

Result Certification Against Massive Attacks in Distributed Computations

Sébastien Varrette¹, Jean-Louis Roch² and Axel Krings³

¹ Computer Science and Communications Unit, University of Luxembourg, Luxembourg

² MOAIS team, LIG Laboratory, Grenoble, France

³ Department of Computer Science, University of Idaho, Moscow, USA

Re-Trust 2008, Trento (Italy), October 15th, 2008




University of Idaho

Summary

- 1 **Context & Motivations**
- 2 **Execution and certification model**
- 3 **Monte-Carlo certification of independent tasks**
- 4 **Monte-Carlo certification of dependent tasks**
- 5 **Conclusion**


Large scale computing platforms

- [Beowulf] Clusters : Chaos.lu




Large scale computing platforms

- Computing grids [Foster&al.97] : Grid5000, Globus etc.




Large scale computing platforms

- « Desktop grid » : Seti@Home, BOINC, XtremWeb etc.




Result-Checking issue

- Falsified result : malicious act or not (cf. Seti@HOME [Molnar00])



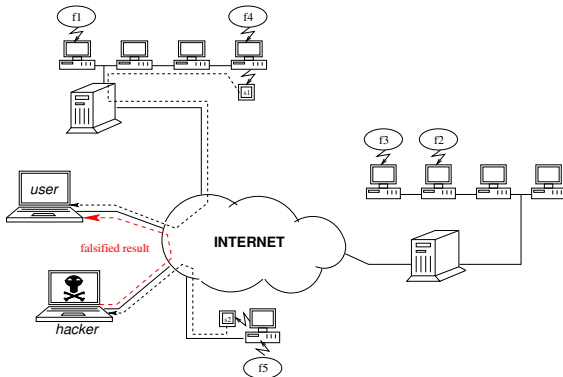
Result-Checking issue

- Falsified result : malicious act or not (cf. Seti@HOME [Molnar00])



Result-Checking issue

- Falsified result : malicious act or not (cf. Seti@HOME [Molnar00])



- Software Counter-measures : prevent before / **control after**

State of the art

- Essentially devoted to batches of **independent** tasks

Specific approach : Simple checker [Blum97]

- check a [cheap] post-condition over computed results
 - ↪ DLP avec $|G| = n : L_n \left[\frac{1}{3}, \left(\frac{64}{9} \right)^{\frac{1}{3}} \right]$ – Simple checker $\mathcal{O}(\log n)$
- **The most efficient approach...** if possible !

State of the art

- Essentially devoted to batches of **independent** tasks

Specific approach : Simple checker [Blum97]

- check a [cheap] post-condition over computed results
 - ↪ DLP avec $|G| = n : L_n \left[\frac{1}{3}, \left(\frac{64}{9} \right)^{\frac{1}{3}} \right]$ – Simple checker $\mathcal{O}(\log n)$
- **The most efficient approach...** if possible !

General approach : duplication

- Direct certification of the batch with sequential tests [Germain-Playez03]
- Batch reinforcement [Sarmenta03]
- **In all case** : attackers modelisation

State of the art

- Essentially devoted to batches of **independent** tasks

Specific approach : Simple checker [Blum97]

- check a [cheap] post-condition over computed results
 - ↳ DLP avec $|G| = n : L_n \left[\frac{1}{3}, \left(\frac{64}{9} \right)^{\frac{1}{3}} \right]$ – Simple checker $\mathcal{O}(\log n)$
- **The most efficient approach...** if possible !

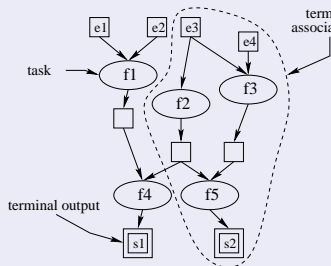
General approach : duplication

- Direct certification of the batch with sequential tests [Germain-Playez03]
- Batch reinforcement [Sarmenta03]
- **In all case** : attackers modelisation

⇒ **What about dependent tasks ?**

Execution model : macro-dataflow graph

Abstract representation of a parallel execution $P(i)$

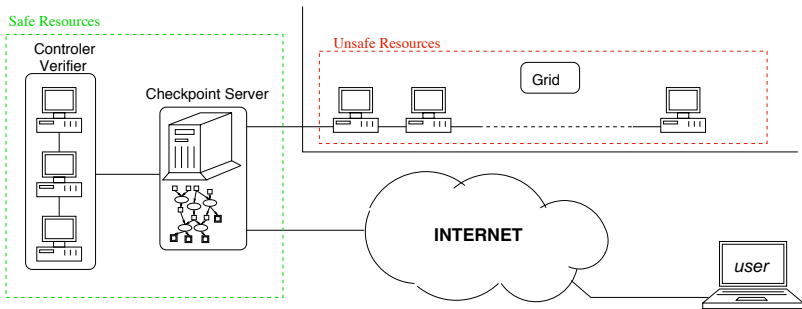


- $G^<(T)$: predecessors of T in G
- $G^{\leq}(T)$: $G^<(T) \cup \{T\}$
- Execution engine : KAAPI

↪ <http://kaapi.gforge.inria.fr/>


↪ C++ library for high performance parallel computing

Execution platform



- Resources partitionning $|R| \ll |U|$
- Reliable system for task re-execution

Impact of the faults (1)



- E : execution of $\mathbf{P}(i)$ over U resources
 - $\hookrightarrow G$ with intermediate values
 - $\hookrightarrow T \in E : i(T, E) \longrightarrow o(T, E)$
- \hat{E} : execution of $\mathbf{P}(i)$ over R resources
 - $\hookrightarrow T \in \hat{E} : \hat{i}(T, \hat{E}) \longrightarrow \hat{o}(T, \hat{E})$

Definition (execution state)

E is **correct** iff $E = \hat{E}$. Otherwise, E is **falsified**.

Task re-execution : compute $\hat{o}(T, E)$ from $i(T, E)$, compare to $o(T, E)$

Impact of the faults (2)

Definition (Correct and Faulty task)

- **Faulty task** $T : o(T, E) \neq \hat{o}(T, E)$
 - ↪ directly detected by controllers
 - ↪ **correct task** T : no task in $G^{\leq}(T)$ are faulty
- **Falsified result** : $o(T, E) \neq \hat{o}(T, \hat{E})$
 - ↪ hard to detect as $\hat{o}(T, \hat{E}) \neq \hat{o}(T, E)$
 - ↪ n_F falsified tasks

Monte-Carlo certification (1)

Definition (certification Monte-Carlo algorithm)

$$\mathcal{A} : (E, \varepsilon) \longrightarrow \begin{cases} \text{CORRECT (with error probability } \leq \varepsilon) \\ \text{FALSIFIED (with falsification proof)} \end{cases}$$

- Cf. Miller-Rabin
- **Interests** :
 - ↪ ε fixed by the user
 - ↪ a limited number of verifier calls (**ideally** $\mathbf{o(n)}$)
 - ↪ **can be done in parallel on R!**

Efficient detection of masive attack ($n_F \geq n_q = \lceil q \cdot n \rceil$)

- ↪ the application **should** tolerate a limited number of faults
- ↪ no assumption on attackers behaviour except n_F

Monte-Carlo certification (2)

Resources	avg. speed/proc	total speed
U	Π_U	Π_U^{tot}
R	Π_R	Π_R^{tot}

- Scheduling by **on-line workstealing**

↪ execution (on U) : $W_1 \gg W_\infty$

↪ certification (on R) : W_1^C and W_∞^C


Theorem (Executing and Certification Time)

w.h.p :

$$T_{EC} \leq \left[\frac{W_1}{\Pi_U^{tot}} + \mathcal{O}\left(\frac{W_\infty}{\Pi_U}\right) \right] + \left[\frac{W_1^C}{\Pi_R^{tot}} + \mathcal{O}\left(\frac{W_\infty^C}{\Pi_R}\right) \right]$$


Independent case

Correct execution :



Independent case

Falsified execution :



Independent case

Monte-Carlo Test $MCT(E)$

Input: Execution E represented by G composed of independent tasks.

Output: The correctness of E (FALSIFIED or CORRECT)

Uniformly choose one task T in G ;

// Re-execution of T on the R resources using the inputs $i(T, E)$

$\hat{o}(T, E) \leftarrow \text{ReexecuteOnVerifier}(T, i(T, E));$

if $o(T, E) \neq \hat{o}(T, E)$ **then**
 return FALSIFIED;


return CORRECT;

Theorem (Probabilistic certification by $MCT(E)$)

- $\mathcal{A}(E, \varepsilon) : N_{\varepsilon, q} = \lceil \frac{\log \varepsilon}{\log(1-q)} \rceil$ calls to $MCT(E)$
- $W_1^C \leq N_{\varepsilon, q} W_\infty$ and $W_\infty^C = W_\infty$
- $T_{EC} \leq \frac{W_1}{\prod_U^{tot}} + \mathcal{O}\left(\frac{W_\infty}{\prod_U}\right) + \frac{N_{\varepsilon, q} W_\infty}{\prod_R^{tot}} + \mathcal{O}\left(\frac{W_\infty}{\prod_R}\right)$


Independent case

Impact of ε over $N_{\varepsilon,q}$



Independent case

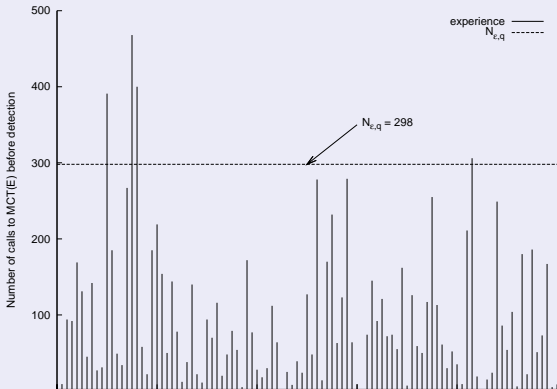
Impact of q over $N_{\epsilon,q}$



Independent case


Non-detection illustration

$$n = 10^6, q = 1\% \text{ and } \varepsilon = 5\%$$




Dependent case

Correct execution :



Dependent case

Falsified execution :



Dependent case

- n_I initiators $\in \mathcal{I}(F) : \begin{cases} i(T, E) = \hat{i}(T, \hat{E}) \\ o(T, E) \neq \hat{o}(T, \hat{E}) \end{cases}$
 - ↳ falsified tasks you are **sure** to detect
- $\mathcal{P}(MCT(E) = CORRECT) \leq 1 - \frac{n_I}{n}$

Theorem (Minimal number of initiators)

For G with height h , maximal out-degree d and $n_F \geq n_q = \lceil q \cdot n \rceil$

$$n_I \geq q \left\lceil \frac{n(d-1)}{d^h - 1} \right\rceil$$

Dependent case

Lemma (Initiators characterization)

- $\mathcal{I}(F) = \{T_i \in F : F \cap G^<(T_i) = \emptyset\}$
- T is falsified $\iff G^{\leq}(T) \cap \mathcal{I}(F) \neq \emptyset$

Extended Monte-Carlo Test $EMCT(E)$

Input: Execution E represented by G composed of dependent tasks.

Output: The correctness of E (FALSIFIED or CORRECT)

Uniformly choose one task T in G ;

// Re-execution of $G^{\leq}(T)$ on R to detect initiators

forall $T_j \in G^{\leq}(T) / T_j$ as not yet been checked **do**
 $\hat{o}(T_j, E) \leftarrow \text{ReexecuteOnVerifier}(T_j, i(T_j, E));$
 if $o(T_j, E) \neq \hat{o}(T_j, E)$ **then**
 return FALSIFIED;

end

return CORRECT;


Dependent case

Theorem (Probabilistic certification by $EMCT(E)$)


- $\mathcal{A}(E, \varepsilon) : N_{\varepsilon, q} = \lceil \frac{\log \varepsilon}{\log(1-q)} \rceil$ calls to $EMCT(E)$
- Expected cost per call : $C_G = \frac{1}{n} \sum_{T \in G} |G^{\leq}(T)|$
- **Worst case** : $W_1^C = \Omega(W_1)$ and $W_\infty^C = \Omega(W_\infty)$

Dependent case on some specific graphs

Trees



Fork-Join graphs



Theorem (Trees and Fork-Join graphs certification)

For G a tree or a Fork-Join graph with height h :

- $C_G \leq h + 3$
- $T_{EC} \leq \frac{W_1}{\prod_U^{tot}} + \mathcal{O}\left(\frac{W_\infty}{\prod_U}\right) + \mathcal{O}\left(\frac{hW_\infty}{\prod_R^{tot}}\right) + \mathcal{O}\left(\frac{W_\infty}{\prod_R}\right)$

$EMCT(E)$ variants to limit worst case cost

- 1 $EMCT_\alpha(E)$: check a proportion α of $G^\leq(T)$
- 2 $EMCT^K(E)$: check $\min(K, |G^\leq(T)|)$ tasks in $G^\leq(T)$

Definition (Minimal number of initiators)

Let $k \leq n_F$ and $V \subset \mathcal{V}_t$.

- *minimum number of initiators* with respect to V and k :

$$\gamma_V(k) = \min |G^\leq(V) \cap \mathcal{I}(F)| \quad \text{for } \begin{cases} |F| \geq k \\ G^\leq(V) \cap \mathcal{I}(F) \neq \emptyset \end{cases}$$

- *minimal initiator ratio* : $\Gamma_V(k) = \frac{\gamma_V(k)}{|G^\leq(V)|}$.

Note : $n_q \leq n_F$ is the smallest number of falsified tasks
 $\implies \gamma_G(n_q)$ is the smallest n_I possible.

$EMCT_\alpha(E)$ variants to limit worst case cost

$EMCT_\alpha(E)$

Input: Execution E represented by G composed of dependent tasks.

Output: The correctness of E (FALSIFIED or CORRECT)

Uniformly choose one task T in G ;

$n_\alpha \leftarrow \lceil \alpha |G^\leq(T)| \rceil$; //number of tasks to re-execute

Define $\mathcal{T}_\alpha \subset G^\leq(T)$ composed of n_α tasks uniformly chosen in $G^\leq(T)$;

// Re-execution of \mathcal{T}_α on R to detect initiators

forall $T_j \in \mathcal{T}_\alpha$ / T_j as not yet been checked **do**

$\hat{o}(T_j, E) \leftarrow \text{ReexecuteOnVerifier}(T_j, i(T_j, E))$;

if $o(T_j, E) \neq \hat{o}(T_j, E)$ **then**

return FALSIFIED;

end

return CORRECT;

$EMCT_\alpha(E)$ variants to limit worst case cost


- Let $N_{\varepsilon,q,\alpha} = \begin{cases} \left\lceil \frac{\log \varepsilon}{\log(1-q\alpha\Gamma_G(n_q))} \right\rceil & \text{if } 0 < \alpha \leq 1 - \Gamma_G(n_q) \\ N_{\varepsilon,q} = \left\lceil \frac{\log \varepsilon}{\log(1-q)} \right\rceil & \text{otherwise.} \end{cases}$

Theorem (Probabilistic certification by $EMCT_\alpha(E)$)

- $\mathcal{A}(E, \varepsilon) : N_{\varepsilon,q,\alpha}$ calls to $EMCT_\alpha(E)$
- Expected cost per call : $C_G = \lceil \frac{\alpha}{n} \sum_{T \in G} |G^\leq(T)| \rceil$
- On average $W_1^C \leq \alpha \frac{N_{\varepsilon,q,\alpha}}{n} W_\infty \sum_{T \in G} |G^\leq(T)|$ and $W_\infty^C = \mathcal{O}(W_\infty)$

$EMCT_{\alpha}(E)$ variants to limit worst case cost

Impact of α on $N_{\varepsilon,q,\alpha}$ (with $\varepsilon = 5\%$ and $q = 1\%$)



Certification algorithms comparison

Test T :	MCT §4	$EMCT$ §5.2	$EMCT_\alpha$ §5.3	$EMCT^1$ §5.4	
# T detected faulty	$n_I \geq \left\lceil \frac{(d-1)n_F}{d^h-1} \right\rceil$	$n_q = \lceil n \cdot q \rceil$	$n_q \alpha \Gamma_T(n_q)$ or n_q	$n_q \Gamma_T(n_q)$	
$P_{error}(T)$	$1 - \Gamma_G(n_q) \leq 1 - \left\lceil q \frac{(d-1)}{d^h-1} \right\rceil$	$1 - q$	$1 - q \alpha \Gamma_T(n_q)$ or $1 - q$	$1 - q \Gamma_T(n_q)$	
N^T : convergence to ϵ	$\left\lceil \frac{\log \epsilon}{\log(1-\Gamma_G(n_q))} \right\rceil$	$\left\lceil \frac{\log \epsilon}{\log(1-q)} \right\rceil$	$\left\lceil \frac{\log \epsilon}{\log(1-q\alpha\Gamma_G(n_q))} \right\rceil$ or $\left\lceil \frac{\log \epsilon}{\log(1-q)} \right\rceil$	$\left\lceil \frac{\log \epsilon}{\log(1-q\Gamma_G(n_q))} \right\rceil$	
exact C_G	1	$ G^{\leq}(T) $	$\lceil \alpha G^{\leq}(T) \rceil$	1	
avg. C_G (n tasks, height h)	G	1	$\overline{ G^{\leq} }$	$\lceil \alpha \overline{ G^{\leq} } \rceil$	1
	Tree	1	$h + 1 = \Theta(\log n)$	$\lceil \alpha(h + 1) \rceil = \Theta(\alpha \log n)$	1
	Fork-Join	1	$h + 3 = \Theta(\log n)$	$\lceil \alpha(h + 3) \rceil = \Theta(\alpha \log n)$	1
W_1^C : N^T calls to T	G	$N^{MCT} W_\infty$	$N^T W_\infty \overline{ G^{\leq} }$	$\alpha N^T W_\infty \overline{ G^{\leq} }$	$N^{EMCT^1} W_\infty$
	Tree	$N^{MCT} W_\infty$	$\mathcal{O}(h W_\infty)$	$\mathcal{O}(\alpha h W_\infty)$	$N^{EMCT^1} W_\infty$
	Fork-Join	$N^{MCT} W_\infty$	$\mathcal{O}(h W_\infty)$	$\mathcal{O}(\alpha h W_\infty)$	$N^{EMCT^1} W_\infty$
W_∞^C	$\mathcal{O}(W_\infty)$	$\mathcal{O}(W_\infty)$	$\mathcal{O}(W_\infty)$	$\mathcal{O}(W_\infty)$	$\mathcal{O}(W_\infty)$

Conclusion

Result-checking for distributed computations

- Approach based on macro-dataflow analysis
 - ↪ deals with task dependencies
- “No” hypothesis on attacker behaviour
- Monte-carlo certification [E]MCT[X]
 - ↪ low overhead for recursive/Fork-Join programs
 - ↪ high overhead in general ($\rightarrow EMCT_{\alpha}(E)$ and $EMCT^K(E)$)
 - ↪ validation on medical application (not presented here)

Perspective/Current work

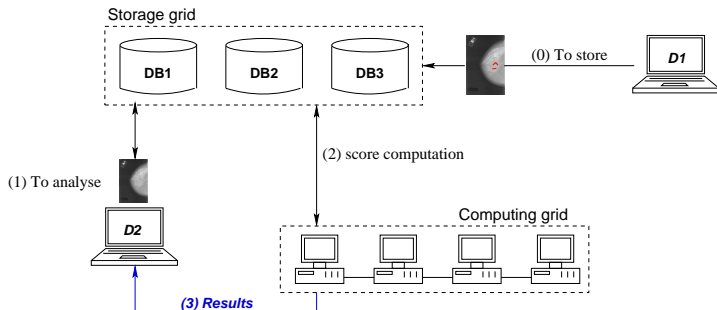
- Atlantic city extension
 - ↪ verifier not so accurate
 - ↪ if test fails, probability to stand below the tolerance threshold?
- Dealing with $n_F < \lceil n.q \rceil$
 - ↪ Algorithm-Based Fault-Tolerance (ABFT)

Thanks for your attention...

Questions ?




Proof of concept



Breast cancer lesions detection in mammograms [Varrette& al.06]

- statistical comparison on a database of studied cases

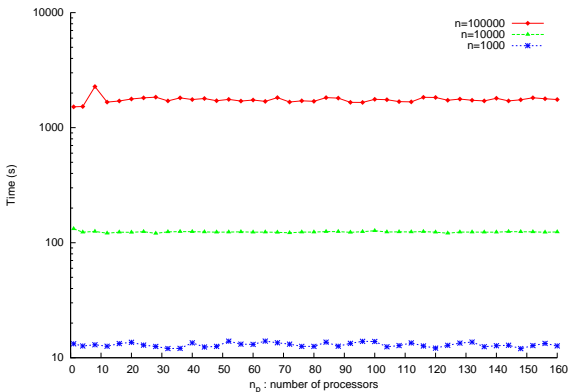
Experimental protocol



Experimentations

Deployment on Grid5000; $\varepsilon = 0.001$, $q = 0.01$ ($N_{\varepsilon,q} = 688$)

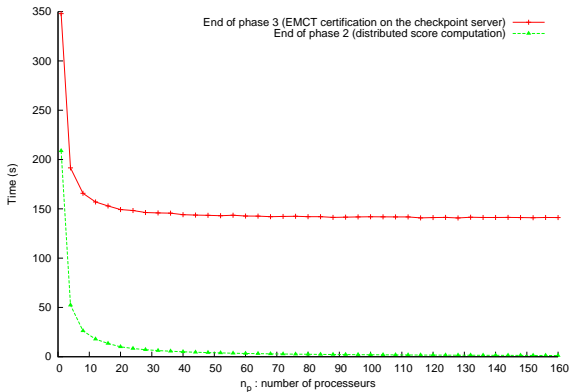
Time required to deploy the images on the grid



Experimentations

Deployment on Grid5000; $\varepsilon = 0.001$, $q = 0.01$ ($N_{\varepsilon,q} = 688$)

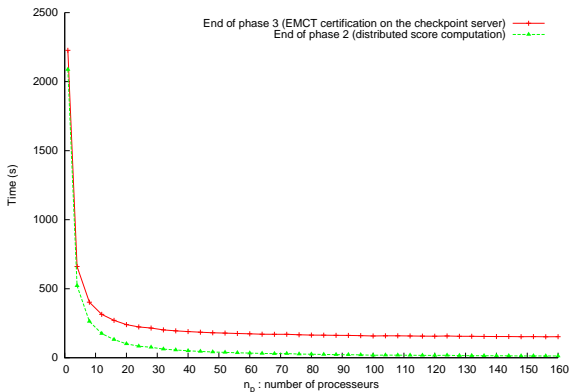
Scores computation + certification : 1000 tasks



Experimentations

Deployment on Grid5000; $\varepsilon = 0.001$, $q = 0.01$ ($N_{\varepsilon,q} = 688$)

Scores computation + certification : 10000 tasks



Experimentations

Deployment on Grid5000; $\varepsilon = 0.001$, $q = 0.01$ ($N_{\varepsilon,q} = 688$)

Scores computation + certification : 100000 tasks

