UNIVERSITY OF WATERLOO

CS 456/656 Computer Networks

Lecture 1: Introduction – Part 1

Mina Tahmasbi Arashloo and Bo Sun

Fall 2024

About this course

- What is a computer network?
- How do modern computer networks work?
- The Internet is one of the most important collections of computer networks.
- So, we will use it as our guiding example throughout the course.

Why learn about computer networks?

- Because they are everywhere!
  - Every time you connect to an online service, you are sending and receiving data through multiple computer networks.
  - Most large-scale online services are distributed applications that use computer networks to communicate and collectively processes millions of requests per second.
- Our world is already relying on computer networks for its day-to-day operation
- Our dependence on computer networks is only going to grow over time.

Today's agenda

- A high-level overview of computer networks
- Course logistics

- But before we start…

Territorial Acknowledgement

The University of Waterloo acknowledges that much of our work takes place on the traditional territory of the Neutral, Anishinaabeg and Haudenosaunee peoples. Our main campus is situated on the Haldimand Tract, the land granted to the Six Nations that includes six miles on each side of the Grand River. Our active work toward reconciliation takes place across our campuses through research, learning, teaching, and community building, and is centralized within the Office of Indigenous Relations.

A note on the slides

Adapted from the slides that accompany this book.

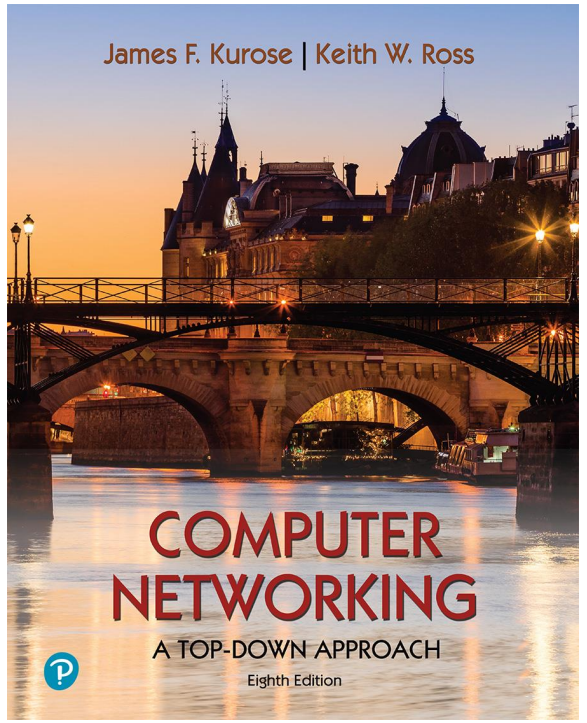

Computer Networking: A Top-Down Approach
8th edition
Jim Kurose, Keith Ross
Pearson, 2020

About the instructors – Mina Tahmasbi Arashloo

- Mina does research on computer networks.
  - Specifically, software-defined and programmable networks, and network verification.
- She got her PhD from Princeton University in 2019
- Then, she did a post-doc at Cornell University until 2022
- She joined University of Waterloo as an assistant professor in July 2022
- A major part of Mina's job is to continue doing research on computer networks
  - Feel free to reach out if you want to talk about research!

About the instructors – Bo Sun

- Bo's research is on algorithms and optimization for networked systems
  - including computer networks and power grids
- Bo got his Ph.D. from the Hong University of Science and Technology (HKUST)
  - Keith Ross is the advisor of his Ph.D. advisor
- Feel free to reach out if you want to talk about research!

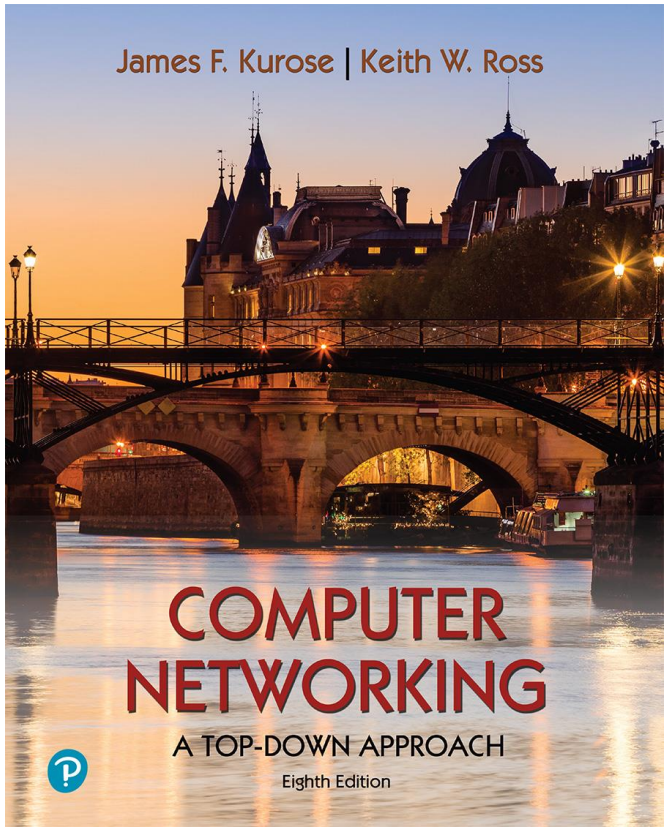

Introduction

What is a computer network?

What is the Internet?

Introduction

What is a computer network?

What is the Internet?

What is a computer network?

- A *system* that enables computers to *exchange information*

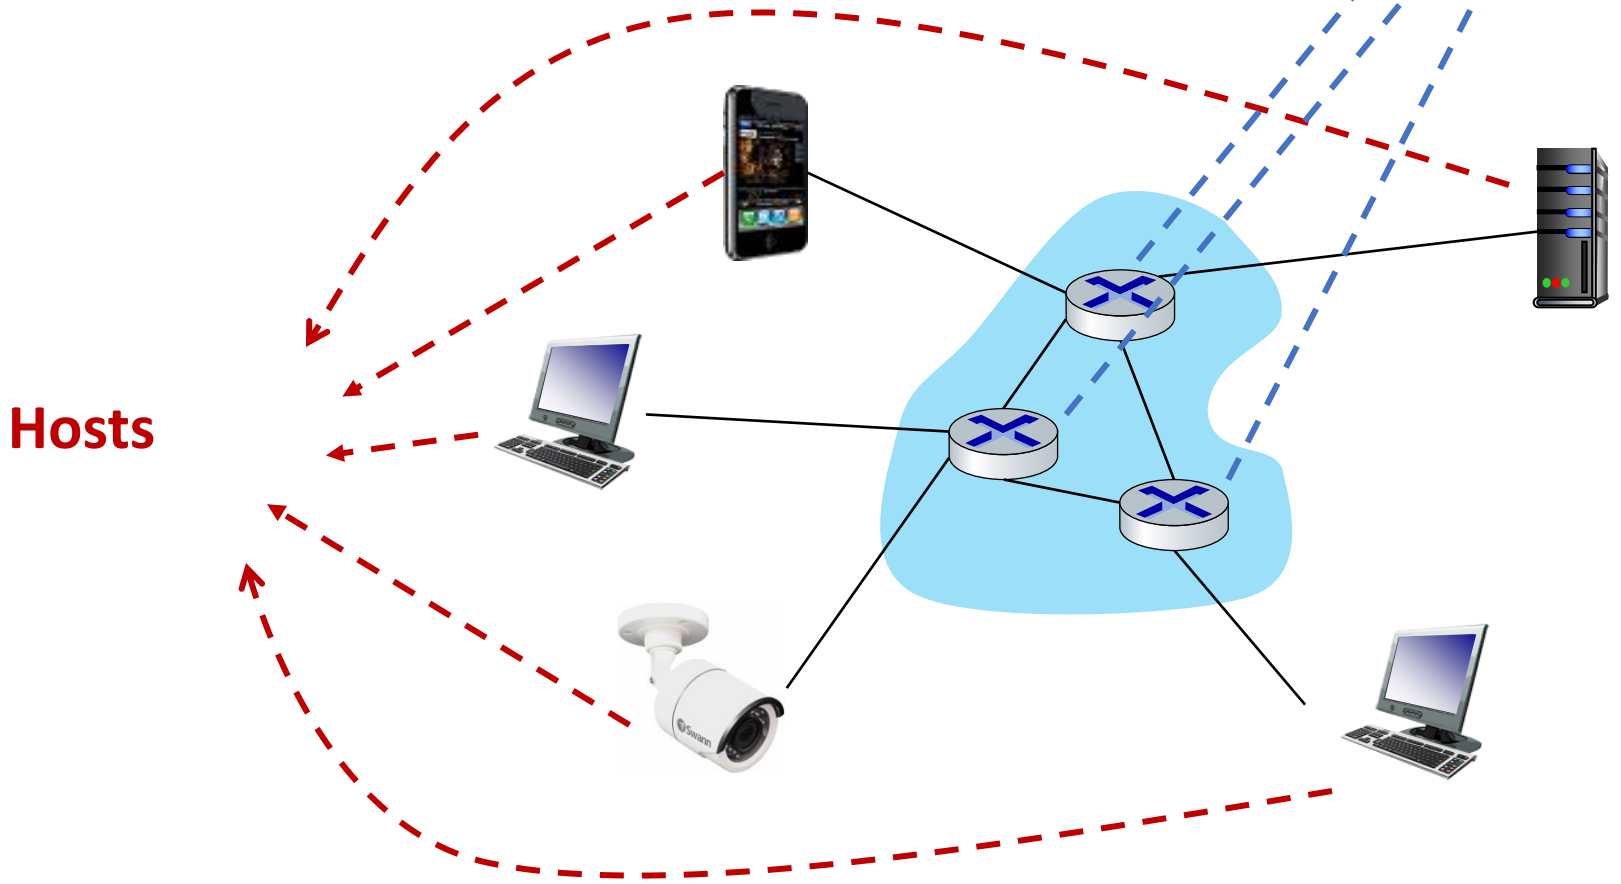

What is a computer network?

- A *system* that enables computers to *exchange information*
- The information exchange is governed by *protocols*.

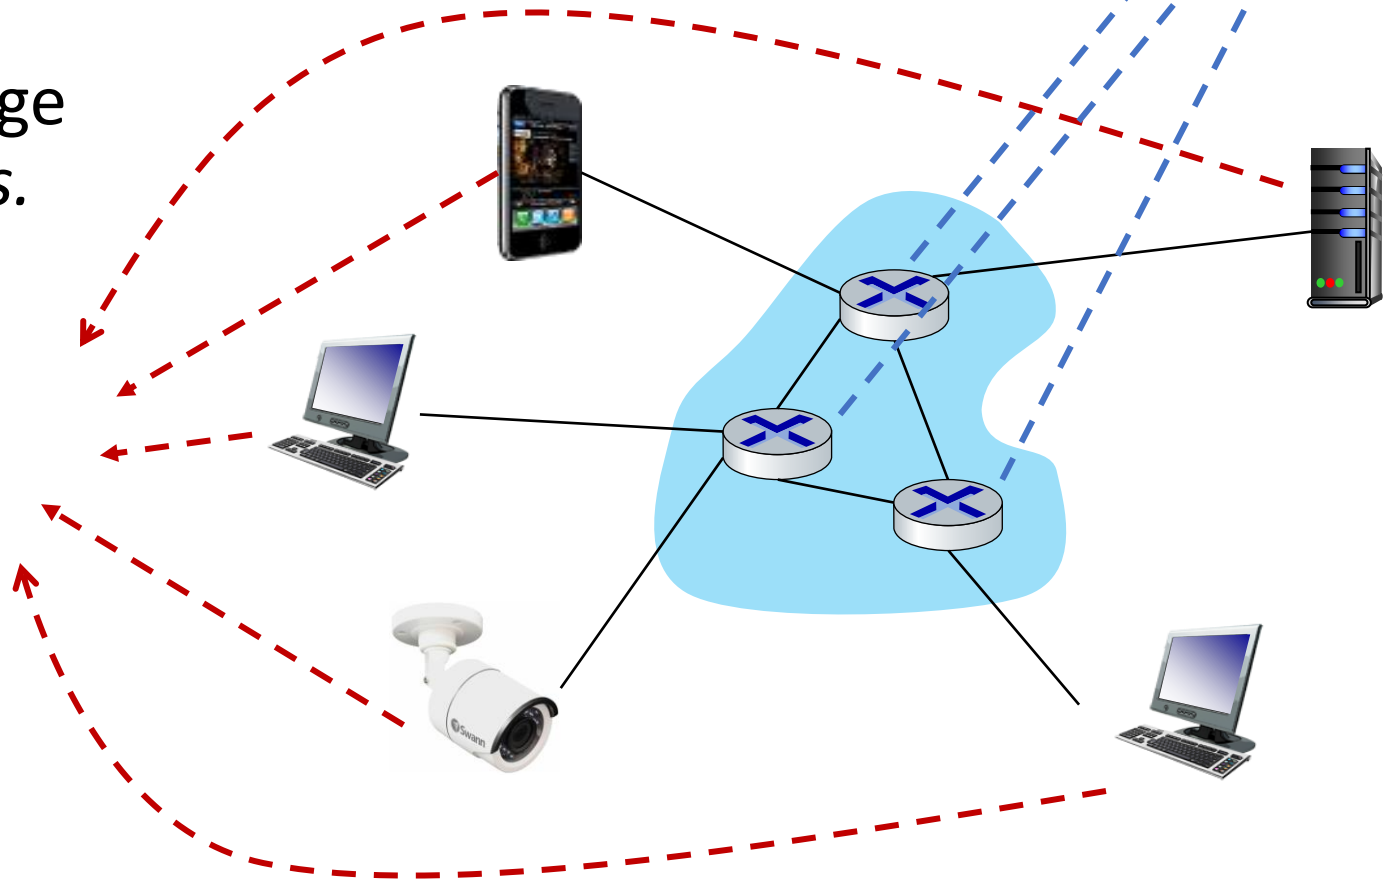

What is a protocol?

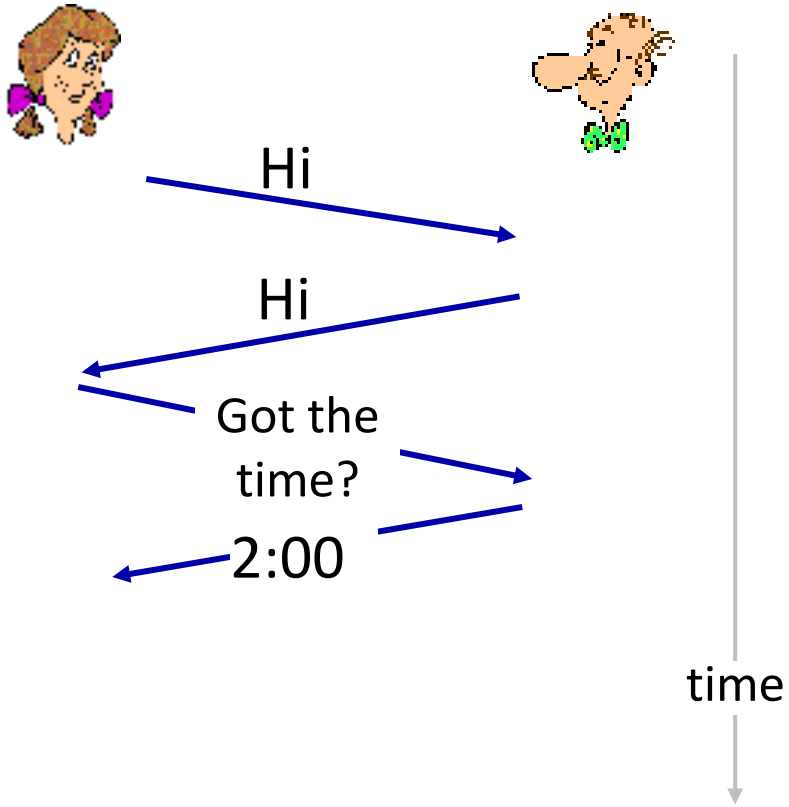

Rules for:

… specific messages sent

… specific actions taken when message received, or other events

Network protocols

- Network protocols are between computers (devices) rather than humans

A protocol defines:

- *the format and order of messages sent and received among network entities, and*
- *actions taken on message transmission and receipt*

Network protocols

- Network protocols are between computers (devices) rather than humans

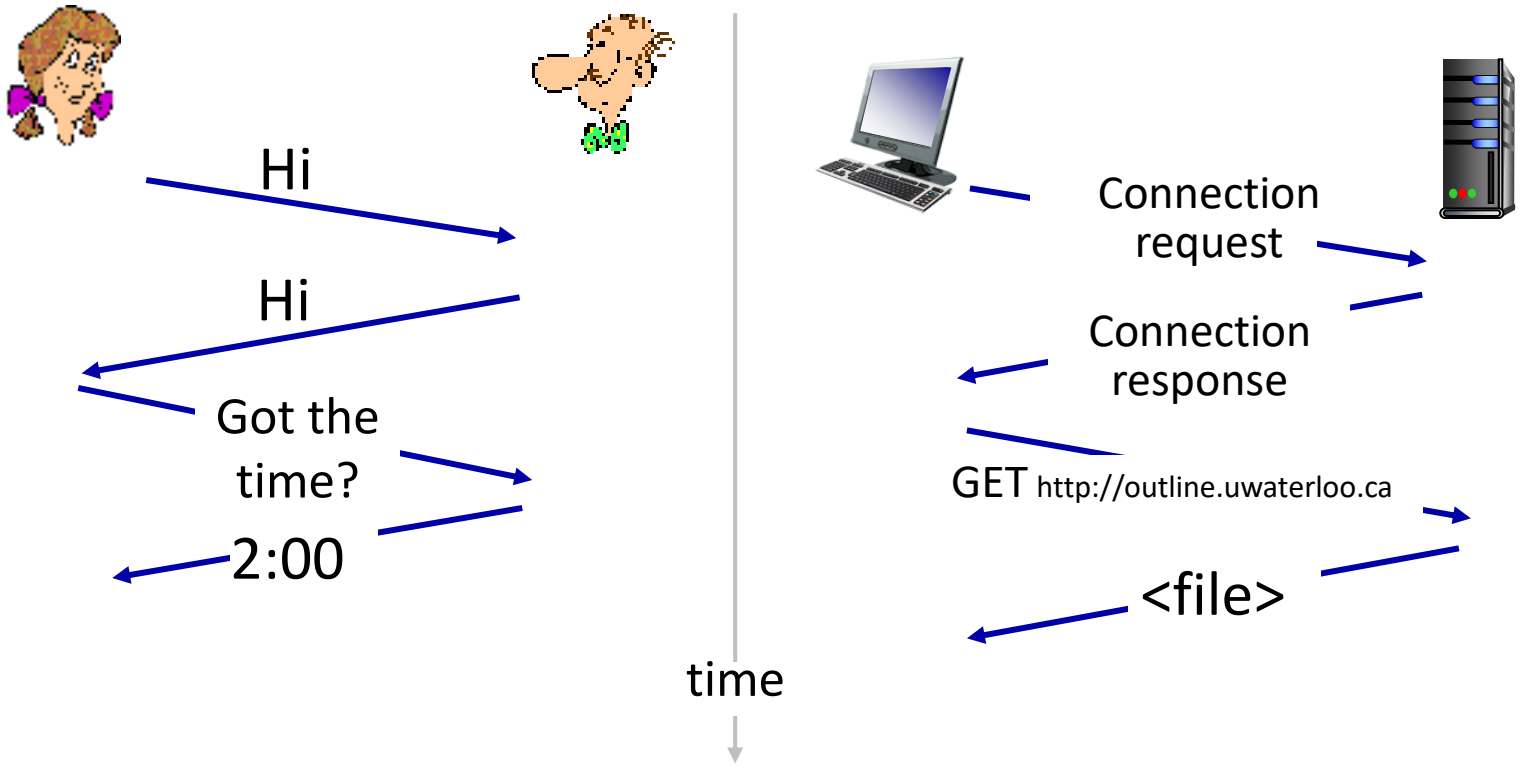

Network protocols

- Network protocols are between computers (devices) rather than humans

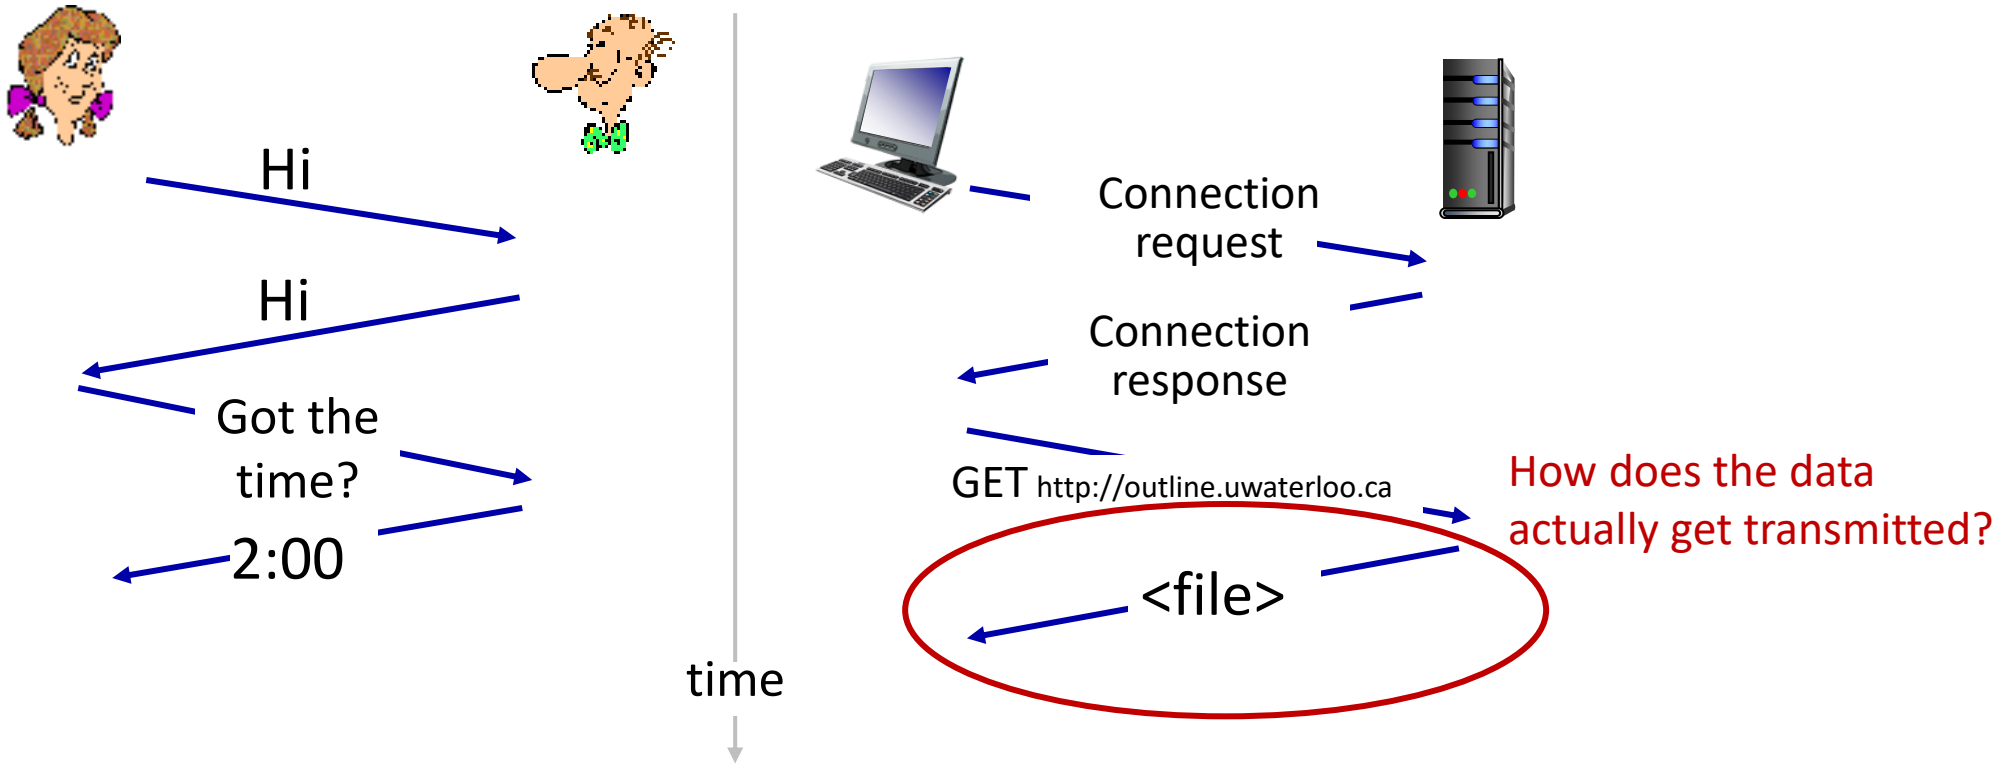

Packets: Data Transmission Units

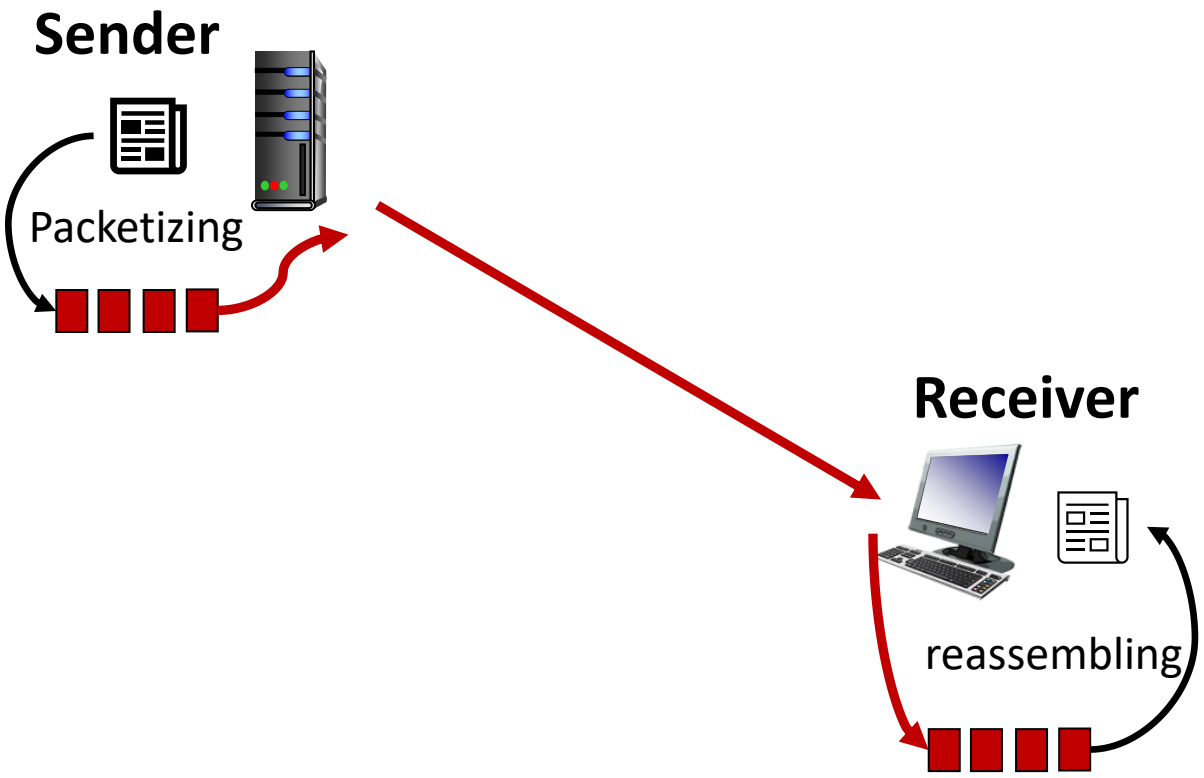

- Sender breaks data into smaller chunks
- Over the Internet, these chunks are known as *packets*
- Receiver receives data chunks and reassembles the data
- Q: What are the benefits of transmitting data in smaller chunks?

What is a computer network?

- A *system* that enables computers to *exchange information*
- The information exchange is governed by *protocols*.
- Data is transmitted in small chunks called *packets*

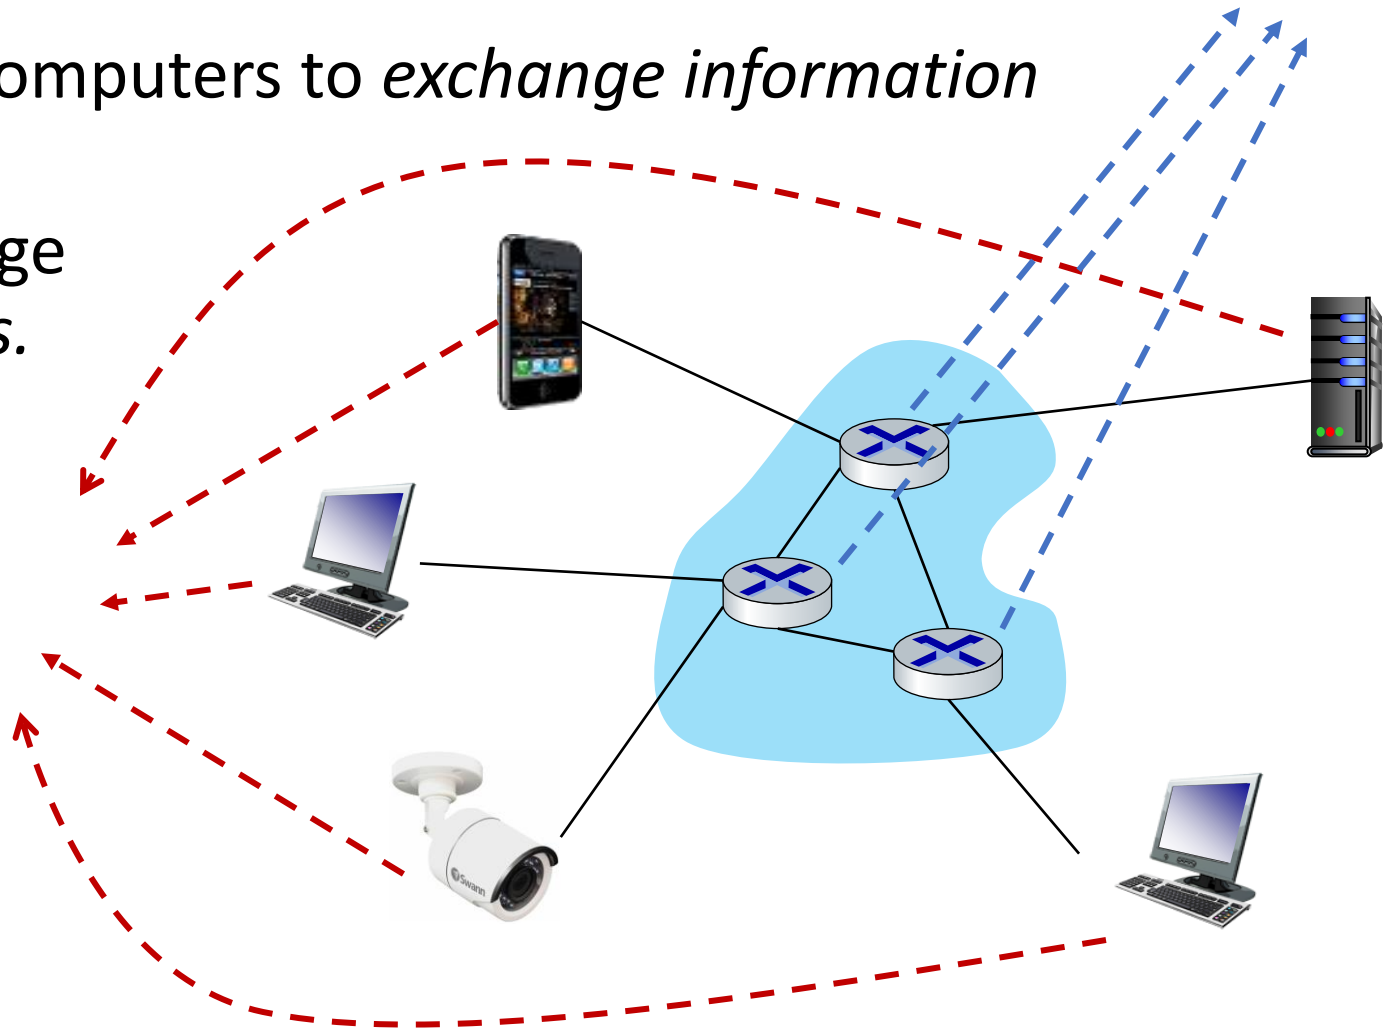

Sounds simple, right?

Turns out, when you want to *connect many computers*, things get complicated very fast.

Connecting many computers

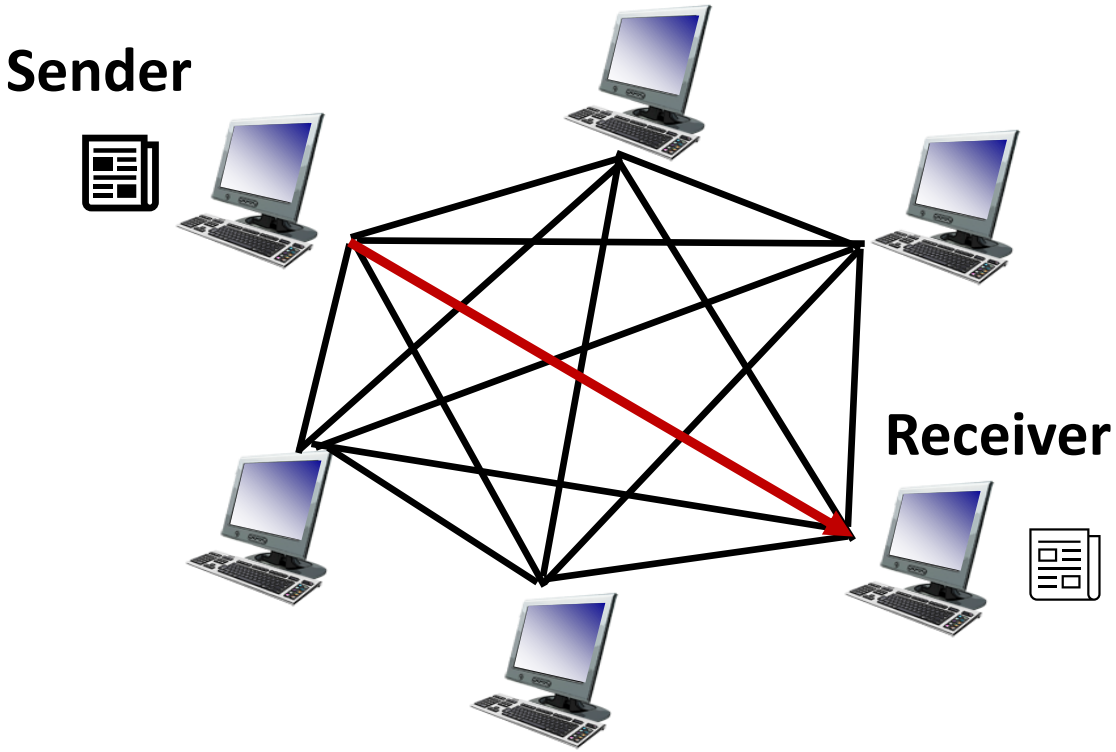

- In principle, if there are N computers, we can use $N(N-1)/2$ dedicated links between each pair
- but this does not scale well
  - N could be thousands or millions of computers!

Connecting many computers – cont.

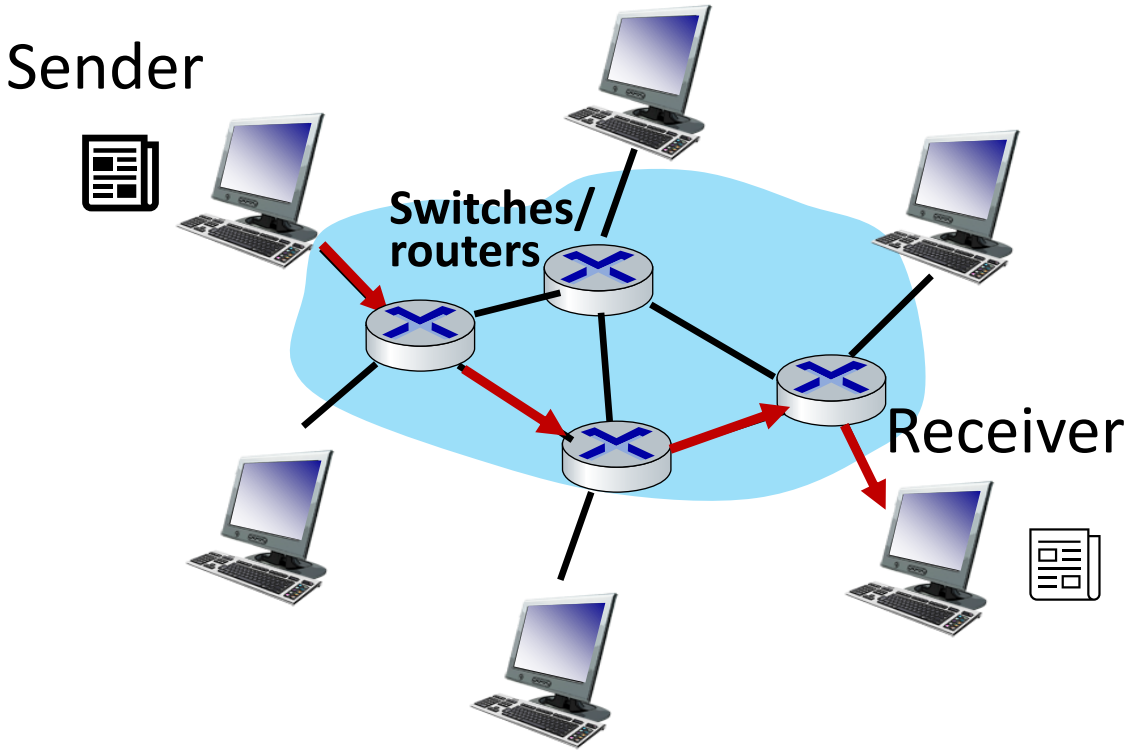

- Solution: Use a collection of shared network devices and links.
- Single-"hop" becomes multi-"hop"
- *Switches/routers's* job is to relay the data
- Q: What are the pros and cons of this approach?

Communicating over a shared network

- Multiple pairs of senders and receivers will use the shared devices and links simultaneously.
- Ideally, we'd like
  - data transfer between each pair to be done within a reasonable time.
  - the network to be utilized well – i.e., not to be overwhelmed with too many simultaneous requests or be left idle for too long.

Communicating over a shared network

- Multiple pairs of senders and receivers will use the shared devices and links simultaneously.
- Ideally, we'd like
  - data transfer between each pair to be done within a reasonable time.
  - the network to be utilized well – i.e., not to be overwhelmed with too many simultaneous requests or be left idle for too long.

Q: what is considered "reasonable time"?

Communicating over a shared network

- To make that happen, networking people have to solve several challenging problems:
  - How to decide when a sender gets to transmit data?
  - How to pick good paths for getting data from its source to its destination?
  - How to adapt when a switch/router or a link fails?
  - …

Communicating over a shared network

- To make that happen, networking people have to solve several challenging problems:
  - How to decide when a sender gets to transmit data?
  - How to pick good paths for getting data from its source to its destination?
  - How to adapt when a switch/router or a link fails?
  - …

Not as straightforward as it may seem.
We'll discuss two common solutions.

Solution 1: Reserve before sending

- D1, D2, D3, and D4 are network devices.
- Suppose that if D_i is connected to D_j, it can send X data units to D_j every second.
- When A wants to communicate with B, it
  - tells the network how many data units per second it needs to send ($\frac{X}{4}$ in this example)
  - asks the network to reserve resources along a path between A and B accordingly.

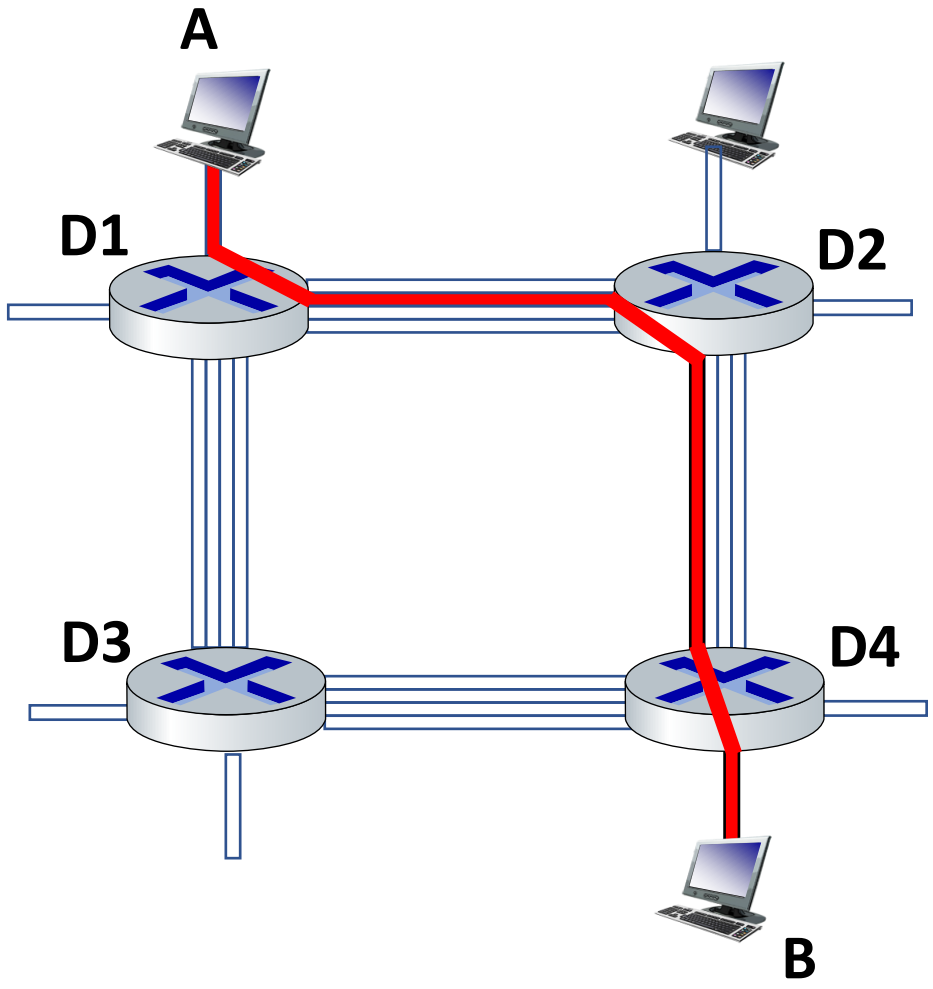

Solution 1: Reserve before sending

- To reserve resources, each device along the path records that it needs to reserve 1/4th of the capacity of its link to the next device for the data from A to B.
- Once the network confirms the reservation, A can start sending.
- Once A is done sending, the network can end the reservation.

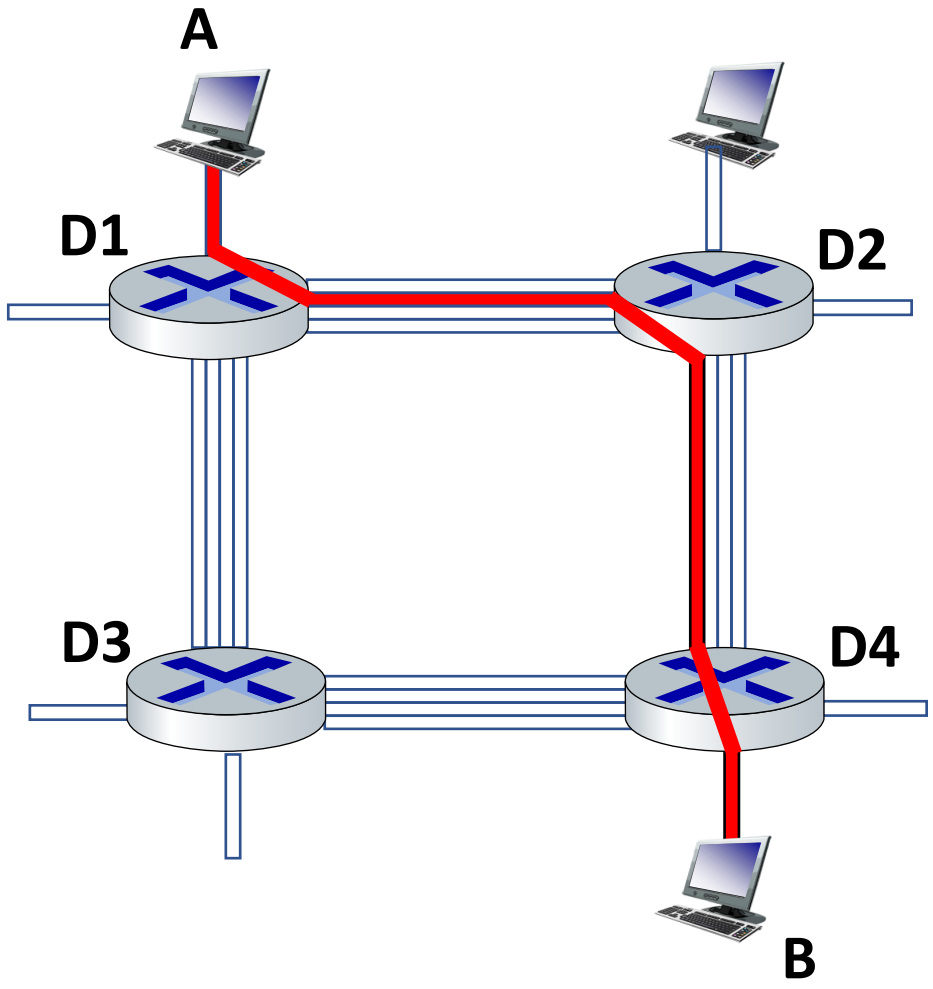

Solution 1: Reserve before sending

- Reservations are great!
- While A is sending data to B, ¼ of the capacity of the top and right links is dedicated to that data transfer.
- Pro: No one else can use it
- Con: No one else can use it
- Why is it both a pro and a con?

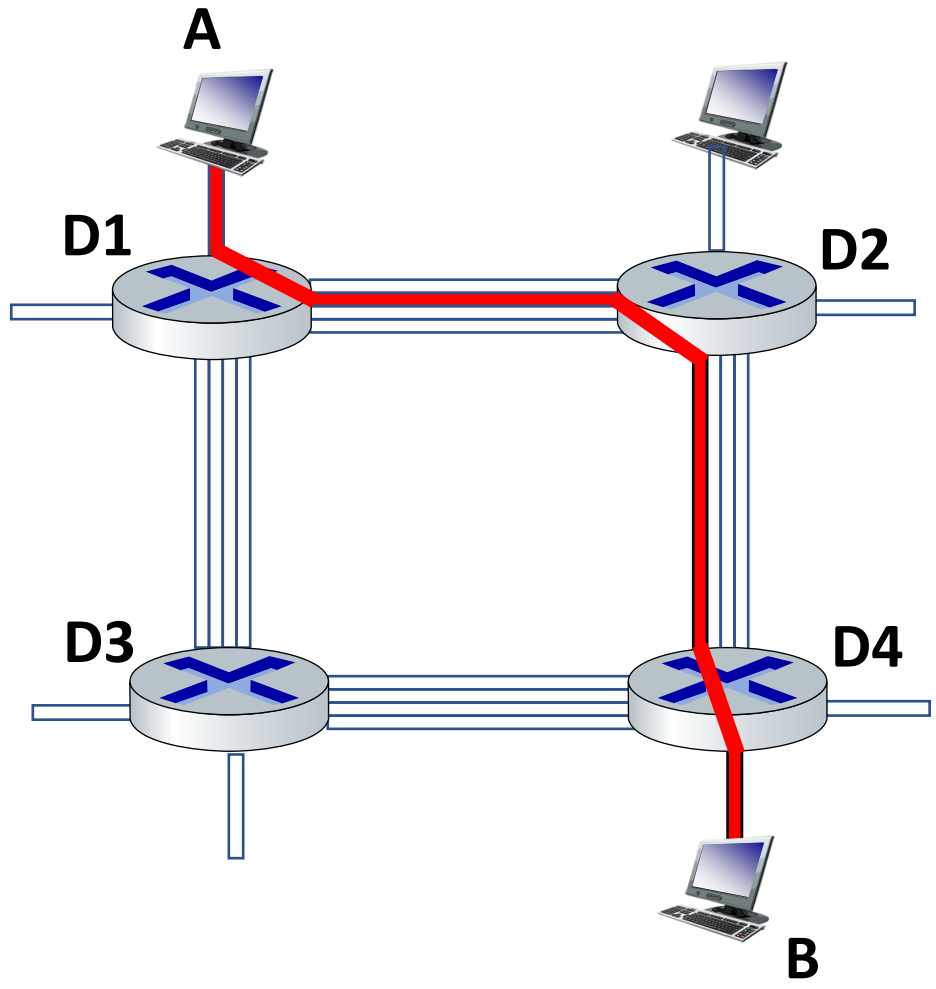

Solution 1: Reserve before sending

- No sharing means dedicated resources and circuit-like (guaranteed) performance.
- No sharing also means if A and B are not using the circuit, it will go unutilized.

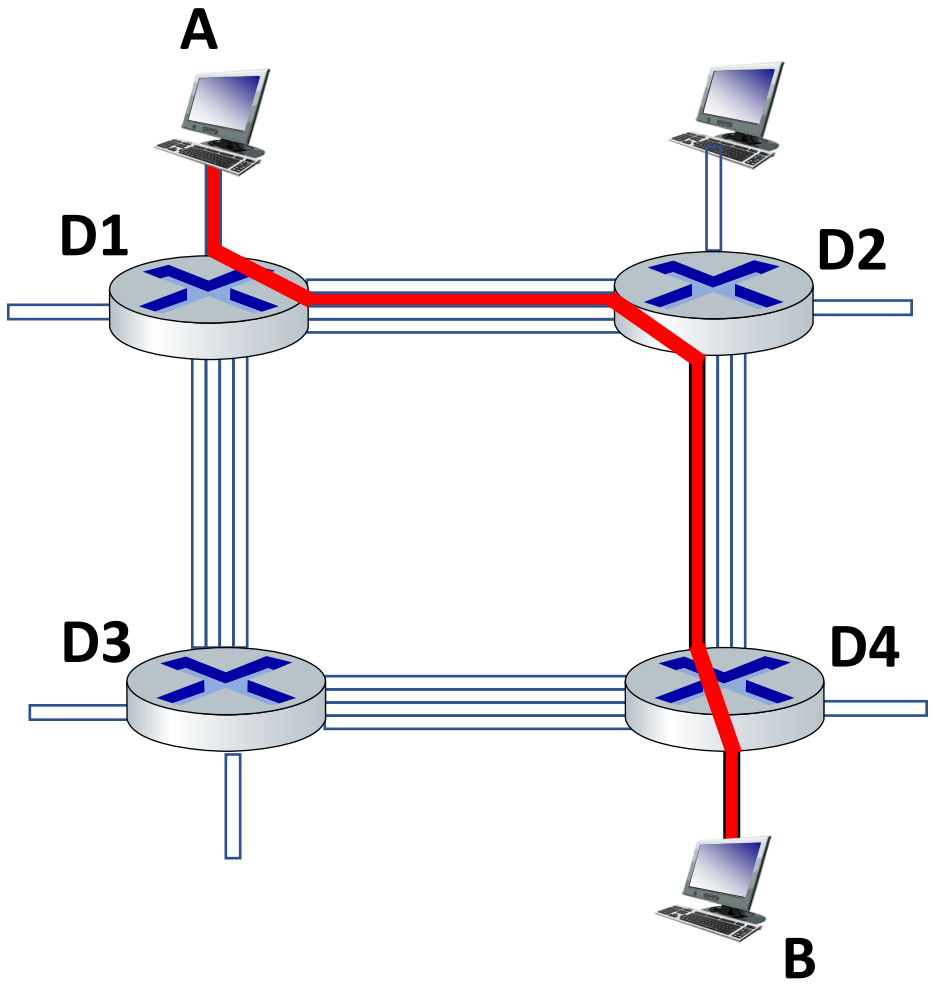

Reserve before sending = Circuit switching

- Circuit: the reserved path over the network
- Switching: moving data from one switch “port” to another (more on this later)
- Circuit switching is commonly used in traditional telephone networks.
- There were some circuit-switching-like proposals in the early days of Internet

Side note – the term switching

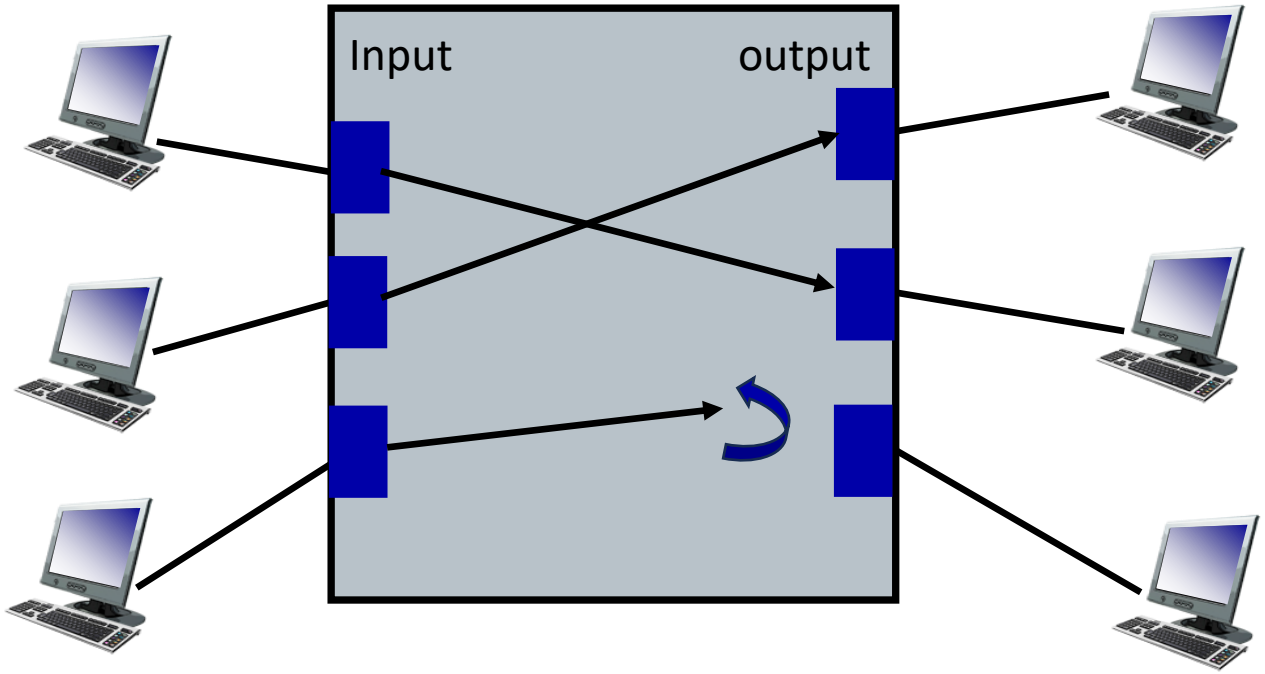

Switching: moving data from one port/channel to another.

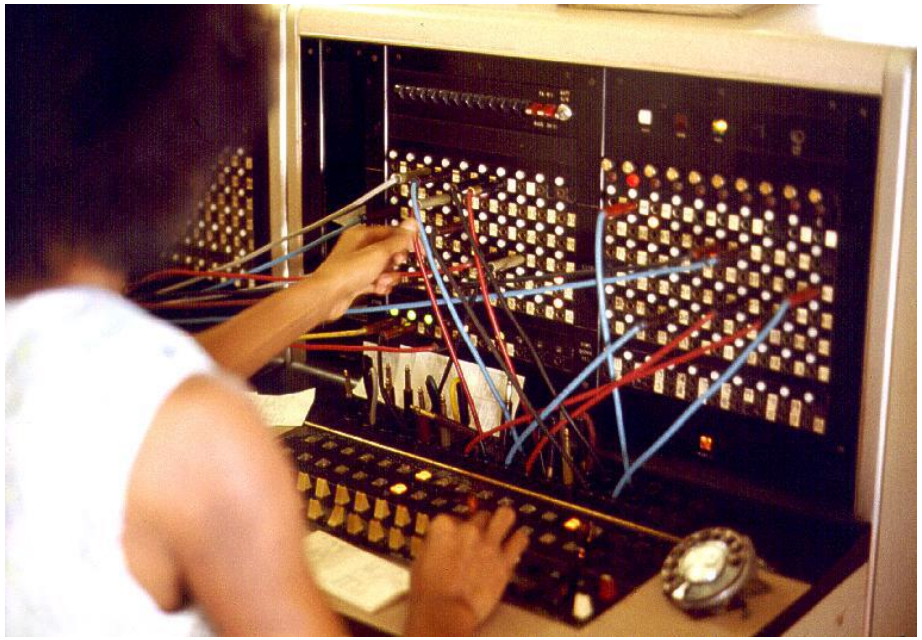

Calls are connected with cord pairs at a telephone switchboard.

Solution 2: Send on demand

- The sender sends as soon as it has data to send.
- Why could this be a good idea?

Solution 2: Send on demand

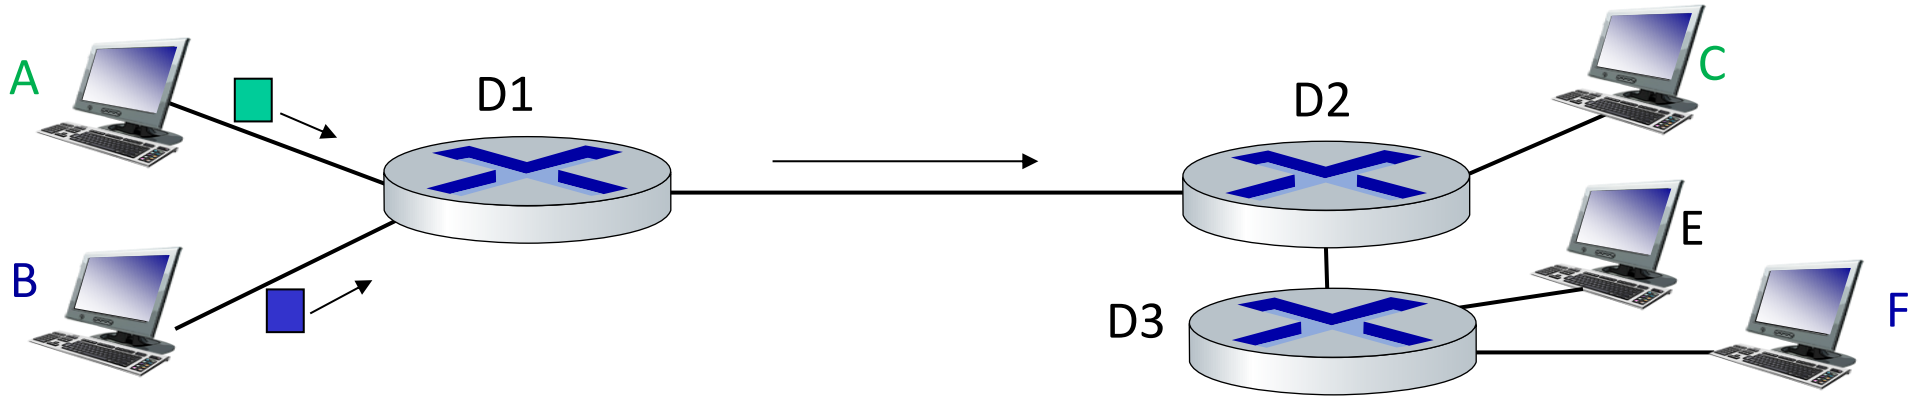

- A is sending data to C
- B is sending data to F
- D1 can send 4 “packets” to D2 every second

Solution 2: Send on demand

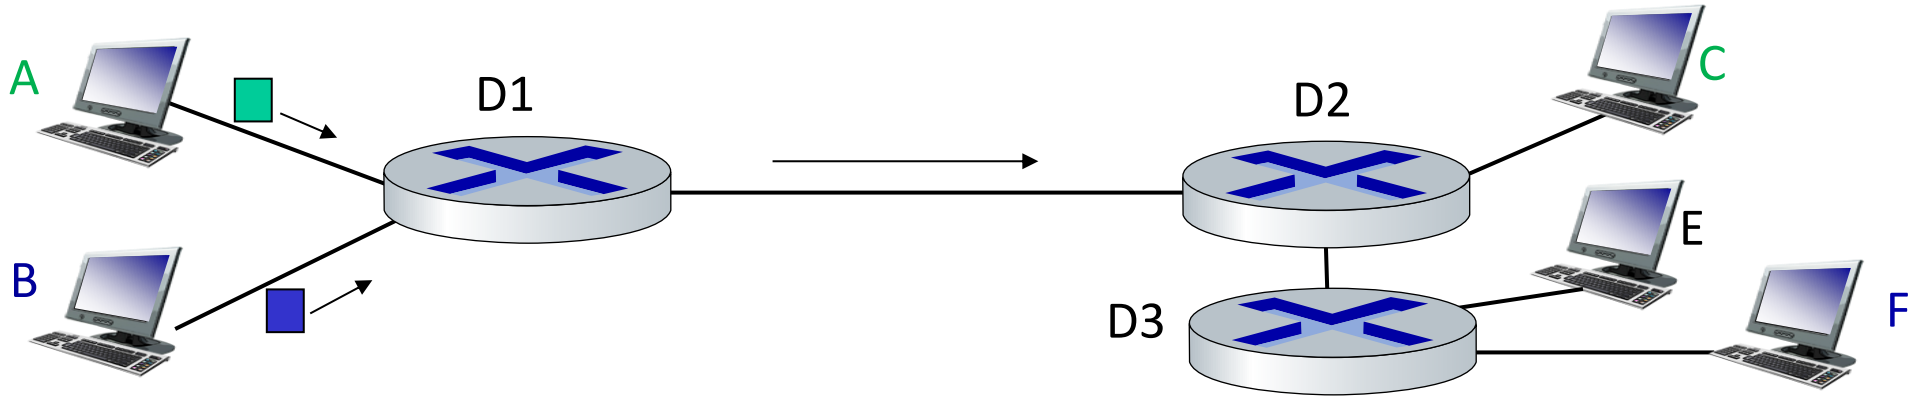

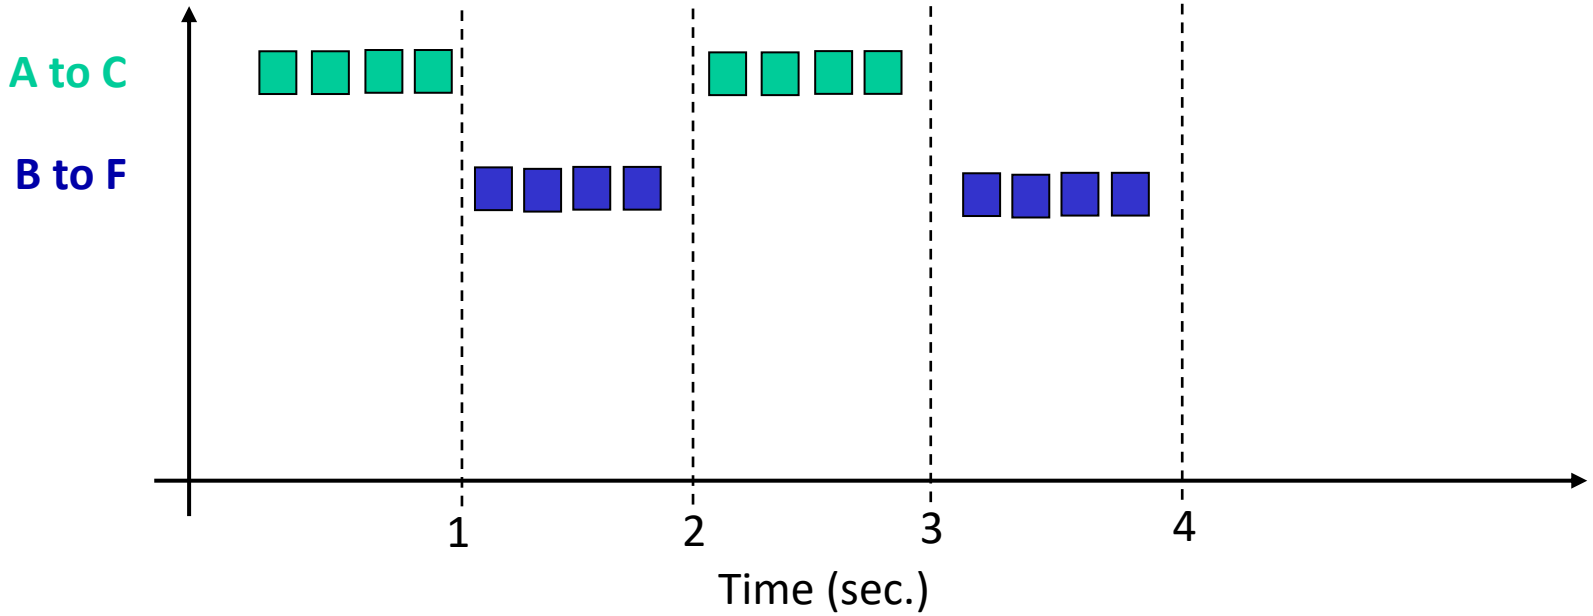

Solution 2: Send on demand

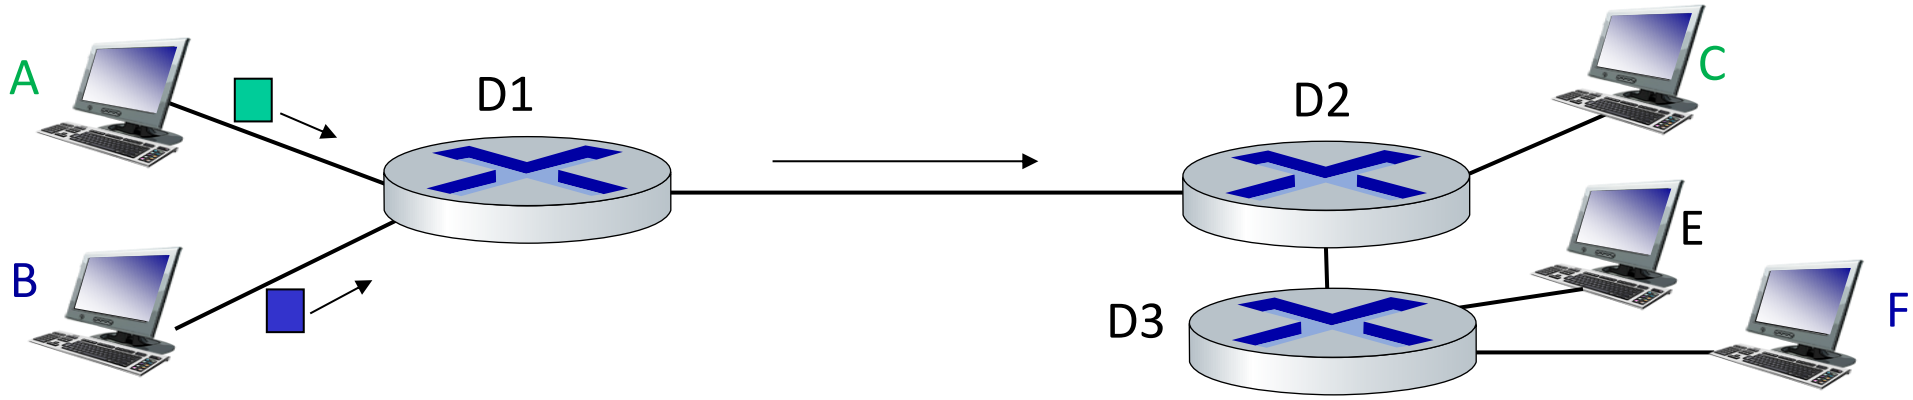

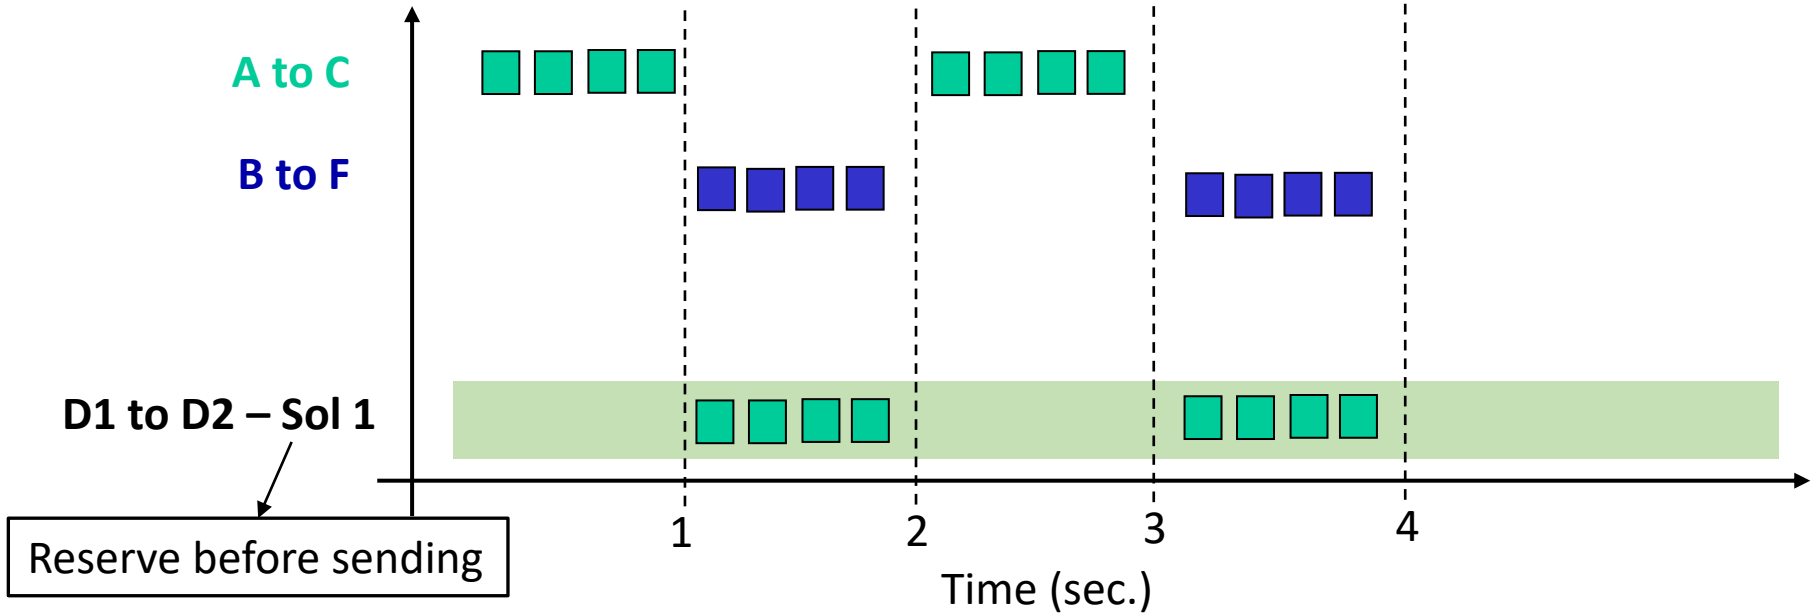

Solution 2: Send on demand

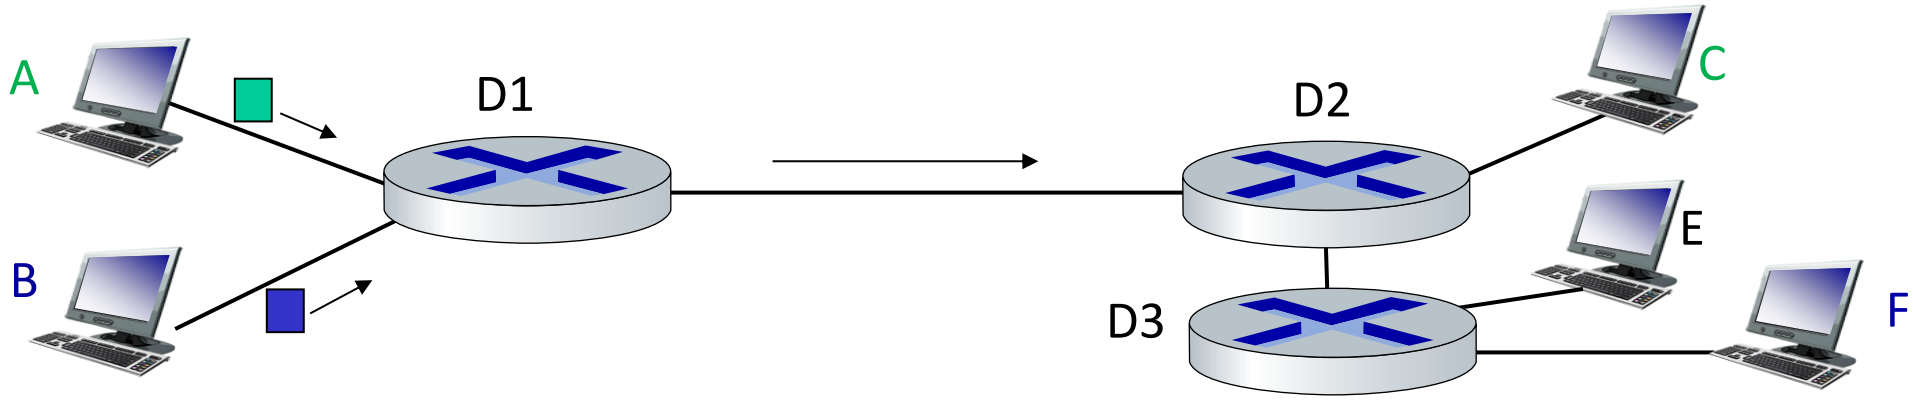

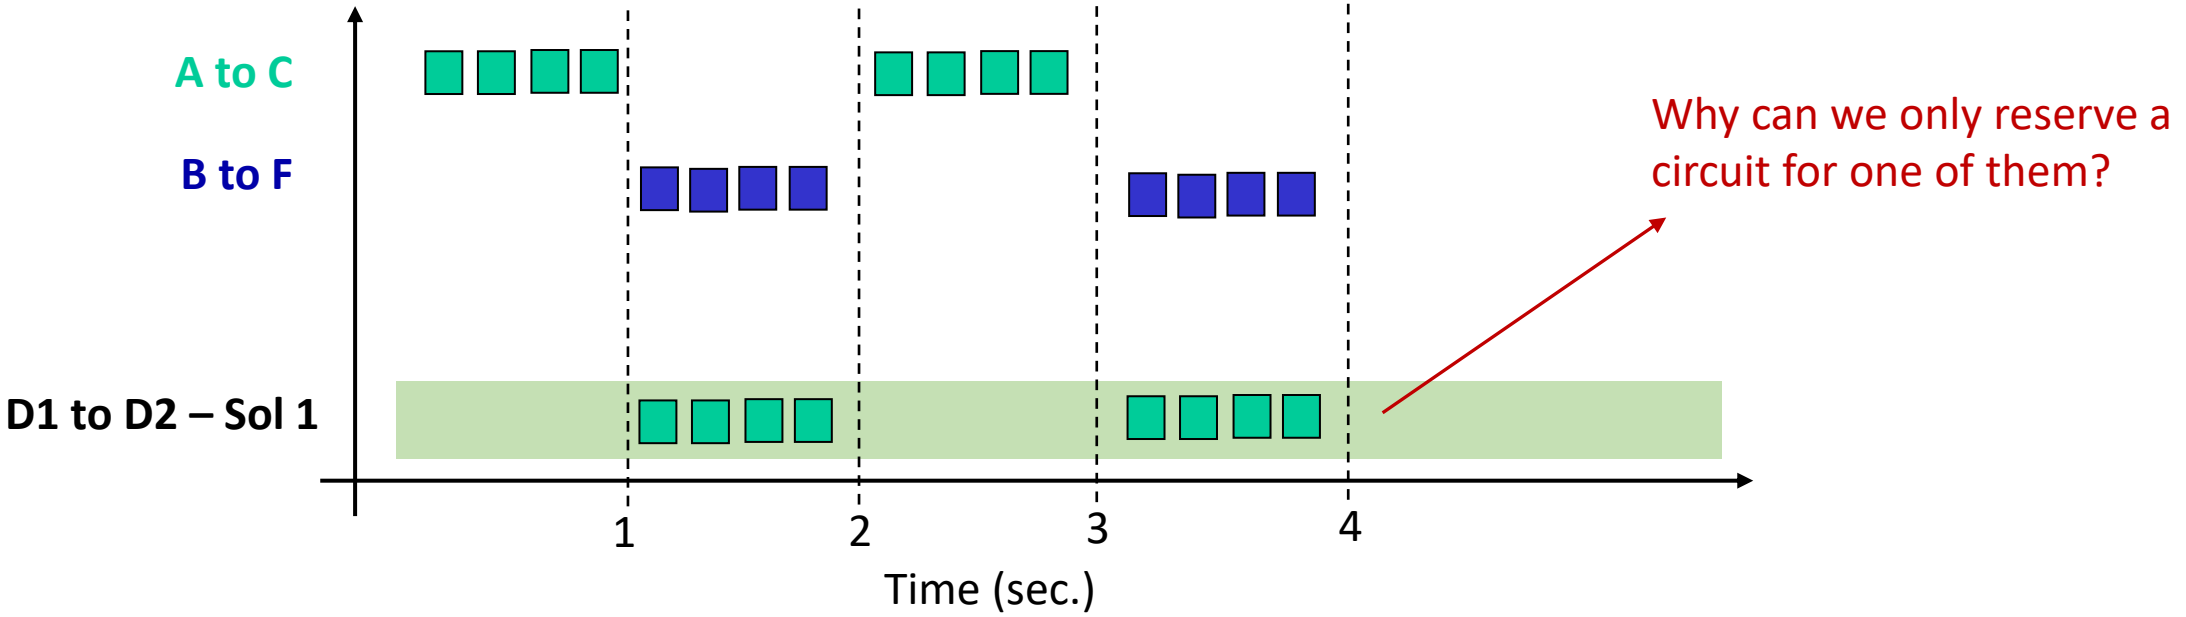

Solution 2: Send on demand

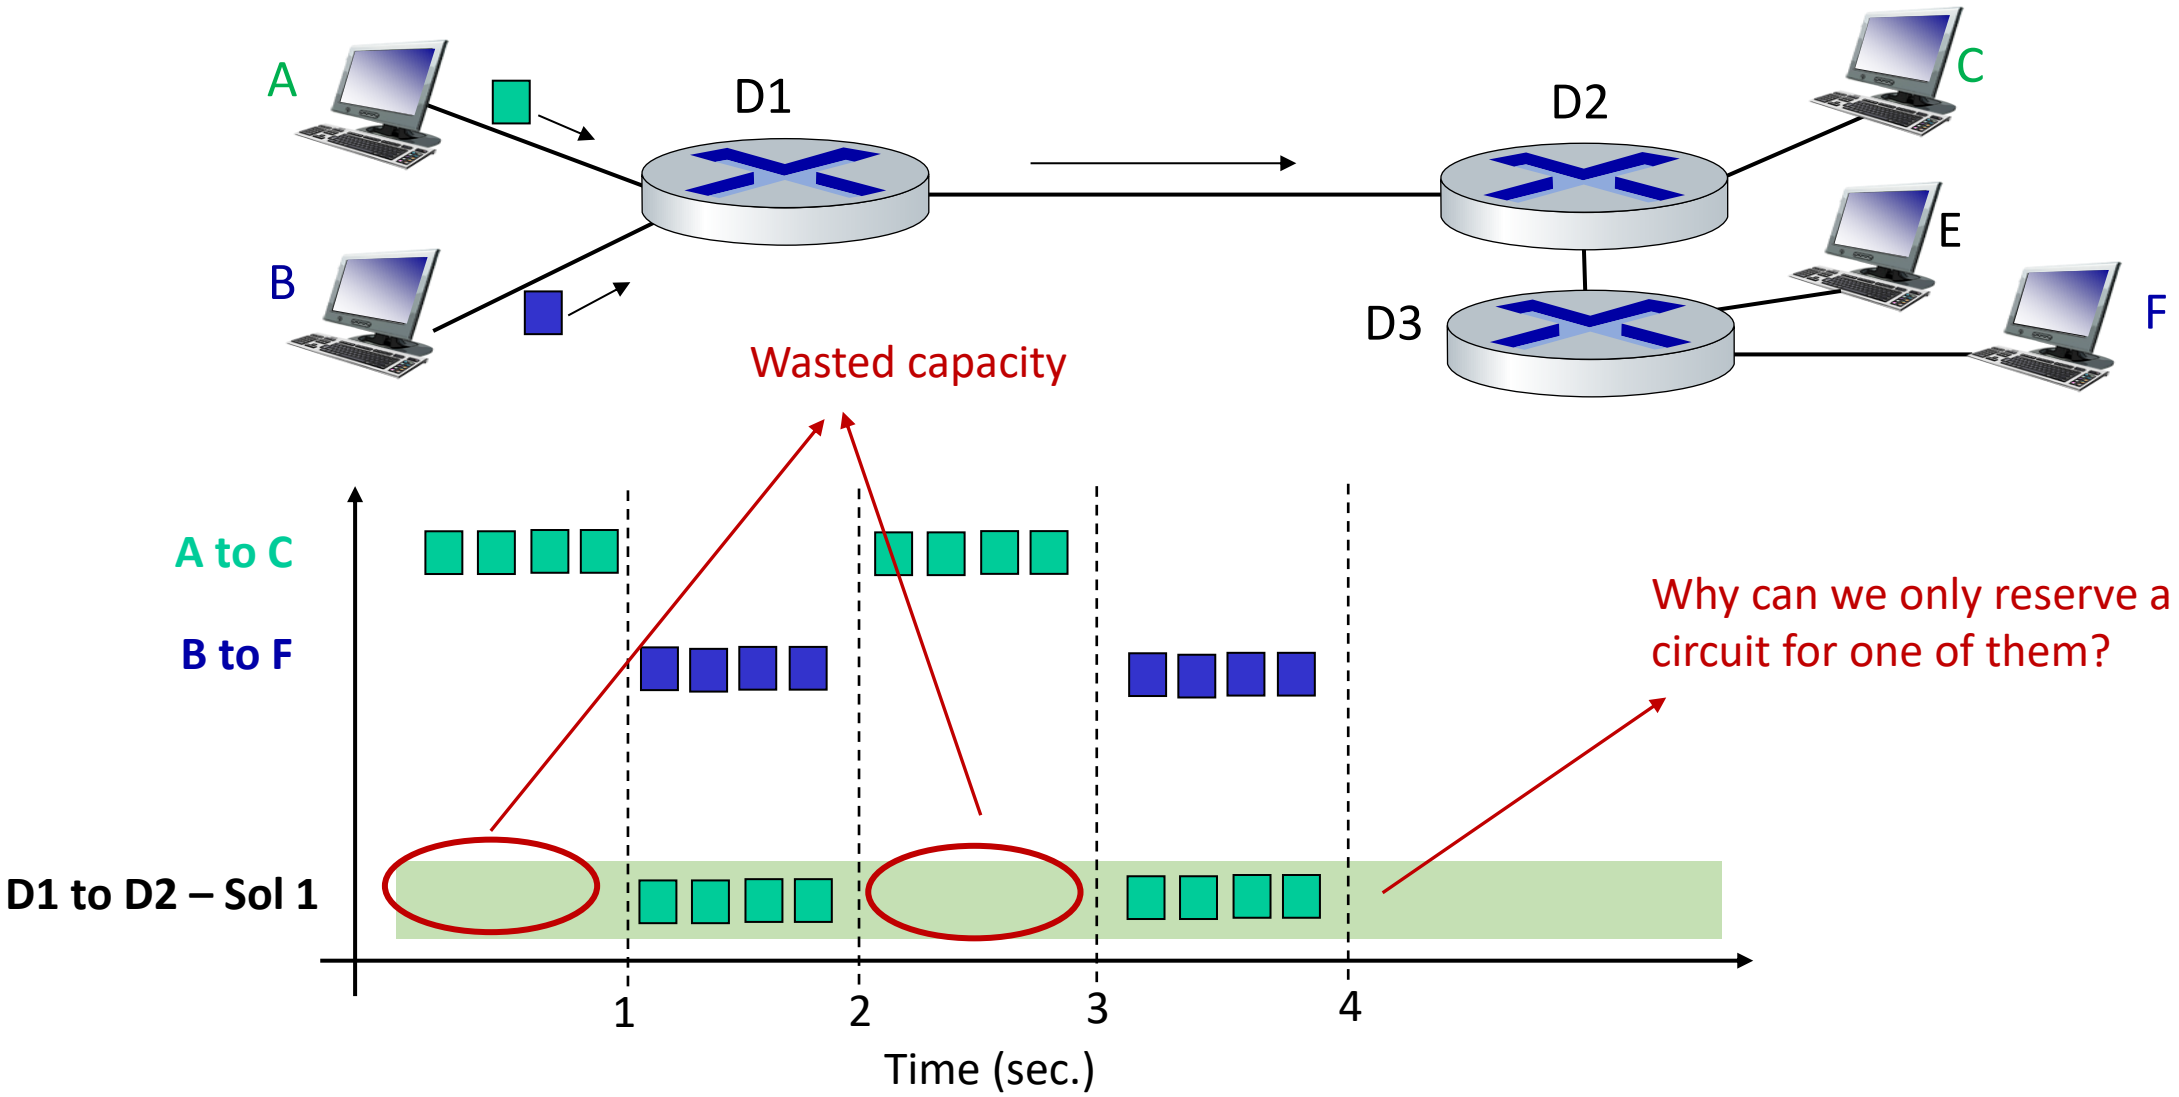

Solution 2: Send on demand

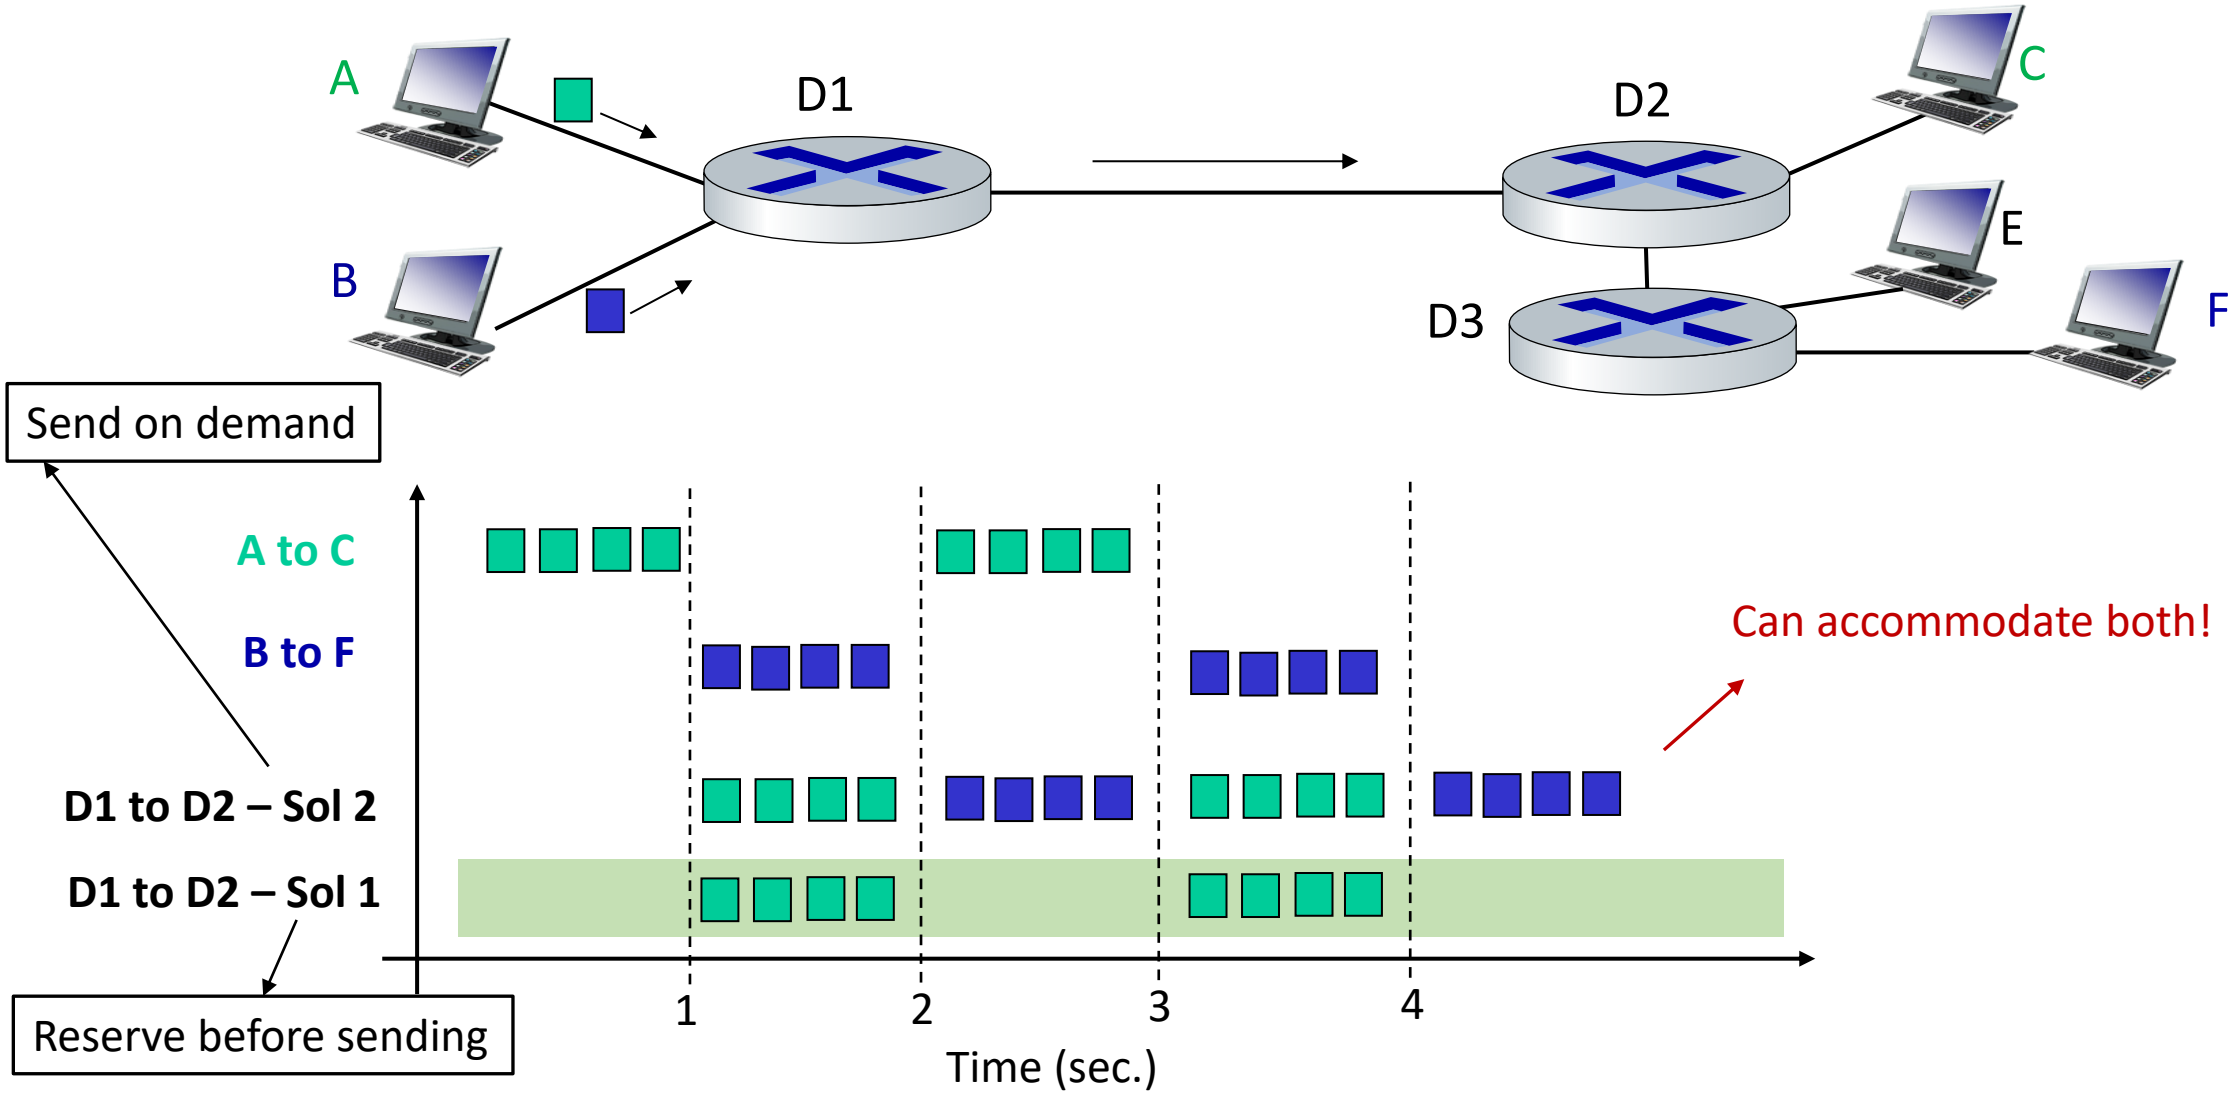

Solution 2: Send on demand

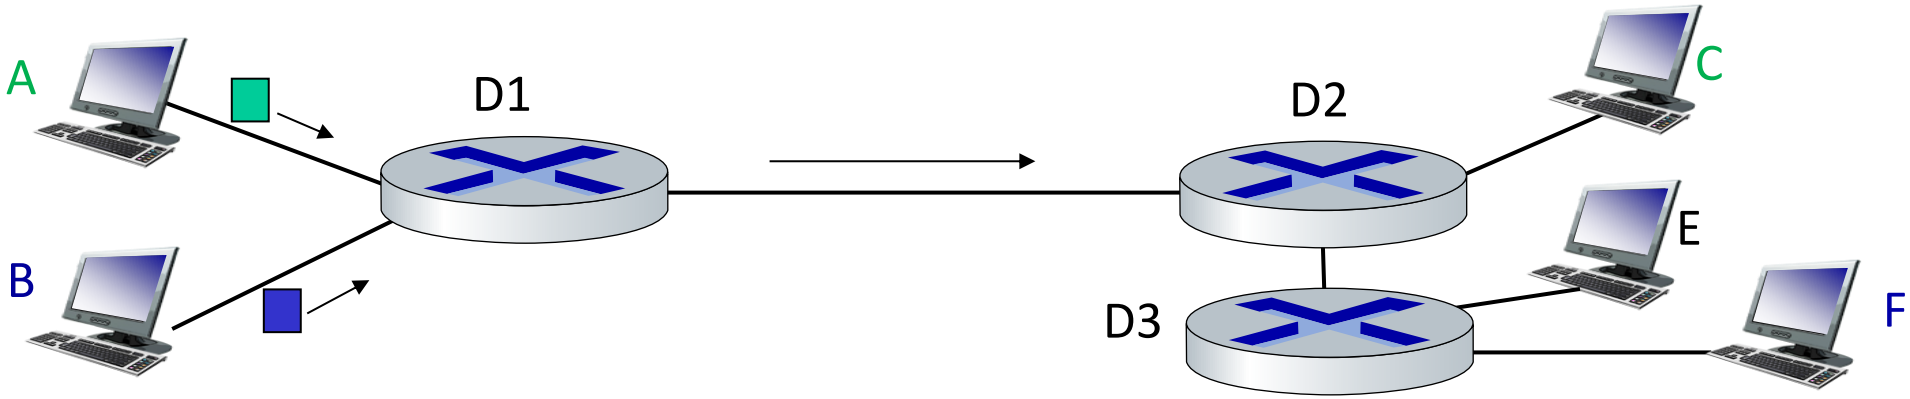

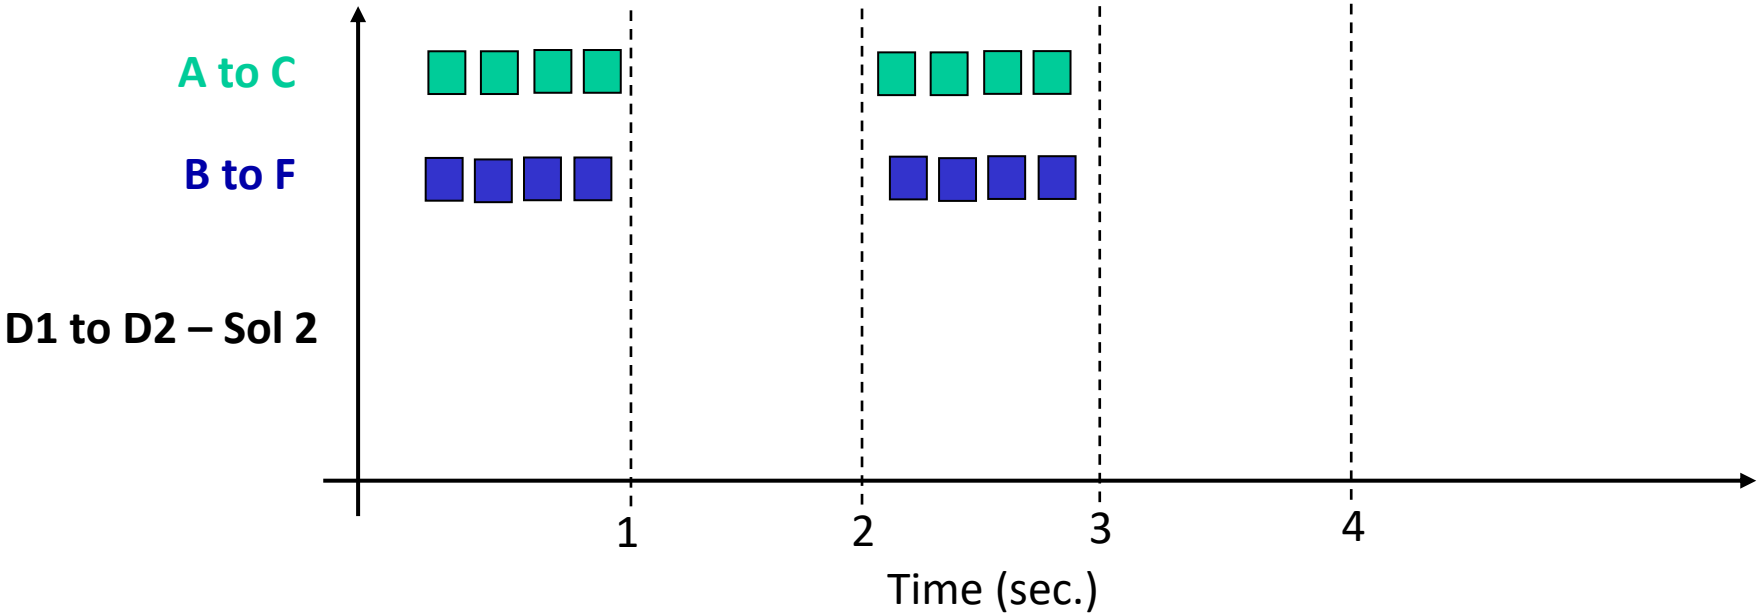

Solution 2: Send on demand

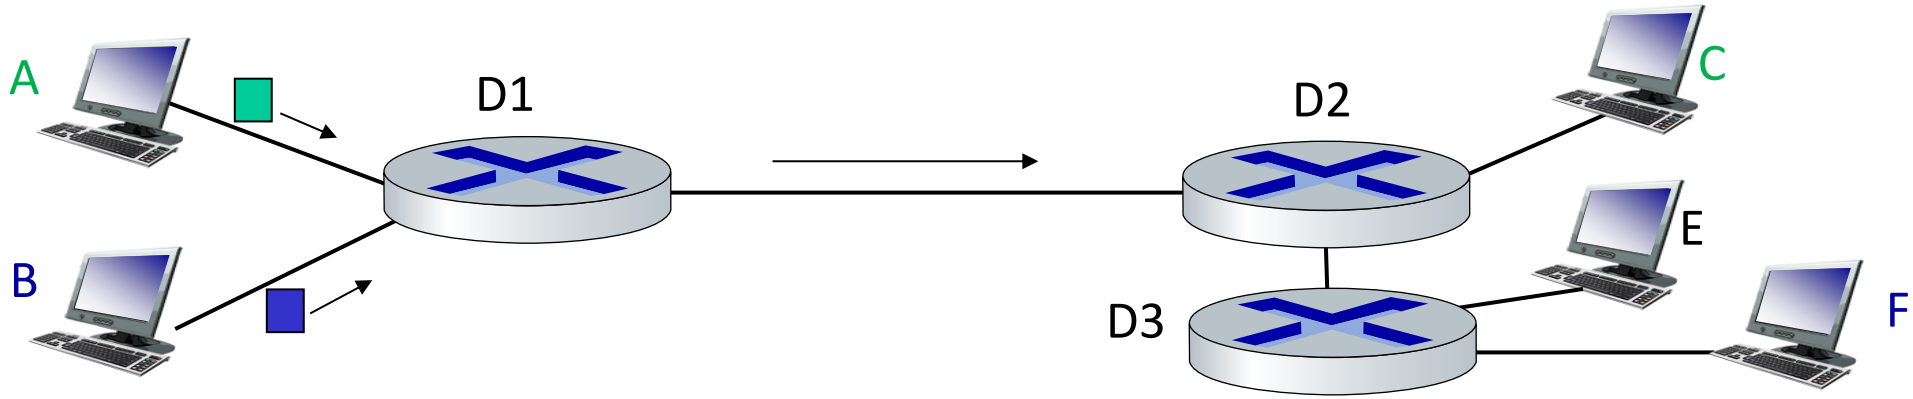

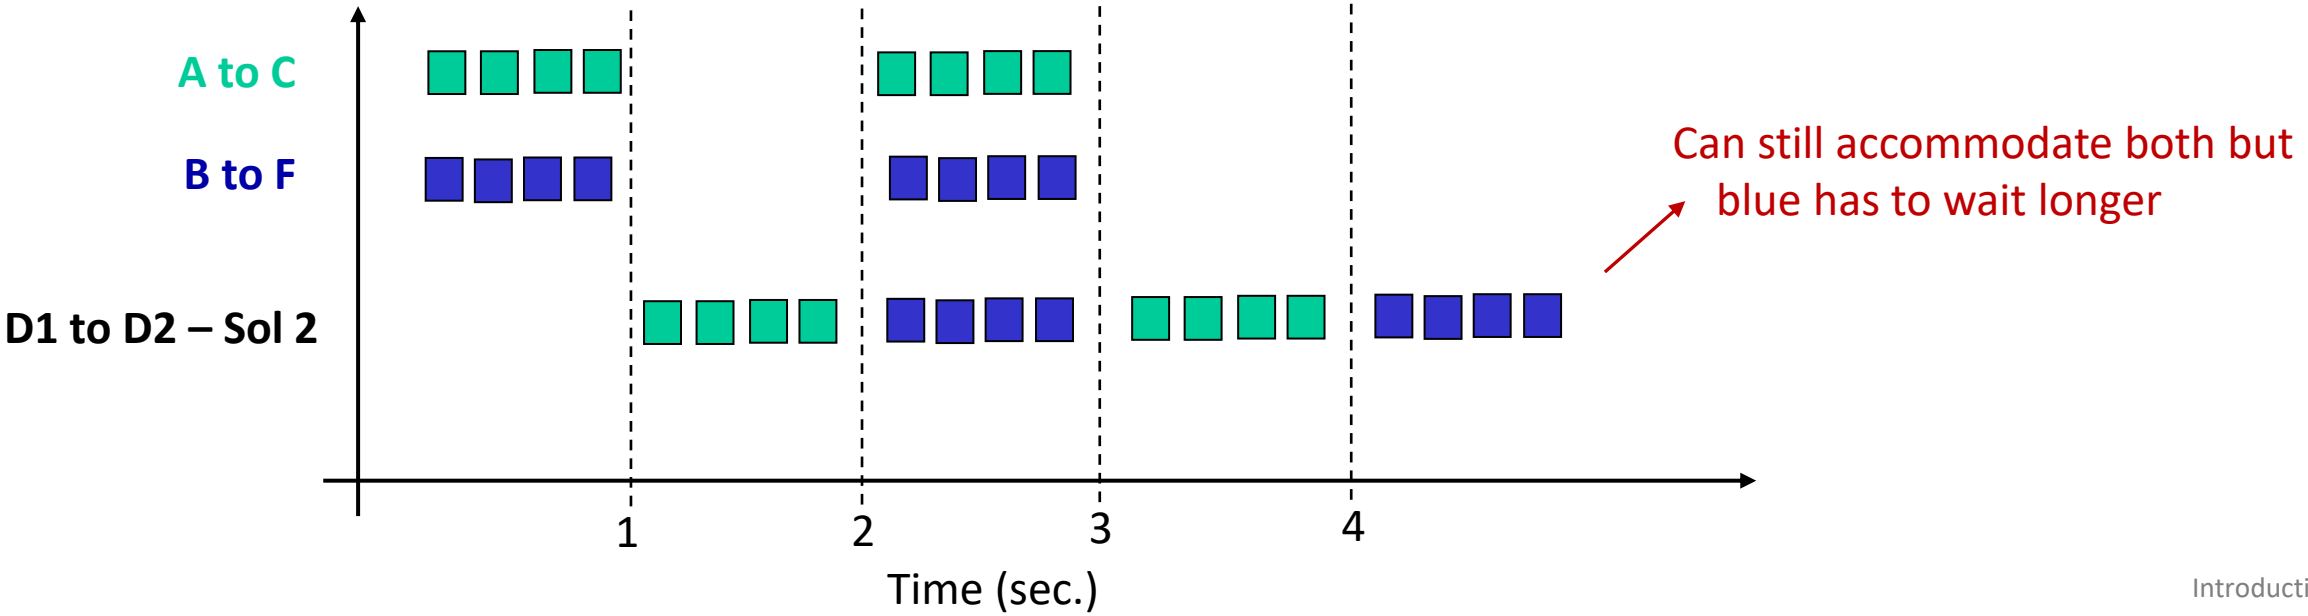

Solution 2: Send on demand

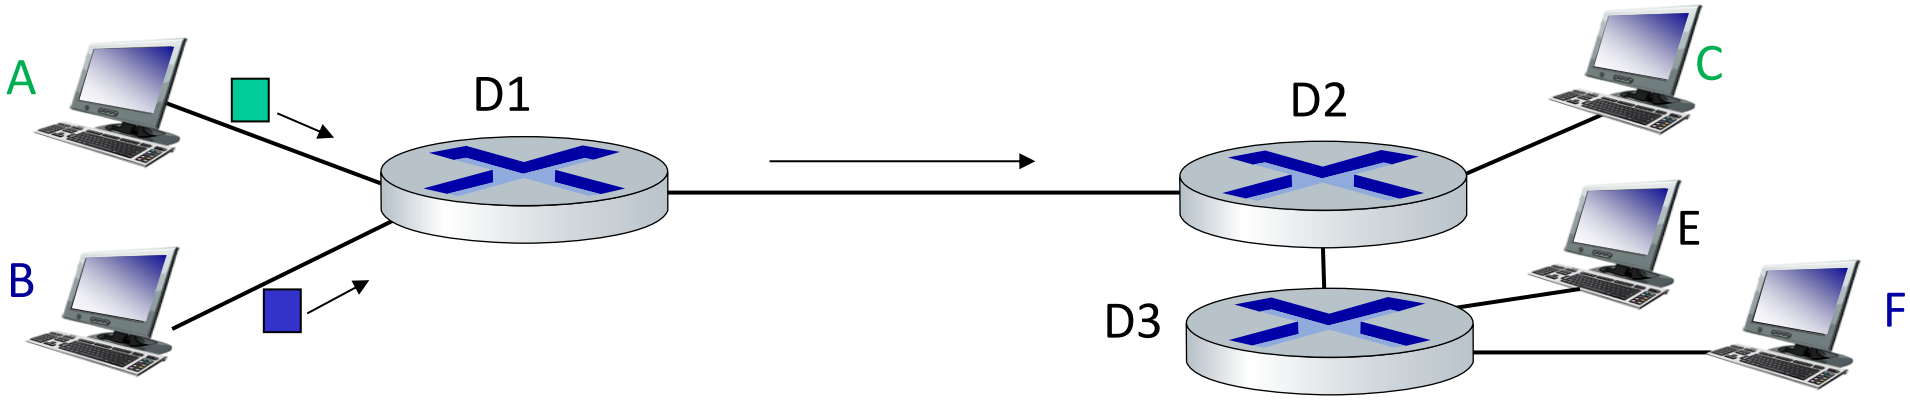

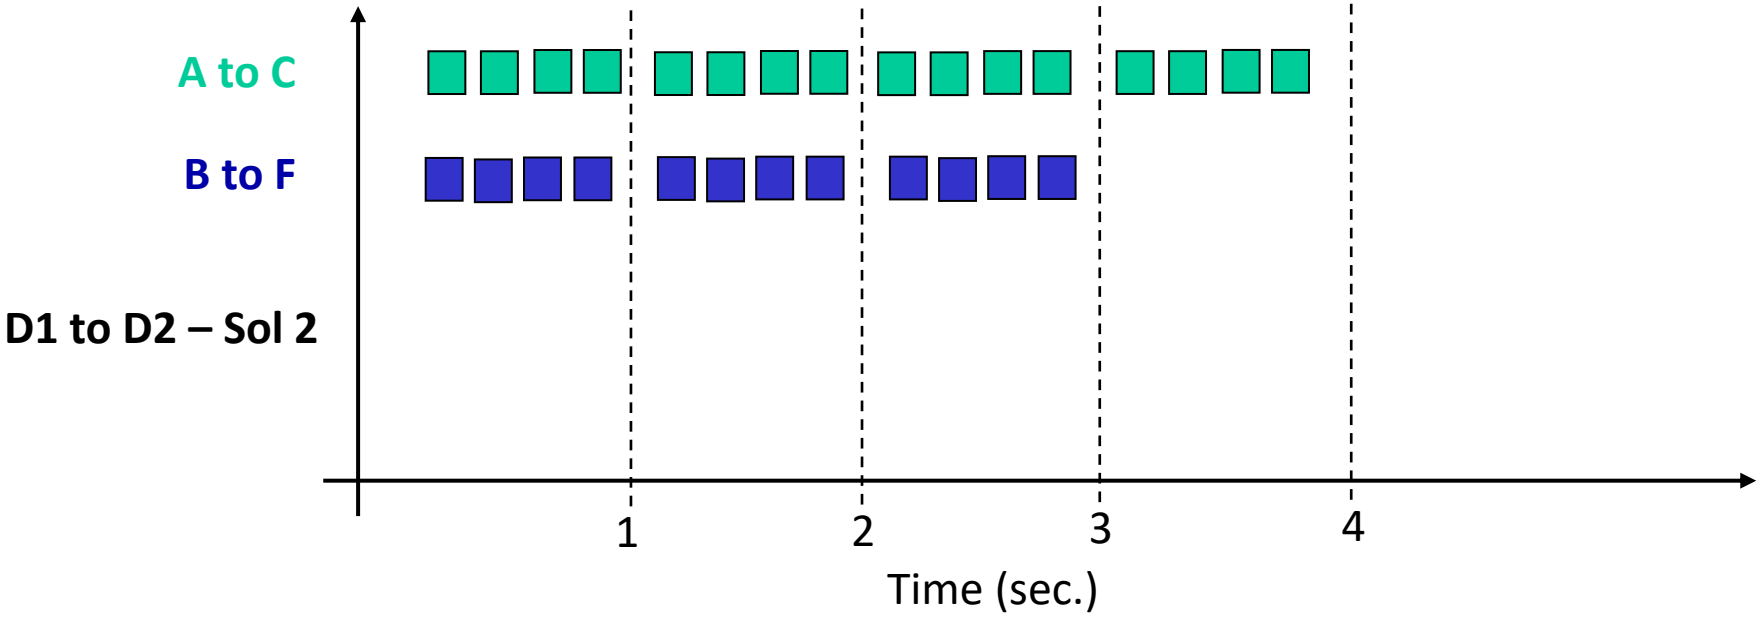

Solution 2: Send on demand

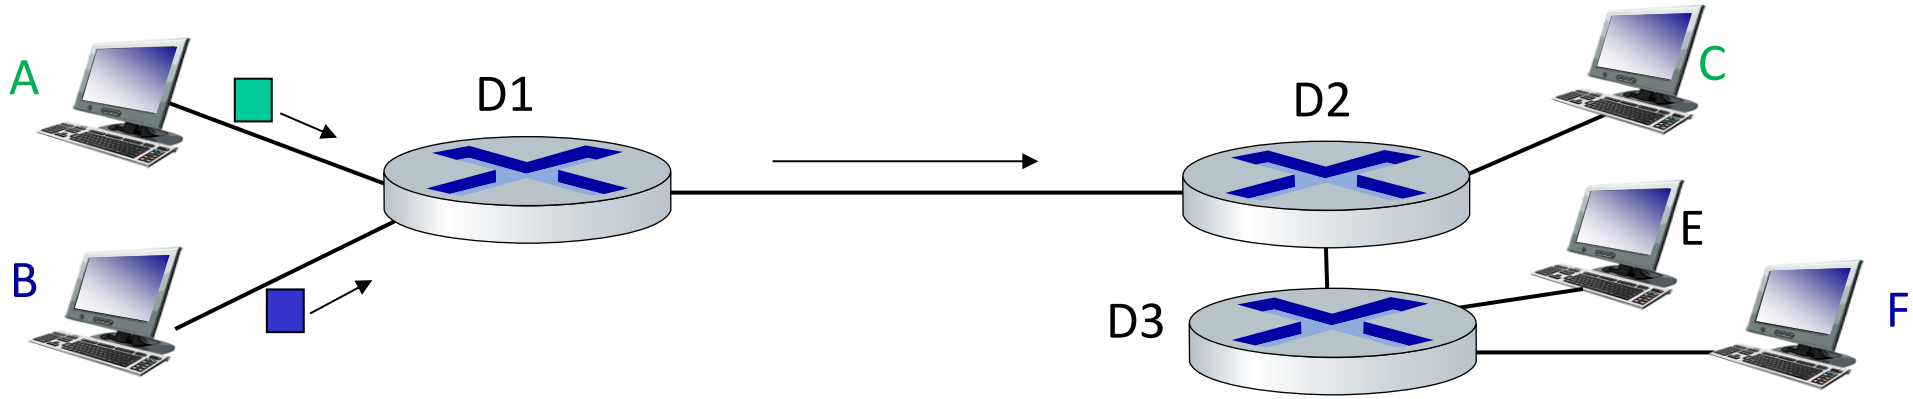

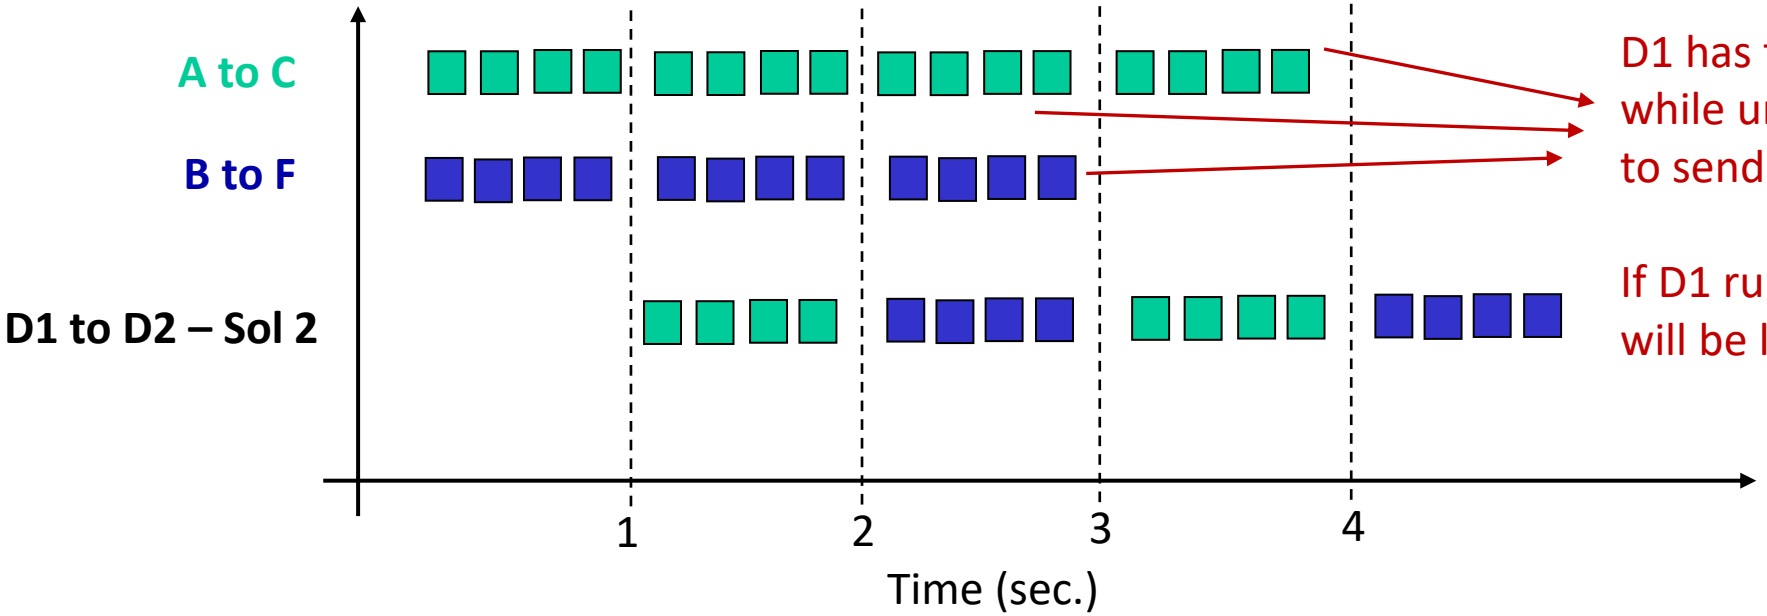

D1 has to hold on to these for a while until it has available slots to send them out.

If D1 runs out of memory, these will be lost…

Send on demand = Packet Switching

- No reservation needed.
- It allows more senders to send data simultaneously over the network.
- As long as their transmission pattern is “bursty” enough, the probability that all of them send at the same time is quite low, so the network should be able to handle that.
  - Statistical multiplexing
- This is the approach used in the Internet.

Circuit switching vs. Packet switching

Circuit switching

- Network resources (e.g., link capacity) divided into “pieces”
- Dedicated end-to-end resources
- Circuit-like guaranteed performance

Packet switching

- Each packet uses full link capacity
- Packets from different users *share* network resources
  - Resources used as needed
- Resource contention
  - Congestion: packets queue in network devices, and wait for link use

What you need to know about packet switching vs circuit switching

- You should understand and be able to explain the difference between packet switching and circuit switching with examples.

From this point on – Packet switching

- Unless stated otherwise, we assume a packet-switched network.
- This is because most networks that you interact with, including the Internet, use packet switching.

Communicating over a shared network

- To make that happen, networking people have to solve several challenging problems:
  - How to decide when a sender gets to transmit data?
  - How to pick good paths for getting data from its source to its destination?
  - How to adapt when a switch/router or a link fails?
  - …

We'll talk about the rest and the Internet, next time

Questions?

Course logistics – where to get what ☺

- The course webpage
  - https://student.cs.uwaterloo.ca/~cs456/
  - Course outline, tentative schedule, slides, policies, and references
- UWaterloo outline
  - Course outline, policies
- LEARN
  - Announcements, slides, quizzes, assignments
- Piazza
  - Questions and discussions

Course logistics – Lectures and office hours

- There are two sections in the course.
  - They will cover the same overall material by the end of the term.
  - But they may not be fully in sync.
  - Try to attend lectures of the same section throughout the term.
- Instructor office hours are generally for questions about the lectures and other topics related to computer networks
  - Mina Tahmasbi Arashloo: Tuesdays, 9:30-10:30am, in DC 3510
  - Bo Sun: Thursdays, 4:30-5:30pm, in DC 1311
- TA office hours are generally for questions about the assignments
  - Time and location will be announced on LEARN

Course logistics – Readings

- Readings are optional and from the following sources
  - Computer Networking, A Top-Down Approach, James Kurose & Keith Ross, Pearson, 8th Edition (7th Edition is also fine).
  - Research papers, blog posts, and platforms related to computer networks.
- Links to the latter will be posted on the course webpage

Course logistics – Assessment

	CS 456	CS 656
Quizzes	10%	10%
Programming Assignments	30%	10%
Midterm	25%	25%
Final	35%	35%
Research Project	-	20%

Course logistics – Quizzes

- Helps you assess your understanding of the course material
- Help us as your instructors to pinpoint subjects that need extra discussion in the class.
- Every week that we have lectures, a quiz will be posted on LEARN on Thursday that week at 8PM EST.
- You have until Sunday at 11:59PM EST to complete the quiz.
  - There are a few exceptions that will be announced the week of the quiz.

Course logistics – Quizzes

- Questions are multiple-choice and are about the material covered in class that week.
- You will have unlimited attempts. Hopefully, this helps you focus on your learning of the material without having to worry about your grade.
- A missed quiz cannot be retaken and doctor's note is not accepted for missed quizzes.
- You can miss two quizzes without it affecting your grade.

Course logistics – Programming Assignments (CS 456)

- There are three programming assignments, each counting as 10% of your final grade.
- Assignment specifications will be posted on LEARN.
- Assignments are to be completed individually and submitted to the appropriate Dropbox on LEARN.
- You have a total of four late days for the whole term.

Course logistics – Programming Assignments (CS 456)

- Use your late days to accommodate unexpected situations. No questions will be asked.
  - Examples of unexpected situations include your internet connection being down right before the deadline or forgetting or missing the deadline.
- Your four late days are for the whole term, not per assignment.
  - E.g.: If you use 2 late days for assignment 1, you have only 2 late days left to use for the rest of the assignments.

Course logistics – Programming Assignments (CS 456)

- Late days are tracked daily, not hourly.
  - E.g.: If an assignment is due on Thursday at 11:59pm, and you turn it in the next day (Friday) at noon, you have used one of your late days.
- The four late days include the self-declared short-term absence
  - E.g. if you have used 1 late day, and declare a two-day short-term absence, you will have 1 late day left for the rest of the term.
- No extensions are granted if you have no late days left, so plan ahead and use them carefully.
  - E.g.: If you have no late days left and submit the assignment after the deadline, you will not receive any points for the assignment

Course logistics – Programming Assignments (CS 656)

- Choose two of the three to complete.
- Each will count as 5% of your grade.
- The rest of the logistics is the same as CS 456.

Course logistics – Exams

- Exams will be solely based on the materials presented in the class
  - this includes the parts that may not necessarily be covered in the textbook or the slides.
- Exams are closed-book.
- The midterm exam is on Thursday, Oct 24, 7 to 8:50 PM.
- Date, time, and location of the final exam will be announced by the Registrar's Office.

Course logistics – Exams

- Midterm and final have to be passed, in the aggregate, in order to pass the course. That is, [(Midterm*25) + (Final*35)] / 60 >= 50%
- In the case of a missed exam, a medical certificate or doctor's note must be uploaded to the university's online portal.
  - Please refer to the course outline for information about acceptable certificates and doctor’s notes.

Course logistics – Midterm and self-declared absence

- If you miss the midterm due to an official self-declared short-term absence, the weight will be shifted to the final exam.
- The final exam will be counted as 60% of your grade.
- You need to score at least 50% in the final exam to pass the course.

Course logistics – Research Project (CS 656)

- Students taking CS 656 are expected to work individually on original research projects related to computer networks.
- There are three deliverables:
  - **Proposal (5%):** One page, due by the end of week 3. Please reach out to set up a time to discuss project ideas.
  - **Progress Report (5%):** Two pages, due a month after the proposal
  - **Final Report (10%):** 6-page conference-style paper, due at the end of the term

Course logistics – Q&A Policy

- We strongly encourage you to ask questions! However, this is a large class.
- We may not be able to answer duplicate questions
  - i.e., topics that are already discussed in the course outline, questions already answered on Piazza, or questions already answered in FAQs.
- Please check the course outline, the course FAQs, assignment FAQs, and Piazza before posting your question.

Course logistics – Q&A Policy continued

- We will do our best to answer your questions as soon as possible, but
  - it may take the teaching team up to 2 business days to answer emails and Piazza questions.
- Please take that into account when planning when to start working on assignments and studying for the exams.
- These two policies are in place to ensure the teaching team can effectively answer everyone's questions in a timely manner.

Course logistics – Attendance

- Attendance is not mandatory but we strongly encourage you to attend the lectures.
- While the slide will be available online, they are mostly intended as teaching aids for the lectures as opposed to detailed lecture notes.
- As such, they do not necessarily include all the details of the topics discussed in this course.

Course logistics – Generative AI

- Like any other tool, it should be used carefully and in a mindful, responsible manner.
- If you decide to use it in your programming assignments, it needs to be with proper documentation, citation, and acknowledgement.
  - Details will be outlines in the assignment instructions.
- You can find pointers to university guidelines and recommendations about using generative AI in the course outline.

Course logistics – Generative AI

- At the end of the day, you are accountable for the content and accuracy of your work, including any supported by generative AI.
- So, be mindful about where, how, and to what extent you use generative AI in your work.

Course logistics - Other

- Please make sure to read the outline carefully regarding
  - Mental health
  - Diversity
  - Academic integrity
  - Special accommodations
- No quiz this week.

Final remarks

- It gets quite boring to have a one-way conversation for 80 minutes ☺
- Your questions and thoughts are always welcome!
- Computer networks are a corner stone of modern society and I'm looking forward to go over the in and outs of how they work with all of you this term.