

CUDA Constraint Programming for AI Gaming in the Cloud

Stéphane Cardon
Eric Jacopin

Ecoles de Saint-Cyr Coëtquidan
France


WHY?

- GPGPU the future of AI in games ^{a,b}
- Games use planning ^{c,d,e}
(plan of 10 actions for 15 bots)
- Constraint Programming can scale up planning ^f
- Validate TRL 3
- Don't touch my graphics card

GOAL

Can we scale up
Constraint Programming
using CUDA ?

Experiments - Random Instances



Intel Xeon E5-2680 (2.70 GHz)
64Gb of RAM
GeForce TITAN
CUDA 6.5 64 bits

Many constraints
few compatible values

Future

- Testing on games' instance (Simple FPS)
- Multi-GPUs, Tesla server
- Reach TRL 4 to 6

Conclusion

- How to modelize planning problem to such constraint satisfaction problem
- GPU-adapted consistency

References

a - Jeff ORKIN - Three States and a Plan : The A.I. Of F.E.A.R. - Proceedings of the Game Developer Conference (2006)
 b - Dana NAU, Yue CAO, Amnon LOTEM & Héctor MUÑOZ-AVILA - The SHOP Planning System - AI Magazine 22(3), AAAI Press (Fall 2001)
 c - William BLEWITT, Gary USHAW & Graham MORGAN - Applicability of GPGPU Computing to Real-Time AI Solutions in Games - Computational Intelligence and AI in Games (2013)
 d - Alex CHAMPANDARD & Andrew RICHARDS - Massively Parallel AI on GPGPUs with OpenCL or C++ - Proceedings of the Game Developer Conference (2014)
 e - Damian SULEWSKI, Stefan EDELKAMP & Peter KISSMANN - Exploiting the Computational Power of the Graphics card : Optimal State Space Planning on th GPU - Proceedings of International Conference on Automated Planning and Scheduling (2011)
 f - Peter VAN BEEK & Xianguang CHEN - CPlan : A Constraint Programming Approach to Planning - Proceedings of National Conference on Artificial Intelligence (1999)