SHARING DATA IN MULTI-PROCESS APPLICATIONS

Professor Ken Birman
CS4414 Lecture 19

IDEA MAP FOR TODAY

Complex Systems often have many processes in them. They are not always running on just one computer.

Modern solutions of this kind often need to run on clusters of computers or in the cloud, and need sharing approaches that work whether processes are local (same machine) or remote.

Linux offers too many choices! They include pipes, mapped files (shared memory), DLLs.
Linux weakness: the "single machine" look and feel.

As a developer, you think of the cloud itself as a kind of distributed operating system kernel, offering tools that work from "anywhere".

LARGE, COMPLEX SYSTEMS

Large systems often involve multiple processes that need to share data for various reasons.

Components may be in different languages: Java, Python, C++, O'CaML, etc…

Big applications are also broken into pieces for software engineering reasons, for example if different teams collaborate

MODERN SYSTEMS DISTINGUISH TWO CASES

Many modern systems use "standard libraries" to interface to storage systems, or for other system services.

You think of the program as an independent agent, but it uses the same library as other programs in the application.

Here, the focus is on how to build libraries that many languages can access. C++ is a popular choice.

LOCAL OPTIONS

These assume that the two (or more) programs or modules live on the same machine.

They might be coded in different languages, which also can mean that data could be represented in memory in different ways (especially for complicated objects or structures – but even an integer might have different representations!)

SINGLE ADDRESS SPACE, TWO (OR MORE) LANGUAGES

Issue: They may not use the same data representations!

EXAMPLE 1: JAVA NATIVE INTERFACE

The Java Native Interface (JNI) allows Java applications to talk to libraries in languages like C or C++.

In effect, you build a Java "wrapper" for each library method.

JNI will load the C++ DLL at runtime and verify that it has the methods you expected to find.

JNI DATA TYPE CONVERSIONS

JNI has special accessor methods to access data in C++, and then the wrapper can create Java objects that match.

For some basic data types, like int or float, no conversion is needed. For complex ones, where conversion does occur, the cost is similar to the cost of copying.

JNI is generally viewed as a high-performance option

EXAMPLE 2: FORTRAN TO C++

Fortran is a very old language, and the early versions made memory structs visible and very easy to access.

This is still true of modern Fortran: the language has evolved enormously, but it remains easy to talk to "native" data types.

So Fortran to C++ is particularly effective.

EXAMPLE 3: PYTHON TO C++ (TRICKY)

There are many Python implementations.

The most widely popular ones are coded in C and can easily interface to C++. There are also versions coded in Java, etc.

But because Python is an interpreter, Python applications can’t just call into C++ without a form of runtime reflection.

HOW PYTHON FINESSES THIS

Python is often used control computations in "external" systems.

For example, we could write Python code to tell a C++ library to load a tensor, multiply it by some matrix, invert the result, then compute the eigenvalues of the inverted matrix…

The data could live entirely in C++, and never actually be moved into the Python address space at all! Or it could even live in a GPU

PYTHON INTEGERS

One example of why it isn't so trivial to just share data is that Python has its own way of representing strings and even integers

A Python integer will use native representations and arithmetic if the integer is small. But Python automatically switches to a larger number of bits as needed and even to a Bignum version.

So… if Python wants to send an integer to C++, we run into the risk that a C++ integer just can't hold the value!

SOLUTION? USE “BINDINGS”

Boost.Python leverages this basic mechanism to let you call Python from C++ or C++ from Python.

1) You need to create a plain C (not C++) “interface” layer. These methods can only take native data types + pointers.

2) Compile it and create a DLL. In Python, load this DLL, then import the interface methods.

4) Now you can call those plain C methods, if you follow certain (well-documented) rules (like: no huge integers!). To call an object instance method, you pass a pointer to the object and then the arguments, as if “this” was a hidden extra argument.

EXAMPLE 4: MICROSOFT DOTNET CLR

Microsoft has many supported languages, including C++ on Ubuntu (just install WSL2 on your laptop)

But C# (a variant of Java) is probably the most popular.

It turns out that ALL of them can talk to C++ via something called the dotnet common language runtime (dotnet CLR).

ISSUE IS SIMILAR TO PYTHON, JAVA

As with those languages, you do need to decide if the memory for objects will be hosted in dotnet or hosted in C++

For objects hosted in dotnet there are methods you call to prevent garbage collection or compaction while your C++ is active. For objects hosted in C++, the dotnet languages can use “unsafe” memory pointers to access them

SHARING WITH DIFFERENT PROCESSES

Issue: They have different address spaces!

SHARING BETWEEN DIFFERENT PROCESSES

Large multi-component systems that explicitly share objects from process to process need tools to help them do this.

Unlike language-to-language, the processes won't be linked together into a single address space. The options depend on where the processes doing this sharing are running.

IF PROCESSES ARE ON A SINGLE (NUMA) MACHINE, WE HAVE A FEW "OLD" SHARING OPTIONS:

1. Single address space, threads share memory directly.
2. Linux pipes. Assumes a "one-way" structure.
3. Shared files. Some programs could write data into files; others could later read those files.
4. Mapped files. Same idea, but now the readers can instantly see the data written by the (single) writer. Also useful as a way to skip past the POSIX API, which requires copying (from the disk to the kernel, then from the kernel into the user's buffer).
5. Shared segment. Same as mapped files, but without a file!

DIMENSIONS TO CONSIDER

Performance, simplicity, security. Some methods have very different characteristics than others.

Ease of later porting the application to a different platform . Some modern systems are built as a collection of processes on one machine, but over time migrate to a cluster of computers.

Standardization. Whatever we pick, it should be widely used.

LET'S LOOK AT SOME EXAMPLES

The C++ command runs a series of sub-programs:

1. The "C preprocessor", to deal with #define, #if, #include
2. The template analysis and expansion stage
3. The compiler, which has a parsing stage, a compilation stage, and an optimization stage.
4. The assembler
5. The linker

… they share data by creating files, which the next stage can read

WHY DOES C++ USE FILE SHARING?

C++ was created as a multi-process solution for a single computer. In the old days we didn't have an mmap system call.

Also, since one process writes a file, and the next one reads it sequentially and "soon", after which it gets deleted, Linux is smart enough to keep the whole file in cache and might never even put it on disk.

There are many such examples on Linux. Most, like C++, have a controlling process that launches subprocesses, and most share files from stage to stage.

MMAP/SHMEN OPTION

We learned about mmap when we first saw the POSIX file system API. At one time people felt that mmap could become the basis for shared objects in Linux.

Linux allocates a segment of memory for the mapped file or shared memory segment. The system call returns a pointer to its base address.

Idea: create a shared memory segment, then allocate objects in it.

A SHARED SEGMENT IS JUST LIKE ANY OTHER SEGMENT IN YOUR ADDRESS SPACE

Only permits a single writer. (But you can always have two shared segments, one for each direction.) **Different processes may see a shared segment at different base addresses.**

If the data being shared is some form of raw information, like pixels in a video display, or numbers in a matrix, it works well.

Mmap starts with a file and changes are saved to disk. Shmem is just a pure shared memory segment.

MAPPED FILES ARE WIDELY USED FOR FINANCIAL TRADING SYSTEMS

Many Wall Street trading firms have real-time ticker feeds of prices for the stocks and bonds and derivatives they trade.

Often this is managed via a daemon that writes into a shared file. The file holds the history of prices.

By mapping the head of the file, processes can track updates. A library accesses the actual data and handles memory fencing.

LINUX ITSELF USES MAPPED FILES

The DLL concept ("linking") is based on a mapped file.

In that case the benefits are these:

- The file actually contains executable instructions. These must be in memory for the CPU to decode and execute.
- But the DLL can be shared between multiple applications, saving memory and improving L3 caching performance.

MAPPING OFTEN ISN'T THE WHOLE STORY

When we share data between processes we might

- Only share data a process has permission to see
- Leverage read-only access to avoid synchronization. For interactive cases, **p** and **q** could share two segments: 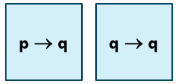

- Force Linux to put the segments at the same address to enable them to hold pointer-based data structures, like trees
- Use **remap**, signal handling to implement dynamic access

SHARING WITH PROCESSES ON DIFFERENT MACHINES

Issue: Now we need to also deal with the network

NETWORKED SETTINGS REQUIRE DIFFERENT APPROACHES

When we run in a networked environment, we need tools that will work seamlessly even if the processes are on different machines.

Mapped files or segments are single-machine solutions. Mmap can be made to work over a network, but performance is disappointing and this option is not common.

CLOUD COMPUTING

In other courses, you'll use modern cloud computing systems.

Those are like a large multicomputer kernel, with services that programs can use *no matter which machine they run on*.

Cloud computing has begun to reshape the ways we develop complex programs even on a single Linux machine.

DIFFERENT MACHINES + INTERNET

1. We will learn about TCP soon… like a pipe, but between machines. This extends the pipe option to the cloud case!
2. We could use a technique called “remote procedure call” where one process can invoke a method in a remote on. We will learn about this soon, too.
3. We could pretend that everything is a web service, and use the same tools that web browsers are built from.

AMAZON.COM

Prior to 2005, Amazon web pages were created by a single server per page. But these servers were just not fast enough.

Famous study: 100ms delay reduces profits by nearly 10%

Today, a request is handled by a “first tier” server supported by a collection of services (as many as 100 per page)

AMAZON INVENTED CLOUD COMPUTING! THEN GOOGLE AND MICROSOFT TOOK IT MUCH FURTHER.

The Amazon services are used by browsers from all over the world: a networked model.

And Amazon's explicit goal was to leverage warehouses full of computers (modern "cloud computing" data centers).

… So Amazon is a great example of a solution that needs to use networking techniques.

INSIDE THE CLOUD?

Users of cloud computing platforms like Amazon's AWS, Microsoft's Azure, or Google Cloud don't need to see the internals.

They see a file system that is available everywhere, as well as other kernel services that look the same from every machine.

The individual machine runs Linux, yet these services make it very easy to spread one application over multiple machines!

FLASK AND WEB SERVICES

Early in the Amazon era, the company needed a lot of machines for peak load bursts but wanted to rent them out.

So they offered a way to

- Copy your web site to AWS, renting their machines for this
- Doing a kind of remote request to a cloud-hosted service through your company web site.

REST RPC

Any program can call any web site that exports a REST API.

- Example: Your company web site could have an AWS service with a method to "Upload the current AQI data"
- Then it could sell AQI sensors that link via the Internet and upload air quality data periodically
- Like calling **upload(sensor-id, PM2.5, O3, ….)** but the request turns into a message. The actual upcall occurs on a cloud server
- The most popular package for using REST is called FLASK. It supports Javascript, Java and C++

PURPLE AIR WEBSITE FOR ITHACA…

Ken’s house is on this!

And in fact there is a global version of this too, called acqicn.org

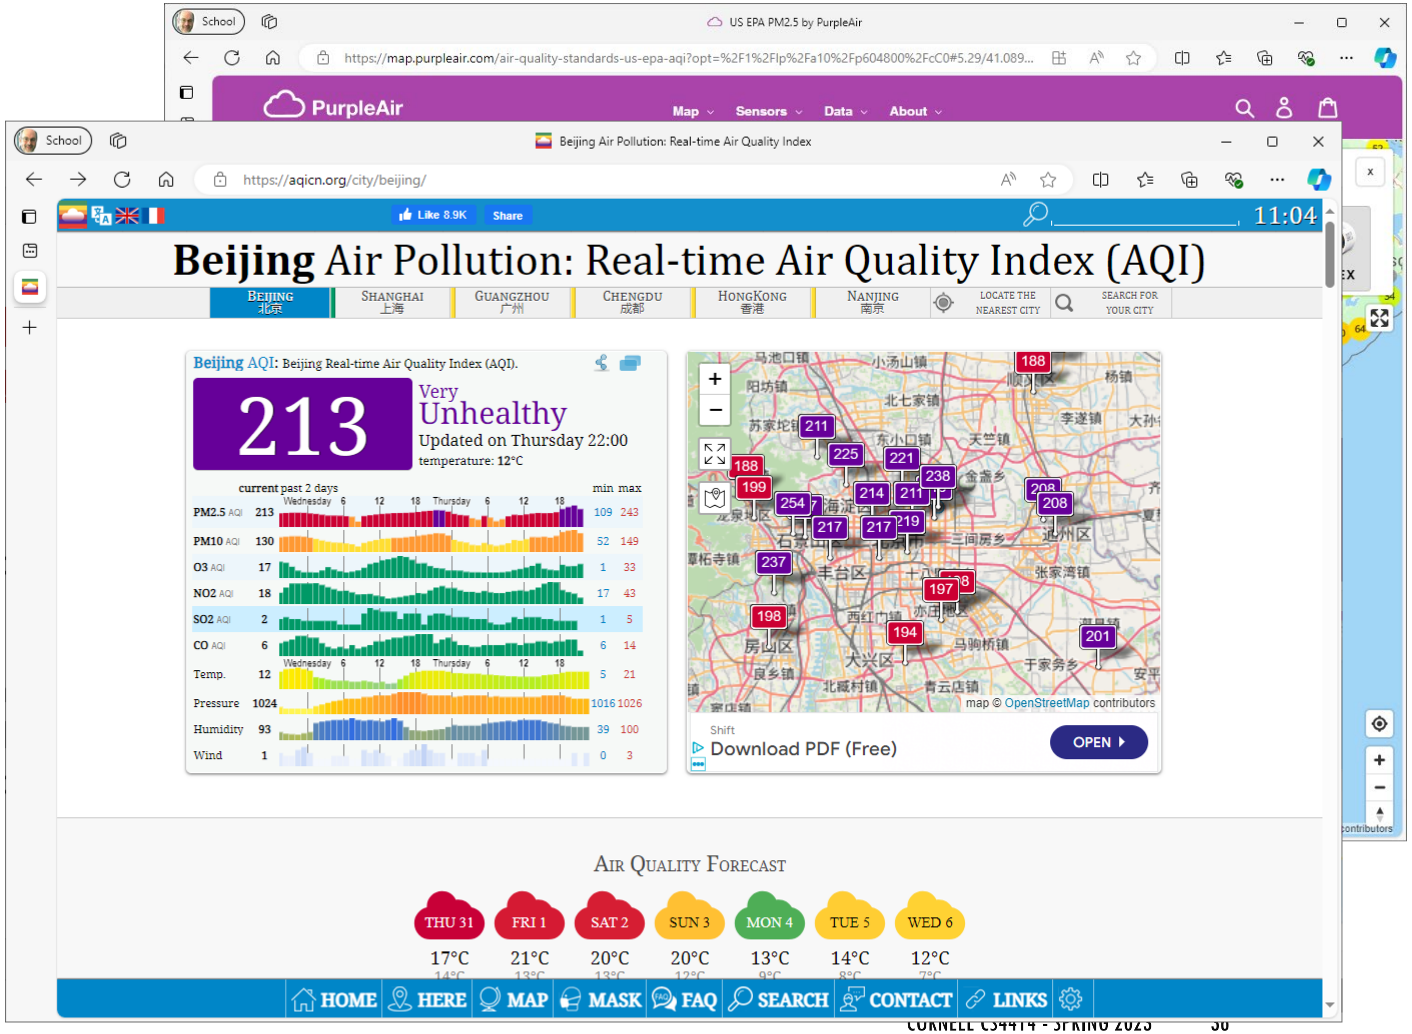

WHAT ABOUT CORBA?

This is an older standard but very widely used

REST encodes RPC requests as web pages and uses HTTP or HTTPS to connect via the same protocol used by browsers

CORBA doesn't use this extra HTTP/HTTPS encoding

AIR TRAFFIC CONTROL

Ken worked on the French ATC solution. It uses CORBA

This system has been continuously active since 1996. It runs on a private cloud, but uses cloud-computing ideas.

ATC systems have many modules that cooperate. The “flight plan” is the most important form of shared information.

AIR TRAFFIC CONTROL SYSTEM

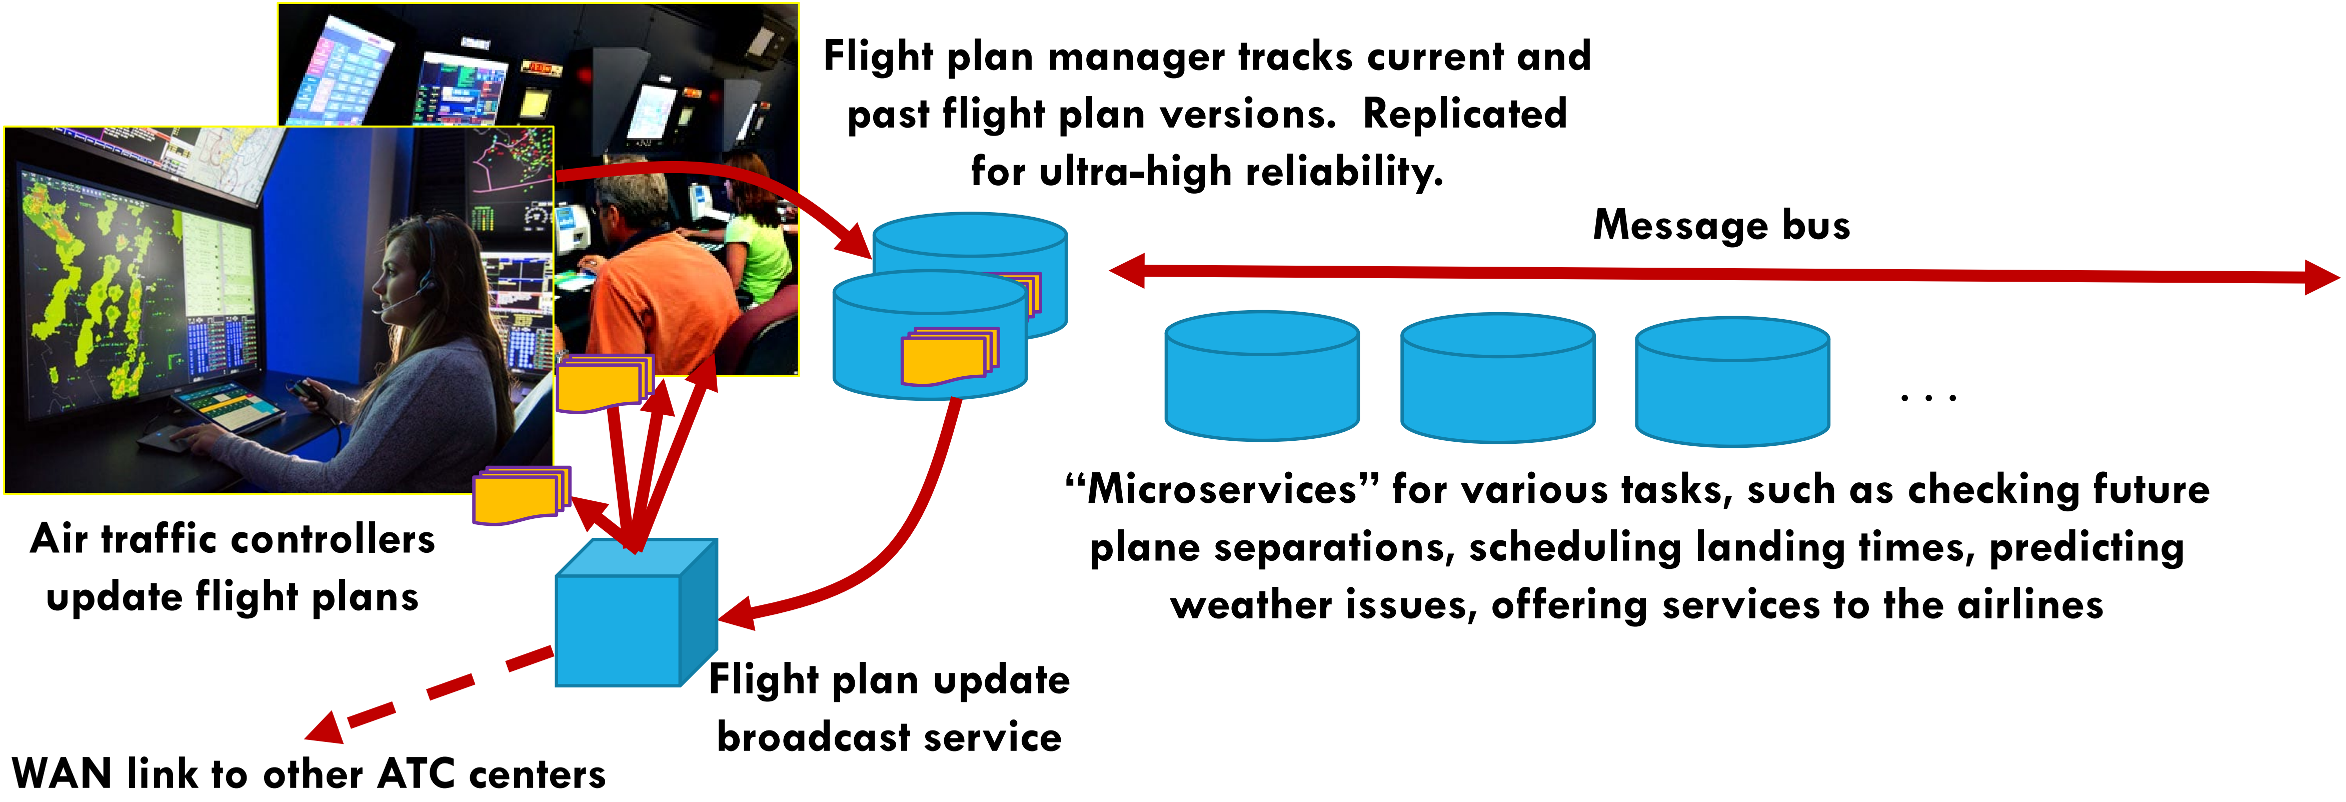

SOFTWARE ENGINEERING AT LARGE SCALE

Big modern applications are created by software teams

They define modular components, which could co-exist in one address space or might be implemented by distinct programs

There is a science of *software engineering* that focuses on best ways of collaborating on big tasks of this kind.

SOFTWARE ENGINEERING AT LARGE SCALE

Each team needs a way to work independently and concurrently.

The teams agree on specifications for each component, then build, debug and unit test their component solutions.

We often pre-agree on some of the unit tests: "release validation" tests and "acceptance" tests. Integration occurs later when all the elements seem to be working.

MULTI-LINGUAL ISSUE

Modularity permits us to use different languages for different tasks. For example, a great deal of existing ATC code is in Fortran 77.

Byte arrays (or text files, character strings) are a least common denominator. Every language has a way to easily access them.

Modern systems have converged around the idea that this matches best with some form of "message passing".

WHERE IS THIS SEEN IN THE ATC SYSTEM?

The flight plan record is a good example.

We want to save and load it from a shared, fault-tolerant database service. Many programs might access it.

… and they might not all be using the same byte order for basic data types!

BYTE ORDER (ENDIAN) ISSUE

Intel and AMD favor Little Endian. Suppose we store 0x12345678 in a 32-bit integer at address 1000:

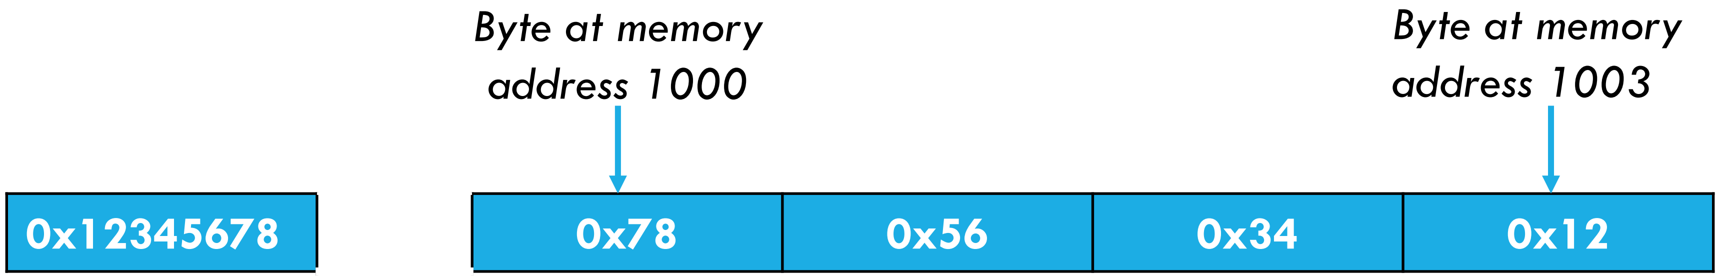

But ARM (used on Raspberry pi) offers *both* options. This example is Big Endian:

ISSUES CAUSED BY PADDING RULES

When we create an object with multiple data fields in it, the C++ compiler has to define a corresponding struct that will hold the data.

The "alignment requirements" for different fields can vary for different hardware. So though Intel and AMD use Little Endian, the same object might have a different layout in memory!

WHAT TO DO?

No need to do anything if we always work in just one hardware architecture, where the same layout and endian policy applies.

But if a system will combine elements from different sources that use different hardware, a second option is to use a “serialization and deserialization” library. These “interoperate”

SERIALIZATION/DESERIALIZATION

Converting an object to a byte array serializes the object and any associated annotations. Deserializing recreates the object.

A serialized object can be stored in a file, or we can use a "message passing" technology to send them from process to process over a network.

ADVANTAGE TO USING A SERIALIZATION PACKAGE: IT CAN DO MORE!

Of course, there are limitations. For example, strings could be plain ascii, but could also encode languages like Chinese, Russian, Japanese… It is not always meaningful to interoperate

Moreover, even if everything else is the same, the OS itself decides if printed output lines should end with \n or \r\n

FULLY ANNOTATED OBJECTS?

Many systems also wish to include information to document application version numbers at patch level, data types in use, sizes of arrays, requirements or assumptions that methods are making, limits on sizes of things, permissions required, etc.

There is very little agreement on how these annotations should look, or what to do if sender and receiver do not match.

APPLICATIONS CODED IN DIFFERENT PROGRAMMING LANGUAGES/SYSTEMS

A system of this kind often has components in C++, components in Java or Python, components in Fortran. One might use Tensor Flow, another could be using PyTorch

(And for historical accuracy: There were periods when many government or military systems favored Cobol, or Ada!)

How are objects like "flight plans" shared between components written in different languages?

NETWORKING STANDARDS AND FLEXIBILITY

If we think about Linux pipes, they are extremely simple and flexible. The main cost is simply that the data itself is a byte stream.

Developers began to question all of these shared memory ideas and complexities. **Are they worth all the trouble?**

LIKE POSIX, A STANDARD SPECIFIES RULES THAT VENDORS AGREE TO RESPECT, TO FACILITATE INTEROPERABILITY.

CORBA: A standard architecture for sharing objects between programs or components from many languages or developers.

Google RPC (GRPC): A faster way for a client program to invoke a method in a server, perhaps over the Internet. Optimized for systems where "sender" and "receiver" are on the same hardware and using the same OS.

Web services: An approach in which web pages in HTML are used to share information between programs. Widely available but slow.

COST ANALYSIS EXAMPLE: AIR TRAFFIC FLIGHT PLAN IN THE ATC SYSTEM WE SAW

In memory, a flight plan is generally no more than 125k bytes.

With CORBA encoding, this grows to between 1MB and 10MB

- All numbers are "printed out", usually in base 10
- CORBA includes details on the way the data types were declared, version information, etc.

Effect? In some ATC settings, the system spends more time encoding and decoding flight plans than controlling aircraft!

WHERE ARE OBJECTS MOVED OR SHARED?

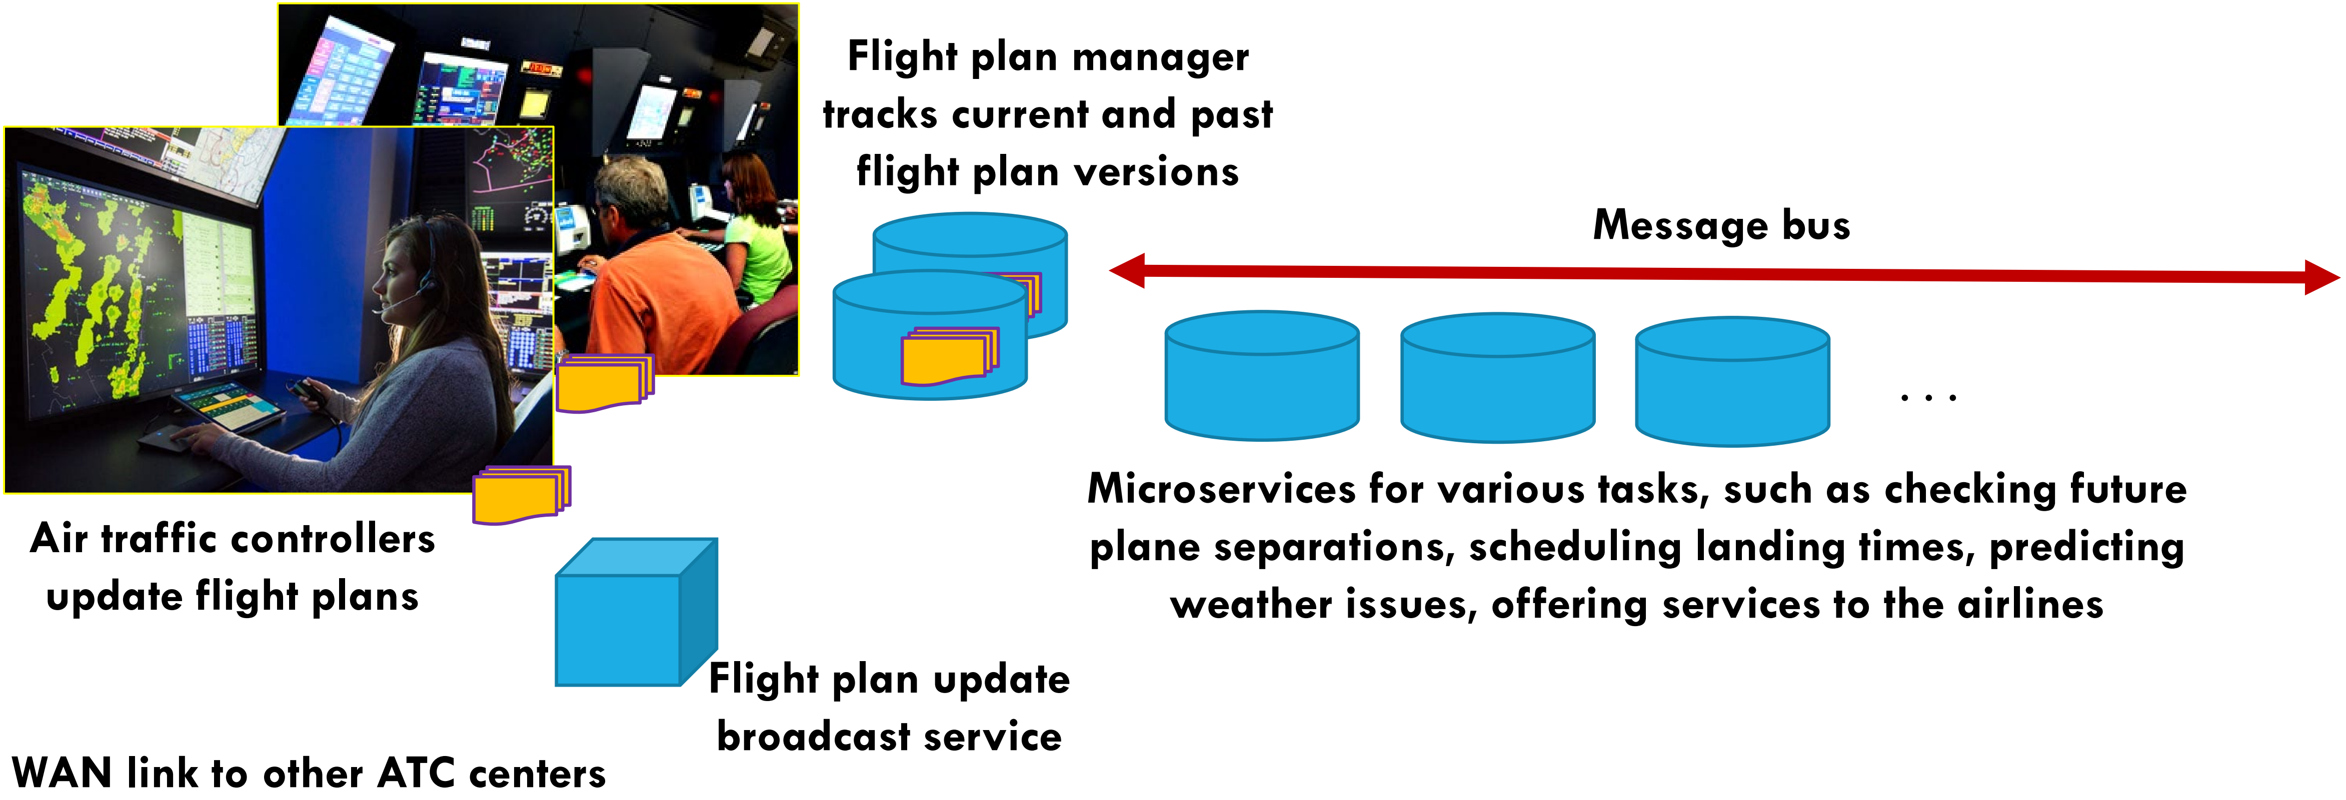

WHERE ARE OBJECTS MOVED OR SHARED?

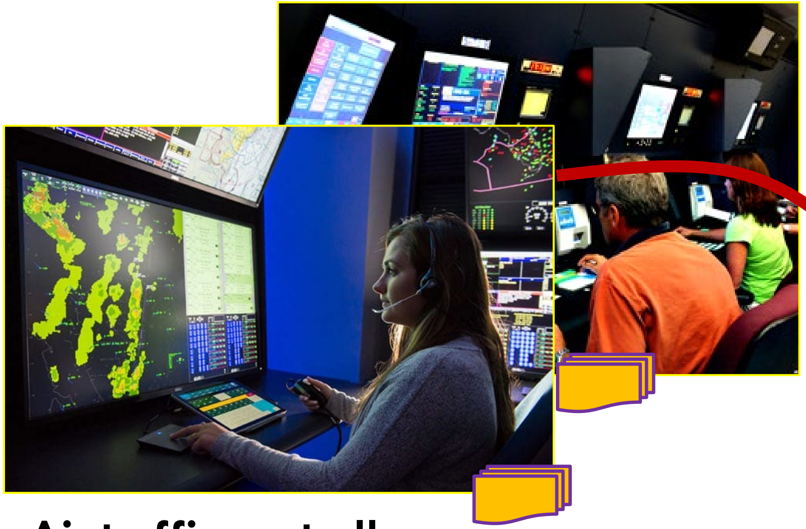

Flight plan manager tracks current and past flight plan versions

Message bus

Microservices for various tasks, such as checking future plane separations, scheduling landing times, predicting weather issues, offering services to the airlines

Air traffic controllers update flight plans

Flight plan update broadcast service

WAN link to other ATC centers

WHERE ARE OBJECTS MOVED OR SHARED?

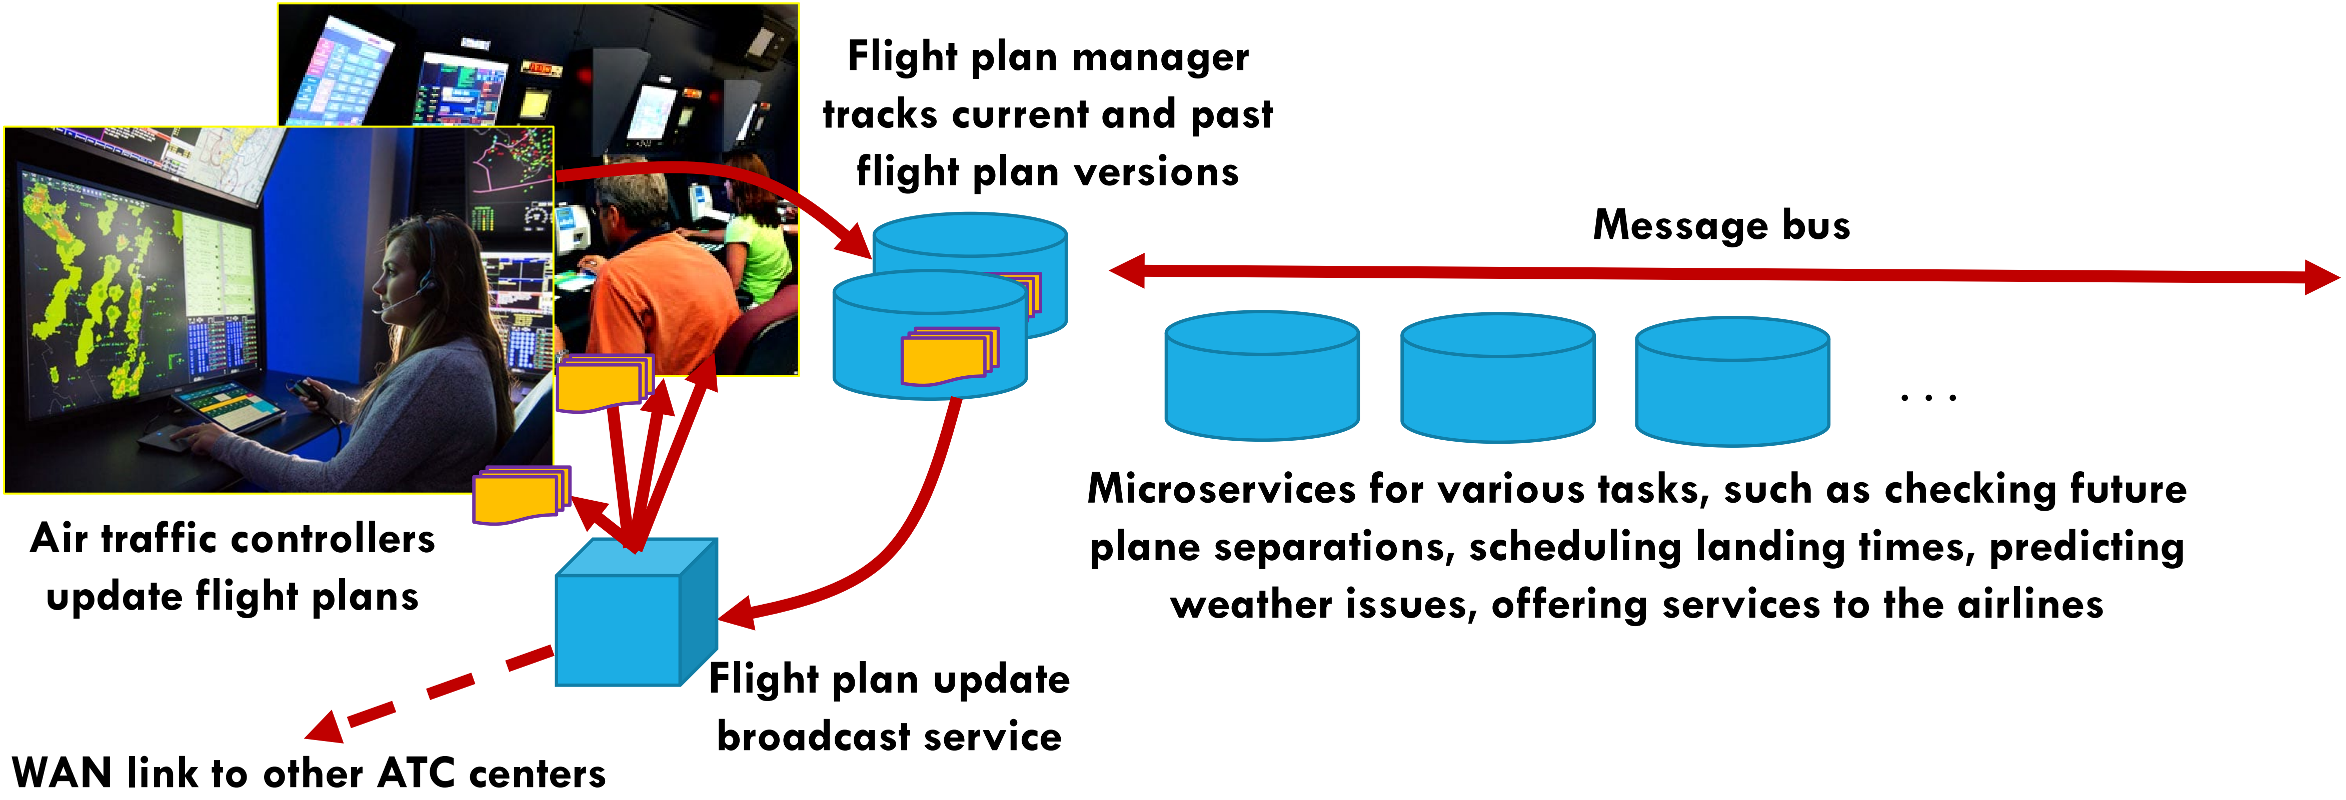

WHEN DO WE SERIALIZE/DESERIALIZE?

Each time an object is read or written (from disk or network)

Each time an object is passed from one module to another

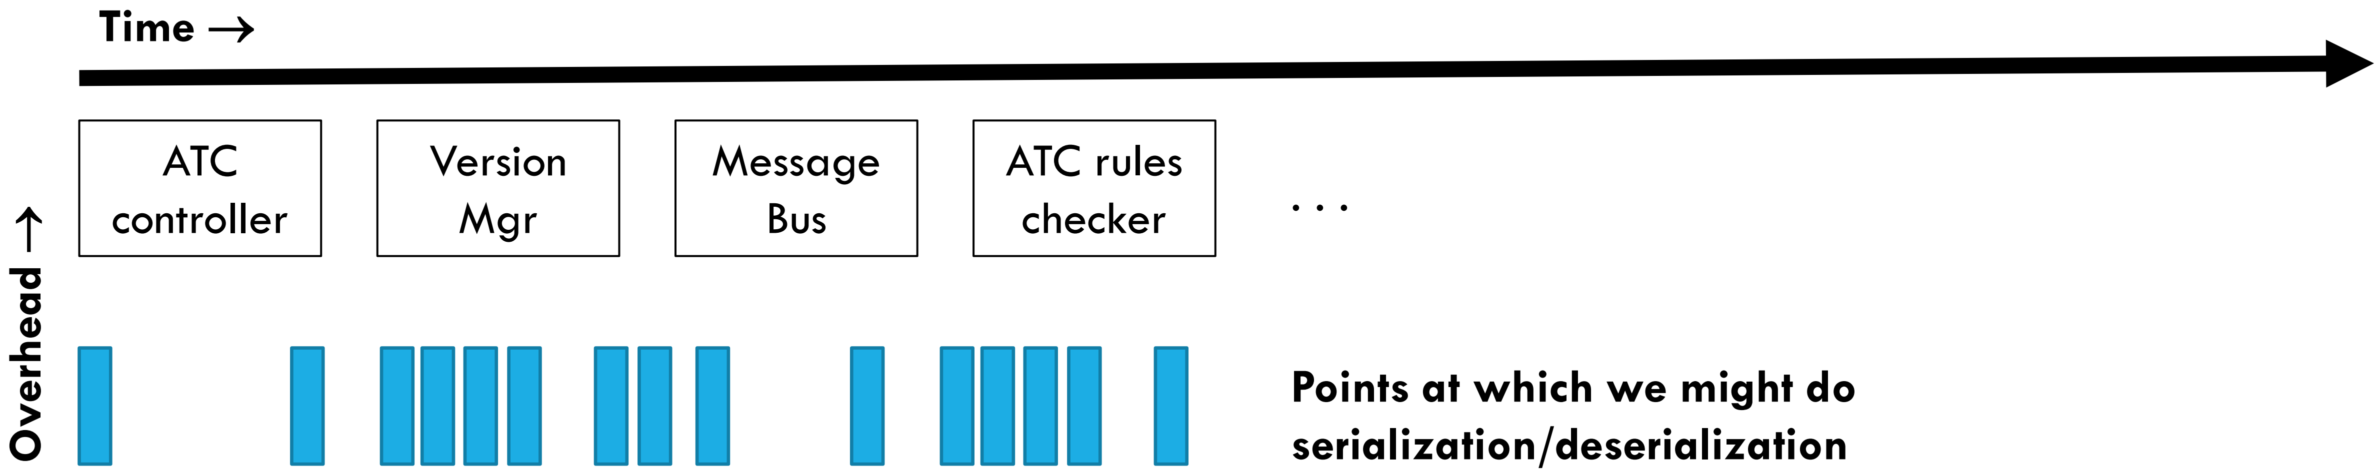

COST IMPLICATIONS

Potentially, a major source of overhead!

Often, it is best to store a complex serialized object in a file, and then just pass the file *name* from place to place. Then the CORBA object just has a few bytes in it (very cheap).

In a complex application where the actual fields in the object aren't needed by many modules, this reduces costs dramatically!

WHY WOULD A MODULE NOT LOOK AT THE DATA?

In the air traffic example, some modules just look at a few fields.

The WAN module is responsible for sharing updates with other air traffic control centers. It doesn't need to actually see the details.

… in fact, several modules simply move objects from process to process.

… all of these would be happy with just sharing the object *name*.

OLD APPROACH

Each time an object is read or written (from disk or network)

Each time an object is passed from one module to another

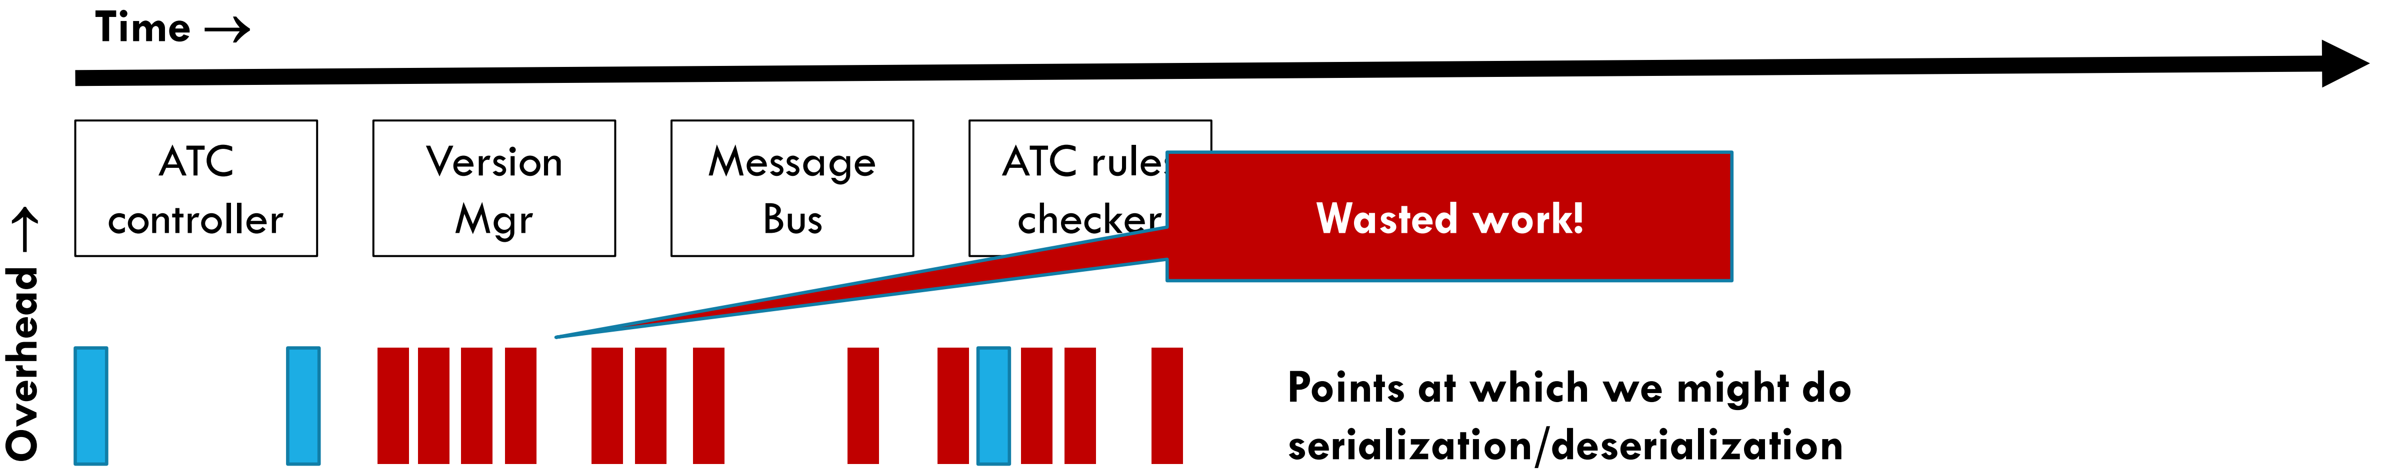

SHARING OBJECT NAMES, ONLY FETCH THE DATA IF THE MODULE REALLY REQUIRES IT

We only do a costly action when the module will actually touch the inner data fields!

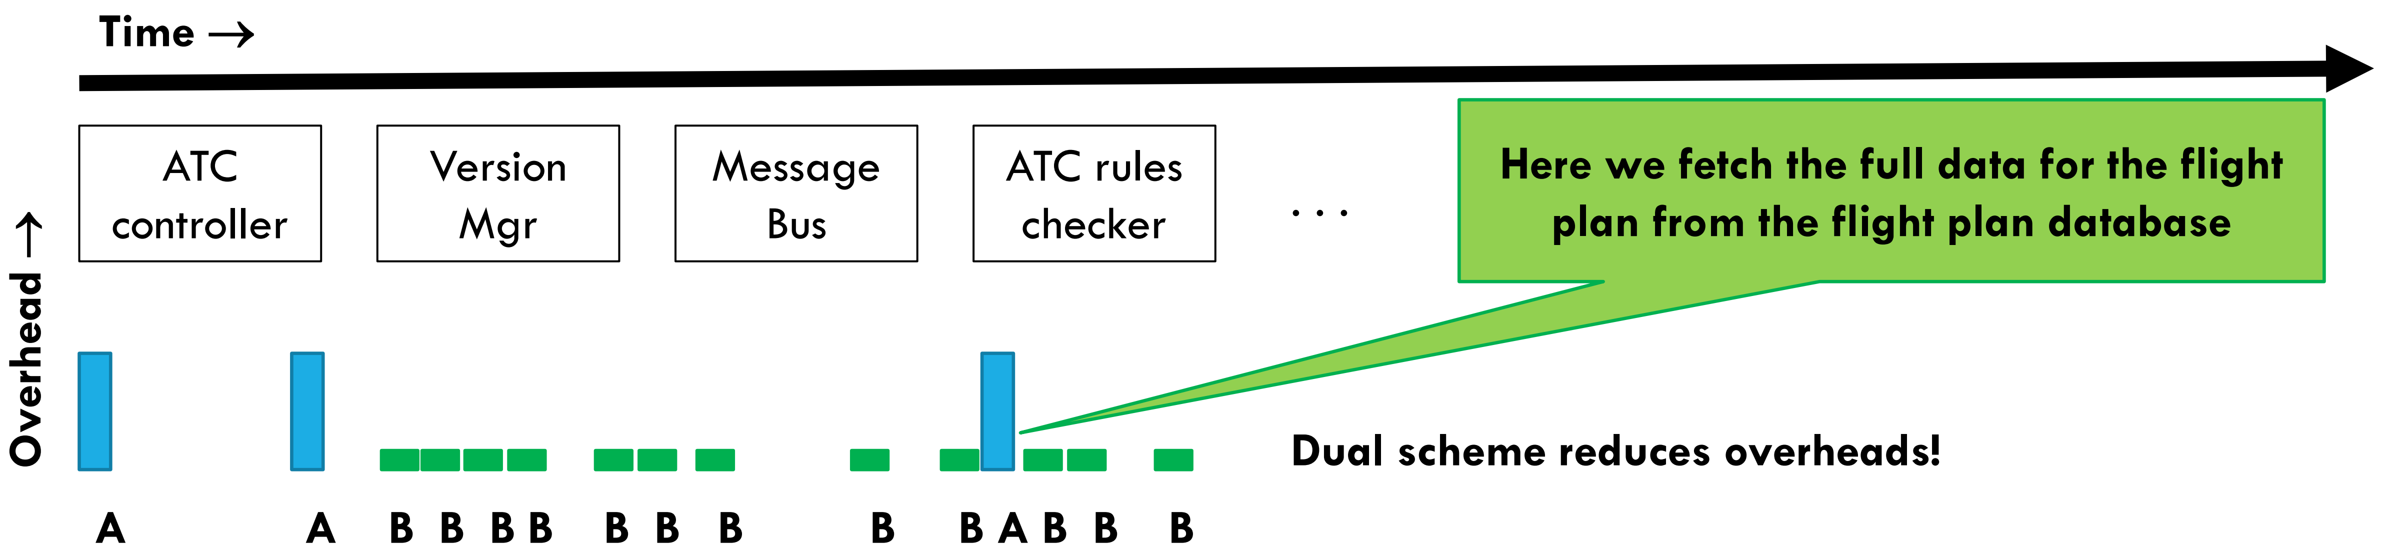

SUMMARY

Modular design creates a need for processes to share data.

In a single Linux system, pipes and file sharing are by far the most common models. But there are some important uses of shared memory.

The options are easy to use, but we need to be very aware of overheads and costs!