

# 1 Pattern: Data flow architecture

## Context

Sometimes the solutions for different problems share several identical steps, yet each problem requires its own solution.

Consider for example multimedia applications that process JPEG still images and MPEG videos. A JPEG viewer applies five transformations and each of them is performed in a separate step: JPEG syntax decoder, dequantizer, inverse DCT transformer, frame decoder and color conversion. Likewise, an MPEG player applies four transformations: MPEG syntax decoder (for simplicity, this also deals with temporal and spatial prediction), inverse DCT transformer, frame decoder and color conversion. For both applications, the transformations are applied sequentially, in the same order.

The inverse DCT transformer and color conversion are identical and operate on the same data types—images represented with one luminance component ( $Y$ ) and two chrominance components ( $C_b$  and  $C_r$ ). The interaction with the other transformations is limited. Control flow is static and the succession of steps rarely changes—usually, only for exceptional circumstances. Each time the inverse DCT transformer is presented a set of coefficients, it computes the corresponding spatial representation. Next this is passed to the frame decoder.

## Problem

Processing within your domain is performed by applying a sequence of operations on similar data elements, in the same order. You want a reusable toolkit that can be used to build software solutions for a wide range of applications within the domain. Your applications are also required to dynamically change their functionality and/or adapt to changing environments, without compromising performance. What architecture can accommodate these requirements?

## Forces

- A high-performance toolkit that is applicable for a wide range of problems from a specific domain is required;
- The application's behavior needs to be modified dynamically or adapt to changing environments at runtime;
- The loose coupling associated with the black-box paradigm has performance penalties;

- Scalability in several dimensions is required.

## Solution

A data flow architecture organizes applications as a network of processing modules<sup>1</sup> that apply a series of transformations to one or several data streams. Each module takes its input from some upstream modules, performs a simple, generic transformation and passes the results to some other downstream modules. Enforcing strict, simple inter-module interfaces yields a large number of possible combinations that provide solutions to many problems within a particular domain. Therefore, in data flow architectures, processing modules are the *computational units*, while the network represents the *operational unit*.

Consequently, the functionality (transfer function) is determined by the:

- types of processing modules within the network; and
- interconnections between them.

Figure 1 shows a data flow JPEG decoder. The decoder reads a JPEG file and shows its contents on a display. Decoders for other formats are implemented by connecting different modules.




Figure 1: A data flow JPEG decoder.

This architecture is applicable only for data-driven applications, where the output is obtained by performing various sequential transformations to the input. It allows applications to scale with the number of data streams. The transfer function is determined by a combination of the individual transfer functions of each processing module. Control flow is distributed across the participating entities and is not explicitly represented at the architectural level.

Modules play a key role in data flow architectures. Applications that follow this pattern manifest an increased degree of modularity. This makes it easy to distribute the development effort among different groups. For example, visualization

<sup>1</sup>In this context, “module” is any processing unit within the application domain.

modules are developed by the visualization group, while MPEG modules are built by the multimedia group, etc. Therefore, the knowledge of domain experts is encapsulated within processing modules that are reusable in many different contexts.

The following design guidelines are essential for generic and reusable processing modules:

- The transfer function should not model the domain, depend on a particular solution or have side effects.
- Each module should be developed independently of its use.
- Processing modules should cooperate only by using the output of one (or several) as the input to another.

Decoupling the algorithmic properties of processing modules from a particular problem allows applications to use them as black-boxes. However, the black-box approach does not take into account context-specific factors and therefore has performance penalties. Good filters balance these two conflicting forces.

The fine granularity decomposition into simple transformations is an incarnation of the “divide and conquer” principle: instead of attempting to solve a complex problem all at once, split it into sub-parts and deal with each separately. This allows applications to scale with the number of transformations and processors—each processing module can run on a separate processor. However, because of the reduced communication overhead, specialized processing modules are more efficient than implementing their functionality with several generic filters. But specialization also limits the number of possible configurations that a filter can be part of. For example, a single specialized JPEG decoder module offers better performance than the decoder from Figure 1. However, in the latter case filters are reusable and can be employed for other applications (e.g., an MPEG decoder). Filter designers have to determine the right balance between granularity and performance.

The fine granularity decomposition has some other advantages. Consider how electrical engineers analyze or test a circuit. They use an oscilloscope to check the signals at different key points. This requires *access to each component*, and the first step is to remove the covers that enclose the circuit. Because of its modular structure, the data flow architecture provides easy access to its computational units, the processing modules. Like for the electrical circuit, this enables software engineers to analyze and test the application. For example, the decoder from Figure 1 provides direct access to the IDCT coefficients, should they be required for analysis or testing. Analyzing a functionally equivalent decoder which does not follow this pattern requires more effort to get to these coefficients. Therefore, from an

electrical engineering perspective, data flow applications do not have “covers” that prevent access to their components.

Easy access to data at different processing stages also enables developers to use signal processing techniques (translation, interpolation, etc.). For fields like multimedia or scientific visualization, these techniques have been shown to improve performance [PLV97].

Inter-module communication is done by passing messages or tokens through unidirectional input and output ports—shown as blue (input) and red (output) plugs in Figure 1. Therefore, data flow systems replace inter-module direct calls with a message passing mechanism. The overhead associated with this mechanism is the main liability of the pattern.

Depending on their ports, modules are classified as follows (a comprehensive classification according to various criteria is available in [Lea96]):

**Sources** Interface with an input device and have one or several output ports.

**Sinks** Interface with an output device and have one or several input ports.

**Filters** Have both input and output ports (not necessarily only one in each direction) and perform processing on the information fed into the input port. When processing completes, the filter writes the results to the output port.

In Figure 1, the file reader is a source, the display adapter is a sink and the other processing modules are filters.


Unidirectional input and output ports are not a limitation. Rather, they increase a component’s autonomy. Provided that there are no feed-back loops, processing is unaffected by the presence or absence of connections at the output port(s). Therefore, because any component depends only on the upstream modules, it is possible to change its output connections at runtime. This is called dynamic composition and determines some of the interesting properties of data flow architectures.

To allow interconnectivity between two modules, the output port of the upstream module and the input port of the downstream module have to be *plug-compatible*. (For simplicity, in Figure 1 all modules have the same type of port.) However, having many port (plug) types limits interconnectivity and requires adapters to connect incompatible ports.


Two priority levels are available for processing at the filter level. High priority processing can be done when the filter receives a message from the upstream module. Once acquired, the data can be processed at a later time, at lower priority. For instances where processing is time-consuming—e.g., scientific visualization applications—the two priority levels are important. Otherwise, if the output

values can be computed relatively fast, all processing is performed at high priority. However, scheduling (i.e., “when” the actual processing takes place) is non-deterministic in both cases.

Applications following the data flow architecture can adapt to heterogeneous and dynamically changing environments at runtime. Different implementations for stateless filters are exchanged to adjust the resource cost and quality characteristics within some user-defined limits. In Figure 1, two different IDCT modules are available and depending on their characteristics, the decoder uses one or the other. Therefore, *dynamic composition* allows an application to (1) adjust the resource consumption of the system at runtime, and (2) adapt to different computing and communications environments, as well as changes in resource availability.



Sometimes the overhead of inter-module communication has unacceptable performance penalties. An effective solution for this problem is to trade flexibility for performance [PLV96]. This is accomplished by replacing the adjacent performance-critical modules with a single optimized module. *Static composition* provides the underlying application (e.g., compiler) with enough information to collapse the sequence of filters into a functionally equivalent primitive filter, reducing the overhead of inter-module communication. However, this cannot be modified at runtime anylonger.



Static binding should be used to improve performance for filters that have low reconfigurability demands and are called with high frequency. Dynamic binding is best suited for filters that have high reconfigurability demands and low invocation frequency. A good balance between the use of static and dynamic composition ensures efficient implementations while maintaining a flexible, configurable architecture. Therefore, data flow applications are scalable with different communication and processing requirements. For the JPEG decoder from Figure 1, the syntax decoder is called many times for each image and therefore is statically bound to the dequantizer [PLV97]. The other components are called with lower frequencies and the message passing overhead is tolerable.

Users familiar only with the application domain (e.g., scientific visualization) can create new applications by simply connecting existing modules, without performing any other programming. This is accomplished with scripting languages [Lin94, JBR96] or visual programming tools [AEW96] that assist the creation of module networks [Foo88].

The network normally triggers recomputations whenever a filter’s output value changes. While manipulated interactively, it should be possible to temporarily disable the network such that no computation triggers until the network is in a valid, stable state. Disabling computations is useful if several parameters need to

be modified without having the filters recompute every time a change is made, or if processing takes a long time.

To summarize, the data flow architecture has the following **benefits** (✓) and **liabilities** (✗):

- ✓ It emphasizes reuse at the processing module level and facilitates the rise of end-user programming, automation and software components.
- ✓ Because the interaction mechanism between modules is simple, the architecture is suitable for domain-specific frameworks and often is the subject of visual programming tools.
- ✓ Data flow applications are scalable in several dimensions—number of transformations, processors, data streams; and communication and processing requirements.
- ✓ The fine granularity decomposition into simple filters provides access to the data stream at various processing stages.
- ✗ A data flow architecture is not a good choice for applications with dynamic control flow or feedback loops.
- ✗ For instances where the overhead of enforcing the inter-module interface is too high, a different architectural choice might be a better solution.
- ✗ It does not deal well with unanticipated conditions. Signaling errors that occur inside filter modules is cumbersome and difficult [BMR<sup>+</sup>96]. Modules do not make any assumptions about their context and communicate only with other data-processing modules. Only when combined with additional knowledge about the network topology and the application's domain could an error message generated by a module be meaningful. The Processing and control partitions pattern provides one possible solution.
- ✗ The architecture works best for applications where modules exchange simple data and do not share state. It is not suitable when the exchanged data is complex or for database applications.

### **Resulting context**

Applications following the data flow architecture can run on heterogeneous distributed systems. Because control flow is not explicitly represented, this change is transparent for sources, filters and sinks.

```

template<class DataType> class PushInput {
public:
    virtual ~PushInput() { };
    virtual void Put(DataType &data) =0;
};

```

Figure 2: The PushInput class.

```

template<class DataType> class PushOutput {
public:
    virtual ~PushOutput() { };
    virtual void Attach(PushInput<DataType> *next)
    {
        _port=new PushPort<DataType>(next);
    };
    virtual void Detach()
    {
        delete _port;
    };
protected:
    inline void Output(DataType &data) { _port->Output(data); };
private:
    PushPort<DataType> *_port;
};

```

Figure 3: The PushOutput class.

### Implementation notes

This section follows the guidelines from [PLV96]. Processing modules inherit an input interface (port) from PushInput and an output interface from PushOutput. PushInput—Figure 2—is an abstract class that determines the input data type. Subclasses implement the Put() method to process the input data and pass it downstream. PushOutput—Figure 3—determines the output data type and defines methods for dynamic composition—Attach() and Detach(). Attachment between two modules is abstracted by the PushPort class—Figure 4. This facilitates establishing one-to-many connections, in which case PushOutput has to be modified accordingly.

C++ templates parameterize the classes presented in Figures 2–4 by the types

```

template<class DataType> class PushPort {
public:
    PushPort(PushInput<DataType> *module)
        : _module(module) { };
    ~PushPort() { };
    inline void Output(DataType &data) { _module->Put(data); };
protected:
    PushInput<DataType> *_module;
};

```

Figure 4: The PushPort class.

```

class Sqrt : public PushInput<int>, public PushOutput<double> {
public:
    Sqrt() { };
    ~Sqrt() { };
    void Put(int &i)
    {
        double d=sqrt(i);
        Output(d);
    };
};

```

Figure 5: An Sqrt processing module.

of data that enters and leaves each processing module. The compiler's type checking system prevents connecting two modules which are not plug-compatible.

Based on these classes, implementing processing modules is straightforward. Figure 5 shows the code corresponding to a module that takes one integer input and outputs its square root. However, this simple implementation has several limitations with respect to how processing is initiated and how data is passed between modules. The following sections cover these aspects and provide alternatives.

In this example, type parameterization—templates—and function inlining are used to implement static composition. Figure 6 shows the code corresponding to a filter that increments its integer input. Processing—incrementing—is performed by the `Transform()` inline method. The `FastInc` class can parameterize the filter from Figure 7. Because `FastFilter` invokes only the inlined `Transform()` method, `FastInc`'s input and output ports are bypassed and the implementation is efficient. Figure 8 shows how to instantiate the static composition of `FastInc`



```

class FastInc : public PushInput<int>, public PushOutput<int> {
public:
    FastInc() { };
    ~FastInc() { };
    virtual void Put(int &i)
        {
            int j=i;
            Transform(j);
            Output(j);
        };
    inline void Transform(int &i) { i++; };
};

```

Figure 6: A filter for static composition. Only the inlined method is used by the composite filter.

with `FastFilter`. The resulting filter first increments the input and then outputs its square root. The idea behind implementing static composition is bypassing the input and output ports to eliminate the communication overhead. This can be accomplished in any language and does not require templates and inline methods.

## Examples

1. **Scientific visualization** is one domain where the data flow architecture is used extensively. Systems like AVS [AVS93] or Iris Explorer [EXP93] allow users to create visualization applications and visualize various data types in many different ways. This is done by connecting modules into a network, with the help of a visual network editor. Because they provide a wide range of generic filters, these systems offer solutions to a large number of visualization problems. Figure 9 shows an example of an AVS network.

Whenever required, the existing set of filters can also be expanded. New filters are created with a module builder tool. This allows end-users to extend functionality with a minimal amount of programming.

2. Data flow architectures have also been investigated for **hardware** systems [Gur85, Pap91]. The execution model offers attractive properties for parallel processing—implicit synchronization of parallel activities and self schedulability. Unlike the von Neumann model which explicitly states the sequence of instructions, in the data flow model the execution of any instruction is driven by

```
template<class T>
class FastFilter : public PushInput<int>, public PushOutput<double> {
public:
    FastFilter(T &component)
        : _component(component) { };
    ~FastFilter() { };
    virtual void Put(int &i)
    {
        int j=i;
        _component.Transform(j);
        double d=sqrt(j);
        Output(d);
    };
private:
    T _component;
};
```

Figure 7: A composite filter.

```
FastInc fi;
FastFilter<FastInc> ff(fi);
```

Figure 8: Instantiating the composite filter.




Figure 9: AVS Network. Top side ports are inputs, while bottom side ports are outputs. The port types are color-coded and the network editor does not permit connecting incompatible ports.

the operand availability. This emphasizes a high degree of parallelism at the instruction level. The Monsoon Project [Pap91] developed by MIT and Motorola produced a data-flow multiprocessor targeted to large-scale scientific and symbolic computation. Its success motivated much of the work on similar projects [NPA92] and contributed to spread the interest in data flow and parallel programming.

3. **Avionics Control Systems** (ACS) also employ this pattern [Lea94]. Because ACS are complex systems, constructing a set of components that merge all possible combinations is not feasible. The data flow architecture provides the means for combining together different types of existing components (filters) to serve a particular purpose.
4. Another domain where the data flow architecture is employed by applications (ActiveMovie [AMS], VuSystem [Lin94]), toolkits (Berkeley Continuous Media Toolkit [JBR96]) and frameworks (MET++ [Ack94], Presentation Processing Engine [PLV97], Java Media Framework [SM97]) is **multimedia** [MN98].

VuSystem applications<sup>2</sup> illustrate very well the flexibility of this architecture.

<sup>2</sup>Some of the VuSystem applications are available on the world-wide web at

In the “room monitor,” the output of a surveillance camera is analyzed such that just the frames that contain motion are recorded. The “joke browser” extracts only selected parts (e.g., jokes) from the complete recording of a late night show. Although these two applications are very different (the former processes real-time video while the latter provides content-based access), both of them have been built with the same tool just by connecting existing modules.

ActiveMovie allows users to play digital movies and sound encoded in various formats. It consists of sources, filters and renderers connected in a filter graph. All graph components are implemented as Component Object Model (COM) objects. Filters have pins which are connected with streams, but other communication channels for specialized communication (e.g., error notifications) are available.

The presence of the data flow architecture within the forthcoming Java Media Framework demonstrates its validity and confirms it as a recurring solution that has passed the test of time.

### Variants<sup>3</sup>

1. Bill Walker’s Ph.D. thesis [Wal94] presents a framework that uses this pattern in the context of **computer-assisted music**. Unlike typical data flow architectures, this system has *feedback loops* and *shared state*. The processing modules are called Composers and Transformers. Global information independent of the components is encapsulated within a PolicyDictionary object.
2. In [Rit84], Dennis Ritchie describes a variant of this pattern and how it is implemented in the **UNIX stream system**. Data flow is *bidirectional* and filters have queues for each direction. The queues are also employed for flow control—Section ??.

### Related Patterns

- Messages encapsulate all types of information exchanged by collaborating modules. The Payloads pattern allows modules to distinguish between messages and the data within them.

---

<http://www.tns.lcs.mit.edu/vs/vusystem.html>

<sup>3</sup>The differences between the data flow architecture and the variants discussed in this section are shown in italics.

- Data flow architectures replace inter-module direct calls with a message passing mechanism. The Payload passing protocol patterns provide several alternatives for this mechanism.
- A distinct part of a data flow application is in charge with dynamic composition and handles error messages. Processing and control partitions describes how to organize such applications.
- Systems following the Layers pattern [BMR<sup>+</sup>96] are organized as layers that exchange data between them. However, in this case each layer corresponds to a different abstraction level where data has different semantics.
- Static composition is an instance of the Composite pattern [GHJV95]. This ensures a consistent interface between primitive (e.g., statically composed filter) and container (e.g., sequence of filters) objects.
- Streams [Edw95], Pipes and filters [Meu95, BMR<sup>+</sup>96] and Pipeline [Sha96] provide similar solutions. However, the data flow architecture is more general. For example, Pipes and filters takes a more static approach, where once the processing pipeline is set up, it is not allowed to change.

## References

- [Ack94] Philipp Ackermann. Design and implementation of an object-oriented media composition framework. In *Proc. International Computer Music Conference*, Aarhus, September 1994.
- [AEW96] Philipp Ackermann, Dominik Eichelberg, and Bernhard Wagner. Visual programming in an object-oriented framework. In *Proc. Swiss Computer Science Conference*, Zurich, Switzerland, October 1996.
- [AMS] Microsoft Corporation, Seattle, WA. *ActiveMovie Software Development Kit version 1.0*. <http://www.microsoft.com/devonly/tech/amov1doc/>.
- [AVS93] CONVEX Computer Corporation, Richardson, TX. *ConvexAVS Developer's Guide*, first edition, December 1993.
- [BMR<sup>+</sup>96] Frank Buschman, Regine Meunier, Hans Rohnert, Peter Sommerlad, and Michael Stal. *Pattern-Oriented Software Architecture—A System of Patterns*. John Wiley & Sons, July 1996. ISBN 0-47195-869-7.
- [Edw95] Stephen Edwards. *Streams: A Pattern for “Pull-Driven” Processing*, volume 1 of *“Pattern Languages of Program Design”*, chapter 21. Addison-Wesley, 1995. Edited by James O. Coplien and Douglas C. Schmidt. ISBN 0-201-60734-4.
- [EXP93] Silicon Graphics, Inc., Mountain View, CA. *IRIS Explorer User's Guide*, 1993.
- [Foo88] Brian Foote. Designing to facilitate change with object-oriented frameworks. Master's thesis, University of Illinois at Urbana-Champaign, 1988.
- [GHJV95] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns—Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1995. ISBN 0-201-63361-2.
- [Gur85] J. R. Gurd. The Manchester prototype dataflow computer. *Communications of the ACM*, 28(1), January 1985.
- [JBR96] MacDonald H. Jackson, J. Eric Baldeschwieler, and Lawrence A. Rowe. Berkeley Continuous Media Toolkit API. Submitted for publication, September 1996.
- [Lea94] Doug Lea. Design patterns for avionics control systems, November 1994. DSSA Adage Project ADAGE-OSW-94-01.
- [Lea96] Doug Lea. *Concurrent Programming in Java*. The Java Series. Addison-Wesley, 1996. ISBN 0-201-69581-2.
- [Lin94] Christopher J. Lindblad. A programming system for the dynamic manipulation of temporally sensitive data. Technical Report 637, Massachusetts Institute of Technology, August 1994. Laboratory for Computer Science.

- [Meu95] Regine Meunier. *The Pipes and Filters Architecture*, volume 1 of “*Pattern Languages of Program Design*”, chapter 22. Addison-Wesley, 1995. Edited by James O. Coplien and Douglas C. Schmidt. ISBN 0-201-60734-4.
- [MN98] Dragoş-Anton Manolescu and Klara Nahrstedt. A scalable approach to continuous-media processing. In *The 8th International Workshop on Research Issues in Data Engineering: Continuous-Media Databases and Applications*, Orlando, FL, February 1998.
- [NPA92] R. S. Nikhil, Gregory M. Papadopoulos, and Arvind. \*T: A multithreaded massively parallel architecture. In *Proc. 19th Annual International Symposium on Computer Architecture*, pages 156–167. ACM, 1992.
- [Pap91] Gregory M. Papadopoulos. *Implementation of a general-purpose dataflow multiprocessor*. MIT Press, 1991.
- [PLV96] Edward J. Posnak, R. Greg Lavender, and Harrick M. Vin. Adaptive pipeline: an object structural pattern for adaptive applications. In *The 3rd Pattern Languages of Programming conference*, Monticello, IL, September 1996.
- [PLV97] Edward J. Posnak, R. Greg Lavender, and Harrick M. Vin. An adaptive framework for developing multimedia software components. *Communications of the ACM*, 40(10), October 1997.
- [Rit84] Dennis M. Ritchie. A stream input-output system. *AT&T Bell Laboratories Technical Journal*, 63(8):1897–1910, October 1984.
- [Sha96] Mary Shaw. *Some Patterns for Software Architecture*, volume 2 of “*Pattern Languages of Program Design*”, chapter 16. Addison-Wesley, 1996. Edited by John M. Vlissides, James O. Coplien and Norman L. Kerth. ISBN 0-201-89527-7.
- [SM97] Inc. SUN Microsystems. Java media players. <http://java.sun.com/products/java-media/jmf>, November 1997.
- [Wal94] William F. Walker. *A Conversation-Based Framework For Musical Improvisation*. PhD thesis, University of Illinois at Urbana-Champaign, 1994.