Speaker: Yehuda Afek, Tel Aviv University, Israel

Title: Asynchrony From Synchrony

Abstract: Synchronous Message Adversary: A Simple Characterization of Asynchronous Computation. In this talk we investigate synchronous message-passing with reliable processors but different scenarios of message loss. In particular, we investigate and present the minimum set of messages whose delivery must be guaranteed to ensure the equivalence of this model to wait-free shared-memory. Delivery guarantees were investigated in the past at the other extreme - delivery guarantees that ensure consensus. Because message failure is a more refined type of fault than processor fail-stop we were able to use in this paper an extremely simple message passing model to characterize exactly what is computable in waitfree read-write shared memory. We use a synchronous complete network where in each round a subset from a defined family of subsets of messages, must be successfully delivered. With this model we obtain an extremely simple characterization derivation of the Herlihy-Shavit condition that equates the wait-free read-write model with a subdivided-simplex. We show how each step in the computation inductively takes a subdividedsimplex and further subdivides it in the simplest way possible, making the characterization of read-write wait-free widely accessible.

Speaker: Zohir Bouzid, ENST, France

Title: Byzantine Processes in a Topological Perspective

Abstract: The topological approach to distributed computing is holistic: it allows to visualize all possible executions of a given model, represent them as an abstract geometrical structure, and derive complexity and computability bounds based on the properties of this structure. However, when it comes to modeling executions with Byzantine processes, the approach is not easy to apply. This is because a Byzantine process may arbitrarily deviate from its protocol, and thus modeling its ßtatecan be tricky. In this talk, we discuss how to account for Byzantine behavior in a topological characterization of task solvability.

Speaker: Nicolas Braud-Santoni, ENS Chachan - Antenne de Bretagne, France

Title: Randomized Byzantine Agreement

Abstrct: Until recently, (probabilistic) Byzantine Agreement algorithms required at least O(n) communication, and polylog time. The algorithm presented here requires polylog time and communication, using reduction from weaker problems

Speaker: Armando Castaneda, Technion, Israel

Title: Pareto Optimal Solutions to Consensus and Set Consensus

Abstract: Abstract: While the very first consensus protocols were designed to match the worst-case lower bound, deciding in exactly t + 1 rounds in all runs, it was soon realized that they could be strictly improved upon by early stopping protocols. These always decide in no more than t + 1 rounds, but often do so faster. For simultaneous consensus and many of its variants, all-case optimal protocols have been presented: For every behavior of the adversary, the all-case optimal protocol decides as soon as any protocol can decide in a run with the same adversarial behavior.

The requirement of simultaneity, however, often imposes a significant penalty. Unfortunately, it is known that all-case optimal protocols for eventual consensus do not exist. Halpern, Moses and Waarts study consensus protocols with a property that we call Pareto-optimality. P is Pareto-optimal if it is not strictly dominated. Namely, no protocol Q can decide as fast as P for all adversary moves, while allowing at least one process to decide strictly earlier, in at least one instance. They present a general logical transformation of any consensus protocol to a Pareto-optimal one that dominates it, and present a particular Pareto-optimal consensus protocol. In this talk we well see Pareto-optimal solutions to consensus and k-set consensus in synchronous message-passing contexts with benign failures. The solution for k-set consensus strictly dominates all known early stopping solutions in this model. Moreover, being Pareto optimal, it cannot be improved on.

Speaker: Eli Gafni, UCLA, USA

Title: A Generalized Asynchronous Computability Theorem

Abstract: The celebrated Asynchronous Computability Theorem (ACT) of two decades ago, characterizes wait-free solvability, by casting it as a question about the existence of a map between complexes. Alas, solvability in other models, e.g. t resiliency, or obstruction-free, were dealt with in ad-hoc manner, through the BG simulation, or using "understanding" of the model.

Here, we derive a general characterization for any model which is a subset of the runs of the wait-free model. We cast our analysis within the iterated model of distributed computing as it is more refined and structured that the non-iterated one. When our characterization is instantiated to the wait-free non-iterated model via a simulation we get ACT.

We exhibit the usefulness of the characterization by providing a task whose operational solution in the t-resilient model is known to us, but is quite involved algorithm. In contrast, deducing its solvability by the characterization is straight forward.

After getting all the definitions and subtleties settled, the theorem is astonishingly easy to derive.

Speaker: Eric Goubault, CEA Saclay, France

Title: Fault-tolerant Protocols and Directed Algebraic Topology

Abstract: I will develop in this talk the first connections between the directed algebraic topological approach to the description of concurrent and distributed systems (in particular the trace space approach of Martin Raussen) and the algebraic topological approach to characterizing fault-tolerant protocols. Much is due to Christine Tasson, Samuel Mimram, and interaction with Maurice Herlihy and Sergio Rajsbaum.

Speaker: Maurice Herlihy, Brown University, USA

Title: The Topology of Asynchronous Byzantine Colorless Tasks

Abstract: Tools adapted from combinatorial topology have been successful in characterizing task solvability in synchronous and asynchronous crash-failure models [2]. We extend the approach to asynchronous Byzantine systems: we give the rst theorem with necessary and su cient conditions to solve arbitrary colorless tasks in such model, capturing the relation between the total number of processes, the number of faulty processes, and the topological structure of the task's simplicial complexes. Joint work with Hammurabi Mendes and , Christine Tasson.

Speaker: Achour Mostefaoui, University Nantes, France

Title: An Online Space Efficient Renaming Algorithm

Abstract: Distributed collaborative editing systems allow users to work distributed in time, space and across organizations. Trending distributed collaborative editors such as Google Docs, Etherpad or Git have grown in popularity over the years. A new kind of distributed editors based on a family of distributed data structure replicated on several sites appeared recently. This talk is about such data structure that represent a distributed sequence of basic elements that can be lines, words or characters. The possible operations on this sequence are the insertion and the deletion of elements. The aim of the approach is to be more decentralized and to better scale in terms of the number of participants and operations. However, its space complexity is linear with respect to the total number of inserts and the insertion points in the document. This makes the overall performance of such editors dependent on the editing behaviour of users. In this talk an adaptive allocation strategy for a sequences will be presented. This problem if abstracted becomes an online renaming problem. The proposed solution achieves in the average a sub-linear spatial-complexity (size of the ids) whatever is the editing behaviour.

Speaker: Thomas Nowak, Ecole Polytechnique, France

Title: The Transient Behavior of Long Walks and Applications

Abstract: What are the possible lengths of walks between two given vertices in a graph? It turns out that the set of possible lengths has a periodic behavior after an initial transient phase. The study of this transient phase has been ongoing research from the '50s until today. Quite a number of results involving certain graph parameters are available. More generally, one can ask for the maximum weight of walks between two vertices with a given length in a weighted digraph. The answer also shows a periodic behavior. The analytic study of this more general case has many applications in distributed systems, like some classical synchronizers, routers, or schedulers. In this talk we give new results on the length of the transient phase and we discuss their applications. Our bounds are the first that are both asymptotically tight and linear in the system size for various classes of systems. The bounds give guidelines for system design to optimize performance.

Speaker: Ami Paz, Technion, Israel

Title: Upper Bound on the Complexity of Solving Hard Renaming

Abstract: The M-renaming task requires n processes, each starting with a unique input name (from an arbitrary large range), to coordinate the choice of new output names from a range of size M. This talk presents the first upper bound on the complexity of hard renaming, i.e., (2n-2)-renaming, when n is not a prime power. It is known that (2n-2)-renaming can be solved if and only if n is not a prime power; however, the previous proof of the "if" part was non-constructive, involving a topological approximation theorem; in particular, it did not yield a concrete upper bound on the complexity of the resulting protocol.

Speaker: Michel Raynal, IRISA, France

Title: Synchrony Weakened by Message Adversaries vs Asynchrony Restricted by Failure Detectors

Abstract: A message adversary is a daemon that suppresses messages in round-based message-passing synchronous systems in which no process crashes. A property imposed on a message adversary defines a subset of messages that cannot be eliminated by the adversary. It has recently been shown that when a message adversary is constrained by a property denoted TOUR (for tournament), the corresponding synchronous system and the asynchronous crash-prone read/write system have the same computability power for task solvability. This talk will introduce new message adversary properties (denoted SOURCE and QUORUM), and shows that the synchronous round-based systems whose adversaries are constrained by these properties are characterizations of classical asynchronous crash-prone systems (1) in which processes communicate through atomic read/write registers or point-to-point message-passing, and (2) enriched with failure detectors such as Omega and Sigma. Hence these properties characterize maximal adversaries, in the sense that they define strongest message adversaries equating classical asynchronous crash-prone systems. They consequently provide strong relations linking round-based synchrony weakened by message adversaries with asynchrony restricted with failure detectors. This not only enriches our understanding of the synchrony/asynchrony duality, but also allows for the establishment of ameaningful hierarchy of property-constrained message adversaries.

Speaker: Stefan Schmid, TU Berlin, Germany

Title: Self-stabilizing and self-optimizing distributed datastructures

Abstract: In this talk, I will present techniques to construct and maintain an efficient distributed datastructure, such as a skip graph. Our distributed algorithm is self-stabilizing in the sense that from any weakly-connected initial state, it converges quickly back to the desired topology. I will prove a time complexity of $O(log^2 n)$, and also show that joins and leaves are handled quickly. Subsequently, I will initiate the discussion of more dynamic distributed datastructures, which even self-adjust to the demand. In the spirit of classical datastructures and splay trees, I introduce the concept of splay networks and provide first insights on how to design such datastructures.

Speaker: Ravi Srivatsan, TU Berlin, Germany

Title: Safety in distributed computing

Abstract: I will revisit the formal notions of safety in distributed computing and techniques for proving safety of correctness properties like linearizability and opacity.

Speaker: Julien Stainer, IRISA, France

Title: Computing in the Presence of Concurrent Solo Executions (joint work with S. Rajsbaum and M. Raynal)

Abstract: In a wait-free model any number of processes may crash. A process runs solo when it computes its local output without receiving any communication from other processes, either because they crashed or they are too slow. While in wait-free shared-memory models at most one process may run solo in an execution, any number of processes may have to run solo in an asynchronous wait-free message-passing model. This talk is on the computability power of models in which several processes may concurrently run solo. It first introduces a family of round-based wait-free models, called d-solo models, 1 $\Leftarrow$ d $\Leftarrow$ n, where up to d processes may run solo. The talk will present a characterization of the colorless tasks that can be solved in each d-solo model. The talk will also introduce the (d,epsilon)-solo approximate agreement problem, which generalizes ϵ -approximate agreement. It will show that (d,ϵ)-solo approximate agreement can be solved in the d-solo model, but it cannot be solved in the (d+1)-solo model. The talk also studies the relation linking d-set agreement and (d,ϵ)-solo approximate agreement in asynchronous wait-free message-passing systems. These results establish for the first time a hierarchy of wait-free models weaker than the basic read/write model, that are nevertheless still strong enough to solve many tasks. The new failure model, based on taking into account solo executions, has message passing as well as shared-memory instantiations.