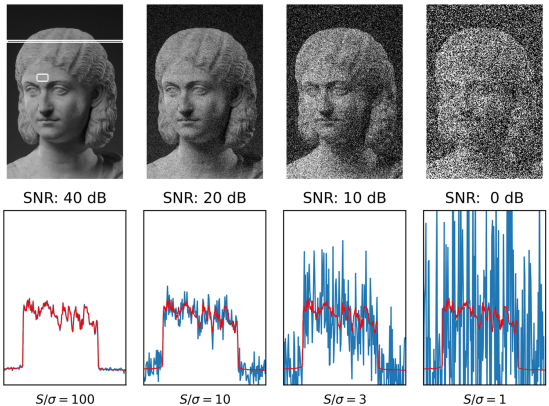

source

- **Signal-to-noise ratio**: the signal's power divided by the noise's. Two powers, so the ratio carries **no unit**: it is a number.
- It runs over many orders of magnitude, so it is written in **decibels**: $10 \log_{10}(S/N)$.

S/N	
1	0 dB
10	10 dB
100	20 dB
10 000	40 dB

The S/σ under each plot is an *amplitude* ratio; square it for the power ratio.

An intuition for S/N

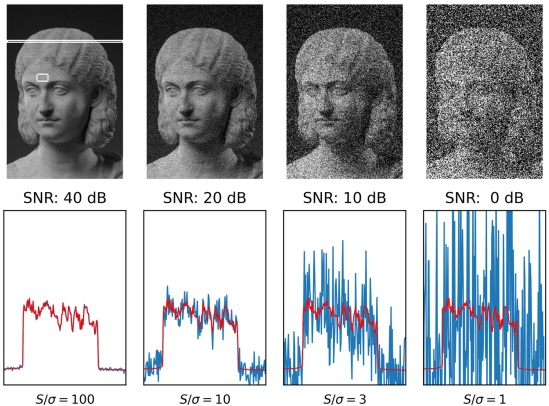

source

- **Signal-to-noise ratio**: the signal's power divided by the noise's. Two powers, so the ratio carries **no unit**: it is a number.
- It runs over many orders of magnitude, so it is written in **decibels**: $10 \log_{10}(S/N)$.

S/N	
1	0 dB
10	10 dB
100	20 dB
10 000	40 dB

The S/σ under each plot is an *amplitude* ratio; square it for the power ratio.

An intuition for S/N

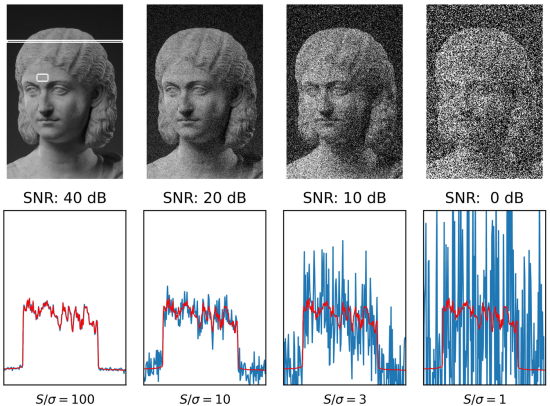

source

- **Signal-to-noise ratio**: the signal's power divided by the noise's. Two powers, so the ratio carries **no unit**: it is a number.
- It runs over many orders of magnitude, so it is written in **decibels**: $10 \log_{10}(S/N)$.

S/N	
1	0 dB
10	10 dB
100	20 dB
10 000	40 dB

The S/σ under each plot is an *amplitude* ratio; square it for the power ratio.

Properties of the medium

- **Attenuation**: signal gets weaker
  fibre: about 0.2 dB/km at 1550 nm;
  copper: far worse
- **Dispersion**: signal morphing
  speed of traversal depends on frequency

One speed in, several through the glass. (source)

- **Interference**: multiple transmitters
  partially overlapping channels
- **Delay**: signal propagation time
  light in optical fibre: about 200 Mm/s
- **Impedance**: coupling loss, from medium to medium
  reflections/absorptions cause unwanted noise

Properties of the medium

- **Attenuation**: signal gets weaker
  fibre: about 0.2 dB/km at 1550 nm;
  copper: far worse
- **Dispersion**: signal morphing
  speed of traversal depends on frequency

One speed in, several through the glass. (source)

- **Interference**: multiple transmitters
  partially overlapping channels
- **Delay**: signal propagation time
  light in optical fibre: about 200 Mm/s
- **Impedance**: coupling loss, from medium to medium
  reflections/absorptions cause unwanted noise

Properties of the medium

- **Attenuation**: signal gets weaker
  fibre: about 0.2 dB/km at 1550 nm;
  copper: far worse
- **Dispersion**: signal morphing
  speed of traversal depends on frequency

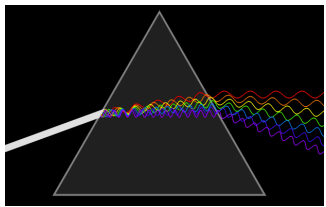

One speed in, several through the glass. (source)

- **Interference**: multiple transmitters
  partially overlapping channels
- **Delay**: signal propagation time
  light in optical fibre: about 200 Mm/s
- **Impedance**: coupling loss, from medium to medium
  reflections/absorptions cause unwanted noise

Properties of the medium

- **Attenuation**: signal gets weaker
  fibre: about 0.2 dB/km at 1550 nm;
  copper: far worse
- **Dispersion**: signal morphing
  speed of traversal depends on frequency

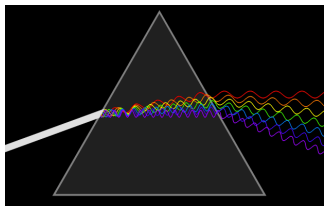

One speed in, several through the glass. (source)

- **Interference**: multiple transmitters
  partially overlapping channels
- **Delay**: signal propagation time
  light in optical fibre: about 200 Mm/s
- **Impedance**: coupling loss, from medium to medium
  reflections/absorptions cause unwanted noise

Properties of the medium

- **Attenuation**: signal gets weaker
  fibre: about 0.2 dB/km at 1550 nm;
  copper: far worse
- **Dispersion**: signal morphing
  speed of traversal depends on frequency

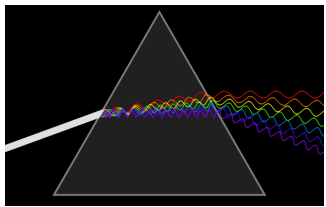

One speed in, several through the glass. (source)

- **Interference**: multiple transmitters
  partially overlapping channels
- **Delay**: signal propagation time
  light in optical fibre: about 200 Mm/s
- **Impedance**: coupling loss, from medium to medium
  reflections/absorptions cause unwanted noise

Multiplexer and demultiplexer

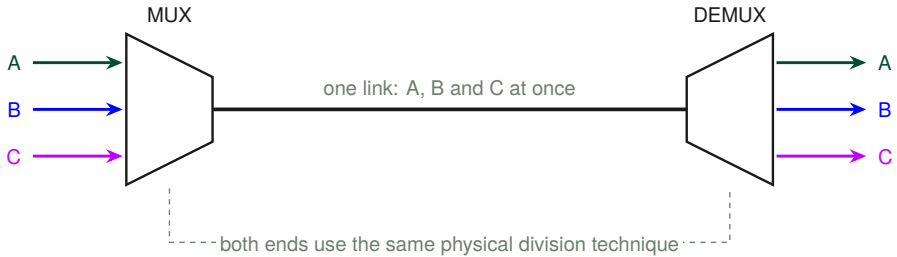

- **Multiplexer**: several channels onto one channel.
- **Demultiplexer**: split one channel into several channels.
- The inner channel carries the **sum** of the capacity.
- Various kinds: by time, by space, by frequency, by polarisation, ...

Multiplexer and demultiplexer

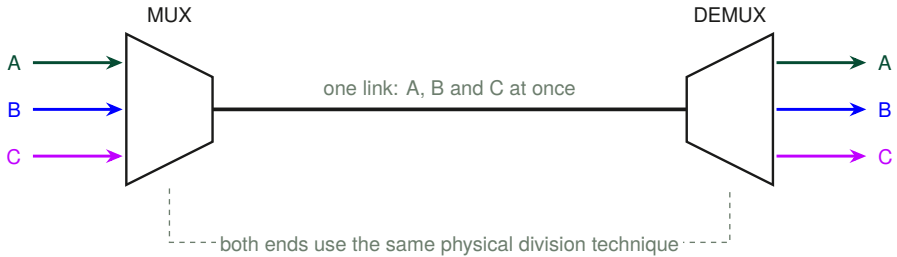

- **Multiplexer**: several channels onto one channel.
- **Demultiplexer**: split one channel into several channels.
- The inner channel carries the **sum** of the capacity.
- Various kinds: by time, by space, by frequency, by polarisation, ...

Multiplexer and demultiplexer

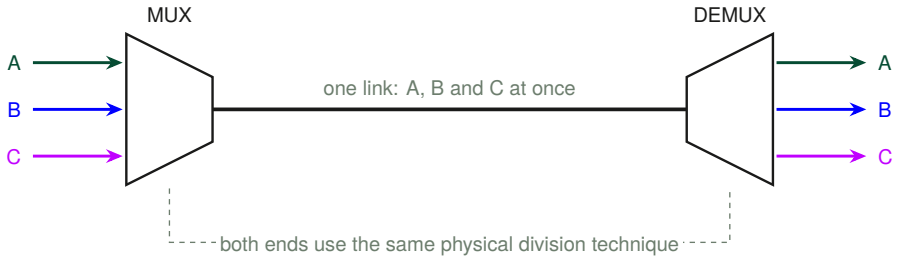

- **Multiplexer**: several channels onto one channel.
- **Demultiplexer**: split one channel into several channels.
- The inner channel carries the **sum** of the capacity.
- Various kinds: by time, by space, by frequency, by polarisation, ...

Multiplexer and demultiplexer

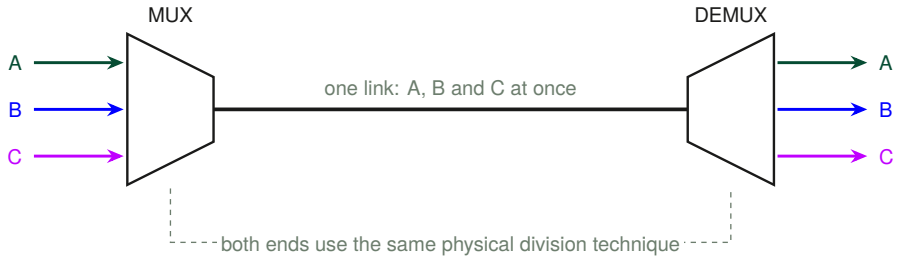

- **Multiplexer**: several channels onto one channel.
- **Demultiplexer**: split one channel into several channels.
- The inner channel carries the **sum** of the capacity.
- Various kinds: by time, by space, by frequency, by polarisation, ...

State-of-the-art physical links

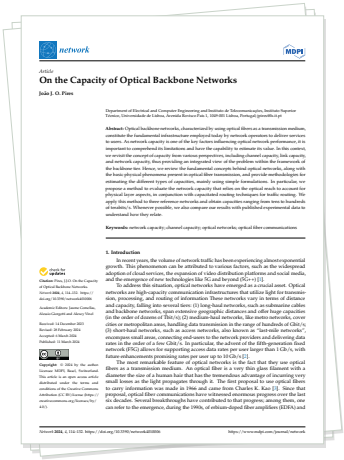
network

MDPI

Article

On the Capacity of Optical Backbone Networks

João J. O. Pires

Department of Electrical and Computer Engineering and Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal; jpires@lx.it.pt

Abstract: Optical backbone networks, characterized by using optical fibers as a transmission medium, constitute the fundamental infrastructure employed today by network operators to deliver services to users. As network capacity is one of the key factors influencing optical network performance, it is important to comprehend its limitations and have the capability to estimate its value. In this context, we revisit the concept of capacity from various perspectives, including channel capacity, link capacity, and network capacity, thus providing an integrated view of the problem within the framework of the backbone tier. Hence, we review the fundamental concepts behind optical networks, along with the basic physical phenomena present in optical fiber transmission, and provide methodologies for estimating the different types of capacities, mainly using simple formulations. In particular, we propose a method to evaluate the network capacity that relies on the optical reach to account for physical layer aspects, in conjunction with capacitated routing techniques for traffic routing. We apply this method to three reference networks and obtain capacities ranging from tens to hundreds of terabit/s. Whenever possible, we also compare our results with published experimental data to understand how they relate.

Keywords: network capacity; channel capacity; optical networks; optical fiber communications

1. Introduction

In recent years, the volume of network traffic has been experiencing almost exponential growth. This phenomenon can be attributed to various factors, such as the widespread adoption of cloud services, the expansion of video distribution platforms and social media, and the emergence of new technologies like 5G and beyond (5G+) [1].

To address this situation, optical networks have emerged as a crucial asset. Optical networks are high-capacity communication infrastructures that utilize light for transmitting information through optical fibers, enabling fast and reliable data transfer with high bandwidth and capacity, falling into several tiers: (1) long-haul networks, such as undersea cables and backbone networks, span extensive geographic distances and offer huge capacities (in the order of dozens of Tbit/s); (2) medium-haul networks, like metro networks, cover cities or metropolitan areas, handling data transmission in the range of hundreds of Gbit/s; (3) short-haul networks, such as access networks, also known as "last-mile networks", encompass small areas, connecting end-users to the network providers and delivering data rates in the order of a few Gbit/s. In particular, the advent of the fifth-generation fixed network (F5G) allows for supporting access data rates per user larger than 1 Gb/s, with future enhancements promising rates per user up to 10 Gb/s [2].

The most remarkable feature of optical networks is the fact that they use optical fibers as the transmission medium. The ability of a fiber to guide light through a path as thin in diameter the size of a human hair that has the tremendous advantage of incurring very small losses as the light propagates through it. The first proposal to use optical fibers to carry information was made in 1966 and came from Charles K. Kao [3]. Since that proposal, optical fiber communications have witnessed enormous progress over the last six decades. Several breakthroughs have contributed to that progress; among them, one can refer to the emergence, during the 1980s, of erbium-doped fiber amplifiers (EDFAs) and

check for updates

Citation: Pires, J.J.O. On the Capacity of Optical Backbone Networks. Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006

Academic Editors: Joanna Conscillo, Alessio Giampietri and Alessia Vinai

Received: 14 December 2023
Revised: 30 February 2024
Accepted: 6 March 2024
Published: 14 March 2024

Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006 https://www.mdpi.com/journal/network

On the capacity of optical backbone networks Pires (2024)

Technology:

- Electronics versus photonics
- Dense Wavelength Division Multiplexing (DWDM)
- Reconfigurable optical add-drop multiplexer (ROADM)

29.25 Tbit/s 1600 km
29.50 Tbit/s 650 km
41.00 Tbit/s 100 km
52.15 Tbit/s 80 km

1	1k	1M	1G	1T	1P
10^0	10^3	10^6	10^9	10^{12}	10^{15}

State-of-the-art physical links

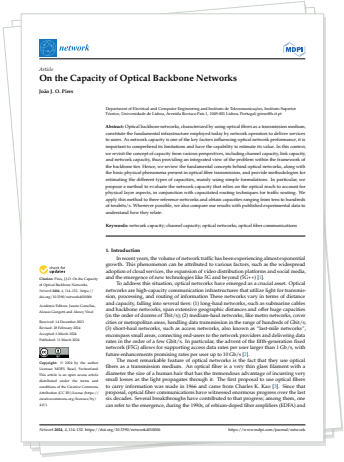
network

MDPI

Article

On the Capacity of Optical Backbone Networks

João J. O. Pires

Department of Electrical and Computer Engineering and Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal; jpires@lx.it.pt

Abstract: Optical backbone networks, characterized by using optical fibers as a transmission medium, constitute the fundamental infrastructure employed today by network operators to deliver services to users. As network capacity is one of the key factors influencing optical network performance, it is important to comprehend its limitations and have the capability to estimate its value. In this context, we revisit the concept of capacity from various perspectives, including channel capacity, link capacity, and network capacity, thus providing an integrated view of the problem within the framework of the backbone tier. Hence, we review the fundamental concepts behind optical networks, along with the basic physical phenomena present in optical fiber transmission, and provide methodologies for estimating the different types of capacities, mainly using simple formulations. In particular, we propose a method to evaluate the network capacity that relies on the optical reach to account for physical layer aspects, in conjunction with capacitated routing techniques for traffic routing. We apply this method to three reference networks and obtain capacities ranging from tens to hundreds of terabits/s. Whenever possible, we also compare our results with published experimental data to understand how they relate.

Keywords: network capacity; channel capacity; optical networks; optical fiber communications

1. Introduction

In recent years, the volume of network traffic has been experiencing almost exponential growth. This phenomenon can be attributed to various factors, such as the widespread adoption of cloud services, the expansion of video distribution platforms and social media, and the emergence of new technologies like 5G and beyond (5G++) [1].

To address this situation, optical networks have emerged as a crucial asset. Optical networks are high-capacity communication infrastructures that utilize light for transmitting information through optical fibers, offering numerous benefits in terms of bandwidth and capacity, falling into several tiers: (1) long-haul networks, such as submarine cables and backbone networks, span extensive geographic distances and offer huge capacities (in the order of dozens of Tbit/s); (2) medium-haul networks, like metro networks, cover cities or metropolitan areas, handling data transmission in the range of hundreds of Gbit/s; (3) short-haul networks, such as access networks, also known as "last-mile networks", encompass small areas, connecting end-users to the network providers and delivering data rates in the order of a few Gbit/s. In particular, the advent of the fifth-generation fixed network (F5G) allows for supporting access data rates per user larger than 1 Gb/s, with future enhancements promising rates per user up to 10 Gb/s [2].

The most remarkable feature of optical networks is the fact that they use optical fibers for transmitting information. This is due to the fact that the transparency of the fiber's diameter the size of a human hair that has the tremendous advantage of incurring very small losses as the light propagates through it. The first proposal to use optical fibers to carry information was made in 1966 and came from Charles K. Kao [3]. Since that proposal, optical fiber communications have witnessed enormous progress over the last six decades. Several breakthroughs have contributed to that progress; among them, one can refer to the emergence, during the 1980s, of erbium-doped fiber amplifiers (EDFAs) and

Citation: Pires, J.J.O. On the Capacity of Optical Backbone Networks. Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006

Academic Editors: Joanna Conscillo, Alvaro Gregori and Alexey Vinel

Received: 14 December 2023
Revised: 30 February 2024
Accepted: 6 March 2024
Published: 14 March 2024

Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006 https://www.mdpi.com/journal/network

On the capacity of optical backbone networks Pires (2024)

Technology:

- Electronics versus photonics
- Dense Wavelength Division Multiplexing (DWDM)
- Reconfigurable optical add-drop multiplexer (ROADM)

29.25 Tbit/s 1600 km
29.50 Tbit/s 650 km
41.00 Tbit/s 100 km
52.15 Tbit/s 80 km

1 10^0 1k 10^3 1M 10^6 1G 10^9 1T 10^{12} 1P 10^{15}

State-of-the-art physical links

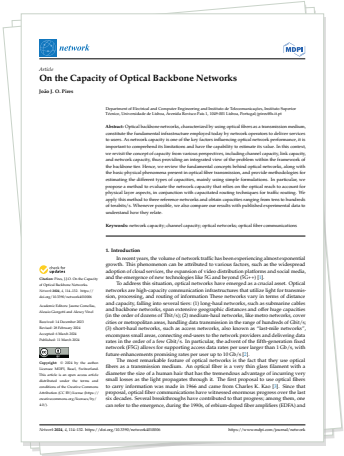
network

MDPI

Article

On the Capacity of Optical Backbone Networks

João J. O. Pires

Department of Electrical and Computer Engineering and Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal; jpires@lx.it.pt

Abstract: Optical backbone networks, characterized by using optical fibers as a transmission medium, constitute the fundamental infrastructure employed today by network operators to deliver services to users. As network capacity is one of the key factors influencing optical network performance, it is important to comprehend its limitations and have the capability to estimate its value. In this context, we revisit the concept of capacity from various perspectives, including channel capacity, link capacity, and network capacity, thus providing an integrated view of the problem within the framework of the backbone tier. Hence, we review the fundamental concepts behind optical networks, along with the basic physical phenomena present in optical fiber transmission, and provide methodologies for estimating the different types of capacities, mainly using simple formulations. In particular, we propose a method to evaluate the network capacity that relies on the optical reach to account for physical layer aspects, in conjunction with capacitated routing techniques for traffic routing. We apply this method to three reference networks and obtain capacities ranging from tens to hundreds of terabit/s. Whenever possible, we also compare our results with published experimental data to understand how they relate.

Keywords: network capacity; channel capacity; optical networks; optical fiber communications

1. Introduction

In recent years, the volume of network traffic has been experiencing almost exponential growth. This phenomenon can be attributed to various factors, such as the widespread adoption of cloud services, the expansion of video distribution platforms and social media, and the emergence of new technologies like 5G and beyond (5G+) [1].

To address this situation, optical networks have emerged as a crucial asset. Optical networks are high-capacity communication infrastructures that utilize light for transmitting information through optical fibers, offering numerous benefits, such as high bandwidth and capacity, falling into several tiers: (1) long-haul networks, such as submarine cables and backbone networks, span extensive geographic distances and offer huge capacities (in the order of dozens of Tbit/s); (2) medium-haul networks, like metro networks, cover cities or metropolitan areas, handling data transmission in the range of hundreds of Gbit/s; (3) short-haul networks, such as access networks, also known as "last-mile networks", encompass small areas, connecting end-users to the network providers and delivering data rates in the order of a few Gbit/s. In particular, the advent of the fifth-generation fixed network (F5G) allows for supporting access data rates per user larger than 1 Gb/s, with future enhancements promising rates per user up to 10 Gb/s [2].

The most remarkable feature of optical networks is the fact that they use optical fibers for transmitting information. This is due to the fact that the transparency of the fiber to the diameter the size of a human hair that has the tremendous advantage of incurring very small losses as the light propagates through it. The first proposal to use optical fibers to carry information was made in 1966 and came from Charles K. Kao [3]. Since that proposal, optical fiber communications have witnessed enormous progress over the last six decades. Several breakthroughs have contributed to that progress; among them, one can refer to the emergence, during the 1980s, of erbium-doped fiber amplifiers (EDFAs) and

Citation: Pires, J.J.O. On the Capacity of Optical Backbone Networks. Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006

Academic Editors: Joanna Conscillo, Alessio Giampietri and Alessio Vecchi

Received: 14 December 2023
Revised: 30 February 2024
Accepted: 6 March 2024
Published: 14 March 2024

Copyright: © 2024 by the author(s). Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006 https://www.mdpi.com/journal/network

On the capacity of optical backbone networks Pires (2024)

Technology:

- Electronics versus photonics
- Dense Wavelength Division Multiplexing (DWDM)
- Reconfigurable optical add-drop multiplexer (ROADM)

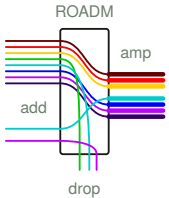

29.25 Tbit/s 1600 km
29.50 Tbit/s 650 km
41.00 Tbit/s 100 km
52.15 Tbit/s 80 km

1 1k 1M 1G 1T 1P
10^0 10^3 10^6 10^9 10^{12} 10^{15}

State-of-the-art physical links

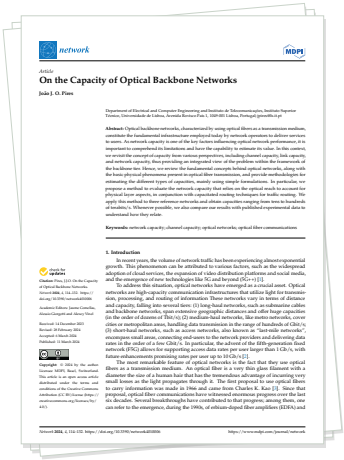
network

MDPI

Article

On the Capacity of Optical Backbone Networks

João J. O. Pires

Department of Electrical and Computer Engineering and Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal; jpires@lx.it.pt

Abstract: Optical backbone networks, characterized by using optical fibers as a transmission medium, constitute the fundamental infrastructure employed today by network operators to deliver services to users. As network capacity is one of the key factors influencing optical network performance, it is important to comprehend its limitations and have the capability to estimate its value. In this context, we revisit the concept of capacity from various perspectives, including channel capacity, link capacity, and network capacity, thus providing an integrated view of the problem within the framework of the backbone tier. Hence, we review the fundamental concepts behind optical networks, along with the basic physical phenomena present in optical fiber transmission, and provide methodologies for estimating the different types of capacities, mainly using simple formulations. In particular, we propose a method to evaluate the network capacity that relies on the optical reach to account for physical layer aspects, in conjunction with capacitated routing techniques for traffic routing. We apply this method to three reference networks and obtain capacities ranging from tens to hundreds of terabits/s. Wherever possible, we also compare our results with published experimental data to understand how they relate.

Keywords: network capacity; channel capacity; optical networks; optical fiber communications

1. Introduction

In recent years, the volume of network traffic has been experiencing almost exponential growth. This phenomenon can be attributed to various factors, such as the widespread adoption of cloud services, the expansion of video distribution platforms and social media, and the emergence of new technologies like 5G and beyond (5G++) [1].

To address this situation, optical networks have emerged as a crucial asset. Optical networks are high-capacity communication infrastructures that utilize light for transmitting data through optical fibers, providing numerous benefits over their electronic counterparts and capacity, falling into several tiers: (1) long-haul networks, such as submarine cables and backbone networks, span extensive geographic distances and offer huge capacities (in the order of dozens of Tbit/s); (2) medium-haul networks, like metro networks, cover cities or metropolitan areas, handling data transmission in the range of hundreds of Gbit/s; (3) short-haul networks, such as access networks, also known as "last-mile networks", encompass small areas, connecting end-users to the network providers and delivering data rates in the order of a few Gbit/s. In particular, the advent of the fifth-generation fixed network (F5G) allows for supporting access data rates per user larger than 1 Gb/s, with future enhancements promising rates per user up to 10 Gb/s [2].

The most remarkable feature of optical networks is the fact that they use optical fibers for transmitting information. Optical fibers are thin, flexible, transparent fibers as thin in diameter the size of a human hair that has the tremendous advantage of incurring very small losses as the light propagates through it. The first proposal to use optical fibers to carry information was made in 1966 and came from Charles K. Kao [3]. Since that proposal, optical fiber communications have witnessed enormous progress over the last six decades. Several breakthroughs have contributed to that progress; among them, one can refer to the emergence, during the 1980s, of erbium-doped fiber amplifiers (EDFAs) and

Check for updates

Citation: Pires, J.J.O. On the Capacity of Optical Backbone Networks. Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006

Academic Editors: Joanna Conscillo, Alvaro Gazquez and Alonso Viteri

Received: 14 December 2023
Revised: 28 February 2024
Accepted: 6 March 2024
Published: 14 March 2024

Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Network 2024, 4, 114–132. https://doi.org/10.3390/network4020006 https://www.mdpi.com/journal/network

On the capacity of optical backbone networks Pires (2024)

Technology:

- Electronics versus photonics
- Dense Wavelength Division Multiplexing (DWDM)
- Reconfigurable optical add-drop multiplexer (ROADM)

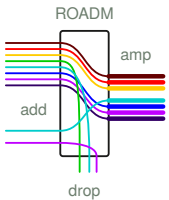

29.25 Tbit/s	1600 km
29.50 Tbit/s	650 km
41.00 Tbit/s	100 km
52.15 Tbit/s	80 km

1 10^0 1k 10^3 1M 10^6 1G 10^9 1T 10^{12} 1P 10^{15}

State-of-the-art physical links

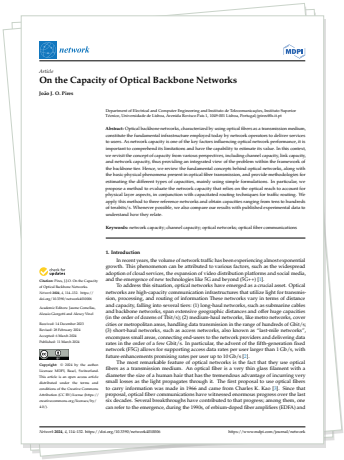
network

MDPI

Article

On the Capacity of Optical Backbone Networks

João J. O. Pires

Department of Electrical and Computer Engineering and Instituto de Telecomunicações, Instituto Superior Técnico, Universidade de Lisboa, Avenida Rovisco Pais 1, 1049-001 Lisboa, Portugal; jpires@lx.it.pt

Abstract: Optical backbone networks, characterized by using optical fibers as a transmission medium, constitute the fundamental infrastructure employed today by network operators to deliver services to users. As network capacity is one of the key factors influencing optical network performance, it is important to comprehend its limitations and have the capability to estimate its value. In this context, we revisit the concept of capacity from various perspectives, including channel capacity, link capacity, and network capacity, thus providing an integrated view of the problem within the framework of the backbone tier. Hence, we review the fundamental concepts behind optical networks, along with the basic physical phenomena present in optical fiber transmission, and provide methodologies for estimating the different types of capacities, mainly using simple formulations. In particular, we propose a method to evaluate the network capacity that relies on the optical reach to account for physical layer aspects, in conjunction with capacitated routing techniques for traffic routing. We apply this method to three reference networks and obtain capacities ranging from tens to hundreds of terabits/s. Whenever possible, we also compare our results with published experimental data to understand how they relate.

Keywords: network capacity; channel capacity; optical networks; optical fiber communications

1. Introduction

In recent years, the volume of network traffic has been experiencing almost exponential growth. This phenomenon can be attributed to various factors, such as the widespread adoption of cloud services, the expansion of video distribution platforms and social media, and the emergence of new technologies like 5G and beyond (5G++) [1].

To address this situation, optical networks have emerged as a critical asset. Optical networks are high-capacity communication infrastructures that utilize light for transmitting data through optical fibers. Compared to traditional copper wires, they offer higher speed and capacity, falling into several tiers: (1) long-haul networks, such as submarine cables and backbone networks, span extensive geographic distances and offer huge capacities (in the order of dozens of Tbit/s); (2) medium-haul networks, like metro networks, cover cities or metropolitan areas, handling data transmission in the range of hundreds of Gbit/s; (3) short-haul networks, such as access networks, also known as "last-mile networks", encompass small areas, connecting end-users to the network providers and delivering data rates in the order of a few Gbit/s. In particular, the advent of the fifth-generation fixed network (F5G) allows for supporting access data rates per user larger than 1 Gb/s, with future enhancements promising rates per user up to 10 Gb/s [2].

The most remarkable feature of optical networks is the fact that they use optical fibers for transmitting information. Optical fibers are thin, flexible, transparent fibers as thin in diameter the size of a human hair that has the tremendous advantage of incurring very small losses as the light propagates through it. The first proposal to use optical fibers to carry information was made in 1966 and came from Charles K. Kao [3]. Since that proposal, optical fiber communications have witnessed enormous progress over the last six decades. Several breakthroughs have contributed to that progress; among them, one can refer to the emergence, during the 1980s, of erbium-doped fiber amplifiers (EDFAs) and

Citation: Pires, J.J.O. On the Capacity of Optical Backbone Networks. *Network* **2024**, *4*, 114–132. https://doi.org/10.3390/network4020006

Academic Editors: Joanna Conscillo, Alvaro Gonzalez and Alexey Vinel

Received: 14 December 2023
Revised: 30 February 2024
Accepted: 6 March 2024
Published: 14 March 2024

Copyright: © 2024 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Network **2024**, *4*, 114–132. https://doi.org/10.3390/network4020006 https://www.mdpi.com/journal/network

On the capacity of optical backbone networks Pires (2024)

Technology:

- Electronics versus photonics
- Dense Wavelength Division Multiplexing (DWDM)
- Reconfigurable optical add-drop multiplexer (ROADM)

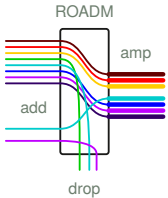

29.25 Tbit/s	1600 km
29.50 Tbit/s	650 km
41.00 Tbit/s	100 km
52.15 Tbit/s	80 km

1	1k	1M	1G	1T	1P
10^0	10^3	10^6	10^9	10^{12}	10^{15}

State-of-the-art: physical networks (1/2)

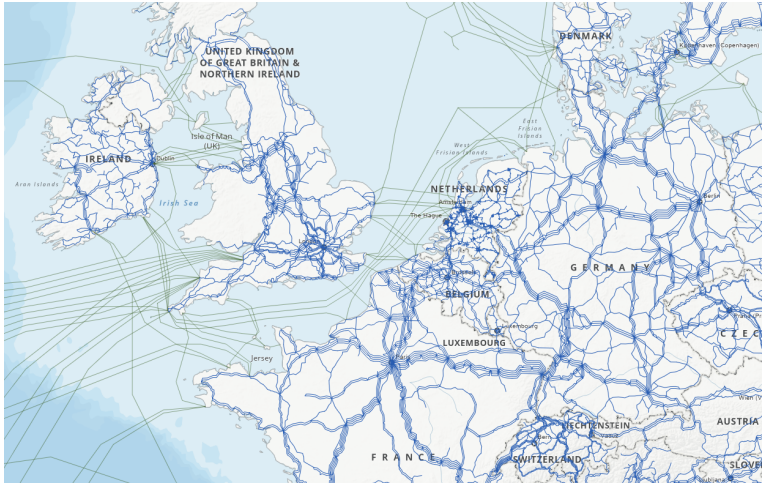

(Source: UN International Telecommunication Union (ITU), 2026)

State-of-the-art: physical networks (2/2)

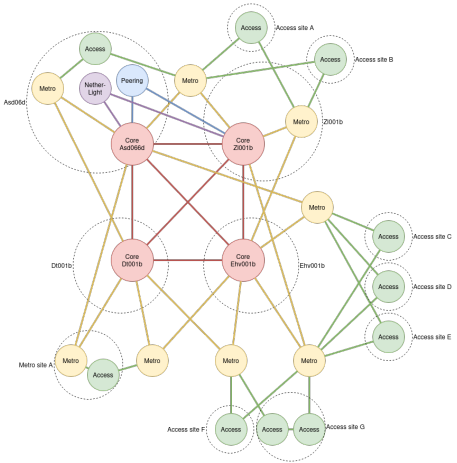

(Source: SURFnet∞, 2026)

Core sites:

- Amsterdam (Nikhef)
- Zwolle (Windesheim)
- Delft (TU Delft)
- Eindhoven (TU Eindhoven)

Core links: 800Gbit/s

Metro links: 400Gbit/s

Access links: 100Gbit/s

Public internet (e.g. AMS-IX)

30+ NRENs (e.g. CERN)

ISP for NATO Communications and Information Agency – The Hague

Glass can be listened to



Optical Fiber Tapping: Methods and Precautions

M. Zafar Iqbal, Habib Fathallah, Nezih Hedhadj

Abstract— Optical fiber communication is seen as serve as potentially the most secure means of communication since any attempt at tapping information into a fiber link, while avoiding detection, has hitherto been been reported as a successfully tapped fiber is difficult to detect. In this paper we present a brief overview of intrusive and non-intrusive tapping methods. We report simulation of optical characteristics of a fiber being tapped by 'bend' method and proof of concept with physical experiment. We also present countermeasures that are available in the commercial arena to support our assumptions security of a fiber link from existing technologies. Some measures to prevent fiber tapping or to detect the unauthorized, surreptitious tapped from fiber were also discussed.

***Index Terms*— Optical Fiber Tapping, Layer 2 Encryption Eavesdropping, Bend tapping.**

I. INTRODUCTION

CONTRARY to the common perception, optical fiber is not inherently secure from tapping or eavesdropping. The enormous amount of revenue critical and sensitive information carried over fiber these days is expected to any unauthorized tapping/eavesdropping [1].

Fiber tapping is a process by which the security of optical fiber is compromised by either extracting or injecting information (or light). Basically fiber tapping can be intrusive and non-intrusive. The former requires the fiber to be broken and connectorized, while the latter requires only the later achieves tapping without cutting the fiber or causing any service disruption. Non-intrusive technique is the focus of this work.

Only a few incidents of fiber tapping could be reported, as it is difficult to detect. However, the process of tapping presents itself is quite simple. Major reported incidents of tapping include the following:

- 2002, three mask-clad Sons of Dastardly Tekkies were branched at Frankfurt Airport in Germany [1].
- 2004, Australian Federal Police (AFP) discovered hushed new Victoria's optical network [1].
- 2005, USA Enemy Carter, submarine specifically retrofitted to conduct tapping into undersea cables [2], [3].

In the following sections we present a brief overview of intrusive and nonintrusive tapping techniques [4]. Then we present simulation of optical characteristics of a fiber being bending followed by a report on physical demonstration of tapping on a prototype developed in our lab. Then we also explain the prototype design, hardware and software. We also discuss possible tapping scenarios in real environments highlighting the measures required to achieve it. Finally we propose some solutions to protect optical fiber links against tapping.

This paper is based on work that was supported by the Royal Saudi Air Force under Project No. 424/329.

M. Zafar Iqbal is working in Prince Sultan Advance Technologies Research Institute (zafar@ksu.edu.sa).

Habib Fathallah is working as Assistant Professor at King Saud University (hfathallah@ksu.edu.sa).

Nezih Hedhadj is graduated researcher at Laval University (nhedhadj@gel.ulaval.ca).

978-1-4577-1149-9/11/$26.00 ©2011 IEEE

II. FIBER TAPPING METHODS

A. Fiber Bending

In this method cable is stripped down to the fiber for bending. This method exploits the principle of propagation of light through an optical fiber better described as the total internal reflection. The light is kept in the core because the light on the core-cladding interface should be greater than the Critical Angle for total internal reflection. Otherwise some light will radiate out of the fiber through its cladding. The critical angle is a function of the refractive indices of the core and cladding, and represented by the following equation

$$\theta_c = \cos^{-1}(n_{cladding}/n_{core}) \qquad (provided\ n_1 > n_2)$$

Here, θ_c is the critical angle, n_2 is refractive index of Cladding, n_1 is refractive index of Cladding.

In this method technique the fiber is bent such that the angle of incidence becomes less than the critical angle and the light radiates. Apparently there are further two types of fiber bending:

1) Macro Bending:

Macro bending of optical fiber results in sharp but wavelength-size curvature resulting in axial displacements of a few microns and spatial wavelength displacements of a few millimeters (Figure 1). The light thus radiated is used for tapping.

Figure 1 Micro Bending

2) Micro Bending:

Micro is a minimum tolerable bend radius associated with each fiber type. A lesser bend radius will result in radiation of Light (figure 2). This property can be used to extract light from fiber for nave shopping. Normally, single mode fibers will not tolerate a minimum bend radius of less than 6.5 to 7.5 cm except some specially developed types. While a

Optical fiber tapping: methods and precautions Iqbal, Pattinson and Kor (2011)

- **Bend it** far enough and light leaves the core. A coupler on the outside of an unbroken fibre is a tap.
- No cut, no splice, no loss of service, and a power drop small enough to look like a dirty connector.
- Countermeasures are **monitoring** and **encryption**, not the medium.

Fibre is harder to tap than copper.
Harder is not the same as **cannot**.

Glass can be listened to



Optical Fiber Tapping: Methods and Precautions

M. Zafar Iqbal, Habib Fathallah, Nezih Hedhadj

Abstract— Optical fiber communication is seen as secure as generally it is difficult to tap optical fiber without detection. However, tapping information into a fiber link, while avoiding detection, has incidents have been reported as a successfully staged fiber in different ways. This paper reviews the different methods of optical fiber tapping methods. We report simulation of optical characteristics of a fiber being tapped by 'bend' method and proof of concept with physical experiment. We also propose countermeasures to detect tapping and to reasonable eavesdropper non-compromise security of a fiber link from existing technologies. Some measures to prevent fiber tapping or to detect tapping, if unfortunate, of information tapped from fiber were also discussed.

***Index Terms*— Optical Fiber Tapping, Layer 2 Encryption Eavesdropping, Bend tapping.**

I. INTRODUCTION

CONTRARY to the common perception, optical fiber is not inherently secure from tapping or eavesdropping. The enormous amount of mission critical and sensitive information carried over fiber these days is expected to any eavesdropper/hacker/intruder/spy/terrorist.

Fiber tapping is a process by which the security of optical fiber is compromised by either extracting or injecting information (or light). Basically fiber tapping can be intrusive and non-intrusive. The former requires the fiber to be cut or spliced, whereas the latter is a cut-free technique. In the later unknown tapping without cutting the fiber or causing any service disruption. Non-intrusive technique is the focus of this work.

Only a few incidents of fiber tapping could be reported, as it is difficult to detect tapping and even if it is detected the process itself is quite simple. Major reported incidents of tapping include the following:

- 2002, three mask-clad Sons of Daniaiche Tekkiem were branched at Frankfurt Airport in Germany [1].
- 2004, Chinese military hackers were discovered hacked into Vietnam's optical network [2].
- 2005, USS Jimmy Carter, submarine specifically retrofitted to conduct tapping into undersea cables [3], [4].

In the following sections we present a brief overview of intrusive and nonintrusive tapping techniques [4]. Then we present our experimental setup and proof of concept for bend bending followed by a report on physical demonstration of tapping on a prototype developed in our lab. Then we also explain the prototype design, hardware and software. We also discuss possible tapping scenarios in real environments highlighting the measures required to achieve it. Finally we propose some solutions to protect optical fiber links against tapping.

II. FIBER TAPPING METHODS

A. Fiber Bending

In this method cable is stripped down to the fiber for bending. This method exploits the principle of propagation of light through an optical fiber better described as the total internal reflection [5]. Whenever the angle of incidence of light on the core-cladding interface should be greater than the Critical Angle for total internal reflection. Otherwise some light will radiate out of the fiber through its cladding. The critical angle is a function of the refractive indices of the core and cladding, and represented by the following equation

$$\theta_c = \cos^{-1}(n_{cladding}/n_{core}) \qquad [provided\ n_1 > n_2]$$

Here, θ_c *is the critical angle,* n_1 *is refractive index of Cladding,* n_2 *is refractive index of Cladding.*

In this bending technique the fiber is bent such that the angle of incidence becomes less than the critical angle and the light radiates. Apparently there are further two types of fiber bending:

1) Macro Bending:

Macro bending is optical force results in sharp but wavelength-size-order resulting in axial displacements of a few microns and spatial wavelength displacements of a few millimeters (Figure 1). The light thus radiated is used for tapping.

Figure 1 Micro Bending

2) Micro Bending:

Micro is a minimum tolerable bend radius associated with each fiber type. A lesser bend radius will result in radiation of Light (figure 2). This property can be used to extract light from fiber for nave shopping. Normally, single mode fibers will not tolerate a minimum bend radius of less than 6.5 to 7.5 cm except some specially developed types. While a

978-1-4577-1149-9/11/$26.00 ©2011 IEEE

This paper is based on work that was supported by the Royal Saudi Air Force through King Saud University.

M. Zafar Iqbal is working in Prince Sultan Advance Technologies Research Institute (zafar@ksu.edu.sa).

Habib Fathallah is working as Assistant Professor at King Saud University (hfathallah@ksu.edu.sa).

Nezih Hedhadj is graduated researcher at Laval University (nhedhadj@gel.ulaval.ca).

Optical fiber tapping: methods and precautions Iqbal, Pattinson and Kor (2011)

- **Bend it** far enough and light leaves the core. A coupler on the outside of an unbroken fibre is a tap.
- No cut, no splice, no loss of service, and a power drop small enough to look like a dirty connector.
- Countermeasures are **monitoring** and **encryption**, not the medium.

Fibre is harder to tap than copper.
Harder is not the same as **cannot**.

Glass can be listened to



Optical Fiber Tapping: Methods and Precautions

M. Zafar Iqbal, Habib Fathallah, Nazih Hellhalj

Abstract— Optical fiber communication is seen as secure as generally it is difficult to tap optical fiber without detection. However, tapping information into a fiber link, while avoiding detection, has incidents have been reported as a successfully staged fiber in different ways. This paper reviews the different methods of optical fiber tapping methods. We report simulation of optical characteristics of a fiber being tapped by 'bend' method and proof of concept with physical experiment. We also discuss countermeasures to detect tapping and the reasonable an advantages/ our components security of a fiber link from existing technologies. Some measures to prevent fiber tapping or to detect tapping, if unfortunate, of information tapped from fiber were also discussed.

Index Terms— **Optical Fiber Tapping, Layer 2 Encryption Eavesdropping, Bend tapping.**

I. INTRODUCTION

CONTRARY to the common perception, optical fiber is not inherently secure from tapping or eavesdropping. The enormous amount of mission critical and sensitive information carried over fiber these days is expected to any unauthorized tapping/eavesdropping [1].

Fiber tapping is a process by which the security of optical fiber is compromised by either extracting or injecting information (or light). Basically fiber tapping can be intrusive and non-intrusive. The former requires the fiber to be cut and splice-in another fiber or a coupler whereas in the later unknown tapping without cutting the fiber or causing any service disruption. Non-intrusive technique is the focus of this work.

Only a few incidents of fiber tapping could be reported, as it is difficult to detect tapping and the tapping process itself is quite simple. Major reported incidents of tapping include the following:

- 2000, three mark Sent of Deutsche Telekom were breached at Frankfurt Airport in Germany [1].
- 2005, Australian Federal Police (AFP) discovered hacked into Victoria's optical network [1].
- 2005, USA Enemy Carter, submarine specifically retrofitted to conduct tapping into undersea cables [2], [3].

In the following sections we present a brief overview of intrusive and nonintrusive tapping techniques [4]. Then we present simulation of optical characteristics of a fiber being bending followed by a report on physical demonstration of tapping on a prototype developed in our lab. Then we also explain the prototype design, hardware and software. We also discuss possible tapping scenarios in real environments highlighting the measures required to achieve it. Finally we propose some solutions to protect optical fiber links against tapping.

This paper is based on work that was supported by the Royal Saudi Air Force through King Saud University.

M. Zafar Iqbal is working in Prince Sultan Advance Technologies Research Institute (zafar@ksu.edu.sa).

Habib Fathallah is working as Assistant Professor at King Saud University (hfathallah@ksu.edu.sa).

Nazih Hellhalj is a graduate/ postdoctoral researcher at Laval University (nhellhalj@gel.ulaval.ca).

978-1-4577-1149-9/11/$26.00 ©2011 IEEE

II. FIBER TAPPING METHODS

A. Fiber Bending

In this method cable is stripped down to the fiber for bending. This method exploits the principle of propagation of light through an optical fiber better described as the total internal reflection. The light signal in the form of pulses of light on the core-cladding interface should be greater than the Critical Angle for total internal reflection. Otherwise some light will radiate out of the fiber through its cladding. The critical angle is a function of the refractive indices of the core and cladding, and represented by the following equation

$$\theta_c = \cos^{-1}(n_{cladding}/n_{core}) \qquad (provided\ n_1 > n_2)$$

Where, θ_c is the critical angle, n_1 is refractive index of Cladding, n_2 is refractive index of Cladding.

In this method technique the fiber is bent such that the angle of incidence becomes less than the critical angle and the light radiates. Apparently there are further two types of fiber bending:

1) Macro Bending

Macro bending is optimal force results in sharp but microscopic curvature resulting in axial displacements of a few microns and spatial wavelength displacements of a few millimeters (Figure 1). The light thus radiated is used for tapping.



Figure 1 Micro Bending

2) Micro Bending

Micro is a minimum tolerable bend radius associated with each fiber type. A lesser bend radius will result in radiation of Light (figure 2). This property can be used to extract light from fiber for micro sharping. Normally, single mode fibers will not tolerate a minimum bend radius of less than 6.5 to 7.5 cm except some specially developed types. While a

- **Bend it** far enough and light leaves the core. A coupler on the outside of an unbroken fibre is a tap.
- No cut, no splice, no loss of service, and a power drop small enough to look like a dirty connector.
- Countermeasures are **monitoring** and **encryption**, not the medium.

Fibre is harder to tap than copper.
Harder is not the same as **cannot**.

Optical fiber tapping: methods and precautions Iqbal, Pattinson and Kor (2011)

Glass can be listened to



Optical Fiber Tapping: Methods and Precautions

M. Zafar Iqbal, Habib Fathallah, Nezih Hedhadj

Abstract— Optical fiber communication is seen as secure as generally it is difficult to tap optical fiber without detection. However, tapping information into a fiber link, while avoiding detection, has incidents have been reported as a successfully tapped fiber is difficult to detect. This paper reviews the optical fiber tapping methods and methods. We report simulation of optical characteristics of a fiber being tapped by 'bend' method and proof of concept with physical experiments. We also propose countermeasures to detect tapping by a commercial-on-the-shelf optical time domain reflectometer. Some measures to prevent fiber tapping or to increase the likelihood of information tapped from fiber were also discussed.

Index Terms— **Optical Fiber Tapping, Layer 2 Encryption Eavesdropping, Bend tapping.**

I. INTRODUCTION

CONTRARY to the common perception, optical fiber is not inherently secure from tapping or eavesdropping. The enormous amount of mission critical and sensitive information carried over fiber fence days is expected to any unauthorized information tapping activities.

Fiber tapping is a process by which the security of optical fiber is compromised by either extracting or injecting information (or light). Basically fiber tapping can be intrusive and non-intrusive. The former requires the fiber to be broken or accessed through splices or couplers whereas in the later unknown tapping without cutting the fiber or causing any service disruption. Non-intrusive technique is the focus of this work.

Only a few incidents of fiber tapping could be reported, as it is difficult to detect a successful tapping. However, the process itself is quite simple. Major reported incidents of tapping include the following:

- 2002, three mark 3nm of Deutsche Telekom were breached at Frankfurt Airport in Germany [1].
- 2004, Russian Lawrence Livermore lab discovered hooked into Sprint's optical network [2].
- 2005, USA Enemy Carter, submarine specifically retrofitted to conduct tapping into undersea cables [3], [4].

In the following sections we present a brief overview of intrusive and nonintrusive tapping techniques [4]. Then we present simulation of optical characteristics of a fiber with bending followed by a report on physical demonstration of tapping on a prototype developed in our lab. Then we also explain the prototype design, hardware and software. We also discuss possible tapping scenarios in real environments highlighting the measures required to achieve it. Finally we propose some solutions to protect optical fiber links against tapping.

II. FIBER TAPPING METHODS

A. Fiber Bending

In this method cable is stripped down to the fiber for bending. This method exploits the principle of propagation of light through an optical fiber better described as the total internal reflection. The light signal in the form of pulses of light on the core-cladding interface should be greater than the Critical Angle for total internal reflection. Otherwise some light will radiate out of the fiber through its cladding. The critical angle is a function of the refractive indices of the core and cladding, and represented by the following equation

$$\theta_c = \cos^{-1}\left(\frac{n_{cladding}}{n_{core}}\right) \qquad (provided\ n_{co} > n_{cl})$$

Here, θ_c is the critical angle, n_{cl} is refractive index of Cladding, n_{co} is refractive index of Cladding.

1) *Micro Bending:* This technique the fiber is bent such that the angle of incidence becomes less than the critical angle and the light radiates. Apparently there are further two types of fiber bending:

i) Macro Bending: Macroscopic external force results in sharp but microscopic curvature resulting in axial displacements of a few microns and spatial wavelength displacements of a few millimeters (Figure 1). The light thus radiated is used for tapping.



Figure 1 Micro Bending

2) Macro Bending: Macro is a minimum tolerable bend radius associated with each fiber type. A lesser bend radius will result in radiation of Light (figure 2). This property can be used to extract light from fiber for nave eavesdropping. Normally, single mode fibers will not tolerate a minimum bend radius of less than 6.5 to 7.5 cm except some specially developed types. While a

978-1-4577-1149-9/11/$26.00 ©2011 IEEE

Optical fiber tapping: methods and precautions Iqbal, Pattinson and Kor (2011)

- **Bend it** far enough and light leaves the core. A coupler on the outside of an unbroken fibre is a tap.
- No cut, no splice, no loss of service, and a power drop small enough to look like a dirty connector.
- Countermeasures are **monitoring** and **encryption**, not the medium.

Fibre is harder to tap than copper.
Harder is not the same as **cannot**.

The clip-on coupler

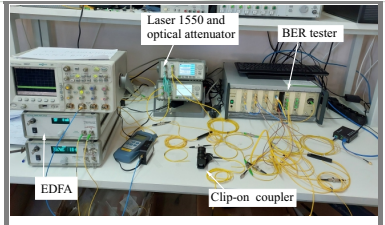

Fig. 3. Image of the test stand and specification of the clip-on coupler used.

The Threat of Optical Transmission Jamming
Kowalczyk et al. (2023), Fig. 3

- A clamp, a bend, and a detector: it is cheap equipment, sold for testing.
- It receives some of the photons (light) without cutting the fibre.
- The same coupler works the other way: interference by pumping photons into the wire.
- One device: intercept and interfere.

The clip-on coupler

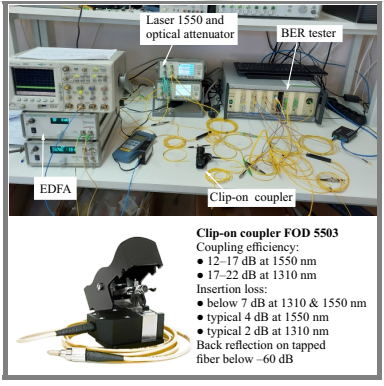

Fig. 3. Image of the test stand and specification of the clip-on coupler used.

The Threat of Optical Transmission Jamming
Kowalczyk et al. (2023), Fig. 3

- A clamp, a bend, and a detector: it is cheap equipment, sold for testing.
- It receives some of the photons (light) without cutting the fibre.
- The same coupler works the other way: interference by pumping photons into the wire.
- One device: intercept and interfere.

The clip-on coupler

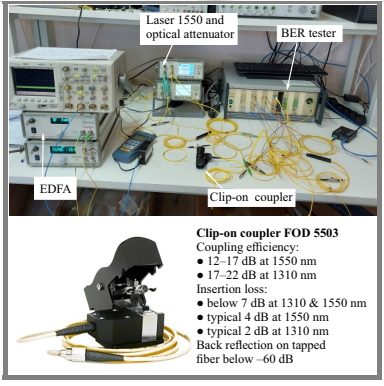

Fig. 3. Image of the test stand and specification of the clip-on coupler used.

The Threat of Optical Transmission Jamming
Kowalczyk et al. (2023), Fig. 3

- A clamp, a bend, and a detector: it is cheap equipment, sold for testing.
- It receives some of the photons (light) without cutting the fibre.
- The same coupler works the other way: interference by pumping photons into the wire.
- One device: intercept and interfere.

The clip-on coupler

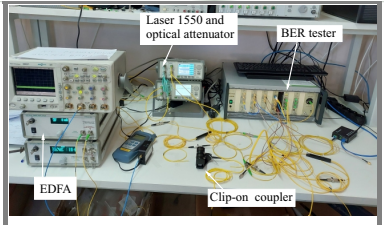

Fig. 3. Image of the test stand and specification of the clip-on coupler used.

The Threat of Optical Transmission Jamming
Kowalczyk et al. (2023), Fig. 3

- A clamp, a bend, and a detector: it is cheap equipment, sold for testing.
- It receives some of the photons (light) without cutting the fibre.
- The same coupler works the other way: interference by pumping photons into the wire.
- One device: intercept and interfere.

Delay in satellite networks



- Delay is geometry. A constellation buys latency back by being close, and pays for it in **numbers of satellites** and in handover.
- **An updated comparison of four low Earth orbit satellite constellation systems** Pachler et al. (2021)

Delay in satellite networks



- Delay is geometry. A constellation buys latency back by being close, and pays for it in **numbers of satellites** and in handover.
- **An updated comparison of four low Earth orbit satellite constellation systems** Pachler et al. (2021)

Free-space optics



An **FSO-R150** terminal: a telescope on a gimbal.



The **fine steering mirror** inside it, which holds the beam on the far end.

(FSO Instruments)

- **Free-space optics**: the same lasers and the same modulation, through air instead of glass.
- 8.21 Tbit/s over 160 m of open atmosphere.
- No cable to dig, no cable to cut, no cable to tap, and a beam that has to **stay pointed** through turbulence.

Terabits without fibres Xie and Yang (2025)

Free-space terabit/s coherent optical links Wang et al. (2025)

Free-space optics



An **FSO-R150** terminal: a telescope on a gimbal.



The **fine steering mirror** inside it, which holds the beam on the far end.

(FSO Instruments)

- **Free-space optics**: the same lasers and the same modulation, through air instead of glass.
- 8.21 Tbit/s over 160 m of open atmosphere.
- No cable to dig, no cable to cut, no cable to tap, and a beam that has to **stay pointed** through turbulence.

Terabits without fibres Xie and Yang (2025)

Free-space terabit/s coherent optical links Wang et al. (2025)

Free-space optics



An **FSO-R150** terminal: a telescope on a gimbal.



The **fine steering mirror** inside it, which holds the beam on the far end.

(FSO Instruments)

- **Free-space optics**: the same lasers and the same modulation, through air instead of glass.
- 8.21 Tbit/s over 160 m of open atmosphere.
- No cable to dig, no cable to cut, no cable to tap, and a beam that has to **stay pointed** through turbulence.

Terabits without fibres Xie and Yang (2025)

Free-space terabit/s coherent optical links Wang et al. (2025)

Overview

Part I. First hour

1. Theory & Practice
2. The physical layer
3. **The data plane**

Part II. Second hour

4. The control plane
5. Learning from 9/11
6. Regaining the cyber high ground

The five layers of the TCP/IP model

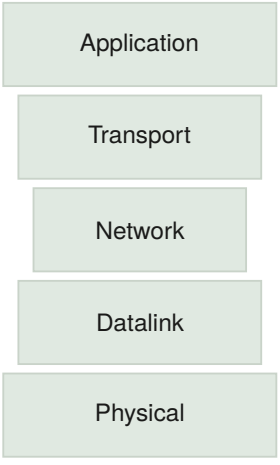

The five layers of the TCP/IP model

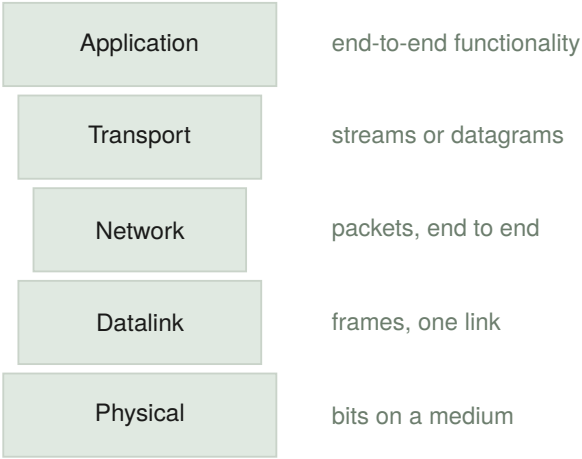

The five layers of the TCP/IP model



end-to-end functionality

a *pipe*

streams or datagrams

a *port*

packets, end to end

a *host* (IP)

frames, one link

an *interface* (MAC)

bits on a medium

A network of networks

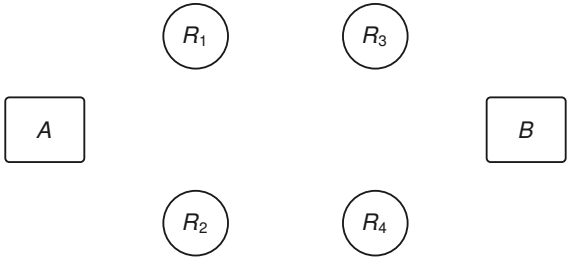

Two computers and four routers.

A network of networks

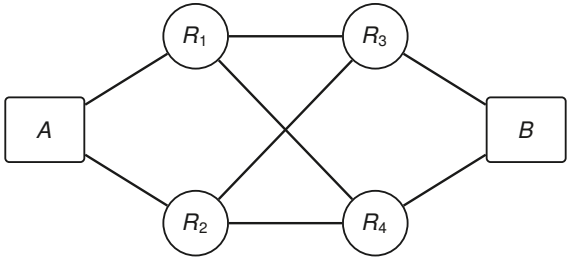

A cannot reach *B* directly. Something has to carry a packet between them, one leg at a time.

A network of networks

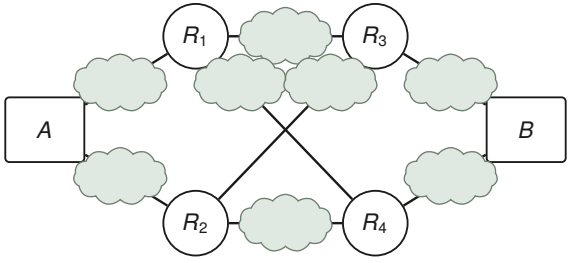

Every leg is a whole **network**, not a wire.

A network of networks

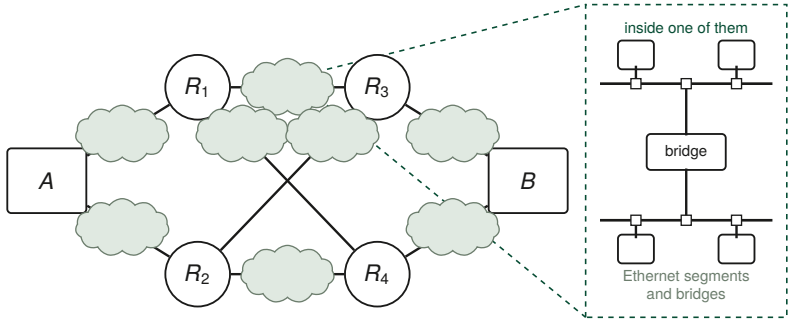

Inside each of them: **datalink layer**.

A network of networks

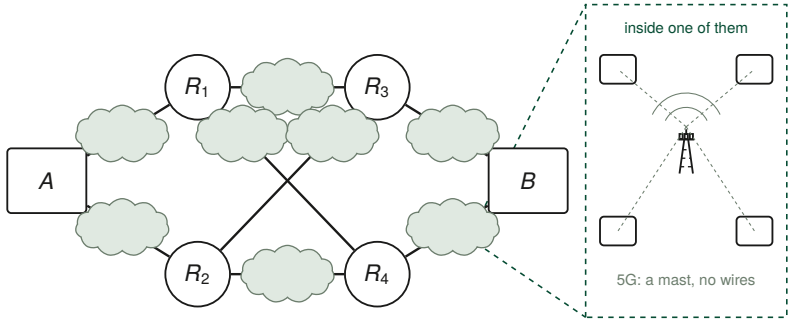

Examples: a cellular network, or a military datalink such as **Link-16**.

A network of networks

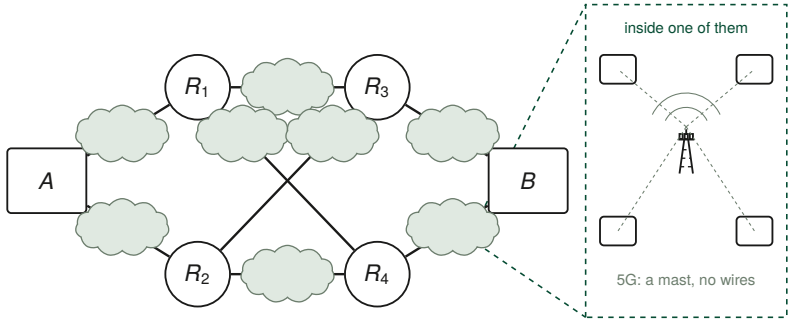

Abstraction: the network layer does not look inside. It only needs to move packet from one side to another.

Survivability of packet-switched networks

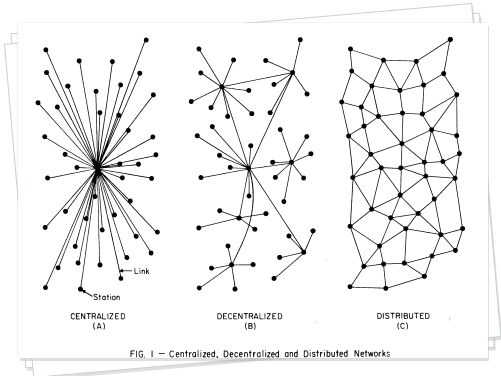

On distributed communications: I. Introduction to distributed communications networks Baran (1964)

- Highly **centralised** networks have single points of failure.
- Highly **distributed** networks have many independent paths between any two points.
- Packets *should* use **any** of them, every packet routed on its own.
  - **Forwarding tables**: each hop decides, but this needs cooperation.
  - **Packet-carrying forwarding state**: the path travels in the packet, and the sender chooses it.

Survivability of packet-switched networks

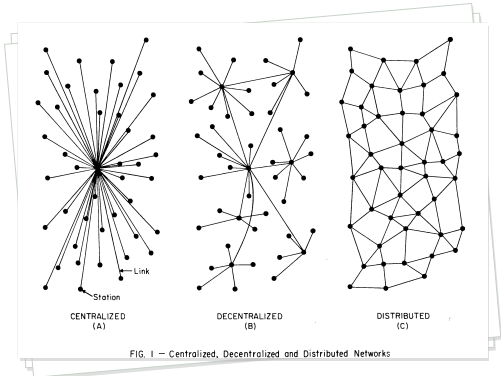

On distributed communications: I. Introduction to distributed communications networks Baran (1964)

- Highly **centralised** networks have single points of failure.
- Highly **distributed** networks have many independent paths between any two points.
- Packets *should* use **any** of them, every packet routed on its own.
  - **Forwarding tables**: each hop decides, but this needs cooperation.
  - **Packet-carrying forwarding state**: the path travels in the packet, and the sender chooses it.

Survivability of packet-switched networks

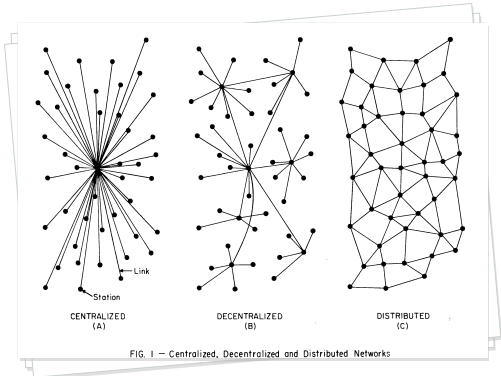

On distributed communications: I. Introduction to distributed communications networks Baran (1964)

- Highly **centralised** networks have single points of failure.
- Highly **distributed** networks have many independent paths between any two points.
- Packets *should* use **any** of them, every packet routed on its own.
  - **Forwarding tables**: each hop decides, but this needs cooperation.
  - **Packet-carrying forwarding state**: the path travels in the packet, and the sender chooses it.

Survivability of packet-switched networks

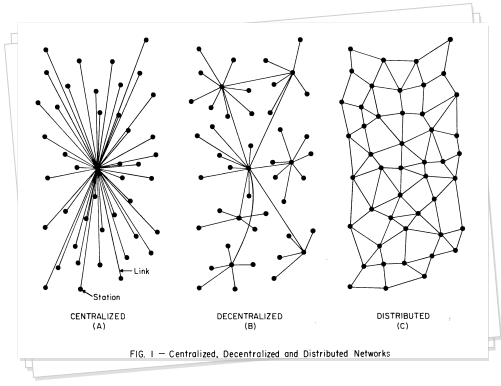

On distributed communications: I. Introduction to distributed communications networks Baran (1964)

- Highly **centralised** networks have single points of failure.
- Highly **distributed** networks have many independent paths between any two points.
- Packets *should* use **any** of them, every packet routed on its own.
  - **Forwarding tables**: each hop decides, but this needs cooperation.
  - **Packet-carrying forwarding state**: the path travels in the packet, and the sender chooses it.

Survivability of packet-switched networks

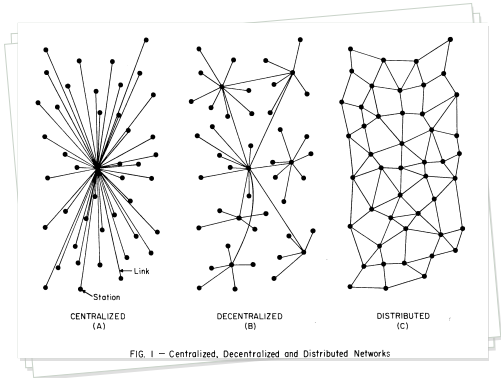

On distributed communications: I. Introduction to distributed communications networks Baran (1964)

- Highly **centralised** networks have single points of failure.
- Highly **distributed** networks have many independent paths between any two points.
- Packets *should* use **any** of them, every packet routed on its own.
  - **Forwarding tables**: each hop decides, but this needs cooperation.
  - **Packet-carrying forwarding state**: the path travels in the packet, and the sender chooses it.

An address for every destination

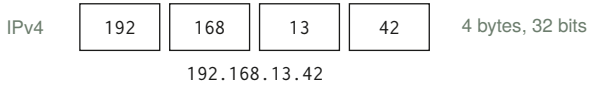

- The address names a **destination** to deliver to, not a piece of hardware.
- Every computer can hold **several**: e.g. one per network it is attached to.
- Assigned **statically**, by whoever runs the network.
- Or **dynamically**, handed out when a computer joins a new network.

An address for every destination

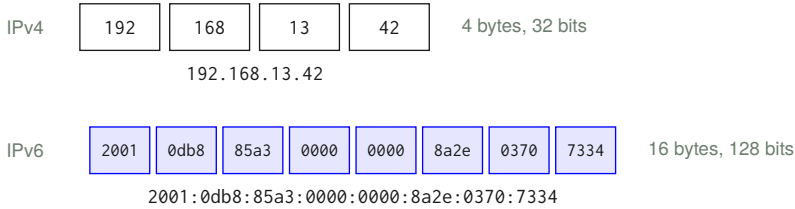

- The address names a **destination** to deliver to, not a piece of hardware.
- Every computer can hold **several**: e.g. one per network it is attached to.
- Assigned **statically**, by whoever runs the network.
- Or **dynamically**, handed out when a computer joins a new network.

An address for every destination

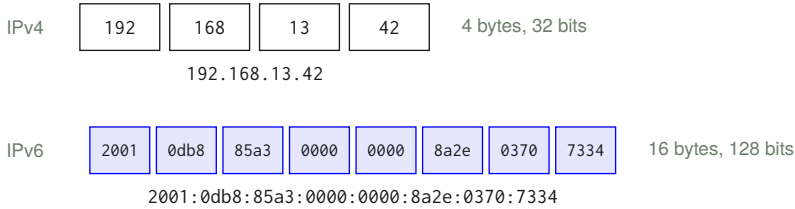

- The address names a **destination** to deliver to, not a piece of hardware.
- Every computer can hold **several**: e.g. one per network it is attached to.
- Assigned **statically**, by whoever runs the network.
- Or **dynamically**, handed out when a computer joins a new network.

An address for every destination

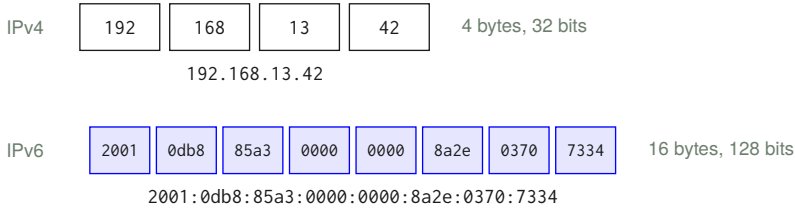

- The address names a **destination** to deliver to, not a piece of hardware.
- Every computer can hold **several**: e.g. one per network it is attached to.
- Assigned **statically**, by whoever runs the network.
- Or **dynamically**, handed out when a computer joins a new network.

An address for every destination

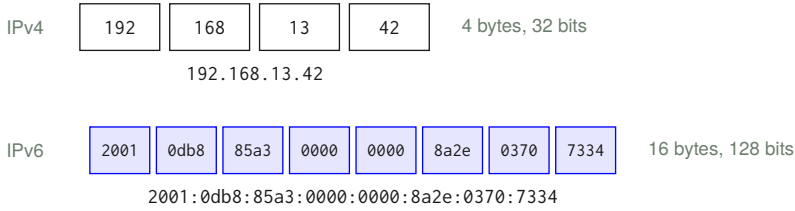

- The address names a **destination** to deliver to, not a piece of hardware.
- Every computer can hold **several**: e.g. one per network it is attached to.
- Assigned **statically**, by whoever runs the network.
- Or **dynamically**, handed out when a computer joins a new network.

An address for every destination

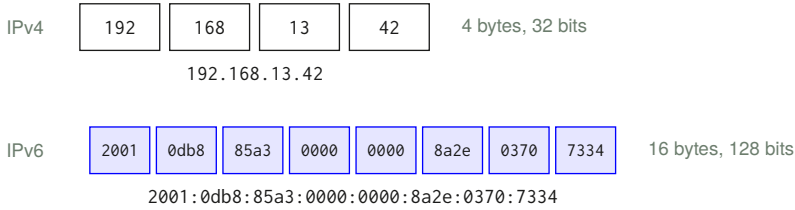

- The address names a **destination** to deliver to, not a piece of hardware.
- Every computer can hold **several**: e.g. one per network it is attached to.
- Assigned **statically**, by whoever runs the network.
- Or **dynamically**, handed out when a computer joins a new network.

Prefix and mask



One address, 32 bits.

Prefix and mask

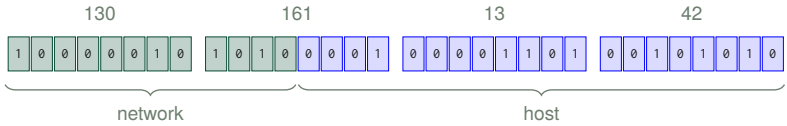

`130.160.0.0/12`: 12 bits of network, $2^{20} - 2 = 1\,048\,574$ usable hosts

A **prefix** length cuts it in two: left of the cut names the **network**, right of it a **host**.

Prefix and mask

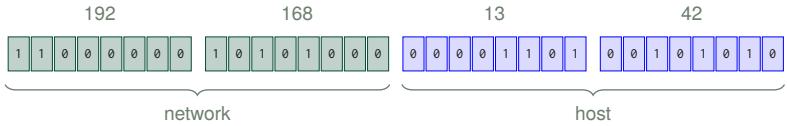

`192.168.0.0/16`: 16 bits of network, $2^{16} - 2 = 65\,534$ usable hosts

A longer prefix, a smaller network.

Prefix and mask

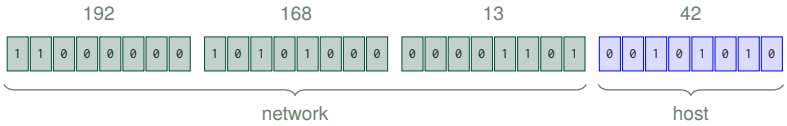

`192.168.13.0/24`: 24 bits of network, $2^8 - 2 = 254$ usable hosts

`192.168.13.0` host bits all zero: the network itself
`192.168.13.255` host bits all one: broadcast

Two addresses in every block are spent before anyone gets one.

Prefix and mask

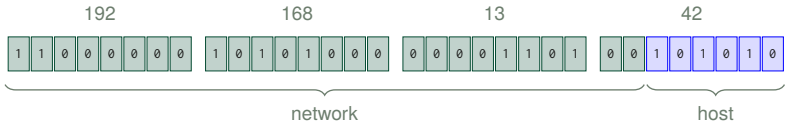

`192.168.13.0/26`: 26 bits of network, $2^6 - 2 = 62$ usable hosts

The cut need not fall on a byte, which is the point of writing the length rather than four numbers.

Prefix and mask

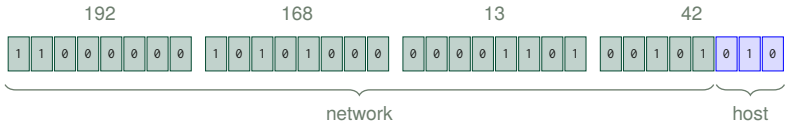

`192.168.13.40/29`: 29 bits of network, $2^3 - 2 = 6$ usable hosts

Six usable addresses: the smallest block worth giving to a network with more than two machines on it.

Network gateways



Two machines, two routers, three networks between them.

Network gateways

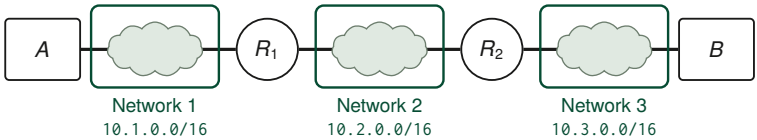

Each network has a block of addresses of its own.

Network gateways

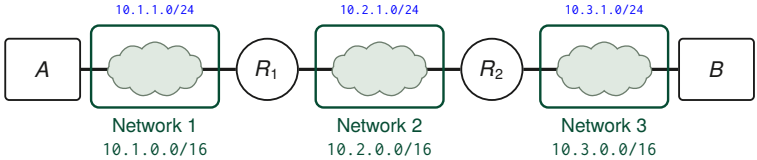

One subnet has been taken out of each.

Network gateways

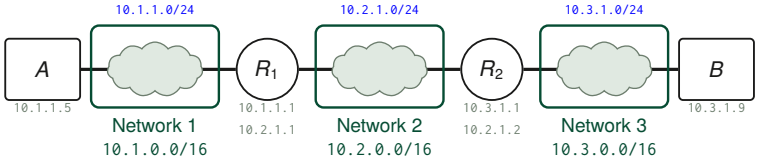

Every computer is assigned one or more address.

Network gateways

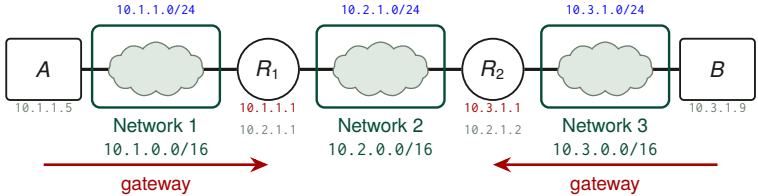

Each hop sends traffic to a next hop.

A packet crossing two routers



Sending a packet from *A* to *B*.

A packet crossing two routers





`10.3.1.9` is not on `10.1.1.0/24`, so A sends to its **default gateway**, R_1.

A packet crossing two routers





On the first network, addressed to *B*.

A packet crossing two routers





In R_1, which is not B.

A packet crossing two routers

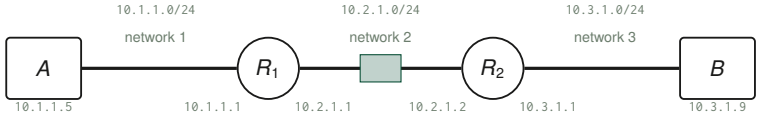



So R_1 **forwards** it.

A packet crossing two routers





In R_2.

A packet crossing two routers





Which **forwards** it onto `10.3.1.0/24`, one of its own networks.

A packet crossing two routers





In *B*: **delivered**.

A packet crossing two routers



The IP header is never changed.

The five layers of the TCP/IP model

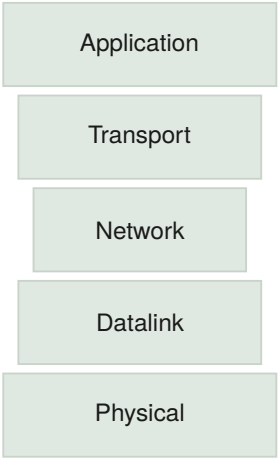

The five layers of the TCP/IP model

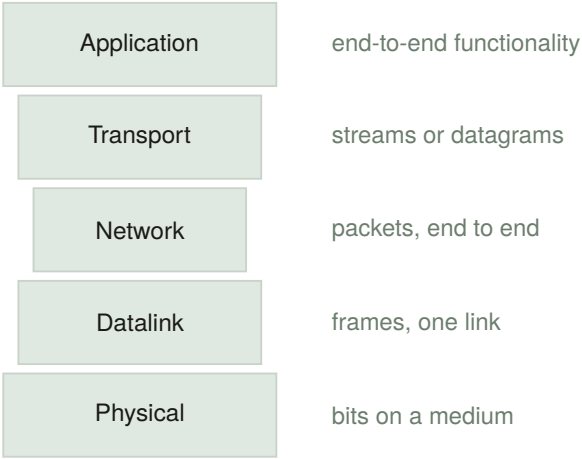

The five layers of the TCP/IP model

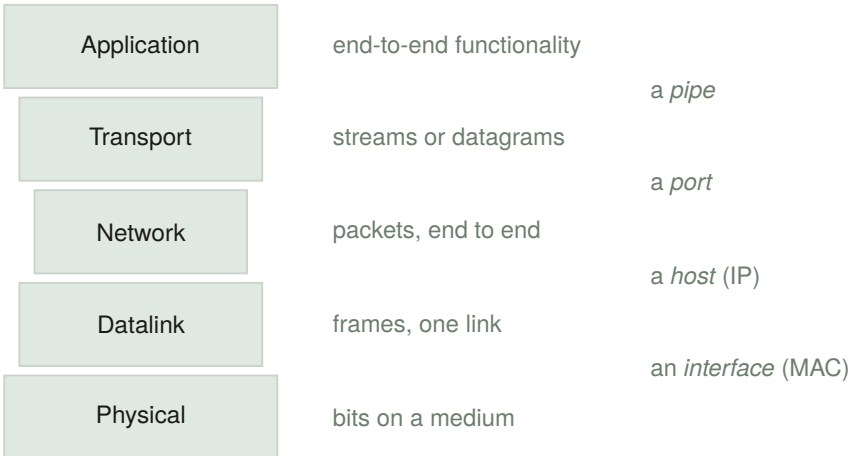

The same crossing, with the frames

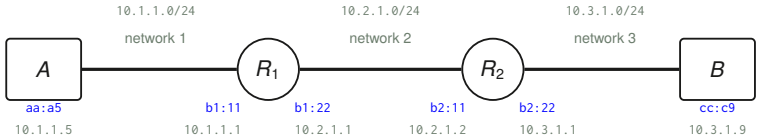

on the wire:

The same three networks, with the hardware addresses on them.

The same crossing, with the frames

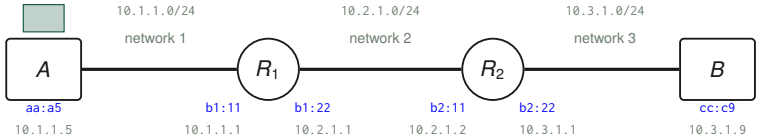

A writes the packet: its own address, and *B*'s. No frame yet.

The same crossing, with the frames

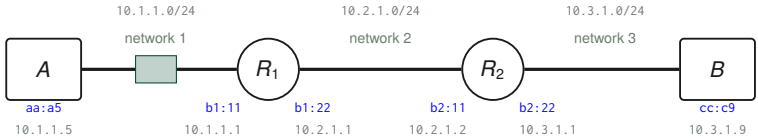

A asked ARP for `10.1.1.1` and got `b1:11`, and builds a frame around the packet. That is the frame's destination; *B* is the packet's.

The same crossing, with the frames

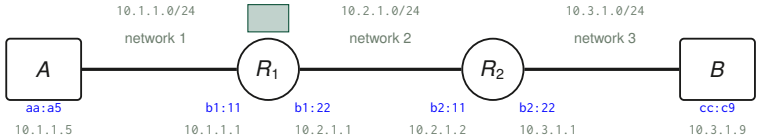

on the wire:

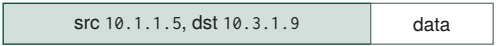

R_1 throws the frame away and keeps the packet. For this moment there is no frame at all.

The same crossing, with the frames

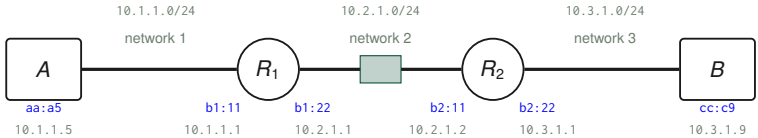

A new frame, on the next network. R_1 asked via ARP what the MAC address of R_2 is, and got `b2:11`.

The same crossing, with the frames

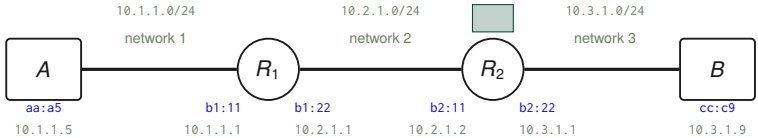

on the wire:

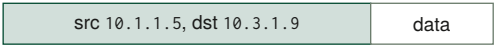

R_2 does the same again.

The same crossing, with the frames

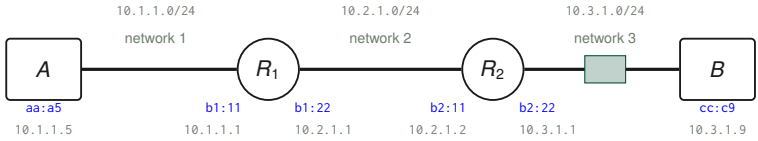

The last frame is addressed to *B*'s own hardware.

The same crossing, with the frames

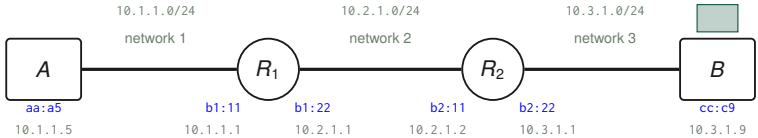

on the wire:

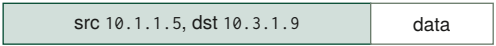

packet header

Delivered.

The same crossing, with the frames

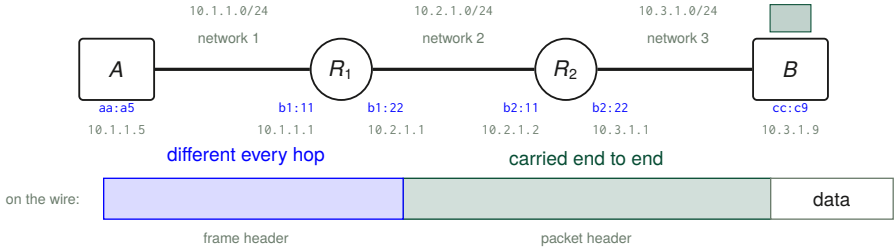

Forwarding, table by table

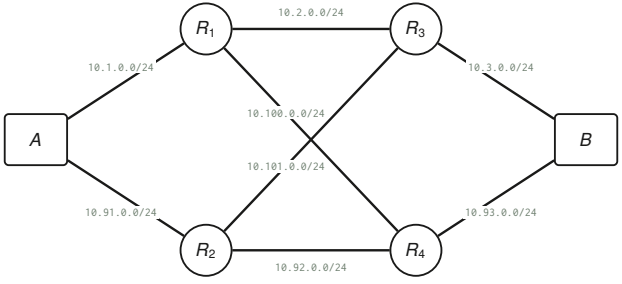

Eight networks. *A* wants to reach `10.93.0.2`, which is *B*, on the network at the bottom right.

Forwarding, table by table

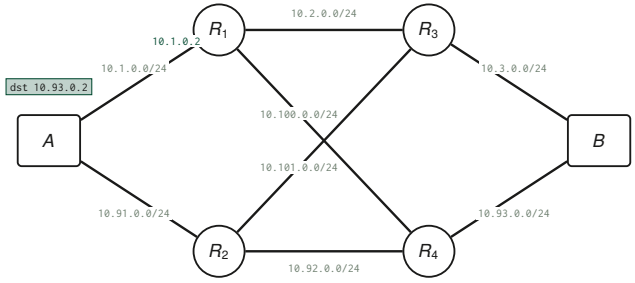

A's forwarding table

destination	next hop
10.1.0.0/24	direct
10.91.0.0/24	direct
0.0.0.0/0	**10.1.0.2**

A is on `10.1.0.0/24` and `10.91.0.0/24`. The destination is on neither, so it goes to a gateway: `10.1.0.2`, which is R_1.

Forwarding, table by table

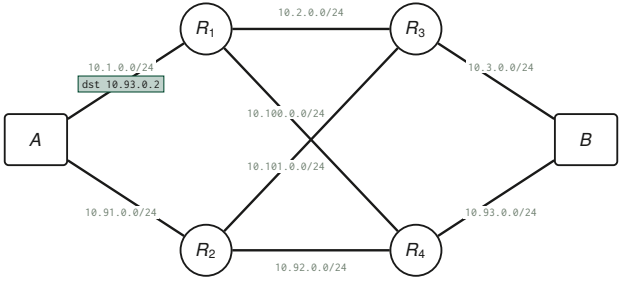

On the wire.

Forwarding, table by table

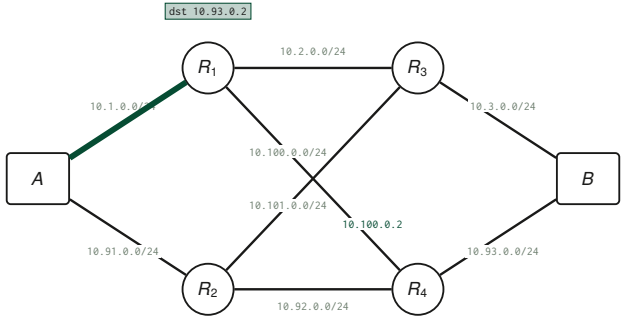

R_1's forwarding table

destination	next hop
10.1.0.0/24	direct
10.2.0.0/24	direct
10.100.0.0/24	direct
10.93.0.0/24	**10.100.0.2**
0.0.0.0/0	10.2.0.2

R_1 looks the destination up. Three of its networks are directly connected; the fourth row says `10.93.0.0/24` is reached through `10.100.0.2`.

Forwarding, table by table

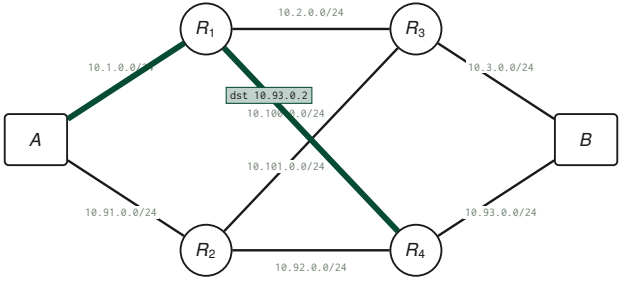

Which is R_4, across the diagonal.

Forwarding, table by table

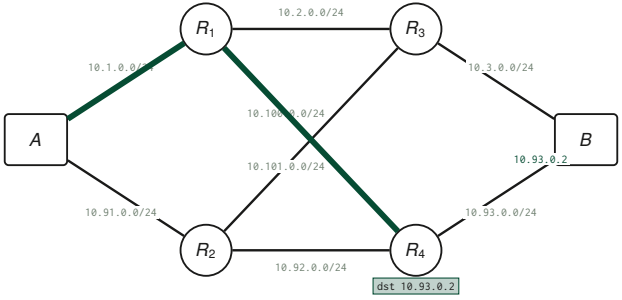

R_4's forwarding table

destination	next hop
10.92.0.0/24	direct
10.93.0.0/24	**direct**
10.100.0.0/24	direct
0.0.0.0/0	10.92.0.1

R_4 finds `10.93.0.0/24` directly connected:
no next hop, it is on that network itself.

Forwarding, table by table

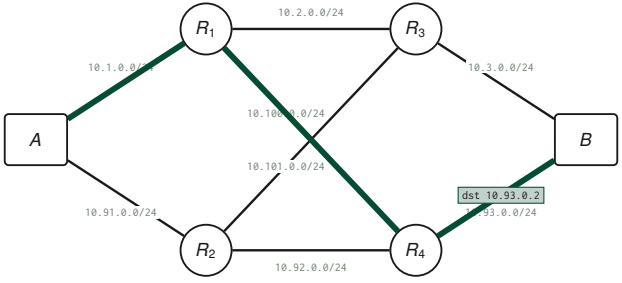

So the last hop is a delivery, not a forward.

Forwarding, table by table

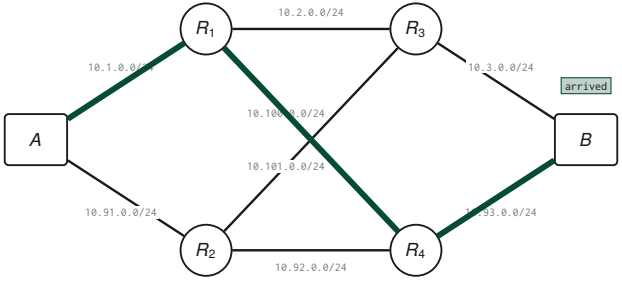

Two tables consulted, one packet delivered.
Neither router knew the whole path.

Longest prefix match

a packet with destination `10.4.7.9` arrives

prefix	next hop	matches?
`0.0.0.0/0`	the default route	
`10.0.0.0/8`	R_a	
`10.4.0.0/16`	R_b	
`10.4.7.0/24`	R_c	

A router does not look for the address. It looks for the longest prefix that covers it.

Longest prefix match

a packet with destination `10.4.7.9` arrives

prefix	next hop	matches?
`0.0.0.0/0`	the default route	yes, 0 bits
`10.0.0.0/8`	R_a	
`10.4.0.0/16`	R_b	
`10.4.7.0/24`	R_c	

Everything matches the default route.

Longest prefix match

a packet with destination `10.4.7.9` arrives

prefix	next hop	matches?
`0.0.0.0/0`	the default route	yes, 0 bits
`10.0.0.0/8`	R_a	yes, 8 bits
`10.4.0.0/16`	R_b	
`10.4.7.0/24`	R_c	

And this one, for eight bits.

Longest prefix match

a packet with destination **`10.4.7.9`** arrives

prefix	next hop	matches?
`0.0.0.0/0`	the default route	yes, 0 bits
`10.0.0.0/8`	R_a	yes, 8 bits
`10.4.0.0/16`	R_b	yes, 16 bits
`10.4.7.0/24`	R_c	

And this one, for sixteen.

Longest prefix match

a packet with destination `10.4.7.9` arrives

prefix	next hop	matches?
`0.0.0.0/0`	the default route	yes, 0 bits
`10.0.0.0/8`	R_a	yes, 8 bits
`10.4.0.0/16`	R_b	yes, 16 bits
`10.4.7.0/24`	R_c	yes, 24 bits

And this one, for twenty-four.

Longest prefix match

a packet with destination `10.4.7.9` arrives

prefix	next hop	matches?
`0.0.0.0/0`	the default route	yes, 0 bits
`10.0.0.0/8`	R_a	yes, 8 bits
`10.4.0.0/16`	R_b	yes, 16 bits
`10.4.7.0/24`	R_c	yes, 24 bits

the most specific match wins: forward to R_c

Four rows matched. The table is a hierarchy, and the deepest entry decides.

Overview

Part I. First hour

1. Theory & Practice
2. The physical layer
3. The data plane

Part II. Second hour

4. **The control plane**
5. Learning from 9/11
6. Regaining the cyber high ground

Exponential growth and deployment

Table 2.1 **Traffic on Internet Backbones in the United States. Data are estimated traffic in terabytes (TB) during December of that year**

Year	*TB/month*
1990	1.0
1991	2.0
1992	4.4
1993	8.3
1994	16.3
1995	?
1996	1,500
1997	2,500–4,000
1998	5,000–8,000
1999	10,000–16,000
2000	20,000–35,000



Growth of the Internet Coffman and Odlyzko (2002)

Exponential growth and deployment

Table 2.1 **Traffic on Internet Backbones in the United States. Data are estimated traffic in terabytes (TB) during December of that year**

Year	*TB/month*
1990	1.0
1991	2.0
1992	4.4
1993	8.3
1994	16.3
1995	?
1996	1,500
1997	2,500–4,000
1998	5,000–8,000
1999	10,000–16,000
2000	20,000–35,000

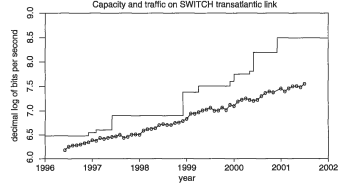

Fig. 2.2 Capacity of link between the Swiss SWITCH network and the United States and traffic on it toward Switzerland.

Growth of the Internet Coffman and Odlyzko (2002)

Early 2000s: network of networks

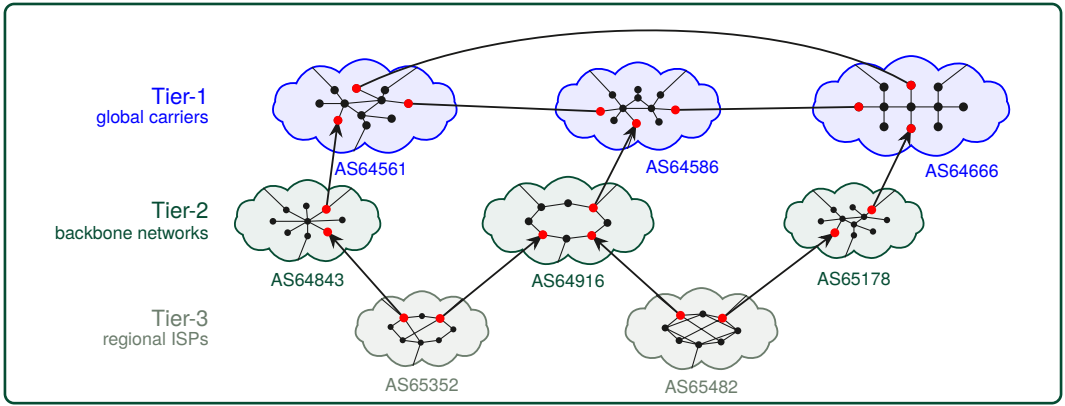

AS numbers are examples (RFC 6996). Real ones are assigned by IANA through the five Regional Internet Registries (RIRs): ARIN, RIPE NCC, APNIC, LACNIC and AFRINIC.

Autonomous Systems

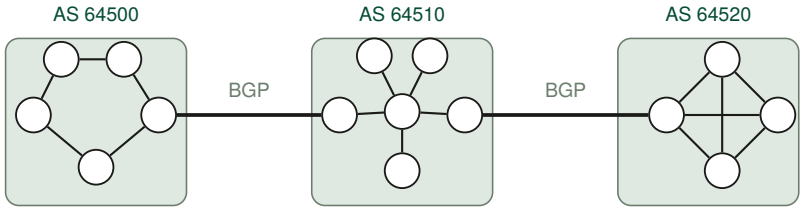

- An **autonomous system** is a network under one routing policy.
- Inside it, an **internal** routing protocol computes shortest paths. Outside them, we need coordination using *de facto* protocol: **Border Gateway Protocol**.
- BGP announces **reachability**, not distance, and every AS decides for itself what to believe.

Autonomous Systems

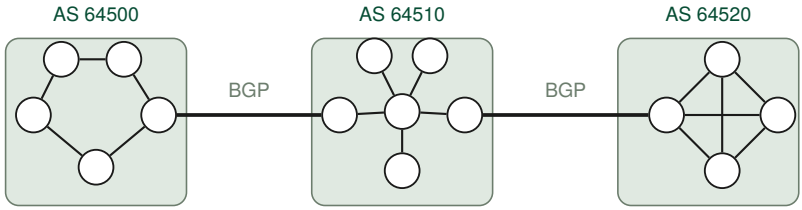

- An **autonomous system** is a network under one routing policy.
- Inside it, an **internal** routing protocol computes shortest paths. Outside them, we need coordination using *de facto* protocol: **Border Gateway Protocol**.
- BGP announces **reachability**, not distance, and every AS decides for itself what to believe.

Autonomous Systems

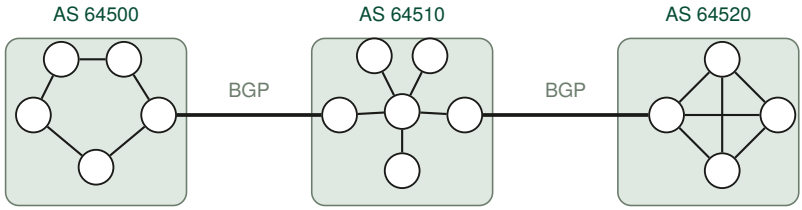

- An **autonomous system** is a network under one routing policy.
- Inside it, an **internal** routing protocol computes shortest paths. Outside them, we need coordination using *de facto* protocol: **Border Gateway Protocol**.
- BGP announces **reachability**, not distance, and every AS decides for itself what to believe.

BGP: announcing what you can reach

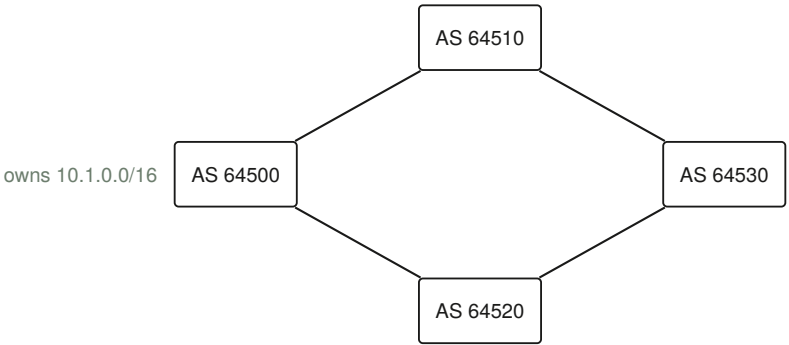

Four autonomous systems. Each own blocks of addresses.

BGP: announcing what you can reach

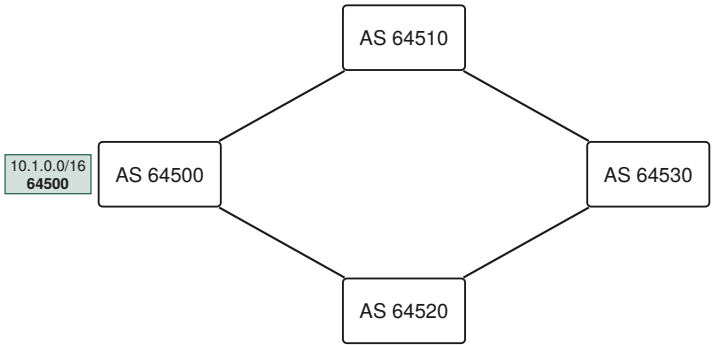

It announces the block to its neighbours: *traffic for this block goes to me*. The announcement carries the block and a path.

BGP: announcing what you can reach

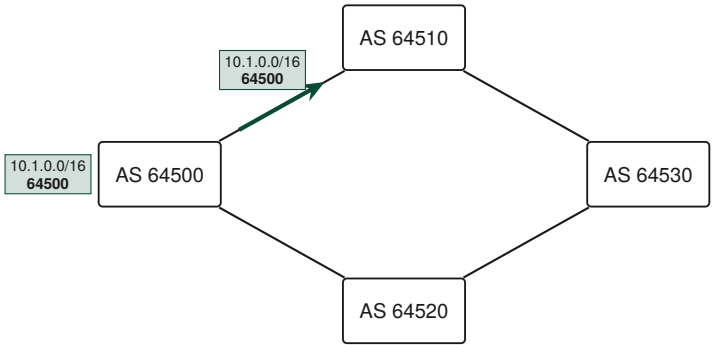

The announcement travels on the BGP sessions.

BGP: announcing what you can reach

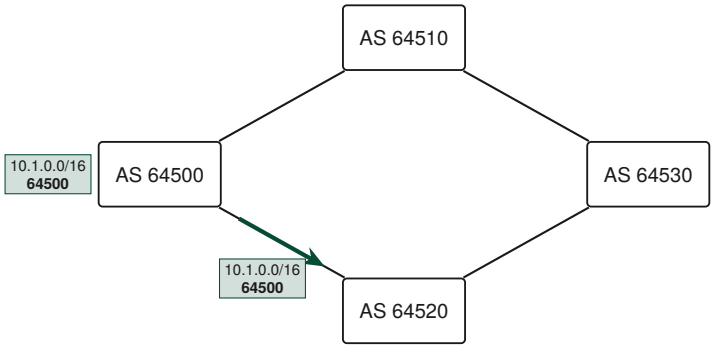

And on the other.

BGP: announcing what you can reach

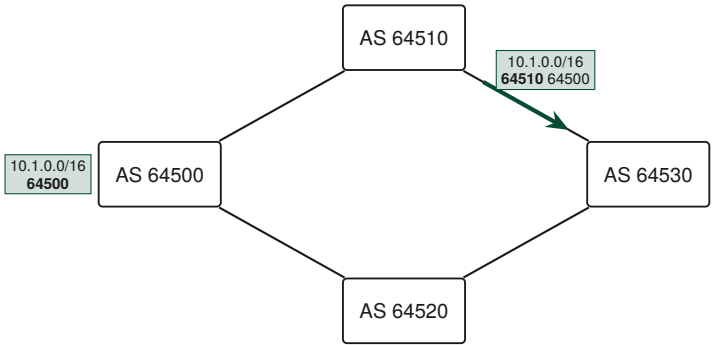

Passing it on, an autonomous system puts its own number on the **front** of the path. So the path lists, right to left, everyone the announcement has been through.

BGP: announcing what you can reach

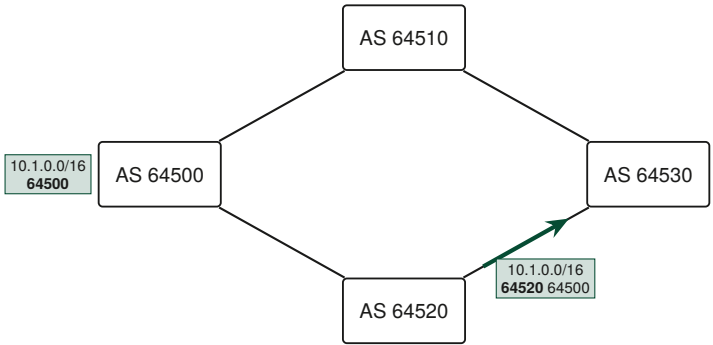

The other one does the same.

BGP: announcing what you can reach

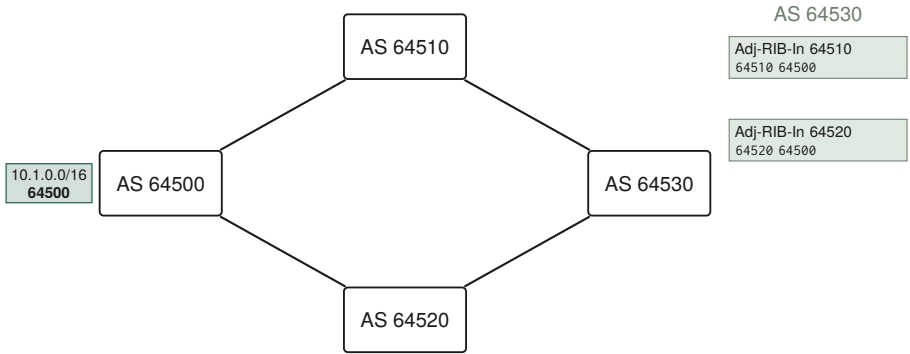

So 64530 hears the block twice. What each neighbour said is kept on its own, in an **Adjacent Routing Information Base**, one per adjacent autonomous system. Nothing is thrown away.

BGP: announcing what you can reach

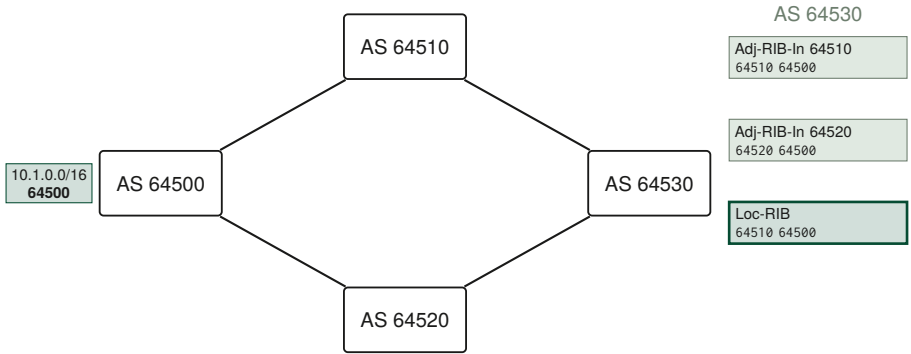

Policy picks one and puts it in the **Local Routing Information Base**. This is what forwarding tables use.

BGP: announcing what you can reach

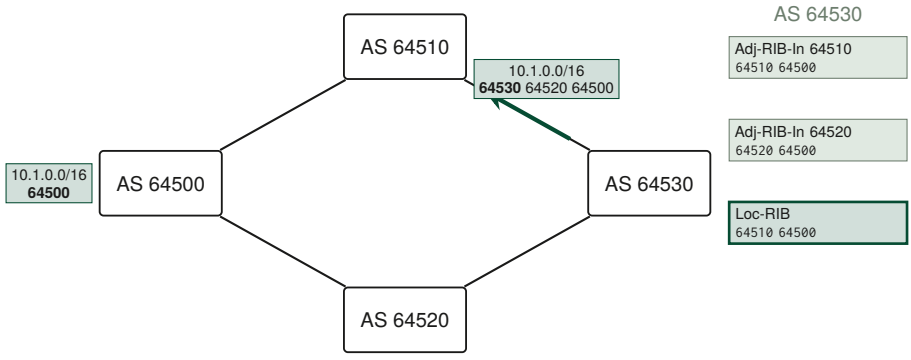

It may still pass on updates, with its own number in front.

BGP: announcing what you can reach

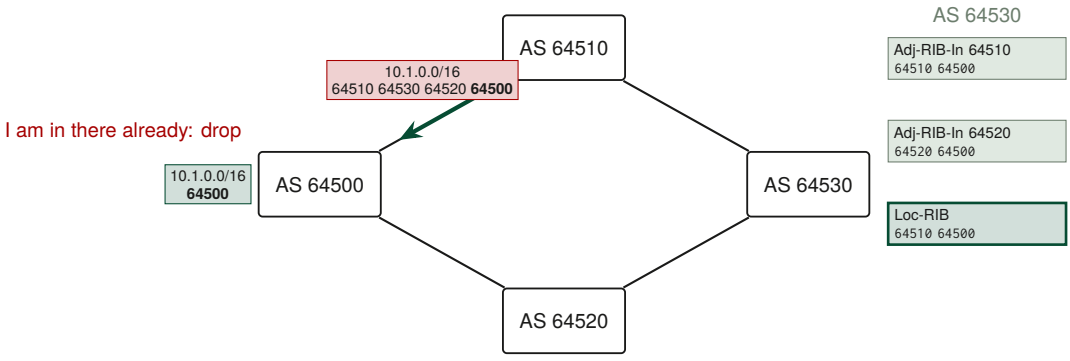

Which comes back to where it started. 64500 finds its own number in the path, so it is dropped: loop prevention.

BGP: announcing what you can reach

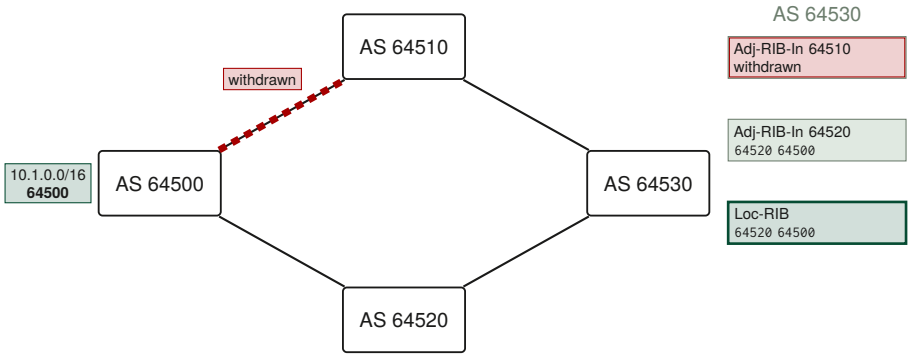

A link goes. The announcement is **withdrawn**. The other neighbour's route was there all along, so the Local Routing Information Base is recomputed from what is still held. That is the failover, and it takes **one to three minutes**.

Overview

Part I. First hour

1. Theory & Practice
2. The physical layer
3. The data plane

Part II. Second hour

4. The control plane
5. **Learning from 9/11**
6. Regaining the cyber high ground



- Biggest effect probably from collapse of Building 7
- Several ISPs, including AT&T, Local Systems, Genuity, Verizon, and WorldCom, had points of presence (POPs)
- Many optical fibres (up to 1300 for some networks) and fibre ring breaks
- Electrical power outages, insufficient backup

(Image: James Nachtwey)

The Internet under crisis conditions

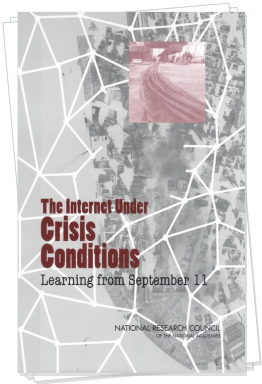

The Internet under crisis conditions: Learning from September 11 National Research Council (2003)

The Internet under crisis conditions

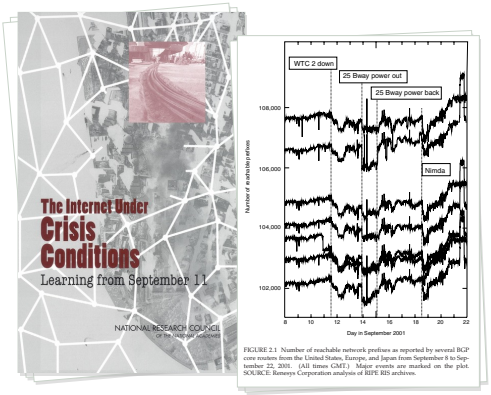

The Internet under crisis conditions: Learning from September 11 National Research Council (2003)

The Internet under crisis conditions

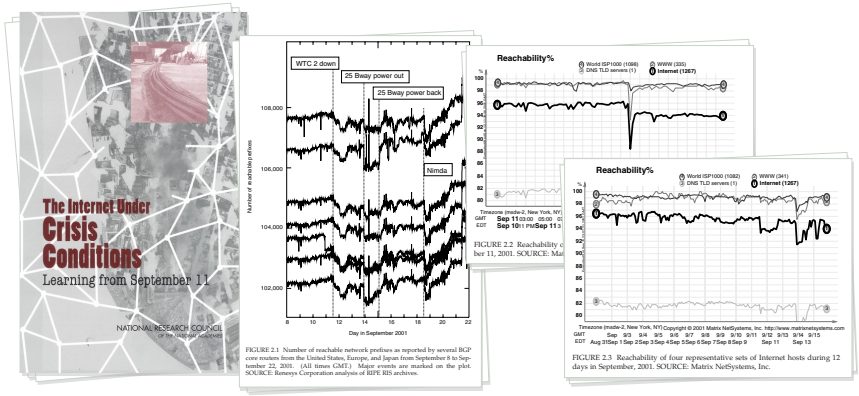

FIGURE 2.1 Number of reachable network prefixes as reported by several BGP core routers from the United States, Europe, and Japan from September 8 to September 22, 2001. (All times GMT.) Major events are marked on the plot. SOURCE: Renesys Corporation analysis of RIPE RIS archives.

FIGURE 2.2 Reachability c
ber 11, 2001. SOURCE: Mat

FIGURE 2.3 Reachability of four representative sets of Internet hosts during 12 days in September, 2001. SOURCE: Matrix NetSystems, Inc.

The Internet under crisis conditions: Learning from September 11 National Research Council (2003)

The Internet under crisis conditions

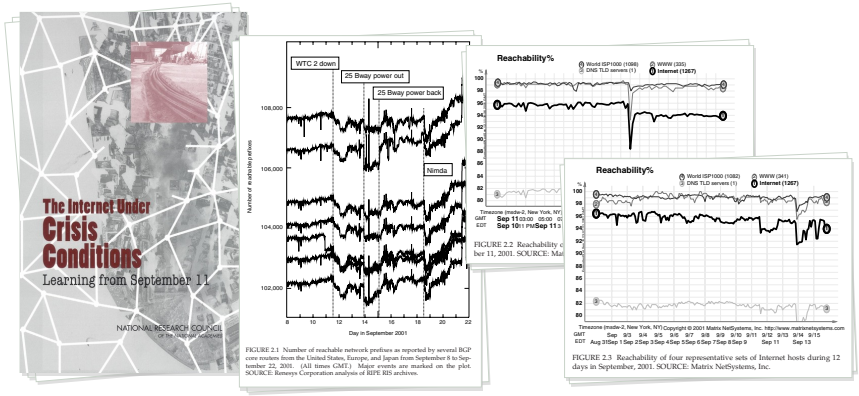

The Internet under crisis conditions: Learning from September 11 National Research Council (2003)

"Network operators often joke that a single misplaced comma in an appropriate configuration file could take down the Internet."

Routing between networks

- **Border Gateway Protocol (BGP)**
- February, 2008: Pakistan Telecom took down YouTube

  “Because the bogus route [...] was always more preferred [than] the legitimate route, and **within minutes**, at least two-thirds of the Internet was sending its YouTube traffic to Pakistan.”
- Attacks: route hijacks, route leaks, ...
- Defences: soBGP (2003), RPKI (2008), BGPsec (2017), ASPA (2025)
- Slow uptake

practice

Routing security incidents can still slip past deployed security defenses.

BY SHARON GOLDBERG

Why Is It Taking So Long to Secure Internet Routing?

Why is it taking so long to secure internet routing? Goldberg (2014)

Routing between networks

- **Border Gateway Protocol (BGP)**
- February, 2008: Pakistan Telecom took down YouTube

  “Because the bogus route [...] was always more preferred [than] the legitimate route, and **within minutes**, at least two-thirds of the Internet was sending its YouTube traffic to Pakistan.”
- Attacks: route hijacks, route leaks, ...
- Defences: soBGP (2003), RPKI (2008), BGPsec (2017), ASPA (2025)
- Slow uptake

practice

Routing security incidents can still slip past deployed security defenses.

BY SHARON GOLDBERG

Why Is It Taking So Long to Secure Internet Routing?

Why is it taking so long to secure internet routing? Goldberg (2014)

Routing between networks

- **Border Gateway Protocol (BGP)**
- February, 2008: Pakistan Telecom took down YouTube

  “Because the bogus route [...] was always more preferred [than] the legitimate route, and **within minutes**, at least two-thirds of the Internet was sending its YouTube traffic to Pakistan.”
- Attacks: route hijacks, route leaks, ...
- Defences: soBGP (2003), RPKI (2008), BGPsec (2017), ASPA (2025)
- Slow uptake

practice

Routing security incidents can still slip past deployed security defenses.

BY SHARON GOLDBERG

Why Is It Taking So Long to Secure Internet Routing?

Why is it taking so long to secure internet routing? Goldberg (2014)

Routing between networks

- **Border Gateway Protocol (BGP)**
- February, 2008: Pakistan Telecom took down YouTube

  “Because the bogus route [...] was always more preferred [than] the legitimate route, and **within minutes**, at least two-thirds of the Internet was sending its YouTube traffic to Pakistan.”
- Attacks: route hijacks, route leaks, ...
- Defences: soBGP (2003), RPKI (2008), BGPsec (2017), ASPA (2025)
- Slow uptake


practice

Routing security incidents can still slip past deployed security defenses.

BY SHARON GOLDBERG

Why Is It Taking So Long to Secure Internet Routing?

Why is it taking so long to secure internet routing? Goldberg (2014)

Awaiting Cyber 9/11



AWAITING CYBER 9/11

By CLIFFORD S. MAGEE

Enemies no longer need to launch missiles or fly airplanes into buildings to attack the United States. A new weapon has been introduced into the fight. Cyberspace has no physical borders or rules, costs little, and has monstrous potential. The weapon is cyber warfare. The Nation's security, economy, and critical infrastructure are under cyber attack. Many of the attacks originate from countries such as China and Russia, while others are from nonstate actors such as terrorist organizations, criminal gangs, teenage hackers, and anarchists. To protect America financial secrets, military plans, intelligence operations, power supplies, intellectual property, and military communications, the U.S. Government needs to designate the Department of Defense (DOD) as the lead organization in preventing, detecting, and defending against cyber attacks.

In 2009, the Wall Street Journal reported that Chinese hackers had gained access to the U.S. electric power grid and created secret openings.[1] These were no reconnoiter value to the average hacker. The only reason for a virus was to use intelligence value that would justify cyber espionage.[2] The only reason to penetrate the grid was to have the capability to conduct widespread cyberattacks on the nation's asymmetrical cyberwar.[3] The Chinese had created a capability that could create power outages across the United States and possibly damage the electrical grid to the point of failure. The cyber-wars assumes their systems were compromised until the intrusions were detected by the U.S. Intelligence Community.[4] What would the U.S. Government have done if it discovered that China had been laying explosive charges throughout the national electrical grid system?

The threats posed to the cyber domain are, in fact, an existential danger to the Nation. Therefore, the DOD must develop and train organizations with the capabilities or authorities to ensure cyber security for the public and private sectors. To develop this capability, the Nation needs to undergo a paradigm shift on how it views the cyber domain.

The Cyber Domain

In 2010, Deputy Secretary of Defense William Lynn III in *Foreign Affairs* Magazine stated that the Royal Navy was necessary because it provided sea power to protect the goods and services that travel on the seas.[5] The cyber domain is analogous to the seas. The Nation has its communications and trade were extraordinarily important for the security and prosperity of Britain. Today, the security and prosperity of the United States depends on cyberspace. The cyber domain is vulnerable to attack. Signals and information can be intercepted, interrupted, and exploited. The Nation must develop a strategy to defend the cyber domain similar to the strategy it has developed for defending land, sea, and air dominance.

Integrated DOD efforts defend these domains. Defending air, trade and commercial routes is not the responsibility of the Federal Aviation Administration or the Coast Guard. Likewise, it is the responsibility of the Defense Department.[6] Similarly, Marcel Lesoi is not responsible for defense of the sea domain, but as the cyber domain, every American company is responsible for it even defense without support from the Government. The U.S. Government does not have a lead organization to defend all government networks from attacks, much less assist with defending the private sector.[7] The Nation must develop the responsibility of defending the cyber domain with assistance from the Department of Homeland Security, Intelligence Community, and private sector.

The DOD must develop an active layered cyber defense with offensive and defensive capabilities. Currently, most cyber defensive strategy rely on firewalls to block attacks. This method is similar to the past. World War II Germany had a similar belief. They felt that it was an expensive defensive measure designed to keep the Germans out of France. But in 1940 the method did not work. The Maginot Line was a single capability; the strategy of the law of the land. The cyber domain requires a layered defensive capability model for cyber defense. To avoid a cyber Maginot Line, the United States needs layered, integrated defenses as well as an offensive capability.

Defining the Battlespace

The cyber domain has been created in a relatively short time. It is not defined by land of location as other battle domains. Land and sea domains have had thousands of years of discussion to create generally accepted definitions. The air domain has

Major Clifford S. Magee, USAR, is a Communications Officer currently serving with 12th Marine Expeditionary Unit and a DOD graduate from the Marine Corps Command and Staff College.

- “To defend cyberspace, the United States needs to develop its **situational awareness** of the cyber domain.”
- “Recent technological innovations such as ‘cloud computing’ must be leveraged to create a more secure, reliable, and cost-effective cyberspace. For example, **collapsing the 15,000 disparate DOD networks** to a cloud environment would provide it the ability to react to threats at ‘net-speed.’ ”



Awaiting cyber 9/11 Magee (2013)

Society-technology dependence

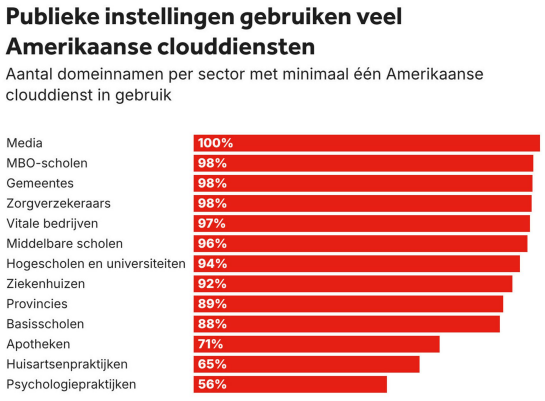

(Source: NOS)

- American technology: great! But is ‘cloud computing’ always such a good idea?

Internet Exchange Points (IXPs)

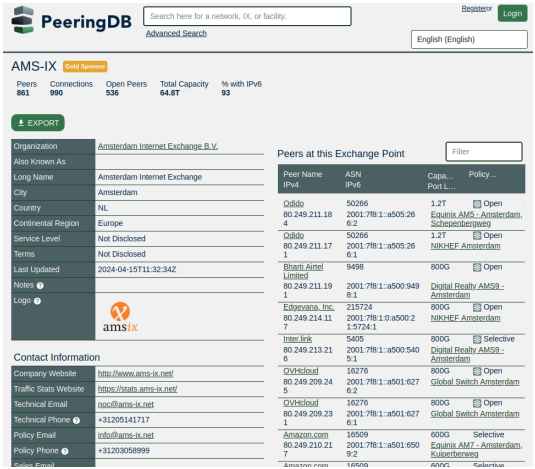

(Source: PeeringDB, 2026)

- 2007 cyberattack on Estonia
- IXPs make ‘traffic washing’ easy
- 2008-2016 the number of IXPs tripled, traffic comparable to Tier-1

The elusive internet flattening: 10 years of IXP growth Böttger et al. (2018)

Internet eXchange points: current challenges and new opportunities
Di Lallo (2018)

One building, one datalink layer



source

AMS-IX: an *internet exchange*. Networks that would otherwise pay someone to carry traffic between them plug into the same switching fabric instead.

- ≈ 15 Tbit/s at peak
- ≈ 880 members on ≈ 1500 ports
- 1, 10, 100 Gbit/s Ethernet, aggregated where one port is not enough
- It is all layer 2: one very large switched network

live statistics

Internet flattening: no more tiers?

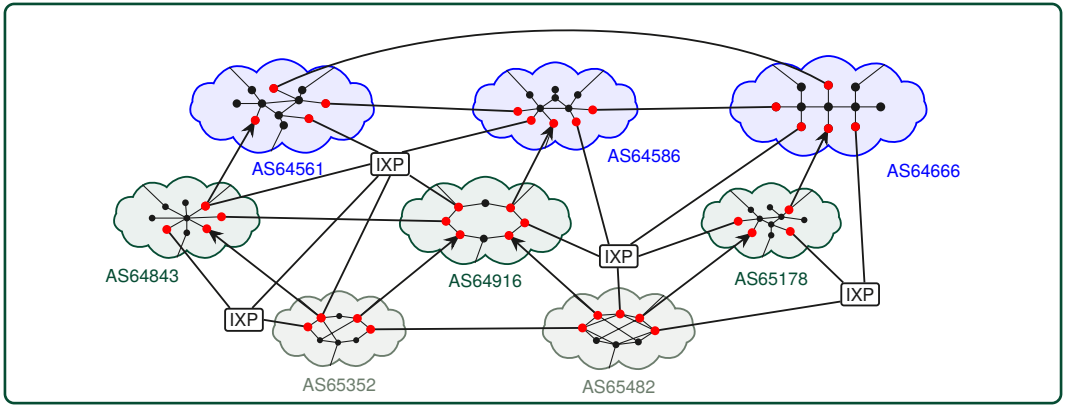

Following the data trail: an analysis of IXP dependencies Tashiro et al. (2024)

Ballistics proves the Earth is not flat. What about ‘Internet ballistics’?

Software-defined networking

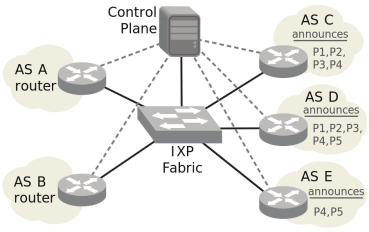

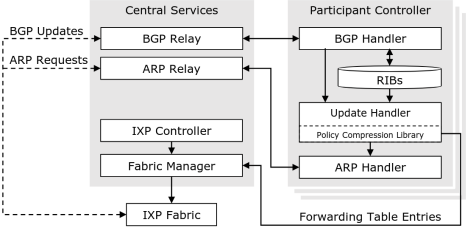

- **An industrial-scale software defined internet exchange point** Gupta et al. (2016)
- **A survey of security in software defined networks** Scott-Hayward, Natarajan and Sezer (2016)

Software-defined networking

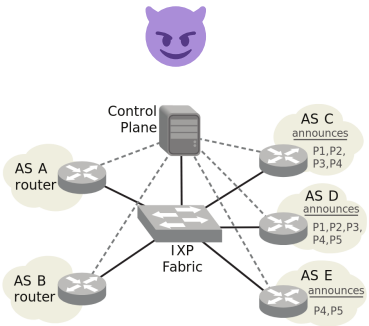

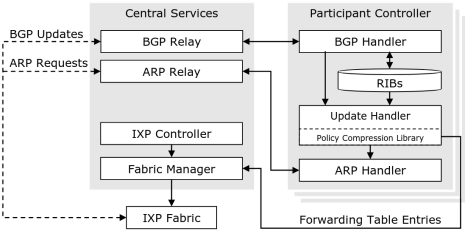

- **An industrial-scale software defined internet exchange point** Gupta et al. (2016)
- **A survey of security in software defined networks** Scott-Hayward, Natarajan and Sezer (2016)

JP 3-12 Cyberspace Operations

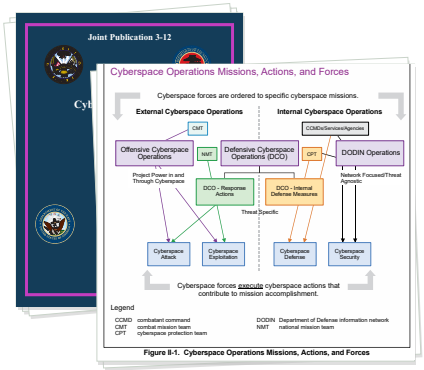

Cyberspace operations Joint Chiefs of Staff (2018)

Figure II-1, page II-3

Cyberspace attack

Actions that “create noticeable denial effects … in cyberspace or manipulation that leads to denial effects in the physical domains.”

Deny: prevent access to, operation of, or availability of a **target function**, by a specified level, for a specified time.

- **Degrade**: deny to a level, given as a % of capacity.
- **Disrupt**: deny completely, but temporarily; degrade 100%.
- **Destroy**: deny completely and irreparably (given enough time and resources, targets can be reconstituted)

Manipulate: change information, systems or networks in gray or red cyberspace to create (physical) denial effects. The adversary’s own resources turned to friendly purposes; the network seems functional until secondary effects are visible.

Border Gateway Protocol (again)

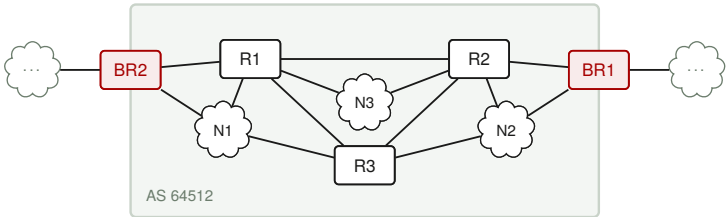

data plane: packets, hop by hop

- **Losing control of the Internet: using the data plane to attack the control plane**
  Schuchard et al. (2010)
- **Internet optometry: assessing the broken glasses in Internet reachability**
  Bush et al. (2009)

Border Gateway Protocol (again)

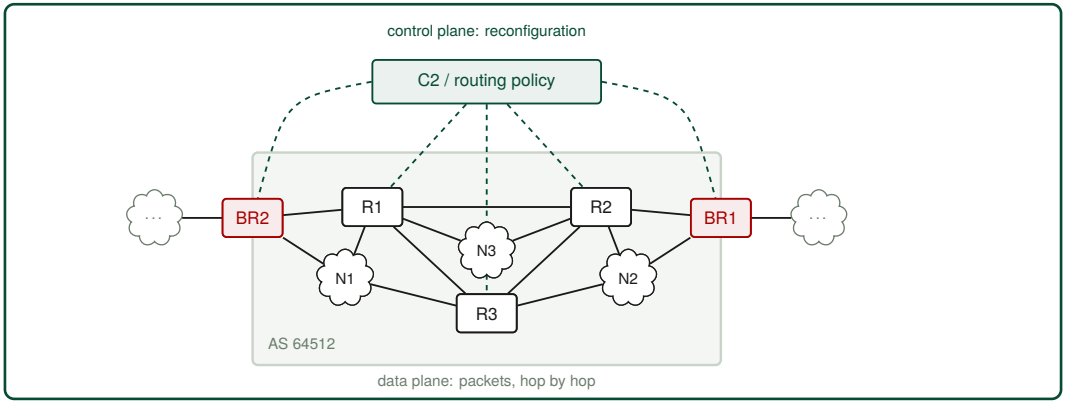

- **Losing control of the Internet: using the data plane to attack the control plane**
  Schuchard et al. (2010)
- **Internet optometry: assessing the broken glasses in Internet reachability**
  Bush et al. (2009)

Border Gateway Protocol (again)

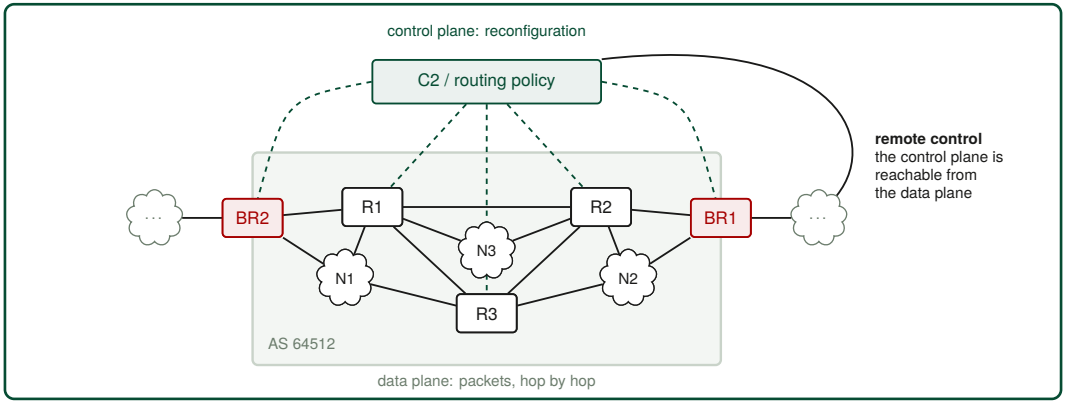

- **Losing control of the Internet: using the data plane to attack the control plane**
  Schuchard et al. (2010)
- **Internet optometry: assessing the broken glasses in Internet reachability**
  Bush et al. (2009)

Border Gateway Protocol (again)

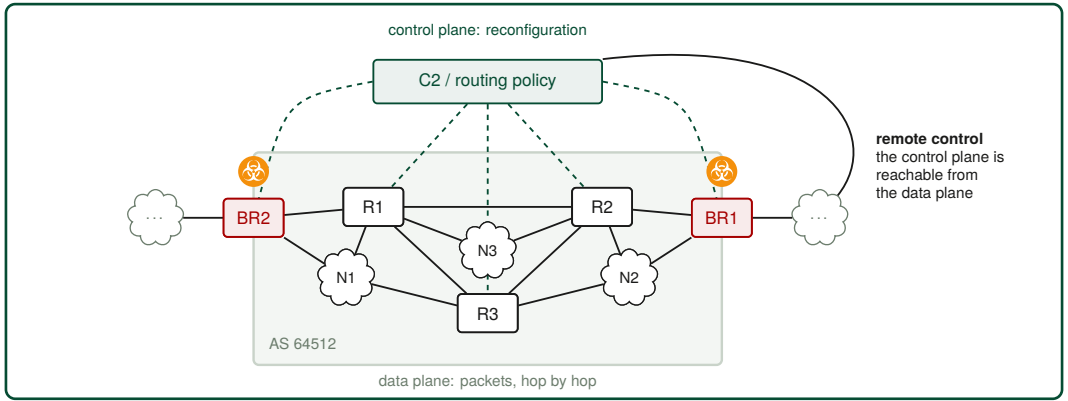

- **Global BGP attacks that evade route monitoring** Birge-Lee, Apostolaki and Rexford (2025)
- **BGP Vortex: update message floods can create internet instabilities** Stöger et al. (2025)

Next Generation Internet

- **Separating routing policy from mechanism in the network layer** Griffioen, Calvert, Ascigil & Yuan (2011)
- **Next Generation Internet 2025** Leenaars et al. (2018) — for the European Commission, DG CONNECT

> *Every technical primitive from the current internet needs to be reevaluated in terms of the complete threat catalog, as we cannot afford our **multi-trillion euro economies and societies** to not be able to trust the internet as we move forward.*

Michiel Leenaars et al. (2018)

Overview

Part I. First hour

1. Theory & Practice
2. The physical layer
3. The data plane

Part II. Second hour

4. The control plane
5. Learning from 9/11
6. **Regaining the cyber high ground**



Joint Cybersecurity Advisory

TLP:CLEAR

Countering Chinese State-Sponsored Actors Compromise of Networks Worldwide to Feed Global Espionage System

Executive summary

People's Republic of China (PRC) state-sponsored cyber threat actors are targeting networks globally, including, but not limited to, telecommunications, government, transportation, lodging, and military infrastructure networks. While these actors focus on large backbone routers of major telecommunications providers, as well as provider edge (PE) and customer edge (CE) routers, they also leverage compromised devices and trusted connections to pivot into other networks. These actors often modify routers to maintain persistent, long-term access to networks.

This activity partially overlaps with cyber threat actor reporting by the cybersecurity industry—commonly referred to as Salt Typhoon, OPERATOR PANDA, RedMike, UNC5807, and GhostEmperor, among others. The authoring agencies are not adopting a particular commercial naming convention and hereafter refer to those responsible for

This report is marked TLP:CLEAR. Recipients may share this information without restriction. Information is subject to standard copyright rules.

UCC/CFIBBS-25 | PP-25-2022 | August 2025 | Ver. 1.0

TLP:CLEAR

📄 **Countering Chinese state-sponsored actors compromise of networks worldwide** NSA et al. (2025)

"Cyber threat actors are targeting networks globally, including, but not limited to, telecommunications, government, transportation, lodging, and military infrastructure networks.

While these actors focus on **large backbone routers** of major telecommunications providers, as well as provider edge (PE) and customer edge (CE) routers, they also leverage compromised devices and trusted connections to pivot into other networks."

The cyber high ground

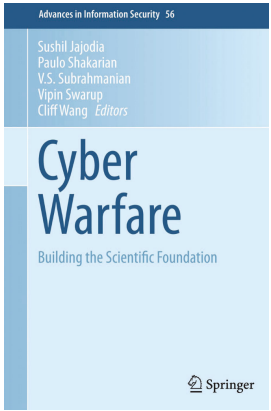

Cyber warfare: building the scientific foundation Jajodia et al. (2015)

- The cyber high ground: who has control?
- **Identifying and exploiting the cyber high ground for botnets** Sweeney and Cybenko (2015)
- **Weapons of mass disruption: how small states use cyber to resist larger powers** Korb (2025)
- Tomorrow: live-fire exercises with disrupting core infrastructure (‘Internet ballistics’)
- **Cyber ranges and testbeds for education, training, and research** Chouliaras et al. (2021)

The cyber high ground

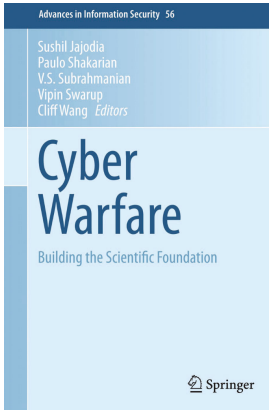

Cyber warfare: building the scientific foundation Jajodia et al. (2015)

- The cyber high ground: who has control?
- **Identifying and exploiting the cyber high ground for botnets** Sweeney and Cybenko (2015)
- **Weapons of mass disruption: how small states use cyber to resist larger powers** Korb (2025)
- Tomorrow: live-fire exercises with disrupting core infrastructure (‘Internet ballistics’)
- **Cyber ranges and testbeds for education, training, and research** Chouliaras et al. (2021)

The cyber high ground

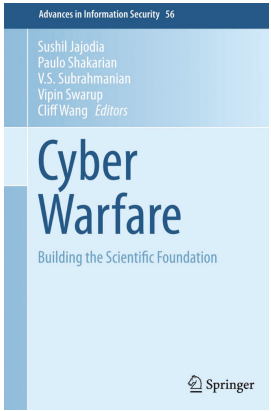

Cyber warfare: building the scientific foundation Jajodia et al. (2015)

- The cyber high ground: who has control?
- **Identifying and exploiting the cyber high ground for botnets** Sweeney and Cybenko (2015)
- **Weapons of mass disruption: how small states use cyber to resist larger powers** Korb (2025)
- Tomorrow: live-fire exercises with disrupting core infrastructure (‘Internet ballistics’)
- **Cyber ranges and testbeds for education, training, and research** Chouliaras et al. (2021)

Network line of defence & strategic weaknesses

Maginot lines and tourniquets: on the defendability of national cyberspace
Gosain, Rawat, Sharma and Acharya (2020)

- **'Line of defence'** small number of ASes covering most traffic
- Large cut set (= **survivability**)
- Authoritarian countries:
  - easy network surveillance
  - small min cut set
- Liberal countries:
  - difficult network surveillance
  - large min cut set

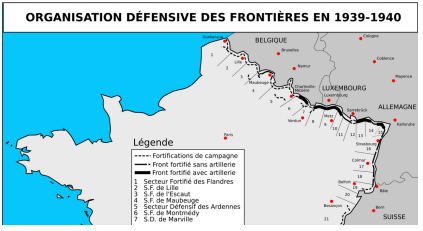



What the routing table records: Ukraine

- After 2014, routes from Crimea and Donbass were moved onto Russian providers.
- **A digital frontline**, drawn in BGP before it was drawn on the ground.
- The lever was not a protocol flaw. It was ownership: whoever runs the AS decides what it announces.

Measuring the fragmentation of the Internet: BGP during the Ukrainian crisis Douzet et al. (2020)

What the routing table records: Ukraine

- After 2014, routes from Crimea and Donbass were moved onto Russian providers.
- **A digital frontline**, drawn in BGP before it was drawn on the ground.
- The lever was not a protocol flaw. It was ownership: whoever runs the AS decides what it announces.

Measuring the fragmentation of the Internet: BGP during the Ukrainian crisis Douzet et al. (2020)

What the routing table records: Ukraine

- After 2014, routes from Crimea and Donbass were moved onto Russian providers.
- **A digital frontline**, drawn in BGP before it was drawn on the ground.
- The lever was not a protocol flaw. It was ownership: whoever runs the AS decides what it announces.

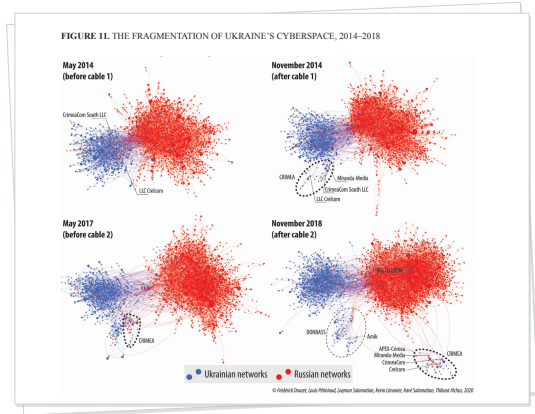

Measuring the fragmentation of the Internet: BGP during the Ukrainian crisis Douzet et al. (2020)

Who the Gulf connects through

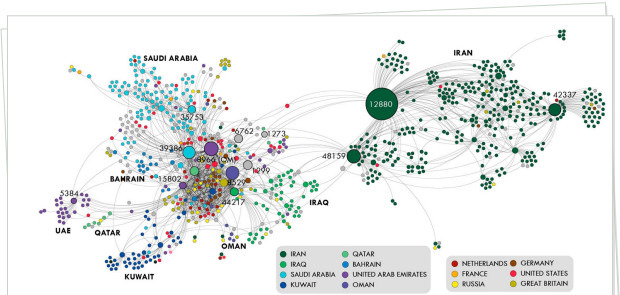

Figure 10. Connectivity of the Gulf Region and their neighbours, 2015.

Digital routes and borders in the Middle East Douzet et al. (2023)

- 2015: a few providers in the middle, none of them in the region.
- In 2021 there are more of them, with more of the region behind them.
- The **Netherlands** is among them. Two Gulf countries may meet each other in Amsterdam.

Who the Gulf connects through

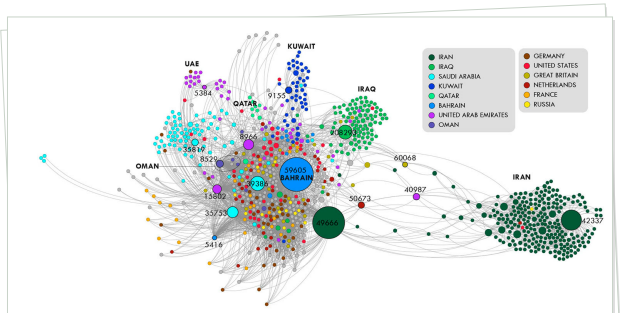

Figure 8. Connectivity of the Gulf Region and their neighbours, 2021.

Digital routes and borders in the Middle East Douzet et al. (2023)

- 2015: a few providers in the middle, none of them in the region.
- In 2021 there are more of them, with more of the region behind them.
- The **Netherlands** is among them. Two Gulf countries may meet each other in Amsterdam.

Who the Gulf connects through

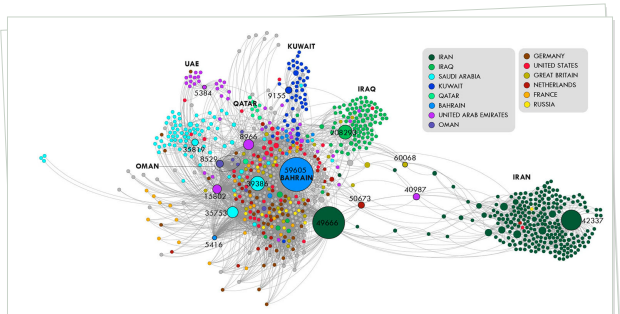

Figure 8. Connectivity of the Gulf Region and their neighbours, 2021.

Digital routes and borders in the Middle East Douzet et al. (2023)

- 2015: a few providers in the middle, none of them in the region.
- In 2021 there are more of them, with more of the region behind them.
- The **Netherlands** is among them. Two Gulf countries may meet each other in Amsterdam.

What each layer gives away

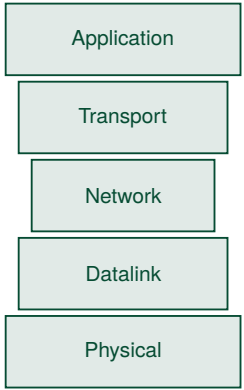

What each layer gives away

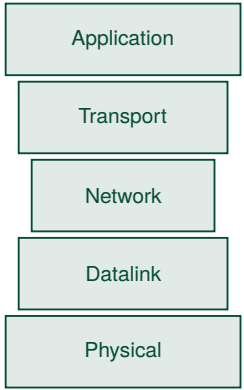

wire tapping

What each layer gives away

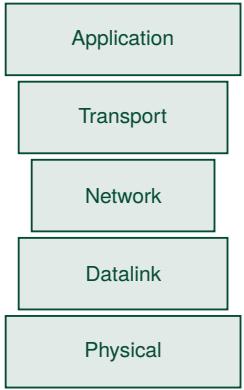

frame sniffing, spoofing

wire tapping

What each layer gives away

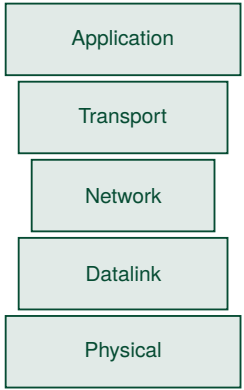

address spoofing, rerouting

frame sniffing, spoofing

wire tapping

What each layer gives away

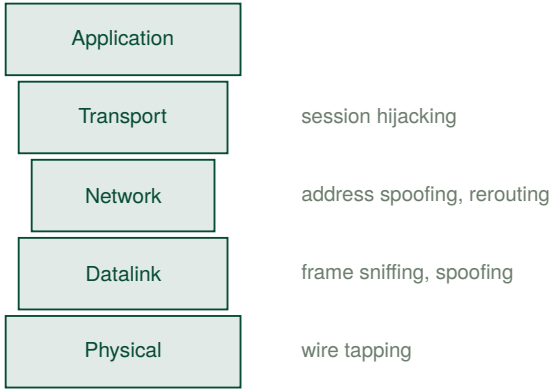

What each layer gives away

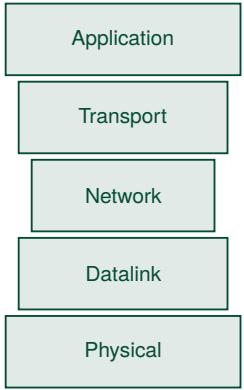

manipulating the information system itself

session hijacking

address spoofing, rerouting

frame sniffing, spoofing

wire tapping

Confidentiality, integrity, availability

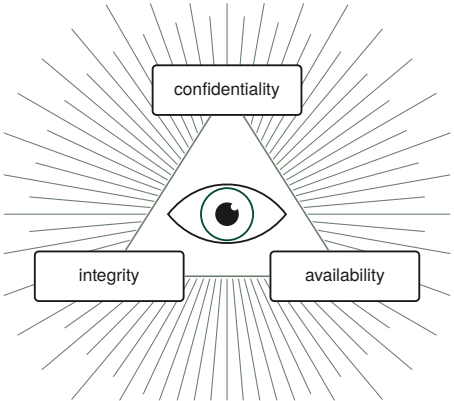

- **Confidentiality**: nobody else can read it.
- **Integrity**: no (undetected) change.
  - **Authenticity**: who really sent it.
  - **Non-repudiation**: cannot be denied.
- **Availability**: usable when needed.

The CIA Strikes Back: Redefining Confidentiality, Integrity and Availability in Security
Samonas and Coss (2014)

Confidentiality, integrity, availability

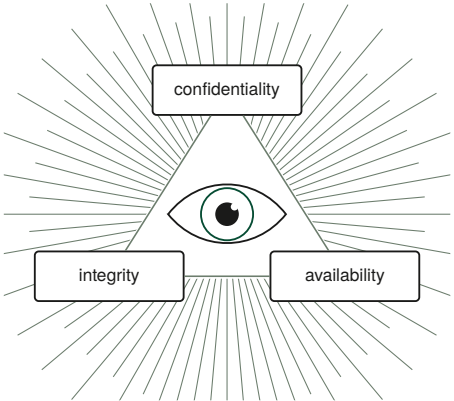

- **Confidentiality**: nobody else can read it.
- **Integrity**: no (undetected) change.
  - **Authenticity**: who really sent it.
  - **Non-repudiation**: cannot be denied.
- **Availability**: usable when needed.

The CIA Strikes Back: Redefining Confidentiality, Integrity and Availability in Security
Samonas and Coss (2014)

Confidentiality, integrity, availability

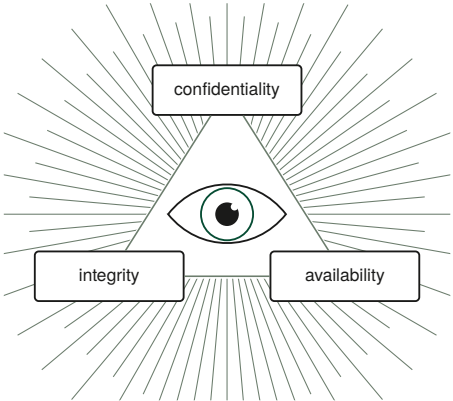

- **Confidentiality**: nobody else can read it.
- **Integrity**: no (undetected) change.
  - **Authenticity**: who really sent it.
  - **Non-repudiation**: cannot be denied.
- **Availability**: usable when needed.

The CIA Strikes Back: Redefining Confidentiality, Integrity and Availability in Security
Samonas and Coss (2014)

Confidentiality, integrity, availability

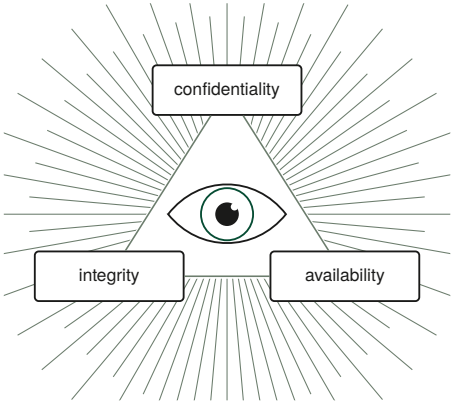

- **Confidentiality**: nobody else can read it.
- **Integrity**: no (undetected) change.
  - **Authenticity**: who really sent it.
  - **Non-repudiation**: cannot be denied.
- **Availability**: usable when needed.

The CIA Strikes Back: Redefining Confidentiality, Integrity and Availability in Security
Samonas and Coss (2014)

Confidentiality, integrity, availability

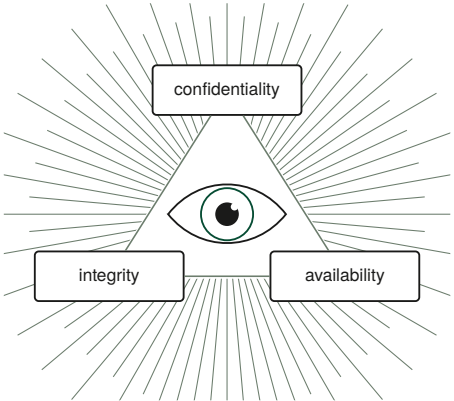

- **Confidentiality**: nobody else can read it.
- **Integrity**: no (undetected) change.
  - **Authenticity**: who really sent it.
  - **Non-repudiation**: cannot be denied.
- **Availability**: usable when needed.

The CIA Strikes Back: Redefining Confidentiality, Integrity and Availability in Security
Samonas and Coss (2014)

Every layer, and what protects it

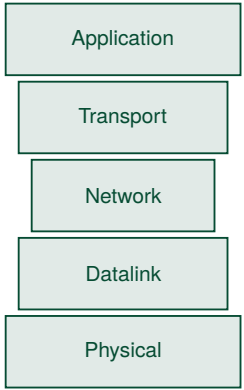

Every layer, and what protects it

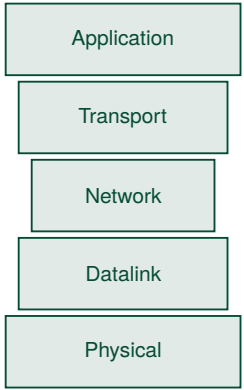

physical security

Every layer, and what protects it

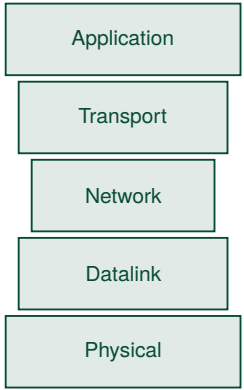

MACsec, 802.1X

physical security

Every layer, and what protects it

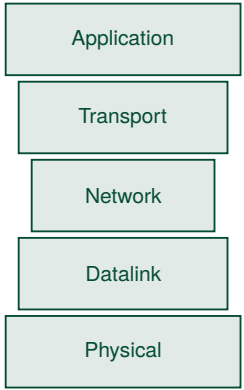

IPsec, and RPKI for the routes

MACsec, 802.1X

physical security

Every layer, and what protects it

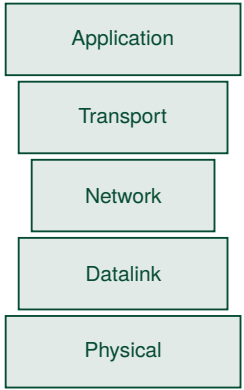

TLS: authenticated, encrypted streams

IPsec, and RPKI for the routes

MACsec, 802.1X

physical security

Every layer, and what protects it

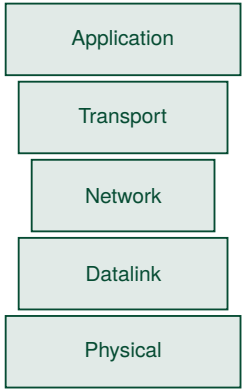

DNSSEC, signed records, and the system's own controls

TLS: authenticated, encrypted streams

IPsec, and RPKI for the routes

MACsec, 802.1X

physical security

Every layer, and what protects it

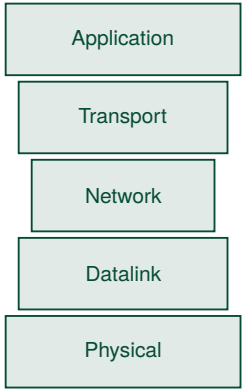

DNSSEC, signed records, and the system's own controls

TLS: authenticated, encrypted streams

IPsec, and RPKI for the routes

MACsec, 802.1X

physical security

Everything rests on protecting **private keys** to stay secret.

Operation Black Tulip



Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach
Fox-IT (2012)

Operation Black Tulip

the first intrusion



Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach
Fox-IT (2012)

Operation Black Tulip



Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach
Fox-IT (2012)

Operation Black Tulip



Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach

Fox-IT (2012)

Operation Black Tulip



Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach
Fox-IT (2012)

Operation Black Tulip



Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach
Fox-IT (2012)

Operation Black Tulip



Six weeks between the issuing and the revocation,
and the certificates were valid the whole time.

Black Tulip: Report of the Investigation into the DigiNotar Certificate Authority Breach
Fox-IT (2012)

What it cost

- Every browser and operating system removed the root. Every certificate DigiNotar had ever issued stopped being believed.
- Including the Dutch government's own: **PKIoverheid** ran on that authority, and the state took over its operation.
- The company was bankrupt within a month.
- One authority, out of a hundred in every trust store, and the weakest one set the strength of all of them.
- Do we learn from our mistakes? In 2026 the **eMudhra** certificate authority had an easy to guess private key.

What it cost

- Every browser and operating system removed the root. Every certificate DigiNotar had ever issued stopped being believed.
- Including the Dutch government's own: **PKIoverheid** ran on that authority, and the state took over its operation.
- The company was bankrupt within a month.
- One authority, out of a hundred in every trust store, and the weakest one set the strength of all of them.
- Do we learn from our mistakes? In 2026 the **eMudhra** certificate authority had an easy to guess private key.

What it cost

- Every browser and operating system removed the root. Every certificate DigiNotar had ever issued stopped being believed.
- Including the Dutch government's own: **PKIoverheid** ran on that authority, and the state took over its operation.
- The company was bankrupt within a month.
- One authority, out of a hundred in every trust store, and the weakest one set the strength of all of them.
- Do we learn from our mistakes? In 2026 the **eMudhra** certificate authority had an easy to guess private key.

What it cost

- Every browser and operating system removed the root. Every certificate DigiNotar had ever issued stopped being believed.
- Including the Dutch government's own: **PKIoverheid** ran on that authority, and the state took over its operation.
- The company was bankrupt within a month.
- One authority, out of a hundred in every trust store, and the weakest one set the strength of all of them.
- Do we learn from our mistakes? In 2026 the **eMudhra** certificate authority had an easy to guess private key.

What it cost

- Every browser and operating system removed the root. Every certificate DigiNotar had ever issued stopped being believed.
- Including the Dutch government's own: **PKIoverheid** ran on that authority, and the state took over its operation.
- The company was bankrupt within a month.
- One authority, out of a hundred in every trust store, and the weakest one set the strength of all of them.
- Do we learn from our mistakes? In 2026 the **eMudhra** certificate authority had an easy to guess private key.

Creating cyber ranges

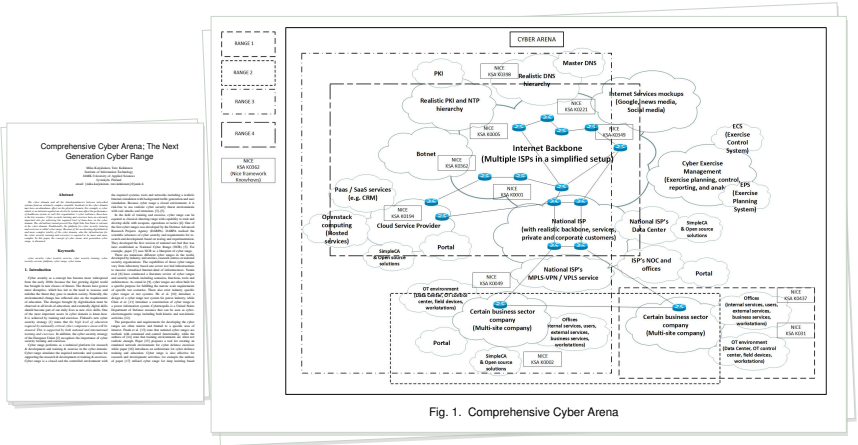

Fig. 1. Comprehensive Cyber Arena

Comprehensive cyber arena: the next generation cyber range Karjalainen & Kokkonen (2020)

SCION: Scalability, Control, Isolation



English *scion*: “a young shoot or twig of a plant”, from Continental Old French *cion*, “offspring, new growth”.

- **S**calable path-based routing
- **C**ontrol for end-users to select their own paths
- **I**solation of failures, limiting impact
- **ON** next-generation networks

Developed at ETH Zürich since 2009: a clean-slate protocol, specified with **formal methods** before it was deployed. Free and open source, and one commercial implementation.

And so much more...



NATO CCDCOE training catalogue, 2026 NATO CCDCOE (2026)

More information

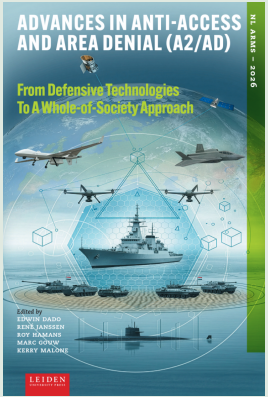

- **The Critical State of Cyberspace**
  Hiep (2026)
- Chapter 13 of **NL ARMS 2026**
  Netherlands Annual Review of Military Studies:

  Advances in Anti-Access and Area Denial (A2/AD): From Defensive Technologies To A Whole-of-Society Approach