



**OST**

Eastern Switzerland  
University of Applied Sciences

# **Distributed Systems (DSy)**

## **Application Protocols**

Thomas Bocek

07.04.2022

# Learning Goals

- Lecture 7 (Application Protocols)
  - What protocols besides HTTP exist and why?
  - What are the advantages and how do they work?

# Protocols

- **Protocols, lecture 3,4**: layer 4
    - TCP, UDP, (QUIC/HTTP/3)
  - Designing custom protocols (e.g. **Kafka**)
  - Needs more time to develop / test
  - + Can be more efficient (space/performance)
  - Protocol generators (binary): Thrift / Avro / Protocol Buffers / (ASN1)
  - + IDL (interface description language) generates code
  - + Standard
  - Has more overhead
- e.g, Avro IDL - higher-level language for authoring Avro schemata → generates Avro schema

```
//Avro IDL
@namespace("ch.hsr.dsl")

protocol MyProtocol{
  record AMessage {
    string request;
    int code;
  }
  record BMessage {
    string reply;
  }


  BMessage GetMessage(AMessage msg);
}

{"namespace": "ch.hsr.dsl",
 "type": "record", "name":
"AMessage",
 "fields": [
  {"name": "request", "type":
"string"},
  {"name": "code", "type": "int"}
 ]
}
```

# Protocols

- Custom encoding/decoding
  - You control every aspect
  - You spend more time on it
- Little-endian / Big-endian
  - sequential order where bytes are converted into numbers
- Networking, e.g. TCP headers: Big-endian
- Most CPUs e.g., x86: Little-endian, RISC-V: Bi-endianness

```
115 public static boolean decodeHeader(final ByteBuf buffer, final InetAddress recipientSocket,
116     final InetAddress senderSocket, final Message message) {
117     LOG.debug("Decode message. Recipient: {}, Sender:{}", recipientSocket, senderSocket);
118     final int versionAndType = buffer.readInt();
119     message.version(versionAndType >>> 4);
120     message.type(Type.values()[versionAndType & Utils.MASK_0F]);
121     message.protocolType(ProtocolType.values()[versionAndType >>> 30]);
122     message.messageId(buffer.readInt());
123     final int command = buffer.readUnsignedByte();
124     message.command((byte) command);
125     final Number160 recipientID = Number160.decode(buffer);
126
127     //we only get the id for the recipient, the rest we already know
128     final PeerAddress recipient = PeerAddress.builder().peerId(recipientID).build();
129     message.recipient(recipient);
130
131
132     final int contentType = buffer.readInt();
133     message.hasContent(contentType != 0);
134     message.contentType(decodeContentType(contentType, message));
```



# Protocols Examples with Golang

- UDP example in repo [DSy](#)

- Why is it failing?

```
func main() {
    fmt.Println("connecting...")
    conn, _ := net.Dial("udp",
"127.0.0.1:7000")
    defer conn.Close()
    buf := make([]byte, 4)
    binary.LittleEndian.PutUint32(buf, 77)
    conn.Write(buf)
}

func main() {
    fmt.Println("listening...")
    inet := &net.UDPAddr{net.IPv4zero, 7000, ""}
    udpConn, _ := net.ListenUDP("udp", inet)
    b := make([]byte, 4)
    n, b2, _ := udpConn.ReadFromUDP(b);
    fmt.Printf("connecting... read: %d, addr: %v, data:
%v," +
    " decoded: %v\n", n, b2, b[:n],
binary.BigEndian.Uint32(b))
}
```

- TCP

- Custom serialization 5,Anybody there?

- 15 bytes

```
func main() {
    fmt.Println("connecting...")
    conn, _ := net.Dial("tcp", "127.0.0.1:7000")
    defer conn.Close()
    buf := make([]byte, 15)
    buf[0]=5
    copy(buf[1:], []byte("Anybody there?"))
    _, _ = conn.Write(buf)
}

func main() {
    fmt.Println("listening...")
    tcpConn, _ := net.Listen("tcp", ":7000")
    conn, _ := tcpConn.Accept() //do this in a go routine
    b := make([]byte, 15)
    n, _ := conn.Read(b)
    fmt.Printf("connecting... read: %d, addr: %v, data: [% x],
decoded: %v\n",
    n, conn.RemoteAddr(), b[:n], string(b[1:]))}
```

# Protocols Example ASN1

- ASN1

- Defined in 1984. Standard interface description language (IDL)
- Define data structures - can be serialized and deserialized
- Used e.g., in: X.509 (hsr.ch)
- Generic binary protocol, [Golang package](#)
- Example: 21 bytes, XML: 48 bytes

```
type TestASN struct {
    Code *big.Int
    Message string
}
...
func main() {
    ...
    var t TestASN
    _, err = asn1.Unmarshal(b, &t)
}
```

30 13 02 01 05 16 0e 41 6e 79 62 6f 64 79 20 74 68 65 72 65 3f

30 – type tag indicating SEQUENCE  
13 – length in octets of value that follows  
02 – type tag indicating INTEGER  
01 – length in octets of value that follows  
05 – value (5)  
16 – type tag indicating IA5String  
(IA5 means the full 7-bit ISO 646 set, including variants,  
but is generally US-ASCII)  
0e – length in octets of value that follows  
41 6e 79 62 6f 64 79 20 74 68 65 72 65 3f – value  
("Anybody there?")

<code>5</code>  
<message>Anybody there?</message>



# Protocols Example Avro

- [Avro](#): data serialization system
  - Remote procedure call and data serialization framework
  - Use: Hadoop (Big-data framework)
- LinkedIn [go-avro](#) library
  - Define a message in JSON ([benchmarks](#)) or IDL (slide 12) – no code generation

```
func main() {
    schema, _ := ioutil.ReadFile("schema.avsc")
    codec, _ := goavro.NewCodec(string(schema))
    map1 := map[string]interface{}{
        "request": "Anybody there?",
        "code": 5,
    }
    binary, _ := codec.BinaryFromNative(nil, map1)
    conn, _ := net.Dial("tcp", "127.0.0.1:7000")
    defer conn.Close()
    n, _ := conn.Write(binary)
    fmt.Printf("wrote %d bytes: [% x]\n", n,
binary)
}
```

- Server
  - Example 16 bytes, assuming both have the same IDL

```
{"namespace": "ch.hsr.dsl",
 "type": "record", "name": "AMessage",
 "fields": [
     {"name": "request", "type": "string"},
     {"name": "code", "type": "int"}
 ]
}
```

```
func main() {
    schema, _ := ioutil.ReadFile("schema.avsc")
    codec, _ := goavro.NewCodec(string(schema))
    tcpConn, _ := net.Listen("tcp", ":7000")
    conn, _ := tcpConn.Accept() //do this in a go routine
    binary := make([]byte, 100)
    n, _ := conn.Read(binary)
    native, _, _ := codec.NativeFromBinary(binary[:n])
    fmt.Printf("read: %v\n", native)
}
```

# Protocol Buffers Example

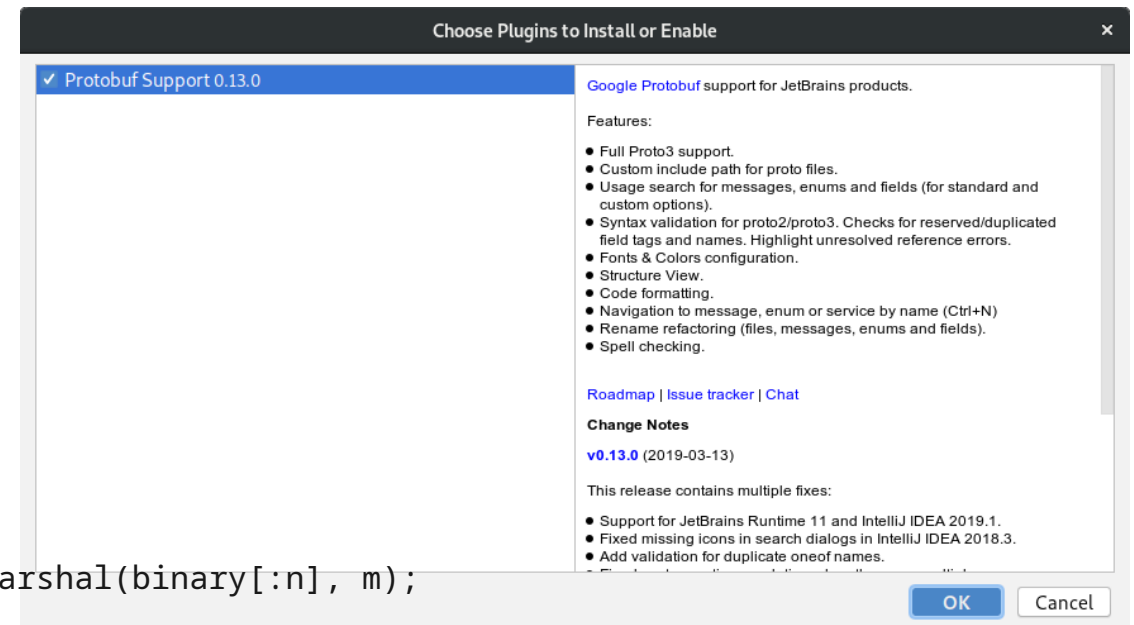
- **Protobuf**: data serialization system from Google
  - Design goals: smaller and faster than XML
  - Use: nearly all inter-machine communication at Google
- Generate 1 go file
  - `protoc schema.proto --go_out .`
  - integers to identify fields. Protocol buffer contains only numbers, not field names
- 18 bytes

```
m := &pb.AMessage{Id: 5, Message: "Anybody there?"}
out, err := proto.Marshal(m)
```

```
m := &pb.AMessage{}
if err := proto.Unmarshal(binary[:n], m);
err != nil {
    panic(err)
}
```

- Not self-describing, but has gzipped description

```
syntax = "proto3";
message AMessage {
    int32 code = 1;
    string message = 2;
}
```







# RPC Example Thrift

- RPC Framework from Facebook

- IDL and binary protocol
- For building cross-platform application in ActionScript, C, C++, C#, Cappuccino, Cocoa, Delphi, Erlang, Go, Haskell, Java, Node.js, Objective-C, OCaml, Perl, PHP, Python, Ruby, Rust, Smalltalk, and Swift

- Installation

- `sudo apt install thrift-compiler`
- `thrift -r --gen go schema.thrift`
- Creates go files (client)

```
service TestService {  
    void AMessage(1:i32 int, 2:string message)  
}
```

- Example generic server

- `go run simple-gen-srv.go`
- `go run simple-thrift.go -p 7000`
- `go run simple-thrift.go -p 7000 AMessage 5 'Anybody there?'`

- 49 bytes transferred

- Thrift also encodes which function to call, larger size

# RPC Example gRPC

- gRPC
  - Uses **HTTP/2** for transport
  - Uses Protocol Buffers
  - Features: authentication, bidirectional streaming and flow control, blocking or nonblocking bindings, and cancellation and timeouts, many **languages**
  - Installation
    - go get -u google.golang.org/grpc
    - go get -u github.com/golang/protobuf/protoc-gen-go
    - protoc schema-grpc.proto --go\_out=plugins=grpc:go-gen3
  - 171 / 124 (wireshark)

```
grpcServer := grpc.NewServer()
s := Server{}
schemagrpc.RegisterMessageServiceServer(grpcServer, &s)
grpcServer.Serve(tcpConn)
```

- Define services and message
  - Generate 1 source file with functions (service)

```
syntax = "proto3";

service MessageService {
  rpc SendMessage (AMessage)
  returns (Empty);
}

message AMessage {
  int32 code = 1;
  string message = 2;
}

message Empty {}

var conn *grpc.ClientConn
conn, _ := grpc.Dial(":7000", grpc.WithInsecure())
if err != nil {
  log.Fatalf("did not connect: %s", err)
}
defer conn.Close()
c := schemagrpc.NewMessageServiceClient(conn)
response, _ := c.SendMessage(context.Background(),
  &schemagrpc.AMessage{Code:5,Message:"Anybody
  there?"})
```

# JSON example

- JSON + REST/HTTP
  - Human-readable text to transmit data
  - Often used for web apps
- 187 bytes

```
func main() {
    fmt.Println("Connecting...")
    req, _ := http.NewRequest("POST",
        "http://localhost:7000",
        strings.NewReader(`{"code": 5,"message": "Anybody
there?"}`))
    req.Header.Set("Content-Type", "application/json")
    client := &http.Client{}
    resp, err := client.Do(req)
    if err != nil {
        panic(err)
    }
    defer resp.Body.Close()
    fmt.Printf("wrote request")
}
```

- Parsing overhead, JSON slower than binary protocol - **benchmarks**

```
[
    {
        "id": "bitcoin",
        "name": "Bitcoin",
        "symbol": "BTC",
        "rank": "1",
        "price_usd": "9324.08",
        "price_btc": "1.0",
        "24h_volume_usd": "9039300000.0",
        "market_cap_usd": "158560288125",
        "available_supply": "17005462.0",
        "total_supply": "17005462.0",
        "max_supply": "21000000.0",
        "percent_change_1h": "0.46",
        "percent_change_24h": "-0.27",
        "percent_change_7d": "4.5",
        "last_updated": "1525011874"
    }, ...
]
```

# Application Protocol: HTTP

- HTTP (**HyperText Transfer Protocol**): foundation of data communication for www
- Started in 1989 by Tim Berners-Lee
  - HTTP/1.1 published in 1997
  - HTTP/2 published in 2015
    - More efficient, header compression, multiplexing
  - HTTP/3 wip (April 2022: HTTP/3 protocol is an Internet Draft – not yet final)
- Request / response (resource)
- HTTP resources identified by URL
  - [https://dsl.hsr.ch/design/hsr\\_logo.svg](https://dsl.hsr.ch/design/hsr_logo.svg)

- Text-based protocol

```
openssl s_client -connect dsl.hsr.ch:443
... TLS handshake ...
GET /
```

- Browser sends a bit more...

```
▼ Request Headers (359 B)
Host: dsl.hsr.ch
User-Agent: Mozilla/5.0 (X11; Linux x86_64; rv:73.0) Gecko/20100101 Firefox/73.0
Accept: text/html,application/xhtml+xml,application/xml;q=0.9,image/webp,*/*;q=0.8
Accept-Language: en-US,en;q=0.5
Accept-Encoding: gzip, deflate, br
DNT: 1
Connection: keep-alive
Upgrade-Insecure-Requests: 1
Cache-Control: max-age=0
TE: Trailers
```

Scheme    User info    Host    Port    Path    Query    Fragment

<http://tbocek:password@dsl.hsr.ch:443/lect/fs21?id=1234&lang=de#topj>

# Application Protocol: HTTP

- Response
  - Header

## ▼ Response Headers (227 B)

```
HTTP/2 200 OK
server: cloudflare-nginx
content-type: text/html; charset=UTF-8
date: Mon, 02 Mar 2020 14:29:39 GMT
x-page-speed: 1.13.35.2-0
cache-control: max-age=0, no-cache
content-encoding: gzip
X-Firefox-Spdy: h2
```

- Status Code: 200
  - **List:** from 1xx (information response), 2xx (success) – 200 OK, 3xx (redirection), 4xx (client error), 404 Not Found, 403 Forbidden (access slides outside HSR), 5xx (server errors)

- Content

```
<!DOCTYPE html>
<html>
<head>
  <title>Distributed Systems and Ledgers Lab</title>
  <link rel="stylesheet" type="text/css"
href="design/layout.css"/>
```

- HTTP is a stateless protocol
  - Server maintains no state
- **Request methods:** GET, HEAD, POST, PUT, DELETE, TRACE, OPTIONS, CONNECT, PATCH
- Web server **one-liner** - with netcat:
  - `while true; do { echo -e "HTTP/1.1 200 OK\n\n<h1> Hallo" } | nc -vl -p 8080 -c; done`
- Every Webbrowser has dev tools to show request / responses
  - Firefox, Chrome: ctrl+shift+i / F12
  - Used regularly

# Protocols Bencoding and Others

- **Bencoding**
  - Integers: i42e, Byte string: 4:test, list: l4:testi42ee
  - Map/dictionary: d4:test3:hsr3:tomi42ee
- Use: BitTorrent
- **UBJSON**
- **Cap'n Proto** , **FlatBuffers**
- **Apache Arrow**
  - Do not serialize, just copy, little-endian
- ... and many **others**
- **Benchmarks**, **benchmarks**, ...

The screenshot shows the GHex application window titled "ubuntu-18.04-desktop-amd64.iso.torrent - GHex". The main area displays a hex dump of a bencoded string. The hex data is: 0000000064 38 3A 61 6E 6E 6F 75 6E 63 65 33 39 3A 68 74 d8:announce39:ht  
0000001074 70 3A 2F 2F 74 6F 72 72 65 6E 74 2E 75 62 75 tp://torrent.ubu  
000000206E 74 75 2E 63 6F 6D 3A 36 39 36 39 2F 61 6E 6E ntu.com:6969/ann  
000000306F 75 6E 63 65 31 33 3A 61 6E 6E 6F 75 6E 63 65 ouncement13:announce  
000000402D 6C 69 73 74 6C 6C 33 39 3A 68 74 74 70 3A 2F -list1139:http:/  
000000502F 74 6F 72 72 65 6E 74 2E 75 62 75 6E 74 75 2E /torrent.ubuntu.  
0000006063 6F 6D 3A 36 39 36 39 2F 61 6E 6E 6F 75 6E 63 com:6969/announc  
0000007065 65 6C 34 34 3A 68 74 74 70 3A 2F 2F 69 70 76 eel44:http://ipv  
0000008036 2E 74 6F 72 72 65 6E 74 2E 75 62 75 6E 74 75 6.torrent.ubuntu  
000000902E 63 6F 6D 3A 36 39 36 39 2F 61 6E 6E 6F 75 6E .com:6969/announ  
000000A063 65 65 65 37 3A 63 6F 6D 6D 65 6E 74 32 39 3A ceee7:comment29:  
000000B055 62 75 6E 74 75 20 43 44 20 72 65 6C 65 61 73 Ubuntu CD releas  
000000C065 73 2E 75 62 75 6E 74 75 2E 63 6F 6D 31 33 3A es.ubuntu.com13:  
000000D063 72 65 61 74 69 6F 6E 20 64 61 74 65 69 31 35 creation datei15  
000000E032 34 37 37 36 33 30 38 65 34 3A 69 6E 66 6F 64 24776308e4:infod  
000000F036 3A 6C 65 6E 67 74 68 69 31 39 32 31 38 34 33 6:lengthi1921843  
000010032 30 30 65 34 3A 6E 61 6D 65 33 30 3A 75 62 75 200e4:name30:ubu  
00001106E 74 75 2D 31 38 2E 30 34 2D 64 65 73 6B 74 6F ntu-18.04-deskto  
000012070 2D 61 6D 64 36 34 2E 69 73 6F 31 32 3A 70 69 p-amd64.iso12:pi  
000013065 63 65 20 6C 65 6E 67 74 68 69 35 32 34 32 38 ece lengthi52428

Below the hex dump, there is a table of decoded values for the selected hex range (0000000064 to 000013065):

Signed 8 bit:	100	Signed 32 bit:	1631205476	Hexadecimal:	64
Unsigned 8 bit:	100	Unsigned 32 bit:	1631205476	Octal:	144
Signed 16 bit:	14436	Signed 64 bit:	1631205476	Binary:	01100100
Unsigned 16 bit:	14436	Unsigned 64 bit:	1631205476	Stream Length:	8 - +
Float 32 bit:	2.146974e+20	Float 64 bit:	4.719431e+257		

At the bottom, there are two checkboxes:  Show little endian decoding and  Show unsigned and float as hexadecimal. The offset is 0x0.