

Embedded Systems

7



BF - ES

- 1 -

Models of computation for embedded systems

Communication/ local computations	Shared memory	Message passing	
		Synchronous	Asynchronous
Communicating finite state machines	StateCharts, StateFlow		SDL, MSCs
Data flow model C			Kahn process networks, SDF
Computational graphs			Petri nets
Von Neumann model	C, C++, Java	C, C++, Java with libraries CSP, ADA	
Discrete event (DE) model	VHDL, Simulink	Only experimental systems, e.g. distributed DE in Ptolemy	

BF - ES


- 2 -

Place/transition nets

REVIEW

Def.: (P, T, F, K, W, M_0) is called a **place/transition net (P/T net)** iff

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




In the following: assume initial marking is finite, capacity ω .

BF - ES

- 3 -

Unbounded Petri net

REVIEW



BF - ES

- 4 -

Invariants & boundedness

REVIEW

- A net is **covered** by place invariants iff every place is contained in some invariant.
- **Theorem 4:**
 - a) If R is a place invariant and $p \in R$, then p is bounded.
 - b) If a net is covered by place invariants then it is bounded.


BF - ES

- 5 -

Deadlock

REVIEW

- A **dead marking (deadlock)** is a marking where no transition can fire.
- A Petri net is **deadlock-free** if no dead marking is reachable.




BF - ES

- 6 -

Structural properties: deadlock-traps REVIEW

- A place set S is a **(static) deadlock** if every transition that adds token to S also removes token from S .
- A place set S is a **trap** if every transition that removes token from S also adds token to S .




BF - ES

- 7 -

Empty structural deadlocks and marked traps REVIEW

REVIEW



- Empty structural deadlocks are never re-marked;
- Marked traps are never emptied.

BF - ES


- 8 -

Sufficiently marked places

REVIEW

A place is called sufficiently marked if there are enough tokens for one of the outgoing transitions:

- Define $W(p) = \min \{ W(p,t) \mid (p,t) \in F \}$ if there exists a $(p,t) \in F$ and 0 otherwise
- Place p is **sufficiently marked** in marking M , if $M(p) \geq W(p)$
- A set of places is sufficiently marked if it contains a sufficiently marked place.




BF - ES

- 9 -

Deadlock-Trap Property

REVIEW

- A P/T has the **deadlock-trap property**, if every (static) deadlock contains a trap that is sufficiently marked in M_0 .



BF - ES

- 10 -

Deadlock-Trap Property

REVIEW

Theorem 5:

Every homogeneous P/T net with non-blocking weights that has the deadlock-trap property is deadlock-free.

Homogeneous: For each place, all outgoing edges have the same weight.

Non-blocking weights: $W^+(p) \geq W^-(p)$

- $W^-(p) = \min \{ W(p,t) \mid (p,t) \in F \}$ if there exists a $(p,t) \in F$ and 0 otherwise
- $W^+(p) = \min \{ W(t,p) \mid (t,p) \in F \}$ if there exists a $(t,p) \in F$ and 0 otherwise

BF - ES


- 11 -

Fairness

REVIEW

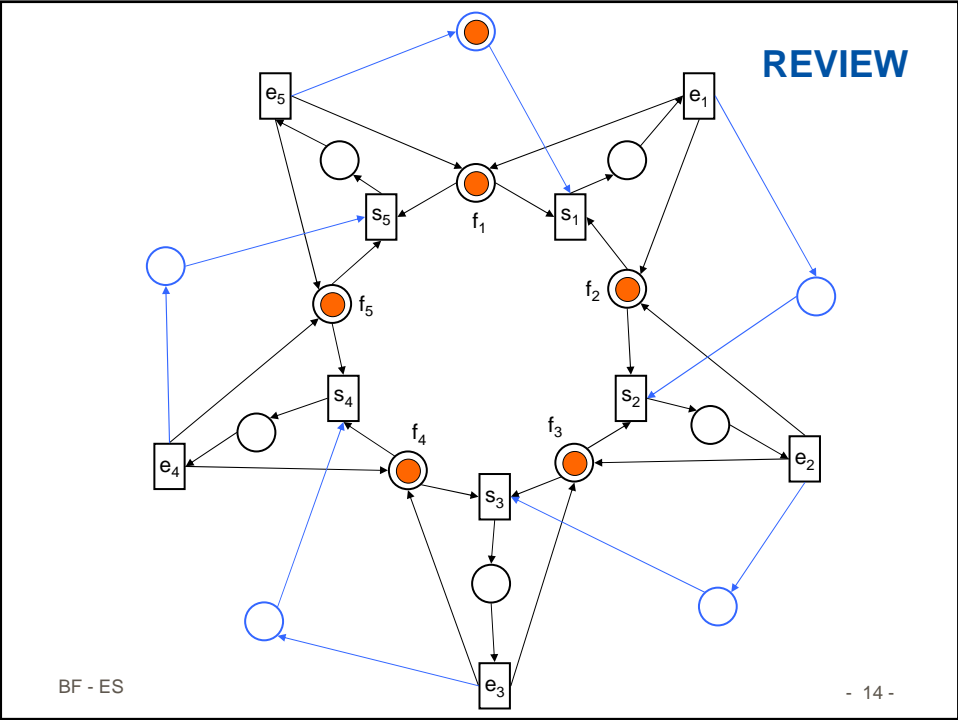

Dining philosophers problem

- $n > 1$ philosophers sitting at a round table;
- n forks,
- n plates with spaghetti;
- philosophers either thinking or eating spaghetti (using left and right fork).
- 2 forks needed!



BF - ES

- 12 -



Fairness

REVIEW

Let N be a Petri net and w an execution of N .

- w is **impartial** with respect to a set of transitions T iff every transition in T occurs infinitely often in w .
 - w is **just** with respect to a set of transitions T iff every transition in T that is enabled in all except finitely many markings occurs infinitely often in w .
 - w is **fair** with respect to a set of transitions T iff every transition in T that is enabled in infinitely many markings occurs infinitely often in w .
- w is **impartial** $\Rightarrow w$ is **fair**
 - w is **fair** $\Rightarrow w$ is **just**

BF - ES

- 15 -

Persistent nets

REVIEW

Theorem 7: If the net is persistent, then every just execution is fair.

BF - ES

- 16 -

State Fairness

REVIEW

Theorem 8: Let N be a bounded net, t a live transition, and w a state-fair execution of N . Then t occurs infinitely often in w .


BF - ES

- 17 -

Extensions: finite capacities

REVIEW

- $K(p)=4$




BF - ES

- 18 -

Extensions: Petri nets with priorities REVIEW

- $t_1 \prec t_2$: t_2 has higher priority than t_1 .



- Petri nets with priorities are Turing-complete.

BF - ES

- 19 -


Predicate/transition model of the dining philosophers problem REVIEW

- Let x be one of the philosophers,
- let $l(x)$ be the left fork of x ,
- let $r(x)$ be the right fork of x .

Token: individuals.

Semantics can be defined by replacing net by equivalent condition/event net.

Model can be extended to arbitrary numbers.



BF - ES

- 20 -

Summary Petri nets

Pros:

- Appropriate for distributed applications,
- Well-known theory for formally proving properties,
- Initially theoretical topic, but now widely adapted in practice due to increasing number of distributed applications.

Cons (for the nets presented) :

- problems with modeling timing,
- no programming elements,
- no hierarchy.

Extensions:

- Enormous amounts of efforts on removing limitations.

Data Flow Models

Data flow modeling

- **Def.:** 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 another),
- *data stores* (the holding areas for data),
- *external entities* (what sends data into a system or receives data from a system, and
- *data flows* (routes by which data can flow).

BF - ES

- 23 -

Data flow as a “natural” model of applications

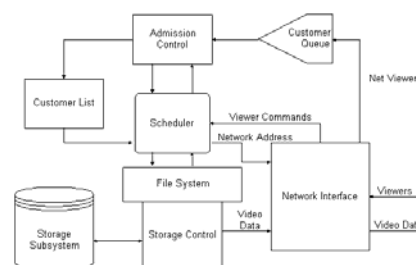
Registering for courses



<http://www.agilemodeling.com/artifacts/dataFlowDiagram.htm>

BF - ES

Video on demand system




www.ece.ubc.ca/~irenek/techpaps/vod/vod.html

- 24 -

Process networks

Many applications can be specified in the form of a set of communicating processes.

Example: system with two sensors:



Alternating read

```
loop
  read_temp; read_humidity
until false;
```

of the two sensors
not the right approach.

BF - ES

- 25 -

The case for multi-process modeling in imperative languages

```
MODULE main;
  TYPE some_channel =
    (temperature, humidity);
  some_sample : RECORD
    value : integer;
    line : some_channel
  END;
  PROCESS get_temperature;
  VAR sample : some_sample;
  BEGIN
  LOOP
    sample.value := new_temperature;
    IF sample.value > 30 THEN ....
    sample.line := temperature;
    to_fifo(sample);
  END
  END get_temperature;
```


```
PROCESS get_humidity;
  VAR sample : some_sample;
  BEGIN
  LOOP
    sample.value := new_humidity;
    sample.line := humidity;
    to_fifo(sample);
  END
  END get_humidity;

BEGIN
  get_temperature; get_humidity;
END;
```

- Blocking calls new_temperature, new_humidity
- **Structure clearer than alternating checks for new values in a single process**

How to model dependencies between tasks/processes?


Dependencies between processes/tasks



BF - ES

- 27 -

Task graphs



Nodes are assumed to be a „program“ described in some programming language, e.g. C or Java.


- **Def.:** A **dependence graph** is a directed graph $G=(V,E)$ in which $E \subseteq V \times V$ is a partial order.
- If $(v1, v2) \in E$, then $v1$ is called an **immediate predecessor** of $v2$ and $v2$ is called an **immediate successor** of $v1$.
- Suppose E^* is the transitive closure of E . If $(v1, v2) \in E^*$, then $v1$ is called a **predecessor** of $v2$ and $v2$ is called a **successor** of $v1$.

BF - ES


- 28 -

Reference model for data flow: Kahn process networks (1974)

For asynchronous message passing:
communication between tasks is buffered




Special case: Kahn process networks:
executable task graphs;
Communication via infinitely large FIFOs



BF - ES

- 29 -



Properties of Kahn process networks (1)

- Each node corresponds to one program/task;
- Communication is only via channels;
- Channels include FIFOs as large as needed;
- Channels transmit information within an unpredictable but finite amount of time;
- Mapping from ≥ 1 input seq. to ≥ 1 output sequence;
- In general, execution times are unknown;
- Send operations are non-blocking, reads are blocking.
- One producer and one consumer;
i.e. there is only one sender per channel;

BF - ES

- 31 -

Properties of Kahn process networks (2)

- There is only one sender per channel.
- A process cannot check whether data is available before attempting a read.
- A process cannot wait for data for more than one port at a time.
- Therefore, the order of reads depends only on data, not on the arrival time.
- Therefore, Kahn process networks are **deterministic** (!); for a given input, the result will always be the same, regardless of the speed of the nodes.


This is the
key beauty
of KPNs!

BF - ES

- 32 -

A Kahn Process

```
process f(in int u, in int v, out int w)
{
  int i; bool b = true;
  for (;;) {
    i = b ? wait(u) : wait(w);
    printf("%i\n", i);
    send(i, w);
    b = !b;
  }
}
```



Process alternately reads from u and v, prints the data value, and writes it to w

Source: Gilles Kahn, The Semantics of a Simple Language for Parallel Programming (1974)

BF - ES

- 33 -

A Kahn Process

```
process f(in int u, in int v, out int w)
{
  int i; bool b = true;
  for (;;) {
    i = b ? wait(u) : wait(w);
    printf("%i\n", i);
    send(i, w);
    b = !b;
  }
}
```

wait() returns the next token in an input FIFO, blocking if it's empty

send() writes a data value on an output FIFO


Source: Gilles Kahn, The Semantics of a Simple Language for Parallel Programming (1974)

BF - ES

- 34 -

A Kahn Process

```
process g(in int u, out int v, out int w)
{
  int i; bool b = true;
  for(;;) {
    i = wait(u);
    if (b) send(i, v); else send(i, w);
    b = !b;
  }
}
```



Process reads from u and alternately copies it to v and w


BF - ES

- 35 -

A Kahn System

- Prints an alternating sequence of 0's and 1's

Emits a 1 then copies input to output



Emits a 0 then copies input to output

BF - ES

- 36 -

Definition: Kahn networks

A **Kahn process network** is a directed graph (V, E) , where


- V is a set of **processes**,
- $E \subseteq V \times V$ is a set of **edges**,
- associated with each edge e is a **domain** D_e
- D^ω : finite or countably infinite sequences over D

D^ω is a complete partial order where
 $X \leq Y$ iff X is an initial segment of Y

BF - ES

- 37 -

Definition: Kahn networks



- associated with each process $v \in V$ with incoming edges e_1, \dots, e_p and outgoing edges e_1', \dots, e_q' is a **continuous function**

$$f_v: D_{e_1}^\omega \times \dots \times D_{e_p}^\omega \rightarrow D_{e_1'}^\omega \times \dots \times D_{e_q'}^\omega$$

(A function $f: A \rightarrow B$ is **continuous** if $f(\lim_A a) = \lim_B f(a)$)

BF - ES

- 38 -

Semantics: Kahn networks

A process network defines for each edge $e \in E$ a **unique** sequence X_e .


X_e is the least fixed point of the equations

$$(X_{e_1}, \dots, X_{e_q}) = f_v(X_{e_1}, \dots, X_{e_q})$$

for all $v \in V$.


Result is independent of scheduling!

Scheduling Kahn Networks



Problem: run processes with finite buffer

Scheduling may be impossible



BF - ES

- 41 -


Parks' Scheduling Algorithm (1995)

- Set a capacity on each channel
- Block a write if the channel is full
- Repeat
 - Run until deadlock occurs
 - If there are no blocking writes → terminate
 - Among the channels that block writes, select the channel with least capacity and increase capacity until producer can fire.

BF - ES

- 42 -

Example



BF - ES

- 43 -

Parks' Scheduling Algorithm

- Whether a Kahn network can execute in bounded memory is undecidable
- Parks' algorithm does not violate this
- It will run in bounded memory if possible, and use unbounded memory if necessary

Disadvantages:


- Requires dynamic memory allocation
- Does not guarantee minimum memory usage
- Scheduling choices may affect memory usage
- Data-dependent decisions may affect memory usage
- Relatively costly scheduling technique
- Detecting deadlock may be difficult

BF - ES

- 44 -

Synchronous data flow (SDF)

- Asynchronous message passing = tasks do not have to wait until output is accepted.
- Synchronous data flow = all tokens are consumed at the same time.



SDF model allows static scheduling of token production and consumption.
In the general case, buffers may be needed at edges.

BF - ES

- 45 -

SDF: restriction of Kahn networks

An **SDF graph** is a tuple $(V, E, \text{cons}, \text{prod}, d)$ where

- V is a set of nodes (activities)
- E is a set of edges (buffers)
- $\text{cons}: E \rightarrow \mathbb{N}$ number of tokens consumed
- $\text{prod}: E \rightarrow \mathbb{N}$ number of tokens produced
- $d: E \rightarrow \mathbb{N}$ number of initial tokens

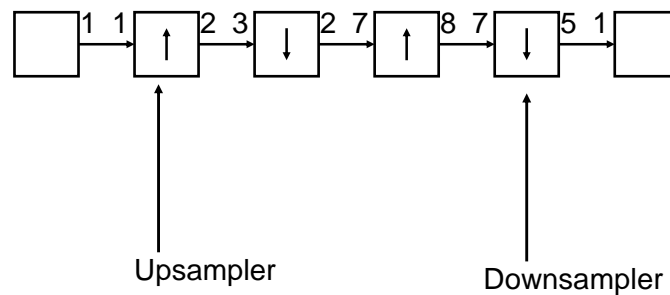
d : „delay“ (sample offset between input and output)

BF - ES

- 46 -

Multi-rate SDF System

- DAT-to-CD rate converter
- Converts a 44.1 kHz sampling rate to 48 kHz



BF - ES

- 47 -

SDF Scheduling Algorithm Lee/Messerschmitt 1987


1. Establish **relative execution rates**
 - Generate balance equations
 - Solve for smallest positive integer vector \mathbf{c}
2. Determine **periodic schedule**
 - Form an arbitrarily ordered list of all nodes in the system
 - Repeat:
 - For each node in the list, schedule it if it is runnable, trying each node once
 - If each node has been scheduled \mathbf{c}_n times, stop.
 - If no node can be scheduled, indicate deadlock.

Source: Lee/Messerschmitt, Synchronous Data Flow (1987)

BF - ES

- 48 -

Example 1




$$d(CA)=6$$

BF - ES

- 49 -


Example 2



BF - ES

- 50 -


Example 3



BF - ES

- 51 -

Example 4



BF - ES

- 52 -

Observations

- Consistent, connected systems have one-dimensional solution
- Disconnected systems have higher-dimensional solution
- Inconsistent systems have 0-schedule; otherwise infinite accumulation of tokens
- Systems may have multiple schedules

BF - ES

- 53 -

Summary dataflow

- Communication exclusively through FIFOs
- Kahn process networks
 - blocking read, nonblocking write
 - deterministic
 - schedulability undecidable
 - Parks' scheduling algorithm
- SDF
 - fixed token consumption/production
 - compile-time scheduling: balance equations

BF - ES

- 54 -