

Implementing S-Net

A Typed Stream Processing Language

— Part I —

Compilation, Code Generation and Deployment

— DRAFT —

Version 0.3

December 8, 2006

Clemens Grellck and Frank Penczek

University of Hertfordshire
Department of Computer Science
Hatfield, Herts, AL10 9AB
United Kingdom

Abstract

We outline the architecture of an S-NET implementation consisting of compiler, code generator, deployer and runtime system. The S-NET compiler actually is a multi-stage compilation framework by itself. For this purpose we define a series of S-NET-like languages that gradually make the shift from full-fledged S-NET to a very simple distant variant of S-NET. The code generator takes this intermediate form to generate proper ISO C code with calls into a comprehensive runtime library. The deployer takes a fully compiled SNet and two interface descriptions for the global input and output stream and forms an executable program. The runtime system actually implements the various language features of S-NET and controls the setup of streaming networks and the orderly cooperative behaviour of asynchronous components. Different runtime system implementations address specific properties of destination architectures.

Contents

1	System Architecture	2
1.1	Overview	2
1.2	Compiler architecture	3
1.3	Code generator architecture	5
1.4	Deployer architecture	6
1.5	Runtime system architecture	6
1.6	Running example	6
2	Compilation	9
2.1	Parsing	9
2.2	Preprocessing	12
2.3	Topology flattening	14
2.4	Type inference	17
2.5	Optimisation	19
2.6	Postprocessing	21
3	Code Generation	31
3.1	Overview	31
3.2	Type representation	31
3.3	Header File	32
3.4	Box	33
3.5	Box Wrapper	33
3.6	Serial Combinator	34
3.7	Parallel Combinator	34
3.8	Split Combinator	35
3.9	Star Combinator	36
3.10	Syncro Cell	37
3.11	Filter	37
4	Deployment	40
5	Runtime System	41

Chapter 1

System Architecture

1.1 Overview

S-NET [1, 2] is a declarative coordination language for describing streaming networks of asynchronous components on a high level of abstraction. Compiling S-NET specifications into efficiently executable code for various hardware architectures is a complex and demanding undertaking. Therefore, we systematically break down this task into smaller, as independent as possible units. As illustrated in Fig. 1.1, we organise the system architecture of an S-NET implementation into four main functional units:

- compiler,
- code generator,
- deployer and
- runtime system.

The S-NET compiler takes an S-NET specification and, in a multi-stage process, transforms it into a substantially different textual representation. This representation is a radically simplified form of S-NET, extended by additional features supported in the runtime system. The S-NET code generator takes these internal textual representations and turns them into proper ISO C code with calls into the runtime system. Technically, the code generator may be considered the final stage of the compiler. However, for the sake of illustration we keep the two issues separate in this document.

The S-NET deployer combines a single (compiled) SNet¹ and two interface descriptions into an executable program. Furthermore, the deployer selects a concrete runtime system implementation.

¹We use different fonts to distinguish between the language S-NET and the SNet networks it describes.

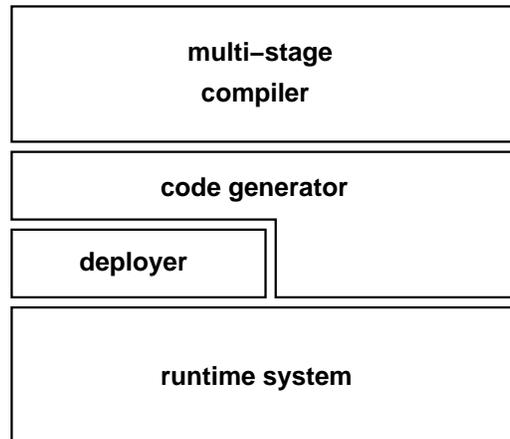


Figure 1.1: System architecture of S-NET implementation

The remainder of this document is organised as follows: The remaining sections of this chapter outline the principles of compiler, code generator, deployer and runtime system before we introduce a simple running example in Section 1.6. We will use this running example to illustrate the details of compiler, code generator, deployer and runtime system in Chapters 2, 3, 4 and 5, respectively.

1.2 Compiler architecture

In order to manage the complexity of compiling fully-fledged declarative S-NET code into a near machine-level representation we define several intermediate variants of S-NET. A multi-stage compilation framework gradually transforms S-NET specifications into less abstract and less declarative code.

Fig. 1.2 shows a sketch of the overall S-NET compiler architecture. We define five compilation stages: preprocessing, topology flattening, type inference, optimisation and postprocessing. In addition we have two auxiliary stages: parsing and printing. The five compilation stages share a common internal representation of S-Nets. The auxiliary stages transform textual representations of S-Nets into this internal representation (parser) and vice versa (printer).

The compilation process may start and stop in any of the compilation stages. The exact stage at which to start is determined by an identifier in the source code: the first line of text must contain a special comment of the form:

```
#!/snet code
```

If this identifier is not present, the compilation process starts at the very beginning. The final compilation stage is determined by a compiler flag. If that stage has been completed, the S-NET compiler prints the intermediate program representation to the standard output stream with the intermediate language identification properly set. The five intermediate languages, S-NET_{core},

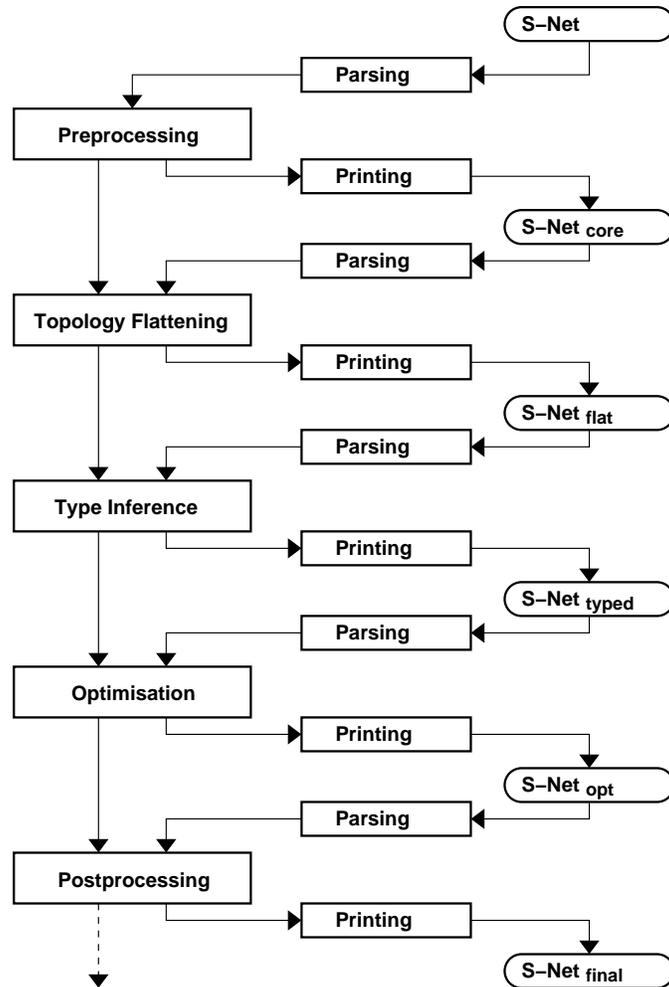


Figure 1.2: Architecture of S-NET compiler

$S\text{-NET}_{\text{flat}}$, $S\text{-NET}_{\text{typed}}$, $S\text{-NET}_{\text{opt}}$ and $S\text{-NET}_{\text{final}}$ are all variants of S-NET itself. Therefore, internal representation, parser and printer can be developed once and parameterised for the various intermediate languages.

The advantage of this multi-stage compiler architecture is that we may develop the individual parts mostly in isolation with well defined interfaces in between them. Ease of use is still achieved by the compiler driver that is responsible for user interaction and the orderly application of the individual compilation stages. Intermediate compiler phases may expect certain side conditions to hold in addition to the purely syntactical restrictions of the intermediate input language. In particular, conditions that have been checked, enforced or

created by preceding compiler phases are not to be checked again. If they are for some reason violated, a compiler phase may arbitrarily fail on the attempt to compile the erroneous code. The feature of stopping and resuming the compilation process is exclusively intended for the sake of compiler development and testing. In a product version it is to be deactivated or entirely removed.

1.3 Code generator architecture

The code generator essentially is the final stage of the S-NET compiler. However, as such it behaves differently from the other stages. It neither transforms the internal intermediate representation like the other stages nor does it merely print the internal representation into a textual format. To emphasise the special role of the code generator we describe it in a separate chapter.

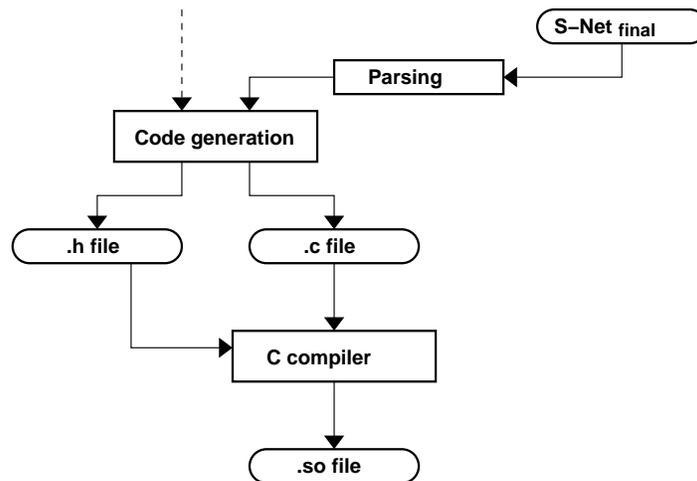


Figure 1.3: Architecture of S-NET code generator

Fig. 1.3 sketches out the architecture of the code generator. The code generator takes an SNet in the final intermediate representation and creates two files: a C source code file and a C header file. The header file contains numerical encodings for all field and tag names used throughout the SNet. Furthermore, it contains an external declaration of the generated function representing the exported network. They are needed for code generation whenever the given SNet is used in the context of another SNet. The C source code file, which includes the header file is then fed into an ISO C compiler that creates the final object code file, which will be taken by the S-NET deployer to form an executable SNet.

1.4 Deployer architecture

... yet to be developed in detail ...

1.5 Runtime system architecture

The runtime system is a rich library of system calls for runtime representations of types and patterns, for setting up S-Nets at runtime and for the dynamic control of asynchronous S-NET components and the communication channels between them.

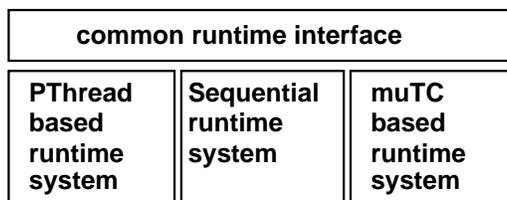


Figure 1.4: Architecture of S-NET runtime system

We show a sketch of the runtime system architecture in Fig. 1.4. The *common runtime interface* is an abstraction layer that allows us to support different target architectures without affecting the compilation and code generation process. For the time being, we envision three destination architectures:

- sequential execution,
- multithreading based on PTHREADS [3] and
- multithreading based on μ TC [4].

The common runtime interface shields the specific properties of these and other target architectures from the S-NET compiler and code generator. Actually, changing the target architecture does not even require the recompilation of an SNet, but merely linking with a different runtime system implementation. Hence, the selection of a concrete target architecture is part of the deployment of S-Nets.

1.6 Running example

We illustrate both the compilation and the code generation process by means of the running example shown in Fig. 1.5. In order to incorporate as many as possible S-NET language features in a single example without making it overly complicated we use an abstract and artificial SNet rather than a concrete S-NET application.

```

type compAB_t = {A} -> {P}, {B} -> {Q};

net compABC ({A} | {C} -> {P}, {B} -> {Q}) {
  box compA ((A) -> (P));
  box compB ((B) -> (Q));
  box compC ((C) -> (P));
}
connect compA || compB || compC;

net example {
  net split
  connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

  box examine( (P,Q) -> (A,B) | (Y,Z));

  net compute {
    net compAB (compAB_t)
    connect compABC;

    net syncPQ
    connect [|{P},{Q}|] ** ({P,Q});
  }
  connect ( [{<T>} -> {}] .. compAB .. syncPQ ) !! (<T>);
}
connect (tag .. split .. compute .. examine) ** ({Y,Z});

```

Figure 1.5: Running example to illustrate compilation and code generation

The SNet `example` in Fig. 1.5 contains two top-level SNETs: the auxiliary network `compABC` and the exported main network `example`. Whereas the former is a rather simple parallel composition of three boxes, the latter contains a serial composition of four subnetworks, `tag`, `split`, `compute` and `examine` embedded within a star combinator. Although this is not annotated in Fig. 1.5, we expect the network `example` to receive incoming records with fields A and B.

There is no definition of the network `tag`. So, we expect another file `tag.so` to contain the definition. Nevertheless, the idea of `tag` is to add a tag T to each record. The network `split`, which is made up of a single filter box, splits each record into two records, one containing field A and tag T and the other one containing field B and a copy of tag T.

The subnetwork `compute` does some computation on the data before the box `examine` checks incoming records for some termination condition. The latter either produces a new record {A,B} or a new record {Y,Z}. Given the termination condition of the star combinator in the connect expression of `example`, {A,B} records are directed into a new incarnation of the `tag..split..compute..examine` sequence while {Y,Z} records are directed to the global output stream.

The interior of the subnetwork `compute` is dynamically replicated using the parallel replication combinator (!!) based on the concrete values of tag T. Within, we first strip the tag T from each record as its sole purpose was to select the proper instance of this parallel replication. The actual computation is

performed by the subnetwork `compAB`. By providing a type signature for `compAB` we effectively specialise the top-level network `compABC` to only handle incoming records `{A}` and `{B}`, but not `{C}`. The resulting `{P}` and `{Q}` records are pairwise synchronised by a “starred” syncrocell in subnetwork `syncPQ`.

Chapter 2

Compilation

2.1 Parsing

The parser is effectively a 2-stage process, as sketched out in Fig. 2.1. For separation of concerns we distinguish between pure parsing and a separate dispatch phase. The pure parser reads in a textual S-NET specification and checks it for lexical and syntactical correctness. However, it does not do any further context checks. Pure parsing, hence, results in an internal representation that keeps all applied occurrences of net or box names and, likewise, all occurrences of record field names or tag names as plain character strings. In particular, the parser does not check context conditions such as the existence of a network definition for each instantiation. Such checks are deliberately organised into a separate dispatch phase to let the parser implementation focus on pure parsing issues.

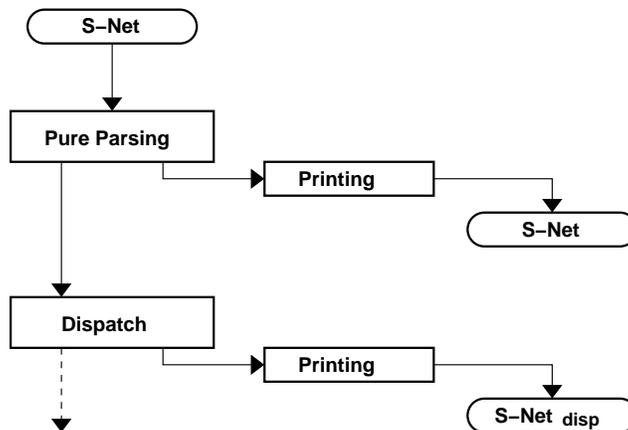


Figure 2.1: Architecture of parser

In analogy, to the overall compiler design, the compilation process may be terminated after pure parsing and the internal representation converted back into a textual specification. This feature allows us to verify the parsing step in isolation during the development process. Apart from source code comments, whitespace characters and text formatting issues this textual representation coincides with the original source S-NET specification.

$$\begin{aligned}
 SNet_{disp} &\Rightarrow [Declaration]^* [Definition]^* \\
 Declaration &\Rightarrow \mathbf{net} \ NetName \ ; \\
 Serial_{disp} &\Rightarrow (\ SNetExpr \ SerialCombinator \ SNetExpr \) \\
 Star_{disp} &\Rightarrow (\ SNetExpr \ StarCombinator \ Terminator \) \\
 Choice_{disp} &\Rightarrow (\ SNetExpr \ ChoiceCombinator \ SNetExpr \) \\
 Split_{disp} &\Rightarrow (\ SNetExpr \ SplitCombinator \ Range \)
 \end{aligned}$$

Figure 2.2: Grammar of S-NET_{disp}

The dispatcher checks context conditions such as the existence of definitions for each instantiation of a locally defined net or box. At each such instantiation the dispatcher replaces the textual specification of the net or box by a reference to its definition. If the dispatcher does not find a matching local definition in the current scope, it considers the name to refer to an external SNet definition and creates a network declaration as stub code. Fig. 2.2 provides a formal definition of the intermediate language S-NET_{disp}. The sole differences with respect to S-NET proper are the network declaration part prior to the standard definitions and the resolution of combinator associativities and priorities in complex nested connect expressions by the parser: S-NET_{disp} expects all topology definitions to be fully parenthesised.

Unlike all other compiler phases, which have a unique place in the overall compiler architecture, the parser may be used to process partially compiled intermediate code as well. In any case but parsing original S-NET source code, the dispatcher must expect external network declarations to be already present. If it still does not find a matching declaration for an otherwise unbound network identifier, the dispatcher issues an error message to properly report the detected inconsistency in the intermediate code.

In analogy to the storage of identifiers for boxes and networks, the dispatcher stores all record field names and tag names occurring in the SNet in a separate data base and replaces their names in all applied occurrences by a reference to the respective data base entry. Among others, these steps facilitate consistent renaming in the subsequent course of compilation.

References cannot be represented in textual representations properly. Therefore, we print plain names rather than references when returning to a textual representation of an SNet after dispatching. The need to restore references in the internal representation from names in the textual specification motivates us to integrate the dispatcher into the parser rather than into the preprocessor.

```

//! snet disp
net tag;
type compAB_t = {A} -> {P}, {B} -> {Q};
net compABC ({A} | {C} -> {P}, {B} -> {Q}) {
  box compA ((A) -> (P));
  box compB ((B) -> (Q));
  box compC ((C) -> (P));
}
connect (compA || (compB || compC));

net example {
  net split
  connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

  box examine( (P,Q) -> (A,B) | (Y,Z));

  net compute {
    net compAB (compAB_t)
    connect compABC;

    net syncPQ
    connect ([{P},{Q}] ** ({P,Q}));
  }
  connect (([<T>] -> []) .. (compAB .. syncPQ)) !! (<T>);
}
connect ((tag .. (split .. (compute .. examine)))**({Y,Z}));

```

Figure 2.3: Running example after parsing

The dispatch needs to be done whenever we convert a textual specification into the compiler-internal representation, regardless of where exactly we are in the overall compilation process otherwise.

Fig. 2.3 illustrates the effects of parsing (i.e. pure parsing and dispatch) on the running example, as introduced in Section 1.6. Those lines of code that show differences with respect to the original example are marked by a grey background. First, we observe the special comment in the first line of code that determines the state of the subsequent code with respect to the overall compilation process. It follows a network declaration for `tag` that is referred to in the connect expression of `example`, but nowhere defined in our example. The other difference we may observe are the additional parentheses in various connect expressions that make the implicitly defined associativities and priorities of combinators explicit.

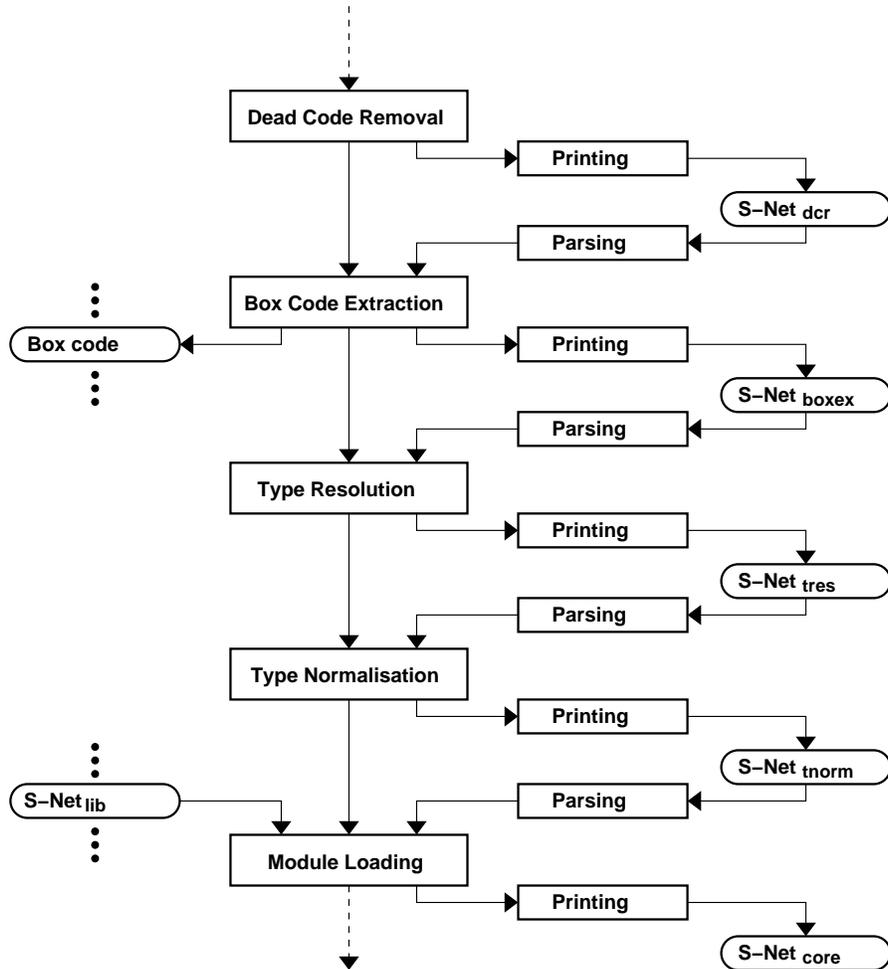


Figure 2.4: Architecture of preprocessor

2.2 Preprocessing

The preprocessor compiles $S\text{-NET}_{\text{disp}}$ code into the intermediate language $S\text{-NET}_{\text{core}}$. Effectively, it pursues a collection of tasks that are partly unrelated to each other, but share the common goal of making $S\text{-NET}_{\text{core}}$ a less complex language as compared with $S\text{-NET}$ proper. Therefore, we adopt the overall compiler architecture for the preprocessor and organise the preprocessing stage again as a sequence of individual transformations with intermediate textual representations in between them. Fig. 2.4 gives an outline.

As a first step, we remove all those SNet definitions that are neither directly

nor indirectly referred to by the externally visible network definition that bears the same name as the file being compiled. Since this is a rather odd situation, we issue a warning to the programmer. The sole purpose of this transformation is to accelerate subsequent transformations by avoiding their application in useless situations. Dead Code Removal does (obviously) not change the intermediate representation. Nevertheless, we formally define the intermediate language S-NET_{dcr} to distinguish intermediate code that has already undergone dead code removal from intermediate code that still needs to be cleaned up.

The second preprocessing step extracts inlined box language code from S-NET specifications. Such code is without inspection written to one external file per box language. Following this step the internal representation of S-Nets no longer allows for inlined box language code.

The next task of the preprocessor is the resolution of type definitions. Type definitions in S-NET (key word `type`) are nothing but placeholders for the respective definitions. Therefore, we replace each applied occurrence of a type name by the corresponding type definition.

Next, the normalisation of type representations means replacing type signatures with multi-variant input types by equivalent type signatures entirely made up of single-variant input types and additional type mappings.

$$\begin{aligned}
 Declaration_{core} &\Rightarrow \mathbf{net} \text{ } NetName \text{ (} TypeSignature \text{) } ; \\
 Definition_{core} &\Rightarrow BoxDef \mid NetDef \\
 BoxDef_{core} &\Rightarrow \mathbf{box} \text{ } BoxName \text{ (} BoxSignature \text{) } ; \\
 TypeMapping_{core} &\Rightarrow RecordType \rightarrow Type \\
 Type_{core} &\Rightarrow RecordType [\mid RecordType]^*
 \end{aligned}$$

Figure 2.5: Grammar of S-NET_{core}

The module loader identifies external network declarations introduced by the dispatcher and identifies the corresponding compiled S-Nets searching in the current directory, in a directory path specified via a command line parameter on initiation of the compilation process and, eventually, in a similar path stored in an environment variable. If the module loader is unable to locate a compiled version of an external network, it issues an error message. Otherwise, it dynamically links with the with the corresponding S-NET library and calls a specific function from that library that recreates the internal representation of the network’s type signature. This type signature is needed to continue with the compilation process.

Fig. 2.5 provides a formal definition of S-NET_{core} as far as it differs from S-NET_{disp}. We see that due to the module loader network declarations now feature a type signature. Type definitions have vanished from the set of symbols that can be defined. Likewise, types may no longer use type names to refer to preceding type definitions. Last but not least, box definitions lack the potential

of inlined box language implementations

```

//! snet core
net tag ({A,B} -> {A,B,<T>});
net compABC ({A} -> {P}, {B} -> {Q}, {C} -> {P}) {
    box compA( (A) -> (P));
    box compB( (B) -> (Q));
    box compC( (C) -> (P));
}
connect (compA || (compB || compC));

net example {
    net split
    connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

    box examine( (P,Q) -> (A,B) | (Y,Z));

    net compute {
        net compAB ({A} -> {P}, {B} -> {Q})

        connect compABC;

        net syncPQ
        connect ([{P},{Q}] ** ({P,Q}));
    }
    connect (([<T>} -> {]} .. (compAB .. syncPQ)) !! (<T>));
}
connect ((tag .. (split .. (compute .. examine))) ** ({Y,Z}));

```

Figure 2.6: Running example after complete preprocessing

The aggregate effect of preprocessing on our running example can be seen in Fig. 2.6 We easily observe the new type signature in the declaration of the external network `tag` and the disappearance of the type definition of `compAB.t`.

2.3 Topology flattening

The topology flatter simplifies complex network topology specifications by systematically abstracting S-NET expressions into additional networks. Formally, the topology flatter turns S-NET_{core} specifications into the S-NET_{flat} intermediate representation format, defined in Fig. 2.7.

The significant difference between S-NET/S-NET_{core} and S-NET_{flat} is the restriction of operand networks of the four network combinators, serial, star, choice and split, to named networks. As a consequence, we may no longer represent nested topology expressions. The connect expression of any network may only contain a single instance of a box, a filter or a synchrocell. Alternatively, it may contain a single application of a network combinator to networks referred to by their names. In particular, each box, each filter and each synchrocell is

$SNetExpr_{flat}$	\Rightarrow	$BoxName$ $NetName$ $Sync$ $Filter$ $Combination$
$Serial_{flat}$	\Rightarrow	$NetName$ $SerialCombinator$ $NetName$
$Star_{flat}$	\Rightarrow	$NetName$ $StarCombinator$ $Terminator$
$Choice_{flat}$	\Rightarrow	$NetName$ $ChoiceCombinator$ $NetName$
$Split_{flat}$	\Rightarrow	$NetName$ $SplitCombinator$ $Range$

Figure 2.7: Grammar of S-NET_{flat}

embedded within its own network. This step prepares the internal representation of S-Nets for the subsequent type inference phase as we may now associate each level of a previously nested topology expression with its type signature.

Fig. 2.8 shows the impact of topology flattening on the running example. We observe the systematic recursive extraction of subexpressions from the network topology specifications of `compABC`, `example` and `compute` into the preceding contexts.

Topology flattening requires the introduction of new networks and, hence,

```

//! snet flat

net tag ({A,B} -> {A,B,<T>});

net compABC ({A} -> {P}, {B} -> {Q}, {C} -> {P}) {
  box compA ((A) -> (P));
  box compB ((B) -> (Q));
  box compC ((C) -> (P));

  net _SL
  connect compA;

  net _SR {
    net _SL
    connect compB;

    net _SR
    connect compC;
  }
  connect _SL || _SR;
}
connect _SL || _SR;

```

Figure 2.8: Running example after topology flattening
(continued on next page)

```

net example {
  net split
  connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

  box examine( (P,Q) -> (A,B) | (Y,Z));

  net compute {
    net compAB ({A} -> {P}, {B} -> {Q})
    connect compABC;

    net syncPQ {
      net _ST
      connect [|{P},{Q}|];
    }
    connect _ST ** ({P,Q});

    net _IS {
      net _SL
      connect [{<T>} -> {}];

      net _SR
      connect compAB .. syncPQ;
    }
    connect _SL .. _SR;
  }
  connect _IS !! (<T>);

  net _ST {
    net _SR {
      net _SR {
        net _SR
        connect examine;
      }
      connect compute .. _SR;
    }
    connect split .. _SR;
  }
  connect tag .. _SR;
}
connect _ST ** ({Y,Z});

```

Figure 2.8: Running example after topology flattening
(continued from previous page)

new network names. The scheme for generation of network names must meet two requirements. Firstly, new network names must not collide with existing network names chosen by the programmer. In Fig. 2.8 all new network names start with an underscore letter; leading underscores are not permitted in language level S-NET identifiers. Secondly, the choice of name should facilitate the interpretation of flattened S-Nets with respect to the original specifications for a human reader. In Fig. 2.8 we use an algorithm that systematically creates names from the original location of a flattened network in the original topology

expression. The acronyms translate as described in Fig. 2.9.

SL	serial left
SR	serial right
PL	parallel left
PR	parallel right
ST	star
IS	index split

Figure 2.9: Creating network names from location in topology expression

2.4 Type inference

The type inference system, as the name suggests, infers a type signature for each $S\text{-NET}_{\text{flat}}$ network following the formal type rules presented in [2]. More formally, the type inference system turns $S\text{-NET}_{\text{flat}}$ code into $S\text{-NET}_{\text{typed}}$ code, as defined in Fig. 2.10. In fact, the syntactic differences between $S\text{-NET}_{\text{flat}}$ and $S\text{-NET}_{\text{typed}}$ are rather small: We only require that each network definition is associated with a type signature.

More precisely, the type inference system may infer a type in addition to a potentially programmer-supplied type signature or input type specification. In fact, existing type information requires the type inference system to do more than just inference. As explained in [2], the annotated type signature must be in box subtype relationship to the inferred type signature. Here, the type inference system effectively becomes a type checker and produces an error message if the condition is not met. If the type check succeeds, type inference continues with the annotated type rather than the inferred type.

$$\begin{aligned}
 \text{NetDef}_{\text{typed}} &\Rightarrow \mathbf{net} \text{ NetName NetTypes } [\text{NetBody}] \text{ Connect} \\
 \text{NetTypes} &\Rightarrow (\text{TypeSignature}) [(\text{NetSignature})]
 \end{aligned}$$

Figure 2.10: Grammar of $S\text{-NET}_{\text{typed}}$

Instead of a full type signature, $S\text{-NET}$ allows the programmer to only annotate an input type to a network definition. If so, it must be in subtype relationship to the input type of the inferred type signature. Otherwise, we produce a type error message. Type inference continues with a type signature amalgamated from the inferred type signature and the annotated input type. If the annotated input type contains less variants than the input type of the inferred type signature, the additional type mappings are eliminated from the type signature. If a variant of the input type contains additional record entries compared with the corresponding variant of the input type of the inferred type signature, the additional record entries are added to the right hand side of the

corresponding type mapping. Last but not least, the output type of the annotated type signature may contain fewer record entries than the output type of the inferred type signature. In this case, the type inference system introduces appropriate filter boxes to adapt the internal (inferred) type signature to the annotated type signature.

If the inferred type signature is not identical to the one with which we continue type inference, the annotated type information is stored as an optional auxiliary type signature in the textual representation of `S-NETtyped`. Differences between annotated and inferred type information may be exploited for optimisation purposes at a subsequent compilation stage.

```


//! snet typed

net tag ({A,B} -> {A,B,<T>});

net compABC ({A} -> {P}, {B} -> {Q}, {C} -> {P}) {
  box compA ((A) -> (P));
  box compB ((B) -> (Q));
  box compC ((C) -> (P));

  net _SL ({A} -> {P})
  connect compA;

  net _SR ({B} -> {Q}, {C} -> {P}) {
    net _SL ({B} -> {Q})
    connect compB;

    net _SR ({C} -> {P})
    connect compC;
  }
  connect _SL || _SR;
}
connect _SL || _SR;


```

Figure 2.11: Running example after type inference
(continued on next page)

Fig. 2.11 demonstrate the effect of type inference on the running example. Each and every network definition is now associated with a type signature. In the case of connect expressions that solely consist of the instance of a box, a filter or a synchronisation pattern, the type signature may be derived rather straightforwardly from the box signature, the filter encoding or the synchronisation pattern, respectively. The other cases involving network combinator applications are handled as described in [2].

In the case of `compABC` the inferred type signature turned out to be identical to the annotated type signature. Hence, we store only one. The situation is different with `compAB`. Here, we infer a type signature that has one additional type mapping when compared to the annotated type signature. As a consequence, we keep both.

```

net example ({A,B} -> {Y,Z}) {
  net split ({A,B,<T>} -> {A,<T>} | {B,<T>})
  connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

  box examine ((P,Q) -> (A,B) | (Y,Z));

  net compute ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q}) {
    net compAB ({A} -> {P}, {B} -> {Q}, {C} -> {P})
      ({A} -> {P}, {B} -> {Q})
    connect compABC;

    net syncPQ ({P} -> {P,Q}, {Q} -> {P,Q}) {
      net _ST ({P} -> {P} | {P,Q}, {Q} -> {Q} | {P,Q})
        connect [| {P}, {Q}|];
    }
    connect _ST ** ({P,Q});

    net _IS ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q}) {
      net _SL ({<T>} -> {})
        connect [{<T>} -> {}];

      net _SR ({A} -> {P,Q}, {B} -> {P,Q})
        connect compAB .. syncPQ;
    }
    connect _SL .. _SR;
  }
  connect _IS !! (<T>);

  net _ST ({A,B} -> {A,B} | {Y,Z}) {
    net _SR ({A,B,<T>} -> {A,B} | {Y,Z}) {
      net _SR ({A,<T>} -> {A,B} | {Y,Z}, {B,<T>} -> {A,B} | {Y,Z}) {
        net _SR ({P,Q} -> {A,B} | {Y,Z})
          connect examine;
      }
      connect compute .. _SR;
    }
    connect split .. _SR;
  }
  connect tag .. _SR;
}
connect _ST ** ({Y,Z});

```

Figure 2.11: Running example after type inference
(continued from previous page)

2.5 Optimisation

The purpose of the optimiser is to reduce resource requirements and to improve the runtime performance of compiled SNETs. For this purpose, the optimiser must make assumptions on both the deployer and on the runtime system. Optimisation naturally is a collection of independent and interdependent code transformations. Optimisation may always be considered optional and may (selectively) switched on and off for evaluation purposes.

```

//! snet opt

net tag ({A,B} -> {A,B,<T>});

net compABC ({A} -> {P}, {B} -> {Q}, {C} -> {P}) {
  box compA ((A) -> (P));

  box compB ((B) -> (Q));

  box compC ((C) -> (P));

  net _SL ({A} -> {P})
  connect compA;

  net _SR__SL ({B} -> {Q})
  connect compB;

  net _SR__SR ({C} -> {P})
  connect compC;
}
connect || _SL _SR__SL _SR__SR;

net example ({A,B} -> {Y,Z}) {
  net split ({A,B,<T>} -> {A,<T>} | {B,<T>})
  connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

  box examine ((P,Q) -> (A,B) | (Y,Z));

```

Figure 2.13: Running example after optimisation
(continued on next page)

else. However, we do keep it to illustrate the further processing of multi-ary choice combinators.

2.6 Postprocessing

In analogy to the preprocessor, the postprocessor again is a collection of several independent measures that prepare the internal representation of intermediate S-NET code for the final code generation step. Fig. 2.14 provides an outline of the individual phases.

The first postprocessing phase is a package translator. Whenever a record crosses the boundary between two S-NET packages (i.e. files), we need to translate the record field and tag names from the domain of one file (or unit of compilation) to the domain of the other file. We explicitly represent this translation in the intermediate code by introducing *translators* and package qualifiers for external symbols.

Fig. 2.15 illustrates the introduction of translators by an excerpt from our running example focussing on the instantiation of the external net `tag`. All symbols in the external network declaration of `tag` are now qualified by the

```

net compute ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q}) {
  net compAB ({A} -> {P}, {B} -> {Q}) {
    net compABC ({A} -> {P}, {B} -> {Q}) {
      box compA ((A) -> (P));

      box compB ((B) -> (Q));

      net _SL ({A} -> {P})
      connect compA;

      net _SR_SL ({B} -> {Q})
      connect compB;
    }
    connect _SL || _SR_SL;
  }
  connect compABC;

  net syncPQ ({P} -> {P,Q}, {Q} -> {P,Q}) {
    net _ST ({P} -> {P} | {P,Q}, {Q} -> {Q} | {P,Q})
    connect [| {P}, {Q}|];
  }
  connect _ST ** ({P,Q});

  net _IS ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q}) {
    net _SL (<T> -> {})
    connect [<T> -> {}];

    net _SR ({A} -> {P,Q}, {B} -> {P,Q})
    connect compAB .. syncPQ;
  }
  connect _SL .. _SR;
}
connect _IS !! (<T>);

net _ST ({A,B} -> {A,B} | {Y,Z}) {
  net _SR ({A,B,<T>} -> {A,B} | {Y,Z}) {
    net _SR ({A,<T>} -> {A,B} | {Y,Z}, {B,<T>} -> {A,B} | {Y,Z}) {
      net _SR ({P,Q} -> {A,B} | {Y,Z})
      connect examine;
    }
    connect compute .. _SR;
  }
  connect split .. _SR;
}
connect tag .. _SR;
}
connect _ST ** ({Y,Z});

```

Figure 2.13: Running example after optimisation
(continued from previous page)

package name `tag`. The translators themselves are embedded within a new wrapper network named `tag` (without package qualifier). All references to the external network `tag` in the file now point to this wrapper rather than to the

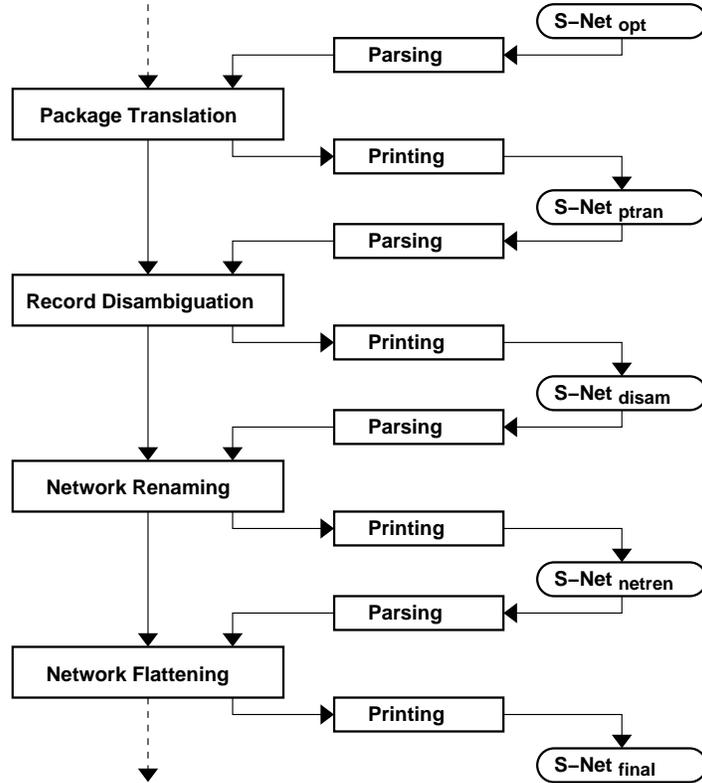


Figure 2.14: Architecture of postprocessor

external declaration. The wrapper network is the flattened and fully typed representation of the S-NET expression

```
[(A,B) -> (tag::A,tag::B)]
.. tag::tag
.. [(tag::A,tag::B,<tag::T>) -> (A,B,<T>)]
```

The syntax of the translators resembles that of filters. However, we use box types with round brackets rather than record types with curly brackets. This emphasises the significance of order as, for example in the first translator, we must map `A` to `tag::A` and `B` to `tag::B`. Fig. 2.16 provides a formal definition of the grammar of $S\text{-NET}_{\text{ptran}}$.

The second postprocessing step is a record entry disambiguation phase. The purpose of this phase is to separate unrelated occurrences of identical field and tag names by renaming. In particular, we must guarantee that the names of record entries that are flow-inherited at some level of our network tree are not used again at a deeper level. Whenever we detect such a case, we rename the

```

//! snet ptran

net tag::tag ({tag::A,tag::B} -> {tag::A,tag::B,<tag::T>});

net tag ({A,B} -> {A,B,<T>}) {
  net translate_in ({A,B} -> {tag::A,tag::B})
  connect [(A,B) -> (tag::A,tag::B)];

  net translate_out ({tag::A,tag::B,<tag::T>} -> {A,B,<T>})
  connect [(tag::A,tag::B,<tag::T>) -> (A,B,<T>)];

  net tag ({tag::A,tag::B} -> {tag::A,tag::B,<tag::T>})
  connect tag::tag;

  net in ({A,B} -> {tag::A,tag::B,<tag::T>})
  connect translate_in .. tag;
}
connect in .. translate_out;

```

Figure 2.15: Running example after package translation (excerpt)

nested occurrence of the name. This step is essential for the efficient implementation of flow inheritance. Since it requires a relatively complex setting of nested network specifications to create the need for record entry disambiguation, we refrain from an illustration in the course of our running example.

The third postprocessing phase consistently renames all networks to reflect the specific location of their definition within the overall tree structure of SNet definitions. Fig. 2.17 demonstrates the effect of network renaming on the running example. Essentially each network name is prefixed with the names of all networks in whose definitions it is embedded. The use of double underscores,

$Declaration_{ptran}$	\Rightarrow	net <i>Netname</i> :: <i>NetName</i> (<i>TypeSignature</i>) ;
$SNetExpr_{ptran}$	\Rightarrow	<i>BoxName</i> <i>[Netname ::] NetName</i> <i>Sync</i> <i>Filter</i> <i>Translator</i> <i>Combination</i>
<i>Translator</i>	\Rightarrow	[<i>BoxType</i> -> <i>BoxType</i>]
<i>Field</i>	\Rightarrow	<i>[Netname ::] FieldName</i>
<i>SimpleTag</i>	\Rightarrow	< <i>[Netname ::] TagName</i> >
<i>BindingTag</i>	\Rightarrow	< # <i>[Netname ::] TagName</i> >

Figure 2.16: Grammar of S-NET_{ptran}

```

//! snet netren

net tag::tag ({tag::A,tag::B} -> {tag::A,tag::B,tag::<T>});

net tag ({A,B} -> {A,B,<T>}) {
  net tag__translate_in ({A,B} -> {tag::A,tag::B})
  connect [(A,B) -> (tag::A,tag::B)];

  net tag__translate_out ({tag::A,tag::B,tag::<T>} -> {A,B,<T>})
  connect [(tag::A,tag::B,tag::<T>) -> (A,B,<T>)];

  net tag__tag ({tag::A,tag::B} -> {tag::A,tag::B,<tag::T>})
  connect tag::tag;

  net tag__in ({A,B} -> {tag::A,tag::B,tag::<T>})
  connect tag__translate_in .. tag__tag;
}
connect tag__in .. tag__translate_out;

net compABC ({A} -> {P}, {B} -> {Q}, {C} -> {P}) {
  box compA ((A) -> (P));
  box compB ((B) -> (Q));
  box compC ((C) -> (P));

  net compABC__SL ({A} -> {P})
  connect compA;

  net compABC__SR__SL ({B} -> {Q})
  connect compB;

  net compABC__SR__SR ({C} -> {P})
  connect compC;
}
connect || compABC__SL compABC__SR__SL compABC__SR__SR;

net example ({A,B} -> {Y,Z}) {

  net example__split ({A,B,<T>} -> {A,<T>} | {B,<T>})
  connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

  box examine ((P,Q) -> (A,B) | (Y,Z));
}

```

Figure 2.17: Running example after network renaming
(continued on next page)

which is ruled out in proper S-NET, to separate the various prefixes from each other and from the original name prevents name clashes. Not that we deliberately do not rename boxes because we consider box names to refer to some external implementation function whose name is fixed.

Network renaming is a preparation step for the final postprocessing phase: the transformation of the hierarchical network structure into a flat sequence of network and box definitions. Separation of renaming and restructuring fa-

```

net example__compute ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q}) {
  net example__compute__compAB ({A} -> {P}, {B} -> {Q}) {
    net example__compute__compAB__compABC ({A} -> {P}, {B} -> {Q}) {
      box compA ((A) -> (P));
      box compB ((B) -> (Q));

      net example__compute__compAB__compABC__SL ({A} -> {P})
      connect compA;

      net example__compute__compAB__compABC__SR__SL ({B} -> {Q})
      connect compB;
    }
    connect example__compute__compAB__compABC__SL
      || example__compute__compAB__compABC__SR__SL;
  }
  connect example__compute__compAB__compABC;

  net example__compute__syncPQ ({P} -> {P,Q}, {Q} -> {P,Q}) {
    net example__compute__syncPQ__ST ({P} -> {P} | {P,Q},
      {Q} -> {Q} | {P,Q})
    connect [| {P}, {Q}|];
  }
  connect example__compute__syncPQ__ST ** ({P,Q});

  net example__compute__IS ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q}) {
    net example__compute__IS__SL ({<T>} -> {})
    connect [{<T>} -> {}];

    net example__compute__IS__SR ({A} -> {P,Q}, {B} -> {P,Q})
    connect example__compute__compAB .. example__compute__syncPQ;
  }
  connect example__compute__IS__SL .. example__compute__IS__SR;
}
connect example__compute__IS !! (<T>);

net example__ST ({A,B} -> {A,B} | {Y,Z}) {
  net example__ST__SR ({A,B,<T>} -> {A,B} | {Y,Z}) {
    net example__ST__SR__SR ({A,<T>} -> {A,B} | {Y,Z},
      {B,<T>} -> {A,B} | {Y,Z}) {
      net example__ST__SR__SR__SR ({P,Q} -> {A,B} | {Y,Z})
      connect examine;
    }
    connect example__compute .. example__ST__SR__SR__SR;
  }
  connect example__split .. example__ST__SR__SR;
}
connect tag .. example__ST__SR;
}
connect example__ST ** ({Y,Z});

```

Figure 2.17: Running example after network renaming
(continued from previous page)

```

//! snet final

net tag::tag ({tag::A,tag::B} -> {tag::A,tag::B,tag::<T>});

net tag__translate_in ({A,B} -> {tag::A,tag::B})
connect [(A,B) -> (tag::A,tag::B)];

net tag__translate_out ({tag::A,tag::B,tag::<T>} -> {A,B,<T>})
connect [(tag::A,tag::B,tag::<T>) -> (A,B,<T>)];

net tag__tag ({tag::A,tag::B} -> {tag::A,tag::B,<tag::T>})
connect tag::tag;

net tag__in ({A,B} -> {tag::A,tag::B,tag::<T>})
connect tag__translate_in .. tag__tag;

net tag ({A,B} -> {A,B,<T>})
connect tag__in .. tag__translate_out;

box compA ((A) -> (P));
box compB ((B) -> (Q));
box compC ((C) -> (P));

net compABC__SL ({A} -> {P})
connect compA;

net compABC__SR__SL ({B} -> {Q})
connect compB;

net compABC__SR__SR ({C} -> {P})
connect compC;

net compABC ({A} -> {P}, {B} -> {Q}, {C} -> {P})
connect || compABC__SL compABC__SR__SL compABC__SR__SR;

net example__split ({A,B,<T>} -> {A,<T>} | {B,<T>})
connect [{A,B,<T>} -> {A,<T>}; {B,<T>}];

box examine ((P,Q) -> (A,B) | (Y,Z));

box compA ((A) -> (P));
box compB ((B) -> (Q));

```

Figure 2.18: Running example after complete postprocessing
(continued on next page)

facilitates the realisation and maintenance of both individual phases. Fig. 2.18 shows the final S-NET representation of our running example. This concludes the compilation process. From this representation we start the generation of ISO C code, as explained in Chapter 3. Fig. 2.19 summarises the complete grammar of S-NET_{final}.

```

net example__compute__compAB__compABC___SL ({A} -> {P})
connect compA;

net example__compute__compAB__compABC___SR__SL ({B} -> {Q})
connect compB;

net example__compute__compAB__compABC ({A} -> {P}, {B} -> {Q})
connect example__compute__compAB__compABC___SL
      || example__compute__compAB__compABC___SR__SL;

net example__compute__compAB ({A} -> {P}, {B} -> {Q})
connect example__compute__compAB__compABC;

net example__compute__syncPQ___ST ({P} -> {P} | {P,Q},
                                   {Q} -> {Q} | {P,Q})
connect [| {P}, {Q}|];

net example__compute__syncPQ ({P} -> {P,Q}, {Q} -> {P,Q})
connect example__compute__syncPQ___ST ** ({P,Q});

net example__compute___IS___SL (<T> -> {})
connect [{<T>} -> {}];

net example__compute___IS___SR ({A} -> {P,Q}, {B} -> {P,Q})
connect example__compute__compAB .. example__compute__syncPQ;

net example__compute___IS ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q})
connect example__compute___IS___SL .. example__compute___IS___SR;

net example__compute ({A,<T>} -> {P,Q}, {B,<T>} -> {P,Q})
connect example__compute___IS !! (<T>);

net example___ST___SR___SR ({P,Q} -> {A,B} | {Y,Z})
connect examine;

net example___ST___SR___SR ({A,<T>} -> {A,B} | {Y,Z},
                             {B,<T>} -> {A,B} | {Y,Z})
connect example__compute .. example___ST___SR___SR___SR;

net example___ST___SR ({A,B,<T>} -> {A,B} | {Y,Z})
connect example__split .. example___ST___SR___SR;

net example___ST ({A,B} -> {A,B} | {Y,Z})
connect tag .. example___ST___SR;

net example ({A,B} -> {Y,Z})
connect example___ST ** ({Y,Z});

```

Figure 2.18: Running example after complete postprocessing
(continued from previous page)

<i>SNet</i>	\Rightarrow	$[Declaration]^* [Definition]^*$
<i>Declaration</i>	\Rightarrow	net <i>Netname</i> :: <i>NetName</i> (<i>TypeSignature</i>) ;
<i>Definition</i>	\Rightarrow	<i>BoxDef</i> <i>NetDef</i>
<i>BoxDef</i>	\Rightarrow	box <i>BoxName</i> (<i>BoxSignature</i>) ;
<i>BoxSignature</i>	\Rightarrow	<i>BoxType</i> \rightarrow <i>BoxType</i> [\ <i>BoxType</i>]*
<i>BoxType</i>	\Rightarrow	([<i>RecordEntry</i> [, <i>RecordEntry</i>]*])
<i>NetDef</i>	\Rightarrow	net <i>NetName</i> (<i>TypeSignature</i>) <i>Connect</i>
<i>TypeSignature</i>	\Rightarrow	<i>TypeMapping</i> [, <i>TypeMapping</i>]*
<i>TypeMapping</i>	\Rightarrow	<i>RecordType</i> \rightarrow <i>RecordType</i> [\ <i>RecordType</i>]*
<i>RecordType</i>	\Rightarrow	{ [<i>RecordEntry</i> [, <i>RecordEntry</i>]*] }
<i>RecordEntry</i>	\Rightarrow	<i>Field</i> <i>Tag</i>
<i>Field</i>	\Rightarrow	[<i>Netname</i> ::] <i>FieldName</i>
<i>Tag</i>	\Rightarrow	<i>SimpleTag</i> <i>BindingTag</i>
<i>SimpleTag</i>	\Rightarrow	< [<i>Netname</i> ::] <i>TagName</i> >
<i>BindingTag</i>	\Rightarrow	< # [<i>Netname</i> ::] <i>TagName</i> >
<i>Connect</i>	\Rightarrow	connect <i>SNetExpr</i> ;
<i>SNetExpr</i>	\Rightarrow	<i>BoxName</i> [<i>Netname</i> ::] <i>NetName</i> <i>Sync</i> <i>Filter</i> <i>Translator</i> <i>Combination</i>
<i>Sync</i>	\Rightarrow	[<i>Pattern</i> [, <i>Pattern</i>]+]
<i>Pattern</i>	\Rightarrow	<i>RecordType</i>
<i>Filter</i>	\Rightarrow	[<i>Pattern</i> \rightarrow [<i>FilterOut</i> [; <i>FilterOut</i>]*]] []
<i>FilterOut</i>	\Rightarrow	{ [<i>FilterOutField</i> [, <i>FilterOutField</i>]*] }
<i>FilterOutField</i>	\Rightarrow	<i>Field</i> [<- <i>FieldInit</i>] <i>Tag</i> [<- <i>TagInit</i>]
<i>FieldInit</i>	\Rightarrow	<i>Field</i>
<i>TagInit</i>	\Rightarrow	<i>Tag</i> <i>IntegerConst</i>

Figure 2.19: Grammar of S-NET_{final}
(continued on next page)

<i>Translator</i>	\Rightarrow	$[\textit{BoxType} \rightarrow \textit{BoxType}]$
<i>Combination</i>	\Rightarrow	$\textit{Serial} \mid \textit{Star} \mid \textit{Choice} \mid \textit{Split}$
<i>Serial</i>	\Rightarrow	$\textit{NetName} \textit{SerialCombinator} \textit{NetName}$
<i>Star</i>	\Rightarrow	$\textit{NetName} \textit{StarCombinator} \textit{Terminator}$
<i>Terminator</i>	\Rightarrow	$(\textit{Pattern} [, \textit{Pattern}]^*)$
<i>Choice</i>	\Rightarrow	$\textit{NetName} \textit{ChoiceCombinator} \textit{NetName}$ \mid $\textit{ChoiceCombinator} \textit{NetName} \textit{NetName} [\textit{NetName}]^+$
<i>Split</i>	\Rightarrow	$\textit{NetName} \textit{SplitCombinator} \textit{Range}$
<i>Range</i>	\Rightarrow	$(\textit{Tag} [: \textit{Tag}])$
<i>SerialCombinator</i>	\Rightarrow	\dots
<i>StarCombinator</i>	\Rightarrow	$* \mid **$
<i>ChoiceCombinator</i>	\Rightarrow	$\mid \mid \mid \mid$
<i>SplitCombinator</i>	\Rightarrow	$! \mid !!$

Figure 2.19: Grammar of S-NET_{final}
 (continued from previous page)

Chapter 3

Code Generation

3.1 Overview

The code generation is performed after completing the postprocessing. To exemplify the process, one entity from the running example (see Fig. 2.18) is translated for each combinator. The generated code contains one function for each entity in the final S-NET code. The combinators and components are provided by a library, i.e. the compiler produces calls to library functions. All entity specific information, such as type signatures (see section 3.2), are created by the compiler and then passed to the library function as parameters. To separate the S-NET namespace, all compiled functions are prefixed with `SNet_` by the compiler.

3.2 Type representation

The symbolic names of fields, tags and binding tags in S-NET code are represented by integers in the generated code. The runtime library provides functions to create type representations. The following shows their signatures and expected parameters. A concrete example for an encoding of the type $\{\{A,B,<T>\},\{D,<\#BT>\}\}$ exemplifies the usage.

Library Function

```
snet_vector_t *SnetTencCreateVector( int A, ...);
```

A number of entries
... entries of the vector

```
snet_variantencoding_t *SnetTencVariant( snet_vector_t *A,  
                                          snet_vector_t *B,  
                                          snet_vector_t *C);
```

A vector containing field names
 B vector containing tag names
 C vector containing binding tag names

```
snet_typeencoding_t *SNetTencTypeEncode( int A, ...);
```

A number of variants
 ... variants of the type

Code example

```
#define A 1
#define B 2
#define D 3
#define T 4
#define BT 5
snet_typeencoding_t *example;

example = SNetTencTypeEncode( 2,
    SNetTencVariantEncode(
        SNetTencCreateVector( 2, A, B),
        SNetTencCreateVector( 1, T),
        SNetTencCreateVector( 0)),
    SNetTencVariantEncode(
        SNetTencCreateVector( 1, D),
        SNetTencCreateVector( 0),
        SNetTencCreateVector( 1, BT)));
```

3.3 Header File

A header file that contains the mapping of tag-, field- and binding tag names to integers is created during the code generation process (see Fig.1.3). The file also contains external declarations of functions for external networks and their fields and tags. External names of fields, tags and binding tags are prefixed to distinguish them from local names. For the running example (Fig. 2.18), the file contains the following.

```
#define A 1
#define B 2
#define C 3
#define D 4
#define P 5
#define Q 6
#define T 7
#define Y 8
#define Z 9
#define tag__A 10
#define tag__B 11
#define tag__T 12

external SNet__tag( snet_buffer_t *inbuf);
```

3.4 Box

S-Net Code Example

```
box compA( {A} -> {P});
```

Generated Code

```
void SNet__compA( snet_handle_t *hnd) {
    snet_record_t *rec;
    void *field_A;
    int tag_Y;

    rec = SNetHndGetRecord( hnd);

    field_A = SNetRecTakeField( rec, A);
    tag_Y = SNetRecTakeTag( rec, Y);

    compA( hnd, field_A, tag_Y);
}
```

3.5 Box Wrapper

S-Net Code Example

```
net compABC__SL( {A} -> {P})
connect compA;
```

Library Function

```
snet_buffer_t *SNetBox( snet_buffer_t *A, snet_buffer_t *B,
                       snet_typeencoding_t *C)
```

- A buffer for incoming records
- B wrapper function
- C type of output

Generated Code

```
snet_buffer_t *SNet__compABC__SL( snet_buffer_t *in_buf) {
    snet_buffer_t *out_buf;
    snet_typeencoding *out_type;

    out_type = SNetTencTypeEncode(
        SNetTencVariantEncode(
            SNetTencCreateVector( 1, P),
            SNetTencCreateVector( 0),
            SNetTencCreateVector( 0)));
```

```

    outbuf = SNetBox( inbuf, &SNet__compA, out_type);

    return( out_buf);
}

```

3.6 Serial Combinator

S-Net Code Example

```

net tag__in ( {A,B} -> {tag::A, tag::B, <tag::T>}
connect tag__translate_in .. tag__tag;

```

Library Function

```

snet_buffer_t *SNetSerial( snet_buffer_t *A, snet_buffer_t *B,
                           snet_buffer_t *C)

```

- A buffer for incoming records
- B component to be connected
- C component to be connected

Generated Code

```

snet_buffer_t *SNet__tag__in( snet_buffer_t *in_buf) {
    snet_buffer_t *out_buf;

    outbuf = SNetSerial( in_buf, &SNet__tag__translate_in, &SNet__tag__tag);

    return( out_buf);
}

```

3.7 Parallel Combinator

S-Net Code Example

```

net example__compute__compAB_compABC ( {A} -> {P} | {B} -> {Q})
connect example__compute__compAB__compABC___SL ||
        example__compute__compAB__compABC___SR__SL;

```

Library Function

```

snet_buffer_t *SNetParallel( snet_buffer_t *A, snet_buffer_t *B ,
                             snet_buffer_t *C, snet_typeencoding_t *D)

```

- A buffer for incoming records
- B component to be connected
- C component to be connected
- D type containing two variants, first variant encodes input type of component A, second variant encodes input type of component B

Generated Code

```
snet_buffer_t *SNet__example__compute__compAB__compABC( buffer_t *in_buf) {
    snet_buffer_t *out_buf;

    out_buf = SNetParallel( in_buf, &SNet__example__compute__compAB__compABC___SL,
                            &SNet__example__compute__compAB__compABC___SR__SL,
                            SNetTencTypeEncode( 2,
                                                  SNetTencVariantEncode(
                                                    SNetTencCreateVector( 1, A),
                                                    SNetTencCreateVector( 0),
                                                    SNetTencCreateVector( 0)),
                                                  SNetTencVariantEncode(
                                                    SNetTencCreateVector( 1, B),
                                                    SNetTencCreateVector( 0),
                                                    SNetTencCreateVector( 0))));

    return( out_buf);
}
```

3.8 Split Combinator

S-Net Code example

```
net example__compute ( {A,<T>} -> {P,Q}, {B,<T>} -> {P,Q})
connect example__compute___IS !! (<T>);
```

Library Function

```
SnetSplit( snet_buffer_t *A, snet_buffer_t *B, int C, int D)
```

- A buffer for incoming records
- B operand component
- C name of tag containing lower value
- D name of tag containing upper value

Generated Code

```
snet_buffer_t *SNet__example__compute( buffer_t *in_buf) {
    snet_buffer_t *out_buf;
```

```

    return( out_buf);
}

```

3.10 Syncro Cell

S-Net Code

```

net example__compute__syncPQ__ST ({P} -> {P} | {P,Q}, {Q} -> {Q} | {P,Q})
connect [| {P}, {Q}|];

```

Library Function

```

SNetSync( snet_buffer_t *A, snet_typeencoding_t *B, snet_typeencoding_t *C)

```

- A buffer for incoming records
- B output type in case of synchronisation
- C patterns to be merged, one pattern per variant

Generated Code

```

snet_buffer_t *SNet__example__compute__syncPQ__ST( buffer_t *in_buf) {
snet_buffer_t *out_buf;

    out_buf = SNetSync( in_buf, SNetTencTypeEncode( 1,
        SNetTencVariantEncode(
            SNetCreateVector( 2, P, Q),
            SNetCreateVector( 0),
            SNetCreateVector( 0))),
        SNetTencTypeEncode( 2,
            SNetTencVariantEncode(
                SNetTencCreateVector( 1, P),
                SNetTencCreateVector( 0),
                SNetTencCreateVector( 0)),
            SNetTencVariantEncode(
                SNetTencCreateVector( 1, Q),
                SNetTencCreateVector( 0),
                SNetTencCreateVector( 0))));

    return( out_buf);
}

```

3.11 Filter

S-Net Code Example

```

net example__split ({A,B,<T>} -> {A,<T>} | {B,<T>})
connect [|{A,B,<T>} -> {A,<T>}; {B,<T>}];

```

Library Function

```
SNetFilter( snet_buffer_t *A, snet_typeencoding_t *B,
           snet_typeencoding_t *C, ...)
```

A buffer for incoming records
 B type of incoming records
 C output type
 ... one instruction set for each variant of the output type,
 all available instructions are listed in the table below

Generated Code

```
snet_buffer_t *SNet__example__split( snet_buffer_t *in_buf) {
snet_buffer_t *out_buf;

    out_buf = SNetFilter( in_buf, SNetTencTypeEncode( 1,
                                                    SNetTencVariantEncode(
                                                        SNetTencCreateVector( 2, A, B),
                                                        SNetTencCreateVector( 1, T),
                                                        SNetTencCreateVector( 0))),
                        SNetTencTypeEncode( 2,
                                                    SNetTencVariantEncode(
                                                        SNetTencCreateVector( 1, A),
                                                        SNetTencCreateVector( 1, T),
                                                        SNetTencCreateVector( 0)),
                                                    SNetTencVariantEncode(
                                                        SNetTencCreateVector( 1, B),
                                                        SNetTencCreateVector( 1, T),
                                                        SNetTencCreateVector( 0))),
                        SNetCreateFilterInstructionSet( 2,
                                                    SNetCreateFilterInstruction( FLT_use_field, A),
                                                    SNetCreateFilterInstruction( FLT_use_tag, T)),
                        SNetCreateFilterInstructionSet( 2,
                                                    SNetCreateFilterInstruction( FLT_use_field, B),
                                                    SNetCreateFilterInstruction( FLT_use_tag, T)));

    return( out_buf);
}
```

Filter Instructions

opcode	parameter 1	parameter 2
FLT_strip_tag	tag name	—
FLT_strip_field	field name	—
FLT_add_tag	tag name	—
FLT_set_tag	tag name	integer value
FLT_rename_tag	old tag name	new tag name
FLT_rename_field	old field name	new field name
FLT_copy_field	name of source field	name of destination field
FLT_use_tag	name of tag	—
FLT_use_field	name of field, note that the field is <i>not</i> copied	—

Chapter 4

Deployment

Chapter 5

Runtime System

Bibliography

- [1] Shafarenko, A., Scholz, S., Grelck, C.: Streaming networks for coordinating data-parallel programs. In Virbitskaite, I., Voronkov, A., eds.: Perspectives of System Informatics, 6th International Andrei Ershov Memorial Conference (PSI'06), Novosibirsk, Russia. Volume 4378 of Lecture Notes in Computer Science., Springer-Verlag, Berlin, Heidelberg, New York (2006) 441–445
- [2] Grelck, C., Shafarenko, A.: Report on S-Net: A Typed Stream Processing Language, Part I: Foundations, Record Types and Networks. Technical report, University of Hertfordshire, Department of Computer Science, Compiler Technology and Computer Architecture Group, Hatfield, England, United Kingdom (2006)
- [3] Institute of Electrical and Electronic Engineers, Inc.: Information Technology — Portable Operating Systems Interface (POSIX) — Part: System Application Program Interface (API) — Amendment 2: Threads Extension [C Language]. IEEE Standard 1003.1c–1995, IEEE, New York City, New York, USA (1995) also ISO/IEC 9945-1:1990b.
- [4] Jesshope, C.: *μtc* — an intermediate language for programming chip multiprocessors. In: Proceedings of the 11th Asia-Pacific Computer Systems Architecture Conference (ACSAC'06), Shanghai, China. (2006)