

Multi-BSP Programming à la ML

MBSML

Victor Allombert

-

Frédéric Gava - Julien Tesson


Laboratory of **A**lgorithms, **C**omplexity and **L**ogic (LACL)
Université Paris Est

Outline

- 1 Introduction
- 2 BSML: Functional BSP Programming
- 3 Hierarchical machines : Multi-BSP
- 4 Conclusion

Parallel Architectures


Distributed computing



Clusters


Parallel Architectures

Distributed computing



Clusters


Shared memory



Multi-core


Parallel Architectures

Distributed computing




Clusters

Shared memory



Multi-core

Hybrid model



(Parallel) Software Errors

Risks

- Over-consumption
- Erroneous results

Typical bugs

Distributed

Shared memory

Deadlocks:

(Parallel) Software Errors

Risks

- Over-consumption
- Erroneous results

Typical bugs

Distributed

Shared memory

Deadlocks:

(Parallel) Software Errors


Risks

- Over-consumption
- Erroneous results

Typical bugs

Deadlocks:

Distributed



Shared memory

(Parallel) Software Errors


Risks

- Over-consumption
- Erroneous results


Typical bugs

Deadlocks:

Distributed



Shared memory



(Parallel) Software Errors


Risks

- Over-consumption
- Erroneous results


Typical bugs

Deadlocks:


Distributed



Shared memory



Race condition:



(Parallel) Software Errors


Risks

- Over-consumption
- Erroneous results


Typical bugs

Deadlocks:



Distributed



Shared memory



Race condition:



(Parallel) Software Errors



Risks

- Over-consumption
- Erroneous results

Considered solutions

- 1 Well **structured** parallelism
- 2 Design a **high-level** language for “**hybrid architectures**”
- 3 Software-hardware **bridging model** \Rightarrow Portability, scalability

Race condition:



Outline

- 1 Introduction
- 2 BSML: Functional BSP Programming**
- 3 Hierarchical machines : Multi-BSP
- 4 Conclusion

Bridging Model: Bulk Synchronous Parallelism (BSP)

The BSP computer

Defined by:

- **p** pairs CPU/memory
- Communication network (**g**)
- Synchronisation unit (**L**)
- Super-steps execution

Properties

Bridging Model: Bulk Synchronous Parallelism (BSP)

The BSP computer

Defined by:

- p pairs CPU/memory
- Communication network (g)
- Synchronisation unit (L)
- Super-steps execution

Properties:

Bridging Model: Bulk Synchronous Parallelism (BSP)

The BSP computer

Defined by:

- **p** pairs CPU/memory
- **Communication** network (**g**)
- **Synchronisation** unit (**L**)
- Super-steps execution

Properties:

- "Confluent"

Bridging Model: Bulk Synchronous Parallelism (BSP)


The BSP computer

Defined by:

- p pairs CPU/memory
- Communication network (g)
- Synchronisation unit (L)
- Super-steps execution

Properties:

- "Confluent"
- "Deadlock-free"
- Predictable performances



local
computations

communication ($\otimes g$)

barrier ($\oplus L$)

next super-step

Bridging Model: Bulk Synchronous Parallelism (BSP)


The BSP computer

Defined by:

- p pairs CPU/memory
- Communication network (g)
- Synchronisation unit (L)
- Super-steps execution

Properties:

- “Confluent”
- “Deadlock-free”
- Predictable performances



local
computations

communication ($\otimes g$)

barrier ($\oplus L$)

next super-step

Bridging Model: Bulk Synchronous Parallelism (BSP)


The BSP computer

Defined by:

- p pairs CPU/memory
- Communication network (g)
- Synchronisation unit (L)
- Super-steps execution

Properties:

- “Confluent”
- “Deadlock-free”
- Predictable performances



local
computations

communication ($\otimes g$)

barrier ($\oplus L$)

next super-step

Bridging Model: Bulk Synchronous Parallelism (BSP)


The BSP computer

Defined by:

- p pairs CPU/memory
- Communication network (g)
- Synchronisation unit (L)
- Super-steps execution

Properties:

- “Confluent”
- “Deadlock-free”
- Predictable performances



local
computations


communication ($\otimes g$)

barrier ($\oplus L$)

next super-step


Examples: broadcasting a values

Direct broadcast (one super-step)



$$\text{Cost} \equiv \mathbf{p} \times \mathbf{g} \times n + \mathbf{L}$$

Broadcast with two super-steps



$$\text{Cost} \equiv 2 \times \mathbf{g} \times n + 2 \times \mathbf{L}$$

The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) → **Coq**

Main Idea

The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) → **Coq**

Main idea

The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) → **Coq**

Main idea

The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) → **Coq**

Main idea

- $(\nu_0, \dots, \nu_{p-1}) \text{ co par} \equiv \nu_j$ on node i
- Four primitives \Rightarrow simple semantics


The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) \rightarrow **Coq**

Main idea

Parallel data structure \Rightarrow vectors:



- 1 $\langle v_0, \dots, v_{p-1} \rangle : \alpha \text{ par} \equiv v_i$ on node i
- 2 Four primitives \Rightarrow simple semantics


The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) \rightarrow **Coq**

Main idea

Parallel data structure \Rightarrow vectors:



- 1 $\langle v_0, \dots, v_{p-1} \rangle : \alpha \text{ par} \equiv v_i$ on node i
- 2 Four primitives \Rightarrow simple semantics


The BSML Language

BSML

- Explicit **BSP** programming with a **functional** approach
- Based upon **ML**; Implemented over **OCaml**
- **Formal semantics** (confluent) \rightarrow **Coq**

Main idea

Parallel data structure \Rightarrow vectors:



- 1 $\langle v_0, \dots, v_{p-1} \rangle : \alpha \text{ par} \equiv v_i$ on node i
- 2 **Four** primitives \Rightarrow simple semantics

The BSML Primitives

Asynchronous operations

- `⟨⟨ ... ⟩⟩`: local execution (vector)
- `v`: element of a parallel vector `v`
- `pid`: id of the processor

The BSML Primitives

Asynchronous operations

- `<< ... >>`: local execution (**vector**)
- `v`: element of a parallel **vector** `v`
- `pid`: id of the processor

Communications

The BSML Primitives

Asynchronous operations

- `⟨⟨ ... ⟩⟩`: local execution (**vector**)
- `v`: element of a parallel **vector** `v`
- `pid`: id of the processor

Communication

The BSML Primitives

Asynchronous operations

- $\langle\langle \dots \rangle\rangle$: local execution (vector)
- $\$v\$$: element of a parallel vector v
- $\$pid\$$: id of the processor

Communication

- $proj: (x_0, \dots, x_{p-1}) \mapsto$

 x_{p-1} x_0 $x_{(p-1)}$ $x_{(p-1)}$

The BSML Primitives

Asynchronous operations

- $\langle\langle \dots \rangle\rangle$: local execution (**vector**)
- $\$v\$$: element of a parallel **vector** v
- $\$pid\$$: id of the processor

Communication

- **proj**: $\langle x_0, \dots, x_{p-1} \rangle \mapsto$

$$\begin{matrix} x_0 \\ \vdots \\ x_{p-1} \end{matrix}$$
- **put**: $\langle f_0, \dots, f_{p-1} \rangle \mapsto$

$$\left\langle \begin{matrix} f_0 0 & & f_0 (p-1) \\ \vdots & \dots & \vdots \\ f_{p-1} 0 & & f_{p-1} (p-1) \end{matrix} \right\rangle$$

The BSML Primitives

Asynchronous operations

- $\langle\langle \dots \rangle\rangle$: local execution (**vector**)
- $\$v\$$: element of a parallel **vector** v
- $\$pid\$$: id of the processor

Communication

- **proj**: $\langle x_0, \dots, x_{p-1} \rangle \mapsto$

$$\begin{matrix} x_0 \\ \vdots \\ x_{p-1} \end{matrix}$$
- **put**: $\langle f_0, \dots, f_{p-1} \rangle \mapsto$

$$\left\langle \begin{matrix} f_0 0 & & f_0 (p-1) \\ \vdots & & \vdots \\ f_{p-1} 0 & , \dots , & f_{p-1} (p-1) \end{matrix} \right\rangle$$

The BSML Primitives


Asynchronous operations

- $\langle\langle \dots \rangle\rangle$: local execution (**vector**)
- $\$v\$$: element of a parallel **vector** v
- $\$pid\$$: id of the processor

Communication

- **proj**: $\langle x_0, \dots, x_{p-1} \rangle \mapsto \begin{matrix} x_0 \\ \vdots \\ x_{p-1} \end{matrix}$
- **put**: $\langle f_0, \dots, f_{p-1} \rangle \mapsto \left\langle \begin{matrix} f_0 & 0 & & & f_0 (p-1) \\ \vdots & & & & \vdots \\ f_{p-1} & 0 & & & f_{p-1} (p-1) \end{matrix} \right\rangle, \dots$

Parallel Sorting by Regular Sampling (PSRS)



Example : PSRS

(psrs: int par \rightarrow 'a list \rightarrow 'a list *)*

let psrs lvlengths lv =

(super-step 1(a): local sorting *)*

let locsort = \ll List.sort compare \$lv\$ \gg **in**

(super-step 1(b): selection of the primary samples *)*

let regsampl = \ll extract_n **P** \$lvlengths\$ \$locsort\$ \gg **in**

(super-step 2(a): total exchange of the primary samples;*)*

let glosampl = List.sort compare (**proj** regsampl) **in**

(super-step 2(b): selection of the secondary samples *)*

let pivots = extract_n **P** (**P***(**P**-1)) glosampl **in**

(super-step 2(c) : building the communicated lists of values *)*

let comm = \ll slice_p \$locsort\$ pivots \gg **in**

(super-step 3: sended them and merging of the received values *)*

let recv = **put** \ll List.nth \$comm\$ \gg **in**

\ll p_merge **P** (List.map \$recv\$ procs_list) \gg

Advantages and Drawbacks

Advantages

- Easy to learn
- “All” OCaml codes can be used
- Easy to get a BSML code from a BSP algorithm
- Pure functional semantics → Coq

Advantages and Drawbacks

Advantages

- Easy to learn
- “All” OCaml codes can be used
- Easy to **get** a BSML code from a BSP algorithm
- Pure functional semantics → **Coq**

Drawbacks

- **Complex** type system to forbid
 - Nested vectors (**let** $v = \ll \dots \gg$ **in** $\ll v \gg$)
 - Replicated **inconsistency** (**if** `rnd()` **then** `sync();0` **else** `0`)
 - **Data-race** using side effects (\ll **if** `pid=0` **then** `v:=2` \gg)
- Hierarchical architecture as a **flat** one
- **Congestion** using network


Outline

- 1 Introduction
- 2 BSML: Functional BSP Programming
- 3 Hierarchical machines : Multi-BSP**
- 4 Conclusion

Multi-BSP Model (1)

What is Multi-BSP? (Valiant)

- 1 A **tree** structure with **nested** components
- 2 Where **nodes** have a storage capacity
- 3 And **leaf** are homogenous processors




Multi-BSP Model (1)

What is Multi-BSP? (Valiant)

- 1 A **tree** structure with **nested** components
- 2 Where **nodes** have a storage capacity
- 3 And **leaf** are homogenous processors

Multi-BSP



BSP




Multi-BSP Model (1)

What is **Multi-BSP**? (Valiant)

- 1 A **tree** structure with **nested** components
- 2 Where **nodes** have a storage capacity
- 3 And **leaf** are homogenous processors

Multi-BSP



BSP




Multi-BSP Model (1)

What is **Multi-BSP**? (Valiant)

- 1 A **tree** structure with **nested** components
- 2 Where **nodes** have a storage capacity
- 3 And **leaf** are homogenous processors

Multi-BSP



BSP




Multi-BSP Model (1)


What is Multi-BSP? (Valiant)

- 1 A **tree** structure with **nested** components
- 2 Where **nodes** have a storage capacity
- 3 And **leaf** are homogenous processors

Multi-BSP



BSP




Multi-BSP Model (1)


What is Multi-BSP? (Valiant)

- 1 A **tree** structure with **nested** components
- 2 Where **nodes** have a storage capacity
- 3 And **leaf** are homogenous processors

Multi-BSP



BSP



Multi-BSP Model (2)

Cost model

A d depth tree is specified by $4 \times d$ parameters:

- p : Number of sub-components
- m : Available memory at a level
- g : Bandwidth with the upper level
- L : Synchronisation

Multi-BSP Model (2)

Cost model

A d depth tree is specified by $4 \times d$ parameters:

- p : Number of sub-components
- m : Available memory at a level
- g : Bandwidth with the upper level
- L : Synchronisation

Example: 16 quad-chips with octo-cores

- Level 4 ($p = 16, g = \infty, L = 1000, m = 16 Tb$) (RAM/IO)
- Level 3 ($p = 4, g = 150, L = 100, m = 64 Gb$) (RAM)
- Level 2 ($p = 8, g = 5, L = 10, m = 2 Mb$) (L2 cache)
- Level 1 ($p = 1, g = 1, L = 1, m = 8 Kb$) (L1 cache)

Multi-BSP Model (3)

Execution model

At a level i , a **super-step** is:

- Each **component** at level $i - 1$ does its own **super-steps**
- Then each **copies** some data to the memory at level i
- Then **synchronisation**
- Finally **copy** of some data from level i to $i - 1$

Multi-BSP Model (3)

Execution model

At a level i , a **super-step** is:

- Each **component** at level $i - 1$ does its own **super-steps**
- Then each **copies** some data to the memory at level i
- Then **synchronisation**
- Finally **copy** of some data from level i to $i - 1$

Advantages and drawbacks

- Implicit **subgroup** synchronisation
- **Recursive** decomposition of problems
- **Harder** to **design/cost** some algorithms

Multi-BSML Language (1)


Syntactic construction

```
let multi f [args] =  
  let cst = CodeOCaml  
  where node [args] = CodeBSML ...  
    << f args >>  
    ... CodeBSML  
  where leaf [args] = CodeOCaml  
in f ...
```

Multi-BSML Language (1)

Syntactic construction


```
let multi f [args] =  
  let cst = CodeOCaml  
  where node [args] = CodeBSML ...  
    << f args >>  
    ... CodeBSML  
  where leaf [args] = CodeOCaml  
in f ...
```



Multi-BSML Language (1)

Syntactic construction

```
let multi f [args] =  
  let cst = CodeOCaml  
  where node [args] = CodeBSML ...  
    << f args >>  
    ... CodeBSML  
  where leaf [args] = CodeOCaml  
  in f ...
```



Limitations and differences

- Nodes are **implicit** computation units
- **Horizontal** communications between level components
- **Garbage collector** ⇒ no L1, L2 caches.

Multi-BSML Language (2)

Semantics of multi

- BSML code to **distribute** values
- `⟨⟨ ... ⟩⟩` and **proj**; **level** changing
- Mutual **recursive** functions of standard OCaml values
- (**Formal**) Big-steps \iff small-steps for a mini-Multi-BSML


Multi-BSML Language (2)

Semantics of multi

- BSML code to **distribute** values
- $\langle\langle \dots \rangle\rangle$ and **proj**; **level** changing
- Mutual **recursive** functions of standard OCaml values
- (**Formal**) Big-steps \iff small-steps for a mini-Multi-BSML

Copy memory values and distribution of values

- **let** $x = 1$ **in** $\langle\langle \#x\# + 1 \rangle\rangle$
- **mkpar** ($\text{fun } i \rightarrow e$)
 $\mapsto \langle v_0, \dots, v_{p-1} \rangle$
 where $(e \ i) \mapsto v_i$



Example : Sum Of The Elements Of A List

```
let multi sum_list l =  
  
  where node l =  
    let v = mkpar (fun i → split i l) in  
      sumSeq (flatten << sum_list $v$ >>) (* flatten uses a proj *)  
            (* sumSeq is List.fold_left *)  
  
  where leaf l = sumSeq l  
  
in ... (sum_list lst) ...
```


Multi With Tree Construction

Goals

- Keep values on each node and leaf
- To program multiple phases of multi

Multi With Tree Construction

Goals


- Keep values on each node and leaf
- To program multiple phases of multi

Extension

- finally; pushes up a value and keeps a value
- If the recursive calls generates partial trees \Rightarrow to raise exceptions “à la BSML”

Example

Keep the intermediate results of the sum



```
let multi sum_list l =
```

```
where node l =
```

```
  let v = mkpar (fun i → split i l) in
```

```
  let s = sumSeq (flatten << sum_list $v$ >>) in
```

```
  finally ~up:s ~keep:s
```

```
where leaf l =
```

```
  let s = sumSeq l in
```

```
  finally ~up:s ~keep:s
```

Semantics

Useful semantics

- ↓ **Big step semantic (with costs)**
- ↓^{co} Coinductive big step semantic
- Small step semantic with explicit substitutions (and costs)
- ↪ Distributed semantic
- ↪_{⊗_i} Distributed semantic with continuations
- ↪^c “Compiled” semantic

Semantics

Useful semantics

- ↓ Big step semantic (with costs)
- ↓^{co} Coinductive big step semantic
- Small step semantic with explicit substitutions (and costs)
- Distributed semantic
- Distributed semantic with continuations
- “Compiled” semantic

Semantics

Useful semantics

- ↓ Big step semantic (with costs)
- ↓^{co} Coinductive big step semantic
- Small step semantic with explicit substitutions (and costs)
- Distributed semantic
- Distributed semantic with continuations
- “Compiled” semantic

Semantics

Useful semantics

- ↓ Big step semantic (with costs)
- ↓^{co} Coinductive big step semantic
- Small step semantic with explicit substitutions (and costs)
- ↳ Distributed semantic
- ↳_i Distributed semantic with continuations
- ↳^L “Compiled” semantic

Semantics

Useful semantics

- ↓ Big step semantic (with costs)
- ↓^{co} Coinductive big step semantic
- Small step semantic with explicit substitutions (and costs)
- ↳ Distributed semantic
- ↳_i Distributed semantic
- ↳ Distributed semantic with continuations
- ↳^L “Compiled” semantic

Semantics

Useful semantics

- ↓ Big step semantic (with costs)
- ↓^{co} Coinductive big step semantic
- Small step semantic with explicit substitutions (and costs)
- ↳ Distributed semantic
- ↳_i Distributed semantic
- ↳ Distributed semantic with continuations
- ↳^L “Compiled” semantic

Implementation (currently test phase)

Sequential

For **debugging** and **toplevel**

- **tree** structure of data
- A global tree of **Hashtables** to represent the memories

Implementation (currently test phase)

Sequential

For **debugging** and **toplevel**

- **tree** structure of data
- A global tree of **Hashtables** to represent the memories

Distributed

Modular (MPI, TCP/IP, etc.) and based on formal semantics.

Implementation (currently test phase)

Sequential


For **debugging** and **toplevel**

- **tree** structure of data
- A global tree of **Hashtables** to represent the memories

Distributed

Modular (MPI, TCP/IP, etc.) and based on formal semantics.

Shared memory



Implementation (currently test phase)

Sequential


For **debugging** and **toplevel**

- **tree** structure of data
- A global tree of **Hashtables** to represent the memories


Distributed

Modular (MPI, TCP/IP, etc.) and based on formal semantics.

Shared memory



Distributed



Outline

- 1 Introduction
- 2 BSML: Functional BSP Programming
- 3 Hierarchical machines : Multi-BSP
- 4 Conclusion**

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow BSML
- Multi-BSP \Rightarrow Multi-BSML

Multi-BSML

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- Recursive functions on different memories of chips
- Structured nesting of BSML codes

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- **Recursive** functions on different memories of chips
- **Structured** nesting of BSML codes
- Big-steps and small-steps **formal** semantics (confluent)
- A skeleton for Coq
- Small number of primitives and little syntax extension

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- **Recursive** functions on different memories of chips
- **Structured** nesting of BSML codes
- Big-steps and small-steps **formal** semantics (confuent)
- A **skeleton** for Coq
- **Small** number of primitives and **little** syntax extension

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- **Recursive** functions on different memories of chips
- **Structured nesting** of BSML codes
- Big-steps and small-steps **formal** semantics (confuent)
- A **skeleton** for Coq
- **Small** number of primitives and **little** syntax extension

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- **Recursive** functions on different memories of chips
- **Structured nesting** of BSML codes
- Big-steps and small-steps **formal** semantics (confluent)
- A **skeleton** for Coq
- **Small** number of primitives and **little** syntax extension

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- **Recursive** functions on different memories of chips
- **Structured nesting** of BSML codes
- Big-steps and small-steps **formal** semantics (confluent)
- A **skeleton** for Coq
- **Small** number of primitives and **little** syntax extension

Conclusion

Multi

- Multi-BSP extension of BSP for **hierarchical** architectures
- BSP \Rightarrow **BSML**
- Multi-BSP \Rightarrow **Multi-BSML**

Multi-BSML

- **Recursive** functions on different memories of chips
- **Structured nesting** of BSML codes
- Big-steps and small-steps **formal** semantics (confluent)
- A **skeleton** for Coq
- **Small** number of primitives and **little** syntax extension

Perspectives (Ongoing/Future Work)

Short term

- Implementation using MPI
- Examples and benchmarks
- Type system for a subpart of OCaml again

Long term

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- Type system for a subpart of OCaml again

Long term

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - ▶ Bad nesting of parallelism
 - ▶ Bad copy of data

Long term

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad use of the parallel operators

Long term

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

- Mechanized semantics in Coq
- Embed in Coq for proofs and extraction of sequential programs as for

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

- **Mechanized semantics** in Coq
- Embed in Coq for **proofs** and **extraction** of multi-BSML programs as for BSML
- **where kernel** \Rightarrow GPU embed in multi-BSML
- **Examples and libraries**

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

- **Mechanized** semantics in Coq
- Embed in Coq for **proofs** and **extraction** of multi-BSML programs as for BSML
- **where kernel** \Rightarrow GPU embed in multi-BSML
- Examples and **libraries**

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

- **Mechanized** semantics in Coq
- Embed in Coq for **proofs** and **extraction** of multi-BSML programs as for BSML
- **where kernel** \Rightarrow GPU embed in multi-BSML
- Examples and **libraries**

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

- **Mechanized** semantics in Coq
- Embed in Coq for **proofs** and **extraction** of multi-BSML programs as for BSML
- **where kernel** \Rightarrow **GPU** embed in multi-BSML
- Examples and **libraries**

Perspectives (Ongoing/Future Work)

Short term

- **Implementation** using MPI
- Examples and **benchmarks**
- **Type system** for a subpart of OCaml again
 - Bad **nesting** of parallelism
 - Bad **copy** of data
 - Bad **use** of the parallel operators

Long term

- **Mechanized** semantics in Coq
- Embed in Coq for **proofs** and **extraction** of multi-BSML programs as for BSML
- **where kernel** \Rightarrow **GPU** embed in multi-BSML
- Examples and **libraries**

Des questions ?