Adaptive Domain-Specific Service Monitoring

Arda Ahmet Ünsal

Görkem Sazara

Barış Aktemur

Hasan Sözer

SERENE 2014

Summary

- An **adaptive** service monitoring approach
- Not only consider generic errors such as *file not found* or *connection timed out*, but also **domain-specific** errors, e.g., codec errors for streamed media
- Evaluation using real-world data from **Smart TV** domain
- **30%** monitoring **cost reduction** with negligible compromise on the quality of monitoring

Motivation: Smart TVs

- Video/Audio content subject to errors
- Provided mainly by third-party services
- Reliability as perceived by the customer

Need for Service Monitoring

- Availability and incompatibility issues are common
  - Interfaces, URLs change without notification
- Domain-specific errors are common
  - Unsupported audio/video codecs
- Services are monitored to avoid errors
  - Do not present the erroneous content
  - Present the content from an alternative provider

Cost of Service Monitoring

- Each service has different error rates for different error types
- Cost of error detection is different for different error types
  - Codec checks are expensive
  - Availability checks are much cheaper
- Idea: Adapt monitoring frequency to the service and the error type to reduce monitoring cost

Contributions

- Adaptation of monitoring frequency for a particular **service** *and* **error type**, based on
  - the temporal history
  - the error rate
- Experimental evaluation on an industrial case study (Smart TVs)

Experimental Setup and Data Collection

- Data Collection Process
  - A service logger that ran on Amazon cloud for 5 weeks, 3 instances at distant locations
  - Logged HTTP status and codec validity
  - 6 providers, 51 services

Collected Data Sets

- Data collection instances
  - from Ireland, USA, and Japan
  - Cosine similarity: Japan-Ireland (0.99), Japan-USA (0.98) and Ireland-USA (0.97)
- Total 132K requests
  - 8K “HTTP 404 not found” errors
  - 9K codec errors

Cost of Monitoring for Smart TV's

- Codec checks
  - Download a piece of the video
  - Play the downloaded piece
  - Uses both network and processor resources
- Cost of monitoring $\propto$ (# of videos checked)
- Quality of monitoring $\propto$ (# of detected errors)

Adaptation of Monitoring Frequency

$$\hat{\mathcal{E}}_i^{codec} = \frac{\mathcal{E}_i^{codec}}{V_i - \mathcal{E}_i^{avail}}$$

$$AER_n^{codec} = \begin{cases} \hat{\mathcal{E}}_0^{codec} & \text{if } n = 0 \\ \alpha \times AER_{n-1}^{codec} + (1-\alpha) \times \hat{\mathcal{E}}_n^{codec} & \text{if } (n > 0) \end{cases}$$

- V_i : total number of video checks during the time period *i*
- $\mathcal{E}_i^{avail}$ the number of availability errors during the time period *i*
- $\mathcal{E}_i^{codec}$: the number of codec errors during the time period *i*
- Accumulated error rate (AER) for codec errors

Frequency Adaptation Schemes

Scheme	Accumulated error rate cutoff values (%)					
F0	-	-	-	-	0	∞
F1	-	-	-	0	0.001	∞
F2	-	-	0	0.001	0.002	∞
F3	-	0	0.001	0.002	0.003	∞
F4	0	0.001	0.002	0.003	0.004	∞
F5	0.001	0.002	0.003	0.004	0.005	∞
F6	0.01	0.02	0.03	0.04	0.05	∞
F7	0.05	0.1	0.15	0.2	0.25	∞
F8	0.1	0.2	0.3	0.4	0.5	∞
F9	1	2	3	4	5	∞
Frequency pattern	1000	100	10	110	1110	1

The monitoring frequency based on accumulated error rates.

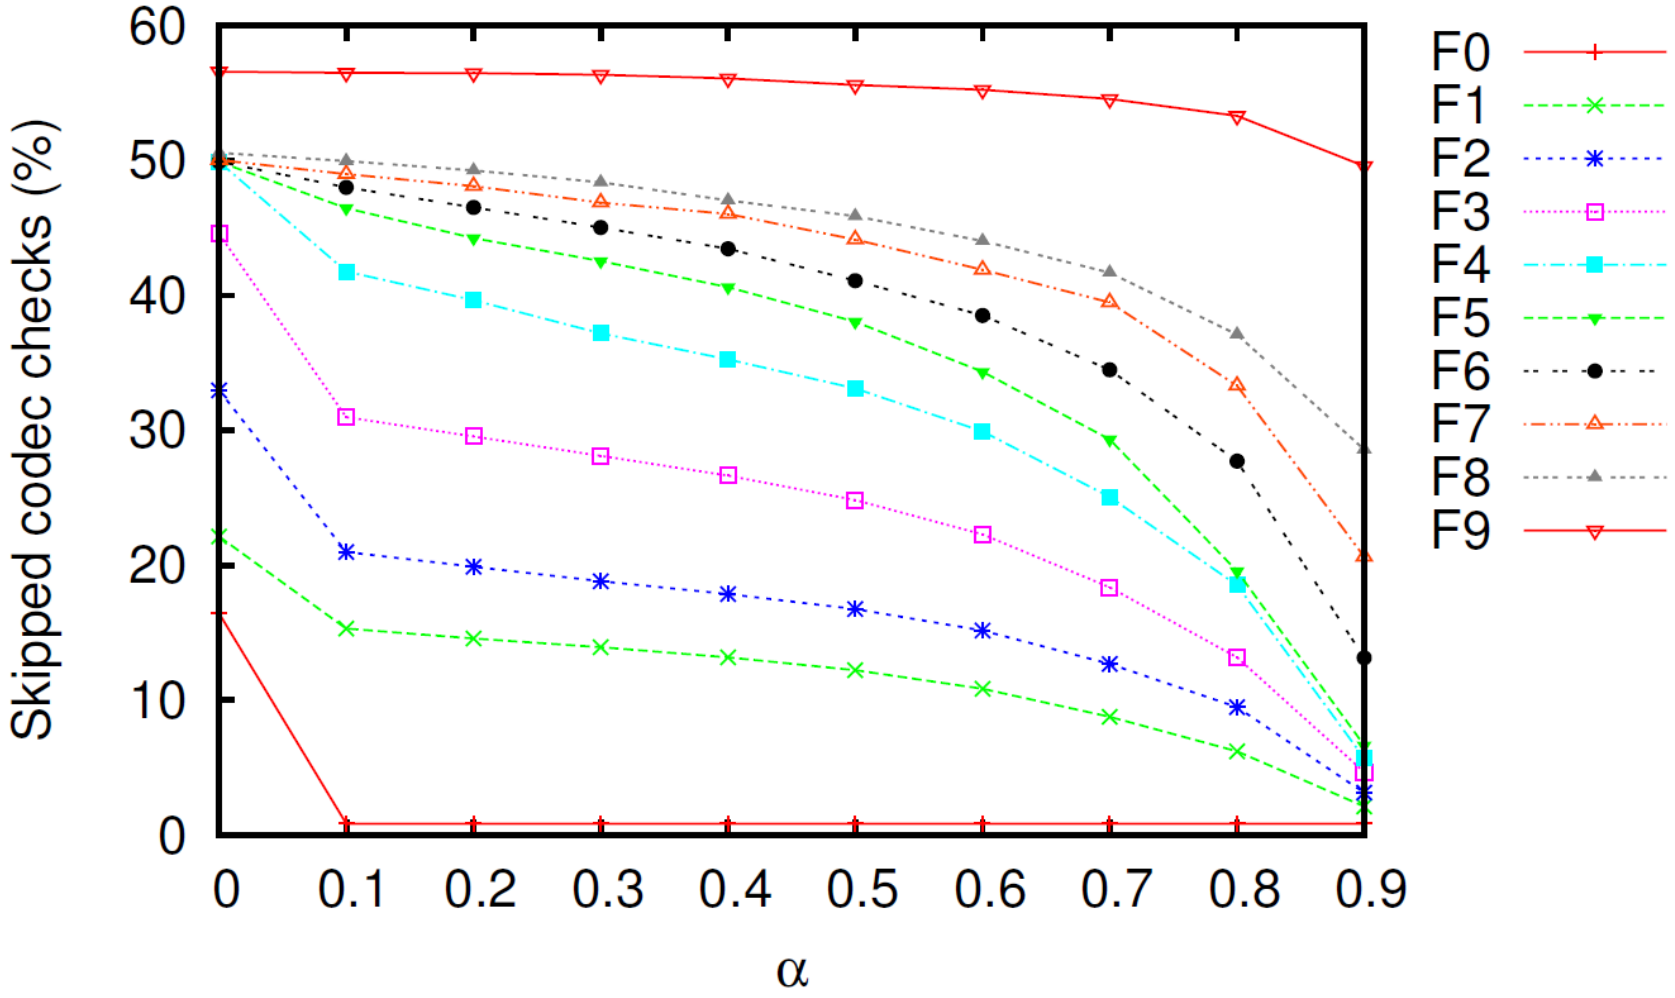
Skipped codec checks (%)
60
50
40
30
20
10
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
α
F0
F1
F2
F3
F4
F5
F6
F7
F8
F9

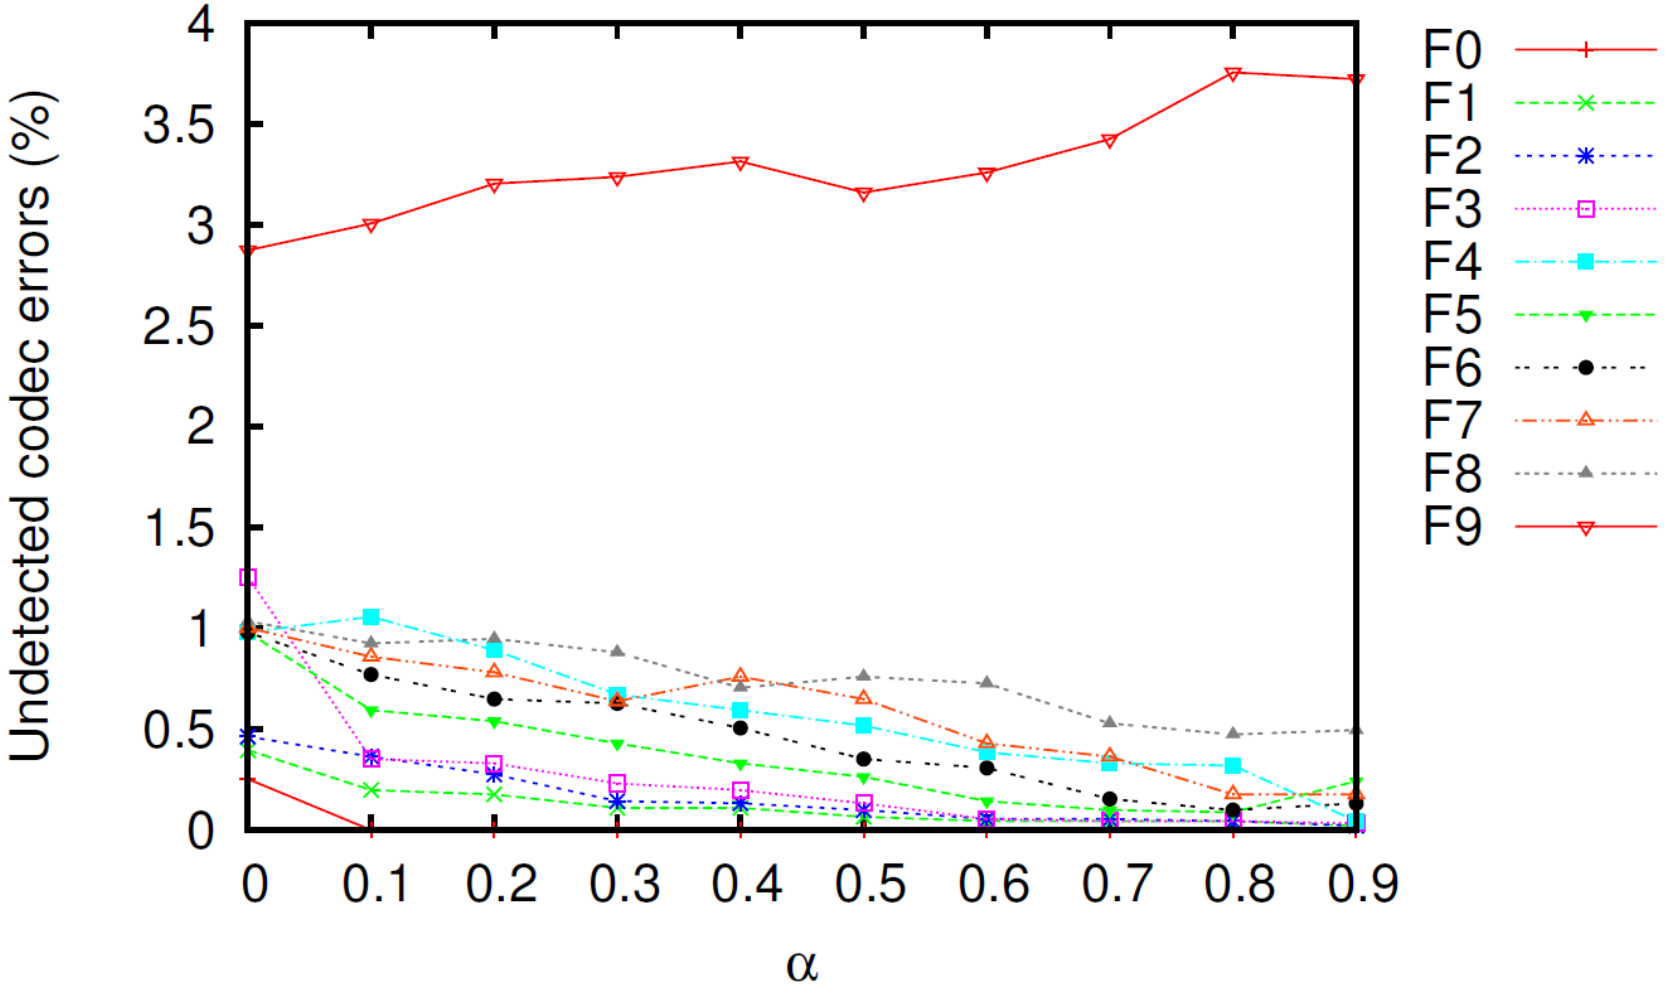
Undetected codec errors (%)
4
3.5
3
2.5
2
1.5
1
0.5
0
0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
α
F0
F1
F2
F3
F4
F5
F6
F7
F8
F9

Effectiveness: Gain vs Loss

$$\mathcal{F} = (\text{rate of skipped checks}) - \beta \times (\text{rate of undetected errors})$$

The effectiveness depends on how much weight, via the β parameter, is given to the undetected errors as opposed to skipped checks.

The value for β can be set by a domain expert or sysadmin.

$$\mathcal{F} = (\text{rate of skipped checks}) - \beta \times (\text{rate of undetected errors})$$

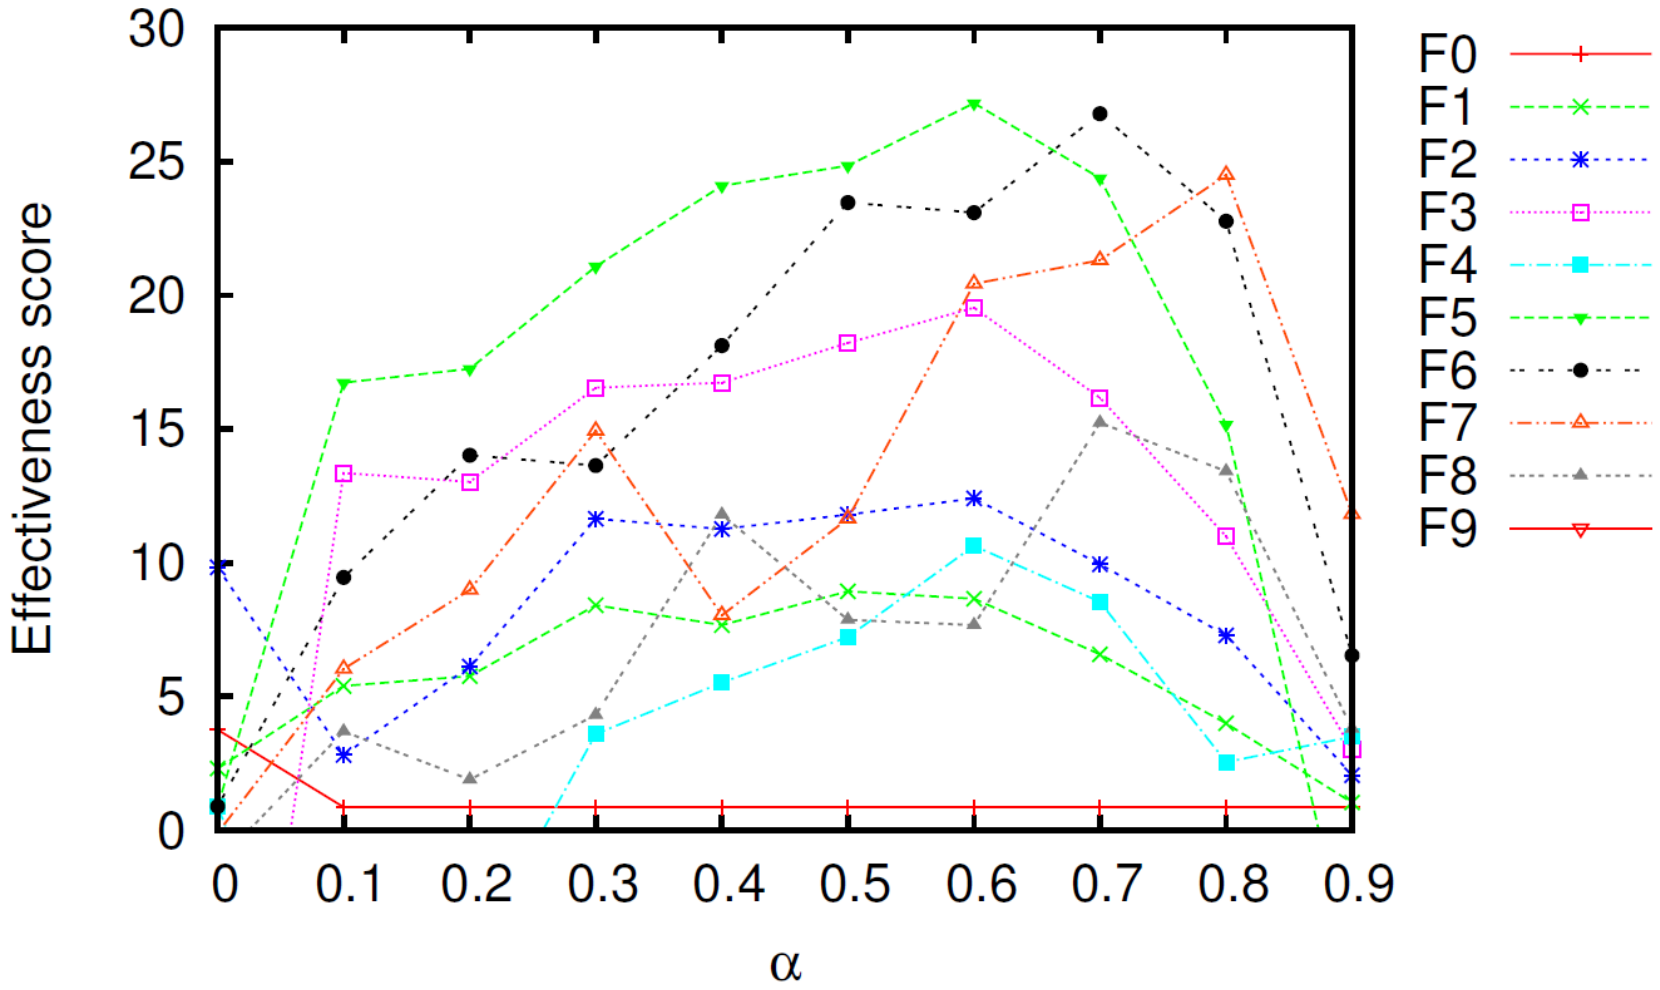

$\beta = 50$

Conclusion

- A novel domain-specific service monitoring approach
- Instance of the approach for the broadcasting and content-delivery domain
- Extensive data set for services used in a commercial Smart TV
- > 30% reduction in monitoring costs without compromising the error detection accuracy significantly

- Applicability to other application domains

$$\mathcal{F} = (\text{rate of skipped checks}) - \beta \times (\text{rate of undetected errors})$$

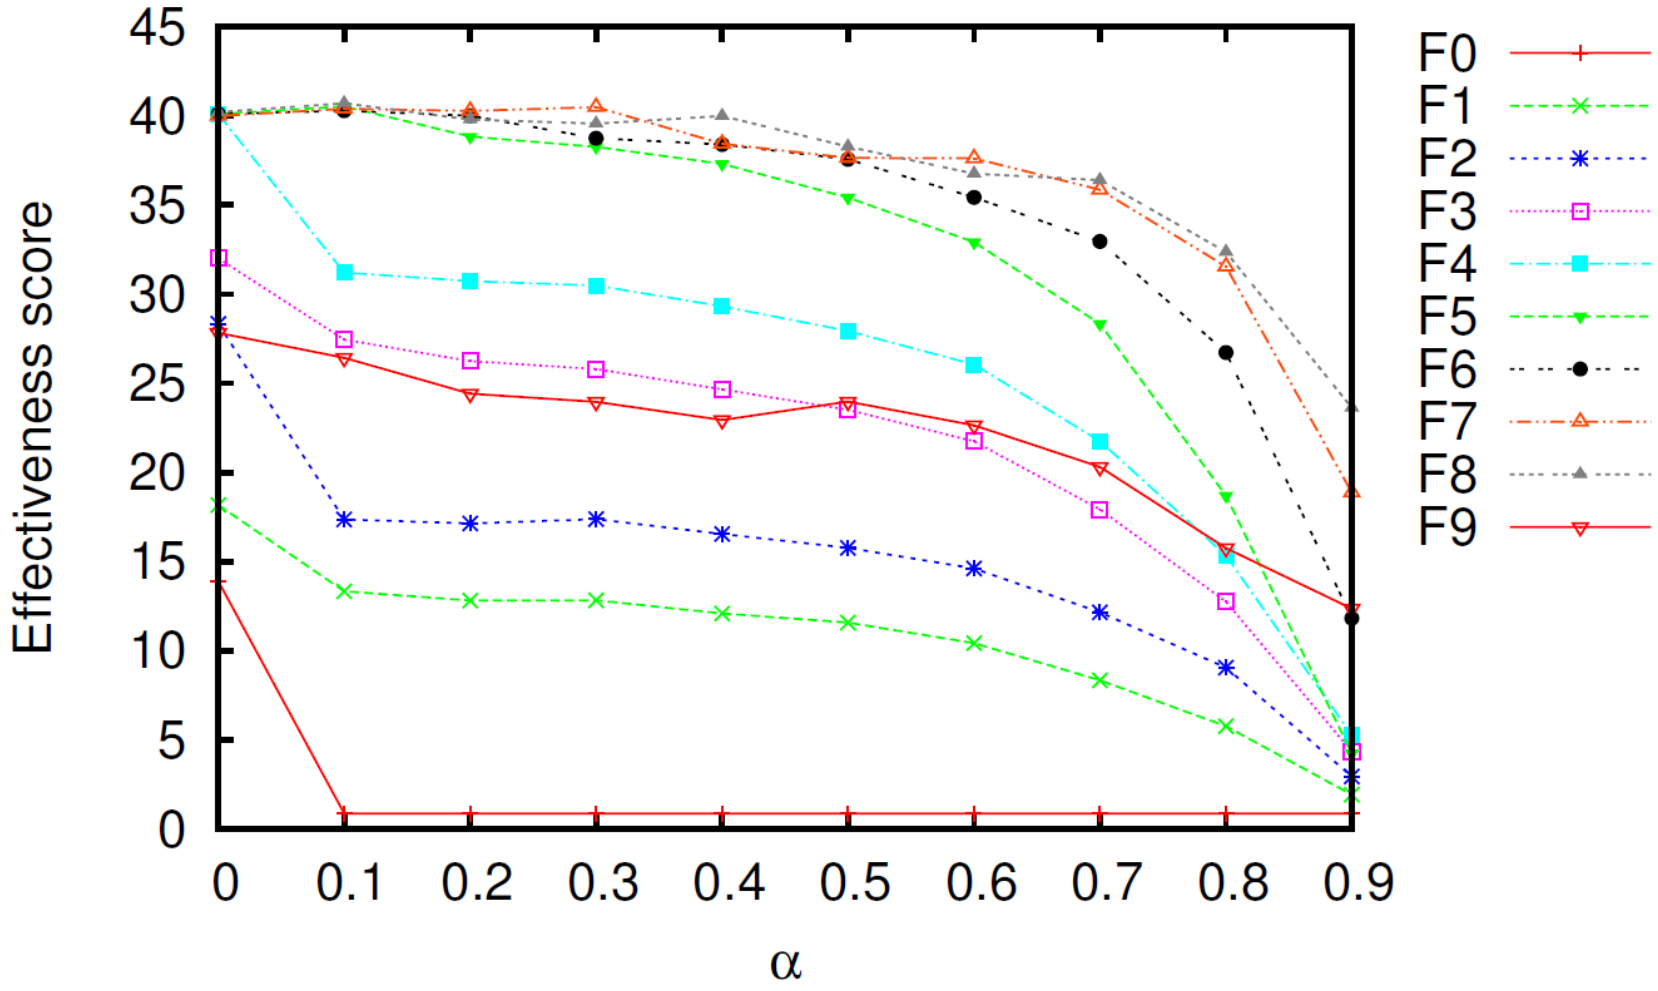

$\beta = 10$

$$\mathcal{F} = (\text{rate of skipped checks}) - \beta \times (\text{rate of undetected errors})$$

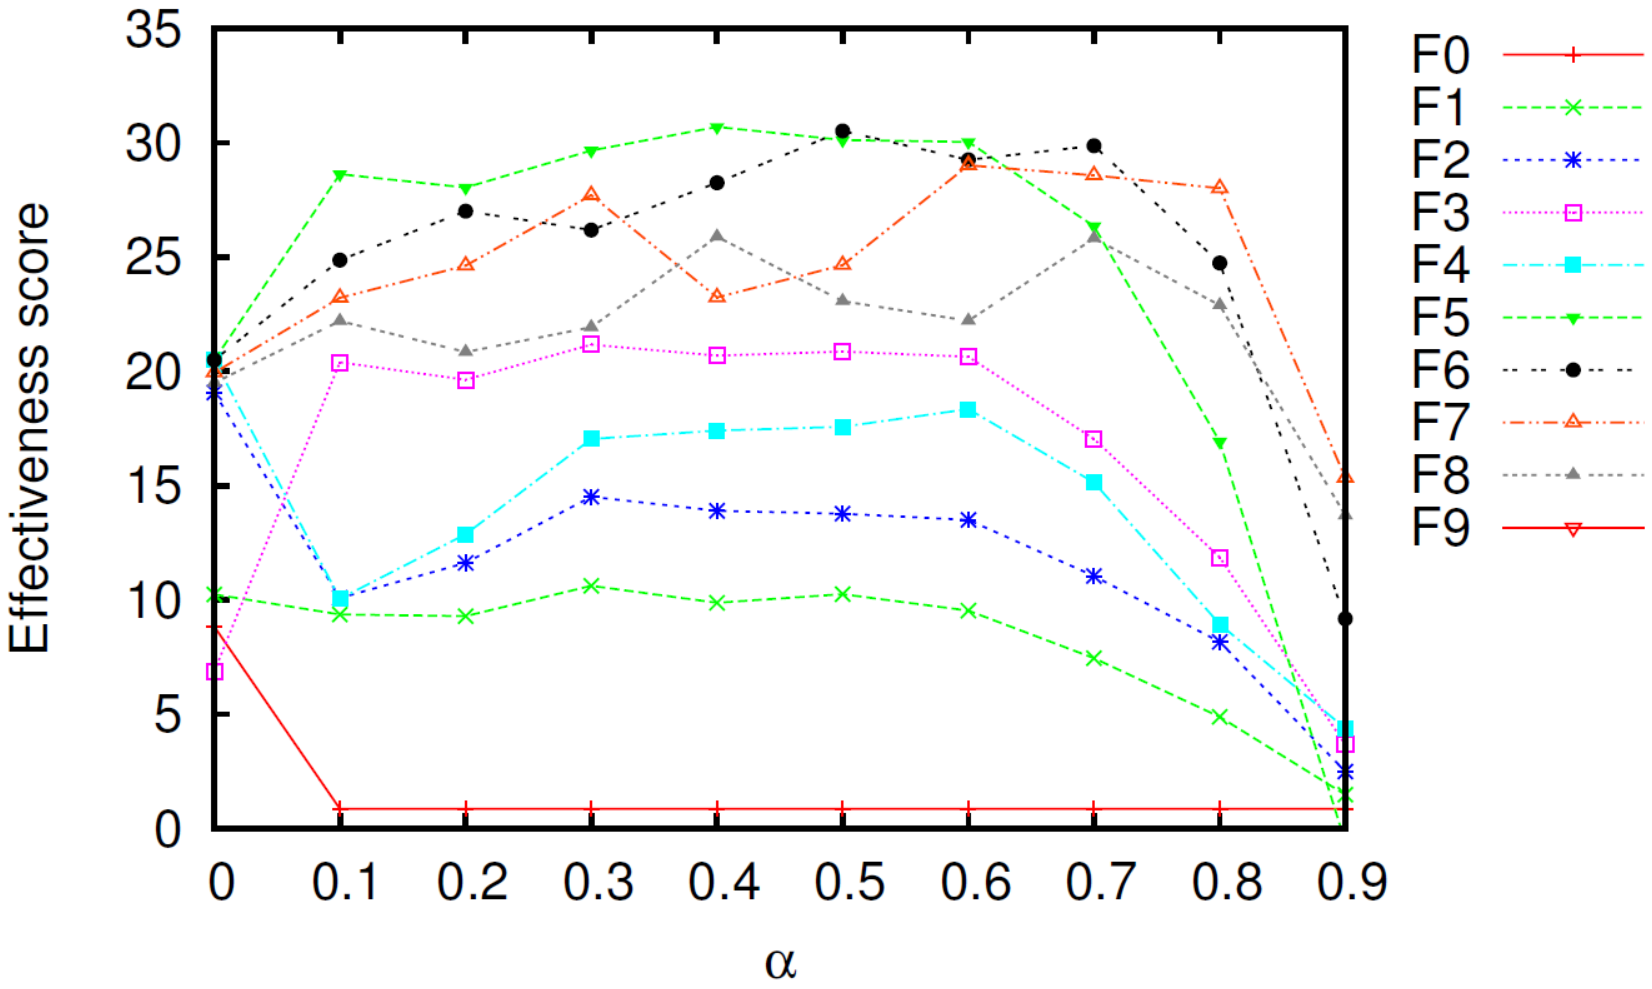

$\beta = 30$