

Verified Parameterized Choreographies

Robert Rubbens, Petra van den Bos, Marieke Huisman

June 18th, 2025



UNIVERSITY
OF TWENTE.

Distributed systems

- Distributed systems are *everywhere*
 - Global internet platforms, embedded settings e.g. cars, software in general
- Distributed systems are *hard*
 - Generic properties: deadlocks, liveness, memory safety, race conditions, ...
 - Highly specific properties
 - Scalability, e.g. increasing number of nodes in a load balancer

Distributed systems

- Distributed systems are *everywhere*
 - Global internet platforms, embedded settings e.g. cars, software in general
- Distributed systems are *hard*
 - Generic properties: deadlocks, liveness, memory safety, race conditions, ...
 - Highly specific properties
 - Scalability, e.g. increasing number of nodes in a load balancer

Establishing correctness of distributed systems is *crucial*

Establishing correctness for distributed systems

- Generic properties, e.g. deadlock freedom
 - Testing, model checking, ...
 - DSL for protocols: *choreographies* [4]
- Highly specific properties: functional correctness verification tool
 - For programs: VerCors, Frama-C, Why3, Verifast, ...
 - For choreographies: *VeyMont* [2, 3, 6]
- Scalability: ...?

Establishing correctness for distributed systems

- Generic properties, e.g. deadlock freedom
 - Testing, model checking, ...
 - DSL for protocols: *choreographies* [4]
- Highly specific properties: functional correctness verification tool
 - For programs: VerCors, Frama-C, Why3, Verifast, ...
 - For choreographies: *VeyMont* [2, 3, 6]
- Scalability: ...?

In this talk:

- Extend *choreographic* verifier & code generator, *VeyMont*
- With support for *parameterized* choreographies
- How?
 - Extend choreographies with parameterization primitives
 - Limitations: one-to-one communication, (kind of) fix the network upfront
 - Encode using “structured parallelism”, conditionals, quantifiers

Talk outline

- Choreographies
 - Case study: distributed summation
- Choreographic Verification
 - With VeyMont (powered by VerCors)
- Parameterization in choreographies
- Choreographic projection
- Endpoint projection
- Distributed summation: correctness

Choreographies

Choreographies

- Constrained DSL for protocols¹
- Essentially:
 - A set of participants: *endpoints*
 - A sequence of interactions: *communications*

¹Fabrizio Montesi. *Introduction to Choreographies*. Cambridge University Press, 2023. DOI: 10.1017/9781108981491

Choreographies

- Constrained DSL for protocols¹
- Essentially:
 - A set of participants: *endpoints*
 - A sequence of interactions: *communications*

```
1  choreography summation2() {  
2      endpoint a = Node(int());  
3      endpoint b = Node(int());  
4      run {  
5          communicate a.sum -> b.in;  
6          communicate b.sum -> a.in;  
7          a.update();  
8          b.update();  
9      }  
10 }
```

¹Fabrizio Montesi. *Introduction to Choreographies*. Cambridge University Press, 2023. DOI: 10.1017/9781108981491

Benefits of choreographies

- Guarantees deadlock-freedom at the message level
 - By virtue of requiring *pairs* of endpoints for communication
- Generate code using the *endpoint projection*, $\llbracket \cdot \rrbracket_r$

Conceptual example of endpoint projection

summation2 ≡

```
1 choreography summation2() {  
2     endpoint a = Node(int());  
3     endpoint b = Node(int());  
4     run {  
5         communicate a.sum -> b.in;  
6         communicate b.sum -> a.in;  
7         a.update();  
8         b.update();  
9     }  
10 }
```

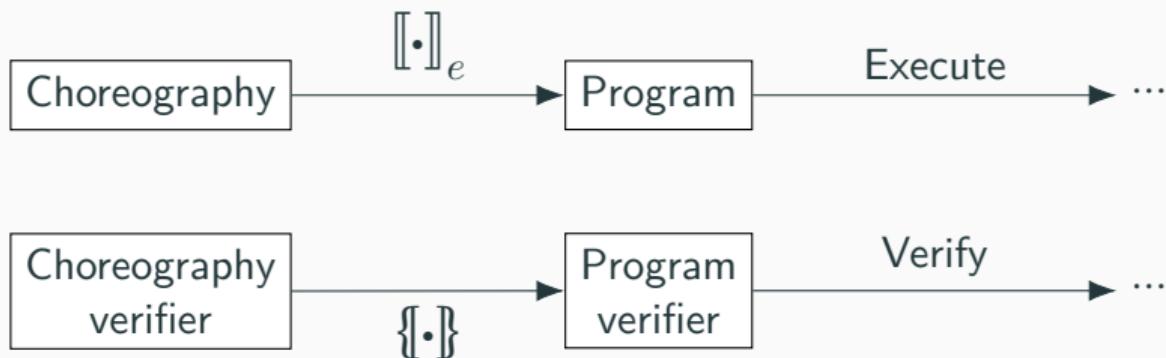
$\llbracket \text{summation2} \rrbracket_a \equiv$

```
1 void summation2_a(  
2     Channel to_b,  
3     Channel from_b) {  
4  
5     to_b.send(a.sum);  
6     a.in = from_b.receive();  
7     a.update();  
8     /* skip */  
9 }
```

Endpoint projection + verification = ...



Endpoint projection + verification = ...



Choreographic Verification

VeyMont & VerCors

VeyMont:

- Choreographic verifier and code generator
- Support for branching, loops [2, 3] and shared memory [6]
- Built on top of the *VerCors* program verifier
- Verifies choreographies with *pre- and postconditions*
- Functional correctness, memory safety, deadlock freedom

VerCors [5]:

- Verifies concurrent programs with shared memory
- Java, C, OpenCL, CUDA, PVL (Prototypal Verification Language)
- Permission-based separation logic [1]

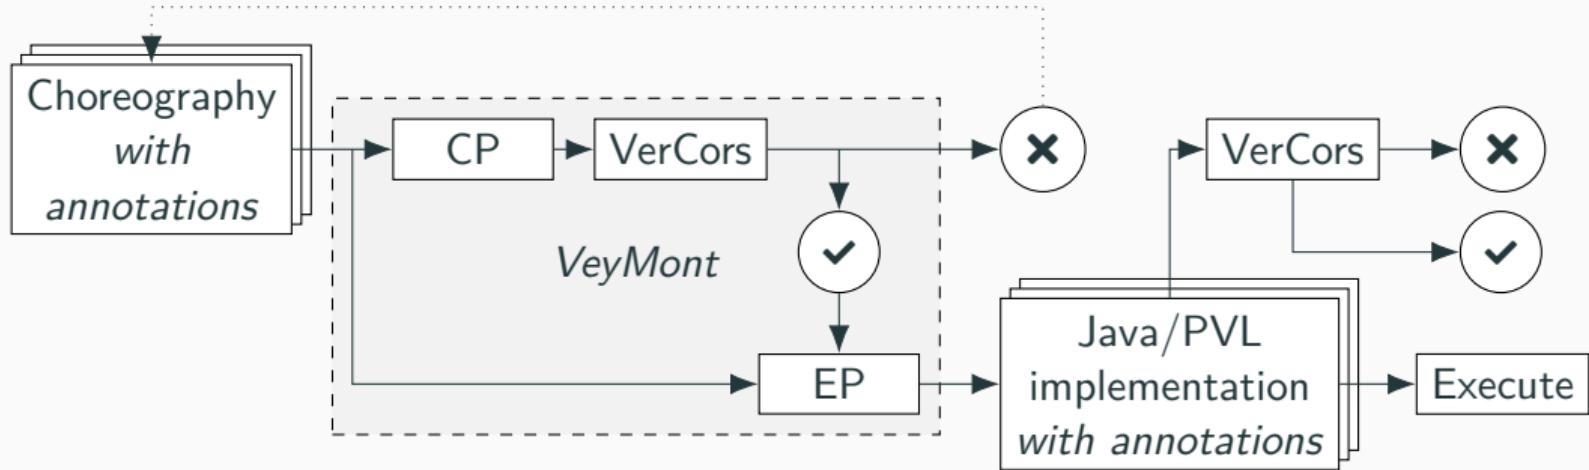
Annotated choreography example

```
1  choreography summation2() {
2      endpoint a = Node(int());
3      endpoint b = Node(int());
4
5
6      ensures a.sum == a.id + b.id; // <--  
7      ensures b.sum == a.id + b.id;
8      run {
9          communicate a.sum -> b.in;
10         communicate b.sum -> a.in;
11         a.update();
12         b.update();
13     }
14 }
```

Annotated choreography example

```
1  choreography summation2() {
2      endpoint a = Node(int());
3      endpoint b = Node(int());
4
5      requires a.sum == a.id && b.sum == b.id;
6      ensures a.sum == a.id + b.id; // <-
7      ensures b.sum == a.id + b.id;
8      run {
9          communicate a.sum -> b.in;
10         communicate b.sum -> a.in;
11         a.update();
12         b.update();
13     }
14 }
```

VeyMont pipeline



- CP: Choreographic Projection, $\{\!\cdot\!\}$
- EP: Endpoint Projection, $\llbracket \cdot \rrbracket_r$

summation2: choreographic projection {}

```
choreography summation2() {  
    endpoint a = Node(int());  
    endpoint b = Node(int());  
  
    requires a.sum == a.id;  
    requires b.sum == b.id;  
    ensures a.sum == a.id + b.id;  
    ensures b.sum == a.id + b.id;  
    run {  
        communicate a.sum -> b.in;  
        communicate b.sum -> a.in;  
        a.update();  
        b.update();  
    }  
}
```

```
1 void summation2_cp() {  
2     Node a = new Node(int());  
3     Node b = new Node(int());  
4  
5     assert a.sum == a.id  
6     assert b.sum == b.id;  
7  
8     b.in = a.sum;  
9     a.in = b.sum;  
10  
11    a.update();  
12    b.update();  
13  
14    assert a.sum == a.id + b.id;  
15    assert b.sum == a.id + b.id;  
16 }
```

Parameterization in choreographies

Parameterization in choreographies

Data parameterization:

```
choreography summationN(int msg) {
    endpoint s = Sender(msg);
    endpoint r = Receiver();
    run {
        communicate s.msg -> r.inbox;
    }
}
```

- Limitation: numbers of endpoints statically known
- Why: natural scalability
 - Sharding, ring networks, users, etc.

Parameterization in choreographies

Data parameterization:

```
choreography summationN(int msg) {
    endpoint s = Sender(msg);
    endpoint r = Receiver();
    run {
        communicate s.msg -> r.inbox;
    }
}
```

- Limitation: numbers of endpoints statically known
- Why: natural scalability
 - Sharding, ring networks, users, etc.

Parameterization in *number of endpoints*

Parameterization primitives

Endpoint family:

```
endpoint ns[i := 0 .. N] = Node(i, int());
```

Endpoint quantification:

```
(\endpoints ns[i := 0..N]; nodes[i].n < N-1)
```

Parameterized communicate:

```
communicate ns[i := 0..N-1].sum -> ns[i+1].in;
```

Parameterization primitives

Endpoint family:

```
endpoint ns[i := 0 .. N] = Node(i, int());
```

Endpoint quantification:

```
(\endpoints ns[i := 0..N]; nodes[i].n < N-1)
```

Parameterized communicate:

```
communicate ns[i := 0..N-1].sum -> ns[i+1].in;
```

Limitations:

1. In $ns[E]$, E cannot read the heap
 - $ns[store.x]$ ✘
2. E must be invertible
 - $i + 1 \Leftrightarrow i - 1$

summation2 → summationN

```
1 endpoint ns[i := 0 .. N] = Node(i, int());
2
3 run:
4 while ((\endpoints ns[i := 0..N]; nodes[i].n < N-1)) {
5   communicate ns[i := 0..N-1].sum -> ns[i+1].in;
6   communicate ns[N-1].sum -> ns[0].in;
7   nodes[i := 0..N].update();
8 }
```

Projections

Choreographic projection: endpoints



Endpoint family:

```
{endpoint ns[i := 0 .. N] = Node(i, int());} =  
    seq<Node> ns = [];  
    for (int i = 0 .. N) {  
        ns = ns + [new Node()];  
    }
```

Choreographic projection: endpoints

{•}

Endpoint family:

```
{endpoint ns[i := 0 .. N] = Node(i, int());} =  
    seq<Node> ns = [];  
    for (int i = 0 .. N) {  
        ns = ns + [new Node()];  
    }
```

Quantification:

```
{(\!endpoints ns[i := 0..N]; nodes[i].n < N-1)} =  
(\!forall int i = 0..N; nodes[i].n < N-1)
```

Choreographic projection: communicate (*)



```
{communicate ns[i := 0..N-1].sum -> ns[i+1].in;} =  
    // Start N threads, and wait until finished  
    par (int i = 0 .. N - 1)  
        // That can read ns[i].sum  
        requires Perm(ns[i].sum, read);  
        // And write to ns[i + 1].in  
        requires Perm(ns[i + 1].in, write);  
    {  
        // Re-use regular encoding  
        {communicate ns[i].sum -> ns[i + 1].in;}  
    }
```

Endpoint projection: endpoints

$\llbracket \dots \rrbracket_{ns[i]}$

Endpoint family:

```
[[endpoint ns[i := 0 .. N] = Node(i, int());]]_{ns[i]} =
void summationN_nodes_i(seq<Node> ns, int i) {
    [[run { ... }]]_{ns[i]}
}
```

Endpoint projection: endpoints

$\llbracket \dots \rrbracket_{ns[i]}$

Endpoint family:

```
[\![ endpoint ns[i := 0 .. N] = Node(i, int()); ]\!]_{ns[i]} =  
    void summationN_nodes_i(seq<Node> ns, int i) {  
        [\![ run { ... } ]\!]_{ns[i]}  
    }
```

Endpoint quantification:

```
[\!(\!\! endpoints ns[k := 0..N]; nodes[k].n < N-1)\!]_{ns[i]} =  
    0 <= i && i < N ==> nodes[i].n < N-1
```

Endpoint projection: parameterized communicate

$\llbracket \dots \rrbracket_{ns[i]}$

```
[communicate ns[i := 0..N-1].sum -> ns[i+1].in;]ns[i] =  
    // i in sending range  
    if (0 <= i && i < N - 1) {  
        ns_ns[i].writeValue(ns[i].sum)  
    }  
    // i in receiving range  
    if (0 + 1 <= i && i < N - 1 + 1) {  
        ns[i].in = ns_ns[i - 1].readValue();  
    }
```

summationN: Functional correctness

summationN correctness

```
1 choreography summationN(int N) {
2     endpoint ns[i := 0 .. N] = Node(i, int());
3     ensures POST;
4     run {
5         loop_invariant INV;
6         while ((\endpoints ns[i := 0..N]; ns[i].n < N-1)) {
7             communicate ns[i := 0..N-1].sum -> ns[i+1].in;
8             communicate ns[N-1].sum -> ns[0].in;
9             nodes[i := 0..N].update();
10        }
11        // Proof step
12        endpoint nodes[i := 0..N].generalizeTo0();
13    }
14 }
```

POST \equiv

```
(\endpoints ns[i := 0..N];
ns[i].sum == sum(ns, 0, ns[i].n));
```

Conclusion

Conclusion

- Extended choreographies with primitives for parameterization
- Two limitations:
 1. Family indexing must be heap-independent
 2. Indexing must be one-to-one
- Verification: sequences & par blocks
- Code generation: conditionals (implication and if)
- Manual encoding of *distributed summation*



Robert Rubbens
Formal Methods & Tools
University of Twente
r.b.rubbens@utwente.nl - <https://bobrubbens.nl>



References

- [1] Afshin Amighi et al. "Permission-Based Separation Logic for Multithreaded Java Programs". In: *Logical Methods in Computer Science* 11.1 (2015). DOI: [10.2168/LMCS-11\(1:2\)2015](https://doi.org/10.2168/LMCS-11(1:2)2015).
- [2] Sung-Shik Jongmans and Petra van den Bos. "A Predicate Transformer for Choreographies - Computing Preconditions in Choreographic Programming". In: *ESOP*. 2022. DOI: [10.1007/978-3-030-99336-8_19](https://doi.org/10.1007/978-3-030-99336-8_19).
- [3] Petra van den Bos and Sung-Shik Jongmans. "VeyMont: Parallelising Verified Programs Instead of Verifying Parallel Programs". In: *FM*. 2023. DOI: [10.1007/978-3-031-27481-7_19](https://doi.org/10.1007/978-3-031-27481-7_19).
- [4] Fabrizio Montesi. *Introduction to Choreographies*. Cambridge University Press, 2023. DOI: [10.1017/9781108981491](https://doi.org/10.1017/9781108981491).
- [5] Lukas Armborst et al. "The VerCors Verifier: A Progress Report". In: *CAV*. 2024. DOI: [10.1007/978-3-031-65630-9_1](https://doi.org/10.1007/978-3-031-65630-9_1).

References (cont.)

- [6] Robert Rubbens, Petra van den Bos, and Marieke Huisman. “VeyMont: Choreography-Based Generation of Correct Concurrent Programs with Shared Memory”. In: *iFM*. 2024. DOI: [10.1007/978-3-031-76554-4_12](https://doi.org/10.1007/978-3-031-76554-4_12).

Bonus

(*)

- Deadlock freedom: branch unanimity
- Communicate: injectivity of E
- Verification only requires injectivity, not invertibility

summationN:

- Ghost code for sending ghost memory back and forth
- Proof steps for generalizing correctness result of one individual endpoint

summationN code

```
choreography summationN(int N) {
    endpoint nodes[i := 0 .. N] = Node(i, int());
    run {
        while ((\endpoints nodes[i := 0..N]; nodes[i].n < N-1)) {
            communicate nodes[i := 0..N-1].sum -> nodes[i+1].in;
            communicate nodes[N-1].sum -> nodes[0].in;
            nodes[i := 0..N].update();
        }
    }
}
```

```
(\endpoints ns[i := 0..N];
  ns[i].total ==
    sum(ns, ns[i].i, ns[i].n))
```