SRVN Input File Format

$Revision : 2.1$

Dorina C. Petriu R. Gregory Franks Alex Hubbard

November 24, 1998

SRVN Input File Format

The Stochastic Rendezvous Network (SRVN) is a modelling method under development, and is currently being upgraded and extendend. Due to this fact, the input file format was chosen to be as flexible as possible, and easy to be expand when the need for new input data occurs.

The SRVN input file has four sections, each consisting of a variable length list, possibly with other lists nested therein. The *end of list* marker is -1.

The input file may be annotated with comments. Comments start with the octothorpe symbol, ⇐ 3
'`#`', and finish at the end of the line. Comments may appear anywhere except within strings.

General information section gives input parameters related to the numerical computation. It consists of five fields:

1. the convergence value,
2. the underrelaxation coefficient,
3. the maximum number of iterations,
4. the number of iterations after which intermediate results are printed,
5. and a field for *comments on model* which is not interpreted by the analysis program.

Processor information section declares the total number of processors in the system, and the corresponding scheduling policy for each processor. The scheduling policy may be First Come First Served (FCFS) (the default value), Priority Preemptive Resume (PPR) or Head Of the Line (HOL)[1]. The processors are named using symbolic identifiers.

Task information section declares the total number of tasks in the system, and the following parameters for each task: task type (reference or non-reference), the list of entries belonging

[1] Also known as Priority, Non-preemptive.

to this task, the processor on which the task is allocated, and task priority (if it applies). The tasks are named using symbolic identifiers. Each task must have at least one entry because the workload parameters are associated with entries. As a particular case, a reference task has one and only one entry, which is never called by other tasks (entries).

Entry information section declares the total number of entries and the workload parameters associated with each. The entries are named using symbolic identifiers. Each entry is associated to a task by way of an *entry list* for the task. (The model description is invalid if an entry is associated to more than one task, or no task at all). In the BNF description below, the non-terminals *from_entry* and *to_entry* must belong to separate tasks.

The entry workload is divided in phases. The input file format does not limit the number of phases: it uses a variable length list with parameters ordered by phase. Default values are assumed for all undeclared parameters. The parameters declared for entries are:

- open arrival rate to entry (default = 0);
- list of mean service times per phase (default = 0);
- list of phase type flags (default = 0):
  1. exponentially distributed service time and geometrically distributed number of calls;
  2. deterministic service time and deterministic number of calls;
  3. exponential service time and deterministic number of calls;
- coefficient of variation of the service time slice between two calls, given by phase (default = 1.0);
- mean number of RNVs per phase from this entry to other entries (default = 0.0);
- mean number of nonblocking sends per phase from this entry to other entries (default = 0.0).
- probability of forwarding a rendezvous to another entry (default = 0.0).
- internal messaging between entries.

BNF Description of SRVN Input File Format

This section gives the formal description of SRVN input file grammar in BNF form. For the nonterminals the notation $\langle nonterminal\_id \rangle$ is used, while the terminals are written without brackets as they appear in the input text.

The notation

$$\{\cdots\}_n^m, \text{ where } n \leq m$$

means that the part inside the curly brackets is repeated at least n times and at most m times. If $n = 0$, then the part may be missing in the input text.

General Information

⟨*SRVN_input_file*⟩ → ⟨*general_info*⟩ ⟨*processor_info*⟩ ⟨*task_info*⟩ ⟨*entry_info*⟩ {⟨*activity_info*⟩}$_0$ ⇐ 9

⟨*general_info*⟩ → `G` ⟨*comment*⟩ ⟨*conv_val*⟩ ⟨*it_limit*⟩ ⟨*print_int*⟩$_{\text{opt}}$ ⟨*underrelax_coeff*⟩$_{\text{opt}}$ ⟨*end_list*⟩ ⇐ 4

⟨*comment*⟩ → ⟨*string*⟩ /* *comment on the model* */

⟨*conv_val*⟩ → ⟨*real*⟩ /* *convergence value* */ ‡

⟨*it_limit*⟩ → ⟨*integer*⟩ /* *max. nb. of iterations* */ ‡

⟨*print_int*⟩ → ⟨*integer*⟩ ‡
/* *intermed. res. print interval* */

⟨*underrelax_coeff*⟩ → ⟨*real*⟩ /* *under_relaxation coefficient* */ ‡

⟨*end_list*⟩ → `-1` /* *end_of_list mark* */

⟨*string*⟩ → `"` ⟨*text*⟩ `"`

Processor Information

⟨*processor_info*⟩ → `P` ⟨*np*⟩ ⟨*p_decl_list*⟩

⟨*np*⟩ → ⟨*integer*⟩ /* *total number of processors* */

⟨*p_decl_list*⟩ → {⟨*p_decl*⟩}$_0^{np}$ ⟨*end_list*⟩

⟨*p_decl*⟩ → `p` ⟨*proc_id*⟩ ⟨*scheduling_flag*⟩ ⟨*quantum*⟩$_{\text{opt}}$ ⟨*multi_server_flag*⟩$_{\text{opt}}$ ⟨*replication_flag*⟩$_{\text{opt}}$ ⟨*proc_rate*⟩$_{\text{opt}}$ ⇐ 1, 7, 8

⟨*proc_id*⟩ → ⟨*integer*⟩ | ⟨*identifier*⟩ ⇐ 2
/* *processor identifier* */

⟨*scheduling_flag*⟩ → `f` /* *First come, first served* */
| `p` /* *Priority, preemeptive* */
| `h` /* *Head Of Line* */
| `r` /* *Random* */ ⇐ 4
| `s` /* *Processor sharing* */

⟨*quantum*⟩ → ⟨*real*⟩

⟨*multi_server_flag*⟩ → `m` ⟨*copies*⟩ /* *number of duplicates* */ ⇐ 1
| `i` /* *Infinite server* */ ⇐ 6

⟨*replication_flag*⟩ → `r` ⟨*copies*⟩ /* *number of replicas* */ ⇐ 8

⟨*proc_rate*⟩ → `R` ⟨*ratio*⟩ /* *Relative proc. speed* */ ⇐ 7, 8

⟨*copies*⟩ → ⟨*integer*⟩

⟨*ratio*⟩ → ⟨*real*⟩ ⇐ 7

Task Information

⟨*task_info*⟩	→	T ⟨*nt*⟩ ⟨*t_decl_list*⟩	
⟨*nt*⟩	→	⟨*integer*⟩ /* *total number of tasks* */	
⟨*t_decl_list*⟩	→	$\{\langle t\_decl \rangle\}_0^{nt}$ ⟨*end_list*⟩	
⟨*t_decl*⟩	→	t ⟨*task_id*⟩ ⟨*task_sched_type*⟩ ⟨*entry_list*⟩ ⟨*proc_id*⟩ $\langle task\_pri \rangle_{\text{opt}}$ $\langle think\_time\_flag \rangle_{\text{opt}}$ $\langle multi\_server\_flag \rangle_{\text{opt}}$ $\langle replication\_flag \rangle_{\text{opt}}$	⇐ 1, 3, 8 †
⟨*t_decl*⟩	→	t ⟨*task_id*⟩ i ⟨*entry_list*⟩ $\langle replication\_flag \rangle_{\text{opt}}$	⇐ 9
⟨*task_id*⟩	→	⟨*integer*⟩ | ⟨*identifier*⟩ /* *task identifier* */	⇐ 2
⟨*task_sched_type*⟩	→	r /* *reference task* */	
	|	n /* *non-reference task* */	
	|	h /* *Head of line* */	⇐ 9
	|	f /* *FIFO Scheduling* */	⇐ 9
	|	p /* *Polled scheduling at entries* */	⇐ 10
	|	b /* *Bursty Reference task* */	⇐ 10
⟨*entry_list*⟩	→	$\{\langle entry\_id \rangle\}_1^{ne_t}$ ⟨*end_list*⟩ /* *task t has* ne_t *entries* */	
⟨*entry_id*⟩	→	⟨*integer*⟩ | ⟨*identifier*⟩ /* *entry identifier* */	⇐ 2
⟨*task_pri*⟩	→	⟨*integer*⟩ /* *task priority, optional* */	

Entry Information

⟨*entry_info*⟩	→	E ⟨*ne*⟩ ⟨*entry_decl_list*⟩	
⟨*ne*⟩	→	⟨*integer*⟩ /* *total number of entries* */	
⟨*entry_decl_list*⟩	→	$\{\langle entry\_decl \rangle\}_1^{ne \times 9}$ ⟨*end_list*⟩ /* k = *maximum nb of phases* */	
⟨*entry_decl*⟩	→	a ⟨*entry_id*⟩ ⟨*arrival_rate*⟩	
	|	c ⟨*entry_id*⟩ $\{\langle coeff\_of\_variation \rangle\}_1^k$ ⟨*end_list*⟩	
	|	f ⟨*entry_id*⟩ $\{\langle ph\_type\_flag \rangle\}_1^k$ ⟨*end_list*⟩	
	|	i ⟨*from_entry*⟩ ⟨*to_entry*⟩ $\{\langle fan\_in \rangle\}_1^k$ ⟨*end_list*⟩	⇐ 8
	|	o ⟨*from_entry*⟩ ⟨*to_entry*⟩ $\{\langle fan\_out \rangle\}_1^k$ ⟨*end_list*⟩	⇐ 8
	|	p ⟨*entry_id*⟩ ⟨*entry_priority*⟩	⇐ 9
	|	s ⟨*entry_id*⟩ $\{\langle service\_time \rangle\}_1^k$ ⟨*end_list*⟩	
	|	y ⟨*from_entry*⟩ ⟨*to_entry*⟩ $\{\langle rendezvous \rangle\}_1^k$ ⟨*end_list*⟩	
	|	z ⟨*from_entry*⟩ ⟨*to_entry*⟩ $\{\langle send\_no\_reply \rangle\}_1^k$ ⟨*end_list*⟩	
	|	A ⟨*entry_id*⟩ ⟨*activity_id*⟩	⇐ 9, †
	|	M ⟨*entry_id*⟩ $\{\langle max\_service\_time \rangle\}_1^k$ ⟨*end_list*⟩	⇐ 10
	|	F ⟨*from_entry*⟩ ⟨*to_entry*⟩ ⟨*p_forward*⟩ ⟨*end_list*⟩	⇐ 1

	|	`Z` ⟨*entry_id*⟩ $\{\langle think\_time \rangle\}_1^k$ ⟨*end_list*⟩	
⟨*arrival_rate*⟩	→	⟨*real*⟩ /* *open arrival rate to entry* */	
⟨*service_time*⟩	→	⟨*real*⟩ /* *mean phase service time* */	
⟨*ph_type_flag*⟩	→	⟨*integer*⟩ /* *0 – stochastic phase* */	
		/* *1 – deterministic phase* */	
⟨*coeff_of_variation*⟩	→	⟨*real*⟩ /* *squared service time coefficient of variation* */	
⟨*fan_in*⟩	→	⟨*integer*⟩ /* *fan in to this entry* */	⇐ 8
⟨*fan_out*⟩	→	⟨*integer*⟩ /* *fan out of this entry* */	⇐ 8
⟨*rendezvous*⟩	→	⟨*real*⟩ /* *mean number of RNVs/ph* */	
⟨*send_no_reply*⟩	→	⟨*real*⟩ /* *mean nb.of non-blck.sends/ph* */	
⟨*p_forward*⟩	→	⟨*real*⟩ /* *probability of forwarding* */	⇐ 1
⟨*rate*⟩	→	⟨*real*⟩ /* *nb. of calls per arrival* */	⇐ 4
⟨*max_service_time*⟩	→	⟨*real*⟩ /* *Max time to flag.* */	⇐ 10
⟨*think_time*⟩	→	⟨*real*⟩ /* *Think time for phase.* */	⇐ 9
⟨*from_entry*⟩	→	⟨*entry_id*⟩ /* *Source of a message* */	
⟨*to_entry*⟩	→	⟨*entry_id*⟩ /* *Destination of a message* */	⇐ 4

Activity Information

⟨*activity_info*⟩	→	⟨*activity_defn_list*⟩ $\langle activity\_connections \rangle_{\text{opt}}$ ⟨*end_list*⟩	⇐ 9
		/* *Activity definition.* */	
⟨*activity_defn_list*⟩	→	$\{\langle activity\_defn \rangle\}_1^{na}$	
⟨*activity_defn*⟩	→	`s` ⟨*activity_id*⟩ ⟨*ph_serv_time*⟩ /* *Service time* */	
	|	`c` ⟨*activity_id*⟩ ⟨*coeff_of_variation*⟩ /* *Sqr. Coeff. of Var.* */	
	|	`f` ⟨*activity_id*⟩ ⟨*ph_type_flag*⟩ /* *Phase type* */	
	|	`y` ⟨*activity_id*⟩ ⟨*to_entry*⟩ ⟨*rendezvous*⟩ /* *Rendezvous* */	
	|	`z` ⟨*activity_id*⟩ ⟨*to_entry*⟩ ⟨*rendezvous*⟩ /* *Send-no-reply* */	
	|	`Z` ⟨*activity_id*⟩ ⟨*think_time*⟩ /* *Think time* */	
		/* *Activity Connections.* */	
⟨*activity_connections*⟩	→	`:` ⟨*activity_conn_list*⟩	
⟨*activity_conn_list*⟩	→	⟨*activity_conn*⟩ $\{$ `;` ⟨*activity_conn*⟩$\}_1^{na}$	
⟨*activity_conn*⟩	→	⟨*join_list*⟩	
	|	⟨*join_list*⟩ `->` ⟨*split_list*⟩	
	|	⟨*repeat_list*⟩ `->` ⟨*split_list*⟩	
⟨*join_list*⟩	→	⟨*reply_activity*⟩	
	|	⟨*and_join_list*⟩	

	|	⟨*or_join_list*⟩
⟨*split_list*⟩	→	⟨*activity_id*⟩
	|	⟨*repeat_list*⟩
	|	⟨*and_split_list*⟩
	|	⟨*or_split_list*⟩
⟨*and_join_list*⟩	→	⟨*reply_activity*⟩ {`&` ⟨*reply_activity*⟩}$_1^{na}$
⟨*or_join_list*⟩	→	⟨*reply_activity*⟩ {`+` ⟨*reply_activity*⟩}$_1^{na}$
⟨*and_split_list*⟩	→	⟨*activity_id*⟩ {`&` ⟨*activity_id*⟩}$_1^{na}$
⟨*or_split_list*⟩	→	⟨*prob_activity*⟩ {`+` ⟨*prob_activity*⟩}$_1^{na}$
⟨*repeat_list*⟩	→	⟨*real*⟩ `*` ⟨*sub_activity_id*⟩ ⟨*next_activity_id*⟩$_{\text{opt}}$
⟨*prob_activity*⟩	→	`(` ⟨*real*⟩ `)` ⟨*activity_id*⟩
⟨*reply_activity*⟩	→	⟨*activity_id*⟩ ⟨*reply_list*⟩$_{\text{opt}}$
⟨*sub_activity_id*⟩	→	⟨*activity_id*⟩ /* *head of nested sub-list* */
⟨*next_activity_id*⟩	→	`,` ⟨*activity_id*⟩ /* *next activity after sublist* */
⟨*reply_list*⟩	→	`[` ⟨*entry_id*⟩ {`,` ⟨*entry_id*⟩ }$_0^{ne}$ `]`

Identifiers

Identifiers may be zero or more leading underscores ('_'), followed by a character, followed by any number of characters, numbers or underscores. Punctuation characters and other special characters such as the dollar-sign ('$') are not permitted. The following, `_p1`, `foo_bar`, and `__P_21_proc` are valid identifiers, while `_21` and `$proc` are not. ⇐ 2

Other Solvers

Parasol

The SRVN input file is also used by the SRVN Parasol simulator. Certain optional nonterminals only apply to the simulator (for example, *cpu_quantum*). These nonterminals are identified by a '†'. If they are present as input to the SRVN solver, they are ignored. Mandatory fields marked by a '‡' are ignored by the simulator. ⇐ 3

PetriSRVN

Mandatory fields marked by a '‡' are ignored by the Petri net solver.

Layered Queueing Network Solver

Bounds Solver

Summary of Changes

1. Multiservers have been added as an optional field for tasks and processors. This change has not been implemented.

2. Symbolic names have been added for tasks, processors and entries. This change has been incorporated into the source.

3. Extra nonterminals for Parasol input. Comments in the input file.

4. Items for entry "queues". "Random" scheduling. Optional specification of underrelaxation and print interval. The latter parameters default to 0.5 for underrelaxation, and 0 for print interval.

5. Queues can have an initial number of tokens.

6. Infinite server capabilty. Semantic analysis will catch bogus infinite server combinations (for example, reference tasks).

7. Added relative processor speed information for processor information.

8. Added replication information for processors, tasks, and entries. Note that the processor rate flag has changed from '`r`' to '`R`'.

9. Priorities. Activities

10. Specialized scheduling types and information. Not supported by most solvers.