



CIS 553: Networked Systems

Network Programming and P4

January 22, 2020



Announcements

- HW0 due next week
 - Start early!

- Extra points for Piazza answers
 - Top 3 students with instructor-endorsed answers get 2 points extra on their final grade.




Agenda

- Networking Basics ← NEXT
- Layout of the Internet
- Programming a switch
- Design goals of the Internet



Networked Systems

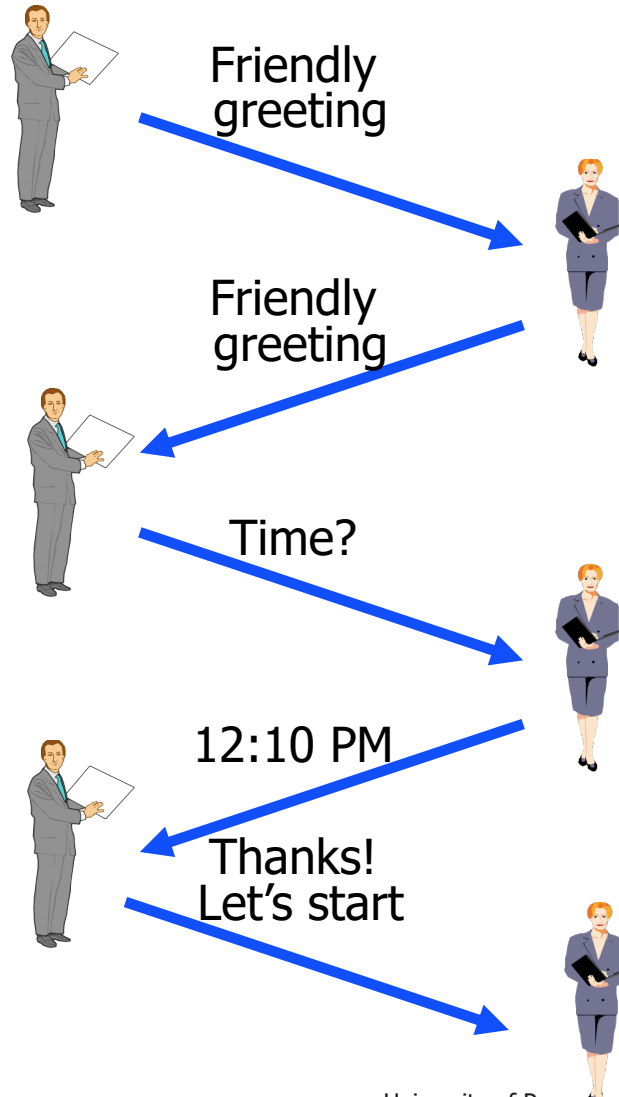
- Study of systems of **links** that interconnect **computers** in order to move **data**



- More simply: How computers communicate with one another



How do we communicate?



What is required to make this work?



Protocols

- An agreement between parties (in the same layer) on how to communicate
- Defines the **syntax** of communication
 - The sound will be broken up into English sentences
 - The sentences will be broken up into words
 - I will pause slightly between words
 - ...



Protocols

- An agreement between parties on how to communicate
- Defines the **syntax** of communication
- And **semantics**
 - “First a hello, then a request...”
- We will study many protocols
 - Exist at many levels (e.g., hardware and software)
 - Exist for many purposes (e.g., data transfer and routing control)
 - Defined by standards bodies like IETF, IEEE, ITU



Designing the Postal Service



- What is the syntax?
- What is the overall architecture?
- What does each member of the system do?




Agenda

- Networking Basics ✓
- Layout of the Internet ← NEXT
- Programming a switch
- Design goals of the Internet



The Internet shares similarities to the Postal Service





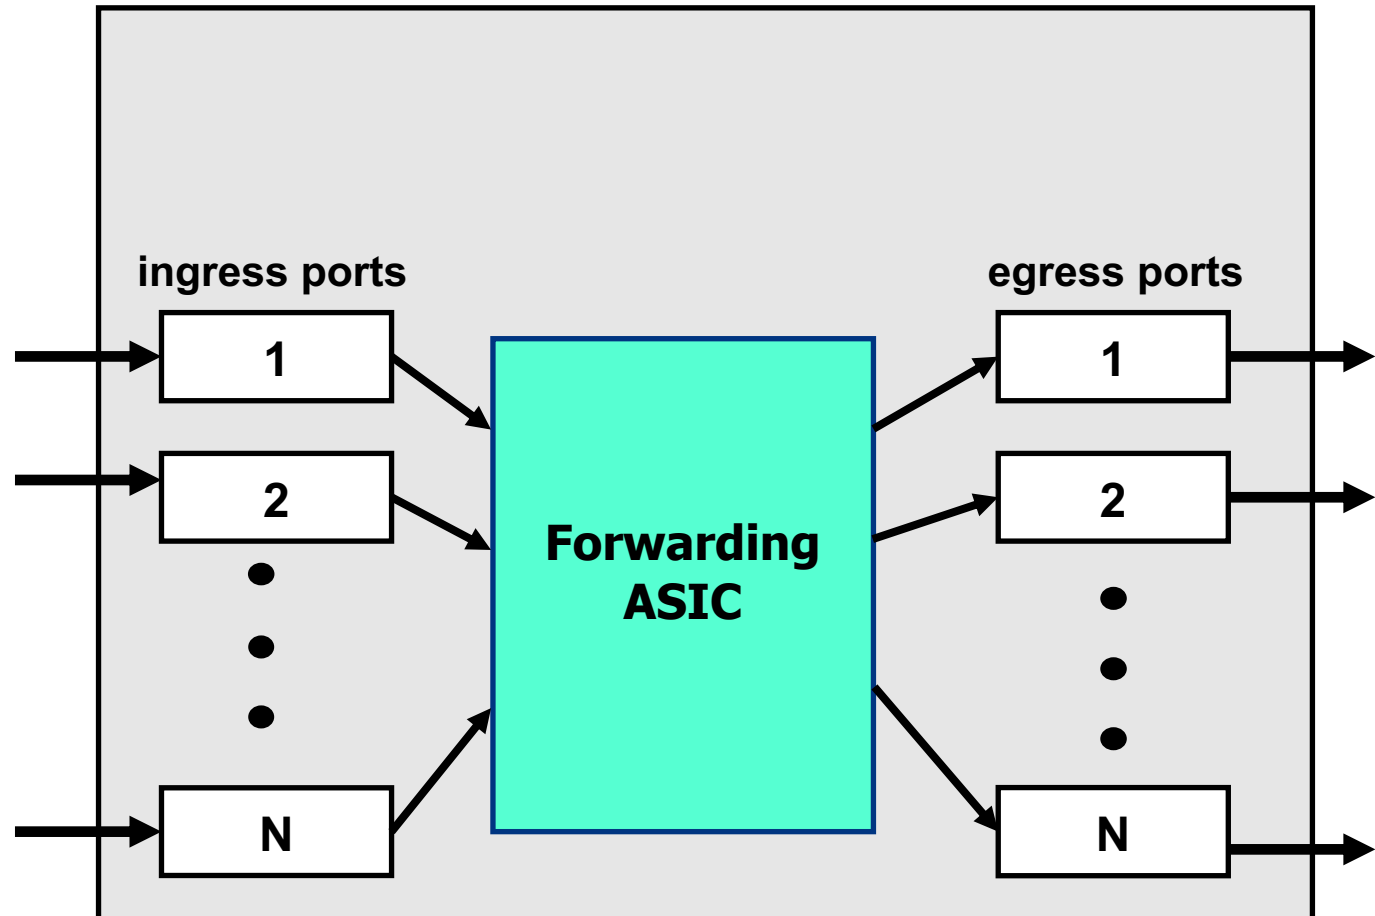
Network devices



- Highly specialized computers
- Connect two or more network **links**
- Forward **messages** from one link to another

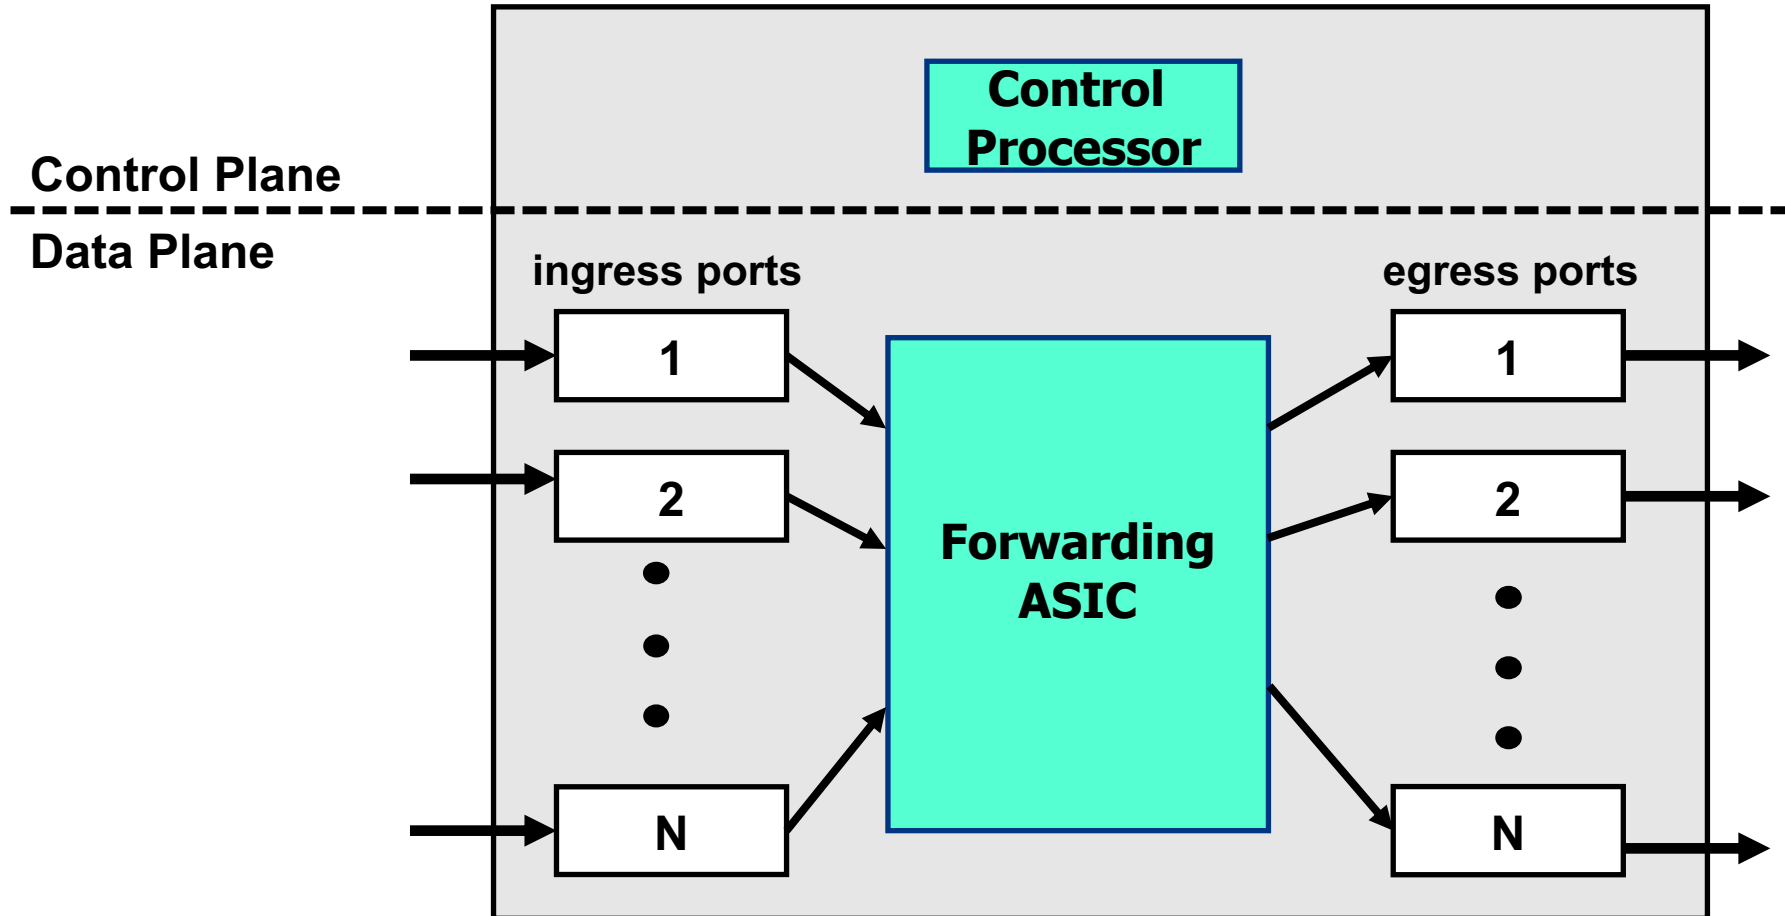


Network devices





Network devices







Different types of network devices

- Operate at different layers
- Ordered by increasing “intelligence”

- Repeaters:



- Hubs: multi-port repeaters




- Switches: can decide on port based on a simple forwarding table
- Routers: can coordinate amongst themselves to decide on a forwarding table for each router




Different types of network devices

- Operate at different layers
- Ordered by increasing “intelligence”

- Repeaters:



- Hubs: multi-port repeaters



- **Switches**: can decide on port based on a simple forwarding table
- **Routers**: can coordinate amongst themselves to decide on a forwarding table for each router



Constructing the Internet

- Project 1:
 - Learning Ethernet switch
- Project 2:
 - IPv4 forwarding
- Project 3/4:
 - IPv4 routing



Agenda

- Networking Basics ✓
- Layout of the Internet ✓
- Programming a switch ← NEXT
- Design goals of the Internet

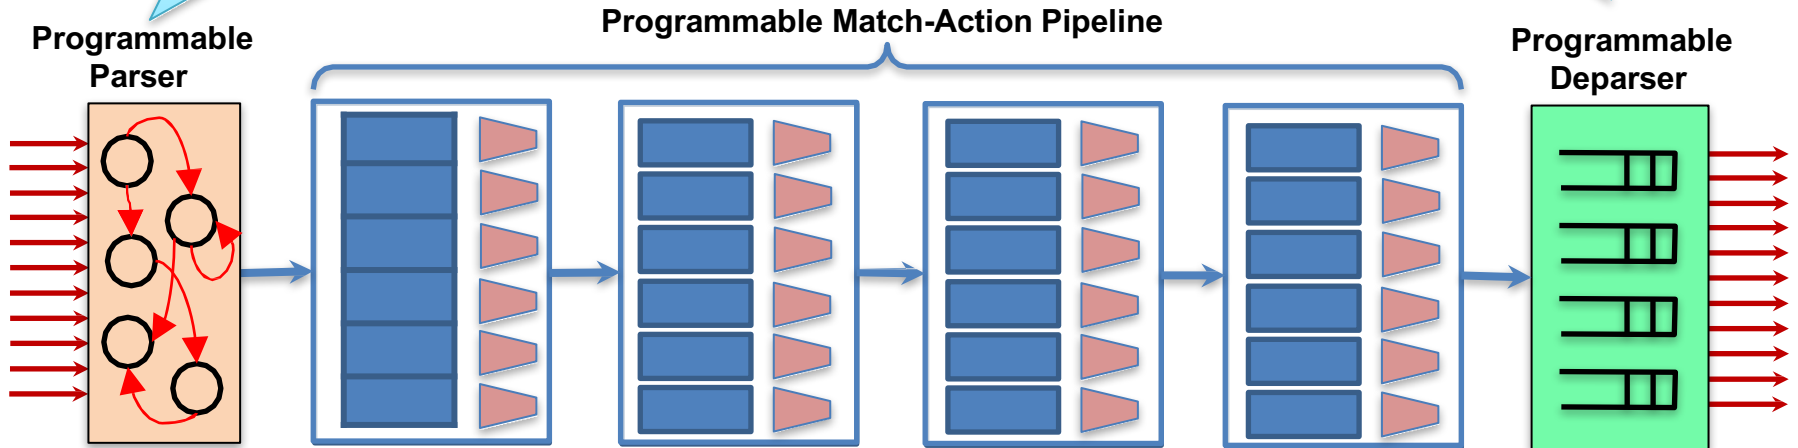


Programmable Switches

Programmer declares the headers that should be recognized and their order in the packet

Programmer defines the tables and the exact processing algorithm

Programmer declares how the output packet will look on the wire



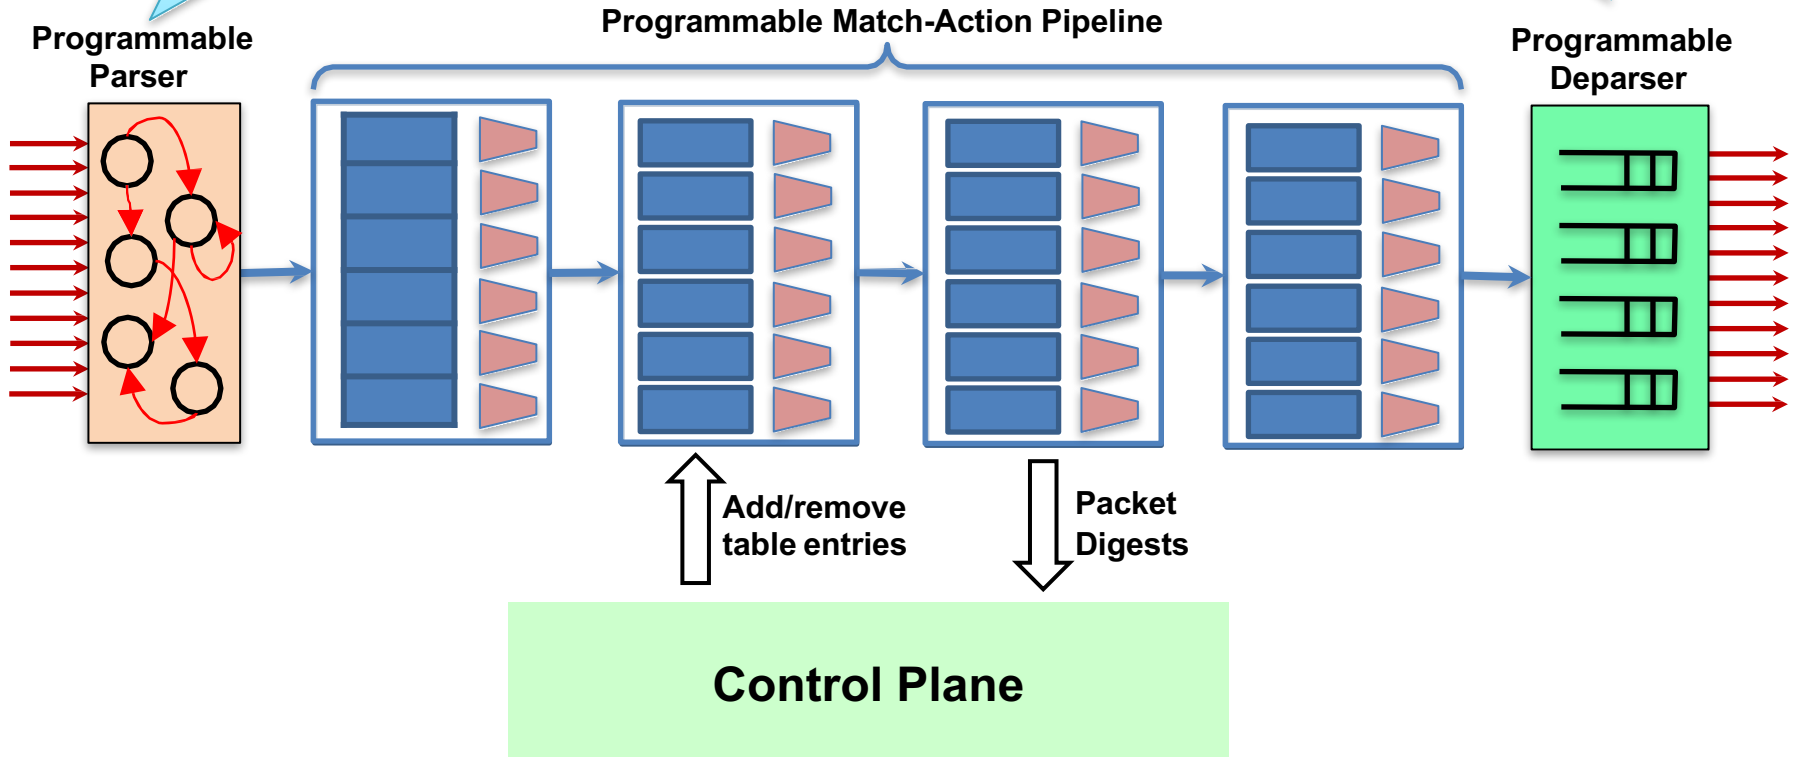


Programmable Switches

Programmer declares the headers that should be recognized and their order in the packet

Programmer defines the tables and the exact processing algorithm

Programmer declares how the output packet will look on the wire





The P4 language

- P4: Specifying packet processing logic
- Basic constructs: Header, Parser, Action, Table

```
/* Header definition */  
  
header ethernet_t {  
    bit<48> dstAddr;  
    bit<48> srcAddr;  
    bit<16> etherType;  
}  
struct headers_t {  
    ethernet_t ethernet;  
}
```

```
action next_hop(arg1, arg2, ...) {  
    // instructions  
}
```

```
/* Parser */  
  
parser parse_ethernet {  
    extract(ethernet);  
    return select(hdr.ethernet  
                  .etherType) {  
        0x800 : parse_ipv4;  
        0x806 : parse_arp;  
        default : ingress; } }
```

```
table forwarding_tbl {  
    key = { hdr.ethernet  
            .dstAddr : exact; }  
    actions { next_hop; _drop; } }
```



The Parser


- Imperative functions for “states”
- Extract header instances
- Select a next “state” by returning a parser function

```
parser parse_ethernet {
  extract(ethernet);
  return select(latest.etherType) {
    ETHERTYPE_CPU    : parse_cpu_header;
    ETHERTYPE_VLAN   : parse_vlan;
    ETHERTYPE_MPLS   : parse_mpls;
    ETHERTYPE_IPV4   : parse_ipv4;
    ETHERTYPE_IPV6   : parse_ipv6;
    ETHERTYPE_ARP    : parse_arp_rarp;
    ETHERTYPE_RARP   : parse_arp_rarp;
    ETHERTYPE_NSH    : parse_nsh;
  }
}
```

- Produces a *Parsed Representation* of the packet



Match+Action Tables

- Specification of:
 - What to examine from each packet 
 - What are the permitted actions that can be applied

```
table forwarding_tbl {  
    key = { hdr.ethernet.dstAddr : exact; }  
    actions { next_hop; _drop; }}
```



Match+Action Tables

- Specification of:
 - What to examine from each packet
 - What are the permitted actions that can be applied ←

```
table forwarding_tbl {  
    key = { hdr.ethernet.dstAddr : exact; }  
    actions { next_hop; _drop; }}
```



Control plane code

Python:

```
table_entry = p4info_helper.buildTableEntry(  
    table_name = "cis553Ingress.tab",  
    match_fields = {"hdr.foo.bar":  
                    0x1000},  
    action_name = "cis553Ingress.act",  
    action_params = {"egress_port": 2})  
  
sw.WriteTableEntry(table_entry)
```





Agenda

- Networking Basics ✓
- Layout of the Internet ✓
- Programming a switch ✓
- Design goals of the Internet ← NEXT
 - Taxonomy of networks
 - Design goals (Clark '88)
 - Layering and the end-to-end principle



Taxonomy of Communication Networks






Broadcast Networks



- Nodes share a common channel
- Everyone can hear everyone else

Switched Networks



- Information is transmitted to a small subset (usually one) of the nodes



Taxonomy of Communication Networks

