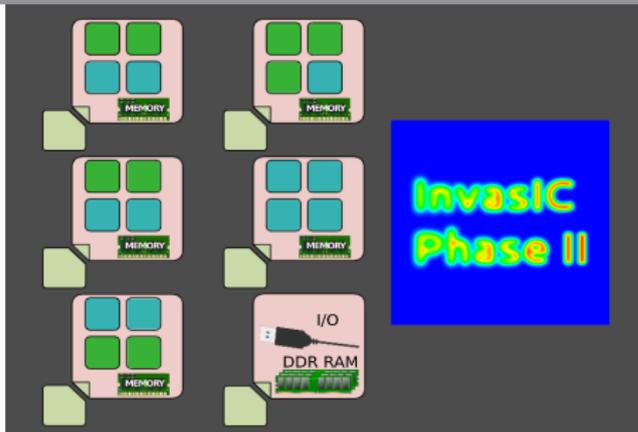


RaPping and Compilation for Highly Dynamic Parallelism

Gregor Snelting

Presentation at ANDREAS-50, 26.2.16



Background

SFB Transregio **InvasIC**: FAU, KIT, TUM

- new language, soft- and hardware for **highly dynamic parallelism**
- 3 universities, 15 PIs, 50 doctoral researchers
- DFG funding 2010–2018 (2022?)

Background

SFB Transregio **InvasIC**: FAU, KIT, TUM

- new language, soft- and hardware for **highly dynamic parallelism**
- 3 universities, 15 PIs, 50 doctoral researchers
- DFG funding 2010–2018 (2022?)

Scenario: big heterogenous **PGAS** architectures, several kCores / Tile

Background

SFB Transregio **InvasIC**: FAU, KIT, TUM

- new language, soft- and hardware for **highly dynamic parallelism**
- 3 universities, 15 PIs, 50 doctoral researchers
- DFG funding 2010–2018 (2022?)

Scenario: big heterogenous **PGAS** architectures, several kCores / Tile

Goal: support highly dynamic parallelism **through all system levels**

Background

SFB Transregio **InvasIC**: FAU, KIT, TUM

- new language, soft- and hardware for **highly dynamic parallelism**
- 3 universities, 15 PIs, 50 doctoral researchers
- DFG funding 2010–2018 (2022?)

Scenario: big heterogenous **PGAS** architectures, several kCores / Tile
Goal: support highly dynamic parallelism **through all system levels**



DFG

Sonderforschungsbereich/Transregio 89
Transregional Collaborative Research Center 89



RaP to the Future

Fundamental idea: **Resource-aware programming**

- applications dynamically determine their resource requirements
- allocate and release resources (in particular CPUs) dynamically
- supported by language, compiler, OS, and dedicated hardware

⇒ *optimize resource utilization, energy consumption, fault tolerance*

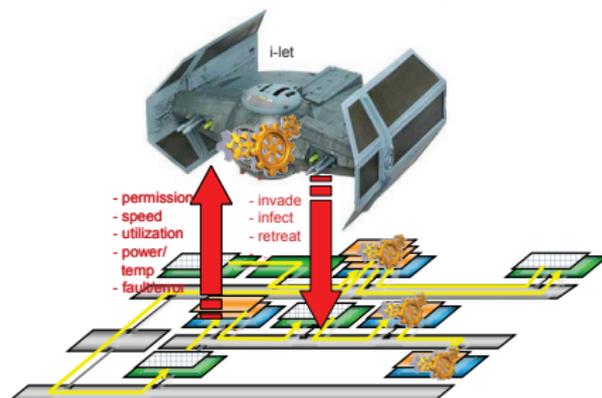
RaP to the Future

Fundamental idea: **Resource-aware programming**

- applications dynamically determine their resource requirements
- allocate and release resources (in particular CPUs) dynamically
- supported by language, compiler, OS, and dedicated hardware

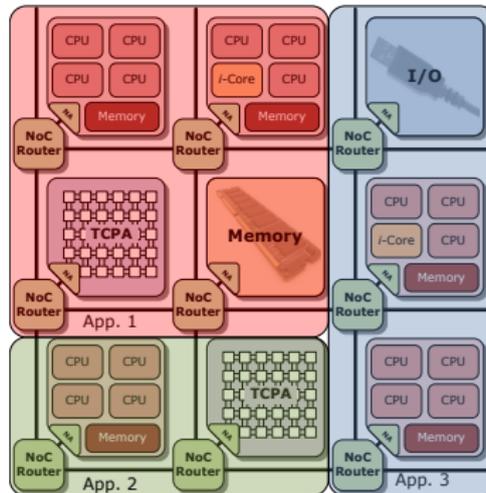
⇒ *optimize resource utilization, energy consumption, fault tolerance*

Slogan (J. Teich): “**Invade, infect, combat, retreat**”



Example: InvasiC Multi-Tile MPSoC

tiles: RISC multicore, memory, I/O, TCPA, ...
3 applications have invaded the example tiles



Demonstrator implemented via FPGAs

Fundamental InvasIC Ingredients

- Dynamic X10 + Compiler: see below
- constraint system: attribute hierarchy for resource selection
- OctoPos OS: dedicated scheduling, load balancing
- Agent system: prioritizes / allocates dynamic resource requests
- CiC: hardware support on local tile for app code dispatch / monitoring
- NoC: network on chip
- i-core: adaptive hardware supporting specific accelerators
- TCPAs: support (numerical) loop parallelization

Fundamental InvasIC Ingredients

- Dynamic X10 + Compiler: see below
- constraint system: attribute hierarchy for resource selection
- OctoPos OS: dedicated scheduling, load balancing
- Agent system: prioritizes / allocates dynamic resource requests
- CiC: hardware support on local tile for app code dispatch / monitoring
- NoC: network on chip
- i-core: adaptive hardware supporting specific accelerators
- TCPAs: support (numerical) loop parallelization

Further aspects:

- applications: HPC, robotics
- dark silicon: dynamic selection of non-overheated cores etc.
- predictability: soft realtime and security
- hardware sensors: temperature, speed, fault status etc

PP Group @ InvasIC

People: Gregor Snelting, (Matthias Braun),
Sebastian Buchwald, Manuel Mohr, Andreas Zwinkau

People: Gregor Snelting, (Matthias Braun),
Sebastian Buchwald, Manuel Mohr, Andreas Zwinkau

Contributions:

- framework design (based on X10)
- fundamental RaP examples
- compiler, library, RTS
- code generation for MPSoCs
- optimizations for InvasIC OS / hardware
- interfaces to simulator, OS, TCPA subcompiler [FAU], ...
- application support

Basic RaP Patterns & Constraints

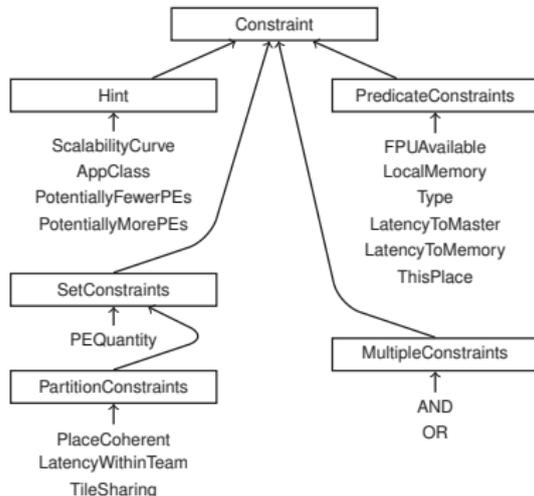
fundamental “invade/infect/retreat” pattern (X10) uses closure for app code:

```
1 val ilet = (id:IncarnationID) => {
2     do_something(id);
3 };
4 claim = Claim.invade(constraints)
5 claim.infect(ilet)
6 claim.retreat()
```

Basic RaP Patterns & Constraints

fundamental “invade/infect/retreat” pattern (X10) uses closure for app code:

```
1 val ilet = (id:IncarnationID) => {  
2     do_something(id);  
3 };  
4 claim = Claim.invade(constraints)  
5 claim.infect(ilet)  
6 claim.retreat()
```



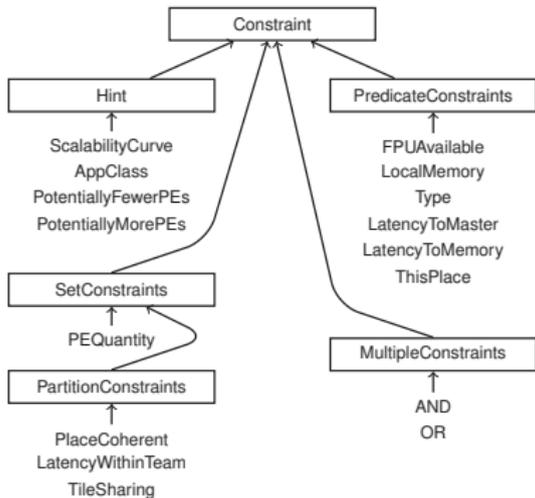
Basic RaP Patterns & Constraints

fundamental “invade/infect/retreat” pattern (X10) uses closure for app code:

```
1 val ilet = (id:IncarnationID) => {  
2     do_something(id);  
3 };  
4 claim = Claim.invade(constraints)  
5 claim.infect(ilet)  
6 claim.retreat()
```

“reinvade()” allows OS to adapt resources
“reinvade(nC)” allows app to reallocate

```
1 claim.infect(ilet)  
2 // optimize resource allocation:  
3 val changed1 = claim.reinvade()  
4 claim.infect(ilet)  
5 // respecify resource needs  
6 val changed2 = claim.reinvade(otherConstraints)
```



Example: Invasive Quicksort

Sequential X10 code:

```
1 public static def sort(data: Array[Int](1)) =
2   qsort(data, 0, data.size-1);
3
4 private static def qsort(
5   data: Array[Int](1),
6   left: Int,
7   right: Int)
8 {
9   if (data.size>1) {
10    val p = partition(data, left, right);
11    val i = p.first;
12    val j = p.second;
13    qsort(data, left, i);
14    qsort(data, j, right);
15  }
16 }
```

Goal: Parallelize Recursion

Pedagogical engineering on this part:

```
1  qsort(data, left, j    );  
2  qsort(data, i,    right);
```

Requirements:

- Invasive
- Multi-core
- Multi-tile
- Dynamically Resource-aware

Disclaimer:

- *Pedagogical* example
e.g. does not parallelize `partition` step

Handling the two parts

- After decision to parallelize: two parallel program paths
 1. `qsort()`
 2. `invade()` + (copying) + `qsort()` + (copying) + `retreat()`
 - Important: *large* and *small* part, not left and right
 - Hope: $t_{large} \approx t_{invade} + t_{copying_1} + t_{small} + t_{copying_2} + t_{retreat}$
- ⇒ Always handle **large part locally**
- ⇒ try to handle small part on new PE



RaP 1: Invasive Single-Tile

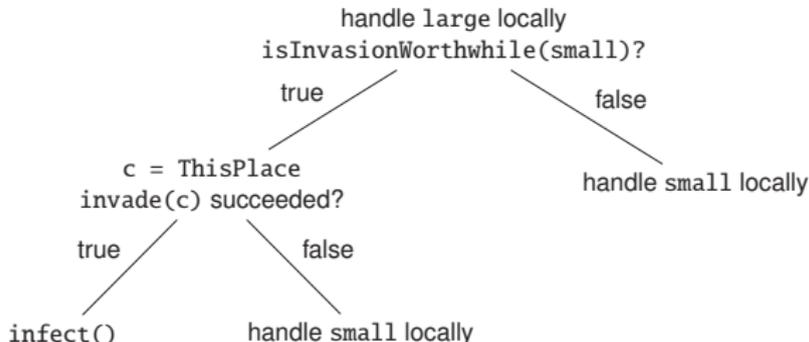
small part runs on new core in local tile (X10: “ThisPlace”)

RaP 1: Invasive Single-Tile

small part runs on new core in local tile (X10: “ThisPlace”)

“isInvasionWorthwhile()”: **user defined**; checks new thread for small part;
based on problem size, processor speed, invade latencies, ...

Decision tree:



RaP 1: X10 Code

```
1  finish {
2    if (large.size > 1) { async qsort(large); }
3    if (isInvasionWorthwhile(small)) {
4      val constraints = new AND(); // constraint system
5      constraints.add(new PEQuantity(0, 1));
6      constraints.add(new ThisPlace());
7      constraints.add(new ScalabilityHint([100, 190]));
8      val claim = Claim.invaade(constraints);
9      if (claim.size() == 1) {
10       val ilet = (id: IncarnationID) => qsort(small);
11       claim.infect(ilet);
12     } else { qsort(small); }
13     claim.retreat();
14   } else if (small.size > 1) {
15     qsort(small);
16   }
17 }
```

RaP 2: Invasive Multi-Tile

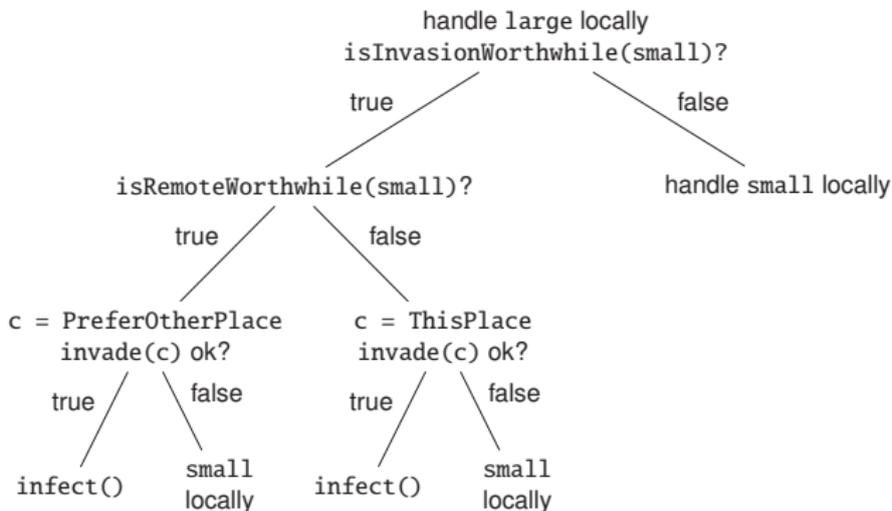
small part runs on new tile

RaP 2: Invasive Multi-Tile

small part runs on new tile

“isRemoteWorthwhile()” checks new tile for small part;
 based on problem size, tile attributes, network latencies, ...

Decision tree:



RaP 2: X10 Code

```
1  finish {
2    if (large.size > 1) { async qsort(large); }
3    if (isInvasionWorthwhile(small)) {
4      val constraints = new AND();
5      constraints.add(new PEQuantity(0, 1));
6      if (isRemoteWorthwhile(small)) {
7        constraints.add(new PreferOtherPlace());
8        constraints.add(new ScalabilityHint([100, 160]));
9      } else {
10       constraints.add(new ThisPlace());
11       constraints.add(new ScalabilityHint([100, 190]));
12     }
13     val claim = Claim.invade(constraints);
14     if (claim.size() == 1) {
15       val ilet = (id: IncarnationID) => qsort(small);
16       claim.infect(ilet);
17     } else {
18       qsort(small);
19     }
20     claim.retreat();
21   } else if (small.size > 1) {
22     qsort(small);
23   }
24 } // finish
```

RaP 3: Fine Tuning

Other sensible constraints in this example:

ThroughputToMaster, ThroughputToMemory, ThroughputWithinTeam,
LatencyToMaster, LatencyToMemory, LatencyWithinTeam, ...

RaP 3: Fine Tuning

Other sensible constraints in this example:

ThroughputToMaster, ThroughputToMemory, ThroughputWithinTeam,
LatencyToMaster, LatencyToMemory, LatencyWithinTeam, ...

General Principle:

- Application dynamically crafts its constraints based on *locally* available data, e.g. the input size
- The *larger* the input size, the *weaker* the constraints
 - Three “levels” in quicksort example
- Application decides when it is certainly *not* sensible to parallelize
- Application supplies scalability hints for `invade()` call
- OS (“Agent system”) decides which application actually gets PEs

The Compiler

working compiler was built with only a few PJ!

- front end: adapted from IBM X10 compiler
- back end: PP@KIT, based on libFIRM (→ Goos, Hack et al.)
“FIRM is like llvm, only better”

The Compiler

working compiler was built with only a few PJ!

- front end: adapted from IBM X10 compiler
- back end: PP@KIT, based on libFIRM (→ Goos, Hack et al.)
“FIRM is like llvm, only better”
- generates code for SPARC (Leon 5)
- “invade” etc implemented as OctoPos calls
- compiler tested and integrated into demonstrator (FPGA system)

The Compiler

working compiler was built with only a few PJ!

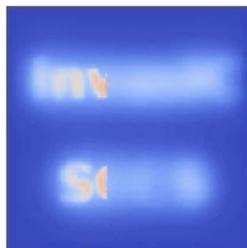
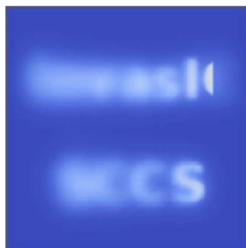
- front end: adapted from IBM X10 compiler
- back end: PP@KIT, based on libFIRM (→ Goos, Hack et al.)
“FIRM is like llvm, only better”
- generates code for SPARC (Leon 5)
- “invade” etc implemented as OctoPos calls
- compiler tested and integrated into demonstrator (FPGA system)

Current work:

- specialised optimizations for InvasIC
inference of app characteristics for better load balancing,
infer minimal execution time for “isInvasionWorthwhile()”,
infer efficient i-core configurations (accelerators, permutation instructions, ...),
efficient copying of data to other places,
...
■ generate code for NoC, i-Core

Case Study: Multigrid Software

- Invasive multigrid: solver for part.diff.equations from TUM+KIT
- discretization with **dynamically varying grid resolutions**
- application: simulation of laser engraving on metal plate
⇒ heat equation (Laplace operator) must be discretized and solved
- grid levels: $N \times N$, $N/2 \times N/2$, $N/4 \times N/4$, ...
- result at timesteps 50/200/450:



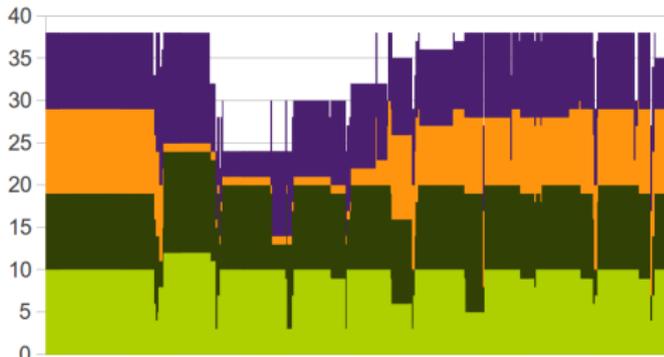
X10 Code + Dynamic Loads

```
1 vcycle(N, x, b):
2 r = computeResidual(N, x, b)
3 while |r| > threshold:
4 vcycleIteration(N, x, b)
5 r = computeResidual(N, x, b)
6
7
8 vcycleIteration(N, x, b):
9 smoother(N, x, b)      # pre-smooth
10 r = residual(N, x, b)  # residual
11
12 nc = reinvade(N, claim) # reinvade claim
13
14 Nr = N/2
15 # restricted level
16 rr = restrict(N, r)
17 # restrict residual to new claim
18 er = (Nr, 0)
19 # setup error with 0 values
20
21 nc2 = vcycleIteration(Nr, er, rr)
22 # vcycle
23
24 # redistribute
25 if (nc != nc2):
26 nc = nc2
27 x.redistribute(Nr, nc)
28 b.redistribute(Nr, nc)
29
30 e = prolongate(Nr, er)
31 # prolongate error
32 x = x + e
33 # apply correction
34
35 smoother(N, x, b)      # post-smooth
36
37 return nc # possibly modified claim
```

X10 Code + Dynamic Loads

```
1 vcycle(N, x, b):
2 r = computeResidual(N, x, b)
3 while |r| > threshold:
4 vcycleIteration(N, x, b)
5 r = computeResidual(N, x, b)
6
7
8 vcycleIteration(N, x, b):
9 smoother(N, x, b)      # pre-smooth
10 r = residual(N, x, b)  # residual
11
12 nc = reinvade(N, claim) # reinvade claim
13
14 Nr = N/2
15 # restricted level
16 rr = restrict(N, r)
17 # restrict residual to new claim
18 er = (Nr, 0)
19 # setup error with 0 values
20
21 nc2 = vcycleIteration(Nr, er, rr)
22 # vcycle
23
24 # redistribute
25 if (nc != nc2):
26 nc = nc2
27 x.redistribute(Nr, nc)
28 b.redistribute(Nr, nc)
29
30 e = prolongate(Nr, er)
31 # prolongate error
32 x = x + e
33 # apply correction
34
35 smoother(N, x, b)      # post-smooth
36
37 return nc # possibly modified claim
```

Four multigrid applications;
combat mode on 38 PEs:



- **InvasIC**: hard- + software support for resource-aware parallel programming
- Dynamic RaPping based on local estimations:
 1. check whether additional resources required,
 2. *invade* resources (= obtain a “claim” of PEs),
 3. *infect* them (= load app code + data to new PEs + execute),
 4. *combat* (= OS balances dynamic resource requests),
 5. *retreat* (= free invaded claims)

Conclusion

- **InvasIC**: hard- + software support for resource-aware parallel programming
- Dynamic RaPping based on local estimations:
 1. check whether additional resources required,
 2. *invade* resources (= obtain a “claim” of PEs),
 3. *infect* them (= load app code + data to new PEs + execute),
 4. *combat* (= OS balances dynamic resource requests),
 5. *retreat* (= free invaded claims)
- PP@KIT: language + compiler + applications
- next step: validate resource utilization / energy consumption / fault tolerance in combat setting on demonstrator system

Thank You!