

Chapter 2

SPECIFICATIONS AND MODELING

2.1 Requirements

Consistent with the simplified design flow (see fig. 1.6), we will now describe requirements and approaches for specifying and modeling embedded systems. Specifications for embedded systems provide **models** of the system under design (SUD). Models can be defined as follows [Jantsch, 2004]:

Definition: “A model is a simplification of another entity, which can be a physical thing or another model. The model contains exactly those characteristics and properties of the modeled entity that are relevant for a given task. A model is minimal with respect to a task if it does not contain any other characteristics than those relevant for the task”.

Models are described in languages. Languages should be capable of representing the following features¹:

- **Hierarchy:** Human beings are generally not capable of comprehending systems containing many objects (states, components) having complex relations with each other. The description of all real-life systems needs more objects than human beings can understand. Hierarchy (in combination with **abstraction**) is a key mechanism helping to solve this dilemma. Hierarchies can be introduced such that humans need to handle only a small number of objects at any time.

There are two kinds of hierarchies:

¹Information from the books of Burns et al. [Burns and Wellings, 1990], Bergé et al. [Bergé et al., 1995] and Gajski et al. [Gajski et al., 1994] is used in this list.

- **Behavioral hierarchies:** Behavioral hierarchies are hierarchies containing objects necessary to describe the system behavior. States, events and output signals are examples of such objects.
- **Structural hierarchies:** Structural hierarchies describe how systems are composed of physical components.

For example, embedded systems can be comprised of components such as processors, memories, actuators and sensors. Processors, in turn, include registers, multiplexers and adders. Multiplexers are composed of gates.

- **Component-based design** [Sifakis, 2008]: It must be “easy” to derive the behavior of a system from the behavior of its components. If two components are connected, the resulting new behavior should be predictable. Example: suppose that we add another component (say, some GPS unit) to a car. The impact of the additional processor on the overall behavior of the system (including busses etc.) should be predictable.
- **Concurrency:** Real-life systems are distributed, concurrent systems composed of components. It is therefore necessary to be able to specify concurrency conveniently. Unfortunately, human beings are not very good at understanding concurrent systems and many problems with real systems are actually a result of an incomplete understanding of possible behaviors of concurrent systems.
- **Synchronization and communication:** Components must be able to communicate and to synchronize. Without communication, components could not cooperate and we would use each of them in isolation. It must also be possible to agree on the use of resources. For example, it is necessary to express mutual exclusion.
- **Timing-behavior:** Many embedded systems are real-time systems. Therefore, explicit timing requirements are one of the characteristics of embedded systems. The need for explicit modeling of time is even more obvious from the term “cyber-physical system”. Time is one of the key dimensions of physics. Hence, timing requirements **must** be captured in the specification of embedded/cyber-physical systems.

However, standard theories in computer science model time only in a very abstract way. The O -notation is one of the examples. This notation just reflects growth rates of functions. It is frequently used to model run-times of algorithms, but it fails to describe real execution times. In physics, quantities have units, but the O -notation does not even have units. So, it would not distinguish between femtoseconds and centuries. A similar remark applies to termination properties of algorithms. Standard theories are concerned

with proving that a certain algorithm *eventually* terminates. For real-time systems, we need to show that an algorithm terminates in a given amount of time.

The resulting problems are very clearly formulated in a statement made by E. Lee: “*The lack of timing in the core abstraction (of computer science) is a flaw, from the perspective of embedded software*” [Lee, 2005].

According to Burns and Wellings [Burns and Wellings, 1990], modeling time must be possible in the following four contexts:

- Techniques for measuring **elapsed time**:
For many applications it is necessary to check, how much time has elapsed since some computation was performed. Access to a timer would provide a mechanism for this.
- Means for **delaying of processes** for a specified time:
Typically, real-time languages provide some delay construct. Unfortunately, typical implementations of embedded systems in software do not guarantee precise delays. Let us assume that task T should be delayed by some amount δ . Usually, this delay is implemented by changing task T 's state in the operating system from “ready” or “run” to “suspended”. At the end of this time interval, T 's state is changed from “suspended” to “ready”. This does not mean that the task actually executes. If some higher priority task is executing or if preemption is not used, the delayed task will be delayed longer than δ .
- Possibility to specify **timeouts**:
There are many situations in which we must wait for a certain event to occur. However, this event may actually not occur within a given time interval and we would like to be notified about this. For example, we might be waiting for a response from some network connection. We would like to be notified if this response is not received within some amount of time, say δ . Real-time languages usually also provide some timeout construct. Implementations of timeouts frequently come with the same problems which we mentioned for delays.
- Methods for specifying **deadlines** and **schedules**:
For many applications it is necessary to complete certain computations in a limited amount of time. For example, if the sensors of some car signal an accident, air-bags must be ignited within about ten milliseconds. In this context, we must guarantee that the software will decide whether or not to ignite the air-bags in that given amount of time. The air-bags could harm passengers, if they go off too late. Unfortunately, most languages do not allow to specify timing constraints. If they can

be specified at all, they must be specified in separate control files, pop-up menus etc. But the situation is still bad even if we are able to specify these constraints: many modern hardware platforms do not have a very predictable timing behavior. Caches, stalled pipelines, speculative execution, task preemption, interrupts etc. may have an impact on the execution time which is very difficult to predict. Accordingly, **timing analysis** (verifying the timing constraints) is a very hard design task.

- **State-oriented behavior:** It was already mentioned in Chapter 1 that automata provide a good mechanism for modeling reactive systems. Therefore, the state-oriented behavior provided by automata should be easy to describe. However, classical automata models are insufficient, since they cannot model timing and since hierarchy is not supported.
- **Event-handling:** Due to the reactive nature of embedded systems, mechanisms for describing events must exist. Such events may be external events (caused by the environment) or internal events (caused by components of the SUD).
- **Exception-oriented behavior:** In many practical systems exceptions do occur. In order to design dependable systems, it must be possible to describe actions to handle exceptions easily. It is not acceptable that exceptions must be indicated for each and every state (such as in the case of classical state diagrams). Example: In fig. 2.1, input k might correspond to an exception.




Figure 2.1. State diagram with exception k

Specifying this exception at each state makes the diagram very complex. The situation would get worse for larger state diagrams with many transitions. We will later show, how all the transitions can be replaced by a single one.

- **Presence of programming elements:** Popular programming languages have proven to be a convenient means of expressing computations that have to be performed. Hence, programming language elements should be available in the specification technique used. Classical state diagrams do not meet this requirement.

- **Executability:** Specifications are not automatically consistent with the ideas in people's heads. Executing the specification is a means of plausibility checking. Specifications using programming languages have a clear advantage in this context.
- **Support for the design of large systems:** There is a trend towards large and complex embedded software programs. Software technology has found mechanisms for designing such large systems. For example, object-orientation is one such mechanism. It should be available in the specification methodology.
- **Domain-specific support:** It would of course be nice if the same specification technique could be applied to all the different types of embedded systems, since this would minimize the effort for developing specification techniques and tool support. However, due to the wide range of application domains, there is little hope that one language can be used to efficiently represent specifications in all domains. For example, control-dominated, data-dominated, centralized and distributed applications-domains can all benefit from language features dedicated towards those domains.
- **Readability:** Of course, specifications must be readable by human beings. Otherwise, it would not be feasible to validate whether or not the specification meets the real intent of the the persons specifying the SUD. All design documents should also be machine-readable into order to process them in a computer. Therefore, specifications should be captured in languages which are readable by humans and by computers.

Initially, such specifications could use a natural language such as English or Japanese. Even this natural language description should be captured in a design document, so that the final implementation can be checked against the original document. However, natural languages are not sufficient for later design phases, since natural languages lack key requirements for specification techniques: it is necessary to check specifications for completeness, absence of contradictions and it should be possible to derive implementations from the specification in a systematic way. Natural languages do not meet these requirements.

- **Portability and flexibility:** Specifications should be independent of specific hardware platforms so that they can be easily used for a variety of target platforms. Ideally, changing the hardware platform should have no impact on the specification. In practice, small changes may have to be tolerated.
- **Termination:** It should be feasible to identify processes that will terminate from the specification. This means that we would like to use specifications

for which the halting problem (the problem of figuring out whether or not a certain algorithm will terminate; see, for example [Sipser, 2006]) is decidable.

- **Support for non-standard I/O-devices:** Many embedded systems use I/O-devices other than those typically found in a PC. It should be possible to describe inputs and outputs for those devices conveniently.
- **Non-functional properties:** Actual SUDs must exhibit a number of non-functional properties, such as fault-tolerance, size, extendibility, expected lifetime, power consumption, weight, disposability, user friendliness, electromagnetic compatibility (EMC) etc. There is no hope that all these properties can be defined in a formal way.
- **Support for the design of dependable systems:** Specification techniques should provide support for designing dependable systems. For example, specification languages should have unambiguous semantics, facilitate formal verification and be capable of describing security and safety requirements.
- **No obstacles to the generation of efficient implementations:** Since embedded systems must be efficient, no obstacles prohibiting the generation of efficient realizations should be present in the specification.
- **Appropriate model of computation (MoC):** The von-Neumann model of sequential execution combined with some communication technique is a commonly used MoC. However, this model has a number of serious problems, in particular for embedded system applications. Problems include:
 - Facilities for describing timing are lacking.
 - Von-Neumann computing is implicitly based on accesses to globally shared memory (such as in Java). It has to guarantee mutually exclusive access to shared resources. Otherwise, multi-threaded applications allowing pre-emptions at any time can lead to very unexpected program behaviors². Using primitives for ensuring mutually exclusive access can, however, very easily lead to deadlocks. Possible deadlocks may be difficult to detect and may remain undetected for many years.

Lee [Lee, 2006] provided a very alarming example in this direction. Lee studied implementations of a simple observer pattern in Java. For this pattern, changes of values must be propagated from some producer to a set of subscribed observers. This is a very frequent pattern in

²Examples are typically provided in courses on operating systems.

embedded systems, but is difficult to implement correctly in a multi-threaded von-Neumann environment with preemptions. Lee's code is a possible implementation of the observer pattern in Java for a multi-threaded environment:

```
public synchronized void addListener(listener) {...}
public synchronized void setValue(newvalue) {
    myvalue=newvalue;
    for (int i=0; i<mylisteners.length; i++) {
        myListeners[i].valueChanged(newvalue)
    }
}
```

Method `addListener` subscribes new observers, method `setValue` propagates new values to subscribed observers. In general, in a multithreaded environment, threads can be pre-empted any time, resulting in an arbitrarily interleaved execution of these threads. Adding observers while `setValue` is already active could result in complications, i.e. we would not know if the new value had reached the new listener. Moreover, the set of observers constitutes a global data structure of this class. Therefore, these methods are synchronized in order to avoid changing the set of observers while values are already partially propagated. This way, only one of the two methods can be active at a given time. This mutual exclusion is necessary to prevent unwanted interleavings of the execution of methods in a multithreaded environment. Why is this code problematic? It is problematic since `valueChanged` could attempt to get exclusive access to some resource (say, R). If that resource is allocated to some other method (say, A), then this access is delayed until A releases R. If A calls (possibly indirectly) `addListener` or `setValue` before releasing R, then these methods will be in a deadlock: `setValue` waits for R, releasing R requires A to proceed, A cannot proceed before its call of `setValue` or `addListener` is serviced. Hence, we will have a deadlock.

This example demonstrates the existence of deadlocks resulting from using multiple threads which can be arbitrarily pre-empted and therefore require mutual exclusion for their access to critical resources. Lee showed [Lee, 2006] that many of the proposed “solutions” of the problem are problematic themselves. So, even this very simple pattern is difficult to implement correctly in a multi-threaded von-Neumann environment. This example demonstrates that concurrency is really difficult to understand for humans and there may be the risk of oversights, even after very rigorous code inspections.

Lee came to the drastic conclusion that “*nontrivial software written with threads, semaphores, and mutexes is incomprehensible to humans*”

and that “*threads as a concurrency model are a poor match for embedded systems. ... they work well only ... where best-effort scheduling policies are sufficient*” [Lee, 2005].

The underlying reasons for deadlocks have been studied in detail in the context of operating systems (see, for example, [Stallings, 2009]). From this context, it is well-known that four conditions must hold at run-time to get into a deadlock: mutual exclusion, no pre-emption of resources, holding resources while waiting for more, and a cyclic dependency between threads. Obviously, all four conditions are met in the above example. The theory of operating systems provides no general way out of this problem. Rare deadlocks may be acceptable for a PC, but they are clearly unacceptable for a safety-critical system.

We would like to specify SUDs such that we do not have to care about possible deadlocks. Therefore, it makes sense to study non-von-Neumann MoCs avoiding this problem. We will study such MoCs from the next section onwards. It will be shown that the observer pattern can be easily implemented in other MoCs.

From the list of requirements, it is already obvious that there will not be any single formal language capable of meeting all these requirements. Therefore, in practice, we must live with compromises and possibly also with a mixture of languages (each of which would be appropriate for describing a certain type of problems). The choice of the language used for an actual design will depend on the application domain and the environment in which the design has to be performed. In the following, we will present a survey of languages that can be used for actual designs. These languages will demonstrate the essential features of the corresponding model of computation.

2.2 Models of computation

Models of computation (MoCs) describe the mechanism assumed for performing computations. In the general case, we must consider systems comprising components. It is now common practice to strictly distinguish between the computations performed in the components and communication. Accordingly, MoCs define (see also [Lee, 1999], [Janka, 2002], [Jantsch, 2004], [Jantsch, 2006]):

- **Components** and the organization of computations in such components: Procedures, processes, functions, finite state machines are possible components.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

- **Communication protocols:** These protocols describe methods for communication between components. Asynchronous message passing and rendez-vous based communication are examples of communication protocols.

Relations between components can be captured in graphs. In such graphs, we will refer to the computations also as processes or tasks. Accordingly, relations between these will be captured by **task graphs** and **process networks**. Nodes in the graph represent components performing computations. Computations map input data streams to output data streams. Computations are sometimes implemented in high-level programming languages. Typical computations contain (possibly non-terminating) iterations. In each cycle of the iteration, they consume data from their inputs, processes the data received, and generate data on their output streams. Edges represent relations between components. We will now introduce these graphs at a more detailed level.

The most obvious relation between computations is their causal dependence: Many computations can only be executed after other computations have terminated. This dependence is typically captured in **dependence graphs**. Fig. 2.2 shows a dependence graph for a set of computations.




Figure 2.2. Dependence graph

Definition: A dependence graph is a directed graph $G = (V, E)$, where V is the set of **vertices** or **nodes** and E is the set of **edges**. $E \subseteq V \times V$ imposes a relation on V . If $(v_1, v_2) \in E$, then v_1 is called an **immediate predecessor** of v_2 and v_2 is called an **immediate successor** of v_1 . Suppose E^* is the transitive closure of E . If $(v_1, v_2) \in E^*$, then v_1 is called a **predecessor** of v_2 and v_2 is called a **successor** of v_1 .

Such dependence graphs form a special case of task graphs. Task graphs may contain more information than modeled in fig. 2.2. For example, task graphs may include the following extensions of dependence graphs:

- 1 **Timing information:** Tasks may have arrival times, deadlines, periods, and execution times. In order to take these into account while scheduling computations, it may be useful to include this information in the graphs. Adopting the notation used in the book by Liu [Liu, 2000], we include possible execution intervals in fig. 2.3. Computations T_1 to T_3 are assumed to

be independent. The first number in brackets is the arrival time, the second the deadline (execution times are not explicitly shown). For example, T_1 is assumed to be available at time 0 and should be completed no later than at time 7.




Figure 2.3. Graphs including timing information

Significantly more complex combinations of timing and dependence relations can exist.

- 2 **Distinction between different types of relations** between computations: Precedence relations just model constraints for possible execution sequences. At a more detailed level, it may be useful to distinguish between constraints for scheduling and communication between computations. Communication can again be described by edges, but additional information may be available for each of the edges, such as the time of the communication and the amount of information exchanged. Precedence edges may be kept as a separate type of edges, since there could be situations in which computations must execute sequentially even though they do not exchange information.

In fig. 2.2, input and output (I/O) is not explicitly described. Implicitly it is assumed that computations without any predecessor in the graph might be receiving input at some time. Also, they might generate output for the successor and that this output could be available only after the computation has terminated. It is often useful to describe input and output more explicitly. In order to do this, another kind of relation is required. Using the same symbols as Thoen [Thoen and Catthoor, 2000], we use partially filled circles for denoting input and output. In fig. 2.4, partially filled circles identify I/O edges.




Figure 2.4. Graphs including I/O-nodes and edges

- 3 **Exclusive access to resources:** Computations may be requesting exclusive access to some resource, for example to some input/output device or some communication area in memory. Information about necessary exclusive access should be taken into account during scheduling. Exploiting this information might, for example, be used to avoid the priority inversion problem (see page 186). Information concerning exclusive access to resources can be included in the graphs.
- 4 **Periodic schedules:** Many computations, especially in digital signal processing, are periodic. This means that we must distinguish more carefully between a task and its execution (the latter is frequently called a **job** [Liu, 2000]). Task graphs for such schedules are infinite. Fig. 2.5 shows a task graph including jobs J_{n-1} to J_{n+1} of a periodic task.




Figure 2.5. Graph including jobs

- 5 **Hierarchical graph nodes:** The complexity of the computations denoted by graph nodes may be quite different. On one hand, specified computations may be quite involved and contain thousands of lines of program code. On the other hand, programs can be split into small pieces of code so that in the extreme case, each of the nodes corresponds only to a single operation. The level of complexity of graph nodes is also called their **granularity**. Which granularity should be used? There is no universal answer to this. For some purposes, the granularity should be as large as possible. For example, if we consider each of the nodes as one process to be scheduled by the RTOS, it may be wise to work with large nodes in order to minimize context-switches between different processes. For other purposes, it may be better to work with nodes modeling just a single operation. For example, nodes will must be mapped to hardware or to software. If a certain operation (such as the frequently used Discrete Cosine Transform, or DCT) can be mapped to special purpose hardware, then it should not be buried in a complex node that contains many other operations. It should rather be modeled as its own node. In order to avoid frequent changes of the granularity, hierarchical graph nodes are very useful. For example, at a high hierarchical level, the nodes may denote complex tasks, at a lower level basic blocks³ and at an even lower level individual arithmetic operations. Fig.

³Basic blocks are code blocks of maximum length not including any branch except possibly at their end and not being branched into.

2.6 shows a hierarchical version of the dependence graph in fig. 2.2, using a rectangle to denote a hierarchical node.




Figure 2.6. Hierarchical task graph

As indicated above, MoCs can be classified according to the models of communication (reflected by edges in the task graphs) and the model of computations within the components (reflected by the nodes in the task graph). In the following, we will explain prominent examples of such models:

■ Models of communication:

We distinguish between two communication paradigms: **shared memory** and **message passing**. Other communication paradigms exist (e.g. entangled states in quantum mechanics [Bouwmeester et al., 2000]), but are not considered in this book.

– **Shared memory:**

For shared memory, communication is carried out by accesses to the same memory from all components.

Access to shared memory should be protected, unless access is totally restricted to reads. If writes are involved, exclusive access to the memory must be guaranteed while components are accessing shared memories. Segments of program code, for which exclusive access must be guaranteed, are called **critical sections**. Several mechanisms for guaranteeing exclusive access to resources have been proposed. These include semaphores, conditional critical regions and monitors. Refer to books on operating systems (e.g. Stallings [Stallings, 2009]) for a description of the different techniques. Shared memory-based communication can be very fast, but is difficult to implement in multiprocessor systems if no common memory is physically available.

- **Message passing:** For message passing, messages are sent and received. Message passing can be implemented easily even if no common memory is available. However, message passing is generally slower than shared memory based communication. For this kind of communication, we can distinguish between the following three techniques:

- * **asynchronous message passing**, also called **non-blocking communication** : In asynchronous message passing, components communicate by sending messages through channels which can buffer the messages. The sender does not need to wait for the recipient to be ready to receive the message. In real life, this corresponds to sending a letter or an e-mail. A potential problem is the fact that messages may must be stored and that message buffers can overflow. There are several variations of this scheme, including communicating finite state machines (see page 52) and data flow models (see page 58).
- * **synchronous message passing or blocking communication, *rendez-vous* based communication**: In synchronous message passing, available components communicate in atomic, instantaneous actions called *rendez-vous*. The component reaching the point of communication first has to wait until the partner has also reached its point of communication. In real life, this corresponds to physical meetings or phone calls. There is no risk of overflows, but the performance may suffer. Examples of languages following this model of computation include CSP (see page 100) and ADA (see page 100).
- * **extended *rendez-vous*, remote invocation**: In this case, the sender is allowed to continue only after an acknowledgment has been received from the recipient. The recipient does not have to send this acknowledgment immediately after receiving the message, but can do some preliminary checking before actually sending the acknowledgment.

■ **Organization of computations within the components:**

- **Von-Neumann model**: This model is based on the sequential execution of sequences of primitive computations.
- **Discrete event model**: In this model, there are events carrying a totally ordered time stamp, indicating the time at which the event occurs. Discrete event simulators typically contain a global event queue sorted by time. Entries from this queue are processed according to this order. The disadvantage is that this model relies on a global notion of event queues, making it difficult to map the semantic model onto parallel implementations. Examples include VHDL (see page 78), SystemC (see page 94), and Verilog (see page 96).
- **Finite state machines (FSMs)**: This model is based on the notion of a finite set of states, inputs, outputs, and transitions between states.

Several of these machines may need to communicate, forming so-called **communicating finite state machines (CFSMs)**.

- **Differential equations:** Differential equations are capable of modeling analog circuits and physical systems. Hence, they can find applications in cyber-physical system modeling.
- **Combined models:** Actual languages are typically combining a certain model of communication with an organization of computations within components. For example, SDL (see page 52) combines finite state machines with asynchronous message passing. StateCharts (see page 40) combines finite state machines with shared memories. ADA (see page 100) and CSP (see page 100) combine von-Neumann execution with synchronous message passing. Fig. 2.7 gives an overview of combined models which we will consider in this chapter. This figure also includes examples of languages for most of the MoCs.

Communication/ Organization of components	Shared memory	Message passing	
		synchronous	asynchronous
Undefined components	Plain text or graphics, use cases (Message) sequence charts		
Communicating finite state machines (CFSMs)	StateCharts		SDL
Data flow	(not useful)		Kahn networks SDF
Petri nets		C/E nets, P/T nets, ...	
Discrete event (DE) model ⁴	VHDL, Verilog SystemC	(Only experimental systems) Distributed DE in Ptolemy	
Von-Neumann model	C, C++, Java	C, C++, Java, ... with libraries CSP, ADA	

Figure 2.7. Overview of MoCs and languages considered

Some MoCs have advantages in certain application areas, while other have advantages in others. Choosing the “best” MoC for a certain application may be difficult. Being able to mix MoCs (such as in the Ptolemy framework [Davis et al., 2001]) can be a way out of this dilemma. Also, models may be translated from one MoC into another one. Non-von-Neumann models are frequently

⁴The classification of VHDL, Verilog and SystemC is based on the implementation of these languages in simulators. Message passing can be modeled in these languages “on top” of the simulation kernel.

translated into von-Neumann models. The distinction between the different models is blurring if the translation between them is easy.

Designs starting from non-von-Neumann models are frequently called **model-based designs**. The key idea of model-based design is to have some abstract model of the system under design. Properties of the system can then be studied at the level of this model, without having to care about software code. Software code is generated only after the behavior of the model has been studied in detail and this software is generated automatically. The term “model-based design” is not precisely defined. It is usually associated with models of control systems, comprising traditional control system elements such as integrators, differentiators etc. However, this view seems to be too restricted, since we could also start with abstract models of consumer systems.

In the following, we will present different MoCs, using existing languages as examples for demonstrating their features. A related (but shorter) survey is provided by Edwards [Edwards, 2006]. For a more comprehensive presentation see [Gomez and Fernandes, 2010].

2.3 Early design phases

The very first ideas about systems are frequently captured in a very informal way, possibly on paper. Frequently, only descriptions of the SUD in a natural language such as English or Japanese exist in the early phases of design projects. They are typically using a very informal style. These descriptions should be captured in some machine-readable document. They should be encoded in the format of some word processor and stored by a tool managing design documents. A good tool would allow links between the requirements, a dependence analysis as well as version management.

DOORS® [IBM, 2010b] exemplifies such a tool.

2.3.1 Use cases

For many applications, it is beneficial to envision potential usages of the SUD. Such usages are captured in **use cases**. Use cases describe possible applications of the SUD. Different notations for use cases could be used.

Support for a systematic approach to early specification phases is the goal of the so-called UML standardization effort [Object Management Group (OMG), 2010b], [Fowler and Scott, 1998], [Haugen and Moller-Pedersen, 2006]. UML stands for “Unified Modeling Language”. UML was designed by leading software technology experts and is supported by commercial tools. UML primarily aims at the support of the software design process. UML provides a standardized form for use cases.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

For use cases, there is neither a precisely specified model of the computations nor is there a precisely specified model of the communication. It is frequently argued that this is done intentionally in order to avoid caring about too many details during the early design phases.

For example, fig. 2.8 shows some use cases for an answering machine⁵.

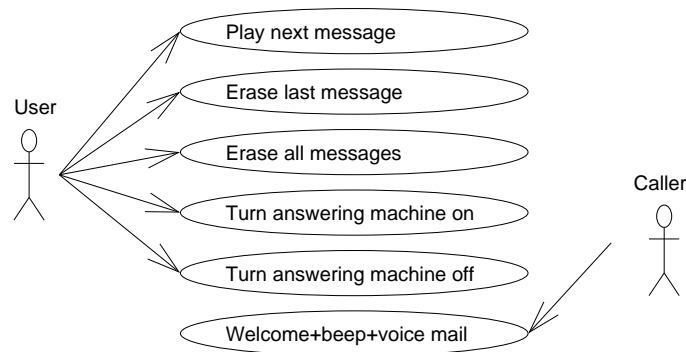


Figure 2.8. Use case example

Use cases identify different classes of users as well as the applications to be supported by the SUD. In this way, it is possible to capture expectations at a very high level.

2.3.2 (Message) Sequence Charts

At a slightly more detailed level, we might want to explicitly indicate the sequences of messages which must be exchanged between components in order to implement some use of the SUD. **Sequence charts** (SCs) -earlier called **message sequence charts** (MSCs)- provide a mechanism for this. Sequence charts use one dimension (usually the vertical dimension) of a 2-dimensional chart to denote sequences and the second dimension to reflect the different communication components. SCs describe partial orders between message transmissions. SCs display a possible behavior of a SUD.

SCs are also standardized in UML. UML 2.0 has extended SCs with elements allowing a more detailed description than UML 1.0. Fig. 2.9 shows one of the use cases of the answering machine as an example.

Dashed lines are so-called “life-lines”. Messages are assumed to be ordered according to their sequence along the life-line. We assume that, in this ex-

⁵We assume that UML is covered in-depth in a software engineering course included in the curriculum. Therefore, UML is only briefly discussed in this book.

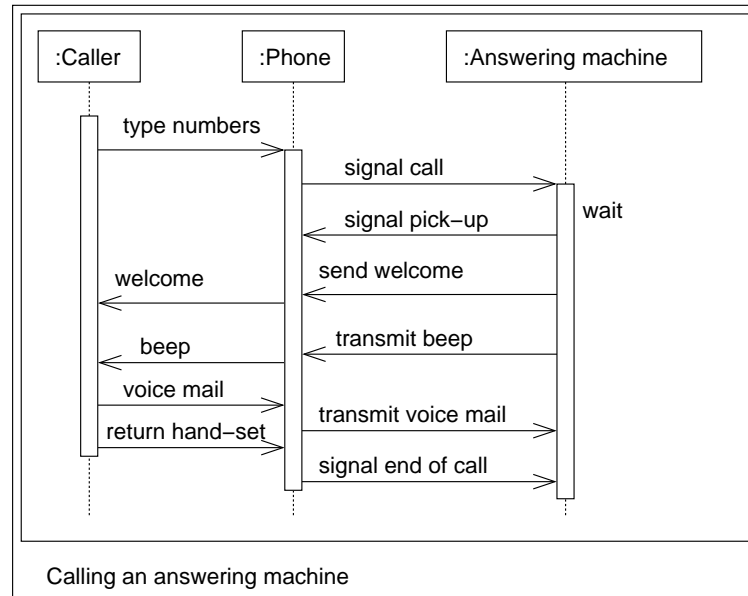


Figure 2.9. Answering machine in UML

ample, all information is sent in the form of messages. Arrows used in this diagram denote asynchronous messages. This means several messages can be sent by a sender without waiting for the receipt to be confirmed. Boxes on top of life-lines represent active control at the corresponding component. In the example, the answering machine is waiting for the user to pick up the phone within a certain amount of time. If he or she fails to do so, the machine signals a pick-up itself and sends a welcome message to the caller. The caller is then supposed to leave a voice-mail message. Alternative sequences (e.g. an early termination of the call by the caller or the callee picking up the phone) are not shown.

Complex control-dependent actions cannot be described by SCs. Other MoCs must be used for this. Frequently, certain preconditions must be met for a SC to apply. Such preconditions, a distinction between sequences which might happen and those which must happen, as well as other extensions are available in the so-called Live Sequence Charts [Damm and Harel, 2001].

Time/distance diagrams (TDDs) are a commonly used variant of SCs. In time/distance diagrams, the vertical dimension reflects real time, not just sequence. In some cases, the horizontal dimension also models the real distance between the components.

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!

TDDs provide the right means for visualizing schedules of trains or buses. Fig. 2.10 is an example. This example refers to trains between Amsterdam, Cologne, Brussels and Paris. Trains can run from either Amsterdam or Cologne to Paris via Brussels. Aachen is included as an intermediate stop between Cologne and Brussels. Vertical segments correspond to times spent at stations. For one of the trains, there is a timing overlap between the trains coming from Cologne and Amsterdam at Brussels. There is a second train which travels between Paris and Cologne which is not related to an Amsterdam train.

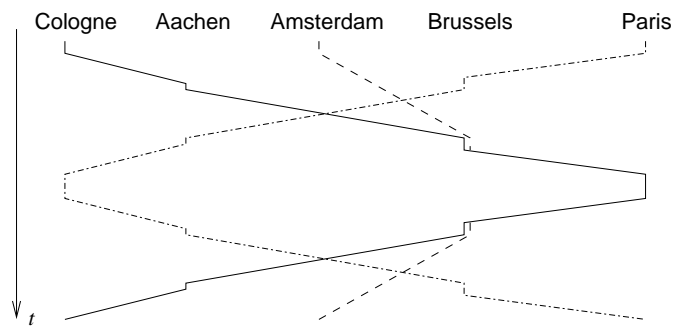


Figure 2.10. Time/distance diagram

This example and other examples can be simulated with the *levi* simulation software [Sirocic and Marwedel, 2007d]. A larger, more realistic example is shown in fig. 2.11. This example [Huerlimann, 2003] describes simulated Swiss railway traffic in the Lötschberg area. Slow and fast trains can be distinguished by their slope in the graph. This modeling technique is very frequently used in practice.

One of the key distinctions between the type of diagrams shown in figs. 2.9 and 2.11 is that fig. 2.9 does not include any reference to real time. UML was initially not designed for real-time applications. UML 2.0 includes **timing diagrams** as a special class of diagrams. Such diagrams enable referring to physical time. Also, certain UML “profiles” (see page 112) allow additional annotations to refer to time [Martin and Müller, 2005], [Müller, 2007].

TDDs are appropriate means for representing typical schedules. However, SCs and TDDs fail to provide information about necessary synchronization. For example, in the presented example of fig. 2.10 it is not known whether the timing overlap at Brussels happens coincidentally or whether some real synchronization for connecting trains is required. Furthermore, permissible deviations from the presented schedule (min/max timing behavior) can hardly be included in these charts.




Figure 2.11. Railway traffic displayed by a TDD (courtesy H. Brändli, IVT, ETH Zürich), ©ETH Zürich

2.4 Communicating finite state machines (CFSMs)

If we start to represent our SUD at a more detailed level, we need more precise models. We mentioned at the beginning of this chapter that we need to describe state-oriented behavior. State diagrams are a classical means of doing this. Fig. 2.12 (the same as fig. 2.1) shows an example of a classical state diagram, representing a **finite state machine (FSM)**.




Figure 2.12. State diagram

Circles denote states. We will consider FSMs for which only one of their states is active. Such FSMs are called **deterministic** FSMs. Edges denote state transitions. Edge labels represent events. Let us assume that a certain state of the FSM is active, and that an event happens which corresponds to one of the out-going edges for the active state. Then, the FSM will change its state from the currently active state to the one indicated by the edge. FSM

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

may be implicitly clocked. Such FSMs are called **synchronous FSMs**. For synchronous FSMs, state changes will happen only at clock transitions. FSMs may also generate output (not shown in fig. 2.12). For more information about classical FSMs refer to, for example, Kohavi [Kohavi, 1987].

2.4.1 Timed automata

Classical FSMs do not provide information about time. In order to model time, classical automata have been extended to also include timing information. Timed automata are essentially automata extended with real-valued variables. *“The variables model the logical clocks in the system, that are initialized with zero when the system is started, and then increase synchronously with the same rate. Clock constraints i.e. guards on edges are used to restrict the behavior of the automaton. A transition represented by an edge can be taken when the clocks values satisfy the guard labeled on the edge. Clocks may be reset to zero when a transition is taken”* [Bengtsson and Yi, 2004].

Fig. 2.13 shows an example.

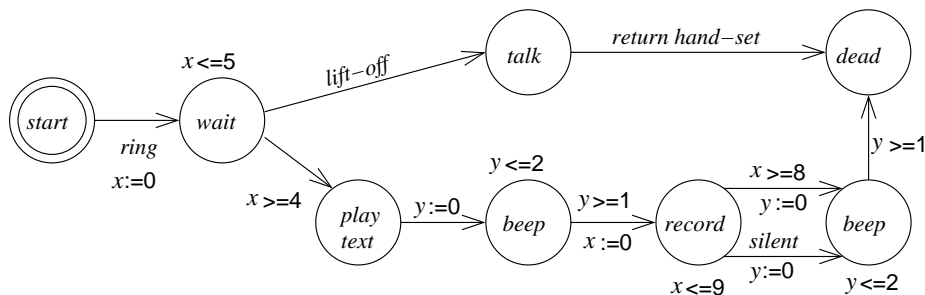


Figure 2.13. Servicing an incoming line in an answering machine

The answering machine is usually in the initial state on the left. Whenever a ring signal is received, clock x is reset to 0 and a transition into a waiting state is made. If the called person lifts off the hand-set, talking can take place until the hand-set is returned. Otherwise, a transition to state *play text* can take place if time has reached a value of 4.

Once the transition took place, a recorded message is played and this phase is terminated with a beep. Clock y ensures that this beep lasts at least one time unit. After the beep, clock x is reset to 0 again and the answering machine is ready for recording. If time has reached a value of 8 or if the caller remains silent, the next beep is played. This second beep again lasts at least one time unit. After the second beep, a transition is made into the final state. In this

example, transitions are either caused by inputs (such as *lift-off*) or by so-called **clock constraints**.

Clock constraints describe transitions which **can** take place, but they do not have to. In order to make sure that transitions actually take place, additional **location invariants** can be defined. Location invariants $x \leq 5$, $x \leq 9$ and $y \leq 2$ are used in the example such that transitions will take place no later than one time unit after the enabling condition became true. Using two clocks is for demonstration purposes only; a single clock would be sufficient.

Formally speaking, timed automata can be defined as follows [Bengtsson and Yi, 2004]:

Let C be a set of real-valued, non-negative variables representing clocks. Let Σ be a finite alphabet of possible inputs.

Definition: A **clock constraint** is a conjunctive formula of atomic constraints of the form $x \circ n$ or $x - y \circ n$ for $x, y \in C$, $\circ \in \{\leq, <, =, >, \geq\}$ and $n \in \mathbb{N}$

Note that constants n used in the constraints must be integers, even though clocks are real-valued. An extension to rational constants would be easy, since they could be turned into integers with simple multiplications. Let $B(C)$ be the set of clock constraints.

Definition [Bengtsson and Yi, 2004]: A **timed automaton** is a tuple (S, s_0, E, I) where:

- S is a finite set of states.
- s_0 is the initial state.
- $E \subseteq S \times B(C) \times \Sigma \times 2^C \times S$ is the set of edges. $B(C)$ models the conjunctive condition which must hold and Σ models the input which is required for a transition to be enabled. 2^C reflects the set of clock variables which are reset whenever the transition takes place.
- $I : S \rightarrow B(C)$ is the set of invariants for each of the states. $B(C)$ represents the invariant which must hold for a particular state S . This invariant is described as a conjunctive formula.

This first definition is usually extended to allow parallel compositions of timed automata. Timed automata having a large number of clocks tend to be difficult to understand. More details about timed automata can be found, for example, in papers by Dill et al. [Dill and Alur, 1994] and Bengtsson et al. [Bengtsson and Yi, 2004].

Timed automata extend classical automata with timing information. However, many of our requirements for specification techniques are not met by timed

automata. In particular, in their standard form, they do not provide hierarchy and concurrency.

2.4.2 StateCharts: implicit shared memory communication

StateCharts is presented here as a very prominent example of a language based on automata and supporting hierarchical models as well as concurrency. It does include a limited way of specifying timing.

The StateCharts language was introduced by David Harel [Harel, 1987] in 1987 and later described more precisely in [Drusinsky and Harel, 1989]. According to Harel, the name was chosen since it was “*the only unused combination of **flow** or **state** with **diagram** or **chart***”.

2.4.2.1 Modeling of hierarchy

StateCharts describes extended FSMs. Due to this, they can be used for modeling state-oriented behavior. The key extension is **hierarchy**. Hierarchy is introduced by means of **super-states**.

Definitions:

- States comprising other states are called **super-states**.
- States included in super-states are called **sub-states** of the super-states.

Fig. 2.14 shows a StateCharts example. It is a hierarchical version of fig. 2.12.




Figure 2.14. Hierarchical state diagram

Super-state S includes states A, B, C, D and E . Suppose the FSM is in state Z (we will also call Z to be an **active state**). Now, if input m is applied to the FSM, then A and S will be the new active states. If the FSM is in S and input k is applied, then Z will be the new active state, regardless of whether the FSM is in sub-states A, B, C, D or E of S . In this example, all states contained in S are non-hierarchical states. In general, sub-states of S could again be

super-states consisting of sub-states themselves. Also, **whenever a sub-state of some super-state is active, the super-state is active as well.**

Definitions:

- Each state which is not composed of other states is called a **basic state**.
- For each basic state s , the super states containing s are called **ancestor states**.

The FSM of fig. 2.14 can only be in one of the sub-states of super-state S at any time. Super states of this type are called **OR-super-states**⁶.

In fig. 2.14, k might correspond to an exception for which state S has to be left. The example already shows that the hierarchy introduced in StateCharts enables a compact representation of exceptions.

StateCharts allows hierarchical descriptions of systems in which a system description comprises descriptions of subsystems which, in turn, may contain descriptions of subsystems. The **hierarchy** of the entire system can be represented by a **tree**. The root of the tree corresponds to the system as a whole, and all inner nodes correspond to hierarchical descriptions (in the case of StateCharts called super-nodes). The leaves of the hierarchy are non-hierarchical descriptions (in the case of StateCharts called basic states).

So far, we have used explicit, direct edges to basic states to indicate the next state. The disadvantage of that approach is that the internal structure of super-states cannot be hidden from the environment. However, in a true hierarchical environment, we should be able to hide the internal structure so that it can be described later or changed later without affecting the environment. This is possible with other mechanisms for describing the next state.

The first additional mechanism is the **default state mechanism**. It can be used in super-states to indicate the particular sub-states that will become active if the super-states become active. In diagrams, default states are identified by edges starting at small filled circles. Fig. 2.15 shows a state diagram using the default state mechanism. It is equivalent to the diagram in fig. 2.14. Note that the filled circle does not constitute a state itself.

Another mechanism for specifying next states is the **history mechanism**. With this mechanism, it is possible to return to the last sub-state that was active before a super-state was left. The history mechanism is symbolized by a circle containing the letter H. In order to define the next state for the very initial

⁶More precisely, they should be called XOR-super-states, since the FSM is in **either** A, B, C, D or E . However, this name is not commonly used in the literature.




Figure 2.15. State diagram using the default state mechanism

transition into a super-state, the history mechanism is frequently combined with the default mechanism. Fig. 2.16 shows an example.




Figure 2.16. State diagram using the history and the default state mechanism

The behavior of the FSM is now somewhat different. If we input m while the system is in Z , then the FSM will enter A if this is the very first time we enter S , and otherwise it will enter the last state that we were in before leaving S . This mechanism has many applications. For example, if k denotes an exception, we could use input m to return to the state we were in before the exception. States A, B, C, D and E could also call Z like a procedure. After completing “procedure” Z , we would return to the calling state.

Fig. 2.16 can also be redrawn as shown in fig. 2.17. In this case, the symbols for the default and the history mechanism are combined.




Figure 2.17. Combining the symbols for the history and the default state mechanism

Specification techniques must also be able to describe concurrency conveniently. Towards this end, the StateChart language provides a second class of super-states, so called **AND**-states.

Definition: Super-states S are called **AND-super-states** if the system containing S will be in all of the sub-states of S whenever it is in S .

An AND-super-state is included in the answering machine example shown in fig. 2.18.

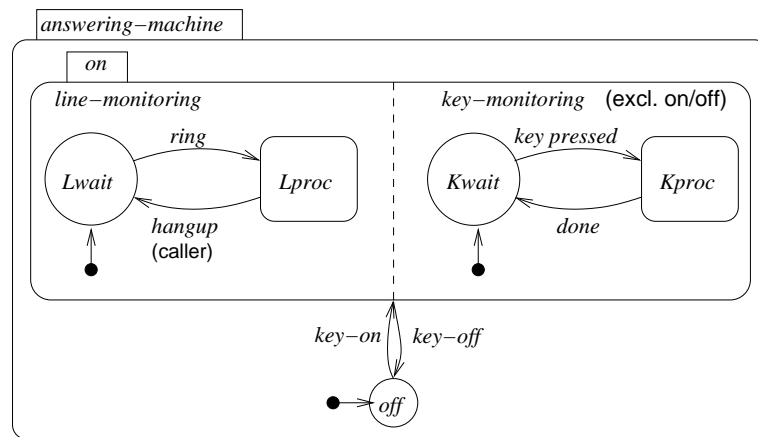


Figure 2.18. Answering machine

An answering machine normally performs two tasks concurrently: it is monitoring the line for incoming calls and the keys for user input. In fig. 2.18, the corresponding states are called *Lwait* and *Kwait*. Incoming calls are processed in state *Lproc* while the response to pressed keys is generated in state *Kproc*. For the time being, we assume that the on/off switch (generating events *key-off* and *key-on*) is decoded separately and pushing it does not result in entering *Kproc*. If this switch is pushed, the line monitoring state as well as the key monitoring state are left and re-entered only if the machine is switched on. At that time, default states *Lwait* and *Kwait* are entered. While switched on, the machine will always be in the line monitoring state as well as in the key monitoring state.

For AND-super-states, the sub-states entered as a result of some event can be defined independently. There can be any combination of history, default and explicit transitions. It is crucial to understand that **all** sub-states will always be entered, even if there is just one explicit transition to one of the sub-states. Accordingly, transitions out of an AND-super-state will always result in leaving **all** the sub-states.

For example, let us modify our answering machine such that the on/off switch, like all other switches, is decoded in state *Kproc* (see fig. 2.19).

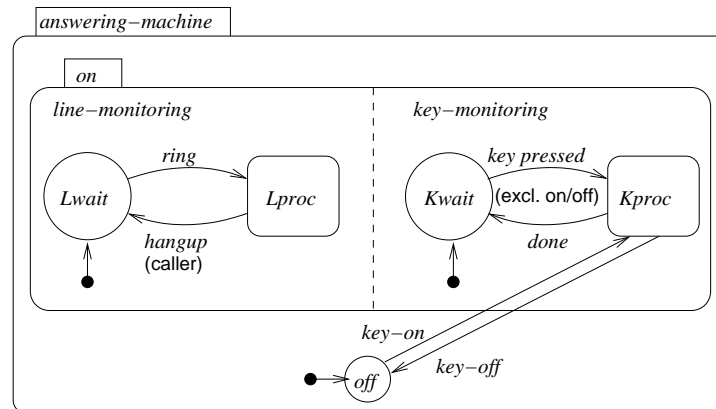


Figure 2.19. Answering machine with modified on/off switch processing

If pushing that key is detected in *Kproc*, a transition is made to the *off* state. This transition results in leaving the line-monitoring state as well. Switching the machine on again results in also entering the line-monitoring state.

AND-super-states provide the key mechanism for describing concurrency. Each sub-state can be considered a state machine by itself. These machines are communicating with each other, forming **communicating finite state machines** (CFSMs). This term has been used as the title of this section.

Summarizing, we can state the following: **States in StateCharts diagrams are either AND-states, OR-states or basic states.**

2.4.2.2 Timers

Due to the requirement to model time in embedded systems, StateCharts also provides timers. Timers are denoted by the symbol shown in fig. 2.20 (left).




Figure 2.20. Timer in StateCharts

After the system has been in the state containing the timer for the specified period, a time-out will occur and the system will leave the specified state. Timers can also be used hierarchically.

Timers can be employed, for example, at the next lower level of the hierarchy of the answering machine in order to describe the behavior of state *Lproc*. Fig. 2.21 shows a possible behavior for that state. The timing specification is slightly different from the one in fig. 2.13.

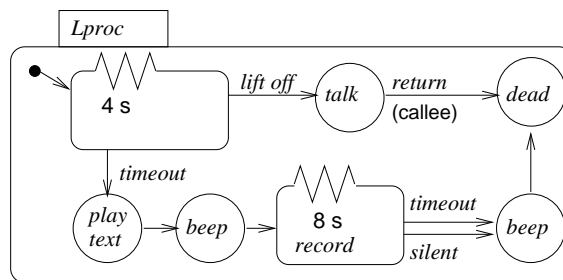


Figure 2.21. Servicing the incoming line in *Lproc*

Due to the exception-like transition for hangups by the caller in fig. 2.18, state *Lproc* is terminated whenever the caller hangs up. For hangups (returns) by the callee, the design of state *Lproc* results in an inconvenience: If the callee hangs up the phone first, the telephone will be dead (and quiet) until the caller has also hung up the phone.

StateCharts do include a number of other language elements. For a full description refer to Harel [Harel, 1987]. A more detailed description of the semantics of StateCharts is described by Drusinsky and Harel [Drusinsky and Harel, 1989].

2.4.2.3 Edge labels and StateMate semantics

Until now, we have not considered outputs generated by our extended FSMs. Generated outputs can be specified using edge labels. The general form of an edge label is “*event*[*condition*]/*reaction*”. All three label parts are optional. The *reaction*-part describes the reaction of the FSM to a state transition. Possible reactions include the generation of events and assignments to variables. The *condition*-part implies a test of the values of variables or a test of the current state of the system. The *event*-part refers to a test of current events. Events can be generated either internally or externally. Internal events are generated as a result of some transition and are described in *reaction*-parts. External events are usually described in the model environment.

Examples:

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

- *on-key* / *on:=1* (Event-test and variable assignment),
- [*on=1*] (Condition test for a variable value),
- *off-key* [not in *Lproc*] / *on:=0* (Event-test, condition test for a state, variable assignment. The assignment is performed if the event has occurred and the condition is true).

The semantics of edge labels can only be explained in the context of the semantics of StateMate [Drusinsky and Harel, 1989], a commercial implementation of StateCharts. StateMate assumes a step-based execution of StateMate-descriptions. Each step consists of three phases:

- 1 In the first phase, the impact of external changes on conditions and events is evaluated. This includes the evaluation of functions which depend on external events. This phase does not include any state changes. In our simple examples, this phase is not actually needed.
- 2 The next phase is to calculate the set of transitions that should be made in the current step. Variable assignments are evaluated, but the new values are only assigned to temporary variables.
- 3 In the third phase, state transitions become effective and variables obtain their new values.

The separation into phases 2 and 3 is especially important in order to guarantee a reproducible behavior of StateMate models. Consider the StateMate model of fig. 2.22.




Figure 2.22. Mutually dependent assignments

In the second phase, new values for a and b are stored in temporary variables, say a' and b' . In the final phase, temporary variables are copied into the user-defined variables:

phase 2: $a' := b; b' := a;$

phase 3: $a := a'; b := b'$

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

As a result, the values of the two variables will be swapped each time an event e happens. This behavior corresponds to that of two cross-coupled registers (one for each variable) connected to the same clock (see fig. 2.23) and reflects the operation of a clocked finite state machine including those two registers⁷.




Figure 2.23. Cross-coupled D-type registers

Without the separation into phases, the same value would be assigned to both variables. The result would depend on the sequence in which the assignments were performed. The separation into (at least) two phases is quite typical for languages that try to reflect the operation of synchronous hardware. We will find the same separation in VHDL (see page 94).

Due to the separation, the results do not depend on the order in which parts of the model are executed by the simulation. This property is extremely important. Otherwise, there could be simulation runs generating different results, all of which would be considered correct. This could be very confusing in all design procedures. This is not what we expect from the simulation of a real circuit with a fixed behavior.

There are different names for this property:

- Kahn [Kahn, 1974] calls this property **determinate**.
- In other papers, this property is called **deterministic**. However, this term is employed with different meanings:
 - This term is used to denote non-deterministic finite state machines, FSMs which can be in several states at the same time [Hopcroft et al., 2006].
 - Languages may have non-deterministic operators. For these operators, different behaviors are legal implementations.
 - Many authors consider systems to be non-deterministic if their behavior depends on some input not known before run-time.

⁷We adopt IEEE standard schematic symbols [IEEE, 1991] for gates and registers for all the schematics in this book. The symbols in fig. 2.23 denote clocked D-type registers. We continue denoting multiplexers, arithmetic units and memories by shape symbols, due to their widespread use in technical documentation.

- In the sense Kahn uses the term “determinate”.

In this book, we prefer to reduce possible confusion by following Kahn⁸. Note that StateMate models can be determinate only if there are no other reasons for an undefined behavior. For example, conflicts between transitions may be allowed (see fig. 2.24).




Figure 2.24. Conflicting StateCharts transitions

Consider fig. 2.24 (a). If event *A* takes place while the system is in the left state, we will must figure out, which transition will take place. If these conflicts would be resolved arbitrarily, then we would have a non-determinate behavior. Typically, priorities are defined such that this type of a conflict is eliminated.

Now, consider fig. 2.24 (b). There will be a conflict for $x=15$. Such conflicts are difficult to detect. Achieving a determinate behavior requires the absence of conflicts that are resolved in an arbitrary manner.

Note that there may be cases in which we would like to describe non-determinate behavior (e.g. if we have a choice to read from two inputs). In such a case, we would typically like to explicitly indicate that this choice can be taken at run-time (see the `select` statement of ADA on page 102).

Implementations of hierarchical state charts other than StateMate typically do not exhibit determinate behavior. These implementations correspond to a software-oriented view onto hierarchical state charts. In such implementations, choices are usually not explicitly described.

The three phases are assumed to be executed for each **step**. Steps are assumed to be executed each time events or variables have changed. The execution of a StateCharts model consists of the execution of a sequence of steps (see fig. 2.25), each step consisting of three phases.

The set of all values of variables, together with the set of events generated (and the current time) is defined as the **status**⁹ of a StateMate model. After

⁸In earlier versions of the book, we used the term “deterministic” together with an additional explanation.

⁹We would normally use the term “state” instead of “status”. However, the term “state” has a different meaning in StateMate.




Figure 2.25. Steps during the execution of a StateMate model

executing the third phase, a new status is obtained. The notion of steps allows us to define the semantics of **events** more precisely. Events are generated, as mentioned, either internally or externally. **The visibility of events is limited to the step following the one in which they are generated.** Thus, events behave like single bit values which are stored in permanently enabled registers at one clock transition and have an effect on the values stored at the next clock transition. They do not live forever.

Variables, in contrast, retain their values, until they are reassigned. According to StateMate semantics, new values of variables are visible to all parts of the model from the step following the step in which the assignment was made onwards. That means, StateMate semantics implies that new values of variables are propagated to all parts of a model between two steps. StateMate implicitly assumes a **broadcast mechanism for updates on variables**. This means that StateCharts or StateMate can be implemented easily for shared memory-based platforms but are less appropriate for message passing and distributed systems. These languages essentially assume shared memory-based communication, even though this is not explicitly stated. For distributed systems, it will be very difficult to update all variables between two steps. Due to this broadcast mechanism, StateMate is not an appropriate language for modeling distributed systems.

2.4.2.4 Evaluation and extensions

StateCharts' main application domain is that of local, control-dominated systems. The capability of nesting hierarchies at arbitrary levels, with a free choice of AND- and OR-states, is a key advantage of StateCharts. Another advantage is that the semantics of StateMate is defined at a sufficient level of detail [Drusinsky and Harel, 1989]. Furthermore, there are quite a number of commercial tools based on StateCharts. StateMate [IBM, 2010a] and StateFlow [MathWorks, 2010] are examples of commercial tools based on StateCharts. Many of them are capable of translating StateCharts into equivalent descriptions in C or VHDL (see page 78). From VHDL, hardware can be generated using synthesis tools. Therefore, StateCharts-based tools provide a complete path from StateCharts-based specifications down to hardware. Generated C

programs can be compiled and executed. Hence, a path to software-based realizations also exists.

Unfortunately, the efficiency of the automatic translation is sometimes a concern. For example, we could map sub-states of AND-states to UNIX-processes. This would hardly lead to efficient implementations on small processors. The productivity gain from object-oriented programming is not available in StateCharts, since it is not object-oriented. Furthermore, the broadcast mechanism makes it less appropriate for distributed systems. StateCharts do not comprise program constructs for describing complex computation and cannot describe hardware structures or non-functional behavior.

Commercial implementations of StateCharts typically provide some mechanisms for removing the limitations of the model. For example, C code can be used to represent program constructs and **module charts** of StateMate can represent hardware structures.

StateCharts allows timeouts. There is no straightforward way of specifying other timing requirements.

UML includes a variation of StateCharts and hence allows modeling state machines. In UML, these diagrams are called **state diagrams** in version 1 of UML and **state machine diagrams** from version 2.0 onwards. Unfortunately, the semantics of state machine diagrams in UML is different from StateMate: the three simulation phases are not included.

2.4.3 Synchronous languages

2.4.3.1 Motivation

Describing complex SUDs in terms of state machine diagrams is difficult. Such diagrams cannot express complex computations. Standard programming languages can express complex computations, but the sequence of executing several threads may be unpredictable. In a multi-threaded environment with preemptive scheduling there can be many different interleavings of the different computations. Understanding all possible behaviors of such concurrent systems is difficult. A key reason for this is that, in general, many different execution orders are feasible, i.e. the execution order is not specified. The order of execution may well affect the result. The resulting non-determinate behavior can have a number of negative consequences, such as, for example, problems with verifying a certain design. For distributed systems with independent clocks, determinate behavior is difficult to achieve. However, for non-distributed systems, we can try to avoid the problems of unnecessary non-determinate semantics.

For synchronous languages, finite state machines and programming languages are merged into one model. Synchronous languages can express complex computations, but the underlying execution model is that of finite automata. They describe concurrently operating automata. Determinate behavior is achieved by the following key feature: “... *when automata are composed in parallel, a transition of the product is made of the “simultaneous” transitions of all of them*” [Halbwachs, 1998]. This means: we do not have to consider all the different sequences of state changes of the automata that would be possible if each of them had its own clock. Instead, we can assume the presence of a single global clock. Each clock tick, all inputs are considered, new outputs and states are calculated and then the transitions are made. This requires a fast broadcast mechanism for all parts of the model. This idealistic view of concurrency has the advantage of guaranteeing **determinate behavior**. This is a restriction if compared to the general communicating finite state machines (CFSM) model, in which each FSM can have its own clock. Synchronous languages reflect the principles of operation in synchronous hardware and also the semantics found in control languages such as IEC 60848 [IEC, 2002] and STEP 7 [Siemens, 2010]. See Potop-Butucaru et al. [Potop-Butucaru et al., 2006] for a survey on synchronous languages.

2.4.3.2 Examples of synchronous languages: Esterel, Lustre and SCADE

Guaranteeing a determinate behavior for all language features has been a design goal for the synchronous languages Esterel [Esterel Technologies Inc., 2010] [Boussinot and de Simone, 1991] and Lustre [Halbwachs et al., 1991].

Esterel is a reactive language: when activated with an input event, Esterel models react by producing an output event. Esterel is a synchronous language: all reactions are assumed to be completed in zero time and it is sufficient to analyze the behavior at discrete moments in time. This idealized model avoids all discussions about overlapping time ranges and about events that arrive while the previous reaction has not been completed. Like other concurrent languages, Esterel has a parallelism operator, written `||`. Similar to StateCharts, communication is based on a broadcast mechanism. In contrast to StateCharts, however, communication is instantaneous. Instantaneous in this context means “within the same clock cycle”. This means that all signals generated in a particular clock cycle are also seen by the others parts of the model in the same clock cycle and these other parts, if sensitive to the generated signals, react in the same clock cycle. Several rounds of evaluations may be required until a stable state is reached. The propagation of values during the same macroscopic instant of time corresponds to the generation of a next status for the same moment in time in StateCharts, except that the broadcast is now instantaneous and not delayed

until the next round of evaluations like in StateMate. For more and updated information about Esterel, refer to the Esterel home page [Esterel Technologies Inc., 2010].

Esterel and Lustre use different syntactic techniques to denote CFSMs. Esterel appears as a kind of imperative language, whereas Lustre looks more like a data flow language (see page 58 for a description of data flow). SyncCharts is a graphical version of Esterel. In all three cases, semantics are explained by the closely-related underlying CFSMs. The commercial graphical language SCADE [Esterel Technologies, 2010] combines elements of all three languages. SCADE is used for a number of safety-critical software components, for example by Airbus.

Due to the three simulation phases in StateMate, StateMate has the key attributes of synchronous languages and it is determinate. According to Halbwachs, “*StateMate is almost a synchronous language and the only feature missing in StateMate is the instantaneous broadcast*” [Halbwachs, 2008].

2.4.4 SDL: A case of message passing

2.4.4.1 Features of the language

StateCharts is not appropriate for modeling distributed communicating finite state machines. For distributed systems, message passing is the better communication paradigm. Therefore, we will now present a second example of a language based on communication finite state machines, an example based on asynchronous message passing.

This language is called SDL (specification and description language). SDL was designed for distributed applications. It dates back to the early seventies. Formal semantics have been available since the late eighties. The language was standardized by the ITU (International Telecommunication Union). The first standards document is the Z.100 Recommendation published in 1980, with updates in 1984, 1988, 1992 (SDL-92), 1996 and 1999. Relevant versions of the standard include SDL-88, SDL-92 and SDL-2000 [SDL Forum Society, 2010].

Many users prefer graphical specification languages while others prefer textual ones. SDL pleases both types of users since it provides textual as well as graphical formats. Processes are the basic elements of SDL. Processes represent components modeled as extended finite state machines. Extensions include operations on data. Fig. 2.26 shows the graphical symbols used in the graphical representation of SDL.

As an example, we will consider how the state diagram in fig. 2.27 can be represented in SDL. Fig. 2.27 is the same as fig. 2.15, except that output has been

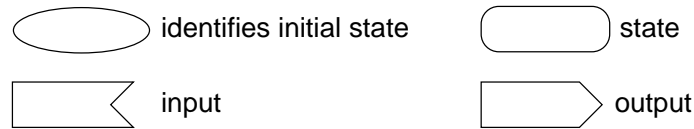


Figure 2.26. Symbols used in the graphical form of SDL




Figure 2.27. FSM to be described in SDL

added, state Z has been deleted, and the effect of signal k has been changed. Fig. 2.28 contains the corresponding graphical SDL representation. Obviously, the representation is equivalent to the state diagram of fig. 2.27.

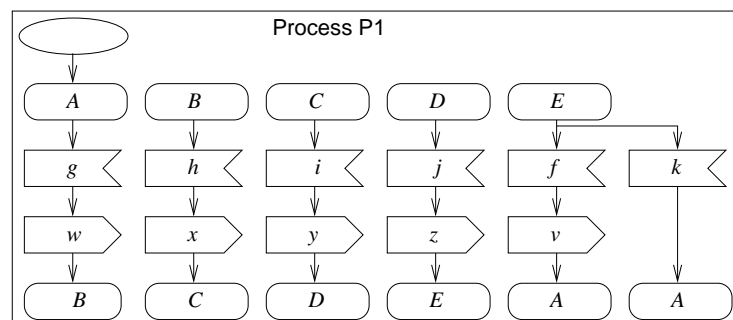


Figure 2.28. SDL-representation of fig. 2.27

As an extension to FSMs, SDL processes can perform operations on data. Variables can be declared locally for processes. Their type can either be pre-defined or defined in the SDL description itself. SDL supports abstract data types (ADTs). The syntax for declarations and operations is similar to that in other languages. Fig. 2.29 shows how declarations, assignments and decisions can be represented in SDL.

SDL also contains programming language elements such as procedures. Procedure calls can also be represented graphically. Object-oriented features became available with version SDL-1992 of the language and were extended with SDL-2000.




Figure 2.29. Declarations, assignments and decisions in SDL

Extended FSMs are just the basic elements of SDL descriptions. In general, SDL descriptions will consist of a set of interacting processes, or FSMs. Processes can send signals to other processes. Semantics of interprocess communication in SDL is based on asynchronous message passing and conceptually implemented through *first-in first-out (FIFO)-queues* associated with processes. There is exactly one queue per process. Signals sent to a particular process will be placed into the corresponding FIFO-queue (see fig. 2.30).




Figure 2.30. SDL interprocess communication

Each process is assumed to fetch the next available entry from the FIFO queue and check whether it matches one of the inputs described for the current state. If it does, the corresponding state transition takes place and output is generated. The entry from the FIFO-queue is ignored if it does not match any of the listed inputs (unless the so-called SAVE-mechanism is used). FIFO-queues are conceptually thought of as being of infinite length. This means: in the description of the semantics of SDL models, FIFO-overflow is never considered. In actual systems, however, infinite FIFO-queues cannot be implemented. They must be of finite length. This is one of the problems of SDL: in order to derive realiza-

tions from specifications, safe upper bounds on the length of the FIFO-queues must be proven.

Process interaction diagrams can be used for visualizing which of the processes are communicating with each other. Process interaction diagrams include **channels** used for sending and receiving signals. In the case of SDL, the term “signal” denotes inputs and outputs of modeled automata. Process interaction diagrams are special cases of **block diagrams** (see below).

Example: Fig. 2.31 shows a process interaction diagram B1 with channels Sw1 and Sw2. Brackets include the names of signals propagated along a certain channel.




Figure 2.31. Process interaction diagram

There are three ways of indicating the recipient of signals:

- 1 **Through process identifiers:** by using identifiers of recipient processes in the graphical output symbol (see fig. 2.32 (left)).



Figure 2.32. Describing signal recipients

Actually, the number of processes does not even need to be fixed at compile time, since processes can be generated dynamically at run-time. OFF-SPRING represents identifiers of child processes generated dynamically by a process.

- 2 **Explicitly:** by indicating the channel name (see fig. 2.32 (right)). Sw1 is the name of a channel.
- 3 **Implicitly:** if signal names imply the channel names, those channels are used. Example: for fig. 2.31, signal B will implicitly always be communicated via channel Sw1.

No process can be defined within any other (processes cannot be nested). However, they can be grouped hierarchically into so-called **blocks**. Blocks at the highest hierarchy level are called **systems**. **Process interaction diagrams** are one level above the leaves of the hierarchical description. B1 can be used within intermediate level blocks (such as within B in fig. 2.33).




Figure 2.33. SDL block

At the highest level in the hierarchy, we have the system (see fig. 2.34). A system will not have any channels at its boundary if the environment is also modeled as a block.




Figure 2.34. SDL system

Fig. 2.35 shows the hierarchy modeled by block diagrams 2.31, 2.33 and 2.34. Process interaction diagrams are next to the *leaves* of the hierarchical description, while system descriptions represent their *root*. Some of the restrictions of modeling hierarchy are removed in version SDL-2000 of the language. With SDL-2000, the descriptive power of blocks and processes is harmonized and replaced by a general *agent* concept.




Figure 2.35. SDL hierarchy

In order to support the modeling of time, SDL includes **timers**. Timers can be declared locally for processes. They can be set and reset using SET and RESET primitives, respectively. Fig. 2.36 shows the use of a timer T. The diagram corresponds to that of fig. 2.28, with the exceptions that timer T is set to the current time plus p during the transition from state D to E. For the transition from E to A we now have a timeout of p time units. If these time units have elapsed before signal f has arrived, a transition to state A is taken without generating output signal v.

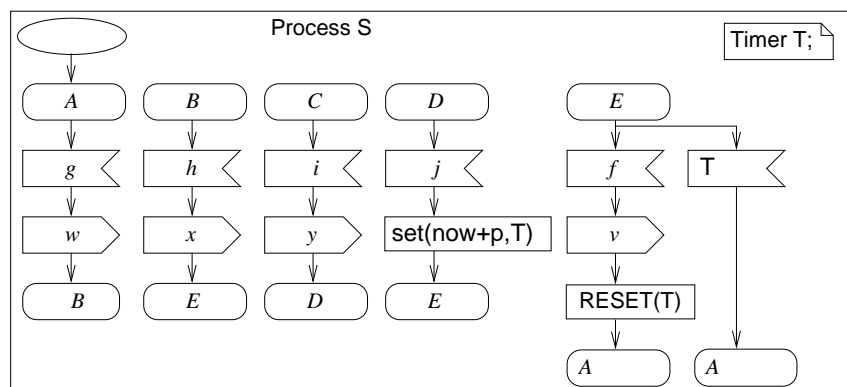


Figure 2.36. Using timer T

SDL can be used, for example, to describe protocol stacks found in computer networks. Fig. 2.37 shows three processors connected through a router. Communication between processors and the router is based on FIFOs.




Figure 2.37. Small computer network described in SDL

The processors as well as the router implement layered protocols (see fig. 2.38).

Each layer describes communication at a more abstract level. The behavior of each layer is typically modeled as a finite state machine. The detailed description of these FSMs depends on the network protocol and can be quite complex. Typically, this behavior includes checking and handling of error conditions, and sorting and forwarding of information packages.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

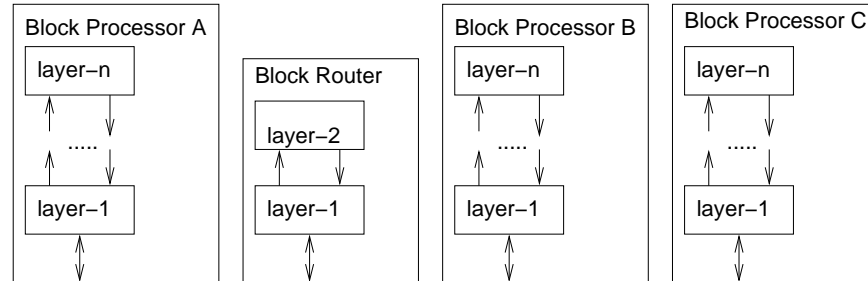


Figure 2.38. Protocol stacks represented in SDL

Available tools for SDL include interfaces to UML (see page 111), and SCs (see page 34). A comprehensive list of tools is available from the SDL forum [SDL Forum Society, 2009].

Estelle is another language which was designed to describe communication protocols. Similar to SDL, Estelle assumes communication via channels and FIFO-buffers. Attempts to unify Estelle and SDL failed.

2.4.4.2 Evaluation of SDL

SDL is excellent for distributed applications and has been used, for example, for specifying ISDN. SDL is not necessarily determinate (the order, in which signals arriving at some FIFO at the same time are processed, is not specified). Reliable implementations require the knowledge of an upper bound on the length of the FIFOs. This upper bound may be difficult to compute. The timer concept is sufficient for soft deadlines, but not for hard ones. Hierarchies are not supported in the same way as in StateCharts. There is no full programming support (but recent revisions of the standard have started to change this) and no description of non-functional properties.

2.5 Data flow

2.5.1 Scope

Data flow is a very “natural” way of describing real life applications. Data flow models reflect the way in which data flows from component to component [Edwards, 2001]. Each component transforms the data in one way or the other. The following is a possible definition of data flow [Wikipedia, 2010]:

Definition: Data flow modeling “*is the process of identifying, modeling and documenting how data moves around an information system. Data flow modeling examines processes (activities that transform data from one form to an-*

other), data stores (the holding areas for data), external entities (what sends data into a system or receives data from a system), data flows (routes by which data can flow)”.

A **data flow program** is specified by a directed graph where the nodes (vertices), called **actors**, represent computations and the arcs represent communication channels. The computation performed by each actor is assumed to be functional, that is, based on the input values only. Each process in a data flow graph is decomposed into a sequence of firings, which are atomic actions. Each firing produces and consumes tokens.

For example, fig. 2.39 describes the flow of data in a video-on-demand system [Ko and Koo, 1996].




Figure 2.39. Video-on-demand system

For unrestricted data flow, it is difficult to prove requested system properties. Therefore, restricted models are commonly used.

2.5.2 Kahn process networks

Kahn process networks (KPN) [Kahn, 1974] are a special case of data flow models. Like other data flow models, KPNs consist of nodes and edges. Nodes correspond to computations performed by some program or task. KPN graphs, like all data flow graphs, show computations to be performed and their dependence, but not the order in which the computations must be performed (in contrast to specifications in von-Neumann languages such as C). Edges imply communication via channels containing potentially infinite FIFOs. Computation times and communication times may vary, but communication is guaranteed to happen within a finite amount of time. Writes are never non-blocking, since the FIFOs are assumed to be as large as needed. Reads must specify a

single channel to be read from. A node cannot check whether data is available before attempting a read. A process cannot wait for data for more than one port at a time. Read operations block whenever an attempt is made to read from an empty FIFO queue. Only a single process is allowed to read from a certain queue and only a single process is allowed to write into a queue. So, if output data has to be sent to more than a single process, duplication of data must be done inside processes. There is no other way for communication between processes except through FIFO-queues.

In the following example, p1 and p2 are incrementing and decrementing the value received from the partner:

```

process p1(in int u, out int v){
    int i;
    i = 0;
    for (;;) {
        send(i,v);
        i = wait(u);
        i = i-1;
    }
}

process p2(in int v, out int u){
    int i;
    for (;;) {
        i = wait(v);
        i = i+1;
        send(i,u);
    }
}

```

Fig. 2.40 shows a graphical representation of this KPN.




Figure 2.40. Graphical representation of KPN

Obviously, we do not really need the FIFOs in this example, since messages cannot accumulate in the channels. This example and other examples can be simulated with the *levi* simulation software [Sirocic and Marwedel, 2007b].

The restrictions are resulting in the **key beauty of KPNs**: the order in which a node is reading data from its channels is fixed by the sequence of read operations and does not depend on the order in which producers are transmitting data over the channels. This means that the sequence of operations is independent of the speed of the nodes producing data. **For a given set of input data, KPNs will always generate the same results, independently of the speed of the nodes.** This property is important, for example, for simulations: it does not matter how fast we are simulating the KPN, the result will always be the same. In particular, the result does not depend on using hardware accelerators for some of the nodes and a distributed execution will give the same result as a centralized one. This property has been called “determinate” and we are following this use. SDL-like conflicts at FIFOs do not exist. Due to this nice property, KPNs are frequently used as an internal representation within a design flow.

Sometimes, KPNs are extended with a “merge”-operator (corresponding to ADA’s `select` statement, see page 102). This operation allows for queuing reads with a list of channels at the same time and waiting for channels to generate data. Such an operator introduces a non-determinate behavior: the order of processing inputs is not specified if both inputs arrive at the same time. This extension is useful in practice, but it destroys the key beauty of KPNs.

In general, Kahn processes require scheduling at run-time, since it is difficult to predict their precise behavior over time. These problems result from the fact that we do not make any assumptions regarding the speed of the channels and the nodes. The question of whether or not finite-length FIFOs are sufficient for an actual KPN model is undecidable in the general case. Nevertheless, execution times are actually unknown during early design phases and therefore this model is very adequate. Useful scheduling algorithms exist [Kienhuis et al., 2000]. For KPNs, the number of processes is fixed, i.e. it does not change at run-time.

2.5.3 Synchronous data flow

Scheduling becomes significantly easier and questions regarding buffer sizes can decidably be answered if we impose restrictions on the timing of nodes and channels. Synchronous data flow (SDF) [Lee and Messerschmitt, 1987] is such a model.

SDF can best be introduced by referring to its graphical notation. Fig. 2.41 (left) shows a synchronous data flow graph. The graph is a directed graph,

nodes A and B denote computations $*$ and $+$. Inputs to SDF graphs are assumed to consist of an infinite stream of samples. Nodes can start their computations when their inputs are available. Edges must be used whenever there is a data dependency between any two nodes.




Figure 2.41. Graphical representations of synchronous data flow

For each execution, the computation in a node is called a firing. For each firing, a number of tokens, representing data, is consumed and produced. In synchronous data flow, the number of tokens produced or consumed in one firing is constant. Constant edge labels denote the corresponding numbers of tokens. These constants facilitate the modeling of **multi-rate** signal processing applications, applications for which certain signals are generated at frequencies that are multiples of other frequencies. For example, in a TV set, some computations might be performed at a rate of 100 Hz while others are performed at a rate of 50 Hz. In general, the number of tokens sent to an edge must be equal to the number of tokens consumed. Let n_s be the number of tokens produced by some sender per firing, and let f_s be the corresponding rate. Let n_r be the corresponding number of tokens consumed per firing at the receiver, and let f_r be the corresponding rate. Then, we must have

$$n_s * f_s = n_r * f_r \quad (2.1)$$

This situation is also visualized in fig. 2.42. The FIFO is needed for buffering if $n_s \neq n_r$. In contrast to Kahn process networks, the size can be computed easily.



Figure 2.42. Multi-rate SDF model

The term **synchronous** data flow reflects the fact that tokens are consumed from the incoming arcs in a synchronous manner (all at the same instant in time). The term **asynchronous** message passing reflects the fact that tokens

can be buffered using FIFOs. The property of producing and consuming a constant number of tokens makes it possible to determine execution order and memory requirements at compile time. Hence, complex run-time scheduling of executions is avoided. SDF graphs may include delays, denoted by the symbol D on an edge (see fig. 2.41 (right)). SDF graphs can be translated into periodic schedules for mono- as well as for multi-processor systems (see e.g. [Pino and Lee, 1995]). A legal schedule for the simple example of fig. 2.41 would consist of the sequence (A, B) (repeated forever). A sequence (A, A, B) (A executed twice as many times as B) would be illegal, since it would accumulate an infinite number of tokens on the implicit FIFO buffer between A and B.

SDF is very useful, for example, in modeling multimedia systems. In this case, each token would correspond to audio or video information, such as an audio sample or a video frame. The observer pattern, mentioned as a problem for modeling with von-Neumann languages on page 25, can be easily implemented correctly in SDF (see fig. 2.43). There is no risk of deadlocks. However, SDF does not allow adding new observers at run-time.




Figure 2.43. Observer pattern in SDF

SDF models are determinate, but they are not appropriate for modeling control flow, such as branches etc. Several extensions and variations of SDF models have been proposed (see, for example Stuijk [Stuijk, 2007]):

- For example, we can have **modes** corresponding to states of an associated finite state machine. For each of the modes, a different SDF graph could be relevant. Certain events could then cause transitions between these modes.
- Homogeneous synchronous data flow (HSDF) graphs are a special case of SDF graphs. For HSDF graphs, the number of token consumed and produced per firing is always 1.
- For cyclo-static data flow (CSDF), the number of tokens produced and consumed per firing can vary over time, but has to be periodic.

Complex SUDs including control flow must be modeled using more general computational graph structures.

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!

2.5.4 Simulink

Computational graph structures are also frequently used in control engineering. For this domain, the Simulink toolbox of MATLAB [The MathWorks Inc., 2010], [Tewari, 2001] is very popular. MATLAB is a modeling and simulation tool based on mathematical models including, for example, partial differential equations. Fig. 2.44 shows an example of a Simulink model [Marian and Ma, 2007].

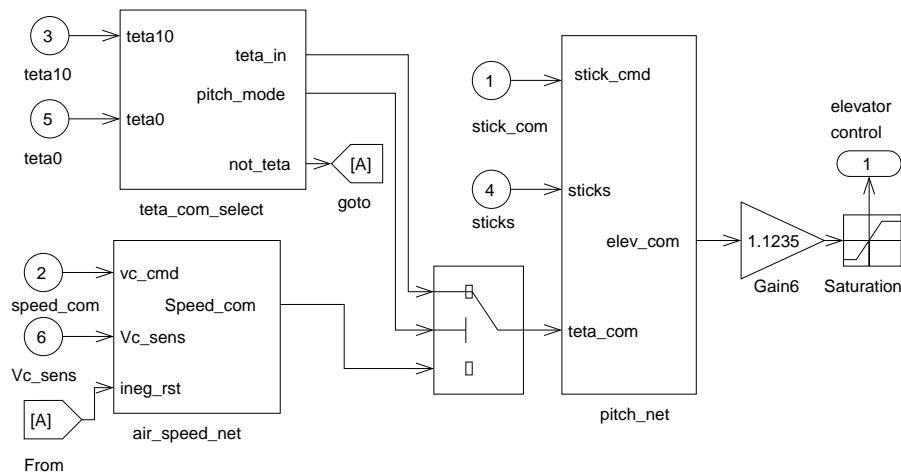


Figure 2.44. Simulink model

The amplifier and the saturation component on the right demonstrate the inclusion of analog modeling. In the general case, the “schematic” could contain symbols denoting analog components such as integrators, differentiators. The switch in the center indicates that Simulink also allows some control flow modeling.

The graphical representation is intuitive and allows control engineers to focus on the control function, without caring about the code necessary to implement the function. The graphical symbols suggest that analog circuits are used as traditional components in control designs. A key goal is to synthesize software from such models. This approach is typically associated with the term **model-based design**, but there is no precise definition for this term.

Semantics of Simulink model reflect the simulation on a digital computer and the behavior may be similar to that of analog circuits, but possibly not quite the same. What is actually the semantics of a Simulink model? Marian and Ma [Marian and Ma, 2007] describe the semantics as follows: “*Simulink uses an idealized timing model for block (node) execution and communication. Both happen infinitely fast at exact points in simulated time. Thereafter, simulated*

time is advanced by exact time steps. All values on edges are constant in between time steps. This means that we execute the model time step after time step. For each step, we compute the function of the nodes (in zero time) and propagate the new values to connected inputs. This explanation does not specify the distance between time steps. Also, it does not immediately tell us how to implement the system in software, since even slowly varying outputs may be recomputed frequently.

This approach is appropriate for modeling physical systems such as cars or trains at a high level and then simulating the behavior of these systems. Also, digital signal processing systems can be conveniently modeled with MATLAB and Simulink. In order to generate implementations, MATLAB/Simulink models first must be translated into a language supported by software or hardware design systems, such as C or VHDL.

Components in Simulink models provide a special case of **actors**. We can assume that actors are waiting for input and perform their operation once all required inputs have arrived. SDF is another case of actor-based languages. In **actor-based languages**, there is no need to pass control to these actors, like in von-Neumann languages.

2.6 Petri nets

2.6.1 Introduction

Very comprehensive descriptions of control flow are feasible with computational graphs known as Petri nets. Actually, Petri nets model **only** control and control dependencies. Modeling data as well requires extensions of Petri nets. Petri nets focus on the modeling of causal dependencies.

In 1962, Carl Adam Petri published his method for modeling causal dependencies, which became known as Petri nets. Petri nets do not assume any global synchronization and are therefore especially suited for modeling distributed systems.

Conditions, **events** and a **flow relation** are the key elements of Petri nets. Conditions are either satisfied or not satisfied. Events can happen. The flow relation describes the conditions that must be met before events can happen and it also describes the conditions that become true if events happen.

Graphical notations for Petri nets typically use circles to denote conditions and boxes to denote events. Arrows represent flow relations. Fig. 2.45 shows a first example.

This example describes mutual exclusion for trains at a railroad track that must be used in both directions. A token is used to prevent collisions of trains going




Figure 2.45. Single track railroad segment

into opposite directions. In the Petri net representation, that token is symbolized by a condition in the center of the model. A partially filled circle (a circle containing a second, filled circle) denotes the situation in which the condition is met (this means: the track is available). When a train wants to go to the right (also denoted by a partially filled circle in fig. 2.45), the two conditions that are necessary for the event “train entering track from the left” are met. We call these two conditions **preconditions**. If the preconditions of an event are met, it can happen. As a result of that event happening, the token is no longer available and there is no train waiting to enter the track. Hence, the preconditions are no longer met and the partially filled circles disappear (see fig. 2.46).




Figure 2.46. Using resource “track”

However, there is now a train going on that track from the left to the right and thus the corresponding condition is met (see fig. 2.46). A condition which is met after an event happened is called a **postcondition**. In general, an event can happen only if all its preconditions are true (or met). If it happens, the preconditions are no longer met and the postconditions become valid. Arrows

identify those conditions which are preconditions of an event and those that are postconditions of an event. Continuing with our example, we see that a train leaving the track will return the token to the condition at the center of the model (see fig. 2.47).




Figure 2.47. Freeing resource “track”

If there are two trains competing for the single-track segment (see fig. 2.48), only one of them can enter.




Figure 2.48. Conflict for resource “track”

In such situations, the next transition to be fired is non-deterministically chosen. Analyses of the net must consider all possible firing sequences. For Petri nets, we are intentionally modeling non-determinism.

A key advantage of Petri nets is that they can be the basis for formal proofs about system properties and that there are standardized ways of generating such proofs. In order to enable such proofs, we need a more formal definition of Petri nets. We will consider three classes of Petri nets: condition/event nets, place/transitions nets, and predicate transition nets.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

2.6.2 Condition/event nets

Condition/event nets are the first class of Petri nets that we will define more formally.

Definition: $N = (C, E, F)$ is called a **net**, iff the following holds:

- 1 C and E are disjoint sets.
- 2 $F \subseteq (E \times C) \cup (C \times E)$ is a binary relation, called flow relation.

The set C is called conditions and the set E is called events.

Definition: Let N be a net and let $x \in (C \cup E)$. Then,

- 1 $\bullet x := \{y | yFx, y \in (C \cup E)\}$ is called the **pre-set** of x . If x denotes an event, $\bullet x$ is also called the set of **preconditions** of x .
- 2 $x^\bullet := \{y | xFy, y \in (C \cup E)\}$ is called the **post-set** of x . If x denotes an event, x^\bullet is also called the set of **postconditions** of x .

The terms preconditions and postconditions are preferred if these sets actually denote conditions $\in C$, that is, if $x \in E$.

Definition: Let $(c, e) \in C \times E$.

- 1 (c, e) is called a **loop**, if $cFe \wedge eFc$.
- 2 N is called **pure**, if F does not contain any loops (see fig. 2.49, left).

Definition: A net is called **simple** if no two transitions t_1 and t_2 have the same set of pre- and postconditions.




Figure 2.49. Nets which are not pure (left) and not simple (center and right)

Simple nets with no isolated elements meeting some additional restrictions are called **condition/event nets**. Condition/event nets are a special case of bipartite graphs (graphs with two disjoint sets of nodes). We will not discuss those additional restrictions in detail since we will consider more general classes of nets in the following.

2.6.3 Place/transition nets

For condition/event nets, there is at most one token per condition. For many applications, it is useful to remove this restriction and to allow more tokens per conditions. Nets allowing more than one token per condition are called place/transition nets. Places correspond to what we so far called conditions and transitions correspond to what we so far called events. The number of tokens per place is called a **marking**. Mathematically, a marking is a mapping from the set of places to the set of natural numbers extended by a special symbol ω denoting infinity.

Let \mathbb{N}_0 denote the natural numbers including 0. Then, formally speaking, place/transition nets can be defined as follows:

Definition: (P, T, F, K, W, M_0) is called a place/transition net \iff

- 1 $N = (P, T, F)$ is a net with places $p \in P$, transitions $t \in T$, and flow relation F .
- 2 Mapping $K : P \rightarrow (\mathbb{N}_0 \cup \{\omega\}) \setminus \{0\}$ denotes the capacity of places (ω symbolizes infinite capacity).
- 3 Mapping $W : F \rightarrow (\mathbb{N}_0 \setminus \{0\})$ denotes the weight of graph edges.
- 4 Mapping $M_0 : P \rightarrow \mathbb{N}_0 \cup \{\omega\}$ represents the initial marking of places.

Edge weights affect the number of tokens that are required before transitions can happen and also identify the number of tokens that are generated if a certain transition takes place. Let $M(p)$ denote a current marking of place $p \in P$ and let $M'(p)$ denote a marking after some transition $t \in T$ took place. The weight of edges belonging to preconditions represents the number of tokens that are removed from places in the precondition set. Accordingly, the weight of edges belonging to the postconditions represents the number of tokens that are added to the places in the postcondition set. Formally, marking M' is computed as follows:

$$M'(p) = \begin{cases} M(p) - W(p, t), & \text{if } p \in \bullet t \setminus t \bullet \\ M(p) + W(t, p), & \text{if } p \in t \bullet \setminus \bullet t \\ M(p) - W(p, t) + W(t, p), & \text{if } p \in \bullet t \cap t \bullet \\ M(p) & \text{otherwise} \end{cases}$$

Fig. 2.50 shows an example of how transition t_j affects the current marking.

By default, unlabeled edges are considered to have a weight of 1 and unlabeled places are considered to have unlimited capacity ω .

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!




Figure 2.50. Generation of a new marking

We now need to explain the two conditions that must be met before a transition $t \in T$ can take place:

- for all places p in the precondition set, the number of tokens must at least be equal to the weight of the edge from p to t and
- for all places p in the postcondition set, the capacity must be large enough to accommodate the new tokens which t will generate.

Transitions meeting these two conditions are called **M-activated**. Formally, this can be defined as follows:

Definition: Transition $t \in T$ is said to be M-activated \iff

$$(\forall p \in {}^\bullet t : M(p) \geq W(p, t)) \wedge (\forall p' \in t^\bullet : M(p') + W(t, p') \leq K(p'))$$

Activated transitions can happen, but they do not need to. If several transitions are activated, the sequence in which they happen is not deterministically defined.

The impact of a firing transition t on the number of tokens can be represented conveniently by a vector \underline{t} associated with t . \underline{t} is defined as follows:

$$\underline{t}(p) = \begin{cases} -W(p, t), & \text{if } p \in {}^\bullet t \setminus t^\bullet \\ +W(t, p), & \text{if } p \in t^\bullet \setminus {}^\bullet t \\ -W(p, t) + W(t, p), & \text{if } p \in {}^\bullet t \cap t^\bullet \\ 0 & \text{otherwise} \end{cases}$$

The new number M' of tokens, resulting from the firing of transition t , can be computed for all places p as follows:

$$M'(p) = M(p) + \underline{t}(p)$$

Using “+” to denote vector addition, we can rewrite this equation as follows:

$$M' = M + \underline{t}$$

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

The set of all vectors \underline{t} form an incidence matrix \underline{N} . \underline{N} contains vectors \underline{t} as columns.

$$\underline{N} : P \times T \rightarrow \mathbb{Z}; \quad \forall t \in T : \underline{N}(p, t) = \underline{t}(p)$$

It is possible to formally prove system properties by using matrix \underline{N} . For example, we are able to compute sets of places, for which firing transitions will not change the overall number of tokens [Reisig, 1985]. Such sets are called **place invariants**. Let us initially consider a single transition t_j in order to find such invariants. Let us search for sets $R \subseteq P$ of places such that the total number of tokens does not change if t_j fires. The following must hold for such sets:

$$\sum_{p \in R} \underline{t}_j(p) = 0 \quad (2.2)$$

Fig. 2.51 shows a transition for which the total number of tokens does not change if it fires.




Figure 2.51. Transition with a constant number of tokens

We are now introducing the characteristic vector \underline{c}_R of some set R of places:

$$\underline{c}_R(p) = \begin{cases} 1 & \text{iff } p \in R \\ 0 & \text{iff } p \notin R \end{cases}$$

With this definition, we can rewrite equation 2.2 as:

$$\sum_{p \in R} \underline{t}_j(p) = \sum_{p \in P} \underline{t}_j(p) * \underline{c}_R(p) = \underline{t}_j \cdot \underline{c}_R = 0 \quad (2.3)$$

\cdot denotes the scalar product. Now, we search for sets of places such that firings of **any** transition will not change the total number of tokens. This means that equation 2.3 must hold for all transitions t_j :

$$\underline{t}_1 \cdot \underline{c}_R = 0$$

$$\begin{aligned}
t_2 \cdot c_R &= 0 \\
&\dots \\
t_n \cdot c_R &= 0
\end{aligned} \tag{2.4}$$

Equations 2.4 can be combined into the following equation by using the transposed incidence matrix \underline{N}^T :

$$\underline{N}^T \underline{c}_R = 0 \tag{2.5}$$

Equation 2.5 represents a system of linear, homogeneous equations. Matrix \underline{N} represents edge weights of our Petri nets. We are looking for solution vectors \underline{c}_R for this system of equations. Solutions must be characteristic vectors. Therefore, their components must be 1 or 0 (integer weights can be accepted if we use weighted sums of tokens). This is more complex than solving systems of linear equations with real-valued solution vectors. Nevertheless, it is possible to obtain information by solving equation 2.5. Using this proof technique, we can for example show that we are correctly implementing mutually exclusive access to shared resources.

Let us now consider a larger example: We are again considering the synchronization of trains. In particular, we are trying to model high-speed Thalys trains traveling between Amsterdam, Cologne, Brussels and Paris. Segments of the train run independently from Amsterdam and Cologne to Brussels. There, the segments get connected and then they run to Paris. On the way back from Paris, they get disconnected at Brussels again. We assume that Thalys trains must synchronize with some other train at Paris. The corresponding Petri net is shown in fig. 2.52.

Places 3 and 10 model trains waiting at Cologne and Amsterdam, respectively. Transitions 2 and 9 model trains driving from these cities to Brussels. After their arrival at Brussels, places 2 and 9 contain tokens. Transition 1 denotes connecting the two trains. The cup symbolizes the driver of one of the trains, who will have a break at Brussels while the other driver is continuing on to Paris. Transition 5 models synchronization with other trains at the Gare du Nord station of Paris. These other trains connect Gare du Nord with some other station (we have used Gare de Lyon as an example, even though the situation at Paris is somewhat more complex). Of course, Thalys trains do not use steam engines; they are just easier to visualize than modern high speed trains. Fig. 2.53 shows matrix \underline{N}^T for this example.

For example, line 2 indicates that firing t_2 will increase the number of tokens on p_2 by 1 and decrease the number of tokens on p_3 by 1. Using techniques

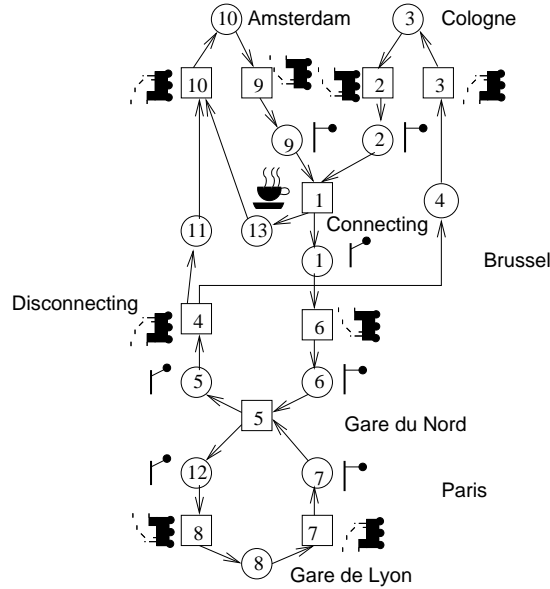


Figure 2.52. Model of Thalys trains running between Amsterdam, Cologne, Brussels, and Paris

	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9	p_{10}	p_{11}	p_{12}	p_{13}
t_1	1	-1							-1				1
t_2		1	-1										
t_3			1	-1									
t_4				1	-1						1		
t_5					1	-1	-1					1	
t_6	-1					1							
t_7							1	-1					
t_8								1				-1	
t_9									1	-1			
t_{10}										1	-1		-1

Figure 2.53. N^T for the Thalys example

from linear algebra, we are able to show that the following four vectors are solutions for this system of linear equations:

$$c_{R,1} = (1, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0, 0, 0, 0)$$

$$c_{R,2} = (1, 0, 0, 0, 1, 1, 0, 0, 0, 1, 1, 1, 0, 0)$$

$$c_{R,3} = (0, 0, 0, 0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 1)$$

$$c_{R,4} = (0, 0, 0, 0, 0, 0, 1, 1, 0, 0, 0, 1, 0, 0)$$

These vectors correspond to the places along the track for trains from Cologne, to the places along the track for trains from Amsterdam, to the places along the path for drivers of trains from Amsterdam, and to the places along the track within Paris, respectively. Therefore, we are able to show that the number of trains and drivers along these tracks is constant (something which we actually expect). This example demonstrates that place invariants provide us with a standardized technique for proving properties about systems.

2.6.4 Predicate/transition nets

Condition/event nets as well as place/transition nets can quickly become very large for large examples. A reduction of the size of the nets is frequently possible with predicate/transition nets. We will demonstrate this, using the so-called “dining philosophers problem” as an example. The problem is based on the assumption that a set of philosophers is dining at a round table. In front of each philosopher, there is a plate containing spaghetti. Between each of the plates, there is just one fork (see fig. 2.54). Each philosopher is either eating or thinking. Eating philosophers need their two adjacent forks for that, so they can only eat if their neighbors are not eating.




Figure 2.54. The dining philosophers problem

This situation can be modeled as a condition/event net, as shown in fig. 2.55. Conditions t_j correspond to the thinking states, conditions e_j correspond to the eating states, and conditions f_j represent available forks.

Considering the small size of the problem, this net is already very large. The size of this net can be reduced by using predicate/transition nets. Fig. 2.56 is a model of the same problem as a predicate/transition net.

With predicate/transition nets, tokens have an identity and can be distinguished from each other¹⁰. We use this in fig. 2.56 in order to distinguish between the three different philosophers p_1 to p_3 and to identify fork f_3 . Furthermore,

¹⁰We could also think of adding a **color** to each of the tokens.

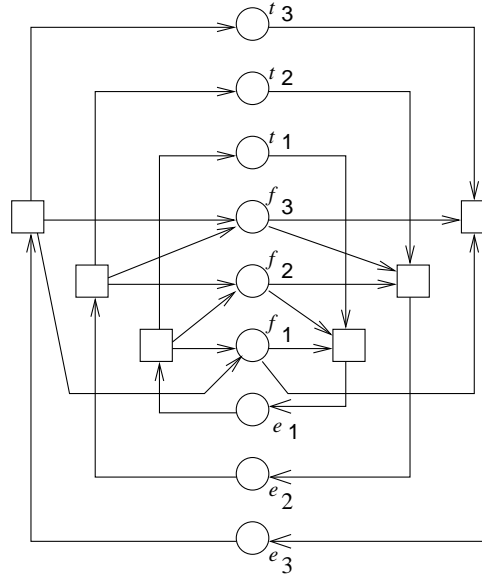


Figure 2.55. Place/transition net model of the dining philosophers problem




Figure 2.56. Predicate/transition net model of the dining philosophers problem

edges can be labeled with variables and functions. In the example, we use variables to represent the identity of philosophers and functions $l(x)$ and $r(x)$ to denote the left and right forks of philosopher x , respectively. These two forks are required as a precondition for transition u and returned as a postcondition by transition v . Note that this model can be easily extended to the case of $n > 3$ philosophers. We just need to add more tokens. In contrast to the net in fig. 2.55, the structure of the net does not have to be changed.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

2.6.5 Evaluation

The key advantage of Petri nets is their power for modeling causal dependencies. Standard Petri nets have no notion of time and all decisions can be taken locally, by just analyzing transitions and their pre- and post-conditions. Therefore, they can be used for modeling geographically distributed systems. Furthermore, there is a strong theoretical foundation for Petri nets, simplifying formal proofs of system properties. Petri nets are not necessarily determinate: different firing sequences can lead to different results. The descriptive power of Petri nets encompasses that of other MoCs, including finite state machines.

In certain contexts, their strength is also their weakness. If time is to be explicitly modeled, standard Petri nets cannot be used. Furthermore, standard Petri nets have no notion of hierarchy and no programming language elements, let alone object oriented features. In general, it is difficult to represent data.

There are extended versions of Petri nets avoiding the mentioned weaknesses. However, there is no universal extended version of Petri nets meeting all requirements mentioned at the beginning of this chapter. Nevertheless, due to the increasing amount of distributed computing, Petri nets became more popular than they were initially.

UML includes extended Petri nets called **activity diagrams**. Extensions include symbols denoting decisions (just like in ordinary flow charts). The placement of symbols is somewhat similar to SDL. Fig. 2.57 shows an example.

The example shows the procedure to be followed during a standardization process. Forks and joins of control correspond to transitions in Petri nets and they use the symbols (horizontal bars) that were initially used for Petri nets as well. The diamond at the bottom shows the symbol used for decisions. Activities can be organized into “swim-lanes” (areas between vertical dotted lines) such that the different responsibilities and the documents exchanged can be visualized. It is interesting to note how a technique like Petri nets was initially certainly not a mainstream technique. Decades after its invention, it has become a frequently applied technique due to its inclusion in UML.

2.7 Discrete event based languages

The discrete event-based model of computation is based on the idea of simulating the generation of events and the processing of events over time. In this model, we are using a queue of future events. These events are sorted by the time at which they should be processed. Semantics is defined by removing the events concerning the current time from the queue, performing the corresponding actions, possibly entering new events into the queue. Time is

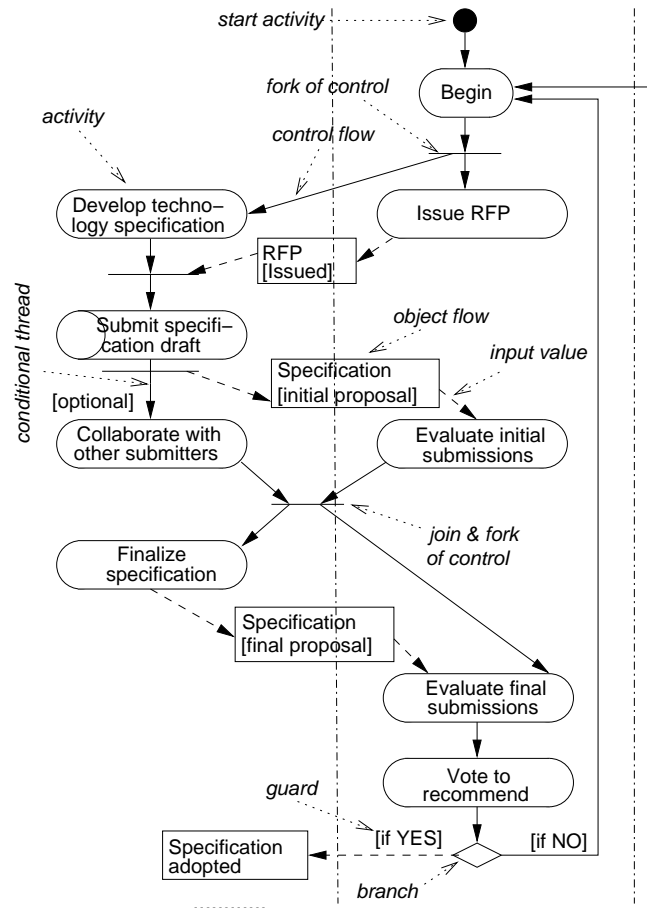


Figure 2.57. Activity diagram [Kobryn, 2001]

advanced whenever no action exists, which should be performed at the current time.

Hardware description languages (HDLs) are designed to model hardware. They are typically based on the discrete event model. We will use HDLs as a prominent example of discrete event modeling. The focus will be on the hardware description language VHDL, and we will briefly cover other HDLs as well.

A key distinction between common software languages and hardware description languages is the need to model time in HDLs. Another distinction comes from the requirement to describe concurrency among different hardware components.

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!

2.7.1 VHDL

2.7.1.1 Introduction

VHDL is a prominent example of HDLs. VHDL uses **processes** for modeling concurrency. Each process models one component of the potentially concurrent hardware. For simple hardware components, a single process may be sufficient. More complex components may need several processes for modeling their operations. Processes communicate through **signals**. Signals roughly correspond to physical connections (wires).

The origin of VHDL can be traced back to the eighties of the last century. At that time, most design systems used graphical HDLs. The most common building block was the gate. However, in addition to using graphical HDLs, we can also use textual HDLs. The strength of textual languages is that they can easily represent complex computations including variables, loops, function parameters and recursion. Accordingly, when digital systems became more complex in the eighties, textual HDLs almost completely replaced graphical HDLs. Textual HDLs were initially a research topic at universities. See Mermet et al. [Mermet et al., 1998] for a survey of languages designed in Europe in the eighties. MIMOLA was one of these languages and the author of this book contributed to its design and applications [Marwedel and Schenk, 1993], [Marwedel, 2008b]. Textual languages became popular when VHDL and its competitor Verilog (see page 96) were introduced.

VHDL was designed in the context of the VHSIC program of the Department of Defense (DoD) in the US. VHSIC stands for *very high speed integrated circuits*¹¹. Initially, the design of VHDL (VHSIC hardware description language) was done by three companies: IBM, Intermetrics and Texas Instruments. A first version of VHDL was published in 1984. Later, VHDL became an IEEE standard, called IEEE 1076. The first IEEE version was standardized in 1987; updates were designed in 1992, in 1997, in 2002 and in 2006 [Lewis et al., 2007]. VHDL-AMS allows modeling analog and mixed-signal systems by including differential equations in the language. The design of VHDL used ADA (see page 100) as the starting point, since both languages were designed for the DoD. Since ADA is based on PASCAL, VHDL has some of the syntactical flavor of PASCAL. However, the syntax of VHDL is much more complex and it is necessary not to get distracted by the syntax. In the current book, we will just focus on some concepts of VHDL which are useful also in other languages. A full description of VHDL is beyond the scope of this book. The standard is available from IEEE (see, for example, [IEEE, 2002]).

¹¹The design of the Internet was also part of the VHSIC program.

2.7.1.2 Entities and architectures

VHDL, like all other HDLs, includes the necessary support for modeling concurrent operation of hardware components. Hardware components are modeled by so-called **design entities** or **VHDL entities**. Entities contain **processes** used to model concurrency. According to the VHDL grammar, design entities are composed of two types of ingredients: an **entity declaration** and one (or several) **architectures** (see fig. 2.58).




Figure 2.58. An entity consists of an entity declaration and architectures

For each entity, the most recently analyzed architecture will be used by default. Using other architectures can be specified. Architectures may contain several processes.

We will discuss a full adder as an example. Full adders have three input ports and two output ports (see fig. 2.59).




Figure 2.59. Full-adder and its interface signals

An entity declaration corresponding to fig. 2.59 is the following:

```

entity full_adder is           -- entity declaration
    port (a, b, carry_in: in Bit;  -- input ports
          sum, carry_out: out Bit); -- output ports
end full_adder;
  
```

Two hyphens (--) are starting a comments. They extend until the the end of the line. Architectures consist of architecture headers and architectural bodies. We can distinguish between different styles of bodies, in particular between structural and behavioral bodies. We will show how the two are different using the full adder as an example. Behavioral bodies include just enough information to compute output signals from input signals and the local state (if any), including the timing behavior of the outputs. The following is an example of this (<= denotes assignments to signals):

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!

```

architecture behavior of full_adder is -- architecture
begin
    sum    <= (a xor b) xor carry_in after 10 ns;
    carry_out <= (a and b) or (a and carry_in) or
                (b and carry_in) after 10 ns;
end behavior;

```

VHDL-based simulators are capable of displaying output signal waveforms resulting from stimuli applied to the inputs of the full adder described above.

In contrast, structural bodies describe the way entities are composed of simpler entities. For example, the full adder can be modeled as an entity consisting of three components (see fig. 2.60). These components are called i1 to i3 and are of type `half_adder` or `or_gate`.




Figure 2.60. Schematic describing structural body of the full adder

In the 1987 version of VHDL, these components must be declared in a so-called component declaration. This declaration is very similar (and it serves the same purpose) as forward declarations in other languages. This declaration provides the necessary information about the component even if the full description of that component is not yet stored in the VHDL database (this may happen in the case of so-called top-down designs). From the 1992 version of VHDL onwards, such declarations are not required if the relevant components are already stored in the component database.

Connections between local component and entity ports are described in **port maps**. The following VHDL code represents the structural body shown in fig. 2.60:

```

architecture structure of full_adder is -- architecture head
component half_adder
    port (in1, in2: in Bit; carry: out Bit;
        sum: out Bit);
end component;

```

```

component or_gate
  port(in1, in2: in Bit; o: out Bit);
end component;

signal x, y, z: Bit;           -- local signals
begin                          -- port map section
i1: half_adder                 -- introduction of half_adder i1
  port map (a, b, x, y); -- connections between ports
i2: half_adder port map (y, carry_in, z, sum);
i3: or_gate port map (x, z, carry_out);
end structure;

```

2.7.1.3 Multi-valued logic and IEEE 1164

In this book, we are restricting ourselves to embedded systems implemented with binary logic. Nevertheless, it may be advisable or necessary to use more than two values for modeling such systems. For example, our systems might contain electrical signals of different strengths and it may be necessary to compute the strength and the logic level resulting from a connection of two or more sources of electrical signals. In the following, we will therefore distinguish between the **level** and the **strength** of a **signal**. While the former is an abstraction of the signal voltage, the latter is an abstraction of the impedance (resistance) of the voltage source. We will be using discrete sets of signal values representing the signal level and the strength. **Using discrete sets of strengths avoids the problems of having to solve Kirchhoff's equations and enables us to avoid analog models used in electrical engineering.** We will also model unknown electrical signals by special signal values.

In practice, electronic design systems use a variety of value sets. Some systems allow only two, while others allow 9 or 46. The overall goal of developing discrete value sets is to avoid the problems of solving network equations (e.g. Kirchhoff's laws) and still model existing systems with sufficient precision. In the following, we will present a systematic technique for building up value sets and for relating these to each other. We will use the strength of electrical signals as the key parameter for distinguishing between various value sets. A systematic way of building up value sets, called CSA-theory, was presented by Hayes [Hayes, 1982]. We will later show how the standard value set used for most cases of VHDL-based modeling can be derived as a special case.

1 signal strength (Two logic values)

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!

In the simplest case, we will start with just two logic values, called '0' and '1'. These two values are considered to be of the same strength. This means: if two wires connect values '0' and '1', we will not know anything about the resulting signal level.

A single signal strength may be sufficient if no two wires carrying values '0' and '1' are connected and no signals of different strength meet at a particular node of electronic circuits.

2 signal strengths (Three and four logic values)

In many circuits, there may be instances in which a certain electrical signal is not actively driven by any output. This may be the case, when a certain wire is not connected to ground, the supply voltage or any circuit node.

For example, systems may contain open-collector outputs (see fig. 2.61, left). If the “pull-down” transistor PD is non-conducting, the output is effectively disconnected. For the tristate outputs (see fig. 2.61, right), an enable signal of '0' will generate a '0' at the outputs of the and-gates (denoted by &), and will make both transistors non-conducting. As a result, output A will be disconnected¹². Hence, using appropriate input signals, such outputs can be effectively disconnected from a wire.

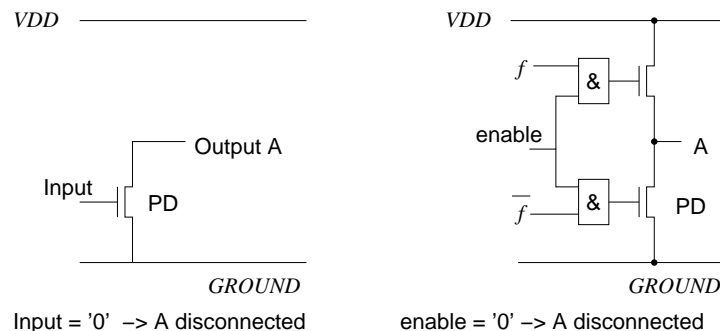


Figure 2.61. Outputs that can be effectively disconnected from a wire

Obviously, the signal strength of disconnected outputs is the smallest strength that we can think of. In particular, the signal strength of Z is smaller than that of '0' and '1'. Furthermore, the signal level of such an output is unknown. This combination of signal strength and signal value is represented by a logic value called 'Z'. If a signal of value 'Z' is connected to another signal, that other signal will always dominate. For example, if two tristate outputs are connected to the

¹²In practice, pull-up transistors may be depletion transistors and the tri-state outputs may be inverting.

same bus and if one output contributes a value of 'Z', the resulting value on the bus will always be the value contributed by the second output (see fig. 2.62).




Figure 2.62. Right output dominates bus

In VHDL, each output is associated with a so-called signal **driver**. Computing the value resulting from the contributions of multiple drivers to the same signal is called **resolution** and resulting values are computed by functions called **resolution functions**.

In most cases, three-valued logic sets $\{ '0', '1', 'Z' \}$ are extended by a fourth value called 'X'. 'X' represents an unknown signal level of the same strength as '0' or '1'. More precisely, we are using 'X' to represent unknown values of signals that can be either '0' or '1' or some voltage representing neither '0' nor '1'¹³.

The resolution that is required if multiple drivers get connected can be computed very easily, if we make use of a partial order among the four signal values '0', '1', 'Z', and 'X'. The partial order is depicted in the **Hasse diagram** in fig. 2.63.




Figure 2.63. Partial order for value set $\{ '0', '1', 'Z', 'X' \}$

Edges in this figure reflect the domination of signal values. Edges define a relation $>$. If $a > b$, then a dominates b . '0' and '1' dominate 'Z'. 'X' dominates all other signal values. Based on the relation $>$, we define a relation \geq . $a \geq b$ holds iff $a > b$ or $a = b$.

¹³There are other interpretations of 'X', but the one presented above is the most useful one in our context.

We define an operation sup on two signals, which returns the **supremum** of the two signal values. The supremum c of the two values a and b is the weakest value for which $c \geq a$ and $c \geq b$ holds. For example, $sup('Z', '0') = '0'$, $sup('Z', '1') = '1'$ etc. **The interesting observation is that resolution functions should compute the sup function according to the above definition.**

3 signal strengths (Seven signal values)

In many circuits, two signal strengths are not sufficient. A common case that requires more values is the use of depletion transistors (see fig. 2.64).

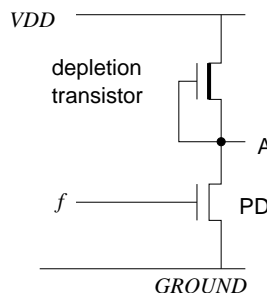


Figure 2.64. Output using depletion transistor

The effect of the depletion transistor is similar to that of a resistor providing a low conductance path to the supply voltage VDD . The depletion transistor as well as the “pull-down transistor” PD act as drivers for node A of the circuit and the signal value at node A can be computed using resolution. The pull-down transistor provides a driver value of '0' or 'Z', depending upon the input to PD. The depletion transistor provides a signal value, which is weaker than '0' and '1'. Its signal level corresponds to the signal level of '1'. We represent the value contributed by the depletion transistor by 'H', and we call it a “weak logic one”. Similarly, there can be weak logic zeros, represented by 'L'. The value resulting from the possible connection between 'H' and 'L' is called a “weak logic undefined”, denoted as 'W'. As a result, we have three signal strengths and seven logic values $\{'0', '1', 'L', 'H', 'W', 'X', 'Z'\}$. Resolution can again be based on a partial order among these seven values. The corresponding partial order is shown in fig. 2.65.

This order also defines an operation sup returning the weakest value at least as strong as the two arguments. For example, $sup('H', '0') = '0'$, $sup('H', 'Z') = 'H'$, $sup('H', 'L') = 'W'$.

'0' and 'L' represent the same signal levels, but a different strength. The same holds for the pairs '1' and 'H'. Devices increasing the signal strength are called **amplifiers**, devices reducing the signal strength are called **attenuators**.




Figure 2.65. Partial order for value set $\{'0', '1', 'L', 'H', 'W', 'X', 'Z'\}$

Ten signal values (4 signal strengths)

In some cases, three signal strengths are not sufficient. For example, there are circuits using charges stored on wires. Such wires are charged to levels corresponding to '0' or '1' during some phases of the operation of the electronic circuit. This stored charge can control the (high impedance) inputs of some transistors. However, if these wires get connected to even the weakest signal source (except 'Z'), they lose their charge and the signal value from that source dominates.

For example, in fig. 2.66, we are driving a bus from a specialized output. The bus has a high capacitive load C . While function f is still '0', we set ϕ to '1', charging capacitor C . Then we set ϕ to '0'. If the real value of function f becomes known and it turns out to be '1', we discharge the bus. The key reason for using pre-charging is that charging a bus using an output such as the one shown in fig. 2.64 is a slow process, since the resistance of depletion transistors is large. Discharging through regular pull-down transistors PD is a much faster process.




Figure 2.66. Pre-charging a bus

In order to model such cases, we need signal values which are weaker than 'H' and 'L', but stronger than 'Z'. We call such values “very weak signal values” and denote them by 'h' and 'l'. The corresponding very weak unknown value is denoted by 'w'. As a result, we obtain ten signal values {'0', '1', 'L', 'H', 'l', 'h', 'X', 'W', 'w', 'Z'}. Using the signal strength, we can again define a partial order among these values (see fig. 2.67).




Figure 2.67. Partial order for value set {'0', '1', 'Z', 'X', 'H', 'L', 'W', 'h', 'l', 'w'}

Five signal strengths

So far, we have ignored power supply signals. These are stronger than the strongest signals we have considered so far. Signal value sets taking power supply signals into account have resulted in the definition of 46-valued value sets [Coelho, 1989]. However, such models are not very popular.

IEEE 1164

In VHDL, there is no predefined number of signal values, except for some basic support for two-valued logic. Instead, the used value sets can be defined in VHDL itself and different VHDL models can use different value sets.

However, portability of models would suffer in a very severe manner if this capability of VHDL was applied in this way. In order to simplify exchanging VHDL models, a standard value set was defined and standardized by the IEEE. This standard is called IEEE 1164 and is employed in many system models. IEEE 1164 has nine values: {'0', '1', 'L', 'H', 'X', 'W', 'Z', 'U', '-'}. The first seven values correspond to the seven signal values described above. 'U' denotes an uninitialized value. It is used by simulators for signals that have not been explicitly defined.

'-' denotes the **input don't care**. This value needs some explanation. Frequently, hardware description languages are used for describing Boolean functions. The VHDL **select** statement is a very convenient means for doing that. The **select** statement corresponds to switch and case statements found in other languages and its meaning is different from the **select** statement in ADA (see page 102).

Example: Suppose that we would like to represent the Boolean function

$$f(a,b,c) = a\bar{b} + bc$$

Furthermore, suppose that f should be undefined for the case of $a = b = c = '0'$. A very convenient way of specifying this function would be the following:

```
f <= select a & b & c -- & denotes concatenation
      '1' when "10-" -- corresponds to first term
      '1' when "-11" -- corresponds to second term
      'X' when "000"
```

This way, functions given above could be easily translated into VHDL. Unfortunately, the **select** statement denotes something completely different. Since IEEE 1164 is just one of a large number of possible value sets, it does not include any knowledge about the “meaning” of '-'. Whenever VHDL tools evaluate select statements such as the one above, they check if the selecting expression ($a \& b \& c$ in the case above) is equal to the values in the when clauses. In particular, they check if e.g. $a \& b \& c$ is equal to "10-". In this context, '-' behaves like any other value: VHDL systems check if c has a value of '-'. Since '-' is never assigned to any of the variables, these tests will never be true. Therefore, '-' is of limited benefit. The non-availability of convenient input don't care values is the price that one has to pay for the flexibility of defining value sets in VHDL itself¹⁴.

The nice property of the general discussion on pages 81 to 86 is the following: it allows us to immediately draw conclusions about the modeling power of IEEE 1164. The IEEE standard is based on the 7-valued value set described on page 84 and, therefore, is capable of modeling circuits containing depletion transistors. It is, however, not capable of modeling charge storage¹⁵.

¹⁴This problem was corrected in VHDL 2006 [Lewis et al., 2007].

¹⁵As an exception, if the capability of modeling depletion transistors or pull-up resistors is not needed, one could interpret weak values as stored charges. This is, however, not very practical since pull-up resistors are found in most actual systems.

2.7.1.4 VHDL processes and assignments

VHDL treats components described above as processes. The syntax used above is just a shorthand for processes. The general syntax for processes is as follows:

```

label : -- optional
process
    declarations -- optional
begin
    statements -- optional
end process ;

```

Assignments are special cases of statements. In VHDL, there are two kinds of assignments:

- **Variable assignments:** The syntax of variable assignments is

```
variable := expression
```

Whenever control reaches such an assignment, the expression is computed and assigned to the variable. Such assignments behave like assignments in common programming languages.

- **Signal assignments:** Signals and signal assignments are introduced in an attempt to model electrical signals in real hardware systems. Signals associate values with instances in time. In VHDL, such a mapping from time to values is represented by **waveforms**. Waveforms are computed from signal assignments. The syntax of signal assignments is

```

signal <= expression;
signal <= transport expression after delay;
signal <= expression after delay;
signal <= reject time inertial expression after delay;

```

Whenever control reaches such an assignment, the expression is computed and used to extend predicted future values of the waveform. In order to compute future values, **simulators are assumed to include a queue of events to happen later than the current simulated time**. This queue is sorted by the time, at which future events (e.g. updates of signals) should happen. Executing a signal assignment results in the creation of entries in this queue. Each entry contains a time for executing the event, the affected signal and the value to be assigned. For signal assignments not containing any **after** clause (first syntactical form), the entry will contain the current

simulation time as the time at which this assignment has to be performed. In this case, the change will take place after an infinitesimally small amount of time, called δ -delay (see below). This allows us to update signals without changing macroscopic time.

For signal assignments containing a **transport** prefix (second syntactical form), the update of the signal will be delayed by the specified amount. This form of the assignment is following the so-called **transport delay model**. This model is based on the behavior of simple wires: wires are (as a first order of approximation) delaying signals. Even short pulses propagate along wires. The transport delay model can be used for logic circuits, even though its main application is to model wires. Suppose that we model a simple or-gate using a transport delay signal assignment:

$c \leq \text{transport } a \text{ or } b \text{ after } 10 \text{ ns};$

Such a model would propagate even short pulses (see fig. 2.68).




Figure 2.68. Gate modeled with transport delay

Transport delay signal assignments will delete all entries in the queue corresponding to the time of the computed update or later times (if we first execute an assignment with a rather large delay and then execute an assignment with a smaller delay, then the entry resulting from the first assignment will be deleted).

For signal assignments containing an **after** clause, but no **transport** clause, **inertial delay** is assumed. The inertial delay model reflects the fact that real circuits come with some “inertia”. This means that short spikes will be suppressed. For the third syntactical form of the signal assignment, all signals changes which are shorter than the specified delay are suppressed. For the fourth form, all signal changes which are shorter than the indicated amount are removed from the predicted waveform. Suppose that we model a simple or-gate using inertial delay:

$c \leq a \text{ or } b \text{ after } 10 \text{ ns};$




Figure 2.69. Gate modeled with inertial delay

For such a model, short spikes would be suppressed (see fig. 2.69).

The implementation of inertial delay relies on the removal of entries in the predicted waveform. The subtle rules for removals are not repeated here.

In addition to assignments, processes may contain **wait** statements. Such statements can be used to suspend a process. There are the following kinds of **wait** statements:

- **wait on** *signal list*; suspend until one of the signals in the list changes;
- **wait until** *condition*; suspend until *condition* is met, e.g. $a = '1'$;
- **wait for** *duration*; suspend for a specified period of time;
- **wait**; suspend indefinitely.

As an alternative to explicit **wait** statements, a list of signals can be added to the process header. In that case, the process is activated whenever one of the signals in that list changes its value. Example: The following model of an and-gate will execute its body once and will restart from the beginning every time one of the inputs changes its value:

```
process(x, y) begin
    prod <= x AND y ;
end process;
```

This model is equivalent to

```
process begin
    prod <= x AND y ;
    wait on x,y;
end process;
```


2.7.1.5 The VHDL simulation cycle

According to the original standards document [IEEE, 1997], the execution of a VHDL model is described as follows: “*The execution of a model consists of an **initialization phase** followed by the **repetitive execution of process statements** in the description of that model. Each such repetition is said to be a **simulation cycle**. In each cycle, the values of all signals in the description are computed. If as a result of this computation an event occurs on a given signal, process statements that are sensitive to that signal will resume and will be executed as part of the simulation cycle.*”

The initialization phase takes signal initializations into account and executes each process once. It is described in the standards as follows¹⁶:

“*At the beginning of initialization, the current time, T_c is assumed to be 0 ns. The initialization phase consists of the following steps:*¹⁷

- *The driving value and the effective value of each explicitly declared signal are computed, and the current value of the signal is set to the effective value. This value is assumed to have been the value of the signal for an infinite length of time prior to the start of the simulation. ...*
- *Each ... process in the model is executed until it suspends. ...*
- *The time of the next simulation cycle (which in this case is the first simulation cycle), T_n is calculated according to the rules of step f of the simulation cycle, below.”*

Each simulation cycle starts with setting the current time to the next time at which changes must be considered. This time T_n was either computed during the initialization or during the last execution of the simulation cycle. Simulation terminates when the current time reaches its maximum, $TIME'HIGH$. According to the original document, the simulation cycle is described as follows: “*A simulation cycle consists of the following steps:*

- a) *The current time, T_c is set equal to T_n . Simulation is complete when $T_n = TIME'HIGH$ and there are no active drivers or process resumptions at T_n .*
- b) *Each active explicit signal in the model is updated. (Events may occur as a result.)” ...*

¹⁶We leave out the discussion of implicitly declared signals and so-called postponed processes introduced in the 1997 version of VHDL.

¹⁷In order not to get lost in the amount of details provided by the standard, some of its sections (indicated by “...”) are omitted in the citation.

This phrase from the document refers to the fact that in the cycle preceding the current cycle, new future values for some of the signals have been computed. If T_c corresponds to the time at which these values become valid, they are now assigned. Note that new values of signals are never immediately assigned while executing a simulation cycle. They are not assigned before the next simulation cycle, at the earliest. Signals that change their value generate so-called events which, in-turn, may enable the execution of processes that are sensitive to that signal.

- c) *“For each process P , if P is currently sensitive to a signal S and if an event has occurred on S in this simulation cycle, then P resumes.*
- d) *Each ... process that has resumed in the current simulation cycle is executed until it suspends.*
- e) *The time of the next simulation cycle, T_n is determined by setting it to the earliest of*
 - 1 *TIME’HIGH* (This is the end of simulation time).
 - 2 *The next time at which a driver becomes active* (this is the next instance in time, at which a driver specifies a new value), *or*
 - 3 *The next time at which a process resumes* (this time is determined by **wait for** statements).

If $T_n = T_c$, then the next simulation cycle (if any) will be a delta cycle.”

The iterative nature of simulation cycles is shown in fig. 2.70.

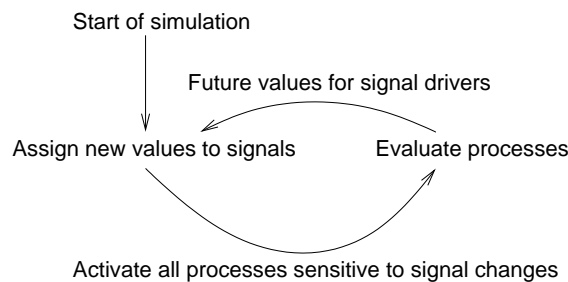


Figure 2.70. VHDL simulation cycles

Delta (δ) simulation cycles have been the source of many discussions. Their purpose is to introduce a infinitesimally small delay even in cases in which the user did not specify any. As an example, we will show the effect of these cycles using a flip-flop as an example. Fig. 2.71 shows the schematic of the flip-flop.

The flip-flop is modeled in VHDL as follows:

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!




Figure 2.71. RS-Flipflop

```

entity RS_Flipflop is
  port (R: in BIT; -- reset
        S: in BIT; -- set
        Q: inout BIT; -- output
        nQ: inout BIT; -- Q-bar
  );
end RS_Flipflop;

architecture one of RS_Flipflop is
  begin
    process: (R,S,Q,nQ)
      begin
        Q <= R nor nQ;
        nQ <= S nor Q;
      end process;
    end one;

```

Ports Q and nQ must be of mode inout since they are also read internally, which would not be possible if they were of mode out. Fig. 2.72 shows the simulation times at which signals are updated for this model. During each cycle, updates are propagated through one of the gates. Simulation terminates after three δ cycles. The last cycle does not change anything, since Q is already 0.

	$< 0ns$	$0ns$	$0ns+\delta$	$0ns+2*\delta$	$0ns+3*\delta$
R	0	1	1	1	1
S	0	0	0	0	0
Q	1	1	0	0	0
nQ	0	0	0	1	1

Figure 2.72. δ cycles for RS-flip-flop

δ cycles correspond to an infinitesimally small unit of time, which will always exist in reality. δ cycles ensure that simulation respects causality.

The results do not depend on the order in which parts of the model are executed by the simulation. This feature is enabled by the separation between the computation of new values for signals and their actual assignment. In a model containing the lines

```
a <= b;
b <= a;
```

signals *a* and *b* will always be swapped. If the assignments were performed immediately, the result would depend on the order in which we execute the assignments (see also page 46). **VHDL models are therefore determinate**. This is what we expect from the simulation of a real circuit with a fixed behavior.

There can be arbitrarily many δ cycles before the current time T_c is advanced. This possibility of infinite loops can be confusing. One of the options of avoiding this possibility would be to disallow zero delays, which we used in our model of the flip-flop.

The propagation of values using signals also allows an easy implementation of the observer pattern (see page 25). In contrast to SDF, the number of observers can vary, depending on the number of processes waiting for changes on a signal.

What is the communication model behind VHDL? The description of the semantics of VHDL relies heavily on a **single, centralized** queue of future events, storing values of all signals in the future. The purpose of this queue is **not** to implement asynchronous message passing. Rather, this queue is supposed to be accessed by the simulation kernel, one entry at a time, in a non-distributed fashion. Attempts to perform distributed VHDL simulations are typically suffering from a poor performance. All modeled components can access values of signals and variables which are in their scope without any message-based communication. Therefore, we tend towards associating VHDL with a shared memory based implementation of the communication. However, FIFO-based message passing could be implemented in VHDL on top of the VHDL simulator as well.

2.7.2 SystemC

Due to the trend of implementing more and more functionality in software, a growing number of embedded systems includes a mixture of hardware and software. Most of the embedded system software is specified in C. For example, embedded systems implement standards such as MPEG 1/2/4 or decoders for mobile phone standards such as GSM or UMTS. The standards are fre-

quently available in the form of “reference implementations”, consisting of C programs not optimized for speed but providing the required functionality. The disadvantage of design methodologies based on VHDL or Verilog is the fact that these standards must be rewritten in order to generate hardware. Furthermore, simulating hardware and software together requires interfacing software and hardware simulators. Typically, this involves a loss of simulation efficiency and inconsistent user interfaces. Also, designers must learn several languages.

Therefore, there has been a search for techniques for representing hardware structures in software languages. Some fundamental problems must be solved before hardware can be modeled with software languages:

- **Concurrency**, as it is found in hardware, has to be modeled in software.
- There has to be a representation for simulation **time**.
- **Multiple-valued logic** and **resolution** as described earlier must be supported.
- The **determinate behavior** of almost all useful hardware circuits must be guaranteed.

SystemCTM [SystemC, 2010], [Open SystemC Initiative, 2005] is a C++ class library designed to solve these problems. With SystemC, specifications can be written in C or C++, making appropriate references to the class libraries.

SystemC comprises a notion of processes executed concurrently. Simulation semantics are similar to VHDL, including the presence of delta cycles. The execution of these processes is controlled via sensitivity lists and calls to wait primitives. The sensitivity list concept of VHDL has been extended to also include dynamic sensitivity lists.

SystemC includes a model of time. Earlier SystemC 1.0 used floating point numbers to denote time. In the current standard, an integer model of time is preferred. SystemC also supports physical units such as picoseconds, nanoseconds, microseconds etc.

SystemC data types include all common hardware types: four-valued logic ('0', '1', 'X' and 'Z') and bitvectors of different lengths are supported. Writing digital signal processing applications is simplified due to the availability of fixed-point data types.

Determinate behavior (see page 47) is not guaranteed in general, unless a certain modeling style is used. Using a command line option, the simulator can be directed to run processes in different orders. This way, the user can check if the simulation results depend on the sequence in which the processes are

executed. However, for models of realistic complexity, only the presence of non-determinate behavior can be proved, not its absence.

Reusing hardware components in different contexts is simplified by the separation of computation and communication. SystemC provides channels, ports and interfaces as abstract components for communication. The introduction of these mechanisms facilitate so-called transaction-level modeling, as defined by Grötter et al. [Grötter et al., 2002]:

Definition: “*Transaction-level modeling (TLM)* is a high-level approach to modeling digital systems where details of communication among modules are separated from the details of the implementation of functional units or of the communication architecture. Communication mechanisms such as buses or FIFOs are modeled as channels, and are presented to modules using SystemC interface classes. Transaction requests take place by calling interface functions of these channel models, which encapsulate low-level details of the information exchange. At the transaction level, the emphasis is more on the functionality of the data transfers - what data are transferred to and from what locations - and less on their actual implementation, that is, on the actual protocol used for data transfer. This approach makes it easier for the system-level designer to experiment, for example, with different bus architectures (all supporting a common abstract interface) without having to recode models that interact with any of the buses, provided these models interact with the bus through the common interface.”

SystemC has the potential for replacing existing VHDL-based design flows. Hardware synthesis starting from SystemC has become available [Herrera et al., 2003a], [Herrera et al., 2003b]. Methodology and applications for SystemC-based design are described in a book on that topic [Müller et al., 2003]. SystemC has been standardized as IEEE standard 1666-2005 [Open SystemC Initiative, 2005].

2.7.3 Verilog and SystemVerilog

Verilog is another hardware description language. Initially it was a proprietary language, but it was later standardized as IEEE standard 1364, with versions called IEEE standard 1364-1995 (Verilog version 1.0) and IEEE standard 1364-2001 (Verilog 2.0). Some features of Verilog are quite similar to VHDL. Just like in VHDL, designs are described as a set of connected design entities, and design entities can be described behaviorally. Also, processes are used to model concurrency of hardware components. Just like in VHDL, bitvectors and time units are supported. There are, however, some areas in which Verilog is less flexible and focuses more on comfortable built-in features. For example, standard Verilog does not include the flexible mechanisms for defin-

ing enumerated types such as the ones defined in the IEEE 1164 standard. However, support for four-valued logic is built into the Verilog language, and the standard IEEE 1364 also provides multiple valued logic with 8 different signal strengths. Multiple-valued logic is more tightly integrated into Verilog than into VHDL. The Verilog logic system also provides more features for transistor-level descriptions. However, VHDL is more flexible. For example, VHDL allows hardware entities to be instantiated in loops. This can be used to generate a structural description for, e.g. n -bit adders without having to specify n adders and their interconnections manually.

Verilog has a similar number of users as VHDL. While VHDL is more popular in Europe, Verilog is more popular in the US.

Verilog versions 3.0 and 3.1 are also known as SystemVerilog. They include numerous extensions to Verilog 2.0. These extensions include [Accellera Inc., 2003], [Sutherland, 2003]:

- additional language elements for modeling behavior,
- C data types such as int and type definition facilities such as typedef and struct,
- definition of interfaces of hardware components as separate entities,
- standardized mechanism for calling C/C++ functions and, to some extent, to call built-in Verilog functions from C,
- significantly enhanced features for describing an environment (called testbench) for the hardware under design (called CUD), and for using the testbench to verify the CUD by simulation,
- classes known from object-oriented programming for use within testbenches,
- dynamic process creation,
- standardized interprocess communication and synchronization, including semaphores,
- automatic memory allocation and deallocation,
- language features that provide a standardized interface to formal verification (see page 201).

Due to the capability of interfacing with C and C++, interfacing to SystemC models is also possible. Improved facilities for simulation- as well as for formal verification-based design validation and the possible interfacing to SystemC will potentially create a very good acceptance. Recently, Verilog and

SystemVerilog have been merged into one standard, IEEE 1800-2009 [IEEE, 2009].

2.7.4 SpecC

The SpecC language [Gajski et al., 2000] is based on the clear separation between communication and computation that should be used for modeling embedded systems. This separation paves the way for re-using components in different contexts and enables *plug-and-play* for system components. SpecC models systems as hierarchical networks of behaviors communicating through channels. SpecC descriptions consist of behaviors, channels and interfaces. Behaviors include ports, locally instantiated components, private variables and functions and a public main function. Channels encapsulate communication. They include variables and functions, which are used for the definition of a communication protocol. Interfaces are linking behaviors and channels together. They declare the communication protocols which are defined in a channel.

SpecC can model hierarchies with nested behaviors. Fig. 2.73 [Gajski et al., 2000] shows a component B including sub-components b1 and b2.




Figure 2.73. Structural hierarchy of SpecC example

The sub-components are communicating through integer c1 and through channel c2. The structural hierarchy includes b1 and b2 as the leaves. b1 and b2 are executed concurrently, denoted by the keyword **par** in SpecC. This structural hierarchy is described in the following SpecC model.

```

interface L { void Write(int x); };
interface R { int Read(void); };
channel C implements L,R
    { int Data; bool Valid;

```

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!


```

void Write(int x) {Data=x; Valid=true;}
int Read (void)
    {while (!Valid) waitfor (10); return (Data);} }
behavior B1(in int p1, L p2, in int p3)
    {void main (void) {/* ...*/ p2.Write(p1);} };
behavior B2 (out int p1, R p2, out int p3)
    {void main(void) {/*...*/ p3=p2.Read(); } };
behavior B(in int p1, out int p2)
    {int c1; C c2; B1 b1(p1, c2, c1); B2 b2(c1, c2, p2);
    void main (void)
        {par {b1.main(); b2.main();}}
    };

```

Note that the interface protocol implemented in channel C, consisting of methods for read and write operations, can be changed without changing behaviors B1 and B2. For example, communication can be bit-serial or parallel and the choice does not affect the models of B1 and B2. This is a necessary feature for IP-reuse.

In order to simplify designs containing software and hardware components, the syntax of SpecC is based on C and C++. In fact, SpecC models are translated into C++ for simulation.

At the specification level, SpecC can model any kind of communication and typically uses message passing. The implementation of simulators is nevertheless typically based on a non-distributed system. The communication model of SpecC has inspired communication in SystemC 2.0.

2.8 Von-Neumann languages

The sequential execution of von-Neumann languages is their common characteristic. Also, such languages allow an almost unrestricted access to global variables. Model-based design using CFSMs and computational graphs is very appropriate for embedded system design. Nevertheless, the use of standard von-Neumann languages is still widespread. Therefore, we cannot ignore these languages.

However, the distinction between KPNs and properly restricted von-Neumann languages is blurring. For KPNs, we do also have sequential execution of the code for each of the nodes. We are still keeping the distinction between KPN and von-Neumann languages since for KPNs, the emphasis of modeling is on

the communication and details of the execution within the nodes are irrelevant. For the first two languages covered in this section, communication is still built into the languages, whereas for the remaining languages, focus is on the computations and communication can be replaced by selecting different libraries.

2.8.1 CSP

CSP (*communicating sequential processes*) [Hoare, 1985] is one of the first languages comprising mechanisms for interprocess communication. Communication is based on channels.

Example:

process A	process B
.....
var a ..	var b ...
a := 3;	...
c!a; -- output to channel c	c?b; -- input from channel c
end ;	end ;

Both processes will wait for the other process to arrive at the input or output statement. This is a case of *rendez-vous*-based or **blocking communication**.

CSP is determinate, since it relies on the commitment to wait for input from a particular channel, like in Kahn process networks.

CSP has laid the foundation for the OCCAM language that was proposed as a programming language of the **transputer** [Thiébaud, 1995]. The focus on communication channels has been picked up again in the design of the XS1 processor [XMOS Ltd., 2010].

2.8.2 ADA

During the eighties of the last century, the Department of Defense (DoD) in the US realized that the dependability and maintainability of the software in its military equipment could soon become a major source of problems, unless some strict policy was enforced. It was decided that all software should be written in the same real-time language. Requirements for such a language were formulated. No existing language met the requirements and, consequently, the design of a new one was started. The language which was finally accepted was based on PASCAL. It was called ADA (after Ada Lovelace, who can be considered being the first (female) programmer). ADA'95 [Kempe, 1995], [Burns and Wellings, 2001] is an object-oriented extension of the original standard.

One of the interesting features of ADA is the ability to have nested declarations of processes (called tasks in ADA). Tasks are started whenever control

passes into the scope in which they are declared. The following is an example (according to Burns et al. [Burns and Wellings, 1990]):

```
procedure example1 is
  task a;
  task b;
  task body a is
    -- local declarations for a
    begin
      -- statements for a
    end a;
  task body b is
    -- local declarations for b
    begin
      -- statements for b
    end b;
  begin
    -- Tasks a and b will start before the 1st statement of example1
  end;
```

The communication concept of ADA is another key concept. It is based on the *rendez-vous* paradigm. Whenever two tasks want to exchange information, the task reaching the “meeting point” first has to wait until its partner has also reached a corresponding point of control. Syntactically, procedures are used for describing communication. Procedures which can be called from other tasks must be identified by the keyword *entry*. Example [Burns and Wellings, 1990]:

```
task screen_out is
  entry call (val : character; x, y : integer);
end screen_out;
```

Task `screen_out` includes a procedure named `call` which can be called from other processes. Some other task can call this procedure by prefixing it with the name of the task:

```
screen_out.call('Z',10,20);
```

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

The calling task has to wait until the called task has reached a point of control, at which it accepts calls from other tasks. This point of control is indicated by the keyword **accept**:

```

task body screen_out is
    ...
    begin
        accept call (val : character; x, y : integer) do
            ...
        end call;
    ...
end screen_out;

```

Obviously, task `screen_out` may be waiting for several calls at the same time. The ADA **select**-statement provides this capability. Example:

```

task screen_output is
    entry call_ch(val:character; x, y: integer);
    entry call_int(z, x, y: integer);
end screen_out;
task body screen_output is
    ...
    select
        accept call_ch ... do...
        end call_ch;
    or
        accept call_int ... do ..
        end call_int;
    end select; ...

```

In this case, task `screen_out` will be waiting until either `call_ch` or `call_int` are called.

Due to the presence of the **select**-statement, ADA is not determinate. ADA has been the preferred language for military equipment produced in the Western hemisphere for some time. Recently produced information about ADA is available from a web sites (see, for example [Kempe Software Capital Enterprises (KSCE), 2010]).

Draft of July 7, 2010; ©P. Marwedel;
 COPYING VERY STRICTLY PROHIBITED!

2.8.3 Java

For Java, communication can be selected by choosing between different libraries. Computation is strictly sequential.

Java was designed as a platform-independent language. It can be executed on any machine for which an interpreter of the internal byte-code representation of Java-programs is available. This byte-code representation is a very compact representation, which requires less memory space than a standard binary machine code representation. Obviously, this is a potential advantage in system-on-a-chip applications, where memory space is limited.

Also, Java was designed as a safe language. Many potentially dangerous features of C or C++ (like pointer arithmetic) are not available in Java. Hence, Java meets the safety requirements for specification languages for embedded systems. Java supports exception handling, simplifying recovery in case of run-time errors. There is no danger of memory leakages due to missing memory deallocation, since Java provides automatic garbage collection. This feature avoids potential problems in applications that must run for months or even years without ever being restarted. Java also meets the requirement to support concurrency since it includes threads (light-weight processes).

In addition, Java applications can be implemented quite fast, since Java supports object orientation and since Java development systems come with powerful libraries.

However, standard Java is not really designed for real-time systems and a number of characteristics which would make it a real-time programming language are missing:

- The size of Java run-time libraries has to be added to the size of the application itself. These run-time libraries can be quite large. Consequently, only really large applications benefit from the compact representation of the application itself.
- For many embedded applications, direct control over I/O devices is necessary (see page 24). For safety reasons, no direct control over I/O devices is available in standard Java.
- Automatic garbage collection requires some computing time. In standard Java, the instance in time at which automatic garbage collection is started cannot be predicted. Hence, the worst case execution time is very difficult to predict. Only extremely conservative estimates can be made.
- Java does not specify the order in which threads are executed if several threads are ready to run. As a result, worst-case execution time estimates must be even more conservative.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

Proposals for solving the problems were made by Nilsen [Nilsen, 1998]. Proposals include hardware-supported garbage-collection, replacement of the run-time scheduler and tagging of some of the memory segments.

Currently (in 2010) relevant Java programming environments include the Java Enterprise Edition (J2EE), the Java Standard Edition (J2SE), the Java Micro Edition (J2ME), and CardJava [Sun, 2010]. CardJava is a stripped-down version of Java with emphasis on security for SmartCard applications. J2ME is the relevant Java environment for all other types of embedded systems. Two library profiles have been defined for J2ME: CDC and CLDC. CLDC is used for mobile phones, using the so-called MIDP 1.0/2.0 as its standard for the application programming interface (API). CDC is used, for example, for TV sets and powerful mobile phones. Currently relevant sources for Java real-time programming include book by Wellings [Wellings, 2004], Dibble [Dibble, 2008] and Bruno [Bruno and Bollella, 2009] as well as web sites [Java Community Process, 2002] and [Anonymous, 2010b].

2.8.4 Pearl and Chill

Pearl [Deutsches Institut für Normung, 1997] was designed for industrial control applications. It does include a large repertoire of language elements for controlling processes and referring to time. It requires an underlying real-time operating system. Pearl has been very popular in Europe and a large number of industrial control projects has been implemented in Pearl. Pearl supports semaphores which can be used to protect communication based on shared buffers.

Chill [Winkler, 2002] was designed for telephone exchange stations. It was standardized by the CCITT and used in telecommunication equipment. Chill is a kind of extended PASCAL.

2.8.5 Communication libraries

Standard von-Neumann languages do not come with built-in communication primitives. However, communication can be provided by libraries. There is a trend towards supporting communication within some local system as well as communication over longer distances. The use of internet protocols is becoming more popular. Libraries will be described in more detail in the section on system software (see page 193).

2.9 Levels of hardware modeling

In practice, designers start design cycles at various levels of abstraction. In some cases, these are high levels describing the overall behavior of the system

to be designed. In other cases, the design process starts with the specification of electrical circuits at lower levels of abstraction. For each of the levels, a variety of languages exists, and some languages cover various levels. In the following, we will describe a set of possible levels. Some lower end levels are presented here for context reasons. Specifications should not start at those levels. The following is a list of frequently used names and attributes of levels:

- **System level models:** The term system level is not clearly defined. It is used here to denote the entire embedded system and the system into which information processing is embedded (“the product”), and possibly also the environment (the physical input to the system, reflecting e.g. the roads, weather conditions etc.). Obviously, such models include mechanical as well as information processing aspects and it may be difficult to find appropriate simulators. Possible solutions include VHDL-AMS (the analog extension to VHDL), SystemC or MATLAB. MATLAB and VHDL-AMS support modeling partial differential equations, which is a key requirement for modeling mechanical systems. It is a challenge to model information processing parts of the system in such a way that the simulation model can also be used for the synthesis of the embedded system. If this is not possible, error-prone manual translations between different models may be needed.

- **Algorithmic level:** At this level, we are simulating the algorithms that we intend to use within the embedded system. For example, we might be simulating MPEG video encoding algorithms in order to evaluate the resulting video quality. For such simulations, no reference is made to processors or instruction sets.

Data types may still allow a higher precision than the final implementation. For example, MPEG standards use double precision floating point numbers. The final embedded system will hardly include such data types. If data types have been selected such that every bit corresponds to exactly one bit in the final implementation, the model is said to be **bit-true**. Translating non-bit-true into bit-true models should be done with tool support (see page 284).

Models at this level may consist of single processes or of sets of cooperating processes.

- **Instruction set level:** In this case, algorithms have already been compiled for the instruction set of the processor(s) to be used. Simulations at this level allow counting the executed number of instructions. There are several variations of the instruction set level:

- In a coarse-grained model, only the effect of the instructions is simulated and their timing is not considered. The information available in assembly reference manuals (instruction set architecture (ISA)) is sufficient for defining such models.
 - **Transaction level modeling:** In transaction level modeling, transactions, such as bus reads and writes, and communication between different components is modeled. Transaction level modeling includes less details than cycle-true modeling (see below), enabling significantly superior simulation speeds [Clouard et al., 2003].
 - In a more fine-grained model, we might have **cycle-true instruction set simulation**. In this case, the exact number of clock cycles required to run an application can be computed. Defining cycle-true models requires a detailed knowledge about processor hardware in order to correctly model, for example, pipeline stalls, resource hazards and memory wait cycles.
- **Register-transfer level (RTL):** At this level, we model all the components at the register-transfer level, including arithmetic/logic units (ALUs), registers, memories, muxes and decoders. Models at this level are always cycle-true. Automatic synthesis from such models is not a major challenge.
 - **Gate-level models:** In this case, models contain gates as the basic components. Gate-level models provide accurate information about signal transition probabilities and can therefore also be used for power estimations. Also delay calculations can be more precise than for the RTL. However, typically no information about the length of wires and hence no information about capacitances is available. Hence, delay and power consumption calculations are still estimates.
- The term “gate-level model” is sometimes also employed in situations in which gates are only used to denote Boolean functions. Gates in such a model do not necessarily represent physical gates; we are only considering the behavior of the gates, not the fact that they also represent physical components. More precisely, such models should be called “Boolean function models”¹⁸, but this term is not frequently used.
- **Switch-level models:** Switch level models use switches (transistors) as their basic components. Switch level models use digital values models (refer to page 81 for a description of possible value sets). In contrast to gate-level models, switch level models are capable of reflecting bidirectional transfer of information.

¹⁸These models could be represented with binary decision diagrams (BDDs) [Wegener, 2000].

- **Circuit-level models:** Circuit theory and its components (current and voltage sources, resistors, capacitances, inductances, and frequently possible macro-models of semiconductors) form the basis of simulations at this level. Simulations involve partial differential equations. These equations are linear if and only if the behavior of semiconductors is linearized (approximated). The most frequently used simulator at this level is SPICE [Vladimirescu, 1987] and its variants.
- **Layout models:** Layout models reflect the actual circuit layout. Such models include **geometric** information. Layout models cannot be simulated directly, since the geometric information does not directly provide information about the behavior. Behavior can be deduced by correlating the layout model with a behavioral description at a higher level or by extracting circuits from the layout, using knowledge about the representation of circuit components at the layout level. In a typical design flow, the length of wires and the corresponding capacitances are extracted from the layout and **back-annotated** to descriptions at higher levels. This way, more precision can be gained for delay and power estimations.
- **Process and device models:** At even lower levels, we can model fabrication processes. Using information from such models, we can compute parameters (gains, capacitances etc) for devices (transistors).

2.10 Comparison of models of computation

2.10.1 Criteria

Models of computation can be compared according to several criteria. For example, Stuijk [Stuijk, 2007] compares MoCs according to the following criteria:

- **Expressiveness** and **succinctness** indicate, which systems can be modeled and how compact they are.
- **Analyzability** relates to the availability of scheduling algorithms and the need for run-time support.
- The **implementation efficiency** is influenced by the required scheduling policy and the code size.

Fig. 2.74 classifies data flow models according to these criteria.

This figure reflects the fact that Kahn process networks are expressive: they are Turing-complete, meaning that any problem which can be computed on a Turing machine can also be computed in a KPN. Turing machines are used




Figure 2.74. Comparison between data flow models

as the standard model of universal computers [Herken, 1995]. However, termination properties and upper bounds on buffer sizes of KPNs are difficult to analyze. SDF graphs, on the other hand, are not Turing-complete. The underlying reason is that they cannot model control flow. However, deadlock properties and upper bounds on buffer sizes of SDF graphs are easier to analyze. Homogeneous SDF (HSDF) graphs (graphs for which all rates are equal to one) are even less expressive, but also easier to analyze.

We could also compare MoCs with respect to the type of processes supported:

- The **number of processes** can be either **static** or **dynamic**. A static number of processes simplifies the implementation and is sufficient if each process models a piece of hardware and if we do not consider “hot-plugging” (dynamically changing the hardware architecture). Otherwise, dynamic process creation (and death) should be supported.
- Processes can either be statically **nested** or all declared at the same level. For example, StateCharts allows nested process declarations while SDL (see page 52) does not. Nesting provides encapsulation of concerns.
- Different techniques for **process creation** exist. Process creation can result from an elaboration of the process declaration in the source code, through the fork and join mechanism (supported for example in Unix), and also through explicit process creation calls.

The expressiveness of different data flow oriented models of computation is also shown in fig. 2.75 [Basten, 2008]. MoCs not discussed in this book are indicated by dashed lines.

Kahn process networks are **Turing complete**. This means that any computation which can be performed by a Turing machine can also be performed by




Figure 2.75. Expressiveness of data flow models

Kahn process networks. Turing machines are the standard reference for evaluating the computability in computer science. In contrast, cyclo-static data flow (CSDF, see page 63) is not Turing complete.

None of the MoCs and languages presented so far meets all the requirements for specification languages for embedded systems. Fig. 2.76 presents an overview over some of the key properties of some of the languages.

Language	Behavioral Hierarchy	Structural Hierarchy	Programming Language Elements	Exceptions Supported	Dynamic Process Creation
StateCharts	+	-	-	+	-
VHDL	+	+	+	-	-
SpecCharts	+	-	+	+	-
SDL	+-	+-	+-	-	+
Petri nets	-	-	-	-	+
Java	+	-	+	+	+
SpecC	+	+	+	+	+
SystemC	+	+	+	+	+
ADA	+	-	+	+	+

Figure 2.76. Language comparison

Interestingly, Spec and SystemC meet all listed requirements. However, some other requirements (like a precise specification of deadlines etc) is not included. It is not very likely that a single MoC or language will ever meet all requirements, since some of the requirements are essentially conflicting. A language supporting hard real-time requirements well may be inconvenient to use for less strict real-time requirements. A language appropriate for distributed control-dominated applications may be poor for local data-flow dominated ap-

plications. Hence, we can expect that we will have to live with compromises and possibly with mixed models.

Which compromises are actually used in practice? In practice, assembly language programming was very common in the early years of embedded systems programming. Programs were small enough to handle the complexity of problems in assembly languages. The next step is the use of C or derivatives of C. Due to the ever increasing complexity of embedded system software (see page xi), higher level languages are to follow the introduction of C. Object oriented languages and SDL are languages which provide the next level of abstraction. Also, languages like UML are required to capture specifications at an early design stage. In practice, these languages can be used like shown in fig. 2.77.

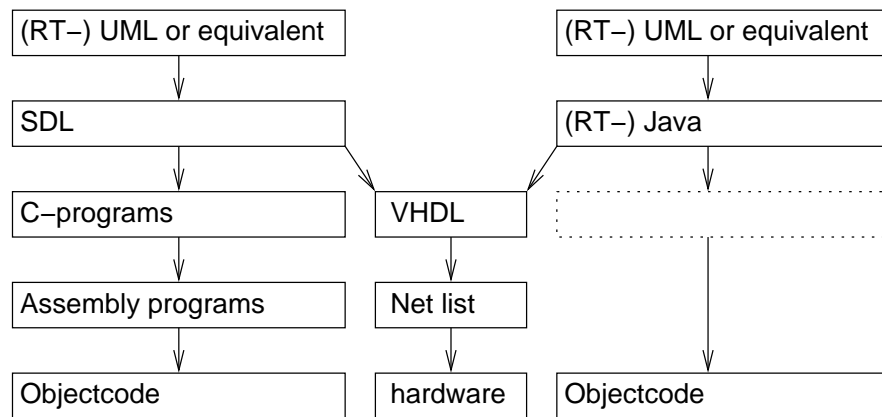


Figure 2.77. Using various languages in combination

According to fig. 2.77, languages like SDL or StateCharts can be translated into C. These C descriptions are then compiled. Starting with SDL or State-Chart also opens the way to implementing the functionality in hardware, if translators from these languages to VHDL are provided. Both C and VHDL will certainly survive as intermediate languages for many years. Java does not need intermediate steps but does also benefit from good translation concepts to assembly languages. In a similar way, translations between various graphs are feasible. For example, SDF graphs can be translated into a subclass of Petri nets [Stuijk, 2007]. Also, they correspond to a subclass of the **computation graph model** proposed by Karp and Miller [Karp and Miller, 1966]. Linking the various models of computation is facilitated by formal techniques [Chen et al., 2007].

Several languages for embedded system design are covered in a book edited by M. Radetzki [Radetzki, 2009]. Popovici et al. [Popovici et al., 2010] use a combination of Simulink and SystemC.

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

2.10.2 UML

UMLTM is a language including diagrams reflecting several MoCs. Fig. 2.78 classifies the UML diagrams mentioned so far with respect to our table of MoCs.

Communication/ Components	Shared memory	Message passing	
		synchronous	asynchronous
Undefined components		Use cases Sequence charts, timing diagrams	
Finite state machines	State diagrams	-	-
Data flow	(not useful)	Data flow diagrams	
Petri nets	(not useful)	Activity charts	
Distributed event model	-	-	
Von-Neumann model	-	-	

Figure 2.78. Models of computation available in UML

This figure shows how UML covers several models of computation, with a focus on early design phases. Semantics of communication is typically imprecisely defined. Therefore, our classification cannot be precise in this respect. In addition to the diagrams already mentioned, the following diagrams can be modeled:

- **Deployment diagrams:** These diagrams are important for embedded systems: they describe the “execution architecture” of systems (hardware or software nodes).
- **Package diagrams:** Package diagrams represent the partitioning of software into software packages. They are similar to module charts in StateMate.
- **Class diagrams:** These diagrams describe inheritance relations of object classes.
- **Communication diagram** (called **Collaboration diagrams** in UML 1.x): These graphs represent classes, relations between classes, and messages that are exchanged between them.
- **Component diagrams:** They represent the components used in applications or systems.
- **Object diagrams, interaction overview diagrams, composite structure diagrams:** This list consists of three types of diagrams which are less frequently used. Some of them may actually be special cases of other types of diagrams.

Available tools provide some consistency checking between the different diagram types. Complete checking, however, seems to be impossible. One reason for this is that the semantics of UML initially was left undefined. It has been argued that this was done intentionally, since one does not like to bother about the precise semantics during the early phases of the design. As a consequence, precise, executable specifications can only be obtained if UML is combined with some other, executable language. Available design tools have combined UML with SDL [IBM, 2009] and C++. There are, however, also some first attempts to define the semantics of UML.

Version 1.4 of UML was not designed for embedded systems. Therefore, it lacks a number of features required for modeling embedded systems (see page 19). In particular, the following features are missing [McLaughlin and Moore, 1998]:

- the partitioning of software into tasks and processes cannot be modeled,
- timing behavior cannot be described at all,
- the presence of essential hardware components cannot be described.

Due to the increasing amount of software in embedded systems, UML is gaining importance for embedded systems as well. Hence, several proposals for UML extensions to support real-time applications have been made [McLaughlin and Moore, 1998], [Douglass, 2000]. These extensions have been considered during the design of UML 2.0. UML 2.0 includes 13 diagram types (up from nine in UML 1.4) [Ambler, 2003]. Special profiles are taking the requirements of real-time systems into account [Martin and Müller, 2005], [Müller, 2007]. Profiles include class diagrams with constraints, icons, diagram symbols, and some (partial) semantics. There are UML profiles for [Müller, 2007]:

- Schedulability, Performance, and Time Specification (SPT) [Object Management Group (OMG), 2005b],
- Testing [Object Management Group (OMG), 2010a],
- Quality of Service (QoS) and Fault Tolerance [Object Management Group (OMG), 2010a],
- a Systems Modeling Language called SysML [Object Management Group (OMG), 2008],
- Modeling and Analysis of Real-Time Embedded Systems (MARTE), [Object Management Group (OMG), 2009]
- UML and SystemC interoperability [Riccobene et al., 2005],

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

- The SPRINT profile for reuse of intellectual property (IP) [Sprint Consortium, 2008].

Using such profiles, we can -for example- attach timing information to sequence charts. However, profiles may be incompatible. Also, UML has been designed for modeling and frequently leaves too many semantical issues open to allow automatic synthesis of implementations [Müller, 2007].

2.10.3 Ptolemy II

The Ptolemy project [Davis et al., 2001] focuses on modeling, simulation, and design of heterogeneous systems. Emphasis is on embedded systems that mix different technologies and, accordingly, also MoCs. For example, analog and digital electronics, hardware and software, and electrical and mechanical devices can be described. Ptolemy supports different types of applications, including signal processing, control applications, sequential decision making, and user interfaces. Special attention is paid to the generation of embedded software. The idea is to generate this software from the MoC which is most appropriate for a certain application. Version 2 of Ptolemy (Ptolemy II) supports the following MoCs and corresponding domains (see also page 31):

- 1 Communicating sequential processes (CSP).
- 2 Continuous time (CT): This model is appropriate for mechanical systems and analog circuits. It is supported through a set of extensible differential equation solvers.
- 3 Discrete event model (DE): this is the model used by many simulators, e.g. VHDL simulators.
- 4 Distributed discrete events (DDE). Discrete event systems are difficult to simulate in parallel, due to the inherent centralized queue of future events. Attempts to distribute this data structure have not been very successful so far. Therefore, this special (experimental) domain is introduced. Semantics can be defined such that distributed simulation becomes more efficient than in the DE model.
- 5 Finite state machines (FSM).
- 6 Process networks (PN), using Kahn process networks (see page 59).
- 7 Synchronous dataflow (SDF).
- 8 Synchronous/reactive (SR) MoC. This model uses discrete time, but signals do not need to have a value at every clock tick. Esterel (see page 51) is a language following this style of modeling.

This list clearly shows the focus on different models of computation in the Ptolemy project.

2.11 Assignments

- 1 Prepare a list of up to 6 requirements for specification languages for embedded systems!
- 2 Suppose the StateCharts in fig. 2.79 model is given.




Figure 2.79. StateCharts example

Also, suppose that we have the following sequence of input events: b c f h g h e a b c. In the diagram in fig. 2.80, mark all the states the StateCharts model will be in after a particular input has been applied! Note that H denotes the history mechanism.

	M	N	P	Q	R	S	T	X	Y	Z
(Reset)							v			
b										
c										
f										
h										
g										
h										
e										
a										
b										
c										

Figure 2.80. States of the StateCharts example

- 3 Are StateCharts determinate models if we follow the StateMate semantics? Please explain your answer!

Draft of July 7, 2010; ©P. Marwedel;
COPYING VERY STRICTLY PROHIBITED!

- 4 Which three types of Petri nets did we discuss in this book?
- 5 One of the types of Petri nets allows several non-distinguishable tokens per place. Which components are used in a mathematical model of such nets?
Hint: $N=(P, \dots\dots\dots)$
- 6 How does a compact model of the dining philosopher's problem look like?
- 7 CSA theory leads to 2, 3 and 4 logic strengths, corresponding to 4, 7 and 10 logic values. How many strengths and values are we using in IEEE 1164? Please show the partial order among the values of IEEE 1164 in a diagram! Which of the values of IEEE 1164 are not included in the partial order and what is the meaning of these values?
- 8 Which of the following circuits can be modeled with IEEE 1164: complementary CMOS outputs, outputs with a depletion transistor, open collector outputs, tristate outputs, pre-charging on buses (if depletion transistors are used as well)?
- 9 Suppose that a bus as shown in fig. 2.81 is given. Rectangles containing an &-sign denote AND-gates.




Figure 2.81. Bus driven by tri-state outputs

Which of the IEEE 1164 values will be on the bus if both enable inputs are set to '0' ($ena1 = ena2 = '0'$)? Which of the IEEE 1164 values will be on the bus if $ena1 = '0'$, $ena2 = '1'$ and $f2 = '1'$?

- 10 Simulate trains between Paris, Brussels, Amsterdam and Cologne, using the levi simulation software [Sirocic and Marwedel, 2007d]! Modify the examples included with the software such that two independent tracks exist between any two stations and demonstrate an (arbitrary) schedule involving 10 trains!

- 11 Simulate a Kahn process network computing Fibonacci numbers, using the levi simulation software [Sirocic and Marwedel, 2007b].
- 12 Which of the following languages are using asynchronous message passing communication: StateCharts, SDL, VHDL, CSP, Petri nets?
- 13 Which of the following languages use a broadcast mechanism for updating variables: StateCharts, SDL, Petri nets?
- 14 Which of the following diagram types are supported by UML: Sequence charts, record charts, Y-charts, use cases, activity diagrams, circuit diagrams?