

## Σ Ε Μ Ι Ν Α Ρ Ι Ο Τ Μ Η Μ Α Τ Ο Σ

**ΟΜΙΛΗΤΗΣ:**



**Αναστασιάδης Στέργιος**  
Επικ. Καθηγητής Τμήματος  
Μηχανικών Η/Υ & Πληροφορικής  
Πανεπιστήμιο Ιωαννίνων

**ΗΜΕΡΟΜΗΝΙΑ:**

Παρασκευή, 16 Ιανουαρίου 2015

**ΩΡΑ:**

12:00

**ΑΙΘΟΥΣΑ:**

Αίθουσα Σεμιναρίων (ισόγειο Ι11)  
Κτήριο Τμήματος Μηχανικών Η/Υ & Πληροφορικής

### Θ έ μ α

## ***«Efficiency and Consistency in Multistream Storage»***

### ***Περίληψη***

Synchronous small writes play critical role in system availability, because they are used to safely log recent state modifications for fast recovery from crashes. Demanding systems typically dedicate separate devices to logging for adequate performance during normal operation and redundancy during state reconstruction. However, storage stacks enforce page-sized granularity in data transfers from memory to disk and consume excessive storage bandwidth to handle small writes. The problem becomes worse in the typical case of multiple concurrent streams, which effectively generate random I/O traffic. In a journaled filesystem, we introduce Wasteless Journaling to coalesce synchronous concurrent small writes of data into full page-sized journal blocks. Additionally, we propose Selective Journaling to automatically activate wasteless journaling on data writes of size below a fixed threshold. We implemented a functional prototype of our design over a widely-used filesystem. We experimentally compare our methods with existing ones using microbenchmarks and application-level workloads on standalone servers and a multi-tier networked system. Coalescing small data updates to the journal preserves filesystem consistency, reduces consumed bandwidth up to several factors, and lowers write latency up to orders of magnitude.