Brains-in-the-Cloud

A New Hope for Pervasive Robotics

Iuliu Vasilescu
Politehnica University of Bucharest,
Computer Science Department
iuliuv@tiaresearch.com — 0724 849 008

Center for Research in Computing

Abstract

Most of the current field-robotics research takes pride in creating fully autonomous robots that are self-contained, including computational resources. We propose to take as much computation as possible off the robots and move it to a computer cloud. The benefits will be threefold: (1) much more computationally-complex tasks can be tackled with a remote, unrestricted computer, (2) algorithms that require training (ex. machine learning) will have real-time access to data streams from multiple robots, (3) the computation resources will be shared, improving their use efficiency and therefore lowering the overall system cost. We will try to demonstrate and make full use of these advantages, while identifying and addressing the specific challenges of this approach: (1) dealing with communication limited bandwidth and latency, (2) achieving tolerance to communication faults, (3) splitting the computation tasks between the robot's and the remote computer, (3) maintaining data anonymity. We will concentrate our efforts in solving real-time tasks such as navigation and object identification in the context of aerial drones. A key component will be using the computation latency as a fundamental primitive in the algorithm design process, the same way one would consider that a robot's mechanical motions are not instant.

Introduction

Despite years of research, autonomous robots failed to enter our lives at any significant scale. Beyond mechanical and electronic parts, a functional autonomous robot requires working solutions to some hard computer science problems: sensing and making sense of the environment, planning actions, dealing with uncertainty and all of this at a reasonable cost for the task performed

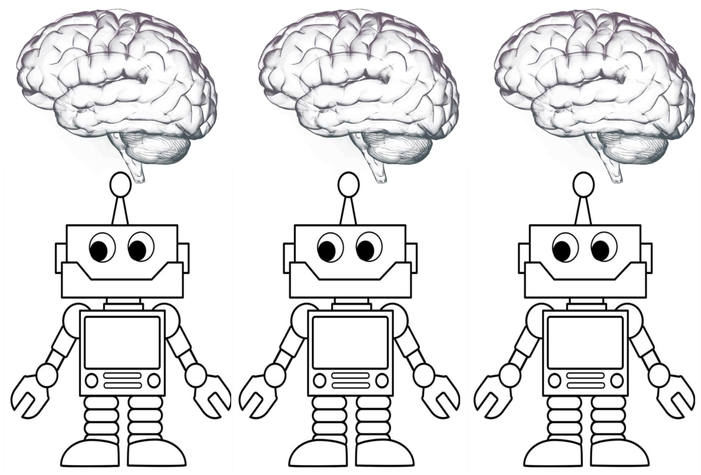

Figure 1: Traditional robots carry enough computation power for all the tasks that they need to solve. This is inefficient: computation needs typically come in bursts and the on-board computation capability is wasted when the robot is not in use.

Traditionally most of the computational tasks are carried out using on board computer, somehow this idea being a staple of having a fully autonomous robot. This poses multiple limitations on the amount of computation power available to the robot. In addition, some of the problems such as object recognition or grasping planning were simply intractable up until recently for the degree of confidence required. Having a full understanding and representation of the environment seems like the only reliable approach, with only a few dissenting opinions[2, 3].

Approach

We propose to move as much computation as possible off the robots, to a computer cloud.

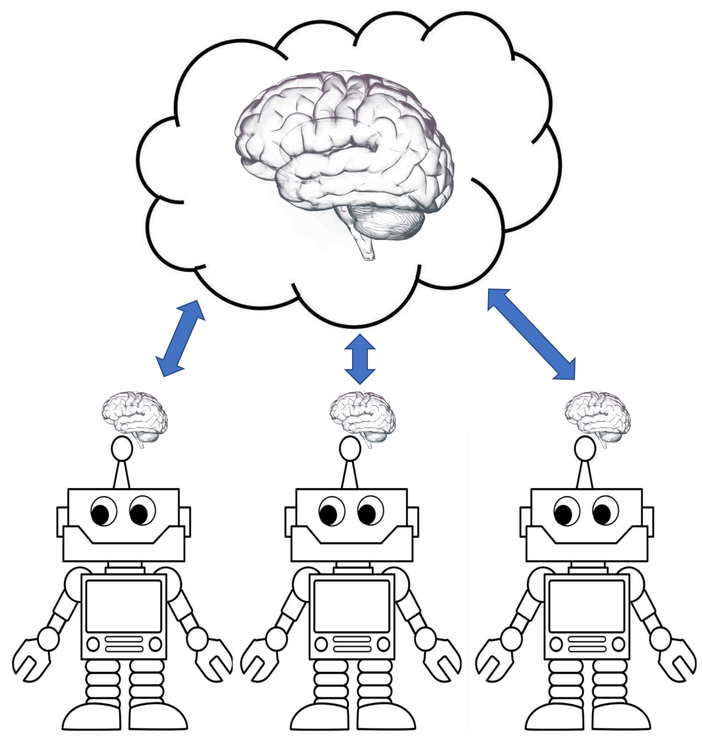

Figure 2: We propose to share the computation power of a computer cloud among several robots, by splitting the computation in parts that are done on board and parts that are done in the cloud. The robots will been a "smaller" on-board brain.

Our plan for this project consists in three activities:

1. We will develop and and build a mobile robotic platform connected to a computer-cloud infrastructure.
2. We will develop new algorithms and we will adapt existing state-of-the-art algorithms so that they can take advantage of the combined capabilities of the robots and computer cloud, while mitigating the inherent challenges of this architecture.
3. We will test the new algorithms on our platform in simulation and in the field, compare the results with the existing state-of-the-art systems.

1 State of the Art

There are a few comprehensive surveys[9, 5, 7] detailing the current state of the art in "cloud robotics", a broad term coined relatively recently.

- **Access to Big Data** such as: ImageNet[4].
- **Access to High Performance Computing**: ex grasping[8].
- **Data Sharing**: collective robot learning[1].
- **Human Computation**: path planning[6].

Research Questions

- **How will computation delay affect the algorithm design process in general?**
- **What types of computation makes sense to move away from the robot?**
- **How can we mitigate the effects of networking constraints: bandwidth, latency, unreliability?**
- **What is the best way to divide a given task between local and remote computation?**
- **What can we gain from doing the processing in a shared computing cloud?**
- **How can we insure data anonymity and security?**

References

[1] Dmitry Berenson, Pieter Abbeel, and Ken Goldberg. A robot path planning framework that learns from experience. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 3671–3678. IEEE, 2012.

[2] Rodney A Brooks. Intelligence without representation. *Artificial intelligence*, 47(1-3):139–159, 1991.

[3] Rodney A Brooks et al. Intelligence without reason. *Artificial intelligence: critical concepts*, 3:107–63, 1991.

[4] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 248–255. IEEE, 2009.

[5] Ken Goldberg and Ben Kehoe. Cloud robotics and automation: A survey of related work. *EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2013-5*, 2013.

[6] Juan Camilo Gamboa Higuera, Anqi Xu, Florian Shkurti, and Gregory Dudek. Socially-driven collective path planning for robot missions. In *Computer and Robot Vision (CRV), 2012 Ninth Conference on*, pages 417–424. IEEE, 2012.

[7] Guoqiang Hu, Wee Peng Tay, and Yonggang Wen. Cloud robotics: architecture, challenges and applications. *IEEE network*, 26(3), 2012.

[8] Ben Kehoe, Akihiro Matsukawa, Sal Candido, James Kuffner, and Ken Goldberg. Cloud-based robot grasping with the google object recognition engine. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 4263–4270. IEEE, 2013.

[9] Ben Kehoe, Sachin Patil, Pieter Abbeel, and Ken Goldberg. A survey of research on cloud robotics and automation. *IEEE Transactions on automation science and engineering*, 12(2):398–409, 2015.

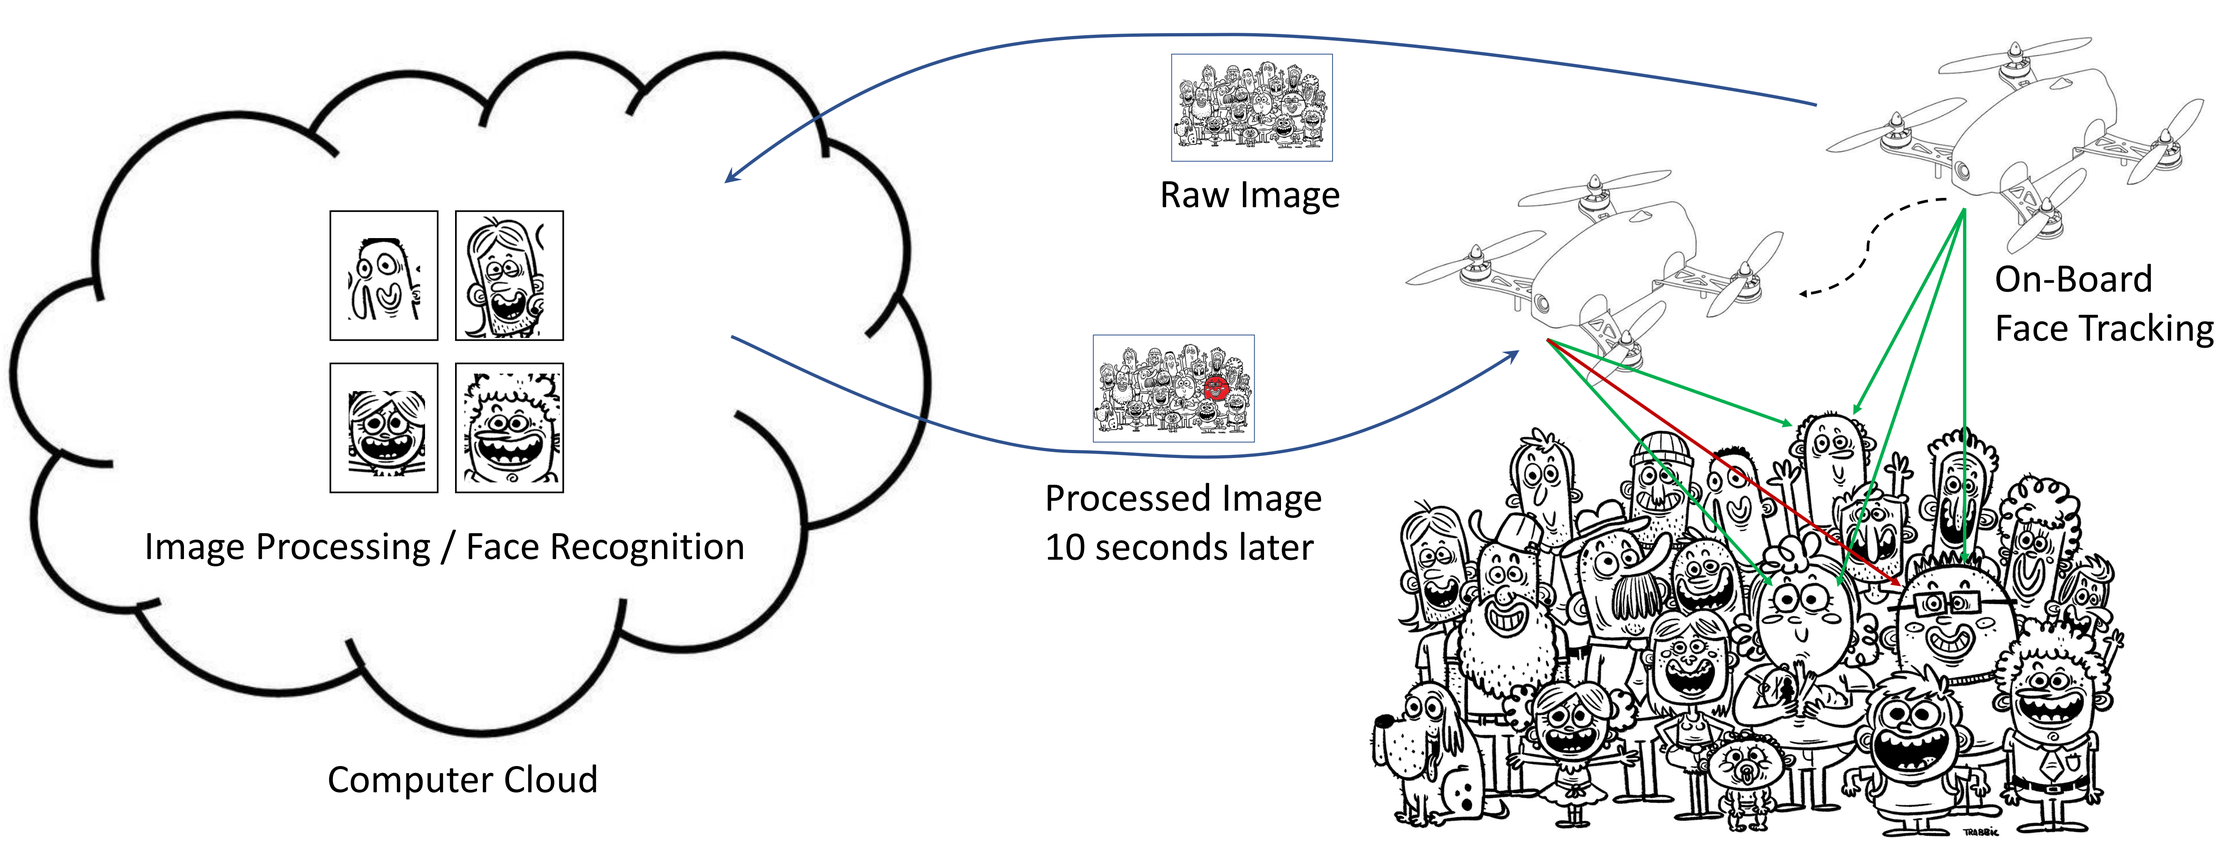

Figure 3: A typical application for brain-in-the-cloud. A drone, searching for a specific person in a crowd, takes pictures and sends them to a computer cloud for processing. While waiting for the identification results keeps track of everybody's movement, using on-board computation. When the results are back from the cloud, the drone identifies the person in the current view using the tracking info. The computer cloud uses its storage and computation power to solve the "diffficult" task of face identification, while the drone only needs to do the "ligher" task of visual tracking.